

R-SCoRe: Revisiting Scene Coordinate Regression for Robust Large-Scale Visual Localization

Xudong Jiang¹ Fangjinhua Wang^{1*} Silvano Galliani² Christoph Vogel² Marc Pollefeys^{1,2}
¹Department of Computer Science, ETH Zurich
²Microsoft Spatial AI Lab, Zurich

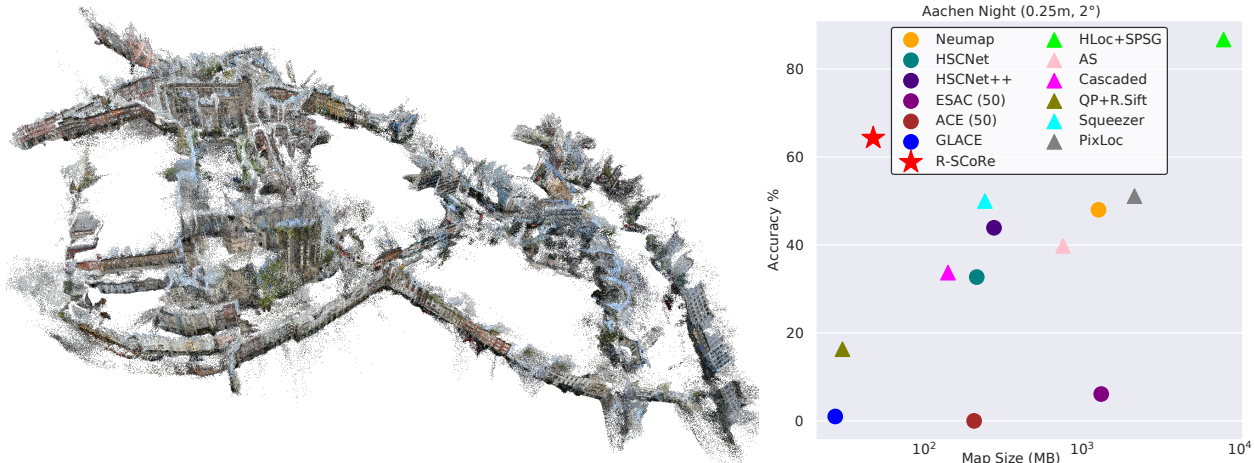


Figure 1. **Robust Visual Localization with R-SCoRe.** *Left:* Point cloud of Aachen reconstructed by R-SCoRe. *Right:* On the large-scale Aachen Day-Night dataset [53, 56] using only daytime training images, R-SCoRe achieves 64.3% accuracy under the (0.25m, 2°) threshold for nighttime query images. It outperforms all previous SCR methods (circles) by a large margin. With a small map size of only 47MB at a comparable accuracy, R-SCoRe is an attractive alternative to traditional methods (triangles).

Abstract

Learning-based visual localization methods that use scene coordinate regression (SCR) offer the advantage of smaller map sizes. However, on datasets with complex illumination changes or image-level ambiguities, it remains a less robust alternative to feature matching methods. This work aims to close the gap. We introduce a covisibility graph-based global encoding learning and data augmentation strategy, along with a depth-adjusted reprojection loss to facilitate implicit triangulation. Additionally, we revisit the network architecture and local feature extraction module. Our method achieves state-of-the-art on challenging large-scale datasets without relying on network ensembles or 3D supervision. On Aachen Day-Night, we are 10× more accurate than previous SCR methods with similar map sizes and require at least 5× smaller map sizes than any other SCR method while still delivering superior accuracy. Code is available at: <https://github.com/cvg/scrstudio>.

*Corresponding author (fangjinhua.wang@inf.ethz.ch).

1. Introduction

Visual localization is the task of estimating the 6-DoF pose of a camera in a known scene with a query image. It is a fundamental problem in computer vision, with applications in augmented reality, autonomous driving, and robotics.

Classical feature matching methods [19, 50, 51, 55] have matured through years of research and now provide robust and accurate localization results. However, these methods typically require explicit 3D scene representations, where a large number of descriptors are stored, leading to substantial map sizes, especially for large-scale scenes. In contrast, pose regression [11, 30, 31, 44, 59, 65, 71, 73] and scene coordinate regression (SCR) [4–6, 8, 9, 21, 36, 68] aim to implicitly encode scene information in neural networks.

SCR methods follow a structure-based paradigm similar to feature matching, *i.e.*, estimating pose from 2D-3D correspondences but replacing explicit matching with regressing the correspondences directly. They are usually limited to small scenes [9] and have yet to match feature matching methods in terms of accuracy and robustness. Recent ad-

vances [68] extend SCR to large-scale scenes using a single model. However, its performance is still not on par with feature matching methods, especially in complex scenes with challenging illumination changes [53, 56].

In this work, we conduct a detailed analysis of the design principles behind the SCR framework, including local and global encoding, network architecture, and training strategies. Based on this analysis, we propose to revisit SCR to enhance the robustness and accuracy of SCR methods for large-scale visual localization tasks. As shown in Fig. 1, our robust SCR (R-SCoRe) improves the night-time localization accuracy for SCR methods to 64.3% under the (0.25m, 2°) threshold on the Aachen Day-Night dataset, all with a map size of only 47MB. R-SCoRe significantly outperforms previous SCR methods and achieves accuracy comparable to feature matching techniques.

Tab. 1 summarizes the practicability of R-SCoRe in complex large-scale scenes. While feature matching methods are also accurate, their map size can be prohibitively large, sometimes more than two orders of magnitude [19, 50, 51]. Compared to SCR methods with similarly small map sizes [9, 68], R-SCoRe is at least one order of magnitude more accurate. While we still clearly outperform other SCR methods, we maintain fast inference and significantly smaller map sizes – all without the need for scene-specific depth supervision. Our contributions are as follows.

- We propose learning a global encoding and performing data augmentation based on the covisibility graph. To address the ambiguity of image retrieval features in complex large-scale scenes, we used multiple global hypotheses during testing.
- To unbias the network from neglecting near points, we introduce a depth-adjusted reprojection loss and show that this allows for accurate localization without scene-specific ground truth coordinate supervision.
- To our knowledge, R-SCoRe is the first attempt of an SCR approach to achieve state-of-the-art performance on complex large-scale scenes without using an ensemble of networks or 3D model supervision.

