

Accurate 3-DoF Camera Geo-Localization via Ground-to-Satellite Image Matching

Yujiao Shi^{ID}, Xin Yu^{ID}, Liu Liu, Dylan Campbell^{ID}, Piotr Koniusz^{ID}, and Hongdong Li^{ID}

Abstract—We address the problem of ground-to-satellite image geo-localization, that is, estimating the camera latitude, longitude and orientation (azimuth angle) by matching a query image captured at the ground level against a large-scale database with geotagged satellite images. Our prior arts treat the above task as pure image retrieval by selecting the most similar satellite reference image matching the ground-level query image. However, such an approach often produces coarse location estimates because the geotag of the retrieved satellite image only corresponds to the image center while the ground camera can be located at any point within the image. To further consolidate our prior research finding, we present a novel geometry-aware geo-localization method. Our new method is able to achieve the fine-grained location of a query image, up to pixel size precision of the satellite image, once its coarse location and orientation have been determined. Moreover, we propose a new geometry-aware image retrieval pipeline to improve the coarse localization accuracy. Apart from a polar transform in our conference work, this new pipeline also maps satellite image pixels to the ground-level plane in the ground-view via a geometry-constrained projective transform to emphasize informative regions, such as road structures, for cross-view geo-localization. Extensive quantitative and qualitative experiments demonstrate the effectiveness of our newly proposed framework. We also significantly improve the performance of coarse localization results compared to the state-of-the-art in terms of location recalls.

Index Terms—Camera geo-localization, cross-view matching, street-view, satellite imagery, geotagging

1 INTRODUCTION

GIVEN an image captured by a camera at the ground level in some large open space, estimating the camera position and the direction it faces is a useful but also challenging problem. This paper addresses the problem of ground-to-satellite image geo-localization which aims to determine the geographical location and azimuth angle of a query image by matching it against a large geo-tagged satellite map covering the region. Due to the accessibility and extensive coverage of satellite imagery, the problem of ground-to-satellite image alignment has been recently noted by researchers as it can feature in a number of computer vision applications e.g., autonomous driving, robot navigation, and way-finding in augmented/virtual reality.

- Yujiao Shi and Hongdong Li are with ANU, Australian National University, Canberra, ACT 2601, Australia. E-mail: {yujiao.shi, hongdong.li}@anu.edu.au.
- Xin Yu is with Australian Artificial Intelligence Institute, University of Technology Sydney, Ultimo, NSW 2007, Australia. E-mail: xin.yu@uts.edu.au.
- Liu Liu is with Cyberverse Lab, Huawei, Shenzhen 518129, China. E-mail: liu.liu@anu.edu.au.
- Dylan Campbell is with the University of Oxford, OX1 2JD Oxford, U.K. E-mail: djcampbell.01@gmail.com.
- Piotr Koniusz is with Machine Learning Research Group (MLRG), Data61/CSIRO (NICTA), Canberra, ACT 2601, Australia, and also with Australian National University, Canberra, ACT 2601, Australia. E-mail: piotr.koniusz@anu.edu.au.

Manuscript received 12 January 2021; revised 25 August 2021; accepted 19 June 2022. Date of publication 11 July 2022; date of current version 3 February 2023. This work was supported in part by ARC-Discovery under Grants DP190102261 and DP220100800. The work of Hongdong Li was supported by a Ford Alliance Grant. The work of Yujiao Shi was supported by China Scholarship Council (CSC). The work of Dylan Campbell was supported by Continental AG. (Corresponding author: Yujiao Shi.)

Recommended for acceptance by J. Hays.

Digital Object Identifier no. 10.1109/TPAMI.2022.3189702

Conventional methods for solving this task often formulate the problem as *image retrieval* (e.g., [1], [2], [3], [4], [5], [6], [7], [8], [9], [10]). They do so by first partitioning a large satellite map to smaller satellite images, to form the reference database. Once a ground-view query image is provided, they compare the query image with the database images to retrieve the most similar one, as shown in Fig. 1a. The gps-tag of the matched satellite image is then used to approximate the location of the query image location. However, oftentimes the geo-tag of the retrieved satellite image corresponds to the centre of the retrieved satellite image from the database, while the true ground camera location can be rather off. Therefore, camera locations estimated by these methods are rather coarse and inaccurate.

To address this problem, we introduce a new two-stage mechanism for *accurate 3-DoF (latitude, longitude, and azimuth angle) camera geo-localization* in this paper. First, we estimate a coarse camera location by searching the most similar satellite image from the database (Fig. 1a). Subsequently, we compute the displacement between the center of the retrieved satellite image and the inquired camera location, achieving fine-grained localization results (Fig. 1b). The orientation alignment (azimuth angle) between the ground and satellite images is estimated in both steps.

At the coarse camera localization stage, we extend the method proposed in our previous work [10]. Apart from a polar transform (Fig. 2b) that roughly bridges the cross-view domain gap, we develop a projective transform in this work to establish geometrically constrained correspondences (Fig. 2c) between the satellite (Fig. 2a) and ground-level (Fig. 2d) images for scene objects on the ground plane. Both the polar-transformed and projective-transformed satellite images are used in our coarse localization pipeline. The former preserves all details from the original satellite

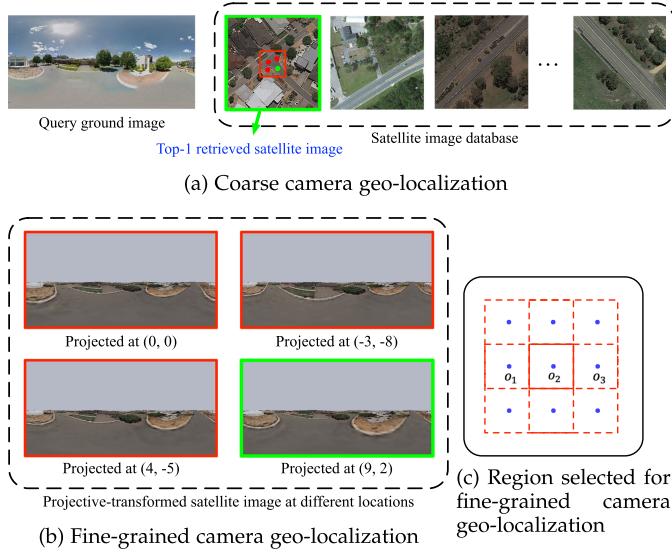


Fig. 1. (a) Given a query ground image, we first retrieve its most similar satellite image from the database. (b) Then, we use a set of candidate locations in the satellite image as projection centers. The fine-grained location of the query ground image is then achieved from the projected satellite image that is most similar to the query image. The estimated orientation is obtained by comparing the selected projected image and the query image. (c) The black box represents a large satellite map covering the whole region, from which the small satellite images in the database (shown in (a)) are cropped for coarse camera localization. The blue dots denote the centers of those cropped images. The red boxes indicate the regions selected for fine-grained camera localization, which cover nearly the entire satellite map.

image while the latter exhibits better visual similarity to the captured ground-level scenes. Then, we employ CNNs to learn feature correspondences between ground-level and transformed satellite images. After satellite images are projected to the corresponding pixels in the ground-level coordinate system, the spatial layout gap between ground and satellite images is significantly reduced. Following our prior works [10], we opt to extract feature volumes as our global descriptors to encode discriminative spatial information.

We note that for both the polar and projective transforms, the horizontal axis corresponds to the azimuth direction. We thus propose a Dynamic Similarity Matching (DSM) module to estimate the orientation of ground images with respect to satellite images. Specifically, DSM computes the correlation between the ground and satellite features in order to generate a similarity score at each angle, denoted by the red curve in Fig. 4. The argument of the similarity score maximum corresponds to the latent orientation of the ground image with respect to the satellite image. If the ground image has restricted FoV, we then extract the appropriate local region from the satellite feature representation for the use in the coarse localization stage. The output of our coarse localization

stage is the satellite image that is the most similar to the query image in the database.

Considering the displacement between the ground camera and the satellite image center, this article further introduces a fine-grained camera localization stage to localize a fine-grained position of the query ground image. Specifically, we project the satellite image to the ground viewpoint at a predetermined set of points of projection, as shown in Fig. 1a. The similarity between the projective-transformed satellite images and the query image is then computed in the same way as in the coarse matching stage. The center of projection of the projective-transformed satellite image that is most similar to the query ground image is taken as the camera location, and the computed relative orientation is taken as the camera azimuth angle. In particular, the precision of the localization result depends on the density of the sampled centers of projection, whereas the precision of the orientation estimation depends on the resolution of the images and the FoV.

Below, we detail novel contributions of this manuscript which are not explored in our earlier work.

- i. Compared to previous methods which formulate the cross-view image-based geo-localization as a pure image retrieval task, this article introduces a new fine-grained localization method to compute the displacement between the retrieved satellite image center and the query ground camera location.
- ii. We further extend our conference work [10] for coarse camera localization. Apart from the polar transform introduced in our conference work, this article proposes a geometry-constrained projective transform to establish more realistic geometric correspondences of points on the ground plane between satellite and ground-level images. This allows our network to focus on informative regions for cross-view image matching.
- iii. We achieve new state-of-the-art cross-view localization performance compared to our conference work. Moreover, we analyze and discuss the different task properties for localizing restricted FoV and panorama images and shed some light on how to exploit the proposed method for different localization situations.

2 RELATED WORK

3D Structure-Based Localization. The task of visual localization is to estimate the 6-DOF camera pose of a query image with respect to a 3D scene model [11], [12], [13], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25]. The key to this problem is to establish efficient, accurate, and robust



Fig. 2. Given a satellite image (a), we explore two transforms, i.e., polar transform (b) and the projective transform (c), to align it to its corresponding ground-view panorama (d).

2D-3D matches. This line of work can predict accurate 6-DoF camera poses. Nevertheless, 3D models are not available everywhere, and they are usually expensive to obtain. This limits the applicability of these methods.

2D Ground-to-Ground Image-Based Localization. Image-based localization aims to estimate the camera pose of a query image by matching it against a large geo-tagged database. It was originally approached as a ground-to-ground image matching task, which is often used for place recognition [26], [27], [28], [29], [30], [31], [32], [33], [34], and loop-closure detection [35], [36], [37]. In the ground-to-ground image-based localization task, both the query and database images are captured at ground level. The challenges are to address the difficulties of large viewpoint differences, illumination differences (e.g., day and night), and weather differences between query and reference images. However, ground-to-ground image matching cannot localize query images where no corresponding reference image is available, since the world is non-uniformly sampled by tourists and ground-level vehicles.

2D Ground-to-Satellite Image-Based Localization. Many recent works [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [38], [39], [40], [41], [42], [43], [44], [45] resort to satellite images as a reference set for image-based camera localization, due to the wide-spread coverage and easy accessibility of satellite imagery. Challenges of ground-to-satellite image matching include the significant visual appearance differences, geometric projection differences, and the unknown relative orientation between the two view images, as well as the limited FoV of query ground images. Existing works have focused on designing powerful network architectures [1], [3], [4], [7], [41], [43], bridging the cross-view domain gaps [6], [8], [9], [10], [45], and learning orientation invariant or equivariant features [4], [5], [10], [43], [44].

