

Scene Coordinate and Correspondence Learning for Image-Based Localization

Mai Bui¹

mai.bui@tum.de

Shadi Albarqouni¹

shadi.albarqouni@tum.de

Slobodan Ilic^{1,2}

slobodan.ilic@siemens.com

Nassir Navab^{1,3}

navab@cs.tum.edu

¹ Chair for Computer Aided Medical

Procedures

Technische Universität München

Munich, Germany

² Siemens AG,

Munich, Germany

³ John Hopkins University,

Baltimore, MD, USA

Abstract

Scene coordinate regression has become an essential part of current camera relocalization methods. Different versions, such as regression forests and deep learning methods, have been successfully applied to estimate the corresponding camera pose given a single input image. In this work, we propose to regress the scene coordinates pixel-wise for a given RGB image by using deep learning. Compared to the recent methods, which usually employ RANSAC to obtain a robust pose estimate from the established point correspondences, we propose to regress confidences of these correspondences, which allows us to immediately discard erroneous predictions and improve the initial pose estimates. Finally, the resulting confidences can be used to score initial pose hypothesis and aid in pose refinement, offering a generalized solution to solve this task.

1 Introduction

Camera re-localization from a single input image is an important topic for many computer vision applications such as SLAM [15], augmented reality and navigation [12]. Due to rapid camera motion or occlusions, tracking can be lost, making re-localization methods an essential component of such applications. Early methods focus on keyframe [7] or descriptor matching by using SIFT or ORB [20] features to obtain point correspondences from which the camera pose could be inferred. However, those methods usually do not perform well under occlusion or in poorly textured environments. On the other side, machine learning methods have recently shown great capabilities in estimating camera poses from a single image. In particular, regression forests have been employed for a robust pixel-to-scene coordinate prediction. Correspondence samples are then used to obtain multiple pose hypotheses, and a robust pose estimate is found using RANSAC [5, 8, 23, 25].

Recently, deep learning methods have emerged, mainly focusing on RGB images as input and directly regressing the camera pose, thus offering fast camera pose estimates [10, 11]. Most of these methods, however, cannot achieve an accuracy similar to the scene coordinate

regression approaches. This leads to the assumption that the intermediate step of regressing scene coordinates plays a crucial role in estimating the camera pose for deep learning algorithms and the generalization of those methods. Moreover, RANSAC [6] usually plays a vital role in achieving good accuracy in any of the methods focusing on camera relocalization via scene coordinate regression.

In this paper, we introduce a method, which as a first step, densely regresses pixel-wise scene coordinates given an input RGB image using deep learning. In addition, we propose a new form of regularization, that smoothes the regressed coordinates and which can be applied to further improve the regressed coordinates. Thus, a detailed analysis of scene coordinate regression and the influence of different loss functions on the quality of the regressed scene coordinates is conducted. Our main contribution is seen in the second step, where the confidences of the obtained image to scene coordinate correspondences are regressed, based on which erroneous predictions can immediately be discarded, which in turn results in a more robust initial pose hypothesis. In addition, the resulting confidence predictions can be used to optimize the estimated camera poses in a refinement step similar to previous works [4]. In contrast to these methods, our approach offers a more general solution by not restricting itself in terms of the optimization function and thresholding, which are typically used to define the inliers and outliers in RANSAC optimizations.

2 Related Work

There exists a vast amount of research focusing on the topic of camera pose estimation. The most related to our work can be divided into two categories: correspondence learning and direct regression. The first focuses on descriptor matching to obtain point correspondences from which the camera pose can be inferred, by using either analytical solutions or learning methods. The second performs direct pose regression by using deep learning methods to obtain a pose estimate from a single input image.

Correspondence Learning. Well-known methods working on the topic of camera relocalization have used random forests for correspondence prediction. Here, the forest is trained to predict pixel to 3D coordinate correspondences, from which the camera pose can then be inferred and iteratively refined in a pre-emptive RANSAC optimization [23]. Several extensions and improvements of the method have been proposed, increasing its performance and robustness [5, 8, 25].

On the other hand, due to the recent success of this method, various related methods have used deep learning approaches. Inspired by the approach of [23], Brachmann et al. [4] use two convolutional neural networks (CNNs) to predict the pose for an RGB image; the first CNN is used to predict point correspondences and is linked to a second CNN by a differentiable version of RANSAC, they name DSAC. Notably, reinforcement learning is used to obtain a probabilistic selection to enable end-to-end learning of the framework.