2. Related Work

Feature Matching. Most state-of-the-art visual localization methods rely on feature matching [19, 50, 51, 55]. These methods typically adopt a structure-based paradigm, establishing 2D-3D correspondences between keypoints in a query image and 3D points in a scene. Camera pose is solved with geometric constraints, often a Perspective-n-Point (PnP) solver [26, 47] within a RANSAC framework [2, 17, 25] to effectively manage outliers. These methods commonly construct a Structure-from-Motion (SfM) map of the scene, which contains both 3D points and their descriptors [19, 22, 24, 49, 58]. To efficiently establish 2D-

	Methods	Size	Time	Acc.	w/o Depth
FM	[19, 50, 51]	✗	✓	✓	✓
PR	[31]	✓	✓	✗	✓
SCR	[64]	✗	✗	✓	✗
	[5, 36, 70]	✗	✓	✓	✗
	[4, 9, 68]	✓	✓	✗	✓
	R-SCoRe	✓	✓	✓	✓

Table 1. **Comparison with other methods on complex large-scale scenes** [53, 56]. Feature matching (FM) methods are accurate but need a large map size. Pose regression (PR) methods are fast but less accurate. We maintain a small map size while achieving remarkable accuracy.

3D matches, they often follow a two-level approach [50]. First, potentially relevant database images are retrieved using image retrieval techniques [1, 48, 75]. This is followed by 2D-2D matching with the query image [39, 50, 51].

However, a significant limitation of these methods is the necessity to store all descriptor vectors of the 3D model, which can lead to storage challenges, particularly in large maps. To address this issue, various compression techniques have been proposed. These include reducing the number of 3D points [12, 13, 23, 37, 72] or compressing descriptors [20, 28, 29, 34, 40, 54, 67, 72]. Recently, several studies [46, 69, 74] have proposed alternative approaches that eliminate the need for explicit descriptor storage. Instead, these methods advocate for direct matching against geometric representations, such as point clouds or meshes.

Pose Regression. These methods [11, 30, 31, 44, 59, 65, 71, 73] directly estimates the camera pose of a query image with a neural network. However, they tend to struggle with generalization and often only achieve an accuracy similar to image retrieval methods [57].

Scene Coordinate Regression. Following a similar structure-based localization paradigm as feature matching methods, SCR methods [4–9, 14, 15, 21, 36, 45, 60, 66, 68, 70] regress 2D-3D correspondences between the query image and the scene, and estimate the camera pose using geometric constraints. Though SCR methods are more accurate than pose regression methods, they usually still struggle with large-scale scenes [9].

Recently, several approaches have been proposed to improve the scalability and performance of SCR in large-scale scenes. These methods often rely on ground truth 3D coordinates and aim to handle large scenes by dividing them into smaller segments, such as spatial regions [5], voxels [64], or hierarchical clusters [36, 70]. Recent advancements have explored alternatives that do not require ground truth 3D supervision. For example, ACE [9] uses reprojection loss only for training. GLACE [68] further introduces a global

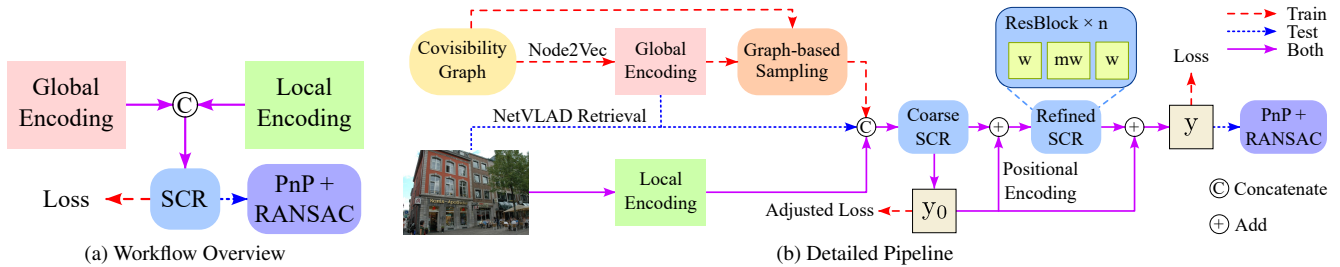


Figure 2. **R-SCoRe pipeline.** (a) Following the SCR workflow in [68], we concatenate patch-level local encodings with image-level global encodings as input to a scene-specific MLP. (b) We learn contrastive global encodings from the covisibility graph using Node2Vec [27]. During training, global encodings are sampled from neighboring nodes for data augmentation. During inference, we retrieve global encodings from the k nearest training images via NetVLAD [1] as hypotheses and select the one yielding the most RANSAC inliers. We enhance the SCR MLP with a refinement module and introduce a depth-adjusted reprojection loss to reduce bias toward distant points.

encoding mechanism that eliminates the need for scene segmentation. Despite these improvements, these methods still encounter limitations under challenging conditions, such as significant changes in illumination.

3. Method

Our Scene Coordinate Regression (SCR) workflow is depicted in Fig. 2a. During training, we have access to a covisibility graph to learn global encodings which are concatenated to the local encodings. During inference, we retrieve global encoding hypotheses from the k nearest training images and predict 2D-3D correspondences: we run PnP for each hypothesis and select the one yielding the most RANSAC inliers. Fig. 2b illustrates our detailed R-SCoRe pipeline, where we split SCR into coarse and refinement blocks and introduce covisibility-based global encoding and data augmentation techniques.

3.1. Preliminaries

Visual Localization. Given a test image I_{test} with known intrinsic K_{test} , the goal is to estimate its extrinsic, i.e., the rigid transformation $[R_{\text{test}} | t_{\text{test}}]$ from world coordinate to camera coordinate. The scene is typically given by a set of training images I_{train} with known ground truth poses $[R_{\text{train}} | t_{\text{train}}]$ and intrinsics K_{train} .

Scene Coordinate Regression. The SCR pipeline employs a neural network f to directly regress the 3D coordinate $y = f(\mathbf{F}(x))$ for each 2D keypoint x with feature $\mathbf{F}(x)$. Without the need to store large point clouds with descriptors, SCR methods implicitly represent the scene with a neural network, which usually results in a smaller map size.

Scalable SCR without 3D ground truth. Recent advances [9, 68] allow SCR to scale to large scenes without scene-specific 3D supervision. To reduce ambiguities in large scenes, GLACE [68] (Fig. 2a) concatenates a local patch-level encoding with an image-level global encoding

as the keypoint feature $\mathbf{F}(x)$. The local encoder is a pre-trained DSAC* [6] backbone, following [9]. The global encoder uses a pre-trained image retrieval model [75] with Gaussian noise augmentation to prevent overfitting to trivial solutions, *cf.* [68]. To accelerate training, all features are precomputed and buffered in GPU memory, from which a random sample is drawn for each batch.

