



HAL
open science

Scene Coordinate Regression with Angle-Based Reprojection Loss for Camera Relocalization

Xiaotian Li, Juha Ylioinas, Jakob Verbeek, Juho Kannala

► **To cite this version:**

Xiaotian Li, Juha Ylioinas, Jakob Verbeek, Juho Kannala. Scene Coordinate Regression with Angle-Based Reprojection Loss for Camera Relocalization. ECCV 2018 - Workshop Geometry Meets Deep Learning, Sep 2018, Munich, Germany. pp.1-16. hal-01867143v1

HAL Id: hal-01867143

<https://inria.hal.science/hal-01867143v1>

Submitted on 4 Sep 2018 (v1), last revised 2 Oct 2018 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Scene Coordinate Regression with Angle-Based Reprojection Loss for Camera Relocalization

Xiaotian Li¹, Juha Ylioinas¹, Jakob Verbeek² and Juho Kannala¹

¹ Aalto University

`firstname.lastname@aalto.fi`

² Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, France

`firstname.lastname@inria.fr`

Abstract. Image-based camera relocalization is an important problem in computer vision and robotics. Recent works utilize convolutional neural networks (CNNs) to regress for pixels in a query image their corresponding 3D world coordinates in the scene. The final pose is then solved via a RANSAC-based optimization scheme using the predicted coordinates. Usually, the CNN is trained with ground truth scene coordinates, but it has also been shown that the network can discover 3D scene geometry automatically by minimizing single-view reprojection loss. However, due to the deficiencies of reprojection loss, the network needs to be carefully initialized. In this paper, we present a new angle-based reprojection loss which resolves the issues of the original reprojection loss. With this new loss function, the network can be trained without careful initialization, and the system achieves more accurate results. The new loss also enables us to utilize available multi-view constraints, which further improve performance.

Keywords: camera relocalization, scene coordinate regression, deep neural networks

1 Introduction

Image-based camera relocalization is the problem of estimating the 6 DoF camera pose in an environment from a single image. It plays a crucial role in computer vision and robotics, and is the key component for a wide range of applications, such as pedestrian localization and navigation of autonomous robots [1], simultaneous localization and mapping (SLAM) [2], and augmented reality [3,4].

Many conventional localization methods proposed in the literature [5,6,7] are based on hand-crafted local image features, such as SIFT, ORB, or SURF [8,9,10]. These methods usually require a 3D point cloud model where each 3D point is associated with the image features from which it was triangulated. Given a query image, its 6D camera pose is recovered by first finding a large set of matches between 2D image features and 3D points in the model via descriptor matching, and then using a RANSAC-based [11] strategy to reject outlier matches and estimate the camera pose on inliers. However, these local image features are limited by their hand-crafted feature detectors and descriptors. They

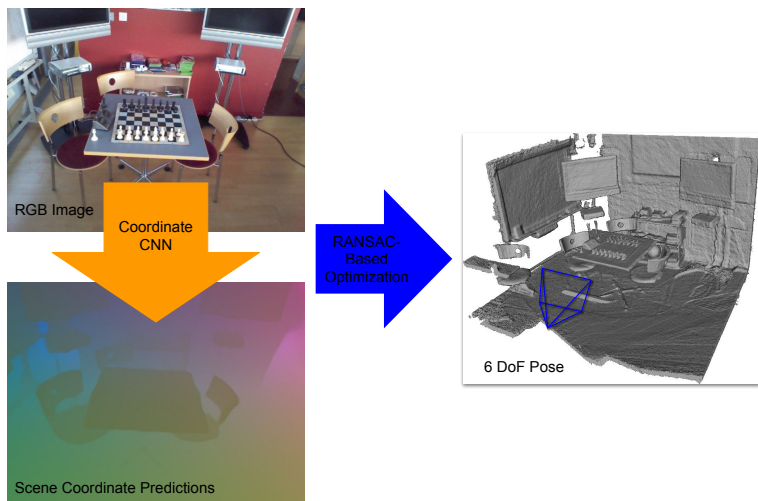


Fig. 1. Localization pipeline. In this two-stage pipeline, a coordinate CNN first produces scene coordinate predictions from an RGB image, and then the predicted correspondences are fed into a RANSAC-based solver to determine the camera pose.

are not robust enough for localization in challenging scenarios, and thus limit the use of these conventional methods.

Neural network based localization approaches have been recently explored. For example, PoseNet [12] utilizes a convolutional neural network that takes a single RGB image as input and directly regresses the 6 DoF camera pose relative to a scene. Since these methods formulate the camera pose estimation task as a regression problem which is solved using neural networks, no conventional hand-crafted features are required. These methods can successfully predict camera pose in very challenging environments where the conventional methods fail, but their overall accuracy still lags behind the conventional ones. An alternative neural network based solution is to keep the two-stage pipeline of conventional methods and formulate the first stage of the pipeline for generating 2D-3D correspondences as a learning problem. For example, the recently presented DSAC pipeline [13] predicts the 6 DoF pose by first regressing for image patches their 3D positions in the scene coordinate frame, and then determining the camera pose via a RANSAC-based scheme using the produced correspondences. The regression step in the first stage is the so-called scene coordinate regression and the 3D positions are the scene coordinates. Results have shown that these methods achieve state-of-the-art localization accuracy.