Recently, Schmidt et al. [21] automatically find correspondences in RGB-D images by relying on learned feature representations for correspondence estimation using a 3D model of the scene. A fully-convolutional neural network is trained on a contrastive loss to produce pixel-wise descriptors. Here, pixels corresponding to the same model coordinate are mapped together whereas the remaining pixels have dissimilar feature descriptors. Despite the lack of complete guarantee that the descriptors learned from one video can be mapped to the features of another video capturing the same scene, the method showed robustness and generalization.

Instead of relying on feature representations for point correspondences, Zamir et al. [30] investigate a generic feature representation based on the viewpoint itself, which could be used to retrieve a pose estimate. A siamese network is trained on a multi-task loss including the pose and a matching function. Image patches are extracted and matched according to their poses, and the network is trained to match patches with similar viewpoints. Additionally, it is shown that the resulting models and features generalize well to novel tasks, such as, scene layout or surface normal estimation.

Capturing both geometric and semantic information, Schöneberger et al. [22] propose to learn 3D descriptors and include the task of 3D semantic scene completion into a variational encoder-decoder network. By incorporating semantic information, robust descriptors are learned and shown to successfully aid in visual localization tasks even under strong viewpoint and illumination changes. Such additional information has been shown to work well in similar scenarios, such as in localization, when a semantic map of the environment is given as prior information [14, 27].

Direct Regression. Lately, direct regression approaches have been emerging, which use deep learning methods to regress camera poses from a single image. Mostly CNNs are used in this context to estimate the camera poses [9, 10, 11, 26]. Therefore, Kendall et al. [11] use a CNN, called PoseNet, to directly predict the six degrees of freedom for camera pose estimation from RGB images. They parameterize rotation by quaternions, which leads to a total of seven parameters to regress for rotation and translation. Although Kendall et al. [11] achieve reasonable performance for several indoor scenes, this method still shows a significant inaccuracy as compared with the random forest approaches. Therefore, we assume that predicting an intermediate representation such as point correspondences is important for inferring the final pose. However, this method relies only on RGB images, without the need for depth information, which makes it easily applicable in indoor and outdoor settings.

Walch et al. [26] extend this approach and connect the last feature vector of the neural network to four LSTM units before concatenating the results and feeding this feature vector to the final regression layer. As in [11], a pre-trained CNN used for classification is adapted and fine-tuned to enable the regression of the camera pose. By connecting LSTMs and thus correlating the spatial information, the receptive field of each pixel is enlarged substantially, improving the final pose estimation.

In [10], the authors of [11] extend their method by introducing novel loss functions, which further reduces the gap in accuracy compared to the state-of-the-art methods. Further, they show that the re-projection error could be used to additionally fine-tune the model and optimize the regression prediction when the depth information is given.

As a first direct regression approach that achieves comparable accuracy with regard to the scene coordinate regression methods, [24] proposes a multi-task learning framework. By combining the global and relative pose regression between the image pairs, the authors present a framework for localization and odometry, which shows great improvements in accuracy.

In comparison, our approach is most related to scene coordinate regression methods using deep learning [4]. Scene coordinates are densely regressed as opposed to patch-based regression proposed in the state-of-the-art methods. Moreover, correspondence confidences are predicted to remove outliers and boost the accuracy of the initial pose hypothesis. Additionally, the resulting confidences can be directly used for hypothesis scoring and pose refinement.




Figure 1: Outline of our re-localization framework. Scene coordinates are densely regressed for each pixel in the input RGB image. Confidences are predicted for point correspondences and finally used to compute the camera pose estimates.

3 Methodology

Given an input RGB image $\mathbf{I} \in \mathbb{R}^{h \times w \times 3}$ of a scene, where h and w are the image height and width, respectively, our goal is to estimate the corresponding camera pose, given by its orientation $R \in \mathbb{R}^{3 \times 3}$ and position $\mathbf{t} \in \mathbb{R}^3$. The camera pose describes the mapping between the camera and the scene coordinates \mathbf{x} and $\mathbf{X} \in \mathbb{R}^3$ as