Although promising results have been achieved, almost all the approaches only estimate the location (latitude and longitude) of a query image but neglect the orientation misalignment. Our prior work [10] is the first attempt to estimate the 3-DOF camera pose (location and estimation) via ground-to-satellite image matching. However, all prior arts (including ours) only retrieve the most similar satellite image for a query image. As a result, the estimates of camera location are quite coarse. In this article, we provide a new mechanism for estimating the fine-grained location and orientation of the query ground image subsequent to the coarse localization stage.

3 IMAGE RETRIEVAL FOR COARSE CAMERA GEO-LOCALIZATION

The approach proposed in this article involves a coarse localization stage in which image retrieval techniques are used to roughly estimate where an image was taken, and a fine-grained localization stage where the displacement between the query ground camera and the center of the retrieved satellite image is computed. In this section, we outline the extended framework for coarse camera geo-localization in Fig. 4.

For the cross-view geo-localization task, query images are captured at ground level, and satellite images in the database are captured from an overhead view. Since there

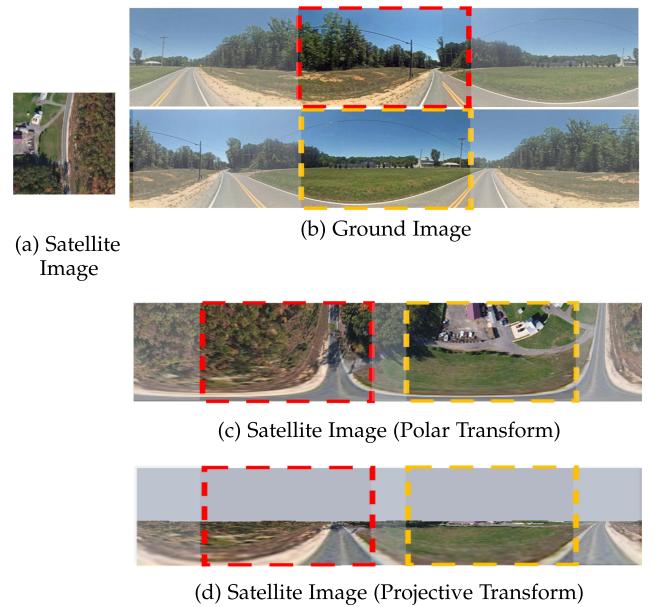


Fig. 3. The challenges of cross-view image matching: the orientation of the query ground image is unknown, and its FoV is restricted. The scene content in panoramas captured at the same location but with different azimuth angles is offset. The image content in an image with a restricted FoV can be entirely different from another image captured from the same location, indicated by different boxes in (b). The polar-transformed satellite image (c) is an approximation to the ground panorama, which preserves all information from the original satellite image. The projective-transformed satellite image (d) loses some information, but preserves the ground-level geometry.

are large appearance variations between these two image domains, our strategy is first to reduce the projection differences between the viewpoints and then to extract discriminative features from the two domains. Furthermore, inspired by how humans localize themselves [46], [47], we exploit the spatial relationships between objects as a critical cue for inferring location and orientation. To this end, we enable our descriptors to encode the spatial relationship among the features, as indicated by F_g and F_s in Fig. 4.

Despite the discriminativeness of the spatially-aware features, they are very sensitive to orientation changes. For instance, when the azimuth angle of a ground camera changes, the scene contents will be shifted in the ground panorama, and the image content may be dramatically different if the camera has a limited FoV, as illustrated in Fig. 3. Therefore, finding the orientation of the ground images is crucial to make the spatially-aware features meaningful. To this end, we propose a Dynamic Similarity Matching (DSM) module, as illustrated in Fig. 4. With this module, we not only estimate the orientation of the ground images but also achieve more accurate feature matching scores, regardless of orientation misalignments and limited FoVs, thus enhancing the performance of geo-localization.

3.1 Bridging the Domain Gap by the Polar Transform

When a scene is planar, a horizontal line in the ground-level panorama corresponds to a circle in the satellite image, and a vertical line in the ground-level panorama corresponds to a ray starting from the center of the satellite image. This layout correspondence motivates us to apply a polar transform to the satellite images. In this way, the spatial layouts of these two domains can be roughly aligned, as illustrated in

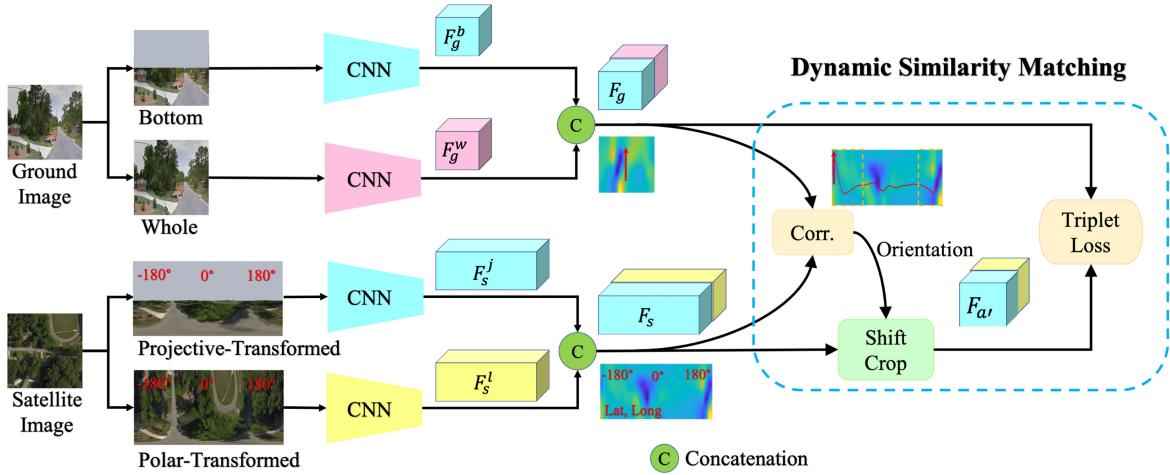


Fig. 4. The overall framework of the proposed method. For the satellite image, we use a two-branch network that first applies a polar transform and a projective transform before extracting features with a CNN. For the ground image, we also use a two-branch network that takes the bottom half of the image corresponding to the projective-transformed satellite image and the whole image corresponding to the polar-transformed image before extracting features. Given the concatenated feature tensors, the correlation between the two streams is used for estimation of the orientation of the ground image with respect to the satellite image. Next, the satellite features are shifted and cropped to obtain the section that (potentially) corresponds to the ground features. The similarity of the resulting features is then used for the retrieval of location.

Figs. 3b and 3c. To be specific, the polar origin is set to the center of each satellite image, corresponding to the geotag location, and the 0° angle is chosen as the northward direction, corresponding to the upwards direction of an aligned satellite image. In addition, we constrain the height (i.e., vertical resolution) of the polar-transformed satellite images to be the same as the ground images, and ensure that the angle subtended by each column of the polar transformed satellite images is the same as in the ground images.

We apply a uniform sampling strategy along rays to the satellite image, such that the innermost and outermost circles of the satellite image are mapped to the bottom and top line of the transformed image respectively. Formally, let $S \times S$ be the size of the satellite image and $H_g \times W_g$ be the target size of polar transform. The polar transform between the original satellite image pixels (u_i^s, v_i^s) and the target polar transformed pixels (u_i^t, v_i^t) is expressed as

$$\begin{cases} u_i^s = u_0 - r(H_g - v_i^s) \cos(2\pi u_i^t / W_g) / H_g, \\ v_i^s = v_0 + r(H_g - v_i^t) \sin(2\pi u_i^t / W_g) / H_g, \end{cases} \quad (1)$$

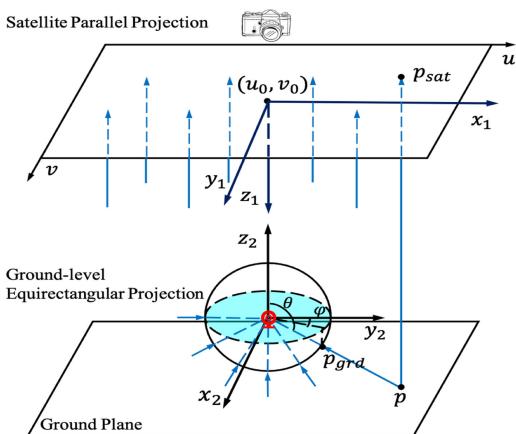


Fig. 5. Illustration of latent geometric correspondences between a satellite image and a ground-level panorama for pixels on the ground plane.

Authorized licensed use limited to: University of Queensland. Downloaded on December 31, 2025 at 13:20:06 UTC from IEEE Xplore. Restrictions apply.

where (u_0, v_0) is the satellite image center, and r is the maximum radius for the polar transform and set to $S/2$.

3.2 Bridging the Domain Gap by the Projective Transform

The polar transform is a simple approximation for the cross-view image transformation. In this article, we further establish more realistic geometric correspondences between the satellite image and the ground-level panorama, especially for scenes that are planar and lie on the ground plane.

Transformation Between Satellite and Ground-Level Cameras. As illustrated in Fig. 5, we use (x_1, y_1, z_1) to represent the satellite camera coordinates and (x_2, y_2, z_2) to denote the ground-level camera coordinates. The transformation between the two camera coordinate systems is expressed as

$$\begin{bmatrix} x_1 \\ y_1 \\ z_1 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & \mathcal{H} \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_2 \\ y_2 \\ z_2 \\ 1 \end{bmatrix}, \quad (2)$$

where \mathcal{H} is the height of the satellite above the ground level.

Satellite Camera Coordinate System. The satellite camera projects a point (x_1, y_1, z_1) to its image coordinates (u_i^s, v_i^s) by a parallel projection

$$\begin{bmatrix} u_i^s \\ v_i^s \end{bmatrix} = \begin{bmatrix} s & 0 & 0 & u_0 \\ 0 & s & 0 & v_0 \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ z_1 \\ 1 \end{bmatrix}, \quad (3)$$

where s is the resolution of the satellite image and (u_0, v_0) is the satellite image center.

Ground-Level Spherical Camera Coordinate System. In the ground-level camera coordinate system, we define θ as the elevation angle with respect to the z_2 axis and ϕ as the azimuth angle. To be consistent with the polar transform, $\phi = 0^\circ$ is the northward direction and it corresponds to the negative direction along x_2 axis. For both θ and ϕ , clockwise is the

positive direction. The mapping between (θ, ϕ) and (x_2, y_2, z_2) is computed as

$$\begin{cases} x_2 = z_2 \tan \theta \cos \phi, \\ y_2 = -z_2 \tan \theta \sin \phi, \end{cases} \quad (4)$$

and the projection between (θ, ϕ) and ground-level panorama image coordinates (u_i^t, v_i^t) is expressed as

$$\begin{cases} \theta = \pi v_i^t / H_g, \\ \phi = 2\pi u_i^t / W_g. \end{cases} \quad (5)$$

Projection Between Satellite and Ground-Level Images. Finally, the mapping between the satellite image coordinates (u_i^s, v_i^s) and ground-level panorama image coordinates (u_i^t, v_i^t) is established as

$$\begin{cases} u_i^s = u_0 + sz_2 \tan(\pi v_i^t / H_g) \cos(2\pi u_i^t / W_g), \\ v_i^s = v_0 - sz_2 \tan(\pi v_i^t / H_g) \sin(2\pi u_i^t / W_g), \end{cases} \quad (6)$$

where z_2 represent the scene height at pixel (u_i^t, v_i^t) . Similar to the polar transform, we choose the satellite image center (u_0, v_0) as the projection point of the projective transform.

