

AtLoc: Attention Guided Camera Localization

Bing Wang, Changhao Chen,* Chris Xiaoxuan Lu, Peijun Zhao,
Niki Trigoni, Andrew Markham

Department of Computer Science, University of Oxford
firstname.lastname@cs.ox.ac.uk

Abstract

Deep learning has achieved impressive results in camera localization, but current single-image techniques typically suffer from a lack of robustness, leading to large outliers. To some extent, this has been tackled by sequential (multi-images) or geometry constraint approaches, which can learn to reject dynamic objects and illumination conditions to achieve better performance. In this work, we show that attention can be used to force the network to focus on more geometrically robust objects and features, achieving state-of-the-art performance in common benchmark, even if using only a single image as input. Extensive experimental evidence is provided through public indoor and outdoor datasets. Through visualization of the saliency maps, we demonstrate how the network learns to reject dynamic objects, yielding superior global camera pose regression performance. The source code is available at <https://github.com/BingCS/AtLoc>.

Introduction

Location information is of key importance to wide variety of applications, from virtual reality to delivery drones, to autonomous driving. One particularly promising research direction is camera pose regression or localization - the problem of recovering the 3D position and orientation of a camera from an image or set of images.

Camera localization has been previously tackled by exploiting the appearance and geometry in a 3D scene, for example, key points and lines, but suffers from performance degradation when deployed in the wild (Brachmann et al. 2017; Walch et al. 2017). This is due to the fact that the hand-crafted features change significantly across different scenarios due to lighting, blur and scene dynamics leading to poor global matches. Recent deep learning based approaches are able to automatically extract features and directly recover the absolute camera pose from a single image, without any hand-engineering effort, as was demonstrated in the seminal PoseNet (Kendall, Grimes, and Cipolla 2015). Extensions include the use of different encoder networks e.g. ResNet in PoseNet Hourglass (Melekhov et al. 2017) or geometric



PoseNet+

AtLoc (Ours)

Figure 1: Saliency maps of one scene selected from Oxford RobotCar (Maddern et al. 2017) indicate that AtLoc is able to force the neural network model to focus on geometrically robust objects (e.g. building structures in the right) rather than environmental dynamics (e.g. moving vehicles in the left) compared with PoseNet+ (Brahmbhatt et al. 2018).

constraints (Kendall and Cipolla 2017). Although these techniques show good performance in general, they are plagued by a lack of robustness when faced with dynamic objects or changes in illumination. This is particularly apparent in outdoor datasets where scenes are highly variable e.g. due to moving vehicles or pedestrians.

To tackle this lack of robustness, further techniques have considered using multiple images as input to the network, with the premise being that the network can learn to reject temporally inconsistent features across frames. Examples include VidLoc (Clark et al. 2017) and the recent MapNet (Brahmbhatt et al. 2018) which achieves state-of-the-art performance in camera pose regression.

In this work we pursue an alternative approach to achieve robust camera localization and ask if we can achieve or even surpass the performance of multi-frame, sequential techniques by learning to attentively focus on parts of the image that are temporally consistent and informative e.g. buildings, whilst ignoring dynamic parts like vehicles and pedestrians, using a single image as input, as shown in Figure 1. We propose **AtLoc - Attention Guided Camera Localization**, an

*Changhao Chen is the corresponding author.
Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

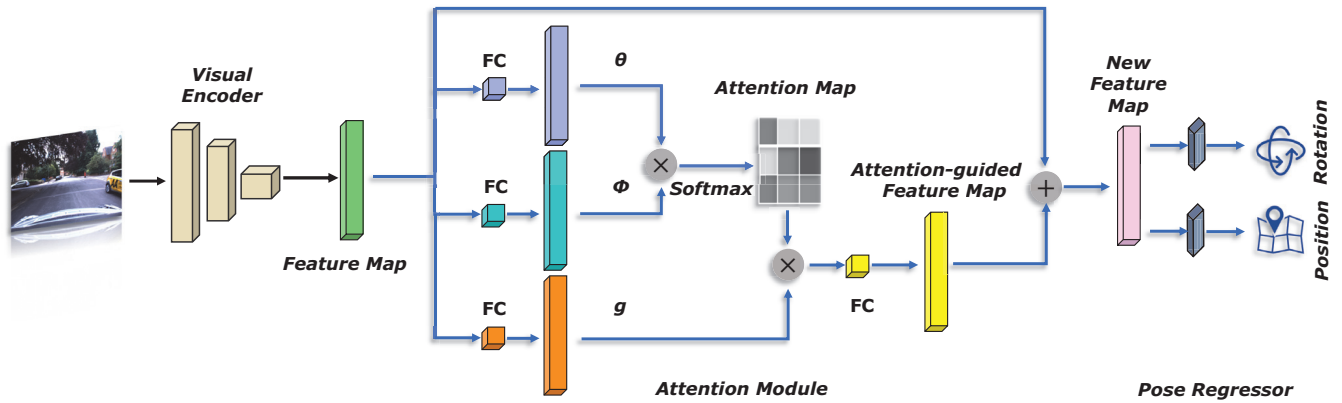


Figure 2: An overview of our proposed AtLoc framework, consisting of Visual Encoder (extracts features from a single image), Attention Module (computes the attention and reweights the features), and Pose Regressor (maps the new features into the camera pose).

attention based pose regression framework to recover camera pose. Unlike previous methods, our proposed AtLoc does not require sequential (multiple) frame nor geometry constraints designed and enforced by humans.

We show that our model outperforms previous techniques, and achieves state-of-the-art results in common benchmarks. It works efficiently across both indoor and outdoor scenarios and is simple and end-to-end trainable without requiring any hand-crafted geometric loss functions. We provide detailed insight into how incorporating attention allows the network to achieve accurate and robust camera localization.

The main contributions of this work are as follows:

- We propose a novel self-attention guided neural network for single image camera localization, allowing accurate and robust camera pose estimation.
- By visualizing the feature saliency map after the attention, we show how our attention mechanism encourages the framework to learn stable features.
- Through extensive experiments in both indoor and outdoor scenarios, we show that our model achieves state-of-the-art performance in pose regression, even outperforming multiple frame (sequential) methods.