$$\mathbf{X} = R\mathbf{x} + \mathbf{t}. \quad (1)$$

The relation between the 3D camera coordinates $\mathbf{x} \in \mathbb{R}^3$ and the image pixels $\mathbf{p}_x \in \mathbb{R}^2$ depends on the camera’s focal length f_x, f_y and the optical center c_x, c_y , and is defined as

$$\mathbf{p}_x = \left(\frac{f_x x}{d} + c_x, \frac{f_y y}{d} + c_y \right)^T, \quad (2)$$

with $\mathbf{x} = (x, y, d)^T$ being a point in the camera coordinate frame, given by its coordinates x, y and its depth value d . In case the camera pose is unknown, it can be retrieved given N_k number of correspondences either, if depth information is available, using the Kabsch algorithm from the 3D-3D correspondences between \mathbf{x} and \mathbf{X} or using the PnP algorithm [13] from the 2D-3D correspondences between the image points \mathbf{p}_x and \mathbf{X} .

For this aim, our proposed framework consists of three steps: (1) scene coordinate regression, in which we densely predict scene coordinates, and in this context, add a novel regularization, (2) confidence prediction, in which we aim to find accurate correspondences in our coordinate predictions, and (3) pose estimation, in which we employ the aforementioned algorithms to compute the camera pose estimate for the most confident predictions N_{best} . An overview of our framework is given in Figure 1.

3.1 Scene Coordinate Regression

Coordinate regression. As the first step, our aim is to model the function $\phi: \mathbf{I} \rightarrow \hat{\mathbf{X}}$, obtaining the predicted scene coordinates $\hat{\mathbf{X}}$, as $\hat{\mathbf{X}} = \phi(\mathbf{I}; \omega)$, where ω are the model parameters. Therefore, we compute the ground truth scene coordinates \mathbf{X} to train a model and regress the scene coordinates of every pixel, obtaining an $\mathbb{R}^{h \times w \times 3}$ output map. For this purpose, we use the Tukey’s Biweight loss function [2] to regress the 3D coordinates, given as

$$\mathcal{L}_{coords} = \frac{1}{h \cdot w} \sum_{i=1}^{h \cdot w} \sum_{s=1}^S \rho(r_{i,s}), \quad \text{with } \rho(r_{i,s}) = \begin{cases} \frac{c^2}{6} [1 - (1 - (\frac{r_{i,s}}{c})^2)^3], & \text{if } |r_{i,s}| \leq c \\ \frac{c^2}{6}, & \text{otherwise} \end{cases} \quad (3)$$

where $r_{i,s} = X_{i,s} - \hat{X}_{i,s}$ is the residual, $S = 3$ is the number of coordinates to regress and $\rho(\cdot)$ is Tukey’s Biweight function. In Tukey’s Biweight function the choice of the tuning constant c plays a crucial role, which is proposed to be chosen according to the median absolute deviation over the residuals assuming a Gaussian distribution. Nevertheless, we propose to choose the parameter c depending on the spatial extent of the current scene, where, after empirical evaluation, we found half of the scenes diameter given in meters to provide better results. In case of missing depth values and thus missing ground truth scene coordinates, we omit these pixels during training in order not to negatively influence the network.

Coordinate smoothing. The graph Laplacian regularization, has been successfully applied for image denoising on image patches [17]; therefore, we consider the scene coordinates in a given neighborhood. Minimizing the graph Laplacian regularizer enables us to smooth the image patches with respect to the given graph. Similarly, we consider the scene coordinates as vertices and compute weights according to the depth value at the corresponding pixel

$$w_{ij} = \frac{e^{-|d_i - d_j|}}{\sum_{k \in K, k \neq i} e^{-|d_i - d_k|}}, \quad (4)$$

where, given a pixel position i , we compute weights w_{ij} for each index j in a given neighborhood K . In this case, d_i represent the depth value at index i . Finally, we obtain the additional smoothing term in our loss function

$$\mathcal{L}_{smooth} = \sum_i^{N_k} \sum_{j \in K} w_{ij} \cdot \|\hat{\mathbf{X}}_i - \hat{\mathbf{X}}_j\|_2, \quad (5)$$

where $\hat{\mathbf{X}}$ corresponds to the predicted scene coordinates at a given pixel index. This term loosely pushes the surrounding points closer together, given the fact that their depth values are similar; otherwise, a larger difference between points is accepted.

Overall, we train the model using our loss function, described as $\mathcal{L} = \mathcal{L}_{coords} + \mathcal{L}_{smooth}$.