In practice, the height maps of satellite images are hard to obtain and the scene heights are much smaller compared to the distance between the ground and the satellite camera. Thus, we assume all the pixels in the satellite image lie on the ground plane. Therefore, z_2 is set to the height of ground plane with respect to the ground camera, and $0.5H_g < v_i^t \leq H_g$ because the ground plane is mostly projected to the bottom half of a ground-view image. In doing so, we construct the geometric correspondences between satellite and ground images for points that lie on the ground plane. For clarity, we call this projection as “projective transform”.

Note that both the polar transform and the projective transform are implemented in an inverse warping manner. Thus, they are differentiable and applicable in end-to-end training. In our implementation, they are applied as a pre-processing step in order to reduce the computation time during both training and testing steps.

3.3 Complementary Between the Transforms

In Figs. 2b and 2c, and Figs. 3c and 3d, we present two examples of the polar-transformed and projective-transformed satellite images. Note that, although the projective transform only retains pixels in a small region of the center of a satellite image, with others pixels being occluded, the scene geometric structure is better preserved by the projective transform than polar transform. This is also reasonable since scene objects far away from the center of a satellite image are unlikely visible in the ground-level panorama.

Furthermore, we illustrate the mapping relationship between the source (satellite) and target (ground-level panorama) coordinates of the two transforms in Fig. 6. As the projection method along the azimuth direction (image columns in target coordinates) is the same in both transforms, we compare their projection difference on mapped image row coordinates. Fig. 6 manifests that most of the pixels are mapped to a few rows in a target image (same color) by the projective transform, while the polar transform maps satellite image pixels to different rows in a target image (indicated by different colors) uniformly. This indicates that the

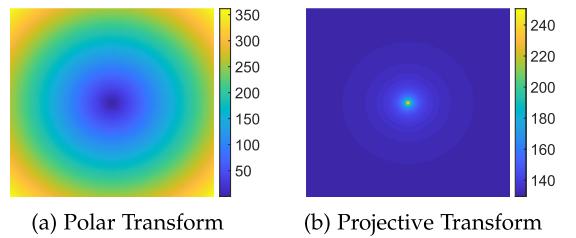


Fig. 6. Visualization of the source (satellite) and target (ground-level panorama) coordinate correspondences by (a) the polar transform and (b) the projective transform. The spatial positions corresponds to the satellite image pixels. Different colors in the two images indicate the row number in the target coordinates.

polar transform retains nearly all the information of a satellite image and the projective transform is able to highlight the ground-level scenes and their structure that would be visible in ground-level panoramas.

As illustrated in Figs. 2b and 2c, and Figs. 3c and 3d, the projective transformed and polar transformed images provide complementary information for the cross-view matching. Hence, we apply both polar and projective transforms to bridge the cross-view domain gap.

3.4 Spatially-Aware Feature Representation

Applying a translation offset along the horizontal axis of a polar-transformed or projective-transformed image is equivalent to rotating a satellite image. Hence, the task of learning the rotation equivariant features for satellite images is re-formulated into learning translation equivariant features on polar or projective transformed images. Doing so significantly reduces the learning difficulty of our network since CNNs are inherently translation equivariant [48], [49], [50]. Because the horizontal axis corresponds to rotation degrees, we need to ensure that the CNN treats the leftmost and rightmost columns of the transformed image as adjacent neighbours. Hence, we employ circular convolutions [51] with periodical padding along the horizontal direction.

Fig. 4 illustrates the pipeline of our coarse localization framework. As the traits of representations resulting from polar- and projective-transformed satellite images are different in nature, we adopt separate CNNs to extract features from them. For the ground images, we also employ separate CNNs. One focuses on the bottom half of the ground image and it is expected to learn similar feature representations to the projective-transformed satellite image. The other one extracts features from the whole ground image and aims to learn matching features with respect to the polar transformed satellite images. The four branches in Fig. 4 have the same architecture. Since the projective-transformed satellite images share the same domain as the ground plane images, we employ the same weights in the first and third branch in Fig. 4 (indicated by the blue color) to extract features and then enforce their similarity. The other two branches, marked with different colors, do not share weights so that they can adapt to their individual domains.

We adopt VGG16 [52] as the network backbone. In particular, the first ten layers of VGG16 are used for features extraction. As the vertical direction of the images may include irrelevant features, such as sky, missing pixels and distortions in the transformed satellite images, we modify the subsequent three layers to decrease the vertical resolution of the

feature maps while maintaining their horizontal resolution. In this manner, our extracted features are more tolerant to distortions along the vertical direction while retaining information along the horizontal direction. We also decrease the number of feature channels to 8. For each branch in Fig. 4, the output feature is of size $4 \times 64 \times 8$. The extracted features in each stream are then concatenated together to form a global feature descriptor. Hence, for a satellite image or a ground-level panorama, the size of their global descriptors is $4 \times 64 \times 16$. Following our previous work [9], the extracted feature volumes are used to preserve the spatial layout of scenes, thus improving the discriminativeness of the descriptors. For query ground images with limited FoV, the width of the extracted features is decreased proportionally.

3.5 Dynamic Similarity Matching (DSM)

When the orientation of ground and transformed satellite features is aligned, their descriptors can be easily compared. However, the orientation of the ground images is not always available, and orientation misalignments increase the difficulty of geo-localization significantly, especially when the ground image has a limited FoV. When humans use a map to localize themselves, they determine their locations and orientation jointly by matching what they have seen to what a map shows [46], [47]. In order to let the network mimic this process, we compute the correlation between the ground and satellite descriptors along the azimuth angle axis. To be specific, we use the ground descriptors as a sliding window and then compute the inner product between then ground and satellite descriptors across all possible orientation.

Let $F_s \in R^{H \times W_s \times C}$ and $F_g \in R^{H \times W_g \times C}$ denote the satellite and ground descriptors, respectively. Let H and C indicate the height and channel number of the descriptors, and W_s and W_g indicate the width of the satellite and ground descriptors. The correlation between F_s and F_g is expressed as

$$[F_s * F_g](i) = \sum_{c=1}^C \sum_{h=1}^H \sum_{w=1}^{W_g} F_s(h, \text{mod}(i+w, W_s), c) F_g(h, w, c), \quad (7)$$

where $F(h, w, c)$ is the feature response at index (h, w, c) while operator mod denotes the modulo operation. Having computed the correlation, we take the argument of the maximum similarity scores as our estimated orientation misalignment of the ground image with respect to the transformed satellite image.

We normalize F_s and F_g by the ℓ_2 norm before calculating the correlation results for panorama images. When a ground image has a limited FoV, we crop the transformed satellite features corresponding to the FoV of the ground image. Then we re-normalize the cropped satellite features and calculate the ℓ_2 distance between the ground and satellite descriptors as the similarity score. Note that if there are multiple maxima in the similarity scores, the satellite image contains indistinguishable symmetries. Thus, we choose one of these maxima at random.

3.6 Training DSM

During the training process, our DSM module is applied to all ground and satellite pairs, regardless of whether they are

matching or not. For matching pairs, DSM forces the network to learn similar feature embeddings for ground and transformed satellite images as well as discriminative feature representations along the horizontal direction (i.e., azimuth). In this way, DSM is able to identify the orientation misalignment as well as find the best feature similarity for matching. For non-matching pairs, we first find the orientation with the highest similarity using the DSM module, and then minimize the maximum similarity score of non-matching pairs to make the descriptors more discriminative. Following traditional cross-view localization methods [4], [5], [8], we employ the weighted soft-margin triplet loss [4] to train our network. Our training loss is expressed as

$$\begin{aligned} \mathcal{L} = & \log \left(1 + e^{\alpha \|F_g^b - F_{s'}^j\|_F - \alpha \|F_g^b - F_{s'}^l\|_F} \right) \\ & + \log \left(1 + e^{\alpha \|F_g^w - F_{s'}^l\|_F - \alpha \|F_g^w - F_{s'}^l\|_F} \right) \\ & + \log \left(1 + e^{\alpha \|F_g - F_{s'}\|_F - \alpha \|F_g - F_{s'}\|_F} \right). \end{aligned} \quad (8)$$

In the above equations, F_g^b and F_g^w denote the extracted features from the bottom half of a query ground image and the whole ground image, $F_{s'}^j$ and $F_{s'}^l$ are the cropped features from the non-matching projective-transformed satellite features and the polar-transformed satellite features, $F_{s'}^j$ and $F_{s'}^l$ denote the cropped features from the matching projective-transformed satellite features and the polar-transformed satellite features. Recall that we do not require ground view query images to be panoramas, and the transformed satellite images will be cropped automatically to fit the resolution of the ground images by the DSM module. $F_g = [F_g^b, F_g^w]$ is the global query ground descriptor, $F_{s'} = [F_{s'}^j, F_{s'}^l]$ and $F_{s'}' = [F_{s'}^j, F_{s'}^l]$ indicate the cropped satellite descriptors of the matching satellite image and a non-matching satellite image aligned by our DSM module, $[\cdot, \cdot]$ is the concatenation operation, and $\|\cdot\|_F$ denotes the Frobenius norm. The parameter α controls the convergence speed of training process. Following precedent [4], [5], [8], we set it to 10.

4 FINE-GRAINED CROSS-VIEW IMAGE MATCHING

After the coarse geo-localization process, we retrieve an satellite image that is the most similar to the query ground image. However, the GPS tag of the satellite image is associated with its center point, whereas the ground image might be captured away from the center of the satellite image. Therefore, in this article we further propose a novel mechanism to compute the displacement between the query ground camera and the retrieved satellite image center.

To achieve this goal, we employ the projective transform explained in Section 3.2. As indicated in Fig. 6, only a small portion of the satellite image is projected to the transformed image. Thus, the projective transform is sensitive to the projection center. This phenomenon is useful in aiding fine-grained localization of where a query image was taken.

