

Full-Frame Scene Coordinate Regression for Image-Based Localization

Xiaotian Li, Juha Ylioinas and Juho Kannala

Aalto University

firstname.lastname@aalto.fi

Abstract—Image-based localization, or camera relocation, is a fundamental problem in computer vision and robotics, and it refers to estimating camera pose from an image. Recent state-of-the-art approaches use learning based methods, such as Random Forests (RFs) and Convolutional Neural Networks (CNNs), to regress for each pixel in the image its corresponding position in the scene’s world coordinate frame, and solve the final pose via a RANSAC-based optimization scheme using the predicted correspondences. In this paper, instead of in a patch-based manner, we propose to perform the scene coordinate regression in a full-frame manner to make the computation efficient at test time and, more importantly, to add more global context to the regression process to improve the robustness. To do so, we adopt a fully convolutional encoder-decoder neural network architecture which accepts a whole image as input and produces scene coordinate predictions for all pixels in the image. However, using more global context is prone to overfitting. To alleviate this issue, we propose to use data augmentation to generate more data for training. In addition to the data augmentation in 2D image space, we also augment the data in 3D space. We evaluate our approach on the publicly available 7-Scenes dataset, and experiments show that it has better scene coordinate predictions and achieves state-of-the-art results in localization with improved robustness on the hardest frames (*e.g.*, frames with repeated structures).

I. INTRODUCTION

Image-based localization is an important problem in computer vision and robotics. It addresses the problem of estimating the 6 DoF camera pose in an environment from an image. It is a key component of many computer vision applications such as navigating autonomous vehicles and mobile robotics [7], simultaneous localization and mapping (SLAM) [8], and augmented reality [4, 18].

Currently, there are plenty of image-based localization methods proposed in the literature. Many state-of-the-art approaches [25, 26] are based on local features, such as SIFT, ORB, or SURF [17, 23, 1], and efficient 2D-to-3D matching. Given a 3D scene model, where each 3D point is associated with the image features from which it was triangulated, localizing a new query image against the model is solved by first finding a large set of matches between 2D image features and 3D points in the model via descriptor matching, and then using RANSAC [9] to reject outlier matches and estimate the camera pose on inliers. Although these local feature based methods have been proven to be very accurate and robust in many situations, due to the limitations of the hand-crafted feature detector and descriptor, extremely large viewpoint changes, repetitive structures and textureless scenes

often produce a large number of false matchings which will lead to localization failure.

Recently, some promising neural network based localization approaches have been proposed. Approaches such as PoseNet [13] formulate 6 DoF pose estimation as a regression problem, and thus no traditional feature extraction, feature description, or feature matching processes are required. It has been shown that these approaches overcome the limitations of the local feature based approaches, *i.e.*, they are able to recover the camera pose in very challenging indoor environments where the traditional methods fail. However, their localization accuracy is still far below traditional approaches in other situations where local features perform well. Unlike PoseNet, approaches such as DSAC [3] obtain the 6 DoF pose by a two-stage pipeline: first, regressing a 3D point position for each pixel in an image, and then determining the camera pose using RANSAC based on these correspondences as in the conventional localization pipeline. Results have shown that these methods are the current state-of-the-art on the standard 7-Scenes benchmark [28].

In this paper, we follow the two-stage RGB-only localization pipeline and present a full-frame Coordinate CNN for the scene coordinate regression in the first stage. This full-frame Coordinate CNN is trained to generate dense scene coordinate predictions from a given RGB image. Compared to the patch-based Coordinate CNN adopted in the DSAC pipeline [3], which is trained to generate a single scene coordinate prediction from an image patch centered around a pixel, our network considers more global information and can be evaluated efficiently at test time. Nevertheless, encoding more global context during regression is more prone to overfitting than using local patches since local patches are more invariant to viewpoint changes. Therefore, we propose to use data augmentation to mitigate the problem. We show with experiments on the widely used 7-Scenes dataset that our approach has improved scene coordinate predictions and achieves state-of-the-art results in localization with strong robustness.

The contributions of this paper can be summarized as follows:

- We adopt a fully convolutional encoder-decoder network for the scene coordinate regression in the two-stage image-based localization pipeline to encode more global information.
- We do data augmentation in both 2D and 3D space for the scene coordinate regression, and we show that data




Fig. 1. Overview of our system. In this two-stage pipeline, the encoder-decoder full-frame Coordinate CNN accepts a whole RGB image as input and produces scene coordinate predictions for all pixels in the image, and then the final pose is solved via a RANSAC-based optimization scheme using the predicted correspondences.

augmentation is necessary for training the full-frame Coordinate CNN to achieve good localization performance.

- Our approach achieves state-of-the-art results and shows improved robustness on the hardest frames.