As the improved version of DSAC, DSAC++ [14] has demonstrated that scene coordinate regression can be learned even without ground truth scene coordinates. A new 3-step training strategy is proposed for training the DSAC++

pipeline. When there are no ground truth scene coordinates available, the first training step is to initialize the network using approximate scene coordinates. In the second training step, the network is further trained via optimizing reprojection loss which is measured by the sum of residuals between projected 3D points in a 2D image plane. Finally, an end-to-end optimization step is performed. The second step is crucial for the network to discover scene geometry and the first step is also necessary for the second step to work. However, initializing the network using approximate scene coordinates that are far from the ground truth ones might also degenerate localization accuracy.

In order to train the network without an initialization step, we propose in this work a novel loss function for learning scene coordinate regression. We call this new loss angle-based reprojection loss. This new loss function has better properties compared to the original reprojection loss, and thus careful initialization is not required. In addition, this new loss allows us to additionally exploit multi-view constraints.

The contributions of this paper can be summarized as follows:

- We present a novel angle-based reprojection loss, which enables the training of coordinate network without careful initialization, and improves localization accuracy.
- We show that based on this new loss function, we can incorporate multi-view constraints to further improve the accuracy.

2 Related Work

Image-based camera localization in large-scale environments is often addressed as an image retrieval problem [15]. These approaches utilize efficient and scalable retrieval approaches to determine the location of a query image. Usually, the query image is matched to a database of geotagged images and the most similar images to it are retrieved. The location of the query image is then estimated according to the known locations of the retrieved images. These approaches can scale to extremely large urban environments and even entire cities. However, all these methods can only produce an approximate location of the camera and are unable to output an exact estimate of 6 DoF camera pose.

To directly obtain 6 DoF camera pose with respect to a scene, approaches based on sparse local features such as SIFT [8] have been proposed. Instead of matching a query image to geotagged database images, these methods usually use a reconstructed 3D point cloud model obtained from Structure-from-Motion [16] to represent the scene. Each 3D point in the model is associated with one or more local image feature descriptors. Thus, a set of 2D-3D correspondences can be established by matching local image features to 3D points in the model. Based on the 2D-3D correspondences, the camera pose of the query image can be determined by a Perspective-n-Point [17] solver inside a RANSAC [11] loop. To make these methods scale to large scale environments, descriptor matching step should be efficient and effective. Therefore, techniques such as prioritized matching [5], intermediate image retrieval [18,19], geometric outlier filtering [20],

and co-visibility filtering [6] have been proposed. However, due to limitations of the hand-crafted features, these approaches can fail in challenging scenarios.

PoseNet [12] was the the first approach to tackle the problem of 6 DoF camera relocalization with deep learning. PoseNet demonstrates the feasibility of directly regressing the 6 DoF camera pose from a query RGB image via a deep CNN. Later works have been proposed to extended this method for better accuracy. For example, in [21] the authors explore Bayesian Neural Networks to produce relocalisation uncertainty of the predicted pose. LSTM-Pose [22] utilizes LSTM units [23] on the output of CNN to extract better features for localization. Another variant, Hourglass-Pose [24], makes use of the encoder-decoder hourglass architecture which can preserve the fine-grained information of input image. Besides, [25] demonstrates that PoseNet’s performance can be improved by leveraging geometric loss functions. However, theses methods are still outperformed by conventional sparse feature based methods. More recently, two multitask models VlocNet [26] and VlocNet++ [27] have been introduced. These models operate on consecutive monocular images and utilize auxiliary learning during training. Remarkably, they can provide very promising localization performance that surpasses the sparse feature based methods.

Unlike PoseNet, Laskar *et al.* introduce a deep learning based localization framework which relies on image retrieval and relative camera pose estimation [28]. This method does not require scene-specific training and can generalize well to previously unseen scenes. Similarly, Taira *et al.* propose an image retrieval based localization system for large-scale indoor environments [29]. After retrieval of candidate poses, the pose estimation is done based on dense matching. A final pose verification step via virtual view synthesis can further improve the robustness of the system. To achieve robust visual localization under a wide range of viewing conditions, semantic information has also been exploited in [30].

Our work is most closely related to methods based on the scene coordinate regression framework. The original scene coordinate regression pipeline is proposed for RGB-D camera relocalization [31]. This method formulates descriptor matching as a regression problem and utilizes a regression forest [32] to produce 2D-3D correspondences from an RGB-D input image. Similar to a sparse feature based pipeline, the final camera pose is recovered from the correspondences via a RANSAC-based solver. Since generating correspondences is directly done by the regression forest, no traditional feature extraction, feature description, or feature matching processes are required. This method has been further extended in later works [33,34,35]. Currently, practical low-cost devices are usually equipped with only RGB cameras. However, these methods still require a depth signal at both training and test time.

To localize RGB-only images, an auto-context random forest is adopted in the pipeline in [36]. In [37], Massiceti *et al.* explored random-forest-to-neural-network mapping strategy. Recently, Brachmann *et al.* proposed a differentiable version of RANSAC (DSAC) [13] and presented a localization system based on it. In the DSAC pipeline, two CNNs are adopted for predicting scene coordinates

and for scoring hypotheses respectively. Since the whole pipeline is differentiable, an end-to-end optimization step can be performed using ground truth camera poses. In contrast to DSAC which utilizes patch-based network for scene coordination regression, a full-frame network which considers global image appearance is presented in [38] and a data augmentation strategy is proposed to ensure the accuracy. However, all of these methods require ground truth scene coordinates for training. DSAC++ [14], the successor version of DSAC, demonstrates the feasibility of learning scene coordinate regression without scene coordinate ground truth. This is achieved by first initializing the predictions using approximate scene coordinates and then optimizing reprojection errors. This method is the current state-of-the-art on the 7-Scenes dataset [31].