In the fine-grained camera geo-localization stage, we first select a central square region in the retrieved satellite image, as shown in Fig. 1b. This central region encloses the possible query camera locations. At each pixel in the selected region,

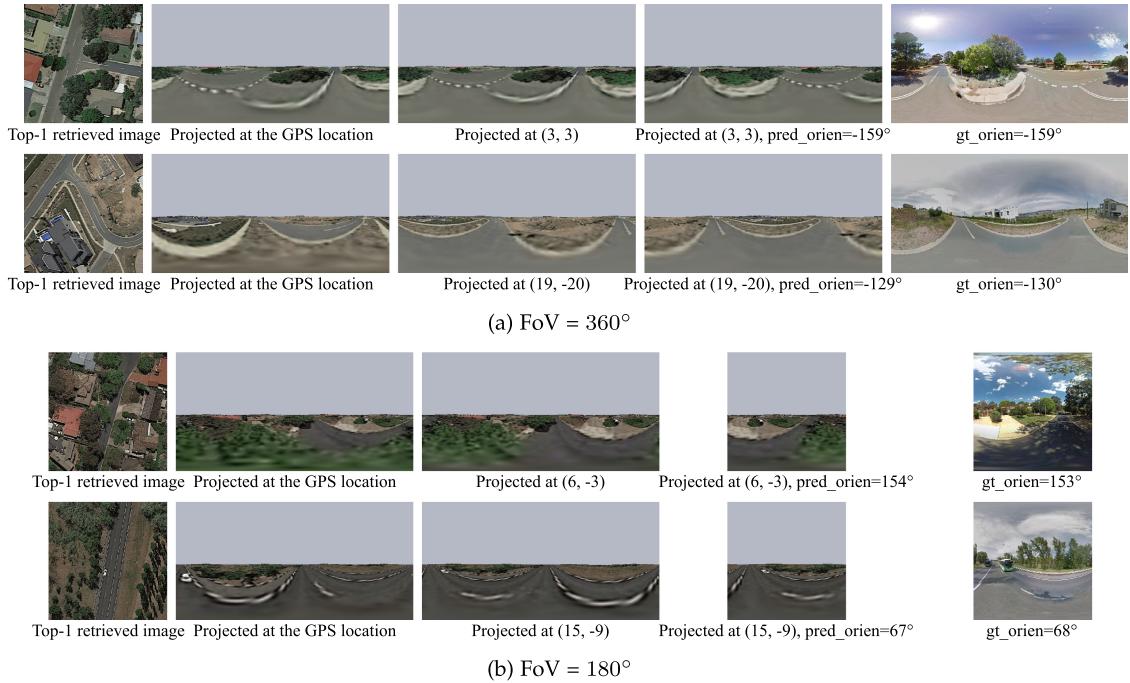


Fig. 7. Qualitative illustration of fine-grained 3-DoF camera localization for query images with unknown orientation and varying FoVs. Given query images (the last column), we first retrieve their most similar satellite images (the first column) from the database. The projective-transformed satellite images according to the query camera GPS locations are presented in the second column. The ground structure of those images (the second column) is significantly different from the query images (the last column), indicating that the GPS locations are not accurate. According to our fine-grained camera geo-localization method, we exhaustively project the retrieved satellite image to their ground panorama coordinates at points in a central square region of the retrieved satellite image. Among the projected images, the most similar ones to the query ground images are presented in the third column. The fourth column shows the shifted and cropped projective-transformed satellite images that align with query images.

we project the satellite image to the corresponding ground panorama coordinates by applying the projective transform. Then, a query image with unknown orientation (or even with limited FoV) is compared with each of the projective-transformed satellite images via the dynamic similarity matching module. Specifically, we circularly shift the projective-transformed satellite image along the azimuth (horizontal) direction. If the query ground image has a restricted FoV, we also crop out a portion of the transformed image according to its FoV. The similarity between the transformed image and the query ground image is then computed. Note that in this fine-grained camera localization process we use the SSIM as the similarity measure instead of the cross-correlation as SSIM is more suitable to evaluate structural differences between the projective-transformed satellite image and the query ground image.

For each projective-transformed satellite image, we record the maximum similarity across different orientations as its similarity to the query image. Among all the projective-transformed satellite images, the most similar one is selected, as marked by the green box in Fig. 1b. Its corresponding projection point, indicated by the blue dot in Fig. 1b, is taken as the query camera location. The computed relative orientation is taken as the camera orientation. Under such a scheme, the precision of the location and orientation estimation depends on the real-world distance of a satellite image pixel and the resolution of the query image, respectively.

Fig. 7 provides some qualitative examples of fine-grained localization of orientation-unknown images. The query images are presented in the last column and their corresponding retrieved top-1 satellite images are shown in the

first column. Since the geotags of the database satellite images are associated with the image centers, we first apply the projective transform to the retrieved satellite images according to their image centers, which is also the GPS location of query ground images provided by the dataset. The projected images are presented in the second column of Fig. 7. They are significantly different from the query ground images. We next apply the proposed fine-grained camera localization method, and visualize the most similar projective-transformed satellite images in the third column. The fourth column provides the shifted and cropped projective-transformed satellite images according to the estimated relative orientation and the FoV of ground images. The estimated camera location with respect to the satellite image center and the relative orientation are presented under each of the images. As can be seen, images presented in the fourth column align with the original query ground images.

Since in the second stage (fine-grained camera localization) we search exhaustively every point in the selected central region, we recommend that the reference satellite images in the first stage (coarse camera localization) should cover the entire region as densely as possible. By doing so, the selected central region used for fine-grained camera geo-localization will be small, reducing the computation complexity.

In our implementation, the satellite image is resized to 256×256 pixels from 1200×1200 pixels. The central square region covers 40×40 pixels in the resized image, corresponding to 11.25×11.25 square meters (± 5.625 meters to the satellite image center). By searching the 40×40 candidate locations and the 512 candidate orientations ($40 \times 40 \times 512 = 614,400$ candidate solutions in total), it takes 15 minutes on



Fig. 8. Cross-view image pairs from the CVUSA (top two rows) and CVACT (bottom two rows) datasets. The satellite images are on the left and the ground panoramas are on the right.

average for the fine-grained localization of a query image on an RTX 2080 Ti GPU.

5 EXPERIMENTS

5.1 Datasets

We carry out the experiments on two standard cross-view datasets, CVUSA [2] and CVACT [5]. They both contain 35,532 training ground and satellite pairs and 8,884 testing pairs. Following an established testing protocol [5], [8], we denote the test sets in CVUSA and CVACT as CVUSA and CVACT_val, respectively. CVACT also provides a larger test set, CVACT_test, which contains 92,802 cross-view image pairs for fine-grained city-scale geo-localization. Note that the ground images in both of the two datasets are panoramas, and all the ground and satellite images are north aligned. Fig. 8 presents samples of cross-view image pairs from the two datasets.

For localizing ground images with unknown orientation and limited (180°) FoV, we use the image pairs in CVUSA and CVACT_val and randomly rotate the ground images along the azimuth direction and crop them according to a predetermined FoV. The source code of this work is available at <https://github.com/shiyujiao/IBL.git>.

Accurate Camera Geo-Localization. In the CVUSA dataset, there is no GPS data available for the ground–satellite image pairs. The CVACT dataset, introduced in our previous work [5], does provide GPS data for every ground–satellite image pair. However, due to GPS drift, the location of a ground camera provided by the dataset is not accurate, which is visualized in Fig. 7. In practice, it is also hard to collect strictly location-aligned satellite–ground image pairs [53]. Hence, we adopt user study for the evaluation of our fine-grained localization method on the CVACT dataset.

5.2 Coarse Camera Geo-Localization

5.2.1 Implementation Details

We use the first ten convolutional layers in VGG16 with pre-trained weights on Imagenet [54] and randomly initialize the parameters in the following three layers for extraction of global feature descriptors. The first seven layers are kept

fixed and the subsequent six layers are learnt. The Adam optimizer [55] with a learning rate of 10^{-5} is employed for training. Following [3], [4], we adopt an exhaustive mini-batch strategy [3] with a batch size of $B = 32$ to create training triplets. Specifically, for each ground image within a mini-batch, there is one matching satellite image and $B - 1$ non-matching images. Thus we construct $B(B - 1)$ triplets. Similarly, for each satellite image, there is one matching ground image and $B - 1$ non-matching images within a mini-batch, and we create another $B(B - 1)$ triplets. In total, we obtain $2B(B - 1)$ triplets in total.

In order to obtain a time-efficient approach, we compute the correlation in our DSM module by the fast Fourier transform during inference. Specifically, we store Fourier coefficients of satellite features in the database, and only calculate Fourier coefficients of the ground descriptors in the forward pass. In doing so, the computation cost yields $13NHW_sC$ flops (including $4NHW_sC$ flops for coefficient multiplication in the spectral domain, and $1.5NHCW_s \log_2 W$ flops for the inverse Fast Fourier Transform), where H , W_s and C are the height, width and channel number of the global feature descriptor of a satellite image, and N is the number of database satellite images. In contrast, performing correlation in the spatial domain requires $2NHW^2 C$ flops. Thus, the computation time by applying the Fourier transform is reduced by a factor of 10 ($\frac{13NHW_sC}{2NHW^2 C} \approx \frac{1}{10}$).

For localizing query images with unknown orientation and a limited FoV, there is a shift-and-crop operation to find the corresponding part of a query image in a satellite image. The flops in the shift operation are NHW_sC , and the flops in the crop operation are NHW_gC , where N is the number of database satellite images, H and C are the height and channel number of satellite (and query ground) features, W_g and W_s are the widths of satellite and query ground features, respectively, and $W_g = \text{FoV} * W_s / 360$. The flops in computing the similarity between a query image and all satellite images are $3NHW_gC$. Hence, the complexity of retrieving a ground image from a database with N satellite images is $O(NHW_sC)$.

Using an RTX 2080 Ti GPU, the feature extraction time for a ground image is 0.01 s, and it takes 0.06 s on average to retrieve its satellite counterpart from a database containing 8884 reference images. On the CVACT_val dataset, the 8884 reference images cover approximately 64 km².

5.2.2 Evaluation Metrics

Location Estimation. Following the standard evaluation procedure for cross-view image localization [3], [4], [5], [6], [7], [8], [43], we use the top K recall as the location evaluation metric to examine the performance of our method and compare it with the state-of-the-art. Specifically, given a ground image, we retrieve the top K satellite images in terms of ℓ_2 distance between their global descriptors. The ground image is regarded as successfully localized if its corresponding satellite image is retrieved within the top K list. The percentage of correctly localized ground images is recorded as recall at K (r@ K).

Orientation Estimation. The predicted orientation of a query ground image is meaningful only when the ground image is localized correctly. Hence, we evaluate the orientation

TABLE 1
Comparison of Our Approach With Existing
Methods on the CVUSA [2] Dataset

Methods	CVUSA			
	r@1	r@5	r@10	r@1%
Workman <i>et al.</i> [1]	—	—	—	34.3
Zhai <i>et al.</i> [2]	—	—	—	43.2
Vo and Hays [3]	—	—	—	63.7
CVM-NET [4]	22.47	49.98	63.18	93.62
Liu and Li [5]	40.79	66.82	76.36	96.12
Regmi and Shah [6]	48.76	73.64	81.27	95.94
GeocapsNet-II [7]	—	—	—	98.07
Siam-FCANet34 [7]	—	—	—	98.3
CVFT [8]	61.43	84.69	90.49	99.02
SAFA [9]	89.84	96.93	98.14	99.64
Shi <i>et al.</i> [10]	91.96	97.50	98.54	99.67
Ours	92.69	97.78	98.60	99.61

estimation accuracy of our DSM only on ground images that have been correctly localized by the top-1 recall. In this experiment, when the differences between the predicted orientation of a ground image and its ground-truth orientation is within $\pm 10\%$ of its FoV, the orientation estimation of this ground image is deemed as a success. We record the percentage of ground images of which the orientation is correctly predicted as the orientation estimation accuracy (Orien_acc).