3.2 Confidence Prediction

The regressed coordinates can be used to obtain a pose estimate. However, these correspondences usually include a large amount of erroneous correspondences, which is most often solved using RANSAC.

Inspired by [29], which classifies point correspondences between image pairs, we train a neural network to estimate the probabilities of 2D-3D correspondence. Instead of solving this as a classification problem, we consider this task as a regression problem. Therefore, we use the model described in the previous section to create probabilities for each scene coordinate prediction and construct the training set $S_{confidence} = \{(\mathbf{p}_{\hat{\mathbf{X}}_1}, \hat{\mathbf{X}}_1, \delta_1), \dots, (\mathbf{p}_{\hat{\mathbf{X}}_{N_k}}, \hat{\mathbf{X}}_{N_k}, \delta_{N_k})\}$, where $\delta_i = e^{-(s \cdot \|\mathbf{X}_i - \hat{\mathbf{X}}_i\|_2)}$. Here, s is used as a scale, so that accurate coordinates are given a high probability. In this step, the objective is, therefore, to compute the function $\Phi: (\mathbf{p}_{\hat{\mathbf{X}}}, \hat{\mathbf{X}}) \rightarrow \delta$, described by the model parameters Ω , so that $\delta = \Phi((\mathbf{p}_{\hat{\mathbf{X}}}, \hat{\mathbf{X}}); \Omega)$. To this end, we feed N_k points containing the image pixel and the predicted scene coordinates to our model. As an output, we obtain a probability for each point according to whether it is likely to be a good

correspondence or not. As a loss function, we use the l_2 loss to train this model,

$$\mathcal{L}_{confidence} = \sum_i^{N_k} \|\delta_i - \hat{\delta}_i\|_2. \quad (6)$$

The pose predictions can then easily be obtained by sampling the most confident point correspondences, while removing the initial erroneous predictions right away.

3.3 Pose Estimation and Refinement

The initial pose hypothesis is refined as a post-processing step. Following previous works [4], h_p pose hypotheses are sampled using N_k number of point correspondences for each. Out of these correspondences, only the 10% points with highest confidence are kept. However, even though we are able to greatly improve the quality of the point correspondences used to compute pose estimates with this step, the initial randomly selected points might still be highly inaccurate, leading to erroneous predictions included in the confidence sampled subset. To overcome this problem, only one hypothesis out of the h_p is chosen, by scoring each hypothesis using the mean confidence over the probabilities of the correspondences used to compute the pose estimate. Then, the best hypothesis is refined by repeatedly sampling N_k randomly selected points and re-running the PnP or Kabsch algorithm, including the additional 10% most confident correspondences in this point set.

3.4 Implementation Details

For scene coordinate regression, U-Net is used as the network architecture, as it could easily be used to regress correspondence maps and has been shown to train well on few input images [19]. Four convolutional layers with pooling layers and dropout applied after each, and four up-sampling layers are used, which gives a final feature map of size $h \times w \times 8$. By applying a last convolutional layer, we obtain the final correspondence map of size $h \times w \times 3$.

For confidence prediction, we adapt PointNet [18] to our specific input and output requirements. As our aim is to predict confidences for a subset of point correspondences, we mimic the input for solving the PnP and keeping our final objective in mind feed randomly selected points to PointNet, which has been shown to work well on unstructured data such as point clouds. This way, the network learns to be independent of the order in which the points are fed and could still be used in scenarios where a dense representation of point correspondences is not given. For the final regression layer, following [29], we first apply a hyperbolic tangent followed by a ReLU activation, so the network can easily predict low confidence for highly inaccurate points. Mostly for evaluation purposes, s is computed, so that the inlier points have a lower bound probability of 0.75, resulting in $s = 2.8768$.

We set $N_k = 500$ out of which only the most confident points are used to estimate a camera pose, resulting in only $N_{best} = 50$ points. For pose refinement, we follow the parameter settings of [4] and sample $h_p = 256$ initial hypotheses and refine the best one for eight iterations on the additional most confident points out of the randomly sampled points.

Both networks are trained separately for 800 epochs with batch size of 20 using RMSprop Optimizer with an initial learning rate of $5 \cdot 10^{-4}$. All experiments were conducted on a Linux-based system with 64-GB RAM and 8-GB NVIDIA GeForce GTX 1080 graphics card and implemented using TensorFlow [1].