II. RELATED WORK

Solving the image-based localization problem consists of estimating the 6 DoF camera pose of a query image in an arbitrary scene. This is also known as camera relocalization. More specifically, the localization system takes only one image as input, and then outputs camera pose according to a given representation of the scene. The representation of the scene depends on the approach used. It can be a database of images, a reconstructed 3D model, or a trained deep neural network. Here we only consider one image as input, although the image-based localization task can be extended by using a sequence of images as input.

In the earliest stages, image-based localization was solved by treating it as a location recognition problem [34]. In these approaches, image retrieval techniques are often applied to determine the location of the query image by matching it to a database of images and finding the most similar images to it. The locations of the database images are known. Thus, the location of the query image can be estimated according to the known locations. Initially, these methods worked on databases consisting of tens of thousands of images and then they were improved to deal with more than a million images. However, the localization performance provided by these approaches is limited by the accuracy of the known locations of the database images [24] and the density of the key-frames.

To obtain better localization performance, more recent techniques [11] are based on more detailed and structured information, *i.e.*, they use a reconstructed 3D point cloud model, usually obtained from Structure-from-Motion, to represent the scene. This is enabled by some powerful SfM approaches such as Bundler [29]. It is even possible to construct huge models with millions of points [30]. Thus, image-based localization systems are enabled to work on a city-scale level. Instead of matching the query image to the database images, these keypoint based methods solve the localization task by finding

correspondences between the query image and the 3D model. This is achieved by the use of the conventional hand-crafted features, *e.g.*, SIFT [17]. The features of the 3D points are computed during the 3D reconstruction, and the features of the query image are detected and extracted at test time. Therefore, the correspondence search is formulated as a descriptor matching problem. After establishing the 2D-3D correspondences, the camera pose is determined by a Perspective-n-Point [14] solver inside a RANSAC [9] loop.

It is obvious that the pose estimation step can only succeed if the quality of the established correspondences is good enough. When the 3D model is too big in size or too complex, an efficient and effective descriptor matching step is essential. There are several techniques in the literature to handle this problem, such as 2D-to-2D-to-3D matching [11], 3D-to-2D-matching [15], prioritized matching [24] and co-visibility filtering [25]. However, even if the matching process is effective and efficient enough, limitations of the hand-crafted features will still cause these keypoint based localization systems to fail.

Recently, it has been shown that machine learning methods have great potential to tackle the problem of image-based localization. PoseNet [13] demonstrates the feasibility of formulating 6 DoF pose estimation as a regression problem. The pose of a query image is directly regressed by a deep CNN with GoogLeNet [31] architecture pre-trained on large-scale image classification data. However, its performance is still below the state-of-the-art keypoint based methods and far from ideal for practical camera relocalization. In order to improve the accuracy of PoseNet, several variants have been proposed in recent papers. For example, LSTM-Pose [33] makes use of LSTM units [10] on the CNN output to exploit the structured feature correlation. The LSTM units play the role of a structured dimensionality reduction on the feature vector and lead to drastic improvements in localization performance. Another variant, Hourglass-Pose [21], is based on hourglass architecture which consists of a chain of convolution and upconvolution layers followed by a regression part. The upconvolution layers are introduced to preserve the fine-grained information of the input image and this

mechanism has been proven to be able to further improve the accuracy of image-based localization using CNN based architectures. Besides, it has been shown that the use of a novel loss function based on scene reprojection error can also significantly improve PoseNet’s performance [12].

Unlike PoseNet, the scene coordinate regression forests (SCoRF) approach [28] adopts a regression forest [6] to generate 2D-3D matches from an RGB-D input image instead of directly regressing the camera pose. The final camera pose is then determined via a RANSAC-based solver. This whole pipeline is similar to the keypoint based localization approaches, but no traditional feature extraction, feature description, or feature matching processes are required. The original SCoRF pipeline is further improved by exploiting uncertainty in the model [32]. Training the random forest to predict multimodal distributions of scene coordinates results in increased pose accuracy.

In order to localize RGB-only images as well, the original SCoRF is extended by utilizing the increased predictive power of an auto-context random forest [2]. In the most recent DSAC paper [3], a neural network based SCoRF pipeline for RGB-only images is proposed. In contrast to previous SCoRF pipelines, two CNNs are adopted for predicting scene coordinates and for scoring hypotheses respectively. Moreover, the conventional RANSAC is replaced by a new differentiable RANSAC, which enables the whole pipeline to be trained end-to-end. In [19], Massiceti *et al.* further explore efficient versus non-efficient and RF- versus non-RF-derived NN architectures for camera localization and propose a new fully-differentiable robust averaging technique for regression ensembles.

III. METHOD

Our CNN based approach presented in this paper is inspired by the state-of-the-art RGB-only DSAC pipeline [3], and it consists of two stages. In the first stage, a Coordinate CNN is adopted to generate 2D-3D correspondences from a given RGB-only image. In the second stage, a RANSAC-based scheme is performed to determine the final pose estimate. The overall pipeline is illustrated in Figure 1.