Combined Measure. Denote the top-1 recall rate for location estimation as Loc_acc. The overall 3-DOF camera localization performance is computed as Loc_acc \times Orien_acc, which we denote as Overall.

5.2.3 Localizing Orientation-Aligned Panoramas

First, we investigate the performance of location estimation of our method and compare it with the state-of-the-art [1], [2], [3], [4], [5], [6], [7], [8], [9], [10] on the standard CVUSA and CVACT datasets, where ground images are orientation-aligned panoramas. The recall results at top-1, top-5, top-10 and top-1% on the CVUSA and CVACT_val datasets are presented in Tables 1 and 2, respectively. They are reported from other works or produced by the re-trained models using source codes provided by the authors. The complete r@K performance curves on CVUSA and CVACT_val are illustrated in Figs. 9a and 9b, respectively.

Among those baseline methods, [1], [2], [3] are the first approaches that utilize deep learning for cross-view related tasks. CVM-NET [4], GeocapsNet-II [7] and Siam-

TABLE 2
Comparison of Our Approach With Existing Methods on the
CVACT_val [5] Dataset by Re-Training Existing Networks

Methods	CVACT_val			
	r@1	r@5	r@10	r@1%
CVM-NET [4]	20.15	45.00	56.87	87.57
Liu and Li [5]	46.96	68.28	75.48	92.01
Regmi and Shah [6]	48.62	72.48	79.65	93.16
CVFT [8]	61.05	81.33	86.52	95.93
SAFA [9]	81.03	92.80	94.84	98.17
Shi <i>et al.</i> [10]	82.49	92.44	93.99	97.32
Ours	82.70	92.50	94.24	97.65

FCANet34 [7] focus on designing powerful feature extraction networks. Liu and Li [5] introduce the orientation information to networks so as to facilitate geo-localization performance. These works do not explicitly address the domain gap between ground and satellite images which leads to their inferior performance.

Regmi and Shah [6] adopt a conditional GAN to generate satellite images from ground panoramas. Although it helps to bridge the cross-view domain gap, undesired scene contents are also introduced in this process. CVFT [8] proposes a Cross-view Feature Transport (CVFT) module to better align ground and satellite features. However, it is hard for networks to learn geometric and feature response correspondences simultaneously. SAFA [9] explores a parameter-free polar transform to bridge the geometric domain gap, and proposes a Spatially-aware Position Embedding (SPE) module to further construct geometric and feature correspondences between the two view images. In [10], we found that further increasing the spatial size of the global image descriptors can increase the localization performance. Similar to SAFA, the polar transform is also employed in [10] to bridge the cross-view domain gap. However, in this journal paper, we further establish the geometric correspondences between a satellite image and its corresponding ground-level panorama by a projective transform. The projective transform and the polar transform cooperate with each other to align the matching ground and satellite image pairs. As shown in Tables 1 and 2, the top-1 recall rate has been further boosted by the newly proposed method compared to our conference work [10]. Although SAFA achieves slightly better performance on the CVACT_val, its performance degrades significantly when orientation of query images is unknown, which will be investigated in later experiments.

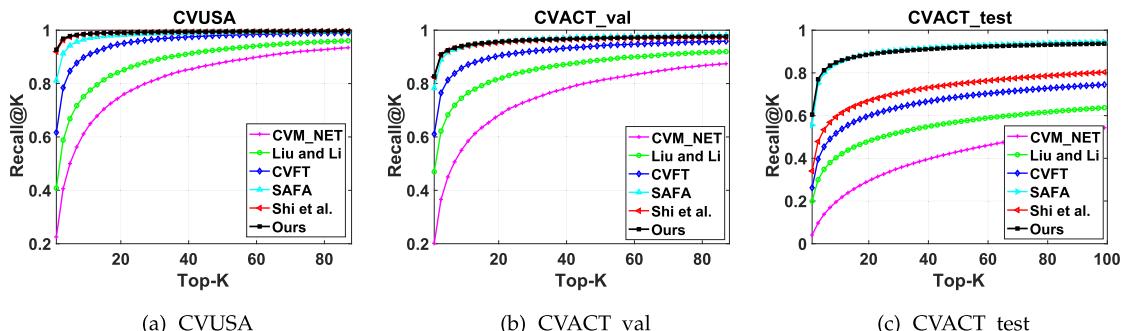


Fig. 9. Evaluations of recall at different values of K on the CVUSA, CVACT_val and CVACT_test datasets.

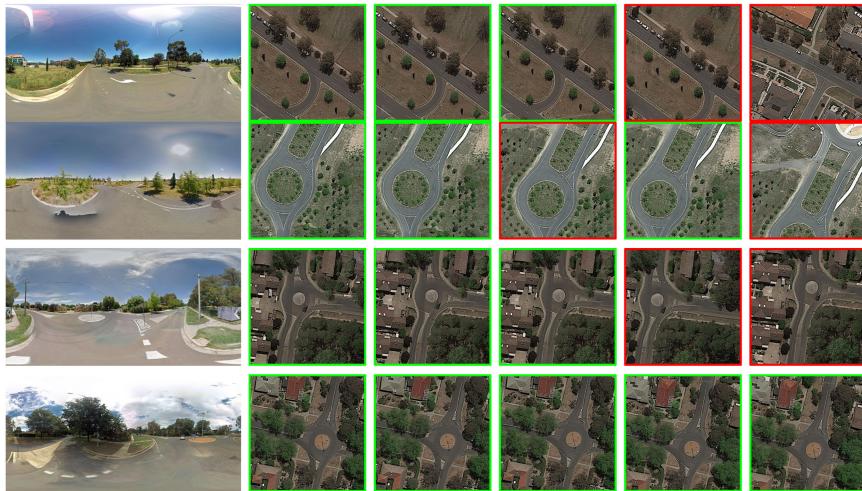


Fig. 10. Visualization of localization results attained by our method on the CVACT_test set. From left to right: ground-level query image and the top 1-5 retrieved satellite candidates. Green and red borders indicate correctly and incorrectly retrieved results, respectively.

Distance-Based Localization. In CVUSA and CVACT_val, there is only one matching satellite image in the database for a query image in the test set. In [5], we introduced CVACT-test to evaluate the performance of different methods on real-world localization scenarios. This test set provides geotagged (GPS) satellite images that densely cover a city, and the localization performance is measured in terms of distance (meters). Specifically, a ground image is considered as successfully localized if one of the retrieved top K satellite images is within 5 meters of the ground-truth location of the query ground image. That is to say, in this test set, there might be several matching satellite images in the database for a query ground image. Following the evaluation protocol in [5], we plot the percentage of correctly localized ground images (recall) at different values of K in Fig. 9c.

Compared to the work [10], SAFA [9] has eight additional Spatially-aware Position Embedding (SPE) modules to construct cross-view correspondences. Thus, the performance of SAFA is significantly better than that of Shi *et al.* [10] in this challenging test set. As a contribution of this paper, we establish more realistic geometric correspondences between satellite and ground-level images compared to our conference version [10] which makes the extracted feature descriptors more informative. Hence, the localization performance on the CVACT_test is further boosted. The top-1 recall rate of SAFA and the newly proposed algorithm on the CVACT_test is 55.50% and 60.46%, respectively. Note that having the SPE module in the framework will prohibit the application of our DSM which is designed to address orientation-unknown and limited FoV query images. Thus we do not retain it in the proposed framework in this article. In the following section, we will demonstrate the superiority of our DSM module on localizing query images with unknown-orientation and limited FoV. Fig. 10 shows a visualization of retrieved satellite images for orientation-aligned panoramas by our method.

5.2.4 Practical Localization: Localizing With Unknown Orientation and Limited FoV

We compare the performance of our newly proposed method with our three previous works [8], [9], [10], on the CVUSA and CVACT_val datasets in a more realistic

localization scenario, where the ground images do not have a known orientation and have a limited FoV. Recall that Liu and Li [5] require the orientation information as an input, so we cannot compare our approach with such a method.

Location Estimation. In order to evaluate the impact of orientation misalignments and limited FoVs on localization performance, we randomly shift and crop the ground panoramas along the azimuth direction for the CVUSA and CVACT_val datasets. In this manner, we mimic the procedure of localizing images with limited FoV and unknown orientation. For completeness, we also report the performance of the models trained without known orientations on orientation-aligned query images. Results are presented in Table 3.

For CVFT [8] and SAFA [9], the difference between the models from Section 5.2.3 and this section is that we apply the random orientation augmentation on training image pairs. By doing so, we expect that the networks can tolerate some orientation misalignments between satellite and ground image pairs. Those two methods do not have the ability to explicitly align the orientation of ground and satellite images. Thus, their performance drops significantly when localizing ground images with unknown orientation, which can be seen by comparing Tables 1, 2 and 3. Furthermore, since the data augmentation on orientation is applied during training, these networks need to spend more capacity on tolerating orientation misalignments. Their performance on localizing orientation-aligned query images drops.

For the method [10] and the newly proposed method, the models from Section 5.2.3 and this section are the same, where random orientation augmentation is applied during training. In these two methods, an explicit parameter-free orientation alignment mechanism (the Dynamic Similarity Matching module) is applied, and thus they outperform CVFT and SAFA by a large margin on localizing orientation-unknown query images. The performance degradation of the two methods between orientation-aligned and orientation-unknown images is mainly due to the orientation ambiguity. Compared to the work [10] that we extend from, the newly proposed method establishes more sensible geometric correspondences between satellite and ground-level images by a projective transform, leading to a better performance.

TABLE 3
Comparison of Recall Rates for Localizing Ground Images With Unknown Orientations and Varying FoVs
(Models Trained With Random Orientation Augmentation)

Dataset	Methods	FoV=360°								FoV=180°							
		Orientation Aligned				Orientation Unknown				Orientation Aligned				Orientation Unknown			
		r@1	r@5	r@10	r@1%												
CVUSA	CVFT [8]	23.13	45.01	55.94	86.93	23.38	44.42	55.20	86.64	7.97	2354	33.32	74.40	8.10	24.25	34.47	75.15
	SAFA [9]	53.33	76.82	84.29	97.48	52.85	76.27	83.78	97.50	27.43	53.85	65.70	92.27	28.71	55.27	66.65	92.78
	Shi <i>et al.</i> [10]	91.96	97.50	98.54	99.67	78.11	89.46	92.90	98.50	75.11	89.72	93.48	98.71	48.53	68.47	75.63	93.02
	Ours	92.69	97.78	98.60	99.61	78.94	90.31	93.42	98.67	75.65	89.17	93.44	98.90	54.27	72.78	79.54	94.73
CVACT_val	CVFT [8]	26.19	46.09	54.52	80.53	26.79	46.89	55.09	81.03	6.56	18.10	26.24	62.78	7.13	18.47	26.83	63.87
	SAFA [9]	44.15	68.02	75.83	93.14	43.54	67.65	75.64	93.05	20.98	43.65	54.58	86.09	21.15	44.60	55.66	86.40
	Shi <i>et al.</i> [10]	82.49	92.44	93.99	97.32	72.91	85.70	88.88	95.28	67.26	83.84	87.57	95.36	49.12	67.83	74.18	89.93
	Ours	82.70	92.50	94.24	97.65	73.06	85.73	88.76	95.44	67.23	83.57	87.81	95.25	52.98	71.18	77.36	91.61

TABLE 4
The Overall Performance of 3-DoF Coarse Camera Localization

Dataset	CVUSA									CVACT_val								
	360°			180°			360°			180°			360°			180°		
	Loc_acc	Orien_acc	Overall	Loc_acc	Orien_acc	Overall	Loc_acc	Orien_acc	Overall									
Shi <i>et al.</i> [10]	78.11	99.41	77.65	48.53	98.54	47.72	72.91	99.84	72.79	49.12	99.10	48.68						
Ours	78.94	99.45	78.51	54.27	96.87	52.57	73.06	99.75	72.88	52.98	98.72	52.30						

Finally, as FoVs increase, one can observe that all of the methods perform better. That is mainly because a larger FoV image provides richer scene contents, making a global descriptor more discriminative.