Without ground truth scene coordinates, the output 3D point y is reprojected with the ground truth pose R, t and intrinsics K , and compared to the keypoint location x in 2D:

$$e_2(x, y) = \|x - \pi(K(Ry + t))\|_2, \quad (1)$$

where π converts homogeneous to Cartesian coordinates.

Instead of explicitly grouping corresponding observations into tracks, this underconstrained supervision is applied to each independent prediction. Prior works [45, 68] suggest that implicit triangulation can still occur as the network tends to produce similar outputs for similar inputs.

The reprojection error is fed into a dynamic robust loss, *cf.* ACE [9], to focus on points accurately regressed:

$$l_{\text{dynamic}}(e_2(x, y)) = \tau(t)\rho\left(\frac{e_2(x, y)}{\tau(t)}\right), \quad (2)$$

where ACE [9] uses tanh as robust loss ($\rho := \tanh$). Based on the relative training time $t \in [0, 1]$, the bandwidth $\tau(t)$ is adjusted dynamically during training:

$$\tau(t) = \sqrt{1 - t^2}\tau_{\text{max}} + \tau_{\text{min}}. \quad (3)$$

Keypoints x whose regressed 3D point y fall outside the valid frustum are penalized differently. Valid points are defined to lie within a valid depth range $[d_{\text{min}}, d_{\text{max}}]$ in front of the camera. Further, their reprojection error $e_2(x, y)$ must be smaller than a threshold e_{max} . For invalid points, we penalize their distance to a pseudo ground truth point \bar{y} :

$$l_{\text{invalid}}(y) = \|y - \bar{y}\|_2, \quad (4)$$

where \bar{y} is computed by the inverse projection of the pixel x using a fixed target depth d_{target} .

3.2. Network Architecture

We adopt the MLP architecture and position decoder from GLACE [68], scaling the network width with scene size. As illustrated in Fig. 2b, we also introduce a refinement module at the end of the network, which adjusts the final output y by predicting an offset from the intermediate prediction y_0 . The coarse coordinate y_0 is reintroduced into the refinement module through positional encoding [43, 63] using sine and cosine functions with periods ranging from 0.5 to 2048, which is added to the intermediate feature. This empirically improves training stability and allows the network to achieve lower training reprojection errors more rapidly.


3.3. Input Encoding

Analysis. In implicit triangulation, reprojection constraints are grouped based on input similarity. Therefore, the desired properties of input encodings are as follows: positive pairs observing the same points should produce similar features, while negative pairs observing distinct points should yield clearly distinguishable features. Additionally, it is preferable for the encodings to be low-dimensional to minimize memory requirements.


Local Encoding. At the local patch level, features should differentiate observations of the same point from those of different points. The requirement aligns with the properties of local descriptors used in traditional feature matching, suggesting that we can directly leverage their local feature extractors. We investigate pretrained feature extractors for both dense and sparse matching methods, such as LoFTR [62] and Dedode [24]. To lower memory consumption during training, we apply PCA to all the features from the training dataset, reducing their dimensionality while retaining most of the variance. We experimentally observe that reducing the dimensionality to 128 dimensions preserves over 90% of the variance on various datasets [35, 53, 56].

Covisibility Graph Based Global Encoding. Image-level global features should distinguish between covisible and non-covisible image pairs, i.e. whether the images are viewing the same part of the scene, to resolve ambiguities in local encodings. Although global encodings with image-level receptive fields can help, they may still be insufficient to resolve ambiguities in complex environments, as shown in Fig. 3. This limitation can lead to imperfect grouping of reprojection constraints during training, thereby impairing the effectiveness of implicit triangulation. Furthermore, we point out that the learned SCR function may lack (Lipschitz) smoothness [32] w.r.t. the global encodings if adapted naively, e.g. minor variations in the global encoding can result in significant shifts in corresponding 3D points and consequently reduce generalization at test time.

To address these issues, we propose to directly learn



(a) Distribution of feature distance for covisible and non-covisible pairs.



(b) Precision Recall Curve of predicting covisibility by feature distance.

Figure 3. Comparison of global encodings. Aligning the learning of global encodings with the covisibility graph topology (Node2Vec [27]) helps distinguish covisible and non-covisible pairs (a) and predict covisibility by feature distance (b).

embeddings aligned with the covisibility graph’s topology using Node2Vec [27], which samples sequences with weighted random walks and optimizes node embeddings with a Skip-gram [42] objective. For training images, the covisibility graph is easily available. It can be estimated from the frustum overlap of ground truth poses (see the supplementary for more details). At test time, however, covisibility information is unknown, so we propose generating multiple global encoding hypotheses by retrieving the nearest training images using NetVLAD [1] features. The global encoding of each retrieved image serves as a hypothesis, and we select the hypothesis yielding the maximum RANSAC inliers for the final localization result.

This approach effectively decouples training-time and test-time ambiguities: during training, the network focuses on learning scene structure without ambiguity, while at test time, multiple hypotheses enable the resolution of complex, often multimodal ambiguities.

Covisibility Graph Based Data Augmentation. Our Covisibility Graph Encoding effectively learns a low-ambiguity global encoding. However, data augmentation is still necessary to prevent the network from distinguishing covisible pairs based on distinct global encodings. Instead of simply adding isotropic Gaussian noise [68], we introduce a graph-based data augmentation strategy. In this approach, rather than applying isotropic noise, we randomly replace an image’s global encoding with that of a neighboring image from the covisibility graph. Specifically, with probabil-




Figure 4. **Statistics of reprojection error for points with different depths.** The kernel density estimation (KDE) of reprojection error distribution conditioned on disparity from SCR models trained with various local encodings across different datasets. We observe that far points (low disparity) exhibit a lower reprojection error. (Detector-free LoFTR [62] with an $8 \times$ downsampled output has a larger 2D keypoint error than Detector-based Dedode [24].)

ity $p = 0.5$, the current image’s global encoding is retained, while with probability $1 - p$, it is replaced by the global encoding of a randomly sampled neighboring image.

3.4. Output Supervision

Depth Bias in Reprojection Loss. Fig. 4 displays statistics collected from training SCR models with various local encodings across different datasets. It indicates that points closer to the camera empirically exhibit higher reprojection errors compared to distant points, hence we observe a bias toward distant points. This bias is magnified by training with the robust loss Eq. (2), as the supervision signal tends to neglect (nearby) points with higher reprojection errors (*cf.* Fig. 5). Assuming the training time distribution of camera poses to be representative for testing, regressing near points fewer and less accurately during testing diminishes the positional localization accuracy. We conjecture that to facilitate implicit triangulation for near points, a higher reprojection error should be allowed to compensate for the reprojection of nearby points being more sensitive to pose variations and coordinate inaccuracies.