Figure 2: Normalized histograms of regressed scene coordinate errors used to compute pose estimates on the 7-Scenes. The points are either randomly sampled or, the most confident points out of the samples are used.

4 Experiments

To evaluate our method, we define baseline models, which are described in Section 4.2, and use the following metrics. For this purpose, inliers and outliers are defined as $\|\mathbf{X} - \hat{\mathbf{X}}\|_2 < t_{inlier}$, with $t_{inlier} = 0.1 m$ being a common threshold chosen to define inliers [4, 23, 25]. Every inlier point is therefore counted as a true positive in our evaluations. For pose estimation we compute the median rotation, the translation error and the pose accuracy, where a pose is considered correct if the rotation and translation error are below 5° and $5 cm$, respectively.

4.1 Dataset

Our method is evaluated on the publicly available 7-Scenes dataset from Microsoft, which consists of seven indoor scenes with varying volumes ranging from 1 to $18 m^3$. RGB and depth images of each scene and their corresponding ground truth camera poses are available. For each scene images in the range of 1K to 7K are provided, including very difficult frames due to motion blur, reflecting surfaces and repeating structures, for example, in case of the *Stairs* scene. The images were captured using a Kinect RGB-D sensor and the ground truth poses obtained using KinectFusion. Following the state-of-the-art, we use the training and test sets specified for this dataset. No augmentation as proposed in [16, 28] was performed and our models were trained individually for each scene.

4.2 Baseline Models

First, we evaluate each individual component of our model and create the baseline models for comparison. To start with, the first step of our pipeline, the scene coordinate regression is evaluated. To this aim, a model is trained on scene coordinate regression by using different loss functions. Mainly we compare between l_1 and the L_{coords} loss as described in section 3.1. In this case, a pose estimate is computed by randomly sampling N_k number of points. Results are given for a single pose estimate. Further, we abbreviate these model as P_{l_1} and




Figure 3: a) ROC comparison between regression and classification for the training images of the *Heads* and *Stairs* scenes, and example images showing b) input RGB, c) corresponding pixel-wise scene coordinate error and d) confidence map predicted by our model.

P_{lukey} . Next, to evaluate the scene coordinate regression quality, regularization is added in the form of the introduced smoothness term L_{smooth} , corresponding to model P_{smooth} .

Further, and more importantly, the second step of our pipeline, the confidence prediction is analyzed. The predicted scene coordinates and associated image points are used to train a model and regress the confidence of each correspondence. For pose estimation, in this case, only the most confident correspondences out of the initial randomly sampled N_k points are kept. We abbreviate this model as $P_{\text{confidence}}$.

4.3 Evaluation of Baseline Models

Each of our models is evaluated, comparing the error between the regressed and ground truth scene coordinates, and the pose error. We compare our models trained on the scene coordinate regression and additionally regularizing the model using the smoothness term. Here, we found a slight improvement in terms of the regressed coordinates as well the pose estimates comparing our models with and without additional regularization.

Using our proposed confidence prediction, Figure 2 shows that the point errors of the points used to compute a pose estimate significantly decrease, successfully eliminating most of the erroneous predictions and greatly boosting the pose estimation accuracy. As a result, the estimated poses also improve significantly, as seen in Table 1. Specifically, the translation error greatly decreases. It should be noted, that only the most confident point correspondences are used to compute a pose estimate for this model. As a result, more accurate poses are obtained using a much smaller number of points. Additionally, the percentage of inliers in the sampled points used to compute initial pose hypothesis significantly increases. For the evaluation of our proposed confidence prediction to RANSAC, we sampled h_p pose hypotheses and keep the pose hypothesis with the highest inlier score as a final result. Due to the very low amount of inliers, it is difficult to apply RANSAC and obtain satisfactory results without pose refinement or additional confidence prediction. Furthermore, we analyze the quality of our model’s coordinate regression and confidence prediction, where example error and probability maps are shown in Figure 3. In our case, since we densely regress the scene coordinates, high error values usually correspond to missing depth values and therefore missing ground truth coordinates in the image. Although it seems difficult for the network to accurately predict low confidences in regions with unusually large error, in regions of inlier predictions, the model is able to predict corresponding high confidence.