Orientation Estimation. As aforementioned, the estimated orientation of a query image is meaningful only if its location has been correctly determined. Thus, the experiment on the orientation estimation is conducted on ground images which are correctly localized in terms of top-1 retrieved candidates. For the fair comparison between different algorithms, we combine the location estimation accuracy and the orientation estimation accuracy together to present the overall performance score in Table 4. As indicated by Table 4, our proposed method achieves higher location estimation accuracy than our previous work [10], whereas the orientation estimation accuracy is slightly lower. This implies that scenes may look similar in multiple directions, even when the location is estimated correctly. For instance, a person standing on a road may be able to localize their position but will find it difficult to determine the orientation if the views are similar along the road in both directions. Fig. 11 contains examples of such symmetric scenes. According to the combined score, our newly proposed method achieves the best performance on 3-DoF camera localization. Fig. 13 visualizes the orientation estimation process.



Fig. 11. Examples of symmetric scenes (satellite images). At these locations, it is hard to determine the orientation (azimuth angle) of a ground image.

Authorized licensed use limited to: University of Queensland. Downloaded on December 31, 2025 at 13:20:06 UTC from IEEE Xplore. Restrictions apply.

Why Does Performance Decrease When the Orientation is Unknown? When the orientation of the camera is unknown, the search space becomes significantly larger. The alignment between the query image with non-matching satellite images increases the number of local maxima in the similarity objective function. Due to occlusions and the definition of the location alignment between the ground and satellite images in two transforms, the matching satellite image projected by the transforms might be not very similar to the query image. When the similarity score of an incorrect location-orientation pair is higher than the true pair, our method will estimate the wrong location and orientation for a query image.

To demonstrate this, we conduct experiments for localizing orientation-unknown query images where the orientation alignment is only performed with respect to the matching satellite image, not the entire database as is standard. Note that this is for illustration only, and is not a practical setting. The results are shown in the third row of Table 5. It can be seen that the performance is significantly better than that of aligning query images with non-matching ones by the DSM (the second row of Table 5). It can be seen that the performance is significantly better than when aligning with respect to the entire database (the second row of Table 5). In fact, the results are even better than when using orientation-aligned query images, since minor orientation misalignments between the query and matching satellite images can be rectified by our DSM module.

5.2.5 Discussion on the Necessity of the Polar Transform and the Projective Transform

In this section, we conduct an ablation study to demonstrate the effectiveness and necessity of the polar transform (polar) and the projective transform (proj). To this end, we remove

TABLE 5
Comparison of Recall Rates When the Query Image is Orientation-Aligned to Every Element
in the Database (Standard Case) or Just the Matching Image

Orien	Align	Negative	CVUSA								CVACT_val							
			FoV=360°				FoV=180°				FoV=360°				FoV=180°			
			r@1	r@5	r@10	r@1%	r@1	r@5	r@10	r@1%	r@1	r@5	r@10	r@1%	r@1	r@5	r@10	r@1%
Aligned	-		92.69	97.78	98.60	99.61	75.65	89.17	93.44	98.90	82.70	92.50	94.24	97.65	67.23	83.57	87.81	95.25
Unknown	✓		78.94	90.31	93.42	98.67	54.27	72.78	79.54	94.73	73.06	85.73	88.76	95.44	52.98	71.18	77.36	91.61
Unknown	✗		93.22	98.22	98.94	99.74	79.44	92.64	95.78	99.46	84.30	93.93	95.52	98.49	73.06	88.61	92.19	97.90

TABLE 6
Comparison of Recall Rates for Localizing Ground Images With Unknown Orientation and Varying FoVs

Dataset	Methods	FoV=360°								FoV=180°							
		Orientation Align				Orientation Unknown				Orientation Align				Orientation Unknown			
		r@1	r@5	r@10	r@1%	r@1	r@5	r@10	r@1%	r@1	r@5	r@10	r@1%	r@1	r@5	r@10	r@1%
CVUSA	Ours w/o polar	84.78	93.99	95.95	99.13	65.33	80.46	85.61	0.9612	60.48	79.78	85.40	96.49	37.27	57.36	65.93	88.15
	Ours w/o proj [10]	91.96	97.50	98.54	99.67	78.11	89.46	92.90	98.50	75.11	89.72	93.48	98.71	48.53	68.47	75.63	93.02
	Ours	92.69	97.78	98.60	99.61	78.94	90.31	93.42	98.67	75.65	89.17	93.44	98.90	54.27	72.78	79.54	94.73
CVACT_val	Ours w/o polar	68.71	84.38	87.94	94.70	50.50	68.83	74.75	89.45	43.88	65.80	73.05	89.46	27.95	46.26	54.10	79.45
	Ours w/o proj [10]	82.49	92.44	93.99	97.32	72.91	85.70	88.88	95.28	67.26	83.84	87.57	95.36	49.12	67.83	74.18	89.93
	Ours	82.70	92.50	94.24	97.65	73.06	85.73	88.76	95.44	67.23	83.57	87.81	95.25	52.98	71.18	77.36	91.61
Dataset	Methods	FoV=90°								FoV=70°							
		Orientation Align				Orientation Unknown				Orientation Align				Orientation Unknown			
		r@1	r@5	r@10	r@1%	r@1	r@5	r@10	r@1%	r@1	r@5	r@10	r@1%	r@1	r@5	r@10	r@1%
CVUSA	Ours w/o polar	11.65	23.73	31.09	59.87	7.02	17.95	25.01	59.43	6.27	14.55	19.57	43.91	4.09	11.81	18.13	50.98
	Ours w/o proj [10]	33.66	51.70	59.68	82.35	16.19	31.44	39.85	71.13	20.88	36.99	44.70	70.95	8.78	19.90	27.30	61.20
	Ours	20.28	36.41	44.47	72.83	11.30	26.00	34.23	69.03	10.50	21.83	28.79	56.07	5.85	15.71	23.13	59.34
CVACT_val	Ours w/o polar	8.13	17.54	23.06	48.50	4.19	11.52	16.34	42.75	3.33	9.31	13.09	33.86	2.04	5.83	9.16	31.89
	Ours w/o proj [10]	31.17	51.44	60.05	82.73	18.11	33.34	40.94	68.65	18.45	35.87	44.39	71.85	8.29	20.72	27.13	57.08
	Ours	21.42	38.43	47.06	73.86	14.34	27.90	35.18	64.86	9.68	20.60	27.68	56.22	4.60	12.73	17.99	46.86

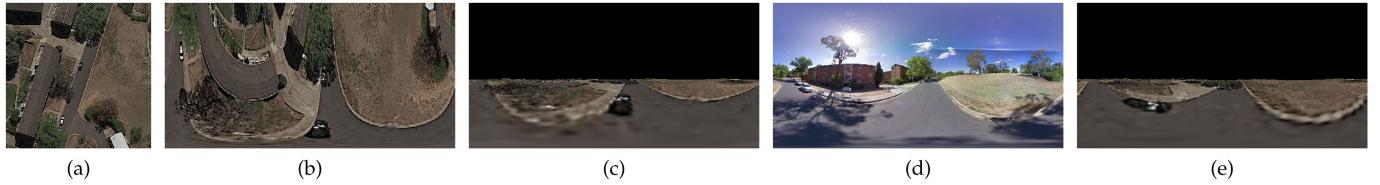


Fig. 12. (a) Satellite image; (b) polar-transformed satellite image at the satellite image center; (c) projective-transformed satellite image at the satellite image center; (d) query ground-level panorama; (e) projective-transformed satellite image at the ground camera location. The ground camera location is determined by our fine-grained camera localization method. All of the images are orientation aligned.

the branch of polar-transformed satellite images or projective-transformed satellite images from our whole pipeline. We denote such settings as “Ours w/o polar” and “Ours w/o proj,” respectively. Accordingly, the corresponding ground branch is also removed from the whole pipeline. In particular, “Ours w/o proj” becomes the method proposed in our conference version [10]. Table 6 reports the localization performance of the three baselines.

When the ground images have a 360° or 180° FoV, the newly proposed method with both polar and projective transform achieves the best performance. That is because the polar transform preserves all of the scene content information and the projective transform recovers the ground structure from a satellite image. As seen in Fig. 12, there is a building in the ground-level panorama (Fig. 12d), and it is also visible from the satellite image (Fig. 12a). The polar transformed satellite image (Fig. 12b) successfully preserves this scene information. However, the building is hard to be recognized from the projective-transformed satellite image

(Fig. 12c). In contrast, the ground structure is better recovered by the projective transform. The two types of transformed images complement with each other to make the satellite image descriptor informative and discriminative. Moreover, we also present the projective-transformed satellite image at the ground camera location in Fig. 12e. It aligns better with the ground-level panorama. The ground camera location is estimated by our fine-grained camera localization method.

When the ground images have a more restricted FoV, i.e. 90° and 70°, the results in Table 6 show that the newly proposed method achieves inferior performance than “Ours w/o proj”. The reason is that the projective-transformed satellite image only preserves a small portion information from the original satellite image. When the query image has a small FoV, e.g., only the building part in Fig. 12d, it will not be matched to the projective-transformed satellite image. Thus its localization result would be incorrect. Not surprisingly, when the polar transform is removed from our pipeline, the

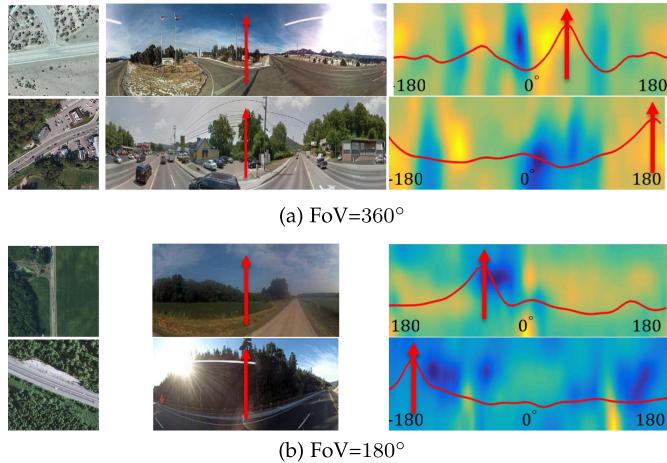


Fig. 13. Visualization of estimated orientation for ground images with $\text{FoV} = 360^\circ$ and 180° . In each of the subfigures, the satellite images are on the left and the ground images are in the middle. We visualize the transformed satellite features and the correlation results (red curves) in the right column. The positions of the correlation maxima in the curves corresponds to the orientation of the ground images.

performance of ‘Ours w/o polar’ decreases significantly. This demonstrate the necessity of the polar transform.

