Weakly-supervised Camera Localization by Ground-to-satellite Image Registration

Yujiao Shi¹⁽⁶⁾, Hongdong Li²⁽⁶⁾, Akhil Perincherry³⁽⁶⁾, and Ankit Vora³⁽⁶⁾

¹ ShanghaiTech University, China
 ² The Australian National University, Australia

 ³ Ford Motor Company, USA
 shiyj2@shanghaitech.edu.cn

Abstract. The ground-to-satellite image matching/retrieval was initially proposed for city-scale ground camera localization. This work addresses the problem of improving camera pose accuracy by ground-tosatellite image matching after a coarse location and orientation have been obtained, either from the city-scale retrieval or from consumer-level GPS and compass sensors. Existing learning-based methods for solving this task require accurate GPS labels of ground images for network training. However, obtaining such accurate GPS labels is difficult, often requiring an expensive Real Time Kinematics (RTK) setup and suffering from signal occlusion, multi-path signal disruptions, etc. To alleviate this issue, this paper proposes a weakly supervised learning strategy for ground-tosatellite image registration when only noisy pose labels for ground images are available for network training. It derives positive and negative satellite images for each ground image and leverages contrastive learning to learn feature representations for ground and satellite images useful for translation estimation. We also propose a self-supervision strategy for cross-view image relative rotation estimation, which trains the network by creating pseudo query and reference image pairs. Experimental results show that our weakly supervised learning strategy achieves the best performance on cross-area evaluation compared to recent state-of-the-art methods that are reliant on accurate pose labels for supervision.

Keywords: Ground-to-satellite image matching \cdot Cross-view image matching \cdot Weakly-supervised camera localization

1 Introduction

Camera localization is pivotal in real-world applications such as autonomous driving, field robotics, and augmented/virtual reality. Recent research has explored various methods to approximate the coarse location and orientation of a ground camera, such as visual retrieval techniques, consumer-level GPS and compass, *etc.* To attain greater pose accuracy, other sensors like Lidar [6,21,41], Radar [37,38], and High Definition (HD) maps, have been investigated. However, many commercial autonomous vehicles at level two/three lack these sensors.



Fig. 1: We derive positive and negative satellite images for each ground-view image based on its coarse location information. A similarity map between the ground view and each satellite image is computed when aligned at different locations. Our training objective aims to maximize the maximum similarity in the positive similarity map while minimizing the maximum similarity in the negative similarity map.

Maintaining and updating high-precision HD maps [18,52] is challenging and expensive. In response, satellite imagery has emerged as an inexpensive alternative reference source due to its wide accessibility and global coverage.

We focus on ground-to-satellite camera localization, aiming to determine a ground camera's location and orientation relative to a geo-tagged satellite map. Prior research [11, 17, 25, 28, 39, 40, 43, 46, 54, 57] has centered on city-scale camera localization, employing image retrieval to match ground and satellite images. However, image retrieval introduces errors that can span tens of meters. Recent efforts have addressed this by enhancing camera pose accuracy through ground-to-satellite image registration, guided by coarse location and orientation estimates. Nonetheless, the significant viewpoint differences between ground and satellite images make handcrafted features fail [31]. Learning-based approaches [7,14,27,50,59] require a large training dataset with accurate ground truth (GT) poses for the ground images. This often requires field surveys or the use of specialized equipment, which can be costly and time-consuming, limiting the availability of GT data, particularly for large-scale or extensive image datasets. Furthermore, even the expensive high-accuracy RTK GPS signals can be affected by factors such as multipath interference, signal blockage, and atmospheric conditions, leading to inaccuracies in location data [8,20]. Therefore, this paper aims to develop a weakly supervised ground-to-satellite image registration strategy to increase the ground cameras' pose accuracy when only coarse pose labels for ground images are provided.

We target the 3-DoF (degree-of-freedom) pose estimation for ground cameras, *i.e.*, 2-DoF location and 1-DoF orientation. Under the weakly-supervised setting, deterministic pose output by a network is suboptimal, as no GT pose is available to supervise the network output. We resort to the recent representation learning approaches by contrastive learning to solve this problem.

Given the coarse pose of a ground camera, we determine a satellite image that covers the local surroundings of the camera and some satellite images that do not cover the local surroundings, which are regarded as positive and negative for this ground image, respectively. We utilize the signal that a ground image is within its positive satellite image while outside its negative satellite images to train a network. The network is trained to learn feature representations such that the similarity between the ground and its positive satellite images at their optimal relative pose is more significant than between the ground and its negative satellite images at their "optimal" relative pose.

The "optimal" relative pose for a ground-and-satellite image pair is determined by the maximum cross-view similarity when aligned at this relative pose (among others). We follow the standard synthesis-and-matching procedure to determine this "optimal" relative pose: first synthesizing an overhead-view feature map from the ground view image and then matching it against the reference satellite feature map. Conventional methods [7, 26] for this matching process rotate and translate the synthesized overhead view feature map according to predetermined candidate poses. Considering the vast search space for the 3-DoF pose, we decouple the rotation and translation estimation and develop different supervision strategies for them, respectively, following Shi et al. [29].

The rotation estimation is conducted first, and the framework is designed by network regression. It takes input as the query and reference images and outputs their relative pose. To supervise this network, we create some "satellite and satellite" image pairs with GT relative poses. Specifically, we randomly rotate and translate a satellite image. The transformed satellite image mimics a synthesized overhead view image from a ground image. We train the network to estimate the relative pose between the original and transformed satellite images. Once the network is trained, it is leveraged to estimate the relative orientation between a query ground image and its positive satellite image.

After this, our translation estimation framework synthesizes an overheadview feature map from the ground image with the orientation aligned with satellite images and then matches it against a satellite feature map. The output is a similarity (/location probability) map of the ground image with respect to the satellite image. This process is implemented between a query image and its positive and negative satellite images, as illustrated in Fig. 1. A contrastive learning supervision strategy is employed to train the network by maximizing the maximum similarity in the positive similarity map while minimizing the maximum similarity in the negative similarity map.

