

SliceMatch: Geometry-Guided Aggregation for Cross-View Pose Estimation

Lentsch, Ted; Xia, Zimin; Caesar, Holger; Kooij, Julian F.P.

DOI

[10.1109/CVPR52729.2023.01652](https://doi.org/10.1109/CVPR52729.2023.01652)

Publication date

2023

Document Version

Final published version

Published in

Proceedings - 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023

Citation (APA)

Lentsch, T., Xia, Z., Caesar, H., & Kooij, J. F. P. (2023). SliceMatch: Geometry-Guided Aggregation for Cross-View Pose Estimation. In L. O'Conner (Ed.), *Proceedings - 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023* (pp. 17225-17234). (Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition; Vol. 2023-June). IEEE.
<https://doi.org/10.1109/CVPR52729.2023.01652>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

SliceMatch: Geometry-guided Aggregation for Cross-View Pose Estimation

Ted Lentsch* Zimin Xia* Holger Caesar Julian F. P. Kooij
Intelligent Vehicles Group, Delft University of Technology, The Netherlands

{T.deVriesLentsch, Z.Xia, H.Caesar, J.F.P.Kooij}@tudelft.nl

Abstract

This work addresses cross-view camera pose estimation, i.e., determining the 3-Degrees-of-Freedom camera pose of a given ground-level image w.r.t. an aerial image of the local area. We propose SliceMatch, which consists of ground and aerial feature extractors, feature aggregators, and a pose predictor. The feature extractors extract dense features from the ground and aerial images. Given a set of candidate camera poses, the feature aggregators construct a single ground descriptor and a set of pose-dependent aerial descriptors. Notably, our novel aerial feature aggregator has a cross-view attention module for ground-view guided aerial feature selection and utilizes the geometric projection of the ground camera's viewing frustum on the aerial image to pool features. The efficient construction of aerial descriptors is achieved using precomputed masks. SliceMatch is trained using contrastive learning and pose estimation is formulated as a similarity comparison between the ground descriptor and the aerial descriptors. Compared to the state-of-the-art, SliceMatch achieves a 19% lower median localization error on the VIGOR benchmark using the same VGG16 backbone at 150 frames per second, and a 50% lower error when using a ResNet50 backbone.

1. Introduction

Cross-view camera pose estimation aims to estimate the 3-Degrees-of-Freedom (3-DoF) ground camera pose, i.e., planar location and orientation, by comparing the captured ground-level image to a geo-referenced overhead aerial image containing the camera's local surroundings. In practice, the local aerial image can be obtained from a reference database using any rough localization prior, e.g., Global Navigation Satellite Systems (GNSS), image retrieval [19], or dead reckoning [13]. However, this prior is not necessarily accurate, for example, GNSS can contain errors up to tens of meters in urban canyons [2, 48, 49]. The cross-view formulation provides a promising alternative to

* indicates equal contribution.

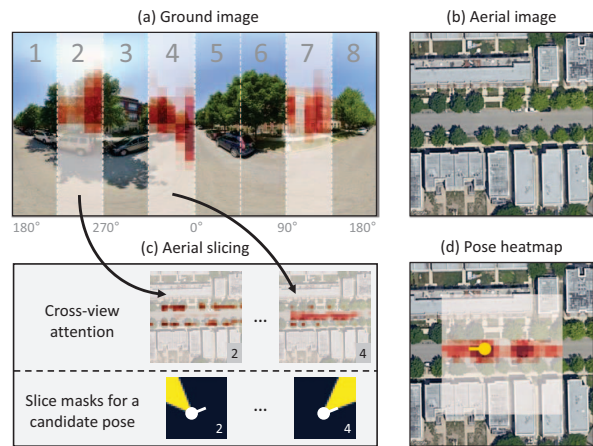


Figure 1. SliceMatch identifies for a ground-level image (a) its camera's 3-DoF pose within a corresponding aerial image (b). It divides the camera's Horizontal Field-of-View (HFOV) into 'slices', i.e., vertical regions in (a). After self-attention, our novel aggregation step (c) applies cross-view attention to create ground slice-specific aerial feature maps. To efficiently test many candidate poses, the slice features are aggregated using pose-dependent aerial slice masks that represent the camera's sliced HFOV at that pose. The slice masks for each pose are precomputed. All aerial pose descriptors are compared to the ground descriptor, resulting in a dense scoring map (d). Our output is the best-scoring pose.

ground-level camera pose estimation techniques that require detailed 3D point cloud maps [31] or semantic maps [3, 43], since the aerial imagery provides continuous coverage of the Earth's surface including the area where accurate point clouds are difficult to collect. Moreover, acquiring up-to-date aerial imagery is less costly than maintaining and updating large-scale 3D point clouds or semantics maps.

Recently, several works have addressed cross-view camera localization [55] or 3-DoF pose estimation [33, 36, 44, 50]. Roughly, those methods can be categorized into global image descriptor-based [50, 55] and dense pixel-level feature-based [33, 36, 44] methods. Global descriptor-based methods take advantage of the compactness of the image representation and often have relatively fast inference time [50, 55]. Dense pixel-level feature-based meth-

ods [33, 36, 44] are potentially more accurate as they preserve more details in the image representation. They use the geometric relationship between the ground and aerial view to project features across views and estimate the camera pose via computationally expensive iterations. Aiming for both accurate and efficient camera pose estimation, in this work, we improve the global descriptor-based approach and enforce feature locality in the descriptor.