The Coordinate CNN used in the original DSAC pipeline [3] is trained and evaluated in a patch-based manner. Similar to the patch-based approaches for semantic segmentation, patches need to be sampled during both training and test time. The patch-based Coordinate CNN is trained using image patches of fixed size (42×42) centered around pixels and the corresponding world coordinates of these pixels. At test time, patches of the same size are used as inputs of the Coordinate CNN to generate 2D-3D correspondences. However, obtaining a set of 2D-3D correspondences in such a way is very inefficient. In the DSAC pipeline [3], 1600 correspondences are needed for pose estimation. Thus 1600 patches should be extracted and fed into the deep neural network. This is obviously time-consuming and also requires much memory. Moreover, the fixed patch size limits the size of the receptive field, and thus limits the information that can be processed by the network. This makes it difficult for the Coordinate CNN

to predict accurate 2D-3D correspondences when the scene exhibits ambiguities (*e.g.*, repeated structures), since the global context cannot be used by the network.

To overcome such limitations, we propose to use an encoder-decoder CNN which accepts a whole image as input and produces scene coordinate predictions for all pixels in the image. This is inspired by recently presented architectures for solving per-pixel tasks, such as semantic segmentation, depth estimation, and disparity estimation. Performing scene coordinate regression in this way, patches are not required to be sampled and the 2D-3D correspondences can be generated efficiently at test time. Since the receptive field of the encoder-decoder architecture covers almost the entire input image, more global context is added to the regression process and more overall information is considered. However, the global image appearance may change a lot due to only small viewpoint changes. That is, it is not as invariant to viewpoint changes as local appearance. Therefore, regressing the scene coordinates using more global information is prone to overfitting. We propose to use data augmentation to mitigate the problem. We call our network full-frame Coordinate CNN.

A. Network Architecture

Our full-frame Coordinate CNN is based on the architecture of DispNet [20]. It is fully convolutional [16] such that dense per-pixel scene coordinate predictions can be generated from arbitrary-sized input images. We call the outputs scene coordinate images. The network consists of a contractive part and an expanding part. The input image is first spatially compressed via the contractive part and then refined via the expanding part. Moreover, shortcut connections are added in between to overcome the data bottleneck. Unlike DispNet [20], there is only one final output layer at the end of the network and no multi-scale side predictions are used. Instead of ReLU [22], we use ELU [5] for the nonlinearity between layers. The details of our network architecture are described in Table I. Layers start with upconv are upconvolutional layers and all other layers are convolutional. Also for each layer, the size of the kernel, the stride, the number of input channels, the number of output channels, and the input (+ is a concatenation) are given.

B. Training Loss

Similar to the patch-based Coordinate CNN used in the original DSAC pipeline, we train our full-frame Coordinate CNN by minimizing the Euclidean distance between the scene coordinate ground truth and the prediction as given in Equation 1. Unlike training in a patch-based manner where patches and single predictions are used and only patches centered at pixels with valid depth values are sampled for training, our network uses pairs of color image and scene coordinate image (dense scene coordinate values) as training samples where the ground truth scene coordinates can be missing for some pixels. Here we simply ignore the pixels without the ground truth scene coordinates and mask out their contributions to the final loss.

TABLE I
Our network architecture.

Name	Kernel	Str.	Ch I/O	Input
conv1a	7×7	2	3/32	image
conv1b	7×7	1	32/32	conv1a
conv2a	5×5	2	32/64	conv1b
conv2b	5×5	1	64/64	conv2a
conv3a	3×3	2	64/128	conv2b
conv3b	3×3	1	128/128	conv3a
conv4a	3×3	2	128/256	conv3b
conv4b	3×3	1	256/256	conv4a
conv5a	3×3	2	256/512	conv4b
conv5b	3×3	1	512/512	conv5a
conv6a	3×3	2	512/512	conv5b
conv6b	3×3	1	512/512	conv6a
conv7a	3×3	2	512/512	conv6b
conv7b	3×3	1	512/512	conv7a
upconv6	3×3	2	512/512	conv7b
iconv6	3×3	1	1024/512	upconv6+conv6b
upconv5	3×3	2	512/512	iconv6
iconv5	3×3	1	1024/512	upconv5+conv5b
upconv4	3×3	2	512/256	iconv5
iconv4	3×3	1	512/256	upconv4+conv4b
upconv3	3×3	2	256/128	iconv4
iconv3	3×3	1	256/128	upconv3+conv3b
upconv2	3×3	2	128/64	iconv3
iconv2	3×3	1	128/64	upconv2+conv2b
upconv1	3×3	2	64/32	iconv2
iconv1	3×3	1	64/32	upconv1+conv1b
upconv0	3×3	2	32/16	iconv1
iconv0	3×3	1	16/16	upconv0
coord_pred	3×3	1	16/3	iconv0

More formally, loss for a training sample can be written as:

$$loss = \sum_{i,j} M_{ij} \|\hat{Y}_{ij} - Y_{ij}\| \quad (1)$$

where \mathbf{Y} and $\hat{\mathbf{Y}}$ are ground truth scene coordinate image and scene coordinate image prediction respectively, \mathbf{M} is a mask and (i, j) is a 2D pixel coordinate. The mask \mathbf{M} has the same resolution as the color image and the scene coordinate image. A pixel of \mathbf{M} is set to 1 if the corresponding scene coordinate ground truth exists and 0 otherwise.