Depth Adjusted Reprojection Error. We propose normalizing the reprojection error in Eq. (1) based on the depth of the predicted scene coordinate. Specifically, the observation standard deviation, σ_o , is defined as:

$$\sigma_o = \sqrt{\left(\frac{\sigma_3}{d}\right)^2 + \sigma_2^2}, \quad (5)$$

where we denote the variance of the noise of the 2D observations by σ_2^2 . The variance of the 3D prediction is denoted by σ_3^2 and by d the depth of the point. The reprojection error is then adjusted accordingly:

$$e_3(x, y) = \frac{e_2(x, y)}{\sigma_o} = \frac{e_2(x, y)}{\sigma_2} \sqrt{\frac{d^2}{d^2 + \left(\frac{\sigma_3}{\sigma_2}\right)^2}}. \quad (6)$$

Selective Application of Depth Adjustment. The bias towards distant points may sometimes actually be beneficial, as underconstrained points can be pushed farther along the ray, making them easier to identify as outliers during test time. Therefore, we apply our depth-adjusted reprojection loss only to the intermediate coarse scene coordinate output y_0 during training (Fig. 2b), and retain the original reprojection loss for the final output. To mitigate the concentration of the supervision signal on points with low projection error, we replace \tanh with the Geman-McClure [3] robust loss function which has a heavier tail than \tanh :

$$\rho(x) = \frac{9x^2}{9x^2 + 4}. \quad (7)$$

In order to guide the convergence of the regressed points, y , by the intermediate output y_0 (affected by the normalized loss from Eq. (6)), we also apply a consistency loss at the beginning of the training:

$$l_{\text{consistency}} = \lambda(t) \|y - y_0\|_2, \quad (8)$$

where $\lambda(t)$ is a dynamic weight that decreases to 0 in a cosine schedule during the first 50% of training time.

$$\lambda(t) = \begin{cases} \frac{1}{2} (1 + \cos 2\pi t), & \text{if } t \in [0, 0.5] \\ 0, & \text{otherwise} \end{cases}, \quad (9)$$

where t is the relative training time.

Optional Depth Supervision. When depth is available, we can also benefit from direct depth supervision. The depth does not need to be accurate since we mainly use the depth for initialization. Specifically, we simply replace the consistency loss between intermediate and final output in Eq. (8) with a ground truth coordinate supervision loss:

$$l_{\text{depth}} = \lambda(t) (\|y - \bar{y}\|_2 + \|y_0 - \bar{y}\|_2), \quad (10)$$

where \bar{y} is the pseudo ground truth computed by the inverse projection of the pixel given the depth, pose, and intrinsic.

4. Experiments

4.1. Datasets

We use the Aachen Day-Night [53, 56] and the Hyundai Department Store dataset [35] to evaluate R-SCoRe on complex large-scale indoor and outdoor scenes.

Aachen Day-Night. It is a large-scale benchmark for outdoor visual localization, covering the historic inner city of Aachen, Germany, over an area of approximately 6 km^2 . It presents significant challenges due to varying illumination conditions, especially between day and night. The dataset includes 4,328 daytime images for training, along with 824 daytime query images and 98 nighttime query images.

Methods	w/o Depth	Size	Aachen Day			Aachen Night		
HLoc+SPSG [19, 50, 51]	Yes	7.82GB	89.6	95.4	98.8	86.7	93.9	100
AS [55]	Yes	750MB	85.3	92.2	97.9	39.8	49.0	64.3
Cascaded [16]	Yes	140MB	76.7	88.6	95.8	33.7	48.0	62.2
QP+R.Sift [41]	Yes	30MB	62.6	76.3	84.7	16.3	18.4	24.5
Squeezer [72]	Yes	240MB	75.5	89.7	96.2	50.0	67.3	78.6
PixLoc [52]	Yes	2.13GB	64.3	69.3	77.4	51.1	55.1	67.3
Neumap [64]	No	1.26GB	80.8	90.9	95.6	48.0	67.3	87.8
HSCNet [36]	No	213MB	71.1	81.9	91.7	32.7	43.9	65.3
HSCNet++ [70]	No	274MB	72.7	81.6	91.4	43.9	57.1	76.5
ESAC ($\times 50$) [5]	No	1.31GB	42.6	59.6	75.5	6.1	10.2	18.4
ACE ($\times 50$) [9]	Yes	205MB	6.9	17.2	50.0	0.0	1.0	5.1
GLACE [68]	Yes	27MB	8.6	20.8	64.0	1.0	1.0	17.3
R-SCoRe (Dedode [24])	Yes	47MB	74.8	86.9	96.4	64.3	89.8	96.9
+ Depth	No	47MB	79.0	88.5	96.4	66.3	89.8	96.9

Table 2. **Aachen Day-Night evaluation.** The map size and percentages of query images within three thresholds: (0.25m, 2°), (0.5m, 5°), and (5m, 10°) and are reported. We report our results with Dedode [24] local encoding and optional depth supervision. Feature matching (FM) methods [19, 50, 51] are more accurate, but the map size is large. R-SCoRe achieves comparable accuracy with a small map size.

Hyundai Department Store. It is a large-scale indoor visual localization benchmark, covering three floors of a department store. Each floor consists of multiple sequences captured over four months, spanning an area of approximately 10,000 m^2 . It presents challenges beyond its large scale, including dynamic objects, illumination changes, and textureless regions. B1 is particularly challenging as the training images are captured under low-lighting conditions, while the query images are brightly illuminated. The dataset includes 44,283 training images and 5,927 test images.

4.2. Benchmark Results

Aachen Day-Night. As shown in Tab. 2, R-SCoRe enhances the performance of scene coordinate regression (SCR) based methods [5, 36, 68, 70], achieving competitive results with a single low-map-size model without the need for scene-specific depth supervision. While R-SCoRe is competitive with the best performing method, HLoc [19, 50, 51], we forfeit some ground at highest accuracy. However, R-SCoRe demands 170 \times less memory to store the map. This huge gap could already render SCR based methods as an attractive alternative for some applications. Most other feature based methods (FM) also deliver significantly larger maps. While delivering comparable performance for the Aachen Day dataset, they all fall behind R-SCoRe on the Aachen Night dataset. The only FM method [41] with a comparable map size is outperformed on all metrics, *e.g.* [41] is 4-5 \times worse in accuracy on the night dataset. Compared to other SCR methods that work without depth supervision [9, 68] R-SCoRe is 10 \times superior in accuracy. The so far most accurate SCR based method [64] produces large maps (27 \times larger) and is prohibitively slow

in inference. The next most accurate SCR method, [70] is outperformed by 46% at night and highest threshold, while R-SCoRe maintains a 6 \times smaller map size – without the need for depth supervision. We observe a small gain in performance if we utilize depth for supervision of R-SCoRe.