We conduct experiments on KITTI [8], a popular autonomous driving dataset where ground images are captured by a pin-hole camera with limited fieldof-view (FoV), and VIGOR [59], a well-known cross-view localization dataset where ground images are panoramas. Experimental results demonstrate that our method achieves the best generalization ability compared to the recent state-ofthe-art despite not requiring accurate pose labels for supervision.

Our contributions are summarized as follows:

- We propose a self-supervised relative rotation estimation strategy between ground-and-satellite image pairs;
- We introduce a weakly-supervised translation estimation strategy for groundto-satellite camera pose estimation;
- Our method achieves the best cross-area generalization ability than the previous fully supervised state-of-the-art that relies on GT labels for training.

4 Shi et al.



2 Related Work

2.1 Cross-view image-based localization

City-scale localization. Ground-to-satellite image-based localization aims to determine the location of a ground camera by matching it with a satellite map covering the region of interest. Initially proposed for city-scale coarse-level localization, it was formulated as an image retrieval problem. Specifically, a large satellite map of the interested region is first split into small patches to construct a geo-referenced satellite image database. For a query image captured at the ground level, its similarity with every database satellite image is computed, and the GPS of the most similar satellite image is taken as the query camera's location. Over the past decades, hand-crafted features [3, 15, 22] have been demonstrated to be a bottleneck for cross-view feature matching due to significant geometric and appearance variations. Seminal works using deep networks [40,46,47] demonstrated that the learned feature descriptors by a metric learning training objective offer better and more reliable performance. Researchers have since investigated learning powerful feature descriptors [2, 4, 53, 57], orientation-aware cross-view representations [11,17,36,58], and different strategies for bridging the cross-view domain gap [25, 28, 32, 39, 54]. Cross-view localization has also been extended from a 2-DoF location estimation task to a 3-DoF joint location and orientation estimation task [30], and from a single image-based localization problem to a video-based localization problem [33, 42, 56]. However, the database images are often discretely sampled, while query images' locations are continuous. Thus, the image retrieval formulation results in a coarse camera pose estimation. Its accuracy is determined by the sample density of ground images.

Increasing localization accuracy. Recently, researchers have started to investigate how to increase the accuracy of a ground camera's pose by ground-to-satellite image matching once the camera's coarse rotation and orientation have been determined. Towards this purpose, network regression [59], pose optimization [27], similarity matching [7,14,26,29,31,49,50], correspondence learning [34], and iterative tomography estimation [44] have been explored. Nonetheless, all these works need sub-meter and sub-degree pose labels for ground images in the training data to train their networks. In this paper, we propose a strategy that

estimates the relative rotation and translation between a ground and a satellite image when such accurate labels are unavailable in the training data.

2.2 Self/weakly-supervised learning

Self- and weakly-supervised learning has been widely explored in image classification [55], object detection [5], semantic segmentation [45], image inpainting [23], point cloud registration [16], human/hand/object pose estimation [1,9,35], the intersection between vision and natural language processing [24] *etc.* Many of them also exploit contrastive learning as supervision signals. Tang *et al.* [37] proposed a self-supervised strategy for localizing outdoor range sensors (*e.g.*, Lidar, Radar). However, little attention has been paid to camera pose estimation with self or weak supervision. Our concurrent work [51] also solves the problem of weakly supervised ground-to-satellite camera localization, but it highlights how to adapt a model trained with GT supervision from one area to another area where only coarse supervision is available. In contrast, this paper does not assume any GT data from any area.

3 Method

Given a coarse location and orientation of a ground camera, our goal is to refine this pose through ground-to-satellite image registration. In contrast to previous works that assume a large training dataset with highly accurate pose labels for ground images, we propose a weakly supervised learning strategy that does not require such labels. Our approach first estimates the relative rotation by network regression and then computes the translation by similarity matching.

3.1 Rotation Estimation by Self-supervision

We draw inspiration from Spatial Transformer Networks [12] to regress relative rotations between cross-view images by a network. The network takes query and reference images as input and outputs their relative pose, as illustrated in Fig. 2.

Training. To supervise the network, we generate some "satellite-and-satellite" image pairs with GT relative poses. This is done by rotating and translating a satellite image using a randomly generated pose, $\mathbf{R}^*(\theta), \mathbf{t}^*$, where the maximum magnitude of the rotation angle θ and translation is based on the statistical error of the ground camera's coarse pose information that we aim to refine during deployment. For example, when the coarse location is achieved by ground-tosatellite image retrieval, the minimum statistical location error (assuming the top-1 retrieval recall rate is 100%) is the half distance between centers of two nearest satellite images, *i.e.*, the half grid size when sampling satellite images for constructing the database [31,33,59]. Furthermore, we apply a mask on the transformed satellite image and extract a triangle region corresponding to the ground camera's horizontal FoV, as shown at the top of Fig. 2a. This is intended to mimic a synthesized overhead view image from a ground-view image. We refer

6 Shi et al.

to the transformed and masked satellite image as the query and the original one as the reference. We then train the network to estimate the relative pose between the two satellite images. The training objective is:

$$\mathcal{L}_{1} = |\theta - \theta^{*}| + |t_{x} - t_{x}^{*}| + |t_{y} - t_{y}^{*}|, \qquad (1)$$

where θ, t_x and t_y denote the network predictions, θ^*, t_x^* and t_y^* indicate the corresponding GT, θ is the 1-DoF inplane rotation, *i.e.*, the yaw angle, t_x and t_y represent the 2-DoF translation, and $|\cdot|$ denote the L_1 norm.

Evaluation. After the network is trained, we substitute the query satellite image with the query ground images during inference for the ground camera's pose estimation, as shown in Fig. 2b. We found the feature extractors for the satellite images and ground images captured by a perspective pin-hole camera are sharable. A potential explanation might be that the projection geometry of them is similar. They both map straight lines in the real world to straight lines in images. We demonstrate this by experiments in the main paper and supplementary material. As the pose regressor accepts query and reference feature maps in the overhead view, we follow previous works [29, 33] to project ground-view feature maps to overhead view by exploiting the ground plane Homography. This projection method does not require additional trainable parameters and thus allows our model trained between satellite-and-satellite images deployable between satellite-and-ground images. Furthermore, since the ground-plane Homohraphy establishes correct correspondences for scene contents on the ground plane between the two views, which occupies a large portion of the field-of-view of the ground camera in autonomous driving scenarios, our pose regressor handles well when the query feature map is the projected one from the ground view.