3 Method

In this work, we follow the two-stage pipeline of DSAC and DSAC++ for RGB-only camera relocalization. In the first stage, given an RGB image, a coordinate CNN is adopted to generate dense 3D scene coordinates predictions to form 2D-3D correspondences. In the second stage, a RANSAC-based scheme is performed to generate pose hypotheses with their scores and determine the final pose estimate which can be further refined. The overall pipeline is illustrated in Figure 1.

DSAC first demonstrates how to incorporate CNNs into the two-stage pipeline and achieves state-of-the-art results. DSAC++ further addresses the main shortcomings of the DSAC pipeline and achieves substantially more accurate results. More importantly, unlike previous works which require scene coordinate ground truth generated using RGB-D training images or available 3D model during training, DSAC++ is able to discover 3D scene geometry automatically. That is, the coordinate CNN can be trained from only RGB images with their ground truth camera poses, and no ground truth scene coordinates or depth information is needed. This is achieved by first using approximate scene coordinates to initialize the network and then optimizing reprojection loss which is calculated using the ground truth poses. In the initialization step, the approximate scene coordinates are generated to have a constant distance d from the camera plane and serve as the dummy ground truth for training. The reprojection loss can then help the network to recover the scene geometry, but it may not work without the initialization step, even though the initialization step itself provides poor localization performance. However, the DSAC++ pipeline still has a main drawback. That is, a proper value for constant distance d should be selected carefully, since a value that is far off the actual range of distances can result in poor accuracy [14].

In the following, we present a new angle-based reprojection loss that allows us to train the coordinate CNN without careful initialization. We explain the deficiencies of the original reprojection loss and present our new loss function in Section 3.1. This new loss function also enables us to utilize available multi-view

reprojection and photo-consistency constraints, which we discuss in Section 3.2 and Section 3.3 respectively.

3.1 Angle-Based Reprojection Loss

For a training image I_i , the reprojection loss is calculated using the scene coordinate predictions and the ground truth pose. It can be written as³:

$$\mathcal{L}_{rep} = \sum_{k \in P_i} \|Ch_i^{-1}y_k(I_i; w) - p_k^i\| \quad (1)$$

where C is the intrinsic matrix, h_i is the camera pose for image I_i , P_i is the set of all points in image I_i , $y_k(I_i; w)$ is the scene coordinate prediction for point k with input image I_i and learnable parameters w , p_k^i is the 2D position of point k in image I_i .

As mentioned above, the DSAC++ pipeline requires a carefully selected distance d for generating approximate scene coordinates when training without ground truth scene coordinates. However, improper selection of d may also lead to poor performance. Therefore, one may wonder if it is possible to train the network without an initialization step. Unfortunately, if we train the network directly using the reprojection loss without the initialization step, the training might not converge due to unstable gradients. The unstable gradients are caused by minimizing the reprojection loss at beginning of training when the predictions of the network could be behind the camera or close to the camera center.

The shortcomings of the reprojection loss are illustrated in Figure 2. As we can see, the reprojection loss does not constrain the predictions to be in front of the camera, and the loss could be 0 even when the predictions are behind the camera. Therefore, in such cases the reprojection loss can not help the network to discover the true geometry of the scene. In addition, when the z coordinate of a prediction in the camera coordinate frame is very close to zero, the corresponding projected point in the image plane could be very far away from the ground truth one, resulting in very large loss value and exploding gradients. More importantly, we believe the ability of the network to discover the scene geometry automatically comes from the patch-based nature of the predictions. The output neurons of the fully-convolutional network used in the DSAC++ pipeline have a limited receptive field. Since local patch appearance is relatively invariant w.r.t. to viewpoint change, the explicitly applied single-view reprojection loss could be considered as implicit multi-view constraints. However, when the predictions are behind camera, minimizing reprojection loss to fulfill multi-view constraints could lead to inconsistent gradients. And this again prevents the network from discovering the true geometry of a scene.

The aforementioned deficiencies restrict the use of reprojection loss. In particular, before using the reprojection loss, we need to initialize the predictions to be in front of the camera using a constant distance d . However, this might lead

³ We omit the conversions between homogeneous coordinates and Cartesian coordinates for simplicity. This also applies to equations in the rest of the paper.