Hyundai Department Store. R-SCoRe again significantly outperforms the SCR based methods [5, 68], including recent ensemble networks [5, 9], see Tab. 3. Compared to the state-of-the-art feature matching based localization [50] we achieve competitive results with a single low-map-size model and forfeit some ground at the highest accuracy threshold. However, our model is at least three orders of magnitude smaller for either feature [22, 49] incorporated into [50], which can be a valuable advantage in practice. Recall that depth supervision is not necessary for R-SCoRe, but if available, it can also enhance performance further. The B1 scene exhibits strong illumination changes and we observe significantly better performance when using local encodings from Dedode instead of LoFTR [62].

4.3. Implementation Details

Most hyperparameters follow default values [68] and extensive tuning is not performed, as we empirically find the approach remains robust within a reasonable range, apart from the trade-off between network size and performance.

Input Encodings. For local encodings, we perform PCA to reduce their dimensionality to 128. The global encodings are represented in 256 dimensions, consistent with the R^2 Former [75] feature dimension used in GLACE [68] for fair comparison. We estimate the covisibility graph based on camera poses, using a weighted frustum overlap method (details provided in the supplementary materials), with a

Methods	Dept. 1F Test	Dept. 4F Test	Dept. B1 Test
HLoc+D2-Net [22, 50]	(78.0 / 82.8 / 88.0) / 398GB	(84.2 / 89.8 / 92.0) / 183GB	(73.7 / 79.3 / 87.2) / 505GB
HLoc+R2D2 [49, 50]	(80.6 / 84.3 / 89.4) / 166GB	(85.3 / 91.0 / 93.1) / 76GB	(75.2 / 80.3 / 87.6) / 210GB
PoseNet [31]	(0.0 / 0.0 / 0.4) / 41MB	(0.0 / 0.0 / 0.1) / 41MB	(0.0 / 0.0 / 0.1) / 41MB
ESAC ($\times 50$) [5]	(43.3 / 66.3 / 77.0) / 1.4GB	(45.2 / 62.5 / 73.1) / 1.4GB	(3.5 / 8.2 / 12.6) / 1.4GB
ACE ($\times 50$) [9]	(14.1 / 54.4 / 75.5) / 205MB	(27.3 / 70.9 / 84.1) / 205MB	(2.7 / 14.4 / 29.3) / 205MB
GLACE [68]	(5.6 / 21.3 / 48.6) / 42MB	(8.4 / 29.8 / 51.6) / 42MB	(0.9 / 4.4 / 11.9) / 42MB
R-SCoRe (LoFTR* [62])	(63.2 / 82.4 / 92.4) / 127MB	(62.2 / 82.7 / 90.9) / 50MB	(26.9 / 50.7 / 69.6) / 130MB
+ Depth	(67.3 / 84.5 / 92.6) / 127MB	(70.5 / 87.0 / 92.9) / 50MB	(30.8 / 53.7 / 72.7) / 130MB
R-SCoRe (Dedode [24])	(61.4 / 80.2 / 90.9) / 127MB	(60.2 / 79.3 / 87.9) / 50MB	(60.1 / 77.3 / 89.6) / 130MB
+ Depth	(63.9 / 83.3 / 90.8) / 127MB	(76.7 / 89.3 / 93.0) / 50MB	(61.5 / 77.6 / 88.8) / 130MB

Table 3. **Hyundai Department Store Test Set evaluation.** The percentages of query images within three thresholds: (0.1m, 1°), (0.25m, 2°), and (1m, 5°) and the map size are reported. R-SCoRe achieves competitive accuracy with a small map size. *We use LoFTR [62] outdoor, trained on MegaDepth [38], instead of the indoor model trained on ScanNet [18] for the B1 scene with strong illumination change.

maximum viewing frustum depth of $d_v = 50$ for outdoor scenes and $d_v = 8$ for indoor scenes.

Output Supervision. The supervision uses a dynamic robust loss bandwidth strategy inspired by ACE [9]. For coarse intermediate outputs, the parameters, see Eq. (3), are set to $\tau_{\min} = 1$ and $\tau_{\max} = 50$. In contrast, $\tau_{\max} = 25$ is used for the final output, which allows the refinement layer to focus on the most accurate predictions while the initial layers do not ignore the optimization of relatively inaccurate predictions. Fixing $\sigma_2 = 1$ in the depth-adjusted reprojection loss, Eq. (6), allows us to control the behavior by adjusting τ and $\frac{\sigma_3}{\sigma_2}$. For indoor scenes, $\sigma_3 = 3$ is applied, while $\sigma_3 = 8$ is used for outdoor scenes to account for different scales. We perform optional depth supervision using depth images rendered from the 3D model for the Hyundai Department Store dataset and Multi-View Stereo depth maps for the Aachen Day-Night dataset.

Network Architecture. We adopt the MLP architecture and position decoder from GLACE [68] with expansion ratio $m = 2$ for the MLP, and 50 clusters for the position decoder. With MLP width $w = 256 \lceil \sqrt{n/1000} \rceil$ for n training images, we scale the parameter count proportionally.

Training. We found that adopting the optimization settings from ACE Zero [10] enhances both stability and convergence speed compared to the original ACE [9]. Specifically, we reduce the warmup ratio of the one-cycle learning rate schedule [61] from 0.25 to 0.04 and lower the peak learning rate from 5×10^{-3} to 3×10^{-3} . For our evaluation we adopt similar training parameters to GLACE [68], including a local feature buffer size of 128M, a large batch size of 320K and a training duration of 100k iterations.

Testing. At test time, we retrieve the 10 nearest training images with NetVLAD [1]. The global encoding and retrieval features for training images are precomputed and compressed using Product Quantization [28]. For final pose

	Dept. 1F Val			Dept. B1 Val		
ACE [9]	68.7	87.5	95.9	14.1	28.3	45.8
LoFTR* [62]	72.3	88.7	95.5	29.4	51.3	69.6
Dedode [24]	70.6	86.6	95.5	57.7	74.7	86.7

Table 4. **Ablation study of local encoders.** Accuracy at (0.1m, 1°), (0.25m, 2°), and (1m, 5°) thresholds are reported. Utilizing pretrained, off-the-shelf feature extractors improves the performance, especially under challenging conditions (B1).