Our pose regressor shares the same pros and cons as the neural pose optimizer in Shi et al. [29]. It leverages the sensitiveness of general neural networks (without specific designs) to rotations on input signals [13, 48], and thus achieves promising results on rotation estimation. However, the translation estimation performance is poor due to the aggregation layers, *e.g.*, max-pooling, inside the pose regressor, making it insensitive to slight translations on input signals. To recover accurate relative translations, we leverage the translational equivariance of spatial correlation (a.k.a. convolution), which is described below.

3.2 Translation Estimation by Deep Metric Learning

As shown in Fig. 3, a two-branch convolutional network is first applied to the ground and satellite image pair. Each branch is a U-Net architecture and extracts multi-level representations of the original images.

For the ground branch, not only the feature representation $\mathbf{F}_g \in \mathbb{R}^{H_g \times W_g \times C}$ but also a confidence map $\mathbf{C}_g \in \mathbb{R}^{H_g \times W_g \times 1}$ is extracted. The confidence map is an additional channel of the feature extractor output followed by a sigmoid layer. It indicates whether features at corresponding spatial pixel positions are trustworthy. For example, dynamic objects (*e.g.*, cars) in the images are detrimental to localization performance, while the static road structures are essential



Fig. 3: Translation estimation framework. During training, the weights of the rotation estimator are fixed, and we only train the two-branch feature extractors. Other uncolored blocks indicate no trainable parameters are involved.

features. The higher the confidence, the more reliable the corresponding features. We should note that no explicit supervision is applied to the confidence map. Instead, it is encoded in the cross-view similarity-matching process and learned statistically from the similarity-matching training objective.

We only extract feature representations $\mathbf{F}_s \in \mathbb{R}^{H_s \times W_s \times C}$ for the satellite branch with no confidence map because we empirically found learning confidence maps for satellite images impairs the performance. The reasons hypothesized are twofold: (i) dynamic objects are fewer on the satellite images, and they occupy a relatively smaller region on the satellite image compared to ground-view images; thus, they have a lower impact on localization performance; (ii) neither the camera poses, nor the confidence maps, have explicit supervision, thus learning a confidence map for satellite images increases the network training difficulty.

We leverage the trained rotation estimator in Sec. 3.1 to estimate the relative rotation between the ground and satellite images. Then, the ground-view features and confidence maps are projected to the overhead view according to the estimated rotation \mathbf{R} and zero translation $\mathbf{t} = \mathbf{0}$, exploiting the ground plane Homography. To maintain claritys, we utilize the original symbols \mathbf{F}_g and \mathbf{C}_g to represent the projected ground features and confidence maps.

Confidence-guided similarity matching. Given that the rotation of the projected overhead-view feature map has been aligned with the observed satellite feature map, the only remaining disparity between them is a translation difference. We explore the normalized spatial correlation to compute this translation difference. Specifically, the projected overhead-view feature map is used as a sliding window, and its normalized inner product with the reference satellite feature map is computed when aligned at varying locations. The mathematical representation, taking into account ground-view confidence maps, is as follows:

$$\mathbf{S}(u,v) = (\mathbf{F}_s * \widehat{\mathbf{F}_g})(u,v) = \frac{\sum_i \sum_j \mathbf{F}_s(u+i,v+j) \widehat{\mathbf{F}_g}(i,j)}{\sqrt{\sum_i \sum_j \mathbf{F}_s^2(u+i,v+j)} \sqrt{\sum_i \sum_j \widehat{\mathbf{F}_g}^2(i,j)}}, \quad (2)$$

where $\widehat{\mathbf{F}}_g = \mathbf{C}_g \mathbf{F}_g$ is to highlight important features while suppressing nonreliable features for localization. The result $\mathbf{S}(u, v)$ is actually the cosine similarity between the two view features when aligned at each location (u, v). Thus, the pixel coordinate corresponding to the maximum similarity, $\arg \max_{(u,v)} \mathbf{S}$, indicates the most likely ground camera's location. This similarity map between the synthesized overhead-view feature map and the reference satellite feature map can also be regarded as the location probability map of the ground camera.

Supervision. We apply deep metric learning for network supervision. For a query ground image, we compute a similarity map based on the possible locations between it and its positive and negative satellite images, denoted as \mathbf{S}_{pos} and \mathbf{S}_{neg} , respectively. We maximize the maximum similarity in \mathbf{S}_{pos} while minimizing the maximum similarity in \mathbf{S}_{neg} :

$$\mathcal{L}_2 = \sum_l \log(1 + e^{\alpha(\max \mathbf{S}_{\text{neg}} - \max \mathbf{S}_{\text{pos}})}), \qquad (3)$$

where α controls the convergence speed and is set to 10.

When the positive image pairs in the training dataset are created in the same way as image pairs during inference, for example, by the same ground-to-satellite retrieval model or the same noisy GPS sensor, the error of location labels in the training set is the same as the error of locations that we aim to refine during deployment. In this case, we only employ Eq. (3) for network training. In another scenario where relatively more accurate location labels for ground images in the training data are available than the poses we aim to refine during employment, we introduce an additional training objective to incorporate this signal:

$$\mathcal{L}_3 = \sum_{l} \left| \max(\mathbf{S}_{\text{pos}}) - \max(\mathbf{S}_{\text{pos}}[u^* - \frac{d}{\gamma} : u^* + \frac{d}{\gamma}, v^* - \frac{d}{\gamma} : v^* + \frac{d}{\gamma}]) \right|, \quad (4)$$

where (u^*, v^*) indicates the location label provided by the training data and has an error of up to d meters, γ denotes the ground resolution of the similarity map in terms of meters per pixel. This training objective forces the global maximum in the similarity map to equal a local maximum, with the local region centered at the location label with a radius of d meters.

The whole training objective is:

$$\mathcal{L} = \mathcal{L}_2 + \lambda \mathcal{L}_3,\tag{5}$$

where $\lambda = 0$ indicates such relatively accurate pose labels are unavailable in the training set, while $\lambda = 1$ suggests such labels are available. In our experiments, we set d = 5 meters.