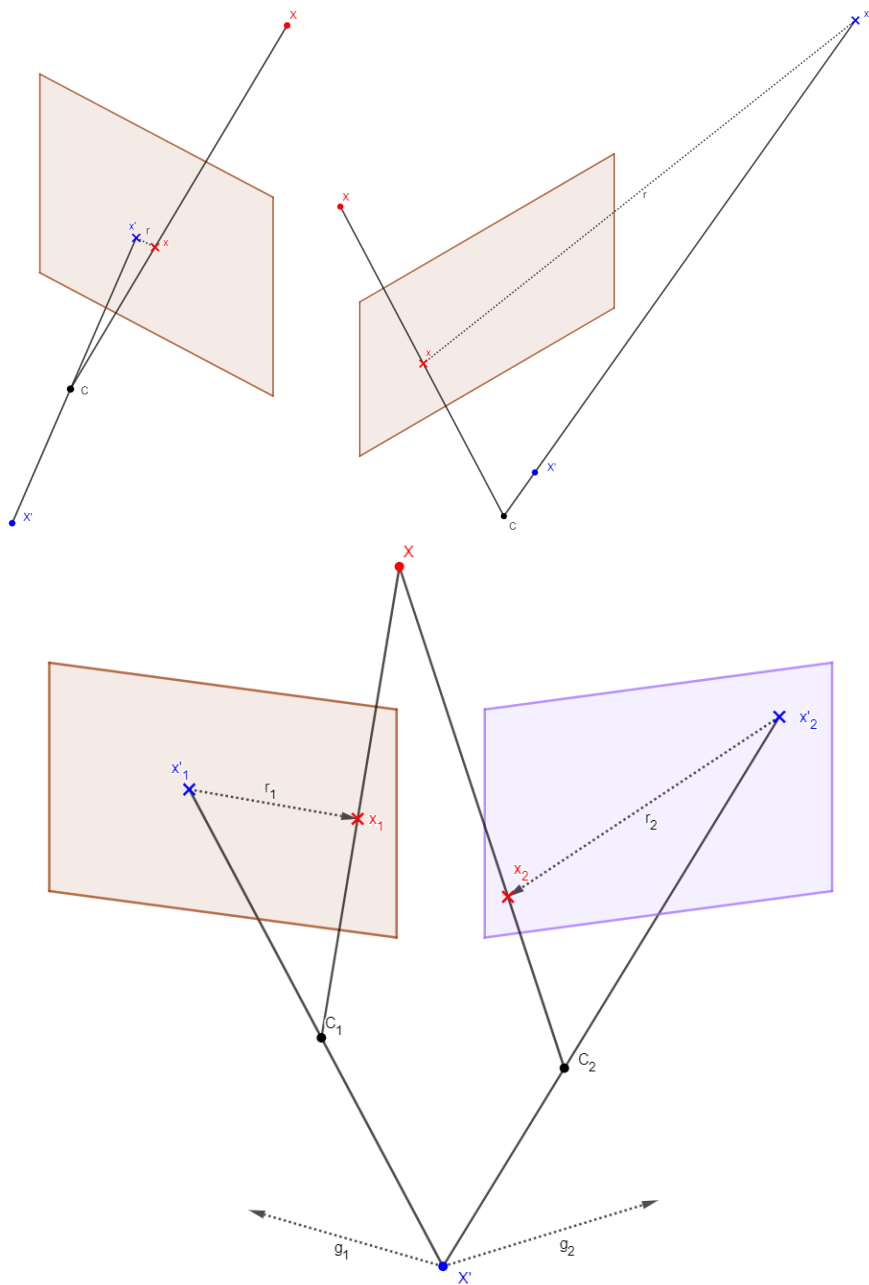


Fig. 2. Top left: when predictions are behind camera, loss can still be 0. Top right: reprojected predictions can easily be far from the ground truth. Bottom: predictions behind camera can cause inconsistent gradients.

to overfitting to the dummy ground truth scene coordinates, and thus affect the localization accuracy.

To address these issues, we present a new angle-based reprojection loss. Instead of minimizing the distance between the projected ground truth and projected prediction in the image plane, we seek to minimize the angle θ between two rays that share the camera center as the common endpoint and pass through the scene coordinate prediction and ground truth respectively. To be specific, we minimize the distance between two points in which the two rays intersect a sphere with center at the camera center and radius equal to the distance between the projected ground truth and the camera center, as shown in Figure 3. More formally, for a training image I_i , the angle-based reprojection loss is defined as⁴:

$$\mathcal{L}_{ang} = \sum_{k \in P_i} \left\| \frac{\|d_k^i\|}{\|D_k^i\|} h_i^{-1} y_k(I_i; w) - f C^{-1} p_k^i \right\| \quad (2)$$

where f is the focal length, $d_k^i = f C^{-1} p_k^i$, and $D_k^i = h_i^{-1} y_k(I_i; w)$.

As we can see, the new angle-based reprojection loss forces the predictions to be in front of camera since the loss function has large value when predictions are behind camera. Moreover, it no longer causes unstable gradients when predictions are behind camera or z coordinates of the predictions in the camera frame are close to zero. In addition, it is easy to see that the new loss function is approximately equal to the original one when predictions are close to ground truth scene coordinates. Assume $h_i^{-1} y_k(I_i; w) = (X_k^i, Y_k^i, Z_k^i)^T$, $p_k^i = (x_k^i, y_k^i)^T$, and the ground truth $\dot{D}_k^i = (\dot{X}_k^i, \dot{Y}_k^i, \dot{Z}_k^i)^T$. When predictions are close to ground truth, we have $\frac{\|d_k^i\|}{\|D_k^i\|} \approx \frac{\|\dot{d}_k^i\|}{\|\dot{D}_k^i\|} = \frac{f}{Z_k^i} \approx \frac{f}{\dot{Z}_k^i}$. Therefore:

$$\begin{aligned} \mathcal{L}_{ang} &= \sum_{k \in P_i} \left\| \frac{\|d_k^i\|}{\|D_k^i\|} h_i^{-1} y_k(I_i; w) - f C^{-1} p_k^i \right\| \\ &= \sum_{k \in P_i} \sqrt{\left(\frac{\|d_k^i\|}{\|D_k^i\|} X_k^i + c_x - x_k^i \right)^2 + \left(\frac{\|d_k^i\|}{\|D_k^i\|} Y_k^i + c_y - y_k^i \right)^2 + \left(\frac{\|d_k^i\|}{\|D_k^i\|} Z_k^i - f \right)^2} \\ &\approx \sum_{k \in P_i} \sqrt{\left(\frac{f}{Z_k^i} X_k^i + c_x - x_k^i \right)^2 + \left(\frac{f}{Z_k^i} Y_k^i + c_y - y_k^i \right)^2} \\ &= \sum_{k \in P_i} \|C h_i^{-1} y_k(I_i; w) - p_k^i\| = \mathcal{L}_{rep} \end{aligned} \quad (3)$$

where c_x and c_y are parameters for principal point offset. This is reason that we call the new loss function angle-based reprojection loss.