C. Data Augmentation

While our network has larger receptive field and can encode more global context information for better understanding of the scene, it is more prone to overfitting compared to a network trained and tested in a patch-based manner. This is because local patches are more invariant to viewpoint changes and we can usually sample sufficiently enough patches for training. However, in our case, we need more training images taken from a lot of different viewpoints to regularize the network, since global image appearance is more sensitive to viewpoint changes. Unfortunately, a common problem for image-based localization is that the training poses in the whole pose space is not complete enough to regularize such a network. Also, image-based localization is typically performed scene-wise, which means that only the images belonging to a particular scene can be used during training. It does not make sense to use data belonging to other scenes, and unlike for other computer vision tasks such as semantic segmentation, it is not




Fig. 2. Data augmentation. Left: original RGB and scene coordinate images. Middle: 2D transformation. Right: 3D transformation.

straightforward to apply transfer learning technique for image-based localization. Hence, overfitting can be a serious problem in training our full-frame Coordinate CNN for solving image-based localization.

To alleviate the overfitting problem, we propose to use data augmentation to generate more data for training. A simple and common way of doing data augmentation in the context of solving per-pixel tasks is to apply 2D affine geometric transformations to the training images. In our case, we apply 2D transformations to both the input RGB images and the ground truth scene coordinate images. The transformations we perform include translation, rotation, and scaling.

However, these 2D affine transformations do not always generate correct input RGB images which are consistent with 3D scenes (although they are always consistent with their transformed ground truth scene coordinate images). Since the camera is in 3D space and 3D points are projected onto the image plane according to a camera model, if camera pose is changed, the transformation between the old and new 2D coordinates of 3D points is much more complex than the combination of 2D transformations in the image plane. Therefore, we propose to augment data also in 3D space. For each training sample, since we already have the ground truth scene coordinate image, we can generate a local 3D point cloud from that. We then add a small random transformation to the ground truth pose to synthesize a new camera pose. Finally, using the new camera pose, the local 3D point cloud and the known camera model, we project the points to the new camera plane to generate a new RGB image and its ground truth scene coordinate image.

D. Camera Pose Optimization

In the second stage of the two-stage pipeline, a RANSAC algorithm is used in combination with a PnP algorithm to determine the final pose estimation from the 2D-3D correspondences generated in the first stage. Here the RANSAC pose optimization algorithm we use is similar to the one used in the DSAC pipeline [3]. The only differences are that firstly we simply score the hypotheses by counting inliers instead of using a Score CNN [3] to regress the scores of hypotheses from their reprojection error images, and secondly we use the conventional argmax hypothesis selection instead of the

probabilistic selection. We believe the Score CNN used in [3] is prone to overfitting.

The detailed steps of the RANSAC-based pose optimization algorithm used in this paper are described in Algorithm 1. Here, $score(H_i)$ is the number of inliers, *i.e.*, the number of correspondences with reprojection errors less than a threshold τ , and the final pose H_f is the one with highest inlier count. Following [3], 40×40 correspondences are sampled ($N = 1600$), 256 hypotheses are sampled per image ($K = 256$), the inlier threshold τ is set to 10 pixels, 8 refinement steps are performed ($R = 8$), in each refinement step at most 100 inliers are chosen ($P = 100$), and the refinement is stopped in case less than 50 inliers have been found ($Q = 50$).

Algorithm 1 Pseudocode for RANSAC optimization

```

01: evaluate the Coordinate CNN to obtain  $N$  correspondences
02:  $i \leftarrow 0$ 
03: while  $i < K$  do
04:   sample 4 2D-3D correspondences
05:   generate a pose hypothesis  $h$  using PnP
06:   if all the 4 correspondences are inliers of  $h$  then
07:      $H_i \leftarrow h$ 
08:      $S_i \leftarrow score(H_i)$ 
09:      $i \leftarrow i + 1$ 
10: select the final pose  $H_f$  according to the scores
11:  $j \leftarrow 0$ 
12: while  $j < R$  do
13:   sample at most  $P$  inliers
14:   if less than  $Q$  inliers are found then
15:     break
16:   recalculate the pose  $H_f$  from the inliers
17: return pose  $H_f$ 

```

IV. EXPERIMENTS

In this section, we evaluate the performance of our system, and compare it with other state-of-the-art methods. We use the 7-Scenes dataset [28] for benchmarking.