Related Work

Deep Neural Networks for Camera Localization Recent attempts have investigated the camera localization using deep neural networks (DNNs). Compared with traditional structure-based methods (Chen et al. 2011; Torii et al. 2013; Liu, Li, and Dai 2017) and image retrieval-based methods (Li et al. 2012; Sattler, Leibe, and Kobbelt 2012; Arandjelovic et al. 2016), DNN-based camera localization methods can automatically learn features from data rather than building a map or a database of landmark features by hand (Sattler et al. 2019). As the seminal work in this vein, PoseNet(Kendall, Grimes, and Cipolla 2015) is the first one to adopt deep neural network to estimate camera pose from a single image. This approach is then extended by leveraging RNNs (e.g. LSTM) to spatially (Walch et al. 2017;

Wang et al. 2018a) and temporally (Clark et al. 2017) improve localization accuracy. Later on, localization performance is further improved by estimating uncertainty of the global camera pose with Bayesian CNN (Kendall and Cipolla 2016; Cai, Shen, and Reid 2018) and replacing feature extraction architecture with Residual Neural Network (Melekhov et al. 2017). However, the aforementioned approaches rely on the hand-tuned scale factor to balance the position and rotation losses during learning process. To address this issue, a learning weighted loss and a geometric reprojection loss (Kendall and Cipolla 2017) are introduced to produce more precise results. Recent efforts additionally leverage the geometric constraints from paired images (Brahmbhatt et al. 2018; Huang et al. 2019), augment training data by synthetic generation(Purkait, Zhao, and Zach 2018) or introduce pose-graph optimization with neural graph model(Parisotto et al. 2018). Instead of imposing the temporal information or geometry constraints as the previous work, we developed an attention mechanism for DNN-based camera localization to self-regulate itself, and automatically learn to constrain the DNNs to focus on geometrically robust features. Our model outperforms previous approaches, and achieves state-of-the-art results in common benchmarks.

Attention Mechanism Our work is related with the self-attention mechanisms, which have been widely embedded in various models capturing long-term dependencies (Bahdanau, Cho, and Bengio 2014; Xu et al. 2015; Yang et al. 2019). Self-attention was initially designed for machine translation (Vaswani et al. 2017; Dou et al. 2019; Cheng, Dong, and Lapata 2016), achieving the state-of-the-art performance. It is also integrated with an autoregressive model to generate image as Image Transformer (Parmar et al. 2018; Kingma and Dhariwal 2018). Another usage is to be formalized as a non-local operation to capture the spatial-temporal dependencies in video sequences (Wang et al. 2018b; Yuan, Mei, and Zhu 2019). A similar non-local architecture was introduced to Generative Adversarial Networks (GANs) for

extracting global long-range dependencies (Zhang et al. 2018; Liu et al. 2019). (Parisotto et al. 2018) used an attention-based recurrent neural network for back-end optimization in a SLAM system, but not for camera relocalization. Despite its successes in a wide range of computer vision (Fu et al. 2019; Chen et al. 2019) and natural language process tasks, self-attention has never been explored in camera pose regression. Our work integrated non-local style self-attention mechanism into the camera localization model to show the effectiveness of correlating robust key features and improve model performance.

Attention Guided Camera Localization

This section introduces Attention Guided Camera Localization (AtLoc), an self-attention based deep neural network architecture to learn camera poses from a single image. Figure 2 illustrates a modular overview of the proposed framework, consisting of a visual encoder, an attention module and a pose regressor. The scene of a single image is compressed into an implicit representation by the visual encoder. Conditioned on the extracted features, the attention module computes the self-attention maps to re-weight the representation into a new feature space. The pose regressor further maps the new features after the attention operators into the camera pose, i.e. the 3-dimensional location and 4-dimensional quaternion (orientation).

Visual Encoder

The visual encoder serves to extract features that are necessary for the pose regression task, from a single monocular image. Previous works (Kendall and Cipolla 2017; Brahmhatt et al. 2018) showed successful applications of the classical convolutional neural network (CNN) architectures in camera pose estimation, e.g. GoogleNet (Szegedy et al. 2015) and ResNet (He et al. 2016). Among them, the ResNet based (Brahmbhatt et al. 2018) frameworks achieved more stable and precise localization results over other architectures, due to the fact that the residual networks allow to train deeper layers of neural networks and reduce the gradient vanishing problems. Therefore, we considered to adopt a residual network with 34 layers (ResNet34) as the foundation for the visual encoder in the proposed AtLoc model.

Here, the weights of ResNet34 were initialized with the ResNet34 pretrained with the image classification on ImageNet dataset (Deng et al. 2009). To encourage learning meaningful features for a pose regression, the base network is further modified by replacing the final 1000 dimensional fully-connected layer with a C dimensional fully-connected layer and removing the Softmax layers used for classification. C is the dimension of the output feature. Considering the efficiency and performance of the model, the dimension is chosen as $C = 2048$. Given an image $\mathbf{I} \in \mathbb{R}^{\hat{C} \times H \times W}$, the features $\mathbf{x} \in \mathbb{R}^C$ can be extracted via the visual encoder f_{encoder} :

$$\mathbf{x} = f_{\text{encoder}}(\mathbf{I}) \quad (1)$$

Attention Module

Although the ResNet34 based visual encoder is capable of automatically learning the necessary features for camera localization, the neural network trained in certain specific scenes can be overfitted into the featureless appearance or the environmental dynamics. This will impact the generalization capacity of the model, and degrade the model performance in testing sets, especially in the outdoor scenarios due to the moving vehicles or weather change. Unlike the previous trials by introducing the temporal information (Clark et al. 2017) or geometric constraints (Brahmbhatt et al. 2018), we propose to adapt a self-attention mechanism into our framework. As the Figure 2 shown, this self-attention module is conditioned the features extracted by the visual encoder, and generates an attention map to enforce the model to focus on stable and geometry meaningful features. It is able to self-regulate itself without any hand-engineering geometry or prior information.

We adopt a non-local style self-attention, which has been applied in video analysis (Wang et al. 2018b) and image generation (Zhang et al. 2018), in our attention module. This aims to capture the long-range dependencies and global correlations of the image features, which will help generate better attention-guided feature maps from widely separated spatial regions (Wang et al. 2018b).