Note that in order to minimize the angle θ , one may also maximize the dot product of two unit vectors of the two rays, but the resulting loss function no longer has the property to be approximately equal to the reprojection loss when predictions are close to ground truth scene coordinates.

⁴ We assume $f_x = f_y = f$ in this paper, but the loss function is also applicable when $f_x \neq f_y$.

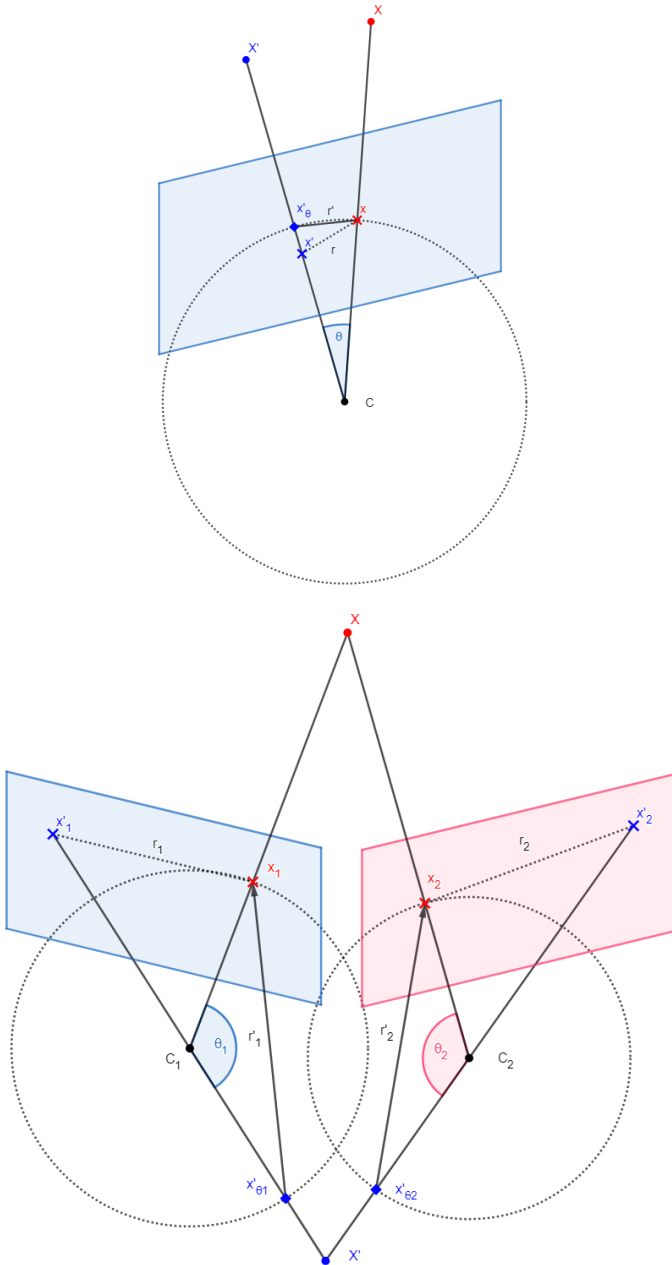


Fig. 3. Top: to minimize θ , we minimize r' . Bottom: the new loss enforces predictions to be in front of camera and does not cause unstable gradients.

3.2 Multi-View Reprojection Loss

Although with only single-view reprojection loss, the network can already approximately recover the scene geometry, sometimes it is still difficult for the network to learn accurate scene coordinate predictions, *e.g.*, when ground truth poses are erroneous or there are ambiguities in the scene. Therefore, if we could exploit explicitly multi-view constraints during training, we might achieve more accurate localization performance.

One way to incorporate multi-view constraints is to utilize available multi-view correspondences to form multi-view reprojection loss. The correspondences might come from a reconstructed 3D model or can be directly generated by correspondence matching.

Since the original reprojection loss is problematic when predictions are behind camera, it is more likely to cause problem in the multi-view cases. Therefore, we formulate the multi-view reprojection loss using our new angle-based reprojection loss. For one training image, the loss function can be written as:

$$\begin{aligned} \mathcal{L}_{multi} = & \sum_{k \in P_i^-} \left\| \frac{\|d_k^i\|}{\|D_k^i\|} h_i^{-1} y_k(I_i; w) - fC^{-1} p_k^i \right\| \\ & + \lambda \sum_{k \in P_i^*} \sum_{j \in \mathcal{I}_k \cup \{I_i\}} \left\| \frac{\|d_k^j\|}{\|D_k^j\|} h_j^{-1} y_k(I_i; w) - fC^{-1} p_k^j \right\| \end{aligned} \quad (4)$$

where λ is the balance weight, P_i^- is the set of all points in image I_i which have no corresponding points, P_i^* is the set of points in image I_i which have corresponding points in other training images, \mathcal{I}_k is the set of images in which point k has a corresponding point. As we can see, $\sum_{k \in P_i^-} \left\| \frac{\|d_k^i\|}{\|D_k^i\|} h_i^{-1} y_k(I_i; w) - fC^{-1} p_k^i \right\|$ encodes single-view constraints, and $\sum_{k \in P_i^*} \sum_{j \in \mathcal{I}_k \cup \{I_i\}} \left\| \frac{\|d_k^j\|}{\|D_k^j\|} h_j^{-1} y_k(I_i; w) - fC^{-1} p_k^j \right\|$ encodes multi-view constraints.