We observe several limitations in existing global descriptor-based cross-view camera pose estimation methods [50, 55]. First, they rely on the aerial encoder to encode all spatial context and the aerial encoder has to learn how to aggregate local information, e.g., via the SAFA module [34], into the global descriptor, without accessing the information in the ground view or exploiting geometric constraints between the ground-camera viewing frustum and the aerial image. Second, existing global descriptor-based methods for cross-view localization [50, 55] do not explicitly consider the orientation of the ground camera in their descriptor construction. As a result, they either do not estimate the orientation [55] or require multiple forward passes on different rotated samples to infer the orientation [50]. Third, existing global descriptors-based methods [50, 55] are not trained discriminatively against different orientations. Therefore, the learned features may be less discriminative for orientation prediction.

To address the observed gaps, we devise a novel, accurate, and efficient method for cross-view camera pose estimation called SliceMatch (see Figure 1). Its novel aerial feature aggregation explicitly encodes directional information and pools features using known camera geometry to aggregate the extracted aerial features into an aerial global descriptor. The proposed aggregation step ‘slices’ the ground Horizontal Field-of-View (HFoV) into orientation-specific descriptors. For each pose in a set of candidates, it aggregates the extracted aerial features into corresponding aerial slice descriptors. The aggregation uses cross-view attention to weigh aerial features w.r.t. to the ground descriptor, and exploits the geometric constraint that every vertical slice in the ground image corresponds to an azimuth range extruding from the projected ground camera position in the aerial image. The feature extraction is done only once for constructing the descriptors for all pose candidates, resulting in fast training and inference speed. We contrastively train the model by pairing the ground image descriptor with aerial descriptors at different locations and orientations. Hence, the model learns to extract discriminative features for both localization and orientation estimation.

Contributions: **i)** A novel aerial feature aggregation step that uses a cross-view attention module for ground-view guided aerial feature selection, and the geometric relationship between the ground camera’s viewing frustum and the aerial image to construct pose-dependent aerial descrip-

tors. **ii)** SliceMatch’s design allows for efficient implementation, which runs significantly faster than previous state-of-the-art methods. Namely, for an input ground-aerial image pair, SliceMatch extracts dense features only once, aggregates aerial descriptors at a set of poses without extra computation, and compares the aerial descriptor of each pose with the ground descriptor. **iii)** Compared to the previous state-of-the-art global descriptor-based cross-view camera pose estimation method, SliceMatch constructs orientation-aware descriptors and adopts contrastive learning for both locations and orientations. Powered by the above designs, SliceMatch sets the new state-of-the-art for cross-view pose estimation on two commonly used benchmarks.

2. Related Work

Here, we review the work most related to SliceMatch.

Cross-view image retrieval is the task of finding matching aerial image patches from a reference database for a query ground image. The location of the retrieved aerial patch can be used as a localization estimate for the query image [17, 46, 47]. In general, this task is done by creating a global image descriptor for the query ground image and each reference aerial patch [10, 18, 21, 34, 37, 40, 51, 54, 55]. Different approaches have been proposed to build discriminative descriptors. CVM-Net [10] uses NetVLAD [1] to build viewpoint invariant descriptors. In [34], spatial attention modules are used to extract corresponding features across views. L2TLR [51] exploits the positional encoding of Transformers [41] to learn geometric correspondences between the ground and aerial view. TransGeo [53] uses attention-guided non-uniform cropping to only pay attention to informative regions in Transformers. Apart from advanced architectures, several works [16, 26, 34, 36, 40] have tried to synthesize one view using another to bridge the domain gap between the ground and aerial images. Besides, some works [34, 38, 45] use the geometric relationship between vertical lines in the ground image and azimuth directions in the aerial image to ease the learning or to estimate the orientation of the ground camera [35, 36, 54]. A few works have tried to explicitly enforce feature locality in global representations [29, 45], but they assume that the camera is located at the center of the aerial image. This limits the generalization of these methods to pose estimation.

Cross-view camera pose estimation works [33, 36, 42, 50, 55] go a step further than retrieval and aim to determine the location and orientation of the ground camera in the matching aerial image. A landmark graph matching-based method is used in [42], but a separate object detector is needed. In [52], the semantic segmentation of the ground-level image is compared to the ground-level semantic map predicted from the aerial image for estimating the location and orientation of the ground camera. Recently, [55] proposes a model that first retrieves an aerial image

for a query image and then uses a multilayer perceptron to regress the query’s location using the global image descriptors. Later, [50] formulates the localization problem as a multi-class classification problem and their model produces a dense multi-modal distribution to capture localization ambiguity. [33] projects the dense aerial features to ground perspective view based on homography and iteratively estimates the ground camera pose using Levenberg–Marquardt algorithm [15, 22]. However, the iterative process is computationally expensive (e.g. ~ 2 frames per second [33]) and it requires an accurate initial estimate to converge to a good local optimum. In [44], the ground image is fused with LiDAR data for iterative pose estimation. [36] samples aerial patches in the retrieved aerial image and applies a projective transformation on each sampled patch. Localization is achieved by selecting the location of the patch with the highest similarity to the ground image in the feature space. However, the computation increases linearly with the number of sampled locations. Lastly, [11] estimates only the orientation using the known location of the ground camera in the aerial image. So far, existing end-to-end methods are either global image descriptor-based [50, 55] or dense local feature-based [11, 33, 36, 44]. We argue that enforcing the right amount of feature locality in global image descriptors can be a promising direction toward the accuracy and runtime requirements of autonomous driving [27].