Table 1: Median rotation and translation error on the *Heads* scene for our baseline models. Results are computed using only one pose hypothesis without any further refinement.

| Model | h_p | P_1 | P_{Tukey} | P_{smooth} | $P_{confidence}$ |
|---------|-------|--------------|--------------|--------------|------------------|
| Kabsch | 1 | 18.0°, 0.29m | 8.77°, 0.15m | 6.67°, 0.12m | 5.77°, 0.07m |
| | 256 | 14.7°, 0.25m | 6.23°, 0.11m | 5.88°, 0.09m | 4.86°, 0.06m |
| PnP | 1 | 50.3°, 0.44m | 41.3°, 0.44m | 44.3°, 0.43m | 10.6°, 0.18m |
| | 256 | 33.7°, 0.46m | 25.9°, 0.43m | 26.3°, 0.37m | 5.09°, 0.10m |
| Inliers | 1 | 5.1% | 7.8% | 11.9% | 50.7% |

4.4 Evaluation of Confidence Prediction

To assess the quality of regressing correspondence probabilities, a model is trained on simple classification, where a point correspondence could be labeled either as an inlier or an outlier depending on the threshold t_{inlier} . The model is trained using cross-entropy loss. As a second step, we train our proposed model, regressing probabilities in the range of $\delta \in [0, 1]$ instead and plot the resulting ROC curves as shown in Figure 3a. As a result, we assess that the performance of regression in this case is more or less equal to classification. However, we do not restrict the model to a specific threshold chosen for the inlier definition, which needs to be adapted for each scene depending on the quality of scene coordinate regression. In comparison, a classification model trained on the challenging *Stairs* scene results in a drop of relative rotation and translation error of 4.4% and 31.5%, respectively, because very few inliers were available during training.

4.5 Comparison to the state-of-the art

Finally, we report the results of our framework using a combination of scene coordinate regression and confidence prediction, described as $P_{confidence}$. We compare our results to the current state-of-the art methods, namely PoseNet [10], which directly regresses the camera poses from the RGB input images and refines the trained models by optimizing on the re-projection error. The median rotation and translation errors evaluated on the 7-Scenes dataset can be found in Table 2, where we report the results obtained using PnP ($Ours_{2D-3D}$) as well as, given depth information is available, using Kabsch algorithm ($Ours_{3D-3D}$). Our model does not depend in any way on the algorithm used to compute pose predictions; therefore, we can easily interchange these algorithms without the need to train additional models.

In most cases, we found a significant improvement in pose accuracy compared to [10] and adaptations of this method [26]. A recent method, namely VLocNet [24], has achieved comparable results to our method, however using direct pose regression. By evaluating the layers up until weights are shared between global pose regression and the odometry stream, they are finally able to achieve an overall median rotation and translation error of 3.1°, 4.2 cm on the 7-Scenes dataset.

In addition, we compare to recent works on scene coordinate regression, [4]. Although, there has been a very recent version of this work [3], we compare to the earlier version, since its framework is more similar to our approach, keeping the hypothesis scoring CNN in mind. With the exception of the challenging *Stairs* scene, the state-of-the art method shows slightly better accuracy in terms of RGB pose estimation considering each scene individually. On

Table 2: Median rotation and translation error on the 7-Scenes dataset. Percentages are given for poses below 5° and $5cm$ threshold.