3.3 Photometric Reconstruction Loss

If training data consists of sequences of images, it is also possible to constrain the scene coordinate predictions using photometric reconstruction loss. Photometric reconstruction loss has recently become a dominant strategy for solving many geometry related learning problems. It enables neural networks to learn many tasks, such as monocular depth [39,40], ego-motion [41], and optical flow [42], in a self-supervised manner.

For learning scene coordinate regression, given a training image I_i and a neighboring image I_j , we can reconstruct the training image by sampling pixels from the neighboring image based on the scene coordinate predictions and the pose of the neighboring image. Specifically, for a point k in the training image I_i , we can project it to the neighboring image and obtain its 2D coordinates in the neighboring image:

$$p_k^j = Ch_j^{-1} y_k(I_i; w) \quad (5)$$

Then, the pixel value for point k in the reconstructed image \hat{I}_i can be sampled using bilinear sampler based on the coordinates p_k^j . That is, the sampled pixel value is the weighted sum of the values of the four neighboring pixels of p_k^j .

Following [40], our photometric reconstruction loss is a combination of an $L1$ loss term and a simplified SSIM loss term with a 3×3 block filter:

$$\mathcal{L}_{pr} = \sum_{k \in P_i} (1 - \alpha) | \hat{I}_i^k - I_i^k | + \alpha \frac{1 - SSIM(\hat{I}_i^k, I_i^k)}{2} \quad (6)$$

where $\alpha = 0.85$.

Since in the beginning of training, the scene coordinate predictions could be quite far from the ground truth ones, the projected 2D coordinates of points in the neighboring image might also be far from valid pixels. In other words, we are unable to obtain valid gradients by minimizing the photometric reconstruction loss. Therefore, using only photometric reconstruction loss for learning scene coordinate regression may not work. Thus, We pair it with our angle-base reprojection loss, and the final loss function can be written as:

$$\mathcal{L}_{ang+pr} = \mathcal{L}_{ang} + \lambda \mathcal{L}_{pr} \quad (7)$$

4 Experiments

We evaluate the performance of our approach on the 7-Scenes dataset [31] which is described in Section 4.1. The implementation details are explained in Section 4.2 and the results are given in Section 4.3.

4.1 Evaluation Dataset

The 7-Scenes dataset [31] is a widely used RGB-D dataset provided by Microsoft Research. It consists of seven different indoor environments, and exhibits several challenges, such as motion blur, illumination changes, textureless surfaces, repeated structures, reflections, and sensor noise. Each scene contains multiple sequences of tracked RGB-D camera frames and these sequences are split into training and testing data. 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. In contrast to some previous works, ground truth scene coordinates are not necessary for learning scene coordinate regression in this work. Therefore, we use only RGB images at both training and test time.

4.2 Implementation Details

We adopt the FCN network architecture used in the DSAC++ pipeline which takes a 640×480 image as input and produces 80×60 scene coordinate predictions. For implementation simplicity, unlike DSAC++, in the second stage of the

pipeline, we use a non-differentiable version of the RANSAC-based scheme for pose estimation, which scores hypotheses using inlier count instead of soft inlier count. The details of the used pose estimation algorithm are explained in [38]. In this paper, we mainly focus on the loss functions for learning scene coordinate regression in the first stage of the pipeline, and thus the entire pipeline does not have to be end-to-end trainable. However, it is also very straightforward to adopt the differentiable optimization strategy and perform an extra end-to-end optimization step at the end of the training as described in [14].

We train the network from scratch without the scene coordinate initialization training step for 600K iterations with a batch size of 1 using the Adam [43] optimizer where $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$. The initial learning rate is set to 0.0001 and is halved after 180K, 360K, 420K, 480K and 540K iterations. Following [14], training images are randomly shifted by a maximum of 8 pixels, horizontally and vertically, to make full use of the training data, since the network produces one prediction for each 8×8 image block.

When training the network using multi-view reprojection loss, for efficiency, we do not use all the correspondences to calculate the latter term of the loss. In particular, we calculate $\sum_{k \in P_i^*} \sum_{j \in \{I_i, I_m\}} \left\| \frac{\|d_k^j\|}{\|D_k^j\|} h_j^{-1} y_k(I_i; w) - fC^{-1} p_k^j \right\|$ instead of $\sum_{k \in P_i^*} \sum_{j \in \mathcal{I}_k \cup \{I_i\}} \left\| \frac{\|d_k^j\|}{\|D_k^j\|} h_j^{-1} y_k(I_i; w) - fC^{-1} p_k^j \right\|$ where I_m is an image randomly selected from \mathcal{I}_k . The co-visibility information used to form the multi-view constraints is extracted from a sparse model constructed using Colmap [44]. And we set the balance weight $\lambda = 60$.

When computing photometric reprojection loss, the neighboring image is randomly selected from the same sequence with difference in index less or equal to 10. Since the network produces 80×60 scene coordinate predictions, we resize both the training images (after shifting) and neighboring images to 80×60 , and the camera intrinsic parameters are also adjusted accordingly. The balance weight λ is set to 20.

4.3 Results

The experimental results of our method on the 7-Scenes dataset are given in Table 1. We report both median localization errors and accuracy measured as the percentage of query images for which the camera pose error is below 5° and 5cm, as the latter one can better represent the localization performance. Our method is compared to the state-of-the-art DSAC++ method trained without an accurate 3D scene model.