Bridging ground and aerial views is relevant in many other research directions. For example, ground-to-Bird’s-Eye-View (BEV) semantic mapping [25, 28, 30] tries to map the semantics in the ground perspective view to BEV given the known ground camera pose and intrinsics. [28] introduces a dense Transformer layer to condense the ground image features along the vertical dimension and then predicts features along the depth axis in a polar coordinate system. [25] lifts the ground features along the depth dimension and then projects them to BEV. In [30], a Transformer is used for the ground image column-to-BEV polar ray mapping. A few works [20, 32] synthesize ground-level panoramas from aerial images. All aforementioned works utilize the geometric relationship between the ground-level camera’s frustum and BEV. Indoor localization using floor maps [8, 9, 23] is also a relevant research direction. LaLaLoc [9] renders the ground view using the 2D floor plan and optimizes the ground camera pose estimation using the global representation of the ground-level query image and the global representation of the rendered view. LaLaLoc++ [8] removes the need for explicit modeling or rendering in [9] by introducing a global floor plan comprehension module. LASER [23] constructs a geometrically-structured latent space by aggregating viewing ray features for Monte Carlo Localization in 2D point cloud floor maps. It needs the occupancy boundaries information (e.g. walls) to form the 2D point cloud input. Thus it is not directly generalizable to aerial imagery.

3. Methodology

We explain the cross-view camera pose estimation task, our SliceMatch method, and its novel aggregation step.

3.1. Cross-View Camera Pose Estimation

Given a ground-level image I_g and a square overhead aerial image I_a that contains the local surroundings of I_g , we aim to determine the 3-DoF pose, $\xi = (u, v, \theta)$, of the ground camera that captured I_g . Here, $(u, v) \in [0, 1]^2$ are the image coordinates in I_a , and $\theta \in [0, 360^\circ)$ is the camera orientation, i.e., the angle from the North direction clockwise to the center line (the ‘front’ direction) of the ground camera projected onto the aerial view. Ground images can either be panoramic or have a limited HFoV. Similar to [33], we assume that the ground camera’s pitch and roll are small.

3.2. SliceMatch Overview

SliceMatch explicitly separates feature extraction and aggregation, where the latter exploits geometric knowledge on how the ground camera’s viewing frustum projects on the aerial image. In SliceMatch, pose estimation is formulated as an efficient process that compares aerial descriptors for a set $\Xi = \{\xi^1, \dots, \xi^K\}$ of K candidate poses to the ground image descriptor. During training, the set consists of K_{train} poses at a fixed uniform grid in 3-DoF pose space. During inference, we use K_{test} poses ($K_{test} > K_{train}$), and the predicted pose is the candidate for which its aerial descriptor is most similar to the ground descriptor. See Figure 2 for an overview of the method. We discuss each step next.

Feature extractor: Input images I_g and I_a are first mapped to feature maps, $z_g = f_g(I_g) \in \mathbb{R}^{H \times W \times C}$ and $z_a = f_a(I_a) \in \mathbb{R}^{L \times L \times C}$, where f_g and f_a can be any convolutional backbone (e.g. VGG [39] or ResNet [7]). We adopt the commonly used setup that f_g and f_a have the same architecture without weight-sharing [50, 55]. We seek translational equivariance in our encoders, and thus do not focus on Vision Transformers [5] in this work.

Feature aggregator: Our novel aggregator step efficiently constructs a single ground and multiple pose-dependent aerial descriptors from the extracted image features through the use of ‘slices’. In our work, each slice represents a non-overlapping range in the azimuth viewing direction, and is used to aggregate the local image features within that azimuth range. In the ground view, a slice thus corresponds to a vertical rectangular region in the image/feature map, and in the aerial view, it is a triangle-shaped region extending from a candidate pose (see Figure 1). This will be explained in more detail in Section 3.3. We refer to an aggregated feature in a single slice as a ground/aerial *slice descriptor*, containing the visual information for that viewing direction. Likewise, we refer to a ground/aerial *global descriptor* as the concatenation of the