Overall Evaluation. After training the networks, the overall evaluation goes through the framework shown in Fig. 3. The input is a query image and its positive satellite image, and the output is an estimated relative rotation between the two images and a location probability map of this query image with respect to this satellite image. The location corresponding to the maximum probability/similarity value is deemed as the query camera location.

4 Experiments

Network Architectures. A UNet-based architecture with a pre-trained VGG16 as the encoder is adopted for feature extraction. The decoder of the UNet is randomly initialized. We empirically found that the feature extractor of satellite images is shareable with ground images captured by a pin-hole camera while not shareable with ground panoramas. This might be because both satellite images and ground images captured by a pin-hole camera map straight lines in the real world to straight lines on images, while panoramas map straight lines in the real world to curves on images. While for different purposes, *i.e.*, rotation and translation estimation, we found non-shareable feature extractors between the two stages help to achieve the best performance. We present detailed analysis and experimental demonstrations in the supplementary material. The pose regressor in Fig. 2 is constructed by two swin transformer layers [19] followed by two FC layers. This is the same as the neural optimizer architecture in [29].

Dataset. Our experiments are conducted on a well-known autonomous driving dataset, KITTI [8], and a cross-view localization dataset, VIGOR [59].

For the KITTI dataset, ground images were captured by a forward-facing pin-hole camera with a limited FoV. The cross-view KITTI dataset includes one training set and two testing sets. Images from Test-1 are from the same region as the training set, while images from Test-2 are from a different region. The location search range for this dataset is around $56 \times 56 \text{ m}^2$, and the orientation noise is 20°, which follows the official setting as in [27]. The performance evaluation on different coarse pose errors is presented in the supplementary material.

The VIGOR dataset contains ground and satellite images from four cities in the US: Chicago, New York, San Francisco, and Seattle. It is divided into same-area and cross-area splits. The same-area split indicates the training and testing images are from the same region (both from the four cities), while the cross-area split adopts images from two cities for training and images from the other two cities for testing. In the original dataset, each ground image has a positive satellite image and several semi-positive satellite images, depending on whether this ground image is within the center $\frac{1}{4}$ region of the satellite image. We follow [14] and use only the positive satellite images.

Since the feature extractors trained for satellite images (Fig. 2a) are not applicable to panoramas due to differences in imaging modality, we only evaluate the translation estimation performance on the VIGOR dataset with known and unknown (360° ambiguity) orientations, respectively. The 3-DoF joint location and orientation pose estimation is performed on the KITTI dataset.

Evaluation Metrics. Following the approach of [27], we decompose the translation along the lateral (orthogonal to driving direction) and longitudinal (along the driving direction) directions for evaluation on the KITTI dataset. Specifically, for a query image, we consider it to be successfully localized along a direction if its estimated location along that direction is within d meters of its ground truth (GT) location. Similarly, we consider the rotation estimation correct if the estimated rotation is within θ° of the GT rotation. We record the



Fig. 4: Visualization of learned confidence maps for ground images and estimated poses by our method.

percentage of successfully localized images along each direction and the percentage of images with correct rotation estimation.

The VIGOR dataset does not provide information on the driving direction of the camera. Therefore, we cannot decompose the translation to be along lateral and longitudinal directions. On this dataset, we report the median and mean errors of comparison algorithms, following the approach of [50] and [14].

Data augmentation and implementation Details. For training the rotation estimator, we use the positive satellite image for each ground image and randomly transform each image once to create training input pairs. We adopt feature size as $\frac{1}{4}$ of the original image size for both rotation and translation estimation. The original image size is not used because of its large memory consumption in the spatial correlation process. We use a batch size of B = 8to train the network. In the translation estimation training, each ground image has one matching satellite image and B-1 non-matching satellite images within each batch. Our experiments are conducted on an RTX 3090 GPU. The network is trained for 3 epochs for both stages on the KITTI dataset and 10 epochs on the VIGOR dataset. For the ground images in the KITTI dataset, we use a resolution of 256×1024 . For the VIGOR dataset, we use a resolution of 320×640 . The satellite image resolution is 512×512 for all datasets. The ground resolution of satellite images is 0.2 meters per pixel in the KITTI dataset, and 0.111, 0.113, 0.118, and 0.101 meters per pixel for the cities Chicago, New York, San Francisco, and Seattle in the VIGOR dataset, respectively. The source code of this paper is available at https://github.com/YujiaoShi/G2SWeakly.git.

4.1 Ablation Study

Below, we demonstrate the necessity and effectiveness of each proposed component. The results are presented in Tab. 1.

(i) The first row shows the performance of our pose regressor (Reg.) in Fig. 2 with the satellite images as queries (note: not ground images). It can be seen that the estimated rotation of almost all the queries has been restricted to within

Table 1: Ablation study on the cross-view KITTI dataset. We report the percentage of query images whose locations are restricted to be within d meters of their GT locations along lateral/longitudinal directions and whose orientations are restricted to be within θ° of their GT orientation, respectively.

													. –		
Query	t Est.	Con	fλ	$\begin{array}{c} \text{Lat} \\ d = 1 \uparrow \end{array}$	$\begin{array}{c} {\rm Te}\\ {\rm eral}\\ d=3\uparrow \end{array}$	est-1 (Sa Longit $d = 1 \uparrow$	ame-are udinal $d = 3 \uparrow$	$\left \begin{array}{c} \mathrm{Azir} \\ \mathrm{Azir} \\ \theta = 1 \uparrow \end{array} \right.$	$\begin{array}{l} \text{muth} \\ \widehat{\theta} = 3 \uparrow \end{array}$	$\begin{array}{c} \text{Lat} \\ d = 1 \uparrow \end{array}$	$\begin{array}{c} {\rm Te}\\ {\rm eral}\\ d=3\uparrow \end{array}$	st-2 (C Longit $d = 1 \uparrow$	ross-are udinal $d = 3 \uparrow$	$\left \begin{array}{c} \operatorname{Azin} \\ \theta = 1 \end{array} \right $	$\theta = 3 \uparrow$
Sat Grd Grd Grd Grd	Reg. Reg. Corr. Corr. Corr.	- N Y Y	$ \begin{bmatrix} - \\ - \\ 0 \\ 0 \\ 1 $	4.88 4.88 45.80 59.58 66.07	15.13 15.11 78.27 85.74 94.22	4.80 4.74 6.18 11.37 16.51	15.64 15.66 16.67 31.94 49.96	99.89 99.66 99.66 99.66 99.66	100.00 100.00 100.00 100.00 100.00	5.06 5.04 45.11 62.73 64.74	15.46 15.46 73.04 86.53 86.18	5.25 5.29 6.13 9.98 11.81	15.79 15.83 18.30 29.67 34.77	99.95 99.99 99.99 99.99 99.99	100.00 100.00 100.00 100.00 100.00

Table 2: Comparison between different overview synthesis methods on KITTI.