A. Evaluation Dataset

For the experiments, we use the 7-Scenes dataset [28] provided by Microsoft Research. The 7-Scenes dataset is a widely used RGB-D dataset which consists of seven different indoor environments. The RGB-D images are captured using a handheld Kinect camera at 640×480 resolution and associated with 6 DoF ground truth camera poses obtained via the KinectFusion system. A dense 3D model is also available for each scene. Each scene contains multiple sequences of tracked RGB-D camera frames and these sequences are split into training and testing data. The dataset exhibits several challenges, namely motion blur, illumination changes, textureless surfaces, repeated structures (*e.g.*, in the Stairs dataset), reflections (*e.g.*, in the Redkitchen dataset), and sensor noise. Following [3], we use the depth images to generate the ground truth scene

coordinates. At test time, we only use RGB images to estimate poses.

B. Implementation Details

We train our network from scratch for 800 epochs with a batch size of 16 (*i.e.*, 200k updates for Chess) using the Adam optimizer where $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$. The loss is computed as described in Section III-B. The initial learning rate is set to 0.0001 and is halved every 200 epochs until the end.

As mentioned in Section III-C, we perform data augmentation online during network training. We perform the 2D affine transformation with a 40% chance, perform the 3D transformation with a 50% chance, and use the original image with a 10% chance. For the 2D transformation, we uniformly sample translation from the range $[-20\%, 20\%]$ of the image width and height for x and y respectively, sample rotation from $[-45^\circ, 45^\circ]$ and sample scaling from $[0.7, 1.5]$. All pixels without valid RGB values are padded with the same random value. For the 3D transformation, we uniformly sample the rotational axis and the rotational angle is uniformly sampled from $[0^\circ, 60^\circ]$. The direction of the translation vector is again uniformly sampled, and its magnitude in mm is sampled from $[0, 200]$. All pixels without valid RGB values are padded with different random values. Figure 2 shows an example of the data augmentation.

At test time, although our full-frame Coordinate CNN can generate 640×480 scene coordinate predictions, we only sample 40×40 of the predictions for the pose estimation stage to make it consistent with the DSAC pipeline [3].

C. Results

We report both median error of camera pose estimations and accuracy measured as the percentage of query images for which the camera pose error is below 5° and 5cm. The results of our method on the 7-Scenes dataset together with two current state-of-the-art methods are shown in Table II. Following [3], Complete denotes the combined set of frames (17000) of all scenes. We also plot the cumulative histograms of localization errors for a more detailed comparison and report the results of [25] and DSAC* (*i.e.*, DSAC [3] patch-based scene coordinate regression + conventional RANSAC). Note that for the DSAC pipeline [3], only the accuracy is reported in the original paper, and we obtain the additional results using the publicly available source code and trained models provided by the authors. For [25], we use Colmap [27] to construct models from scratch and register the models against the ground truth poses of the training images, and again we use the publicly available implementation of [25].

As we can see, the overall performance of DSAC* is better than the original DSAC [3]. While achieving almost the same results on the easy frames, DSAC* has superior performance on the harder ones, *i.e.*, DSAC* provides better 0.95 quantiles, especially for scenes that have fewer training images (Fire, Heads). More importantly, DSAC* is able to produce reasonable median localization error for Stairs, though

TABLE II

Results of our approach on the 7-Scenes dataset. Our approach outperforms two state-of-the-art methods and has significantly improved performance on Stairs. Training the full-frame Coordinate CNN without our proposed data augmentation dramatically degenerates the localization performance.

Scene	5°, 5cm				Median Error			
	Brachmann <i>et al.</i> [2]	Brachmann <i>et al.</i> [3]	Ours noaug	Ours	Brachmann <i>et al.</i> [2]	Brachmann <i>et al.</i> [3]	Ours noaug	Ours
Chess	94.9%	94.6%	43.2%	88.5%	1.5cm, 1.3°	2.1cm, 0.7°	5.9cm, 1.9°	2.4cm, 0.8°
Fire	73.5%	74.3%	10.7%	62.3%	3.0cm, 1.4°	2.8cm, 1.0°	13.2cm, 4.5°	3.7cm, 1.4°
Heads	48.1%	71.7%	14.1%	75.1%	5.9cm, 3.4°	2.0cm, 1.3°	12.6cm, 9.4°	2.4cm, 1.7°
Office	53.2%	71.2%	30.6%	73.1%	4.7cm, 1.7°	3.4cm, 1.0°	7.0cm, 2.0°	3.5cm, 1.0°
Pumpkin	54.5%	53.6%	18.5%	51.4%	4.3cm, 2.1°	4.7cm, 1.3°	10.5cm, 2.7°	4.9cm, 1.3°
Red Kitchen	42.2%	51.2%	38.8%	60.4%	5.8cm, 2.2°	5.0cm, 1.5°	6.1cm, 1.7°	4.2cm, 1.2°
Stairs	20.1%	4.5%	0.2%	29.5%	17.4cm, 7.0°	1.9m, 49.4°	43.4cm, 10.1°	7.9cm, 2.1°
Average	55.2%	60.1%	22.3%	62.9%	6.1cm, 2.7°	30cm, 8.0°	14.1cm, 4.6°	4.1cm, 1.4°
Complete	55.2%	62.5%	28.0%	64.9%	-	3.9cm, 1.6°	8.0cm, 2.3°	3.8cm, 1.2°