	Dept. 1F Val		
R^2 Former [75] w/ Gaussian	34.1	60.1	78.3
+ Multi Hypotheses	42.1	74.5	92.2
+ Covis Augmentation	62.0	83.8	94.8
+ Covis Encoding	72.3	88.7	95.5

Table 5. **Ablation study of global encodings.** We experiment with using multiple hypotheses at test time, applying covisibility graph-based data augmentation during training, and learning global encodings directly from the covisibility graph. Accuracy at (0.1m, 1°), (0.25m, 2°), and (1m, 5°) thresholds.

estimation, we utilize PoseLib [33] with a maximum reprojection error of 10 pixels and up to 10,000 RANSAC iterations. On our PC (NVIDIA RTX 2080 GPU & Intel i7-9700K CPU), the average inference time for a 640×480 query image is 140 to 270 ms in total.

- Global: NetVLAD (20ms), Retrieval (<1ms)
- Local: LoFTR (7ms) or DeDoDe (50ms)
- MLP: $w = 768$ (70ms) or 1280 (160ms)
- Pose Solving: 40ms

4.4. Ablation Study Results

In our ablation studies, we investigate the impact of the different components in R-SCoRe. We evaluate on the validation split for the Hyundai Department Store dataset. Since the Aachen Day-Night dataset does not provide a validation

	Aachen Day			Aachen Night		
FM Covis	75.4	87.6	95.8	64.3	89.8	95.9
Pose Covis	74.8	86.9	96.4	64.3	89.8	96.9

Table 6. **Ablation study of covisibility graph.** Building the covisibility graph using frustum overlap performs similarly to utilizing feature matching. Accuracy at (0.25m, 2°), (0.5m, 5°), and (5m, 10°) thresholds.

Supervision	Dept. 1F Val			Dept. 4F Val		
Original	62.3	82.2	93.7	59.1	82.2	97.6
Adjusted	70.6	86.6	95.5	63.9	84.2	98.3
Depth	76.8	88.1	95.6	68.5	84.9	98.5

Table 7. **Ablation study of supervision methods.** *Original* refers to the original reprojection error supervision, *Adjusted* refers to our depth-adjusted reprojection error supervision, and *Depth* uses ground truth depth for supervision. Accuracy at (0.1m, 1°), (0.25m, 2°), and (1m, 5°).

split, we evaluate on the test set.

Local Encoding. As shown in Tab. 4, for large scale indoor scenes with small illumination changes, alternative off-the-shelf local feature extractors [24, 62] achieve similar or even superior performance compared to the original ACE [9]. Note that this finding contradicts earlier investigations [9] that prefer a specifically trained backbone in their work. Additionally, local descriptors trained on MegaDepth [38], especially Dedode [24], demonstrate greater robustness in scenes with significant illumination changes.

Global Encoding. Retrieving global encodings from training images avoids the domain gap. Better retrieval method and multiple hypotheses verification can help resolve ambiguities. Without retraining (Tab. 5), utilizing multiple global hypotheses at test time (+ *Multi Hypotheses*) results in a direct performance improvement in complex scenes. The performance improves significantly, once we incorporate our covisibility graph-based data augmentation during training. In particular, we replace isotropic Gaussian noise [68] with our more precise covisibility-based technique (+ *Covis Augmentation*). Finally, learning the global encoding directly from the covisibility graph (+ *Covis Encoding*) reduces the interference between non-covisible training images and thereby facilitates implicit triangulation, especially in indoor scenes with significant ambiguity.

Finally, we also explore the effect of computing the covisibility graph via feature matching [19, 51]. As shown in Tab. 6, using a more accurate graph yields no significant improvement, indicating that R-SCoRe is robust to the quality of the covisibility graph. Therefore, our simple frustum overlap-based graph is sufficient for effective performance.

Supervision. Our depth-adjusted supervision effectively




Figure 5. **Ablation study of depth distribution after training with different supervision methods.** Our depth-adjusted supervision matches the distribution of ground truth depth for supervision as compared to the original.

mitigates the bias towards distant points and enhances the implicit triangulation of nearby points. As demonstrated in Fig. 5, depth-adjusted supervision significantly alters the depth distribution of predicted points, alleviating the previous ignorance of nearby points. This adjustment brings the distribution closer to that achieved with ground truth depth supervision, demonstrating a substantial reduction in the bias inherent in the original supervision approach.

In Tab. 7, we observe that depth-adjusted supervision also leads to notable improvements in localization accuracy, particularly under stricter thresholds, where accurate translation estimation relies heavily on near points. Even without ground truth depth supervision, depth-adjusted supervision enables the model to achieve competitive performance.

5. Conclusion

In this work, we revisited scene coordinate regression (SCR) methods for robust visual localization in large-scale, complex environments. We analyzed the design principles of input encoding and training strategies, identifying several areas for enhancement. Our proposed R-SCoRe includes a covisibility graph-based global encoding learning and data augmentation strategy, a depth-adjusted reprojection loss to improve the implicit triangulation, and also other improvements including better architecture and local feature. Our contributions advance the state-of-the-art in SCR and demonstrate that SCR-based localization methods can achieve competitive performance in large-scale applications. While operating at comparably very small map sizes, R-SCoRe trails the state-of-the-art FM-based localization methods only at the strictest error thresholds. Although out-of-distribution generalization remains challenging, and gaps persist in handling extreme cases, given the relatively small history of SCR, we are positive the accuracy gap can be closed completely in the near future.