The features $\mathbf{x} \in \mathbb{R}^C$ extracted by the visual encoder are first used to compute the dot-product similarity between two embedding spaces $\theta(\mathbf{x}_i)$ and $\phi(\mathbf{x}_j)$:

$$S(\mathbf{x}_i, \mathbf{x}_j) = \theta(\mathbf{x}_i)^T \phi(\mathbf{x}_j), \quad (2)$$

where embeddings $\theta(\mathbf{x}_i) = \mathbf{W}_\theta \mathbf{x}_i$ and $\phi(\mathbf{x}_j) = \mathbf{W}_\phi \mathbf{x}_j$ linearly transform features at the position i and j into two feature spaces respectively.

The normalization factor C is defined as the $C(\mathbf{x}_i) = \sum_{\forall j} S(\mathbf{x}_i, \mathbf{x}_j)$ with all feature position j . Given another linear transformation $g(\mathbf{x}_j) = \mathbf{W}_g \mathbf{x}_j$, the output attention vector \mathbf{y} is calculated via:

$$\mathbf{y}_i = \frac{1}{C(\mathbf{x}_i)} \sum_{\forall j} S(\mathbf{x}_i, \mathbf{x}_j) g(\mathbf{x}_j), \quad (3)$$

where the attention vector \mathbf{y}_i indicates to what extent the neural model focuses on the features \mathbf{x}_i at the position i . Finally, the self-attention of input features \mathbf{x} can be written as:

$$\mathbf{y} = \text{Softmax}(\mathbf{x}^T \mathbf{W}_\theta^T \mathbf{W}_\phi \mathbf{x}) \mathbf{W}_g \mathbf{x} \quad (4)$$

Furthermore, we add a residual connection back to a linear embedding of the self-attention vectors:

$$\text{Att}(\mathbf{x}) = \alpha(\mathbf{y}) + \mathbf{x}, \quad (5)$$

where the linear embedding $\alpha(\mathbf{y}) = \mathbf{W}_\alpha \mathbf{y}$ outputs a scaled self-attention vectors with learnable weights \mathbf{W}_α .

In our proposed model, fully-connected layers are implemented to generate learned weight matrices \mathbf{W}_θ , \mathbf{W}_ϕ , \mathbf{W}_g and \mathbf{W}_α in space (C/n) , where C is the number of channels of the input feature x and n is the downsampling ratio for the attention maps. Based on extensive experiments, we found that $n = 8$ performs best across different datasets.

Scene	PoseNet	Bayesian PoseNet	PoseNet Spatial LSTM	Hourglass	PoseNet17	AtLoc (Ours)
Chess	0.32m, 6.60°	0.37m, 7.24°	0.24m, 5.77°	0.15m, 6.17°	0.13m, 4.48°	0.10m, 4.07°
Fire	0.47m, 14.0°	0.43m, 13.7°	0.34m, 11.9°	0.27m, 10.8°	0.27m, 11.3°	0.25m, 11.4°
Heads	0.30m, 12.2°	0.31m, 12.0°	0.21m, 13.7°	0.19m, 11.6°	0.17m, 13.0°	0.16m, 11.8°
Office	0.48m, 7.24°	0.48m, 8.04°	0.30m, 8.08°	0.21m, 8.48°	0.19m, 5.55°	0.17m, 5.34°
Pumpkin	0.49m, 8.12°	0.61m, 7.08°	0.33m, 7.00°	0.25m, 7.01°	0.26m, 4.75°	0.21m, 4.37°
Kitchen	0.58m, 8.34°	0.58m, 7.54°	0.37m, 8.83°	0.27m, 10.2°	0.23m, 5.35°	0.23m, 5.42°
Stairs	0.48m, 13.1°	0.48m, 13.1°	0.40m, 13.7°	0.29m, 12.5°	0.35m, 12.4°	0.26m, 10.5°
Average	0.45m, 9.94°	0.47m, 9.81°	0.31m, 9.85°	0.23m, 9.53°	0.23m, 8.12°	0.20m, 7.56°

Table 1: Camera localization results on 7 Scenes (Without temporal s). For each scene, we compute the median errors in both position and rotation of various single-image based baselines and our proposed method.

Table 2: Training and testing Sequences of Oxford RobotCar. LOOP is a relatively shorter subset (1120m in total length) and FULL covers a length of 9562m.

Sequence	Time	Tag	Mode
–	2014-06-26-08-53-56	overcast	Training
–	2014-06-26-09-24-58	overcast	Training
LOOP1	2014-06-23-15-41-25	sunny	Testing
LOOP2	2014-06-23-15-36-04	sunny	Testing
–	2014-11-28-12-07-13	overcast	Training
–	2014-12-02-15-30-08	overcast	Training
FULL1	2014-12-09-13-21-02	overcast	Testing
FULL2	2014-12-12-10-45-15	overcast	Testing

Learning Camera Pose

The pose regressor maps the attention guided features $\text{Att}(\mathbf{x})$ to location $\mathbf{p} \in \mathbb{R}^3$ and quaternion $\mathbf{q} \in \mathbb{R}^4$ respectively through Multilayer Perceptrons (MLPs):

$$[\mathbf{p}, \mathbf{q}] = \text{MLPs}(\text{Att}(\mathbf{x})) \quad (6)$$

Given training images \mathbf{I} and their corresponding pose labels $[\hat{\mathbf{p}}, \hat{\mathbf{q}}]$ represented by the camera position $\hat{\mathbf{p}} \in \mathbb{R}^3$ and a unit quaternion $\hat{\mathbf{q}} \in \mathbb{R}^4$ for orientation, the parameters inside the neural networks are optimized with L1 Loss via the following loss function (Brahmbhatt et al. 2018):

$$\text{loss}(\mathbf{I}) = \|\mathbf{p} - \hat{\mathbf{p}}\|_1 e^{-\beta} + \beta + \|\log \mathbf{q} - \log \hat{\mathbf{q}}\|_1 e^{-\gamma} + \gamma \quad (7)$$

where β and γ are the weights that balance the position loss and rotation loss. $\log \mathbf{q}$ is the logarithmic form of an unit quaternion \mathbf{q} , which is defined as:

$$\log \mathbf{q} = \begin{cases} \frac{\mathbf{v}}{\|\mathbf{v}\|} \cos^{-1} u, & \text{if } \|\mathbf{v}\| \neq 0 \\ \mathbf{0}, & \text{otherwise} \end{cases} \quad (8)$$

Here, \mathbf{u} denotes the real part of an unit quaternion while \mathbf{v} is its imaginary part. For all scenes, both β and γ are simultaneously learned during training with approximate initial values of β_0 and γ_0 . In camera pose regression tasks, quaternions are widely used to represent the orientation due their ease of formulation in a continuous and differentiable way. By normalizing any 4D quaternions to unit length, we can easily

Table 3: Camera localization results on 7 Scenes (With temporal Constraints). For each scene, we compare the median errors in both position and rotation of VidLoc, MapNet and our approach.