To summarize, our newly proposed method is more suitable for localizing larger FoV images, and when the FoVs of query images are small, removing the projective transform branch and its corresponding ground image branch will be more recommended. In real-world applications, it would be better to employ different branches in our proposed pipeline according to different deployment situations.

Orientation Estimation w.r.t. the Two Transforms. We also study the influence on localization performance when only using the polar transform or the projective transform in orientation estimation, denoted as ‘polar’ and ‘proj,’ respectively. The results are given in Table 7. The ablation labelled as ‘combined’ is our whole method where both of the transforms are applied to orientation estimation. It can be seen that the polar transform contributes more to the final performance, but both help boost performance.

5.3 Fine-Grained Camera Geo-Localization

Here, we conduct three groups of experiments on fine-grained camera geo-localization: (1) location estimation with known orientation; (2) orientation estimation with known location; and (3) joint location and orientation estimation.

Location Estimation. We first evaluate the location estimation accuracy of our method when the orientation of a query

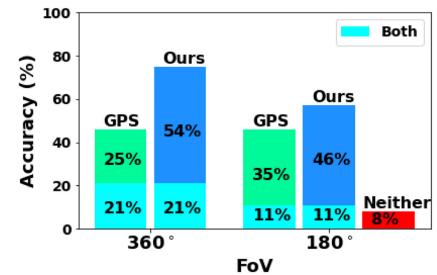


Fig. 14. User study results for fine-grained camera localization (orientation aligned). In this evaluation, users are asked to determine whether a location is correct or not. The color cyan indicates the portion of data where both GPS and our estimated locations are correct. The red bar ‘Neither’ indicates the portion of data that cannot be localized when the FoV decreases from 360° to 180° .

image is given. Since ground-truth labels of the precise camera locations are not available, we conduct a user study with five participants for quantitative evaluation (two of them are the authors of this paper). As shown in Fig. 7, given a query ground image (the last column) and its corresponding satellite image (the first column), it is very difficult for a human to localize which pixel in the satellite image corresponds to the query camera location. However, when the projected satellite image at a given location (the second or third column) is provided, it is much easier to determine whether this location is correct or not by comparing the projected satellite image with the query image. Therefore, the participants are provided with a query image, a projected satellite image at the GPS location, and a projected satellite image at the location estimated by our method. They are required to decide whether the camera location of each projected image is correct or not (binary classification). Moreover, when the query image is different from both projected images, we mark this case as unknown. This case occurs when there is a severe occlusion in the vertical dimension. For example, tree canopies entirely cover the road underneath, which makes the ground-level localization extremely difficult.

We randomly select 150 image pairs from the CVACT_val set for this user study, and we note that participants in this user study are unaware which image is projected according to the GPS location and which is projected according to the estimated location by our method. The user study is first conducted on 360° query images. We exclude the rare (#) ‘unknown’ samples and use the remaining images for the evaluation. From the results in Fig. 14, it can be seen that our method helps to correct the GPS drifting problem. The experiments on localizing 180° FoV images

TABLE 7
Comparison of Localization Performance When Ablating the Transforms for Orientation Estimation

Orien. Est.	FoV=360°				FoV=180°				FoV=90°				FoV=70°				
	r@1	r@5	r@10	r@1%	r@1	r@5	r@10	r@1%	r@1	r@5	r@10	r@1%	r@1	r@5	r@10	r@1%	
CVUSA	Polar	78.61	90.32	93.34	98.56	53.88	72.43	79.10	94.48	10.85	24.16	32.11	66.06	5.49	14.61	21.15	55.47
	Proj	78.57	89.90	93.00	98.37	52.00	69.82	76.28	92.28	10.74	23.69	31.37	65.49	5.61	15.31	22.20	56.89
	Combined	78.94	90.31	93.42	98.67	54.27	72.78	79.54	94.73	11.30	26.00	34.23	69.03	5.85	15.71	23.13	59.34
CVACT_val	Polar	73.11	85.56	88.74	95.41	52.71	70.45	76.83	91.11	13.45	26.31	33.26	61.91	3.80	10.65	15.26	40.54
	Proj	72.73	85.02	88.12	94.65	50.54	67.83	73.67	87.97	12.13	23.12	29.19	55.31	4.12	11.19	15.62	41.90
	Combined	73.06	85.73	88.76	95.44	52.98	71.18	77.36	91.61	14.34	27.90	35.18	64.86	4.60	12.73	17.99	46.86

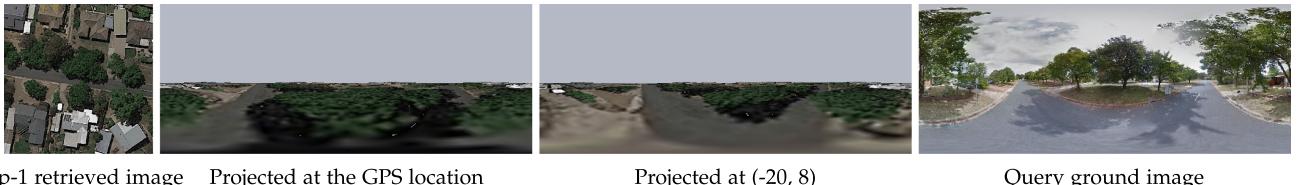


Fig. 15. A failure case of the location estimation produced by our method in fine-grained 360°-FoV camera geo-localization (given orientation). In this case, our method mistakenly regards the satellite image projected at the coordinate (-20, 8) rather than the GPS location as the most similar to the query image.

are conducted on the same image set. We found that around 8% of the query images cannot be localized when their FoV decreases from 360° to 180°, marked as “Neither” in Fig. 14. This is due to the increasing ambiguity when FoV decreases.

We present a failure case of our method in Fig. 15. As seen in the middle of the projected satellite image at the GPS location, the pixels from the tree canopies occlude the road pixels underneath. Those pixels are not similar to the pixels in the query image at corresponding positions. However, humans still can infer that the two images are at the same location, by leveraging the visible surrounding content and the geometric correspondences. In contrast, the SSIM metric only measures the overall pixel-wise similarity between the two images. It cannot handle pixels that are visually different but geometrically consistent. Thus, it yields a low similarity value between the image projected at the GPS location and the query image. Instead, it regards the image projected at a different coordinate (-20, 8) as the most similar one to the query image, since the road pixels in the two images align well and they account for a significant fraction of the image. As a result, our method fails to localize the correct location in this case.

Orientation Estimation. We also investigate the accuracy of the orientation estimation of our method when the camera location is known. Here, the ground-truth location information is obtained from the previous stage of the user study, rather than the noisy GPS provided by the dataset. We generate query images by rotating the aligned images randomly along the azimuth direction and storing the rotation as the ground-truth orientation. The results in the second row of Table 8 indicate that the accuracy of orientation estimation is satisfactory when the query images have a 360° FoV. However, the accuracy decreases significantly when the FoV of the query images decreases to 180°. Fig. 16 presents a failure case of the orientation estimation for 180° FoV images. As shown in the figure, the scene contents at the estimated orientation are very similar to those at the ground-truth orientation (i.e., the query image), making the orientation estimation rather ambiguous.

Joint Location and Orientation Estimation. Here we evaluate the overall performance of our method on joint location and orientation estimation for fine-grained camera localization. The

experiments are conducted on the images of which users can confidently determine whether the GPS location or our estimated location is correct, and we use the ground-truth provided by the user study participants for this evaluation. For location estimation, we consider localization within ± 5 pixels (± 1.4 meters) as a success. The results are presented in the third row of Table 8. The first row of Table 8 shows the localization performance of our method when orientation is given. It can be seen that the joint estimation task is much more challenging, since the ambiguity of the problem becomes even severe.

6 LIMITATIONS

Our method assumes that the image plane of the query camera is perpendicular to the ground plane. The tilt (elevation) and roll angles of the ground cameras in the current datasets are approximately zero. Thus, the sensitivity of our method to these angles is not investigated.

For the fine-grained 3-DoF camera localization from satellite images, there is considerable space for our method to improve. The pixel-wise SSIM similarity used by our method is not robust to severe occlusions between the ground and satellite images. It could be replaced by a higher-level content similarity metric. Improving the runtime is also an important aspect. For example, a more sophisticated searching strategy instead of the exhaustive searching is required to find the camera location within a selected region.

7 CONCLUSION

In this paper, we proposed an effective two-stage algorithm for ground-to-satellite image geo-localization, which can handle complex cases where neither location nor orientation are known. In contrast to many existing methods, our algorithm provides accurate 3-DoF (location and orientation) camera localization results. Specifically, our method includes a coarse localization step which retrieves the most similar satellite image from the database given a query image, and a fine-grained

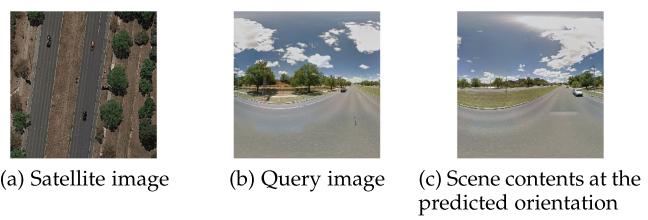


Fig. 16. A failure case of the orientation estimation produced by our method in fine-grained 180°-FoV camera localization (given location). The scene structure at the predicted orientation is very similar to that of the query image.

localization step which computes the displacement between the query ground camera location and the retrieved satellite image center. The orientation alignment between satellite and ground images in both steps is evaluated. In contrast to our previous works, we established more authentic geometric correspondences between satellite and ground images using a projective transform. By exploring the geometric correspondences, we successfully boost the coarse localization performance in terms of higher location recalls, and provide a novel method for accurate 3-DOF camera localization.

ACKNOWLEDGMENTS

The authors would like to thank AE and all reviewers for their time and patience in reviewing this paper and their constructive suggestions. Yujiao Shi and Xin Yu contributed equally to this work.