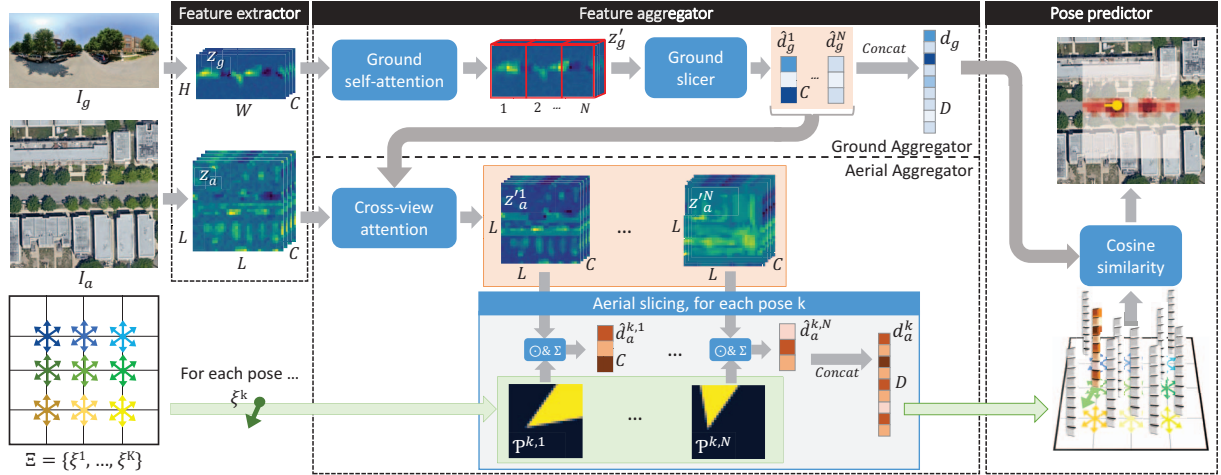


Figure 2. **The SliceMatch pipeline.** The input to SliceMatch is a pair a ground-aerial image pair and a set of K candidate ground camera poses. SliceMatch consists of ground and aerial feature extractors, feature aggregators, and a pose predictor. In the shown output image, we overlaid the matching scores for all poses on the input aerial image. The predicted pose is the one with highest score.

slice descriptors of all pose-relative orientations, representing the full HFoV of the ground camera.

Concretely, the extracted feature maps z_g and z_a are fed into our heterogeneous feature aggregators, as shown in Figure 2. The ground aggregator $agg_g(z_g)$ generates a set¹ $\hat{\Delta}_g = \{\hat{d}_g^1, \dots, \hat{d}_g^N\}$ of C -dimensional slice descriptors \hat{d}_g^n for N azimuth directions, where N is a hyperparameter for the number of slices. The ground global descriptor $d_g = \text{Concat}(\hat{d}_g^1, \dots, \hat{d}_g^N)$ is thus a vector of length $D = N \cdot C$. The aerial aggregator $agg_a(z_a, \hat{\Delta}_g, \Xi)$ receives the aerial features z_a , the ground slice descriptors $\hat{\Delta}_g$, and the set of K poses Ξ . It generates $\Delta_a = \{d_a^1, \dots, d_a^K\}$, the set of K pose-dependent aerial global descriptors $d_a^k \in \mathbb{R}^D$. Section 3.3 will discuss both aggregators in detail.

Pose predictor: The pose predictor receives the ground global descriptor d_g , and the set Δ_a that contains the K aerial global descriptors corresponding to the candidate poses in set Ξ . We compute the cosine similarity c^k between d_g and all $d_a^k \in \Delta_a$ and, during inference, use ξ^k corresponding to the highest similarity value $c_{max}^k = \max(c^1, \dots, c^K)$ as the predicted pose. Note that similar to [50], we obtain a heatmap that can express multimodal pose estimation ambiguity, which can be beneficial for downstream fusion.

Loss Function: We modify the infoNCE loss [24] from contrastive representation learning [12] to train SliceMatch. Using $K = K_{train}$ training poses, our loss \mathcal{L} is defined as,

$$\mathcal{L} = -\log \left(\frac{\exp(c^{GT}/\tau)}{\frac{\alpha}{K} \sum_{k=1}^K \exp(c^k/\tau) + \exp(c^{GT}/\tau)} \right). \quad (1)$$

¹Note that we use \hat{d} and $\hat{\Delta}$ (with hat) to indicate *slice* descriptors/sets, and d and Δ (without hat) to indicate *global* descriptors/sets.

In Equation (1), α is our introduced hyperparameter that weighs the contribution of K poses to the learning. Variable c^{GT} is the cosine similarity between d_g and d_a^{GT} at ξ^{GT} , and c^k is that between d_g and d_a^k at ξ^k . Hyperparameter τ is proposed in [24]. The original infoNCE loss in [24] can be acquired using $\alpha = K$. With \mathcal{L} , we contrast the ground truth pose with K_{train} other poses at different locations and orientations, thus the model learns to extract discriminative features for both location and orientation prediction.

3.3. Geometry-Guided Cross-View Aggregation

Here, we describe the novel aggregation step in more detail. Unlike the SAFA module [34] used in [50, 55], our aggregation uses geometric knowledge on how the views should spatially relate. Ground-to-aerial attention further improves quality, as the visual information in each ground slice informs what aerial features are relevant to produce the corresponding aerial slice descriptors, thus promoting shared features specific to each viewing direction.

3.3.1 Ground Feature Aggregator

To summarize the important features in each vertical slice in the ground camera’s viewing frustum, we construct our ground feature aggregator $agg_g(z_g)$ with a self-attention module and a feature slicer. Since not all information in ground image I_g will be present in the aerial image I_a (e.g. sky and transient objects), the self-attention module re-weights z_g along the spatial dimensions H and W ,

$$z'_g = \mathcal{M}_g \odot z_g, \quad \mathcal{M}_g = \text{Sigmoid}(\text{Conv}_{1 \times 1}(z_g)). \quad (2)$$

Here, \mathcal{M}_g is a learned mask with shape $H \times W \times 1$ that re-weights the ground feature map z_g into z'_g . The *Sigmoid*

operation enforces the weights in \mathcal{M}_g are between 0 and 1. The \odot denotes element-wise multiplication, with the ability to broadcast the mask \mathcal{M}_g over all channels of z_g .