Scene	VidLoc	MapNet	AtLoc+ (Ours)
Chess	0.18m, NA	0.08m , 3.25°	0.10m, 3.18°
Fire	0.26m, NA	0.27m, 11.7°	0.26m, 10.8°
Heads	0.14m, NA	0.18m, 13.3°	0.14m, 11.4°
Office	0.26m, NA	0.17m, 5.15°	0.17m, 5.16°
Pumpkin	0.36m, NA	0.22m, 4.02°	0.20m, 3.94°
Kitchen	0.31m, NA	0.23m, 4.93°	0.16m, 4.90°
Stairs	0.26m, NA	0.30m, 12.1°	0.29m, 10.2°
Average	0.25m, NA	0.21m, 7.77°	0.19m, 7.08°

map any rotations in 3D space to valid unit quaternions. But this has one main issue: quaternions are not unique. In practice, both $-\mathbf{q}$ and \mathbf{q} can represent the same rotation because a single rotation can be mapped to two hemispheres. To ensure that each rotation only has a unique value, all quaternions are restricted to the same hemisphere in this paper.

Temporal Constraints

Sharing the same flavor with geometry-aware learning methods (Brahmbhatt et al. 2018; Xue et al. 2019; Huang et al. 2019), we extend our proposed **AtLoc** to **AtLoc+** by incorporating temporal constraints between image pairs. Intuitively, temporal constraints can enforce the learning of globally consist features, and thereby improve the overall localization accuracy. In this work, the loss considering temporal constraints is defined as:

$$\text{loss}(\mathbf{I}_{total}) = \text{loss}(\mathbf{I}_i) + \alpha \sum_{i \neq j} \text{loss}(\mathbf{I}_{ij}) \quad (9)$$

where \mathbf{i} and \mathbf{j} indicate the index of images. $\mathbf{I}_{ij} = (\mathbf{p}_i - \mathbf{p}_j, \mathbf{q}_i - \mathbf{q}_j)$ represents the relative pose between images \mathbf{I}_i and \mathbf{I}_j . α denotes the weight coefficient between the loss of the absolute pose from a single image and the relative pose from image pairs.



PoseNet

AtLoc (Ours)

Figure 3: Saliency maps of two scenes selected from Chess. Each scene contains the saliency maps generated by PoseNet (left) and AtLoc (right) using attention.

Experiments

To train the proposed network consistently on different datasets, we rescale the images such that the shorter side is of length 256 pixels. The input images are then normalized to have pixel intensities within the range -1 to 1. The ResNet34 (He et al. 2016) component in our network is initialized by using a pretrained model on the ImageNet dataset while the remaining components follow random initialization. 256×256 pixels images are cropped for our network during both the training and testing phase with random and central cropping strategy respectively. For the training on Oxford RobotCar dataset, random ColorJitter is additionally applied when performing data augmentation, with values of 0.7 for brightness, contrast and saturation setting and 0.5 for hue. We note that this augmentation step is essential to improve the generalization ability of model over various weather and time-of-day conditions. We implement our approaches with PyTorch, using the ADAM solver (Kingma and Ba 2014) and an initial learning rate of 5×10^{-5} . The network is trained on a NVIDIA Titan X GPU with the following hyperparameters: mini-batch size of 64, dropout rate probability of 0.5 and weight initializations of $\beta_0 = 0.0$ and $\gamma_0 = -3.0$. When introducing temporal constraints, we sample consecutive triplets every 10 frames with $\alpha_0 = 1.0$ and initialize weight coefficient $\alpha_0 = 1.0$.

Datasets and Baselines

7 Scenes (Shotton et al. 2013) is a dataset consisting of RGB-D images from seven different indoor scenes captured



PoseNet+

AtLoc (Ours)

Figure 4: Saliency maps of two scenes selected from Oxford RobotCar generated from models without attention (left: PoseNet+) and with attention (right: AtLoc). Note how AtLoc learns to ignore visually uninformative features e.g. the road in the top figure and instead focus on more distinctive objects e.g. the skyline in the distance. AtLoc also learns to reject affordable objects e.g. the bicycles in the bottom figure, yielding more robust global localization.

by a handheld Kinect RGB-D camera. The corresponding ground truth camera poses were calculated using KinectFusion. All images were captured in a small-scale indoor office environment at the resolution of 640×480 pixels. Each scene contains two to seven sequences in a single room for training/testing, with 500 or 1000 images for each sequence. As a popular dataset for visual relocalization, the sequences contained in this dataset were recorded under various camera motion status and different conditions, e.g. motion blur, perceptual aliasing and textureless features in the room.

Oxford RobotCar (Maddern et al. 2017) was recorded by an autonomous Nissan LEAF car in Oxford, UK over several periods for a year. This dataset exhibits substantial observations in the presence of various weather conditions, such as sunny and snowy days, as well as different lighting conditions, e.g., dim and glare roadworks. Moreover, we also found many dynamic or affordable objects in the scenes (e.g., parked/moving vehicles, cyclists and pedestrians), making this dataset particularly challenging for vision-based relocalization tasks. For a fair comparison, we follow the same evaluation strategy of MapNet (Brahmbhatt et al. 2018; Xue et al. 2019) and use two subsets of this dataset in our experiments, labelled as LOOP and FULL (length-based)