Accordingly to the results, when training the network using our new angle-based reprojection loss without initializing the network predictions using constant distance value, the fraction of accurately localized test images is improved by 8.8 percentage points compared to DSAC++. Moreover, the accuracy can be further improved by utilizing either multi-view correspondences or photometric reconstruction metric. Compared to using single-view constraints only, we observed that the additional loss terms could make the network to produce

Table 1. Results of our method on the 7-Scenes dataset. We compare the performance of our method to DSAC++. Note that Complete denotes the accuracy measured by the 5° , 5cm criterion on the combined set of test frames (17000) of all scenes.

Scene	5° , 5cm (Median Error)			
	DSAC++ [14]	Ours	Ours+multi-view	Ours+photometric
Chess	93.8% (2cm, 0.7°)	95.1% (2cm, 0.6°)	96.1% (2cm, 0.6°)	96.0% (2cm, 0.6°)
Fire	75.6% (3cm, 1.1°)	84.0% (3cm, 1.0°)	88.6% (2cm, 1.0°)	86.4% (2cm, 1.0°)
Heads	18.4% (12cm, 6.7°)	80.5% (2cm, 1.1°)	86.9% (2cm, 1.0°)	83.2% (2cm, 1.1°)
Office	75.4% (3cm, 0.8°)	80.4% (3cm, 0.8°)	80.6% (3cm, 0.8°)	81.6% (3cm, 0.8°)
Pumpkin	55.9% (5cm, 1.1°)	56.8% (4cm, 1.1°)	60.3% (4cm, 1.1°)	59.2% (4cm, 1.1°)
Red Kitchen	50.7% (5cm, 1.3°)	59.9% (4cm, 1.2°)	61.9% (4cm, 1.2°)	60.0% (4cm, 1.2°)
Stairs	2.0% (29cm, 5.1°)	2.9% (25cm, 4.5°)	11.3% (13cm, 3.4°)	4.7% (22cm, 4.9°)
Complete	60.4%	69.2%	71.8%	70.4%

more accurate scene coordinate predictions during training. This means that incorporating multi-view constraints could help the network better discover scene geometry, and thus leads to improved localization performance.

We also attempted to train the network using the original reprojection loss without the scene coordinate initialization step. However, we found that for most of the scenes, training could not converge in both single-view and multi-view cases. In addition, when using photometric reconstruction loss, for all the scenes, we observed that training of the network without the angle-based reprojection loss term would always get stuck in a local minimum with very large scene coordinate prediction error, resulting in completely failed localization system.

When training using either multi-view reprojection loss or photometric reconstruction loss, the balance weight λ is important for achieving good localization results. If the weight is set too small, the multi-view constraints will not have effect on training. However, if it is set too large, we found that the localization accuracy could also drop. For example, when training using photometric reconstruction loss with $\lambda = 50$, test accuracy for Office decreases to 79.5%. It would be interesting to explore learnable weighting strategy presented in [25] in future work.

Note that our method is still less accurate compared to the DSAC++ method trained with a 3D model (76.1%), as our method is trained without access to adequate scene geometry information. When training with multi-view reprojection loss, only a sparse set of points have corresponding points in other images (usually about 1000 points for a 640×480 image). That is, the multi-view constraints are active for only a small portion of the points in a training image. We believe that with denser correspondence information our method could achieve better results. For photometric reconstruction loss, adding an additional loss term similar to the left-right consistency proposed in [40] which enforces predictions to be consistent between different views might be also helpful to further improve the accuracy but we did not explore it.

5 Conclusion

In this work, we have presented a new angle-based reprojection loss for learning scene coordinate regression for image-based camera relocalization. Our novel loss function makes it possible to train the coordinate CNN without a scene coordinate initialization step, resulting in improved localization accuracy. This novel loss function further allows us to explore available multi-view constraints which can additionally improve performance.

Acknowledgements

This work has been partially supported by the grant Deep in France (ANR16-CE23-0006) and LabEx PERSYVAL (ANR-11-LABX0025-01).

References

1. Cummins, M., Newman, P.: Fab-map: Probabilistic localization and mapping in the space of appearance. *The International Journal of Robotics Research* **27**(6) (2008) 647–665
2. Eade, E., Drummond, T.: Scalable monocular slam. In: *CVPR*. (2006)
3. Castle, R.O., Klein, G., Murray, D.W.: Video-rate localization in multiple maps for wearable augmented reality. In: *ISWC*. (2008)
4. Lynen, S., Sattler, T., Bosse, M., Hesch, J.A., Pollefeys, M., Siegwart, R.: Get out of my lab: Large-scale, real-time visual-inertial localization. In: *RSS*. (2015)
5. Sattler, T., Leibe, B., Kobbelt, L.: Fast image-based localization using direct 2d-to-3d matching. In: *ICCV*. (2011)
6. Sattler, T., Leibe, B., Kobbelt, L.: Improving image-based localization by active correspondence search. In: *ECCV*. (2012)
7. Sattler, T., Leibe, B., Kobbelt, L.: Efficient effective prioritized matching for large-scale image-based localization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **39**(9) (2017) 1744–1756
8. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* **60**(2) (2004) 91–110
9. Rublee, E., Rabaud, V., Konolige, K., Bradski, G.R.: ORB: an efficient alternative to SIFT or SURF. In: *ICCV*. (2011)
10. Bay, H., Tuytelaars, T., Gool, L.V.: Surf: Speeded up robust features. In: *ECCV*. (2006)
11. Fischler, M.A., Bolles, R.C.: Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* **24**(6) (1981) 381–395
12. Kendall, A., Grimes, M., Cipolla, R.: PoseNet: A convolutional network for real-time 6-dof camera relocalization. In: *ICCV*. (2015)
13. Brachmann, E., Krull, A., Nowozin, S., Shotton, J., Michel, F., Gumhold, S., Rother, C.: Dsac - differentiable ransac for camera localization. In: *CVPR*. (2017)
14. Brachmann, E., Rother, C.: Learning less is more - 6d camera localization via 3d surface regression. In: *CVPR*. (2018)
15. Zhang, W., Kosecka, J.: Image based localization in urban environments. In: *3DPVT*. (2006)