| Scene | Volume | RGB information | | | | | | RGB-D information | |
|---------|----------|-------------------|-------------------|------------------|------------------|------------------|-----------------------------|-----------------------------|--------------|
| | | PoseNet LSTM [26] | PoseNet 2017 [10] | VLocNet STL [24] | VLocNet MTL [24] | DSAC [4] | <i>Ours_{2D-3D}</i> | <i>Ours_{3D-3D}</i> | SCoRe [23] |
| Chess | $6m^3$ | 5.7°, 24cm | 4.8°, 13cm | 1.7°, 4cm | 1.7°, 3cm | 1.2°, 2cm | 1.3°, 3cm (83.0%) | 1.2°, 3cm (85.7%) | 92.6% |
| Fire | $2.5m^3$ | 11.9°, 34cm | 11.3°, 27cm | 5.3°, 4cm | 4.9°, 4cm | 1.5°, 4cm | 2.9°, 6cm (42.4%) | 2.7°, 5cm (48.8%) | 82.9% |
| Heads | $1m^3$ | 13.7°, 21cm | 13°, 17cm | 6.6°, 5cm | 5.0°, 5cm | 2.7°, 3cm | 3.2°, 4cm (59.6%) | 3.1°, 3cm (60.1%) | 49.4% |
| Office | $7.5m^3$ | 8.0°, 30cm | 5.5°, 19cm | 2.0°, 4cm | 1.5°, 3cm | 1.6°, 4cm | 2.1°, 6cm (42.5%) | 2.0°, 5cm (49.2%) | 74.9% |
| Pumpkin | $5m^3$ | 7.0°, 33cm | 4.7°, 26cm | 2.3°, 4cm | 1.9°, 4cm | 2.0°, 5cm | 2.9°, 4cm (62.2%) | 1.3°, 3cm (66.6%) | 73.7% |
| Kitchen | $18m^3$ | 8.8°, 37cm | 5.3°, 23cm | 2.3°, 4cm | 1.7°, 3cm | 2.0°, 5cm | 2.7°, 5cm (58.2%) | 1.3°, 4cm (66.4%) | 71.8% |
| Stairs | $7.5m^3$ | 13.7°, 40cm | 12.4°, 35cm | 6.5°, 10cm | 5.0°, 7cm | 33.1°, 1.17m | 6.3°, 13cm (9.9%) | 6.1°, 13cm (11.6%) | 27.8% |
| Average | | 9.8°, 31.3cm | 8.1°, 22.9cm | 3.8°, 5cm | 3.1°, 4cm | 6.3°, 20cm | 3.1°, 5.8cm (51.1%) | 2.5°, 5.2cm (55.5%) | 67.6% |

average our method shows good performance compared to the state-of-the art.

Although our confidence prediction significantly improves the results, the initial scene coordinate regression still seems erroneous, which will be further explored in future work considering optimizations in handling missing depth and thus ground truth scene coordinates as well. Given that the depth information is available, improvements of the accuracy using RGB-D information can easily be obtained since neither the models nor the pose refinement rely on these algorithms. The results including RGB-D information can be seen in Table 2, for which we give a comparison to the state-of-the art scene coordinate regression forest approach [23]. Also, since we only depend on the most confident points, our results were obtained using a smaller number of points. Further, RANSAC based optimization, as applied in most state-of-the art methods, could be easily applied to obtain more accurate pose estimates. For evaluation and comparison to current RGB methods, we keep the parameter settings for pose refinement as proposed in [4].

In terms of computational time, our method is evaluated on an Intel Core i7 4.2 GHz CPU. The only part, which is running on GPU is the evaluation of the networks. To solve the PnP, we use OpenCV’s implementation of [13]. Our framework runs in 1.06 s, most of which is due to hypothesis sampling (0.8 s) and refinement (0.26 s). In comparison, DSAC, implemented in C++, reports a run-time of approximately 1.5 s, whereas our method, implemented in Python, already performs well.

5 Conclusion

In this work, we present a framework for dense scene coordinate regression in the context of camera re-localization using a single RGB image as the input. When the depth information is available for obtaining the camera coordinates, the corresponding scene coordinates could be regressed and used to obtain a camera pose estimate. We incorporate this information into the network and analyze how the scene coordinate regression can be further optimized using a smoothing term in the loss function. In addition and more importantly, we predict confidences for the resulting image point to scene coordinate correspondences, from which the camera pose can be inferred, thus eliminating most of the outliers in advance and greatly improving the accuracy of the estimated camera poses. As a final step, the resulting confidences can be used to refine the initial pose estimates, which further improves the accuracy of our method.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
- [2] Vasileios Belagiannis, Christian Rupprecht, Gustavo Carneiro, and Nassir Navab. Robust optimization for deep regression. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2830–2838, 2015.
- [3] Eric Brachmann and Carsten Rother. Learning less is more-6d camera localization via 3d surface regression. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2018.
- [4] Eric Brachmann, Alexander Krull, Sebastian Nowozin, Jamie Shotton, Frank Michel, Stefan Gumhold, and Carsten Rother. Dsac-differentiable ransac for camera localization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 3, 2017.
- [5] 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 relocalisation. In *CVPR*, volume 2, page 3, 2017.
- [6] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. In *Readings in computer vision*, pages 726–740. Elsevier, 1987.
- [7] Ben Glocker, Jamie Shotton, Antonio Criminisi, and Shahram Izadi. Real-time rgb-d camera relocalization via randomized ferns for keyframe encoding. *IEEE transactions on visualization and computer graphics*, 21(5):571–583, 2015.
- [8] Abner Guzman-Rivera, Pushmeet Kohli, Ben Glocker, Jamie Shotton, Toby Sharp, Andrew Fitzgibbon, and Shahram Izadi. Multi-output learning for camera relocalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1114–1121, 2014.
- [9] Alex Kendall and Roberto Cipolla. Modelling uncertainty in deep learning for camera relocalization. In *Robotics and Automation (ICRA), 2016 IEEE International Conference on*, pages 4762–4769. IEEE, 2016.
- [10] Alex Kendall and Roberto Cipolla. Geometric loss functions for camera pose regression with deep learning. In *Proc. CVPR*, volume 3, page 8, 2017.
- [11] Alex Kendall, Matthew Grimes, and Roberto Cipolla. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *Computer Vision (ICCV), 2015 IEEE International Conference on*, pages 2938–2946. IEEE, 2015.
- [12] Young Hoon Lee and Gérard Medioni. Rgb-d camera based wearable navigation system for the visually impaired. *Computer Vision and Image Understanding*, 149:3–20, 2016.