The ground slicer then divides z'_g into N vertical slices, cutting the feature map along the horizontal (azimuth) direction. For each slice, a normalized slice descriptor is computed by averaging all features within the slice and applying L2 normalization. This results in the set $\hat{\Delta}_g = \{\hat{d}_g^1, \dots, \hat{d}_g^N\}$ of N ground slice descriptors. Each slice local descriptor thus represents the model's attended feature in the corresponding vertical slice (i.e. an azimuth range) in the ground camera's viewing frustum. The ground global descriptor is obtained by concatenating all N ground slice descriptors, i.e. $d_g = \text{Concat}(\hat{d}_g^1, \dots, \hat{d}_g^N)$.

3.3.2 Aerial Feature Aggregator

The aerial aggregator $agg_a(z_a, \hat{\Delta}_g, \Xi)$ has a similar role as the ground aggregator, but its feature selection is also conditioned on the ground slice descriptors $\hat{\Delta}_g$ using a cross-view attention module and the set of poses Ξ for geometry-guided feature aggregation.

Cross-view attention: Since in the ground view most content that is seen in the aerial view will be occluded, we propose a cross-view attention module to specifically extract the aerial features that should match the visible content of each ground slice. In detail, we match the C -dimensional aerial feature $z_a^{i,j}$ at each spatial location (i, j) with $1 \leq i \leq L, 1 \leq j \leq L$ in the aerial feature map z_a to each ground slice descriptor $\hat{d}_g^n \in \hat{\Delta}_g$ to acquire a similarity score map S^n of size $L \times L$, where $S^{n,i,j} = \text{Sim}(\hat{d}_g^n, z_a^{i,j})$. In total, there are N similarity score maps, i.e. one for each ground slice descriptor. Then, we treat each S^n as extra features and concatenate it along the feature dimension with aerial feature map z_a [50], and use these extended features to produce a cross-view attention mask,

$$\mathcal{M}_a^n = \text{Sigmoid}(\text{Conv}_{1 \times 1}(\text{Concat}(z_a, S^n))). \quad (3)$$

We thus get in total N cross-view masks \mathcal{M}_a^n . Each of these denotes the importance of the aerial features w.r.t. the n -th ground slice descriptor \hat{d}_g^n . Finally, we re-weigh z_a for each ground slice descriptor, giving us N re-weighted aerial feature maps $z_a'^n$ of size $L \times L \times C$, i.e. $z_a'^n = \mathcal{M}_a^n \odot z_a$.

Geometry-guided feature aggregation: Finally, the K pose-dependent aerial descriptors d_a^k can be constructed for the candidate poses in Ξ . For each pose ξ^k , we can precompute N aerial slice masks $\mathcal{P}^{k,n} \in [0, 1]^{L \times L}$, $1 \leq n \leq N$. The slice mask $\mathcal{P}^{k,n}$ expresses the geometry of the ground camera's viewing frustum in the aerial feature map for the n -th orientation slice, assuming that the camera would have the k -th pose. Each cell in the slice mask contains a value in the range $[0, 1]$ proportional to how much of that cell intersects this frustum, so 1.0 for fully contained cells, 0.0 for

cells fully outside the frustum, and an intermediate value for cells that partially overlap.

With the slice masks, the n -th aerial slice descriptor at pose k can be computed efficiently. For each of the C channels, we compute a weighted average over all of the $L \times L$ spatial locations (i, j) in the feature map $z_a'^n$, using the elements of slice mask $\mathcal{P}^{k,n}$ as weights. After L2 normalization, we obtain aerial slice descriptor $\hat{d}_a^{k,n}$,

$$\hat{d}_a^{k,n} = \text{Norm}\left(\frac{1}{\sum_{i,j} \mathcal{P}_{i,j}^{k,n}} \sum_{i,j} (\mathcal{P}_{i,j}^{k,n} \odot z_a'^n)_{i,j}\right). \quad (4)$$

Analogous to the ground view, the k -th pose's global descriptor is obtained using $d_a^k = \text{Concat}(\hat{d}_a^{k,1}, \dots, \hat{d}_a^{k,N})$.

Efficient implementation: A benefit of our proposed architecture is that the computational complexity of most operations is independent of the number of candidate poses K . The main cost to increase K , and therefore improve accuracy by testing more diverse poses at inference time, is to add more precomputed slice masks, and perform the additional multiplications and normalizations for Equation (4) and the final cosine similarity comparison. These are simple operations that can be highly optimized and parallelized in the implementation, and we will show that testing more candidate poses does not increase our runtime.

4. Experiments

We first introduce the used datasets and the evaluation metrics. After that, our implementation details and ablation studies are presented. Finally we quantitatively and qualitatively compare SliceMatch to state-of-the-art baselines.