Sequence	PoseNet+		AtLoc (Ours)		MapNet		AtLoc+ (Ours)	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median
–								
LOOP1	25.29m, 17.45°	6.88m, 2.06°	8.61m, 4.58°	5.68m, 2.23°	8.76m, 3.46°	5.79m, 1.54°	7.82m, 3.62°	4.34m, 1.92°
LOOP2	28.81m, 19.62°	5.80m, 2.05°	8.86m, 4.67°	5.05m, 2.01°	9.84m, 3.96°	4.91m, 1.67°	7.24m, 3.60°	3.78m, 2.04°
FULL1	125.6m, 27.10°	107.6m, 22.5°	29.6m, 12.4°	11.1m, 5.28°	41.4m, 12.5°	17.94m, 6.68°	21.0m, 6.15°	6.40m, 1.50°
FULL2	131.1m, 26.05°	101.8m, 20.1°	48.2m, 11.1°	12.2m, 4.63°	59.3m, 14.8°	20.04m, 6.39°	42.6m, 9.95°	7.00m, 1.48°
Average	77.70m, 22.56°	55.52m, 11.7°	23.8m, 8.19°	8.54m, 3.54°	29.8m, 8.68°	12.17m, 4.07°	19.7m, 5.83°	5.38m, 1.74°

Table 4: Camera localization results on the LOOP and FULL of the Oxford RobotCar. For each sequence, we calculate the median and mean errors of position and rotation of Posenet+, MapNet and our approaches. Posenet and AtLoc leverage a single image while MapNet and AtLoc+ utilize sequential ones.

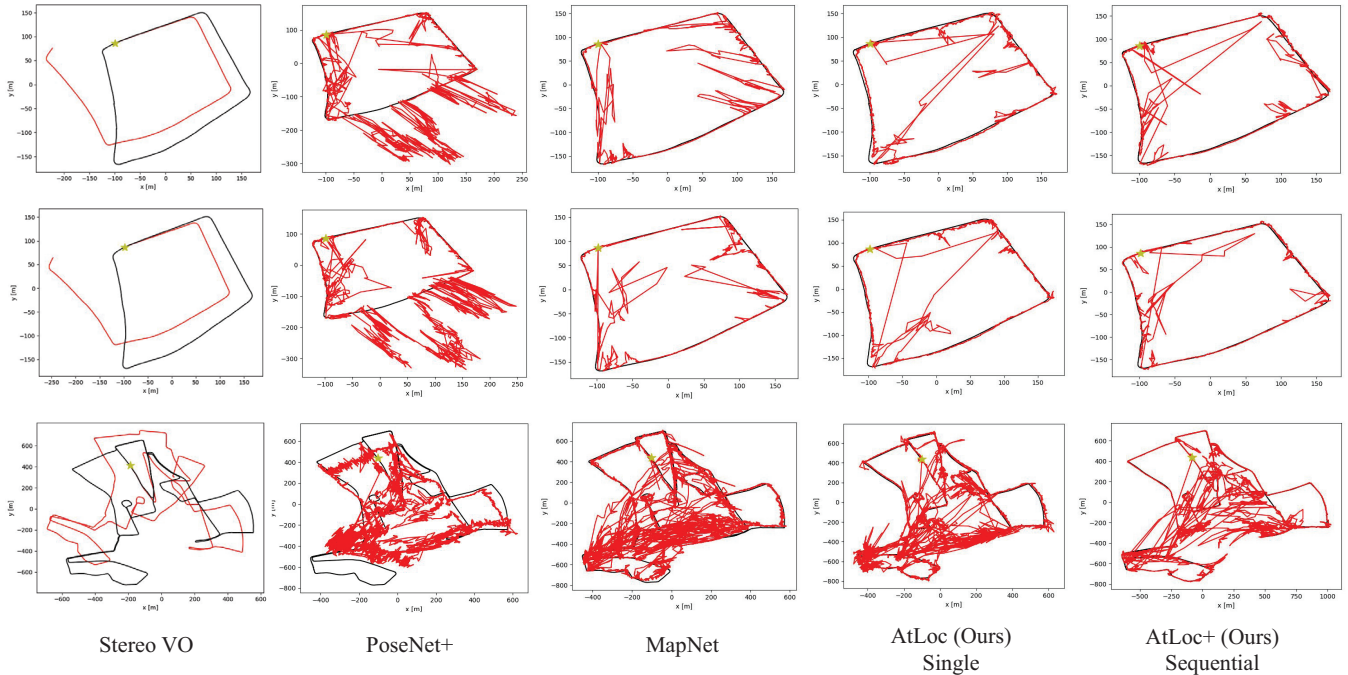


Figure 5: Trajectories on LOOP1 (top), LOOP2 (middle) and FULL1 (bottom) of Oxford RobotCar. The ground truth trajectories are shown in black lines while the red lines are the predictions. The star in the trajectory represents the starting point.

respectively. More details about these two sequences can be found in Table 2. In terms of implementation, we take the images recorded by the centre camera at a resolution of 1280×960 as the input to our network. The corresponding ground truth poses are labelled by the interpolations of INS measurements.

Baselines To validate the performance of our proposed network, we compare the results of several competing approaches. For experiments on 7 Scenes, we choose the following mainstream single-image-based methods: PoseNet (Kendall, Grimes, and Cipolla 2015), Bayesian PoseNet (Kendall and Cipolla 2016), PoseNet Spatial-LSTM (Walch et al. 2017), Hourglass (Melekhov et al. 2017) and PoseNet17 (Kendall and Cipolla 2017). Moreover, we also report the results of temporal approaches, VidLoc (Clark et al. 2017) and MapNet, for comparison. For the outdoor Oxford RobotCar dataset, Stereo VO (Maddern et al. 2017) and PoseNet+ (aka. ResNet34+log q) (Brahmbhatt et al. 2018) are selected

as our baselines. It is worth mentioning that PostNet+ is the best variant of PoseNet (Kendall, Grimes, and Cipolla 2015) on the RobotCar dataset (Brahmbhatt et al. 2018). Lastly, we also report the performance of MapNet (Brahmbhatt et al. 2018), the state-of-art method on this dataset using a sequence of images for relocalization. Note that as sequence based methods can exploit temporal constraints, they generally perform better than single-image based approaches. We nevertheless still compare with MapNet in evaluation to examine how accurate our single-image based AtLoc is.

Experiments on 7 Scenes

7Scenes Dataset contains 7 static indoor scenes with a large number of images captured in an office building. We take all scenes for comprehensive performance evaluation.