Ove Fe	rhead-view $\begin{vmatrix} \lambda \\ \text{Lat} \end{vmatrix}$	Test-1 (Sateral Longit $d = 3 \uparrow d = 1 \uparrow$	ame-area) udinal Azin $d = 3 \uparrow \theta = 1 \uparrow$	muth $\begin{bmatrix} L\\ \theta = 3 \uparrow d = 1 \end{bmatrix}$	$\begin{array}{c} \text{Tes} \\ \text{ateral} \\ \uparrow d = 3 \uparrow d \end{array}$	t-2 (Cross-a Longitudina $d = 1 \uparrow d = 3$	$ \begin{array}{l} \text{rea}) \\ 1 & \text{Azin} \\ \uparrow \theta = 1 \uparrow \end{array} $	$\hat{\theta} = 3 \uparrow$
Geo.	$\begin{array}{c c} {\rm Trans.} \ [29] \\ {\rm Ours} \end{array} \Big \begin{smallmatrix} 0 \\ 59.58 \\ \end{smallmatrix} \Big \begin{array}{c} 48.74 \\ 59.58 \\ \end{array} \Big $	$\begin{array}{c c c} 83.14 & 8.27 \\ 85.74 & 11.37 \end{array}$	$\begin{array}{c c c} 24.20 & 99.66 \\ 31.94 & 99.66 \end{array}$	$\begin{array}{c c} 100.00 & 52.6 \\ 100.00 & 62.7 \end{array}$	9 79.99 3 86.53	8.46 24.00 9.98 29.67) 99.99 7 99.99	$100.00 \\ 100.00$
Geo.	$\begin{array}{c c} {\rm Trans.} & [29] \\ {\rm Ours} & 1 \end{array} \begin{vmatrix} 62.02 \\ 66.07 \end{vmatrix}$	$\begin{array}{c c c} 92.13 & 16.01 \\ 94.22 & 16.51 \end{array}$	$\begin{array}{c c c c c c c c c c c c c c c c c c c $	$\begin{array}{c c} 100.00 & 62.9 \\ 100.00 & 64.7 \end{array}$	$\begin{array}{c ccc}7 & 87.54 \\ 4 & 86.18 \end{array}$	$\begin{array}{rrrr} 11.36 & 33.70 \\ 11.81 & 34.77 \end{array}$) 99.99 7 99.99	$100.00 \\ 100.00$

1° of its GT rotation, while the translation estimation performance is poor, even though there is no domain gap between reference and query images. This supports our intuition that a neural network-based regressor's ability to provide accurate translation estimation is limited.

(ii) Then, we verify the generalization ability of the pose regressor on real ground images. The results in the second row show that the rotation estimation accuracy for real ground images is very close to that for satellite images, demonstrating the effectiveness of our self-supervision training strategy by satellite-and-satellite image pairs.

(iii) The third row presents the performance of our method by using spatial correlation for translation estimation and pose regressor for rotation estimation. The performance of translation estimation is significantly boosted.

(iv) In what follows, we encode the confidence in the spatial correlation process. The results in the fourth row demonstrate the effectiveness of the confidence-guided similarity matching. Fig. 4 provides visualizations of the learned confidence maps and the probability maps of query cameras' location with respect to their matching satellite images. It can be seen that the learned confidence maps are able to ignore dynamic objects and highlight reliable features (e.g., lane lines, road edges).

(v) Finally, when relatively accurate (but still noisy) pose labels are available in the training dataset, *i.e.*, the pose error is up to d = 5 meters, we set $\lambda =$ 1 and use Eq. 4 to incorporate this information in training. The last row of Tab. 1 demonstrates that it successfully boosts the network learning process and improves the performance on both same-area and cross-area evaluation.

4.2 Different overhead-view projections

In this section, we ablate the performance of our method with different overheadview projections. Among existing works, OrienterNet [26] and Shi et al. [29] both propose geometry-guided cross-view transformers (Geo. Trans.) for overheadview feature synthesis. Shi et al. [29] achieves slightly better performance, as illustrated in Tab. 3. Thus, we evaluate the performance of the Geo. Trans. from Shi et al. [29] in our framework. From the comparison results in Tab. 2, it can be seen the performance of Geo. Trans is inferior to the simple ground plane Homography-based overhead-view projection under the weakly-supervised setting. We hypothesize two main reasons for this:

(i) Smaller batch size. In this paper, we leverage deep metric learning for network supervision. It is well-recognized that a small batch size in metric learning affects performance negatively. SimpleBEV [10] also confirms that a larger batch size improves the overall performance more significantly compared to different overhead-view projection methods. With our simple Homography-based projection method, we achieved a batch size of 8 on an RTX 3090 GPU, whereas the batch size is 4 when using the Geo. Trans. due to its complexity.

(ii) Weak supervision. The previous geometry-guided cross-view transformers are usually constructed with complex designs to make the learned overhead view feature representations invariant to parallel, occlusion, and appearance changes. This works well when the supervision signal is strong, *i.e.*, GT pose labels are available. However, in our scenario, the pose labels are noisy, and only weak supervision is applied. This limits the learning ability of geometry-guided cross-view transformers. Instead, our method explores the explicit ground plane Homography for the overhead-view feature projection. It requires no additional trainable parameters and reduces the learning burden of deep networks. The ground plane Homography handles the co-visible regions for the overhead-view feature synthesis, *i.e.*, scenes on the ground plane. For not co-visible regions, our confidence map lowers their weights in the similarity matching process for camera localization.