4.1. Datasets

VIGOR dataset [55] contains geo-tagged ground-level panoramas and aerial images collected in 4 cities in the US. As defined in [55], each ground panorama has 1 positive and 3 semi-positive aerial images. An aerial image is positive if the ground camera's location is within the aerial image's center quarter area, otherwise, it is semi-positive. Importantly, we found that the original ground truth locations in [55] can contain errors up to 3 meters due to the use of wrong ground resolutions (0.114m/pixel) of the aerial images, thus we created and use here corrected labels, see details in Supplementary Material. For training and testing our method and baselines, we use positive aerial images and corrected ground truth (we reran all baselines since quantitative results with the new labels differ slightly from those reported in the literature). We adopt the same-area and cross-area splits from [55] to test the model's generalization to new measurements in the same cities and across different cities. Besides, we use the same-area training dataset of New York as a tuning split for the ablation study.

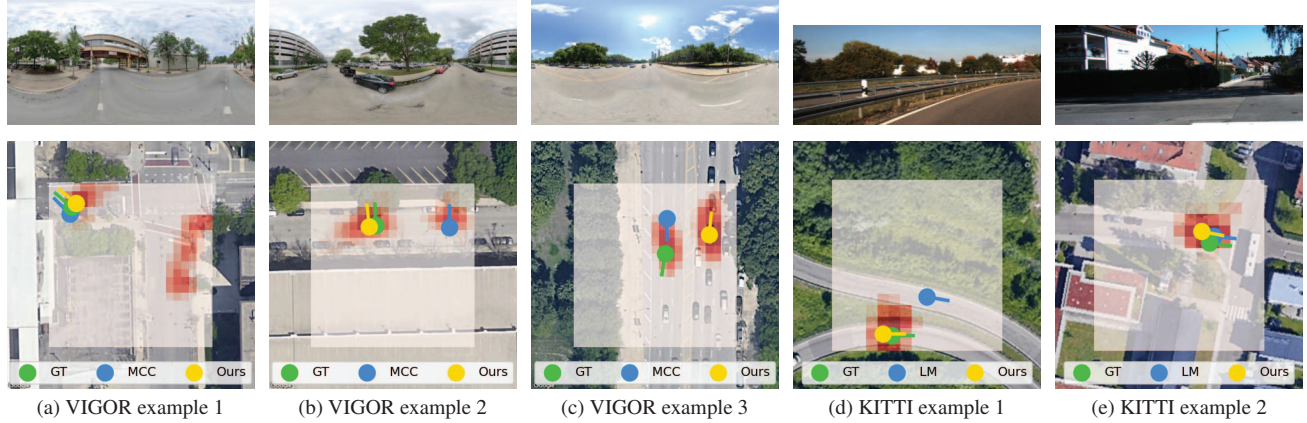


Figure 3. **Qualitative evaluation of SliceMatch on VIGOR [55] and KITTI [6,33].** Top row: input ground image. Bottom row: GT and pose estimation results overlaid on input aerial image. Red shading indicates highest similarity score between the ground descriptor and the aerial descriptors among all orientations at that location. (c) shows a SliceMatch failure: the best match is in the wrong mode.

KITTI dataset [6] contains ground-level images with a limited HFoV taken by a moving vehicle from different trajectories at different times and [33] augmented the dataset with aerial images. We use their split. The Training and Test1 sets are different measurements from the same region, while the Test2 set has been captured in a different region.

4.2. Evaluation Metrics

We follow the convention of [50] and report the mean and median error in meters between the predicted and ground truth location over all test image pairs. Similarly, for orientation prediction, we report the mean and median absolute angular difference between the predicted and ground truth orientation in degrees. Following [33], for the KITTI dataset, we additionally include the recall under a certain threshold for longitudinal (driving direction) and lateral localization error, and orientation estimation error. Our thresholds are set to 1m and 5m for localization and to 1° and 5° for orientation estimation.

4.3. Implementation Details

As in [50,55], we use VGG16 [39] up to stage 5 for the feature extractors f_g and f_a . The pooling operation of the last layer is removed. The spatial size of I_g is 320×640 on VIGOR dataset and 256×1024 on KITTI dataset, and that of I_a is 512×512 on both datasets. This results in feature maps with $H \times W = 20 \times 40 / 16 \times 64$ on VIGOR / KITTI, $L \times L = 32 \times 32$, and $C = 512$. The feature extractors do not share their weights and are pre-trained on ImageNet [4]. In Equation (2) and (3), $Conv_{1 \times 1}$ consists of two sequential convolution layers with a kernel size of 1 and a ReLU activation in between. During training, SliceMatch is trained end-to-end using Adam optimizer [14] with a learning rate of 1×10^{-5} , and we use a batch size of 4. To get a set of candidate camera poses Ξ , we use poses at a uniform grid of

7×7 locations \times 16 orientations on VIGOR, and 5×5 locations \times 16 orientations on KITTI during training. For inference, we use $21 \times 21 \times 64$ and $15 \times 15 \times 64$ poses, respectively. This results in $K_{train} = 784$ and $K_{test} = 28224$ on VIGOR, and $K_{train} = 400$ and $K_{test} = 14400$ on KITTI.