the accuracy on the hardest frames of Stairs is still terrible. The results of DSAC* suggest that the use of the Score CNN [3] could make the entire localization pipeline less robust.

According to the results, our proposed method outperforms the two state-of-the-art methods [2, 3] and the traditional keypoint based method [25], and its overall localization performance is more robust. To be specific, our method significantly improves the performance on the hardest frames. We find that while DSAC [3] (as well as DSAC*) and [25] often have extreme localization errors, *i.e.*, rotational errors close to 180° and translational errors larger than 2-5m, the maximum errors of our methods are always reasonable. As shown in Figure 3, the *Ours* curves always reach 1 very quickly compared to others. The results suggest that the global information could help the system to perform better in challenging scenarios. Remarkably, our method is able to provide significantly better localization performance for Stairs. Even the hardest frames can be localized with the errors less than 0.41m and 13°. It shows that the full-frame Coordinate CNN with enlarged receptive field can better cope with the repetitive structures, while the patch-based Coordinate CNN [3] and the local keypoint based method [25] are limited due to their local nature. Our method is also on-par with [19]. However, in [19], only the overall accuracy is reported and this prevents us from carrying out a more detailed comparison.

To show that our full-frame Coordinate CNN is more efficient at test time, we present the runtimes of the CNNs in Table III. We see that our full-frame Coordinate CNN is one order of magnitude faster than the patch-based one [3]. And this comparison is even not fair since our full-frame Coordinate CNN produces 640×480 predictions while only 40×40 are generated by the patch-based one [3]. In addition, our full-frame Coordinate CNN has a relatively small model size compared to the regression forest based model [19]. Our network has a model size of ~ 120 MB, while the most compact one from [19] has a model size of ~ 250 MB.

Following [19], in addition to the direct measure of localization performance, we present the accuracy of the intermediate scene coordinate predictions on the test images. In Table IV, we report the percentage of scene coordinate inliers and the mean Euclidean distance between the inliers and their ground

TABLE III

The runtimes of the full-frame and patch-based Coordinate CNNs.

GPU	Full-frame	Patch-based
NVIDIA GeForce GT 750M	~ 0.3 s	~ 5 s
NVIDIA GeForce GTX 1080	~ 0.02 s	~ 0.3 s

TABLE IV

The percentages of the scene coordinate prediction inliers and the mean errors of the inliers. Our predictions are clearly the most accurate.

DSAC [3]	Best from [19]	Ours
59.4%, 41mm	45.3%, 47mm	90.3%, 28mm

truth scene coordinate labels. A prediction is considered as an inlier if its Euclidean distance to its ground truth label is less than 10mm [19]. The normalized histograms of scene coordinate errors are illustrated in Figure 4. As we can see, our full-frame Coordinate CNN is able to produce significantly better scene coordinate predictions. This shows why it is more robust than the patch-based one. However, this does not directly lead to equally better localization accuracy. This is because the RANSAC-based optimizer is highly robust and non-deterministic [19].

To evaluate the effectiveness of data augmentation, we also report the performance of our method without the proposed data augmentation. According to the results, training the full-frame Coordinate CNN without data augmentation dramatically degenerates the localization performance on all the scenes. Without data augmentation, the full-frame Coordinate CNN is unable to generalize well to unseen images. This proves the importance of data augmentation which is used for regularizing our full-frame Coordinate CNN during training.