References

- [1] R. Arandjelović, P. Gronat, A. Torii, T. Pajdla, and J. Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. In *CVPR*, 2016. 2, 3, 4, 7
- [2] Daniel Barath, Jana Noskova, Maksym Ivashechkin, and Jiri Matas. MAGSAC++, a fast, reliable and accurate robust estimator. In *CVPR*, 2020. 2
- [3] Jonathan T Barron. A general and adaptive robust loss function. In *CVPR*, 2019. 5
- [4] Eric Brachmann and Carsten Rother. Learning Less is More - 6D Camera Localization via 3D Surface Regression. In *CVPR*, 2018. 1, 2
- [5] Eric Brachmann and Carsten Rother. Expert sample consensus applied to camera re-localization. In *ICCV*, 2019. 2, 6, 7
- [6] Eric Brachmann and Carsten Rother. Visual camera re-localization from RGB and RGB-D images using DSAC. *TPAMI*, 2021. 1, 3
- [7] Eric Brachmann, Frank Michel, Alexander Krull, Michael Y. Yang, Stefan Gumhold, and Carsten Rother. Uncertainty-driven 6D pose estimation of objects and scenes from a single RGB image. In *CVPR*, 2016.
- [8] Eric Brachmann, Alexander Krull, Sebastian Nowozin, Jamie Shotton, Frank Michel, Stefan Gumhold, and Carsten Rother. DSAC-Differentiable RANSAC for camera localization. In *CVPR*, 2017. 1
- [9] Eric Brachmann, Tommaso Cavallari, and Victor Adrian Prisacariu. Accelerated coordinate encoding: Learning to relocalize in minutes using rgb and poses. In *CVPR*, 2023. 1, 2, 3, 6, 7, 8
- [10] Eric Brachmann, Jamie Wynn, Shuai Chen, Tommaso Cavallari, Áron Monzpart, Daniyar Turmukhambetov, and Victor Adrian Prisacariu. Scene coordinate reconstruction: Posing of image collections via incremental learning of a relocalizer. In *ECCV*, 2024. 7
- [11] Samarth Brahmabhatt, Jinwei Gu, Kihwan Kim, James Hays, and Jan Kautz. Geometry-aware learning of maps for camera localization. In *CVPR*, 2018. 1, 2
- [12] Federico Camposco, Andrea Cohen, Marc Pollefeys, and Torsten Sattler. Hybrid scene compression for visual localization. In *CVPR*, 2019. 2
- [13] Song Cao and Noah Snavely. Minimal scene descriptions from structure from motion models. In *CVPR*, 2014. 2
- [14] Tommaso Cavallari, Stuart Golodetz, Nicholas A Lord, Julien Valentin, Luigi Di Stefano, and Philip HS Torr. On-the-fly adaptation of regression forests for online camera re-localisation. In *CVPR*, 2017. 2
- [15] Tommaso Cavallari, Stuart Golodetz, Nicholas A. Lord, Julien Valentin, Victor A. Prisacariu, Luigi Di Stefano, and Philip H. S. Torr. Real-time rgb-d camera pose estimation in novel scenes using a relocalisation cascade. *TPAMI*, 2019. 2
- [16] Wentao Cheng, Weisi Lin, Kan Chen, and Xinfeng Zhang. Cascaded parallel filtering for memory-efficient image-based localization. In *CVPR*, 2019. 6
- [17] Ondřej Chum, Jiří Matas, and Josef Kittler. Locally optimized RANSAC. In *Pattern Recognition*. Springer Berlin Heidelberg, 2003. 2
- [18] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In *CVPR*, 2017. 7
- [19] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *CVPRW*, 2018. 1, 2, 6, 8
- [20] Hao Dong, Xieyuanli Chen, Mihai Dusmanu, Viktor Larsson, Marc Pollefeys, and Cyrill Stachniss. Learning-based dimensionality reduction for computing compact and effective local feature descriptors. In *ICRA*. IEEE, 2023. 2
- [21] Siyan Dong, Shuzhe Wang, Yixin Zhuang, Juho Kannala, Marc Pollefeys, and Baoquan Chen. Visual localization via few-shot scene region classification. In *3DV*, 2022. 1, 2
- [22] Mihai Dusmanu, Ignacio Rocco, Tomas Pajdla, Marc Pollefeys, Josef Sivic, Akihiko Torii, and Torsten Sattler. D2-Net: A Trainable CNN for Joint Detection and Description of Local Features. In *CVPR*, 2019. 2, 6, 7
- [23] Marcin Dymczyk, Simon Lynen, Titus Cieslewski, Michael Bosse, Roland Siegwart, and Paul Furgale. The gist of maps - summarizing experience for lifelong localization. In *ICRA*, 2015. 2
- [24] Johan Edstedt, Georg Bökman, Mårten Wadenbäck, and Michael Felsberg. DeDoDe: Detect, Don't Describe — Describe, Don't Detect for Local Feature Matching. In *3DV*. IEEE, 2024. 2, 4, 5, 6, 7, 8
- [25] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 1981. 2
- [26] Xiao-Shan Gao, Xiao-Rong Hou, Jianliang Tang, and Hang-Fei Cheng. Complete solution classification for the perspective-three-point problem. *IEEE TPAMI*, 25(8), 2003. 2
- [27] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016. 3, 4
- [28] Herve Jégou, Matthijs Douze, and Cordelia Schmid. Product quantization for nearest neighbor search. *IEEE TPAMI*, 33(1), 2011. 2, 7
- [29] Yan Ke and Rahul Sukthankar. PCA-SIFT: A more distinctive representation for local image descriptors. In *CVPR*, 2004. 2
- [30] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. *CVPR*, 2017. 1, 2
- [31] Alex Kendall, Matthew Grimes, and Roberto Cipolla. PoseNet: A convolutional network for real-time 6-dof camera relocalization. In *ICCV*, 2015. 1, 2, 7
- [32] Grigory Khromov and Sidak Pal Singh. Some fundamental aspects about lipschitz continuity of neural networks. In *ICLR*, 2024. 4
- [33] Viktor Larsson and contributors. PoseLib - Minimal Solvers for Camera Pose Estimation, 2020. 7
- [34] Zakaria Laskar, Iaroslav Melekhov, Assia Benbihi, Shuzhe Wang, and Juho Kannala. Differentiable product quantiza-