Quantitative Results Table 1 and Table 3 summarize the performance of all methods. Clearly, we can see that our method outperforms other single-image-based methods, with

a **13%** improvement in position accuracy and a **7%** improvement in rotation than the best single-image based baseline. In particular, AtLoc achieves the best performance gain in large texture-less (such as whiteboard) and highly texture-repetitive (such as stairs) scenarios. AtLoc reduces the position error from $0.35m$ to $0.26m$ and the rotation error from 12.4° to 10.5° in the scene of *Stairs*, which is a significant improvement over prior arts. In other regular scenes, AtLoc still reaches a comparable accuracy against baselines. By using only a single image, AtLoc achieves a superior accuracy compared with MapNet, despite the uses of image sequences and handcrafted geometric constraints in the MapNet design. Last but not least, after incorporating temporal constraints, AtLoc+ further narrows the median position and rotation errors to $0.19m$ and 7.08° respectively, outperforming MapNet by a large margin.

Qualitative Results To deeply understand the reasons behind these improvements, we visualize the attention maps of some scenes from 7Scenes. As shown in Figure 3, by using attention, AtLoc focuses more on geometrically meaningful areas (e.g. key points and lines) rather than feature-less regions and shows better consistency over time. In contrast, the saliency maps of PoseNet are relatively scattered and tend to focus on random regions in the view. A video that compares the saliency map between PoseNet and AtLoc in detail can be found at https://youtu.be/_xObJ1xwt94.

Localization Results on Oxford RobotCar

We next evaluate our approach on Oxford RobotCar dataset. Due to the substantial dynamics over the long collection period, this dataset is very challenging and strictly demands high robustness and adaptability of a relocalization model.

Quantitative Results Table 4 shows the comparison of our methods against PoseNet+ and MapNet. Compared with PoseNet+, AtLoc presents significant improvements on both LOOP trajectories and FULL trajectories. The mean position accuracy is improved from $25.29m$ to $8.61m$ on LOOP1, and $28.81m$ to $8.86m$ on LOOP2. The largest performance gains are observed on FULL1 and FULL2, where our approach outperforms PoseNet+ by 76.5% and 63.3%. When compared against the sequence-based MapNet, our AtLoc has obvious accuracy gain on all cases. Even for the unfavorable routes (FULL1 and FULL2), AtLoc still provides 28.5% and 30.3% improvements over MapNet. Furthermore, AtLoc+ significantly improves the position accuracy to $19.7m$ and the rotation accuracy to 5.83° after temporal constraints introduced.

Qualitative Results We now investigate why AtLoc significantly outperforms baselines on the Oxford RoboCar dataset. In Figure 5, we plot the predictions of LOOP1 (top), LOOP2 (middle) and FULL1 (bottom) by Stereo VO, Posenet+, MapNet, AtLoc and AtLoc+. Stereo VO is the official baseline from Oxford RobotCar. Although Stereo VO has very smooth predicted trajectories, it suffers from significant drifts as route length increases. Due to strong local similarity, there are

Scene	7Scenes			
	AtLoc Basic	AtLoc Basic+LSTM	AtLoc	AtLoc+
Chess	0.11m, 4.29	0.13m, 4.26	0.10m, 4.07	0.10m, 3.18
Fire	0.29m, 12.1	0.27m, 11.7	0.25m, 11.4	0.26m, 10.8
Heads	0.19m, 12.2	0.16m, 12.3	0.16m, 11.8	0.14m, 11.4
Office	0.19m, 6.35	0.20m, 5.74	0.17m, 5.34	0.17m, 5.16
Pumpkin	0.22m, 5.05	0.26m, 4.19	0.21m, 4.37	0.20m, 3.94
Kitchen	0.25m, 5.27	0.19m, 4.63	0.23m, 5.42	0.16m, 4.90
Stairs	0.30m, 11.3	0.29m, 12.1	0.26m, 10.5	0.29m, 10.2
Average	0.22m, 8.07	0.21m, 7.83	0.20m, 7.56	0.19m, 7.08
RobotCar				
LOOP1	25.29m, 17.45	29.71m, 15.72	8.61m, 4.58	7.82m, 3.62
LOOP2	28.81m, 19.62	32.79m, 17.76	8.86m, 4.67	7.24m, 3.60
FULL1	125.6m, 27.10	48.29m, 17.18	29.6m, 12.4	21.0m, 6.15
FULL2	131.1m, 26.05	67.62m, 11.40	48.2m, 11.1	42.6m, 9.95
Average	77.70m, 22.56	44.60m, 15.52	23.8m, 8.19	19.7m, 5.83

Table 5: Ablation study of AtLoc on 7 Scenes and Oxford RobotCar. AtLoc (Basic) denotes the model without using attention.

many outliers predicted by PoseNet+. These outliers, however, are significantly reduced in by AtLoc and AtLoc+. By looking into the saliency maps (Figure 4), we found PoseNet+ heavily relies on texture-less regions, such as local road surface (top), dynamic cars (middle) and affordance objects such as bicycles (bottom). These regions are either too similar in appearance or unreliable due to changes overtime, making pose estimation difficult. By contrast, our attention-guided AtLoc is able to automatically focus on unique, static and stable areas/objects, including vanishing lines and points (top), buildings (middle and bottom). These areas are tightly related to the latent geometric features of an environment, enabling robust pose estimation in the wild.

To further understand the efficacy of the attention mechanism, we depict the feature distances for a sequence of images. Specifically, we select a starting frame in the trajectory and then calculate feature distances (L_2) of subsequent frames to the starting frame. Features are extracted by PoseNet+ and AtLoc respectively, with the intention to understand to what extent the attention mechanism can help extract robust features. For experiments, we plot the distance profile under two cases: (i) dynamic vehicles and (ii) changing illumination. As we can see in Fig.6 (left), when the camera is static (i.e., the data-collection car is not moving), PoseNet+ is sensitive to dynamic objects entering the scene, resulting in a large variation of distances. In contrary, thanks to the adopted attention mechanism, AtLoc is robust to these moving vehicles and provides more stable features overall. Distance spikes are only observed when a large truck enters/leaves the scene, in which a substantial portion of view is blocked/revealed to the camera. On the right side of Fig.6, the extracted features of PoseNet+ suffer from illumination changes and gives abrupt shifts under different levels of glare. The features extracted by AtLoc, however, consistently change as the camera moves forward, agnostic to various lighting conditions.