4.4. Baselines

We compare SliceMatch to state-of-the-art global descriptor-based methods Cross-View Regression (CVR) [55] and Multi-Class Classification (MCC) [50] on the VIGOR dataset². Since CVR does localization with known orientation and, in [50], MCC mainly focuses on localization, we compare SliceMatch to baselines for localization with known orientation and also for 3-DoF pose estimation. Following [50], we train CVR [55] for localization only (not retrieval) as it gives better localization results. On the KITTI dataset, SliceMatch is compared to dense local feature-based fine-grained image retrieval method DSM [35], and to iterative camera pose estimation method LM [33]. In [33], the LM method is trained and tested with a 20° prior on the ground camera’s orientation. We adopt the same setting and additionally provide the results with LM and SliceMatch trained and tested with unknown orientation. On both datasets, baselines are trained with inputs with the same size as used for SliceMatch.

4.5. Ablation Study

Before other experiments, we test on the VIGOR tuning set using $\alpha \in \{2, 4, 8, 16, K\}$ for the loss of Equation (1), and tune the number of slices N . We find $\alpha = 4$ gives the best result, yielding 0.48m improvement on the mean

²We re-trained and evaluated the existing baselines on our corrected ground truth locations (see Section 4.1 and our Supplementary Material). The improved ground truth and code for our model are available at <https://github.com/tudelft-iv/SliceMatch>.

localization error for our model compared to $\alpha = K$ in the original infoNCE loss [24]. If the number of slices N is small, the mean and median localization and orientation estimation errors increase (see Table 1). The model with $N = 1$ cannot infer orientation. When the width of the ground feature map W is not a multiple of N , we interpolate the ground feature map z_g to acquire \hat{d}_g^n . However, it can be seen that the performance saturates above 16 slices. Next, we tested SliceMatch without cross-view attention by dropping the concatenated S^n in Equation (3). Table 1 shows that including our proposed cross-view attention module brings a boost to both localization and orientation estimation performance. Thus, we include cross-view attention and use $\alpha = 4$ and $N = 16$ in our main experiments.

N	Cross-View Attention	↓ Location (m)		↓ Orientation (°)	
		Mean	Median	Mean	Median
1	X	12.73	11.51	-	-
4	✓	9.47	7.47	51.49	32.96
8	✓	9.16	6.81	37.68	15.58
16	✓	7.60	5.23	29.27	9.22
32	✓	8.14	5.31	32.01	10.31
16 [‡]	✓	8.08	5.44	31.05	11.02
16	X	7.93	5.81	29.50	12.32

Table 1. Location and orientation error for different slice number N values on the VIGOR tuning split. ‡ indicates model trained with original infoNCE loss [24]. Best performance in bold.

4.6. Same-Area Generalization

We test model generalization to new panoramic and limited HFoV ground images within the same area on VIGOR and KITTI. As shown in Table 2 Same-Area, SliceMatch surpasses CVR [55] and MCC [50] in terms of both localization with known orientation and 3-DoF camera pose estimation. Compared to MCC, in which location-wise discriminative features are learned, SliceMatch contrasts the learned global descriptors with aerial descriptors at different locations and orientations. Hence, it is more discriminative especially w.r.t. orientations, and has a 19% and 68% reduction in the median localization and median orientation error when the orientation of test ground images is unknown. We use VGG16 as our main backbone for a fair comparison to the baselines, though we note our localization and orientation error decreases even further when using ResNet50 as backbone. We show in Figure 3b that SliceMatch can express its multimodal uncertainty when the observed scene has a symmetric layout. However, it sometimes picks candidate poses at a wrong mode, resulting in large errors (see Figure 3c). Over all test samples, SliceMatch has a substantially lower median error than its mean error for both local-

ization and orientation estimation, indicating that the mean is skewed by such outliers. In practice, SliceMatch’s multimodal uncertainty could be resolved by applying downstream a probabilistic temporal filter on its output [50].

As shown in Table 3 Same-Area, on the KITTI dataset, both camera pose estimation methods, LM [33] and SliceMatch, surpass the fine-grained image retrieval-based method DSM [35]. When the orientation prior is present, SliceMatch has 34% and 62% lower mean and median localization error than LM [33], and its recall@1m and recall@5m is higher than that of LM [33] for localization in both lateral and longitudinal directions. Notably, since the ground images in KITTI view in the driving direction with a limited HFoV, finding the location along the longitudinal direction is more challenging than that for the lateral direction. Thus, recall for longitudinal direction is considerably lower than that for lateral direction, and this trend applies to all compared methods. The iterative refinement LM method [33] shows its advantage in orientation prediction when the strong orientation prior is present. We highlight that SliceMatch can work without this prior. In contrast, LM [33] relies on the projection of dense local features from the aerial view to ground view [33] and does not work when there is no same scene captured in the projected view and the ground view (see Table 3).

4.7. Cross-Area Generalization

Generalization to new ground images in different areas is a more difficult task than that in the same area since the test area can look very different from the training area (e.g. different cities in the VIGOR dataset). As shown in Table 2 Cross-Area, SliceMatch generalizes well under this challenging setting in terms of both localization and orientation estimation, while we observe more degeneration in the cross-area test performance of MCC [50]. MCC’s feature decoder receives the full scene information from its encoder, while SliceMatch divides the observed scene into slices and seeks per-slice discriminative features, resulting in more robustness against the change of the scene. Again, using a ResNet50 backbone further improves our results.