V. CONCLUSION

In this paper, we present a image-based localization approach using fully convolutional encoder-decoder network to predict scene coordinates of RGB image pixels in a full-frame manner. The proposed full-frame Coordinate CNN can process more global information to better understand the scene. It takes a whole image as input and produces scene coordinate predictions for all pixels in the image and can be efficiently

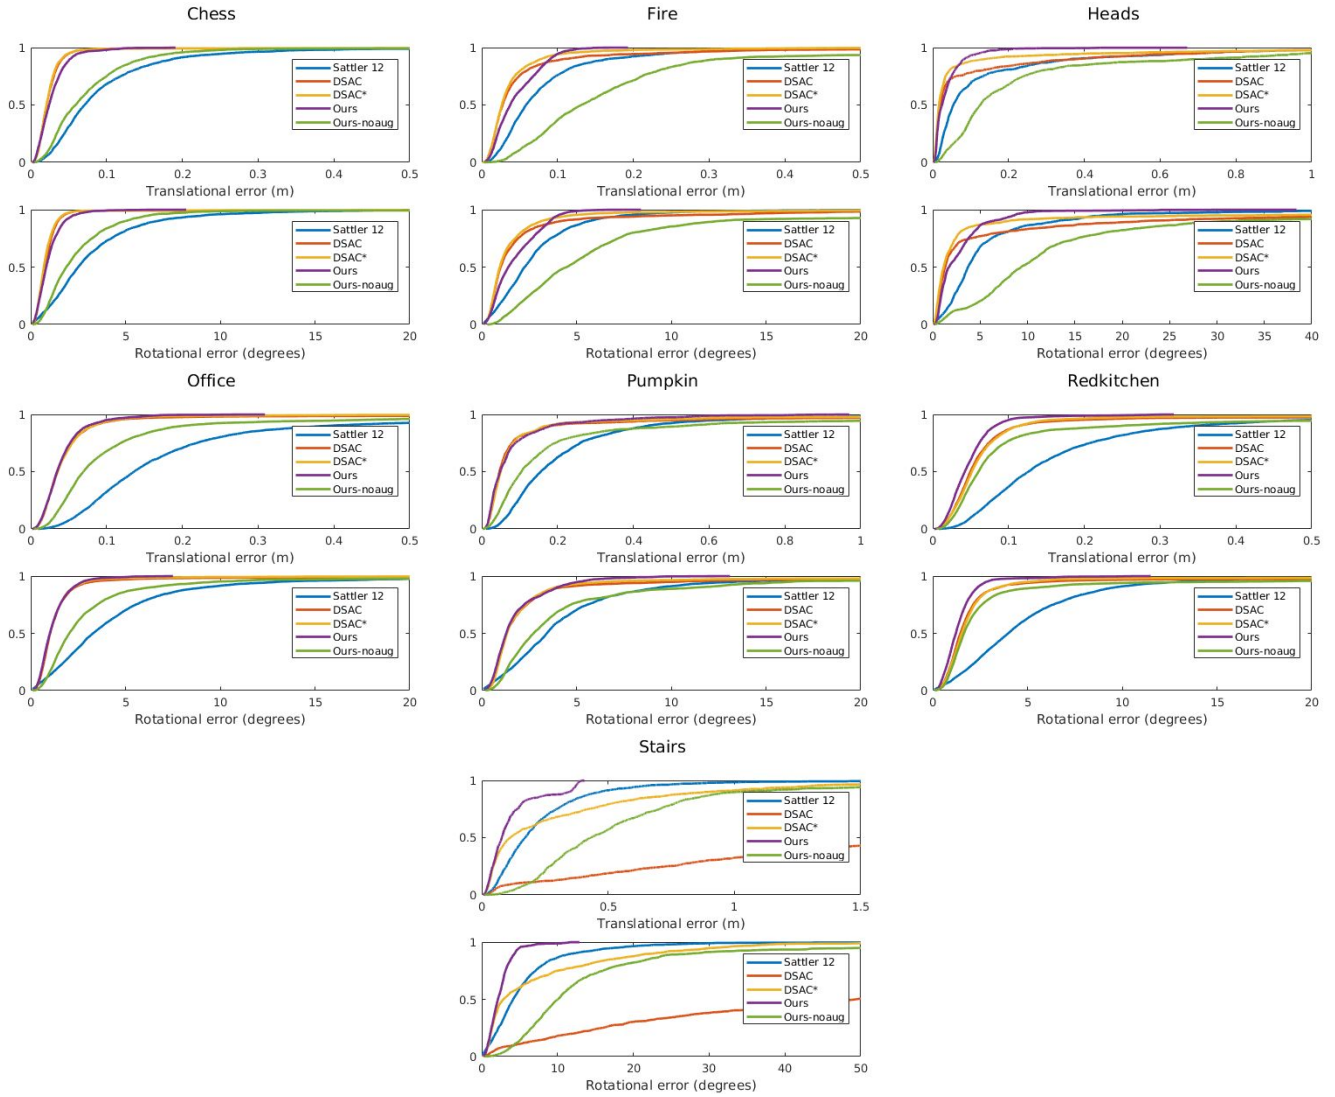


Fig. 3. Localization performance presented by cumulative histograms (normalized) of errors. Compared to other methods, our method is more robust, *i.e.*, it has better performance on the hardest frames, since the curves always reach 1 very quickly.

evaluated at test time. The proposed data augmentation is crucial for our approach to achieve good localization performance. The results show that our method provides more accurate scene coordinate predictions, achieves state-of-the-art localization performance and has improved robustness in challenging scenarios.