We should note that we do not want to claim our solution is optimal for overhead-view synthesis in the weakly-supervised ground-to-satellite camera localization task. We believe there should be more advanced approaches that make the best balance between capability, complexity, and the easiness of training under the weak supervision scenario. The main purpose of this paper is to introduce training strategies for pose estimation under the weak supervision scenario. In order to not dilute our contributions, we leave the exploration of a better overhead-view feature synthesis module as our future work.

Visualization of learned overhead-view features. We visualize the synthesized overhead view feature maps from ground images in Fig. 5 (b). The center of the synthesized overhead view feature map corresponds to the ground camera location, and the right direction corresponds to the heading direction (z-axis) of the ground camera. The corresponding satellite image features are presented in Fig. 5 (c). Using our framework, we register the query feature map Fig. 5 (b) to the reference feature map Fig. 5 (c). The resulting camera



Fig. 5: (a) Ground image; (b) Synthesized overhead view feature map from the ground image; (c) Satellite image feature map; (d) Estimated location probability map and camera orientation indicated by the red arrow.

Table 3: Comparison with the state-of-the-art on KITTI. * indicates full supervision is adopted.

Algorithms	$\begin{array}{c} \text{Lat} \\ d = 1 \uparrow \end{array}$	$\begin{array}{c} \mathrm{Te}\\ \mathrm{eral}\\ \mathrm{d}=3\uparrow \end{array}$	st-1 (S Longit $d = 1 \uparrow$	ame-are to the second	$\left \begin{array}{c} \mathrm{Azin} \\ \mathrm{Azin} \\ \theta = 1 \end{array} \right $	$\begin{array}{l} \text{muth} \\ \uparrow \theta = 3 \uparrow \end{array}$	$Lat \\ d = 1 \uparrow$	Tes eral $d = 3 \uparrow$	t-2 (C Longit $d = 1 \uparrow$	ross-ar udinal $d = 3 \uparrow$	$\theta = 1 \uparrow$	$\hat{\theta} = 3 \uparrow$
DSM [30]*	10.12	30.67	4.08	12.01	3.58	13.81	10.77	31.37	3.87	11.73	3.53	14.09
CVR [59]*	18.61	49.06	4.29	13.01	-	-	17.38	48.20	4.07	12.52	-	-
Shi and Li $[27]^*$	35.54	70.77	5.22	15.88	19.64	51.76	27.82	59.79	5.75	16.36	18.42	49.72
SliceMatch [14]*	49.09	91.76	14.19	49.99	13.41	42.62	32.43	78.98	8.30	24.48	46.82	46.82
OrienterNet [26]*	-	-	-	-	-	-	51.26	84.77	22.39	46.79	20.41	52.24
Xia et al. [49] *	97.35	98.65	77.13	96.08	77.39	99.47	44.06	81.72	23.08	52.85	57.72	92.34
Shi et al. [29]*	76.44	96.34	23.54	50.57	99.10	100.00	57.72	86.77	14.15	34.59	98.98	100.00
Song et al. [34]*	95.47	_	87.89	_	89.40	-	54.19	—	23.10	—	43.44	_
Ours $(\lambda = 0)$	59.58	85.74	11.37	31.94	99.66	100.00	62.73	86.53	9.98	29.67	99.99	100.00
Ours $(\lambda = 1)$	66.07	94.22	16.51	49.96	99.66	100.00	64.74	86.18	11.81	34.77	99.99	100.00

orientation and location probability maps are shown in Fig. 5 (d). To reduce ambiguity, we only consider the scene content within 20m of the query camera location for localization, marked as red circles in Fig. 5. It can be seen that the scene contents on the ground plane are recovered faithfully. A visual illustration of this spatial correlation process is provided in the supplementary material.

4.3 Comparison with the state-of-the-art

In this section, we compare the performance of our method with the recent state-of-the-art, including DSM [30], CVR [59], [27], MCC [50], SliceMatch [14], OrienterNet [26], Xia *et al.* [49], Shi *et al.* [29], and Song *et al.* [34]. All these state-of-the-art algorithms adopt full supervision and rely on accurate pose labels for their network, and their results are either taken from their original papers or re-evaluated using the author-released models.

The performance comparison between our method and the state-of-the-art on the KITTI dataset is presented in Tab. 3. It can be seen that among all the comparison algorithms, our method achieves the best cross-area evaluation performance, and the performance discrepancy between same-area and crossarea evaluation of our method is the smallest. This is because our method has

14 Shi et al.

Algorithms	Aligned Mean↓	Sar -orientatio Median↓	ne-area nUnknow Mean↓	n-orientatio Median↓	nAligned Mean↓	Cross-area Aligned-orientationUnknown-orientation Mean↓ Median↓ Median↓				
CVR [59]*	8.99	7.81	-	_	8.89	7.73	-	_		
MCC [50]*	6.94	3.64	9.87	6.25	9.05	5.14	12.66	9.55		
SliceMatch [14]*	5.18	2.58	8.41	5.07	5.53	2.55	8.48	5.64		
Shi et al. [29]*	4.12	1.34	-	_	5.16	1.40	_	_		
Xia et al. [49]*	3.60	1.36	3.74	1.42	4.97	1.68	5.41	1.89		
Song et al. [34]*	3.03	0.97	4.97	1.90	5.01	2.42	7.67	3.67		
$\begin{array}{l} \mathbf{Ours} \ (\lambda=0) \\ \mathbf{Ours} \ (\lambda=1) \end{array}$	5.22 4.19	1.97 1.68	5.33 4.18	$2.09 \\ 1.66$	5.37 4.70	1.93 1.68	5.37 4.52	1.93 1.65		

Table 4: Comparison with the state-of-the-art on VIGOR. * indicates full supervision is adopted.

no information on GT poses, and it is trained to leverage similarity matching for location estimation, preventing itself from overfitting on the GT poses. The performance of our method on the same-area evaluation is close to Shi *et al.* [29] and slightly inferior to three recent state-of-the-art methods. However, in the supplementary material, we demonstrate that the performance of our method on both same-area and cross-area evaluation can be potentially improved when more training image pairs (with pose errors up to tens of meters) are available. Our method does not rely on accurate pose labels of training data. This leads to significant cost and effort savings by eliminating the need for high-precision pose label acquisition.