Ablation Study and Efficiency Evaluation

We conduct an ablation study on the introduced attention module above 7 scenes and Oxford RoboCar datasets. In Table 5, AtLoc is compared with a basic version without

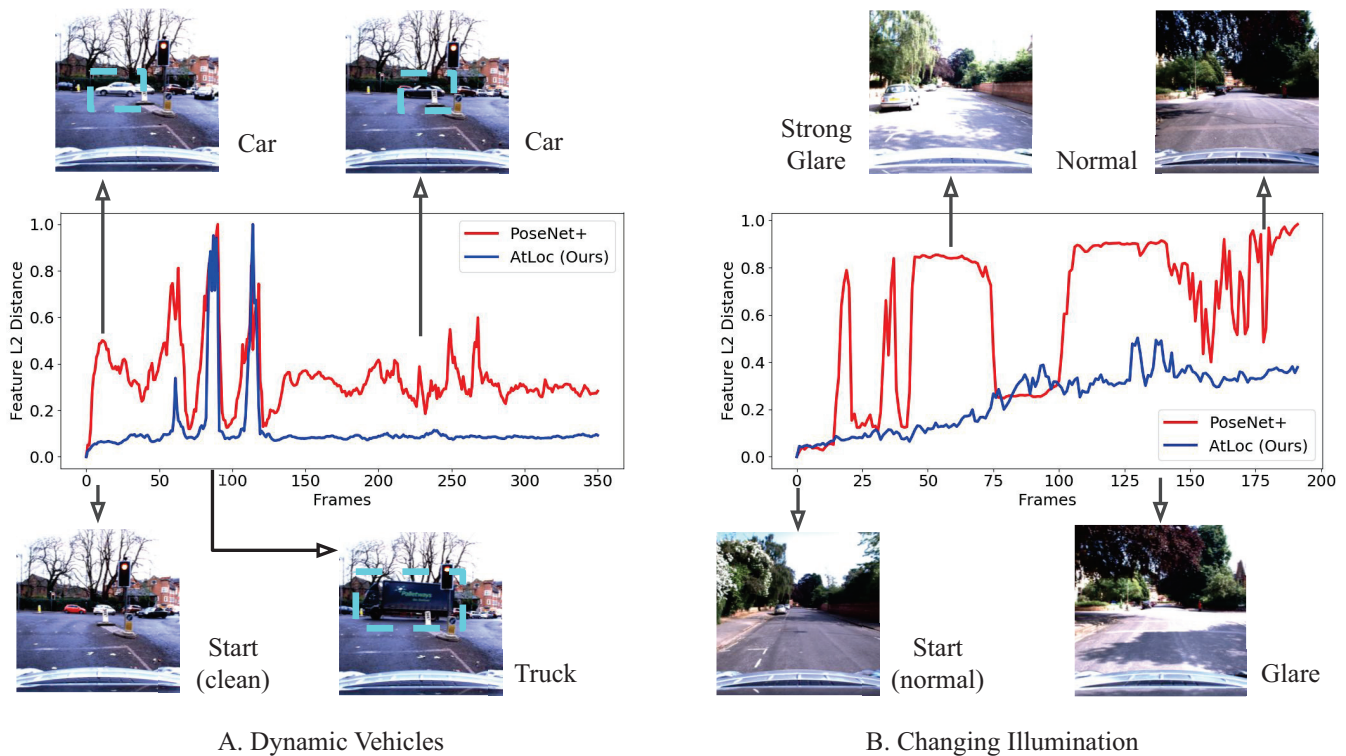


Figure 6: Feature distance comparisons under different dynamic disturbances. (a) Dynamic vehicles and (b) Changing illumination. Feature distances of AtLoc reasonably change with the motion status of the camera and are agnostic to various dynamics, while PoseNet suffers in both experiments.

the attention module, a LSTM version by replacing the attention module with the LSTM module, and a sequence enhanced version with temporal constraints. The rest modules are kept as the same for a fair comparison. The comparison of AtLoc (Basic) and AtLoc indicates that the model performance clearly increases on both datasets by adopting the self-attention into the pose regression model: it shows a 9% improvement in location accuracy and 6% in rotation accuracy on the 7 Scenes dataset; AtLoc achieves an average localization accuracy of $23.8m$ and an average rotation accuracy of only 8.19° on Oxford RobotCar dataset. The LSTM based AtLoc is only better than Basic AtLoc. With temporal constraints, AtLoc+ yields the best performance, obtaining a median error of $0.19m$, 7.08° on 7 Scenes, and $19.7m$, 5.83° on RobotCar.

To evaluate the efficiency of our proposed AtLoc, we analyze the average running time of three models - MapNet, PoseLSTM and AtLoc. Among the three models, MapNet consumes the longest running time of $9.4ms$ per frame, as it needs to process additional data from other sensory inputs and a sequence of images to apply geometric constraints. Due to the time-consuming recursive operations in LSTMs, PoseLSTM takes a running time of $9.2ms$ per frame, $3.7ms$ higher than its corresponding basic model PoseNet. In contrast, our proposed AtLoc achieves an ideal balance between the computational efficiency and localization accuracy, consuming only $6.3ms$ per frame while obtaining the best localization

performance.

Conclusion and Discussion

Camera localization is a challenging task in computer vision due to scene dynamics and high variability of environment appearance. In this work, we presented a novel study of self-attention guided camera localization from a single image. The introduced self-attention can encourage the framework to learn geometrically robust features, mitigating the impacts from dynamic objects and changing illumination. We demonstrate state-of-the-art results, even surpassing sequential based techniques in challenging scenarios. Further work includes refining the attention module and determining whether it can improve multi-frame camera pose regression.

References

- Arandjelovic, R.; Gronat, P.; Torii, A.; Pajdla, T.; and Sivic, J. 2016. Netvlad: Cnn architecture for weakly supervised place recognition. In *CVPR*.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Brachmann, E.; Krull, A.; Nowozin, S.; Shotton, J.; Michel, F.; Gumhold, S.; and Rother, C. 2017. Dsac-differentiable ransac for camera localization. In *CVPR*.
- Brahmbhatt, S.; Gu, J.; Kim, K.; Hays, J.; and Kautz, J. 2018.