REFERENCES

- [1] S. Workman, R. Souvenir, and N. Jacobs, "Wide-area image geo-localization with aerial reference imagery," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 3961–3969.
- [2] M. Zhai, Z. Bessinger, S. Workman, and N. Jacobs, "Predicting ground-level scene layout from aerial imagery," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 4132–4140.
- [3] N. N. Vo and J. Hays, "Localizing and orienting street views using overhead imagery," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 494–509.
- [4] S. Hu, M. Feng, R. M. H. Nguyen, and G. H. Lee, "CVM-Net: Cross-view matching network for image-based ground-to-aerial geo-localization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7258–7267.
- [5] L. Liu and H. Li, "Lending orientation to neural networks for cross-view geo-localization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 5617–5626.
- [6] K. Regmi and M. Shah, "Bridging the domain gap for ground-to-aerial image matching," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 470–479.
- [7] S. Cai, Y. Guo, S. Khan, J. Hu, and G. Wen, "Ground-to-aerial image geo-localization with a hard exemplar reweighting triplet loss," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 8390–8399.
- [8] Y. Shi, X. Yu, L. Liu, T. Zhang, and H. Li, "Optimal feature transport for cross-view image geo-localization," in *Proc. AAAI Conf. Artif. Intell.*, 2020, pp. 11990–11997.
- [9] Y. Shi, L. Liu, X. Yu, and H. Li, "Spatial-aware feature aggregation for image based cross-view geo-localization," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 10090–10100.
- [10] Y. Shi, X. Yu, D. Campbell, and H. Li, "Where am I looking at? Joint location and orientation estimation by cross-view matching," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 4063–4071.
- [11] L. Svärm, O. Enqvist, F. Kahl, and M. Oskarsson, "City-scale localization for cameras with known vertical direction," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 7, pp. 1455–1461, Jul. 2017.
- [12] M. Donoser and D. Schmalstieg, "Discriminative feature-to-point matching in image-based localization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 516–523.
- [13] Y. Li, N. Snavely, D. Huttenlocher, and P. Fua, "Worldwide pose estimation using 3D point clouds," in *Proc. Eur. Conf. Comput. Vis.*, 2012, pp. 15–29.
- [14] Y. Li, N. Snavely, and D. P. Huttenlocher, "Location recognition using prioritized feature matching," in *Proc. Eur. Conf. Comput. Vis.*, 2010, pp. 791–804.
- [15] S. Lynen, T. Sattler, M. Bosse, J. A. Hesch, M. Pollefeys, and R. Siegwart, "Get out of my lab: Large-scale, real-time visual-inertial localization," *Robot. Sci. Syst.*, vol. 1, 2015, Art. no. 1.
- [16] T. Sattler, M. Havlena, F. Radenovic, K. Schindler, and M. Pollefeys, "Hyperpoints and fine vocabularies for large-scale location recognition," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 2102–2110.
- [17] T. Sattler, B. Leibe, and L. Kobbelt, "Efficient & effective prioritized matching for large-scale image-based localization," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 9, pp. 1744–1756, Sep. 2017.
- [18] H. Taira *et al.*, "InLoc: Indoor visual localization with dense matching and view synthesis," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 7199–7209.
- [19] B. Zeisl, T. Sattler, and M. Pollefeys, "Camera pose voting for large-scale image-based localization," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2015, pp. 2704–2712.
- [20] T. Sattler *et al.*, "Benchmarking 6DOF outdoor visual localization in changing conditions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 8601–8610.
- [21] C. Toft *et al.*, "Long-term visual localization revisited," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 4, pp. 2074–2088, Apr. 2022.
- [22] L. Liu, H. Li, and Y. Dai, "Stochastic attraction-repulsion embedding for large scale image localization," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 2570–2579.
- [23] S. Cao and N. Snavely, "Minimal scene descriptions from structure from motion models," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 461–468.
- [24] V. Larsson, J. Fredriksson, C. Toft, and F. Kahl, "Outlier rejection for absolute pose estimation with known orientation," in *Proc. Brit. Mach. Vis. Conf.*, 2016.
- [25] Y. Zhou *et al.*, "DA4AD: End-to-end deep attention-based visual localization for autonomous driving," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 271–289.
- [26] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "NetVLAD: CNN architecture for weakly supervised place recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 5297–5307.
- [27] Z. Chen *et al.*, "Deep learning features at scale for visual place recognition," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 3223–3230.
- [28] S. Lowry *et al.*, "Visual place recognition: A survey," *IEEE Trans. Robot.*, vol. 32, no. 1, pp. 1–19, Feb. 2016.
- [29] T. Sattler, M. Havlena, K. Schindler, and M. Pollefeys, "Large-scale location recognition and the geometric burstiness problem," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 1582–1590.
- [30] A. Torii, R. Arandjelovic, J. Sivic, M. Okutomi, and T. Pajdla, "24/7 place recognition by view synthesis," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 1808–1817.
- [31] T. Sattler *et al.*, "Are large-scale 3D models really necessary for accurate visual localization?," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1637–1646.
- [32] H. J. Kim, E. Dunn, and J.-M. Frahm, "Learned contextual feature reweighting for image geo-localization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 3251–3260.
- [33] H. Noh, A. Araujo, J. Sim, T. Weyand, and B. Han, "Large-scale image retrieval with attentive deep local features," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 3456–3465.
- [34] Y. Ge, H. Wang, F. Zhu, R. Zhao, and H. Li, "Self-supervising fine-grained region similarities for large-scale image localization," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 369–386.
- [35] M. Cummins and P. Newman, "FAB-MAP: Probabilistic localization and mapping in the space of appearance," *Int. J. Robot. Res.*, vol. 27, no. 6, pp. 647–665, 2008.
- [36] D. Gálvez-López and J. D. Tardos, "Bags of binary words for fast place recognition in image sequences," *IEEE Trans. Robot.*, vol. 28, no. 5, pp. 1188–1197, Oct. 2012.
- [37] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos, "ORB-SLAM: A versatile and accurate monocular SLAM system," *IEEE Trans. Robot.*, vol. 31, no. 5, pp. 1147–1163, Oct. 2015.
- [38] F. Castaldo, A. Zamir, R. Angst, F. Palmieri, and S. Savarese, "Semantic cross-view matching," in *Proc. IEEE Int. Conf. Comput. Vis. Workshops*, 2015, pp. 1044–1052.
- [39] T.-Y. Lin, S. Belongie, and J. Hays, "Cross-view image geo-localization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2013, pp. 891–898.
- [40] A. Mousavian and J. Kosecka, "Semantic image based geolocation given a map," 2016, *arXiv:1609.00278*.
- [41] S. Workman and N. Jacobs, "On the location dependence of convolutional neural network features," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, 2015, pp. 70–78.
- [42] Y. Tian, C. Chen, and M. Shah, "Cross-view image matching for geo-localization in urban environments," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 3608–3616.
- [43] B. Sun, C. Chen, Y. Zhu, and J. Jiang, "GeoCapsnet: Aerial to ground view image geo-localization using capsule network," 2019, *arXiv:1904.06281*.
- [44] S. Zhu, T. Yang, and C. Chen, "Revisiting street-to-aerial view image geo-localization and orientation estimation," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, 2021, pp. 756–765.

[45] A. Toker, Q. Zhou, M. Maximov, and L. Leal-Taixé, "Coming down to earth: Satellite-to-street view synthesis for geo-localization," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 6484–6493.

[46] P. W. Halligan, G. R. Fink, J. C. Marshall, and G. Vallar, "Spatial cognition: Evidence from visual neglect," *Trends Cogn. Sci.*, vol. 7, no. 3, pp. 125–133, 2003.

[47] K. Mania, S. Badariah, M. Coxon, and P. Watten, "Cognitive transfer of spatial awareness states from immersive virtual environments to reality," *ACM Trans. Appl. Percep.*, vol. 7, no. 2, pp. 1–14, 2010.

[48] K. Lenc and A. Vedaldi, "Understanding image representations by measuring their equivariance and equivalence," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 991–999.

[49] T. Cohen and M. Welling, "Group equivariant convolutional networks," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 2990–2999.

[50] D. E. Worrall, S. J. Garbin, D. Turmukhambetov, and G. J. Brostow, "Harmonic networks: Deep translation and rotation equivariance," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 7168–7177.

[51] C. Esteves, C. Allen-Blanchette, X. Zhou, and K. Daniilidis, "Polar transformer networks," in *Proc. Int. Conf. Learn. Representations*, 2018.

[52] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. Int. Conf. Learn. Representations*, 2014.

[53] X. Lu, Z. Li, Z. Cui, M. R. Oswald, M. Pollefeys, and R. Qin, "Geometry-aware satellite-to-ground image synthesis for urban areas," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 856–864.

[54] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2009, pp. 248–255.

[55] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.



Yujiao Shi received the BE and MS degrees in automation from the Nanjing University of Posts and Telecommunications, Nanjing, China, in 2014 and 2017, respectively. She is currently working toward the PhD degree in the College of Engineering and Computer Science (CECS), Australian National University. Her research interests include satellite image based geo-localization, novel view synthesis, and scene understanding.



Xin Yu received the BS degree in electronic engineering from the University of Electronic Science and Technology of China, Chengdu, China, in 2009, the PhD degree from the Department of Electronic Engineering, Tsinghua University, Beijing, China, in 2015, and the PhD degree from the College of Engineering and Computer Science, Australian National University, Canberra, Australia, in 2019. He is currently a senior lecturer with the University of Technology Sydney. His interests include computer vision and image processing.



Liu Liu received the BE degree from the School of Automation, Northwestern Polytechnical University, in 2012. He is currently working toward the PhD degree in the College of Engineering and Computer Science (CECS), Australian National University. He is also a member of the Australian Centre for Robotic Vision (ACRV) Group. He is from the honors college. His research interests include robotic vision, multi-view geometry, and visual localization.



Dylan Campbell received the BE degree in mechatronic engineering from the University of New South Wales, and the PhD degree in engineering from the Australian National University, in 2018 while concurrently working as a research assistant with the Cyber-Physical Systems Group, Data61–CSIRO. He is a research fellow with the Visual Geometry Group (VGG), University of Oxford. Prior to that, he was a research fellow with the Research School of Computer Science, Australian National University and the Australian Research Council Centre of Excellence in Robotic Vision. His research interests cover a range of computer vision and machine learning topics, including visual geometry, and differentiable optimization.



Piotr Koniusz received the BSc degree in telecommunications and software engineering from the Warsaw University of Technology, Poland, in 2004, and the PhD degree in computer vision from CVSSP, University of Surrey, U.K., in 2013. He is currently a senior research scientist with Machine Learning Research Group, Data61/CSIRO (former NICTA), and a senior honorary lecturer with the Australian National University (ANU). He was a postdoctoral researcher with the team LEAR, INRIA, France. His interests include few-shot learning, graph/tensor learning, RKHS, contrastive learning, etc. He has been recognised as an outstanding area chair for ICLR 2021/2022. He has served on the best paper committee of ICLR 2021, as an area chair for recent NeurIPS, ICML, CVPR, ICCV, ECCV, AAAI, etc., as a workshop co-chair for ACM Multimedia 2021 and ACCV 2022, and as a guest editor of the *Image and Vision Computing*.



Hongdong Li is currently a professor with ANU. He is also a founding chief investigator for the Australia Centre of Excellence for Robotic Vision (ACRV). His research interests include 3D vision reconstruction, structure from motion, multi-view geometry, as well as applications of optimization methods in computer vision. Prior to 2010, he was with NICTA working on the "Australia Bionic Eyes" project. He is an associate editor of the *IEEE Transactions on Pattern Analysis and Machine Intelligence*, guest editor of the *International Journal of Computer Vision*, and area chair in recent year ICCV, ECCV, and CVPR conferences. He was a program chair for ACRA 2015 – Australia Conference on Robotics and Automation, and a program co-chair for ACCV 2018 – Asian Conference on Computer Vision. He won a number of paper awards in computer vision and pattern recognition, and was the receipt for the CVPR 2012 Best Paper Award, the ICCV Marr Prize Honorable Mention in 2017, and a shortlist of the CVPR 2020 Best Paper Award.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.