The performance comparison between our method and the state-of-the-art on the VIGOR dataset is presented in Tab. 4. The training images in the VIGOR dataset are about $4.5 \times$ larger than those in the KITTI dataset, and our method achieves comparable performance with [29] on same-area evaluation. Furthermore, similar to the observations on the KITTI dataset, our method's generalization ability from same-area to cross-area is also the best on the VIGOR dataset, with the cross-area evaluation performance surpassing almost all fully supervised methods. Due to the space limit, we provide more discussions in the supplementary material.

5 Conclusion

This paper has introduced the first weakly supervised ground camera pose refinement strategy by ground-to-satellite image registration. Given a coarse location and orientation of a ground camera obtained from noisy sensors or visual retrieval techniques, our method is able to refine this pose by ground-to-satellite image registration using a training dataset without accurate pose labels for ground images. Key components of our approach include a training scheme for groundand-satellite rotation alignment using satellite-and-satellite image pairs and a deep metric learning supervision mechanism that trains the network for translation estimation. Benefiting from these two innovations, our method, without requiring accurate labels, achieves comparable or superior performance to the recent fully supervised state-of-the-art.

Acknowledgements

This research is funded in part by an ARC Discovery Grant DP220100800.

References

- Bouazizi, A., Wiederer, J., Kressel, U., Belagiannis, V.: Self-supervised 3d human pose estimation with multiple-view geometry. In: 2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021). pp. 1–8. IEEE (2021)
- Cai, S., Guo, Y., Khan, S., Hu, J., Wen, G.: Ground-to-aerial image geo-localization with a hard exemplar reweighting triplet loss. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)
- Castaldo, F., Zamir, A., Angst, R., Palmieri, F., Savarese, S.: Semantic cross-view matching. In: Proceedings of the IEEE International Conference on Computer Vision Workshops. pp. 9–17 (2015)
- Chu, M., Zheng, Z., Ji, W., Wang, T., Chua, T.S.: Towards natural language-guided drones: Geotext-1652 benchmark with spatial relation matching. In: Proceedings of the European Conference on Computer Vision (ECCV) (2024)
- Dang, T., Kornblith, S., Nguyen, H.T., Chin, P., Khademi, M.: A study on selfsupervised object detection pretraining. In: Computer Vision–ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IV. pp. 86–99. Springer (2023)
- De Paula Veronese, L., de Aguiar, E., Nascimento, R.C., Guivant, J., Auat Cheein, F.A., De Souza, A.F., Oliveira-Santos, T.: Re-emission and satellite aerial maps applied to vehicle localization on urban environments. In: 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 4285–4290 (2015). https://doi.org/10.1109/IROS.2015.7353984
- Fervers, F., Bullinger, S., Bodensteiner, C., Arens, M., Stiefelhagen, R.: Uncertainty-aware vision-based metric cross-view geolocalization. arXiv preprint arXiv:2211.12145 (2022)
- 8. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. International Journal of Robotics Research (IJRR) (2013)
- 9. Gharaee, Z., Lawin, F.J., Forssén, P.E.: Self-supervised learning of object pose estimation using keypoint prediction. arXiv preprint arXiv:2302.07360 (2023)
- Harley, A.W., Fang, Z., Li, J., Ambrus, R., Fragkiadaki, K.: Simple-bev: What really matters for multi-sensor bev perception? In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 2759–2765. IEEE (2023)
- 11. Hu, S., Feng, M., Nguyen, R.M.H., Hee Lee, G.: Cvm-net: Cross-view matching network for image-based ground-to-aerial geo-localization. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018)
- Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: Advances in neural information processing systems. pp. 2017–2025 (2015)
- Lee, J., Kim, B., Kim, S., Cho, M.: Learning rotation-equivariant features for visual correspondence. In: CVPR. pp. 21887–21897 (June 2023)
- Lentsch, T., Xia, Z., Caesar, H., Kooij, J.F.: Slicematch: Geometry-guided aggregation for cross-view pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17225–17234 (2023)

- 16 Shi et al.
- Lin, T.Y., Belongie, S., Hays, J.: Cross-view image geolocalization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 891–898 (2013)
- Liu, D., Chen, C., Xu, C., Qiu, R.C., Chu, L.: Self-supervised point cloud registration with deep versatile descriptors for intelligent driving. IEEE Transactions on Intelligent Transportation Systems (2023)
- 17. Liu, L., Li, H.: Lending orientation to neural networks for cross-view geolocalization. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
- Liu, R., Wang, J., Zhang, B.: High definition map for automated driving: Overview and analysis. The Journal of Navigation 73(2), 324–341 (2020)
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10012–10022 (2021)
- Maddern, W., Pascoe, G., Linegar, C., Newman, P.: 1 year, 1000 km: The oxford robotcar dataset. The International Journal of Robotics Research 36(1), 3–15 (2017)
- Mishra, S., Parchami, A., Corona, E., Chakravarty, P., Vora, A., Parikh, D., Pandey, G.: Localization of a smart infrastructure fisheye camera in a prior map for autonomous vehicles. In: 2022 International Conference on Robotics and Automation (ICRA). pp. 5998–6004 (2022). https://doi.org/10.1109/ICRA46639. 2022.9811793
- Mousavian, A., Kosecka, J.: Semantic image based geolocation given a map. arXiv preprint arXiv:1609.00278 (2016)
- Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2536–2544 (2016)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
- Regmi, K., Shah, M.: Bridging the domain gap for ground-to-aerial image matching. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)
- 26. Sarlin, P.E., DeTone, D., Yang, T.Y., Avetisyan, A., Straub, J., Malisiewicz, T., Bulò, S.R., Newcombe, R., Kontschieder, P., Balntas, V.: Orienternet: Visual localization in 2d public maps with neural matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21632–21642 (2023)
- Shi, Y., Li, H.: Beyond cross-view image retrieval: Highly accurate vehicle localization using satellite image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17010–17020 (2022)
- Shi, Y., Liu, L., Yu, X., Li, H.: Spatial-aware feature aggregation for image based cross-view geo-localization. In: Advances in Neural Information Processing Systems. pp. 10090–10100 (2019)
- 29. Shi, Y., Wu, F., Perincherry, A., Vora, A., Li, H.: Boosting 3-dof ground-to-satellite camera localization accuracy via geometry-guided cross-view transformer. arXiv preprint arXiv:2307.08015 (2023)
- 30. Shi, Y., Yu, X., Campbell, D., Li, H.: Where am I looking at? joint location and orientation estimation by cross-view matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4064–4072 (2020)