- [13] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epn: An accurate o (n) solution to the pnp problem. *International journal of computer vision*, 81(2):155, 2009.
- [14] Oscar Mendez Maldonado, Simon Hadfield, Nicolas Pugeault, and Richard Bowden. Sedar—semantic detection and ranging: Humans can localise without lidar, can robots? In *Proceedings of the 2018 IEEE International Conference on Robotics and Automation, May 21-25, 2018, Brisbane, Australia*, 2018.
- [15] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: a versatile and accurate monocular slam system. *IEEE Transactions on Robotics*, 31(5): 1147–1163, 2015.
- [16] Tayyab Naseer and Wolfram Burgard. Deep regression for monocular camera-based 6-dof global localization in outdoor environments. In *Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on*, pages 1525–1530. IEEE, 2017.
- [17] Jiahao Pang and Gene Cheung. Graph laplacian regularization for image denoising: analysis in the continuous domain. *IEEE Transactions on Image Processing*, 26(4): 1770–1785, 2017.
- [18] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 1(2):4, 2017.
- [19] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [20] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *Computer Vision (ICCV), 2011 IEEE international conference on*, pages 2564–2571. IEEE, 2011.
- [21] Tanner Schmidt, Richard Newcombe, and Dieter Fox. Self-supervised visual descriptor learning for dense correspondence. *IEEE Robotics and Automation Letters*, 2(2):420–427, 2017.
- [22] Johannes L Schönberger, Marc Pollefeys, Andreas Geiger, and Torsten Sattler. Semantic visual localization. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2018.
- [23] 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 *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 2930–2937. IEEE, 2013.
- [24] Abhinav Valada, Noha Radwan, and Wolfram Burgard. Deep auxiliary learning for visual localization and odometry. In *International Conference on Robotics and Automation (ICRA 2018)*. IEEE, 2018.
- [25] Julien Valentin, Matthias Nießner, Jamie Shotton, Andrew Fitzgibbon, Shahram Izadi, and Philip HS Torr. Exploiting uncertainty in regression forests for accurate camera relocalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4400–4408, 2015.

- [26] Florian Walch, Caner Hazirbas, Laura Leal-Taixe, Torsten Sattler, Sebastian Hilsenbeck, and Daniel Cremers. Image-based localization using lstms for structured feature correlation. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)*, volume 1, page 3, 2017.
- [27] Shenlong Wang, Sanja Fidler, and Raquel Urtasun. Lost shopping! monocular localization in large indoor spaces. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2695–2703, 2015.
- [28] Jian Wu, Liwei Ma, and Xiaolin Hu. Delving deeper into convolutional neural networks for camera relocalization. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pages 5644–5651. IEEE, 2017.
- [29] Kwang Moo Yi, Eduard Trulls, Yuki Ono, Vincent Lepetit, Mathieu Salzmann, and Pascal Fua. Learning to find good correspondences. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2018.
- [30] Amir R Zamir, Tilman Wekel, Pulkit Agrawal, Colin Wei, Jitendra Malik, and Silvio Savarese. Generic 3d representation via pose estimation and matching. In *European Conference on Computer Vision*, pages 535–553. Springer, 2016.