- Geometry-aware learning of maps for camera localization. In *CVPR*.
- Cai, M.; Shen, C.; and Reid, I. D. 2018. A hybrid probabilistic model for camera relocalization. In *BMVC*.
- Chen, D. M.; Baatz, G.; Köser, K.; Tsai, S. S.; Vedantham, R.; Pylvänäinen, T.; Roimela, K.; Chen, X.; Bach, J.; Pollefeys, M.; et al. 2011. City-scale landmark identification on mobile devices. In *CVPR*.
- Chen, C.; Rosa, S.; Miao, Y.; Lu, C. X.; Wu, W.; Markham, A.; and Trigoni, N. 2019. Selective sensor fusion for neural visual-inertial odometry. In *CVPR*.
- Cheng, J.; Dong, L.; and Lapata, M. 2016. Long short-term memory-networks for machine reading. In *EMNLP*.
- Clark, R.; Wang, S.; Markham, A.; Trigoni, N.; and Wen, H. 2017. Vidloc: A deep spatio-temporal model for 6-dof video-clip relocalization. In *CVPR*.
- Deng, J.; Dong, W.; Socher, R.; Li, L.; Kai Li; and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*.
- Dou, Z.-Y.; Tu, Z.; Wang, X.; Wang, L.; Shi, S.; and Zhang, T. 2019. Dynamic layer aggregation for neural machine translation with routing-by-agreement. In *AAAI*.
- Fu, Y.; Wang, X.; Wei, Y.; and Huang, T. 2019. Sta: Spatial-temporal attention for large-scale video-based person re-identification. In *AAAI*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Huang, Z.; Xu, Y.; Shi, J.; Zhou, X.; Bao, H.; and Zhang, G. 2019. Prior guided dropout for robust visual localization in dynamic environments. In *ICCV*.
- Kendall, A., and Cipolla, R. 2016. Modelling uncertainty in deep learning for camera relocalization. In *ICRA*.
- Kendall, A., and Cipolla, R. 2017. Geometric loss functions for camera pose regression with deep learning. In *CVPR*.
- Kendall, A.; Grimes, M.; and Cipolla, R. 2015. Posenet: A convolutional network for real-time 6-dof camera relocalization. In *CVPR*.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. In *ICLR*.
- Kingma, D. P., and Dhariwal, P. 2018. Glow: Generative flow with invertible 1x1 convolutions. In *NIPS*.
- Li, Y.; Snavely, N.; Huttenlocher, D.; and Fua, P. 2012. Worldwide pose estimation using 3d point clouds. In *ECCV*.
- Liu, A.; Liu, X.; Fan, J.; Ma, Y.; Zhang, A.; Xie, H.; and Tao, D. 2019. Perceptual-sensitive gan for generating adversarial patches. In *AAAI*.
- Liu, L.; Li, H.; and Dai, Y. 2017. Efficient global 2d-3d matching for camera localization in a large-scale 3d map. In *ICCV*.
- Maddern, W.; Pascoe, G.; Linegar, C.; and Newman, P. 2017. 1 year, 1000 km: The oxford robotcar dataset. *The International Journal of Robotics Research* 36(1):3–15.
- Melekhov, I.; Ylioinas, J.; Kannala, J.; and Rahtu, E. 2017. Image-based localization using hourglass networks. In *ICCV*.
- Parisotto, E.; Singh Chplot, D.; Zhang, J.; and Salakhutdinov, R. 2018. Global pose estimation with an attention-based recurrent network. In *CVPR Workshops*.
- Parmar, N.; Vaswani, A.; Uszkoreit, J.; Kaiser, Ł.; Shazeer, N.; Ku, A.; and Tran, D. 2018. Image transformer. In *ICML*.
- Purkait, P.; Zhao, C.; and Zach, C. 2018. Synthetic view generation for absolute pose regression and image synthesis. In *BMVC*.
- Sattler, T.; Zhou, Q.; Pollefeys, M.; and Leal-Taixe, L. 2019. Understanding the limitations of cnn-based absolute camera pose regression. In *CVPR*.
- Sattler, T.; Leibe, B.; and Kobbelt, L. 2012. Improving image-based localization by active correspondence search. In *ECCV*.
- Shotton, J.; Glocker, B.; Zach, C.; Izadi, S.; Criminisi, A.; and Fitzgibbon, A. 2013. Scene coordinate regression forests for camera relocalization in rgb-d images. In *CVPR*, 2930–2937.
- Szegedy, C.; Liu, W.; Jia, Y.; Sermanet, P.; Reed, S.; Anguelov, D.; Erhan, D.; Vanhoucke, V.; and Rabinovich, A. 2015. Going deeper with convolutions. In *CVPR*.
- Torii, A.; Sivic, J.; Pajdla, T.; and Okutomi, M. 2013. Visual place recognition with repetitive structures. In *CVPR*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *NIPS*.
- Walch, F.; Hazirbas, C.; Leal-Taixe, L.; Sattler, T.; Hilsenbeck, S.; and Cremers, D. 2017. Image-based localization using lstms for structured feature correlation. In *CVPR*.
- Wang, S.; Clark, R.; Wen, H.; and Trigoni, N. 2018a. End-to-end, sequence-to-sequence probabilistic visual odometry through deep neural networks. *The International Journal of Robotics Research* 37(4-5):513–542.
- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018b. Non-local neural networks. In *CVPR*.
- Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A.; Salakhutdinov, R.; Zemel, R.; and Bengio, Y. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*.
- Xue, F.; Wang, X.; Yan, Z.; Wang, Q.; Wang, J.; and Zha, H. 2019. Local supports global: Deep camera relocalization with sequence enhancement. In *ICCV*.
- Yang, B.; Li, J.; Wong, D. F.; Chao, L. S.; Wang, X.; and Tu, Z. 2019. Context-aware self-attention networks. In *AAAI*.
- Yuan, Y.; Mei, T.; and Zhu, W. 2019. To find where you talk: Temporal sentence localization in video with attention based location regression. In *AAAI*.
- Zhang, H.; Goodfellow, I.; Metaxas, D.; and Odena, A. 2018. Self-attention generative adversarial networks. In *ICML*.