16. Snavely, N., Seitz, S.M., Szeliski, R.: Modeling the world from internet photo collections. *International Journal of Computer Vision* **80**(2) (2008) 189–210
17. Kneip, L., Scaramuzza, D., Siegwart, R.: A novel parametrization of the perspective-three-point problem for a direct computation of absolute camera position and orientation. In: *CVPR*. (2011)
18. Irschara, A., Zach, C., Frahm, J.M., Bischof, H.: From structure-from-motion point clouds to fast location recognition. In: *CVPR*. (2009)
19. Sattler, T., Torii, A., Sivic, J., Pollefeys, M., Taira, H., Okutomi, M., Pajdla, T.: Are large-scale 3d models really necessary for accurate visual localization? In: *CVPR*. (2017)
20. Zeisl, B., Sattler, T., Pollefeys, M.: Camera pose voting for large-scale image-based localization. In: *ICCV*. (2015)
21. Kendall, A., Cipolla, R.: Modelling uncertainty in deep learning for camera relocalization. In: *ICRA*. (2016)
22. Walch, F., Hazirbas, C., Leal-Taixe, L., Sattler, T., Hilsenbeck, S., Cremers, D.: Image-based localization using lstms for structured feature correlation. In: *ICCV*. (2017)
23. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Computation* **9**(8) (1997) 1735–1780
24. Melekhov, I., Ylioinas, J., Kannala, J., Rahtu, E.: Image-based localization using hourglass networks. In: *ICCVW*. (2017)
25. Kendall, A., Cipolla, R.: Geometric loss functions for camera pose regression with deep learning. In: *CVPR*. (2017)
26. Valada, A., Radwan, N., Burgard, W.: Deep auxiliary learning for visual localization and odometry. In: *ICRA*. (2018)
27. Radwan, N., Valada, A., Burgard, W.: Vlocnet++: Deep multitask learning for semantic visual localization and odometry. *CoRR* **abs/1804.08366** (2018)
28. Laskar, Z., Melekhov, I., Kalia, S., Kannala, J.: Camera relocalization by computing pairwise relative poses using convolutional neural network. In: *ICCVW*. (2017)
29. Taira, H., Okutomi, M., Sattler, T., Cimpoi, M., Pollefeys, M., Sivic, J., Pajdla, T., Torii, A.: Inloc: Indoor visual localization with dense matching and view synthesis. In: *CVPR*. (2018)
30. Schönberger, J.L., Pollefeys, M., Geiger, A., Sattler, T.: Semantic visual localization. In: *CVPR*. (2018)
31. Shotton, J., Glocker, B., Zach, C., Izadi, S., Criminisi, A., Fitzgibbon, A.W.: Scene coordinate regression forests for camera relocalization in rgb-d images. In: *CVPR*. (2013)
32. Criminisi, A., Shotton, J.: *Decision Forests for Computer Vision and Medical Image Analysis*. Springer Publishing Company, Incorporated (2013)
33. Guzmán-Rivera, A., Kohli, P., Glocker, B., Shotton, J., Sharp, T., Fitzgibbon, A.W., Izadi, S.: Multi-output learning for camera relocalization. In: *CVPR*. (2014)
34. Valentin, J.P.C., Nießner, M., Shotton, J., Fitzgibbon, A.W., Izadi, S., Torr, P.H.S.: Exploiting uncertainty in regression forests for accurate camera relocalization. In: *CVPR*. (2015)
35. Cavallari, T., Golodetz, S., Lord, N.A., Valentin, J.P.C., di Stefano, L., Torr, P.H.S.: On-the-fly adaptation of regression forests for online camera relocalisation. In: *CVPR*. (2017)
36. Brachmann, E., Michel, F., Krull, A., Yang, M.Y., Gumhold, S., Rother, C.: Uncertainty-driven 6d pose estimation of objects and scenes from a single RGB image. In: *CVPR*. (2016)

37. Massiceti, D., Krull, A., Brachmann, E., Rother, C., H. S. Torr, P.: Random forests versus neural networks - what's best for camera relocalization? In: ICRA. (2017)
38. Li, X., Ylioinas, J., Kannala, J.: Full-frame scene coordinate regression for image-based localization. In: RSS. (2018)
39. Garg, R., G, V.K.B., Reid, I.D.: Unsupervised CNN for single view depth estimation: Geometry to the rescue. In: ECCV. (2016)
40. Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: CVPR. (2017)
41. Zhou, T., Brown, M., Snavely, N., Lowe, D.G.: Unsupervised learning of depth and ego-motion from video. In: CVPR. (2017)
42. Yu, J.J., Harley, A.W., Derpanis, K.G.: Back to basics: Unsupervised learning of optical flow via brightness constancy and motion smoothness. In: ECCVW. (2016)
43. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR. (2015)
44. Schönberger, J.L., Frahm, J.: Structure-from-motion revisited. In: CVPR. (2016)