REFERENCES

- [1] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *ECCV*, 2006.
- [2] Eric Brachmann, Frank Michel, Alexander Krull, Michael Ying Yang, Stefan Gumhold, and Carsten Rother. Uncertainty-driven 6d pose estimation of objects and scenes from a single RGB image. In *CVPR*, 2016.
- [3] Eric Brachmann, Alexander Krull, Sebastian Nowozin, Jamie Shotton, Frank Michel, Stefan Gumhold, and Carsten Rother. Dsac - differentiable ransac for camera localization. In *CVPR*, 2017.
- [4] Robert Oliver Castle, Georg Klein, and David W. Murray. Video-rate localization in multiple maps for wearable augmented reality. In *ISWC*, 2008.
- [5] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). In *ICLR*, 2016.
- [6] Antonio Criminisi and Jamie Shotton. *Decision Forests for Computer Vision and Medical Image Analysis*. Springer Publishing Company, Incorporated, 2013. ISBN 1447149289, 9781447149286.
- [7] Mark Cummins and Paul Newman. Fab-map: Probabilistic localization and mapping in the space of appearance. *The International Journal of Robotics Research*, 27(6): 647–665, 2008.
- [8] Ethan Eade and Tom Drummond. Scalable monocular




Fig. 4. Histograms (normalized) of scene coordinate errors. Our method provides better scene coordinate predictions.

- slam. In *CVPR*, 2006.
- [9] 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*, 24(6):381–395, 1981.
- [10] Sepp Hochreiter and Juergen Schmidhuber. Long Short-Term Memory. 9(8):1735–1780, 1997.
- [11] Arnold Irschara, Christopher Zach, Jan-Michael Frahm, and Horst Bischof. From structure-from-motion point clouds to fast location recognition. In *CVPR*, 2009.
- [12] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. In *CVPR*, 2017.
- [13] Alex Kendall, Matthew Grimes, and Roberto Cipolla. PoseNet: A convolutional network for real-time 6-dof camera relocalization. In *ICCV*, 2015.
- [14] Laurent Kneip, Davide Scaramuzza, and Roland Siegwart. A novel parametrization of the perspective-three-point problem for a direct computation of absolute camera position and orientation. In *CVPR*, 2011.
- [15] Yunpeng Li, Noah Snavely, and Daniel P. Huttenlocher. Location recognition using prioritized feature matching. In *ECCV*, 2010.
- [16] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.
- [17] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.
- [18] 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 *RSS*, 2015.
- [19] Daniela Massiceti, Alexander Krull, Eric Brachmann, Carsten Rother, and Philip H. S. Torr. Random forests versus neural networks - what’s best for camera relocalization? In *ICRA*, 2017.
- [20] Nikolaus Mayer, Eddy Ilg, Philip Häusser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *CVPR*, 2016.
- [21] Iaroslav Melekhov, Juha Ylioinas, Juho Kannala, and Esa Rahtu. Image-based localization using hourglass networks. In *ICCVW*, 2017.
- [22] Vinod Nair and Geoffrey E. Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, 2010.
- [23] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary R. Bradski. ORB: an efficient alternative to SIFT or SURF. In *ICCV*, 2011.
- [24] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Fast image-based localization using direct 2d-to-3d matching.

In *ICCV*, 2011.

- [25] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Improving image-based localization by active correspondence search. In *ECCV*, 2012.
- [26] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Efficient effective prioritized matching for large-scale image-based localization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(9):1744–1756, 2017.
- [27] Johannes L. Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, 2016.
- [28] Jamie Shotton, Ben Glocker, Christopher Zach, Shahram Izadi, Antonio Criminisi, and Andrew W. Fitzgibbon. Scene coordinate regression forests for camera relocalization in rgb-d images. In *ICCV*, 2013.
- [29] Noah Snavely, Steven M. Seitz, and Richard Szeliski. Modeling the world from internet photo collections. *International Journal of Computer Vision*, 80(2):189–210, 2008.
- [30] Christoph Strecha, Timo Pylvninen, and Pascal Fua. Dynamic and scalable large scale image reconstruction. In *CVPR*, 2010.
- [31] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.
- [32] Julien P. C. Valentin, Matthias Nießner, Jamie Shotton, Andrew W. Fitzgibbon, Shahram Izadi, and Philip H. S. Torr. Exploiting uncertainty in regression forests for accurate camera relocalization. In *CVPR*, 2015.
- [33] Florian Walch, Caner Hazirbas, Laura Leal-Taixe, Torsten Sattler, Sebastian Hilsenbeck, and Daniel Cremers. Image-based localization using lstms for structured feature correlation. In *ICCV*, 2017.
- [34] Wei Zhang and Jana Kosecka. Image based localization in urban environments. In *3DPVT*, 2006.