- Shi, Y., Yu, X., Liu, L., Campbell, D., Koniusz, P., Li, H.: Accurate 3-dof camera geo-localization via ground-to-satellite image matching. IEEE Transactions on Pattern Analysis and Machine Intelligence (2022)
- Shi, Y., Yu, X., Liu, L., Zhang, T., Li, H.: Optimal feature transport for cross-view image geo-localization. In: AAAI. pp. 11990–11997 (2020)
- Shi, Y., Yu, X., Wang, S., Li, H.: Cvlnet: Cross-view semantic correspondence learning for video-based camera localization. In: Computer Vision-ACCV 2022: 16th Asian Conference on Computer Vision, Macao, China, December 4–8, 2022, Proceedings, Part I. pp. 123–141. Springer (2023)
- Song, Z., Lu, J., Shi, Y., et al.: Learning dense flow field for highly-accurate crossview camera localization. Advances in Neural Information Processing Systems 36 (2024)
- Spurr, A., Dahiya, A., Wang, X., Zhang, X., Hilliges, O.: Self-supervised 3d hand pose estimation from monocular rgb via contrastive learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11230–11239 (2021)
- Sun, B., Chen, C., Zhu, Y., Jiang, J.: Geocapsnet: Ground to aerial view image geo-localization using capsule network. In: 2019 IEEE International Conference on Multimedia and Expo (ICME). pp. 742–747. IEEE (2019)
- 37. Tang, T.Y., De Martini, D., Wu, S., Newman, P.: Self-supervised learning for using overhead imagery as maps in outdoor range sensor localization. The International Journal of Robotics Research 40(12-14), 1488–1509 (2021)
- Tang, T.Y., De Martini, D., Barnes, D., Newman, P.: Rsl-net: Localising in satellite images from a radar on the ground. IEEE Robotics and Automation Letters 5(2), 1087–1094 (2020)
- 39. Toker, A., Zhou, Q., Maximov, M., Leal-Taixé, L.: Coming down to earth: Satelliteto-street view synthesis for geo-localization. CVPR (2021)
- Vo, N.N., Hays, J.: Localizing and orienting street views using overhead imagery. In: European Conference on Computer Vision. pp. 494–509. Springer (2016)
- 41. Vora, A., Agarwal, S., Pandey, G., McBride, J.: Aerial imagery based lidar localization for autonomous vehicles. arXiv preprint arXiv:2003.11192 (2020)
- Vyas, S., Chen, C., Shah, M.: Gama: Cross-view video geo-localization. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVII. pp. 440–456. Springer (2022)
- 43. Wang, C., Zheng, Z., Ruijie, Q., Yang, Y.: Depth-aware blind image decomposition for real-world adverse weather recovery. In: Proceedings of the European Conference on Computer Vision (ECCV) (2024)
- 44. Wang, X., Xu, R., Cui, Z., Wan, Z., Zhang, Y.: Fine-grained cross-view geolocalization using a correlation-aware homography estimator. Advances in Neural Information Processing Systems 36 (2024)
- 45. Wang, Y., Zhuo, W., Li, Y., Wang, Z., Ju, Q., Zhu, W.: Fully self-supervised learning for semantic segmentation. arXiv preprint arXiv:2202.11981 (2022)
- 46. Workman, S., Jacobs, N.: On the location dependence of convolutional neural network features. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 70–78 (2015)
- Workman, S., Souvenir, R., Jacobs, N.: Wide-area image geolocalization with aerial reference imagery. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3961–3969 (2015)
- Worrall, D.E., Garbin, S.J., Turmukhambetov, D., Brostow, G.J.: Harmonic networks: Deep translation and rotation equivariance. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5028–5037 (2017)

- 18 Shi et al.
- Xia, Z., Booij, O., Kooij, J.F.: Convolutional cross-view pose estimation. arXiv preprint arXiv:2303.05915 (2023)
- Xia, Z., Booij, O., Manfredi, M., Kooij, J.F.: Visual cross-view metric localization with dense uncertainty estimates. In: European Conference on Computer Vision. pp. 90–106. Springer (2022)
- 51. Xia, Z., Shi, Y., Li, H., Kooij, J.F.: Adapting fine-grained cross-view localization to areas without fine ground truth. In: Proceedings of the European Conference on Computer Vision (ECCV) (2024)
- Xiao, Z., Yang, D., Wen, T., Jiang, K., Yan, R.: Monocular localization with vector hd map (mlvhm): A low-cost method for commercial ivs. Sensors 20(7), 1870 (2020)
- Yang, H., Lu, X., Zhu, Y.: Cross-view geo-localization with layer-to-layer transformer. Advances in Neural Information Processing Systems 34, 29009–29020 (2021)
- Zhai, M., Bessinger, Z., Workman, S., Jacobs, N.: Predicting ground-level scene layout from aerial imagery. In: IEEE Conference on Computer Vision and Pattern Recognition. vol. 3 (2017)
- Zhai, X., Oliver, A., Kolesnikov, A., Beyer, L.: S4l: Self-supervised semi-supervised learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (October 2019)
- Zhang, X., Sultani, W., Wshah, S.: Cross-view image sequence geo-localization. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2914–2923 (2023)
- 57. Zhu, S., Shah, M., Chen, C.: Transgeo: Transformer is all you need for cross-view image geo-localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1162–1171 (June 2022)
- Zhu, S., Yang, T., Chen, C.: Revisiting street-to-aerial view image geo-localization and orientation estimation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 756–765 (2021)
- 59. Zhu, S., Yang, T., Chen, C.: Vigor: Cross-view image geo-localization beyond oneto-one retrieval. CVPR (2021)