- tion for memory efficient camera relocalization. In *ECCV*, 2024. 2
- [35] Donghwan Lee, Soohyun Ryu, Suyong Yeon, Yonghan Lee, Deokhwa Kim, Cheolho Han, Yohann Cabon, Philippe Weinzaepfel, Guérin Nicolas, Gabriela Csurka, and Martin Humenberger. Large-scale localization datasets in crowded indoor spaces. In *CVPR*, 2021. 4, 5
- [36] Xiaotian Li, Shuzhe Wang, Yi Zhao, Jakob Verbeek, and Juho Kannala. Hierarchical scene coordinate classification and regression for visual localization. In *CVPR*, 2020. 1, 2, 6
- [37] Yunpeng Li, Noah Snavely, and Daniel P Huttenlocher. Location recognition using prioritized feature matching. In *ECCV*, 2010. 2
- [38] Zhengqi Li and Noah Snavely. Megadepth: Learning single-view depth prediction from internet photos. In *CVPR*, 2018. 7, 8
- [39] Philipp Lindenberger, Paul-Edouard Sarlin, and Marc Pollefeys. LightGlue: Local Feature Matching at Light Speed. In *ICCV*, 2023. 2
- [40] Simon Lynen, Torsten Sattler, Michael Bosse, Joel A Hesch, Marc Pollefeys, and Roland Siegwart. Get out of my lab: Large-scale, real-time visual-inertial localization. In *Robotics: Science and Systems*, 2015. 2
- [41] Marcela Mera-Trujillo, Benjamin Smith, and Victor Fragoso. Efficient scene compression for visual-based localization. In *3DV*, 2020. 6
- [42] Tomas Mikolov, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. In *ICLR*, 2013. 4
- [43] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 4
- [44] Tayyab Naseer and Wolfram Burgard. Deep regression for monocular camera-based 6-dof global localization in outdoor environments. In *IROS*, 2017. 1, 2
- [45] Son Tung Nguyen, Alejandro Fontan, Michael Milford, and Tobias Fischer. FocusTune: Tuning Visual Localization through Focus-Guided Sampling. In *2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. IEEE Computer Society, 2024. 2, 3
- [46] Vojtech Panek, Zuzana Kukelova, and Torsten Sattler. MeshLoc: Mesh-Based Visual Localization. In *ECCV*, 2022. 2
- [47] Mikael Persson and Klas Nordberg. Lambda twist: An accurate fast robust perspective three point (p3p) solver. In *ECCV*, 2018. 2
- [48] Filip Radenović, Giorgos Tolias, and Ondřej Chum. Fine-tuning cnn image retrieval with no human annotation. *IEEE TPAMI*, 41(7), 2018. 2
- [49] Jerome Revaud, Cesar De Souza, Martin Humenberger, and Philippe Weinzaepfel. R2d2: Reliable and repeatable detector and descriptor. In *Advances in Neural Information Processing Systems*, 2019. 2, 6, 7
- [50] Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, and Marcin Dymczyk. From coarse to fine: Robust hierarchical localization at large scale. In *CVPR*, 2019. 1, 2, 6, 7
- [51] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning feature matching with graph neural networks. In *CVPR*, 2020. 1, 2, 6, 8
- [52] Paul-Edouard Sarlin, Ajaykumar Unagar, Mans Larsson, Hugo Germain, Carl Toft, Viktor Larsson, Marc Pollefeys, Vincent Lepetit, Lars Hammarstrand, Fredrik Kahl, et al. Back to the feature: Learning robust camera localization from pixels to pose. In *CVPR*, 2021. 6
- [53] Torsten Sattler, Tobias Weyand, Bastian Leibe, and Leif Kobbelt. Image Retrieval for Image-Based Localization Revisited. In *BMVC*, 2012. 1, 2, 4, 5
- [54] Torsten Sattler, Michal Havlena, Filip Radenovic, Konrad Schindler, and Marc Pollefeys. Hyperpoints and fine vocabularies for large-scale location recognition. In *ICCV*, 2015. 2
- [55] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Efficient & effective prioritized matching for large-scale image-based localization. *IEEE TPAMI*, 39(9), 2016. 1, 2, 6
- [56] Torsten Sattler, Will Maddern, Carl Toft, Akihiko Torii, Lars Hammarstrand, Erik Stenborg, Daniel Safari, Masatoshi Okutomi, Marc Pollefeys, Josef Sivic, Fredrik Kahl, and Tomas Pajdla. Benchmarking 6DOF Outdoor Visual Localization in Changing Conditions. In *CVPR*, 2018. 1, 2, 4, 5
- [57] Torsten Sattler, Qunjie Zhou, Marc Pollefeys, and Laura Leal-Taixe. Understanding the limitations of cnn-based absolute camera pose regression. In *CVPR*, 2019. 2
- [58] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016. 2
- [59] Yoli Shavit, Ron Ferens, and Yosi Keller. Learning multi-scene absolute pose regression with transformers. In *ICCV*, 2021. 1, 2
- [60] Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew Fitzgibbon. Scene coordinate regression forests for camera relocalization in RGB-D images. In *CVPR*, 2013. 2
- [61] Leslie N. Smith and Nicholay Topin. Super-convergence: very fast training of neural networks using large learning rates. In *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications*, 2019. 7
- [62] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. LoFTR: Detector-free local feature matching with transformers. *CVPR*, 2021. 4, 5, 6, 7, 8
- [63] Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *NeurIPS*, 2020. 4
- [64] Shitao Tang, Sicong Tang, Andrea Tagliasacchi, Ping Tan, and Yasutaka Furukawa. Neumap: Neural coordinate mapping by auto-transdecoder for camera localization. In *CVPR*, 2023. 2, 6
- [65] Mehmet Özgür Türkoğlu, Eric Brachmann, Konrad Schindler, Gabriel Brostow, and Áron Monszpart. Visual Camera Re-Localization Using Graph Neural Networks and Relative Pose Supervision. In *3DV*, 2021. 1, 2

- [66] Julien Valentin, Matthias Nießner, Jamie Shotton, Andrew Fitzgibbon, Shahram Izadi, and Philip H. S. Torr. Exploiting uncertainty in regression forests for accurate camera relocalization. In *CVPR*, 2015. [2](#)
- [67] Ricardo Eugenio González Valenzuela, William Robson Schwartz, and Helio Pedrini. Dimensionality reduction through PCA over SIFT and SURF descriptors. In *11th International Conference on Cybernetic Intelligent Systems (CIS)*. IEEE, 2012. [2](#)
- [68] Fangjinhua Wang, Xudong Jiang, Silvano Galliani, Christoph Vogel, and Marc Pollefeys. Glace: Global local accelerated coordinate encoding. In *CVPR*, 2024. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [69] Shuzhe Wang, Juho Kannala, and Daniel Barath. Dgc-gnn: Leveraging geometry and color cues for visual descriptor-free 2d-3d matching. In *CVPR*, 2024. [2](#)
- [70] Shuzhe Wang, Zakaria Laskar, Iaroslav Melekhov, Xiaotian Li, Yi Zhao, Giorgos Toliás, and Juho Kannala. Hscnet++: Hierarchical scene coordinate classification and regression for visual localization with transformer. *International Journal of Computer Vision*, 2024. [2](#), [6](#)
- [71] Dominik Winkelbauer, Maximilian Denninger, and Rudolph Triebel. Learning to localize in new environments from synthetic training data. In *ICRA*, 2021. [1](#), [2](#)
- [72] Luwei Yang, Rakesh Shrestha, Wenbo Li, Shuaicheng Liu, Guofeng Zhang, Zhaopeng Cui, and Ping Tan. Scenesqueezer: Learning to compress scene for camera relocalization. In *CVPR*, 2022. [2](#), [6](#)
- [73] Qunjie Zhou, Torsten Sattler, Marc Pollefeys, and Laura Leal-Taixe. To learn or not to learn: Visual localization from essential matrices. In *ICRA*, 2020. [1](#), [2](#)
- [74] Qunjie Zhou, Sérgio Agostinho, Aljoša Ošep, and Laura Leal-Taixé. Is geometry enough for matching in visual localization? In *ECCV*, 2022. [2](#)
- [75] Sijie Zhu, Linjie Yang, Chen Chen, Mubarak Shah, Xiaohui Shen, and Heng Wang. R2former: Unified retrieval and reranking transformer for place recognition. In *CVPR*, 2023. [2](#), [3](#), [6](#), [7](#)