On KITTI Test2 set (Table 3 Cross-Area), SliceMatch achieves a lower median localization error than LM [33] when the 20° orientation prior is present in both training and testing. But our mean error is higher than LM [33] by 0.92m and LM [33] surpasses SliceMatch in orientation prediction when a strong prior is available. SliceMatch performs considerably better when no orientation prior is available as LM [33] gets stuck in local optima.

4.8. Runtime Analysis

We compare the runtime of SliceMatch to that of baselines on the same hardware, a single NVIDIA Tesla V100 GPU. For all baselines, we use the released code from their

Model	Backbone	Aligned Images	Same-Area				Cross-Area			
			↓ Location (m)		↓ Orientation (°)		↓ Location (m)		↓ Orientation (°)	
			Mean	Median	Mean	Median	Mean	Median	Mean	Median
CVR [55]	VGG16	✓	8.99	7.81	-	-	8.89	7.73	-	-
MCC [50]	VGG16	✓	6.94	3.64	-	-	9.05	5.14	-	-
SliceMatch (ours)	VGG16	✓	5.18	2.58	-	-	5.53	2.55	-	-
MCC [50]	VGG16	X	9.87	6.25	56.86	16.02	12.66	9.55	72.13	29.97
SliceMatch (ours)	VGG16	X	8.41	5.07	28.43	5.15	8.48	5.64	26.20	5.18
SliceMatch (ours)	ResNet50	X	6.49	3.13	25.46	4.71	7.22	3.31	25.97	4.51

Table 2. **Location and orientation estimation errors on VIGOR [55]**. *Aligned Images* means the ground image orientation is known. For *unaligned images*, the models estimate the 3-DoF ground camera pose. Best performance in **bold**.

Model	Area	Prior	↓ Location (m)		↑ Lateral (%)		↑ Long. (%)		↓ Orien. (°)		↑ Orien. (%)	
			Mean	Median	r@1m	r@5m	r@1m	r@5m	Mean	Median	r@1°	r@5°
DSM [35]	Same	20°	-	-	10.12	48.24	4.08	20.14	-	-	3.58	24.44
LM [33]	Same	20°	12.08	11.42	35.54	80.36	5.22	26.13	3.72	2.83	19.64	71.72
SliceMatch (ours)	Same	20°	7.96	4.39	49.09	98.52	15.19	57.35	4.12	3.65	13.41	64.17
LM [33]	Same	X	15.51	15.97	5.17	25.44	4.66	25.39	89.91	90.75	0.61	2.89
SliceMatch (ours)	Same	X	9.39	5.41	39.73	87.92	13.63	49.22	8.71	4.42	11.35	55.82
DSM [35]	Cross	20°	-	-	10.77	48.24	3.87	19.50	-	-	3.53	23.95
LM [33]	Cross	20°	12.58	12.11	27.82	72.89	5.75	26.48	3.95	3.03	18.42	71.00
SliceMatch (ours)	Cross	20°	13.50	9.77	32.43	86.44	8.30	35.57	4.20	6.61	46.82	46.82
LM [33]	Cross	X	15.50	16.02	5.60	25.60	5.64	25.76	89.84	89.85	0.60	2.65
SliceMatch (ours)	Cross	X	14.85	11.85	24.00	72.89	7.17	33.12	23.64	7.96	31.69	31.69

Table 3. **Location and orientation estimation error and recall on KITTI [6,33]**. *Prior* means the orientation is known with a certain amount of noise. *Long.* and *Orien.* are abbreviations for *Longitudinal* and *Orientation*, respectively. Best performance in **bold**. The results for DSM [35] are taken from [33] and we used the trained LM model provided by [33] for its evaluation.

authors. CVR [55] and MCC [50] are implemented in TensorFlow, LM [33] and our SliceMatch in PyTorch. The frames per second (FPS) are calculated by taking the average inference time per input pair over all test samples. On VIGOR, SliceMatch achieves an FPS of 167, which is considerably faster than global descriptor-based baselines: 50 FPS for CVR [55] for localization only, 29 FPS / 3 FPS for MCC [50] for localization only / pose estimation. On KITTI, SliceMatch runs at 156 FPS, while the local feature-based iterative method, LM [33] has 0.59 FPS. Importantly, the runtime of SliceMatch remains nearly constant as the number of used candidate poses K increases (we tested K up to 1×10^6), see details in Supplementary Material.

5. Conclusion

We have introduced SliceMatch, a novel, accurate, and efficient method for cross-view 3-DoF camera pose estimation. By splitting the HFoV into slices, our architecture can learn discriminative features in terms of both localization and orientation estimation. Our proposed aggregation can select the relevant aerial image features for each ground

view slice through cross-view attention, and we observe further accuracy gains by reweighing the terms in the infoNCE loss. With the same VGG backbone, SliceMatch achieves 19% and 62% lower median localization error than the previous state-of-the-art on the VIGOR and KITTI datasets. A better backbone improves SliceMatch’s performance even further, e.g. with ResNet50 its 50% lower median error on VIGOR sets a new state-of-the-art. To construct the global descriptor for a candidate pose, only an efficient weighted averaging over the aerial features is needed using precomputed masks (which represent the ground camera’s frustum geometry in the aerial view), achieving inference at more than 150 FPS. SliceMatch can include available priors in its candidate poses, e.g. for an initial orientation estimate, but does not require it. Future work will explore adapting pose candidates for temporal filtering and sensor fusion.

Acknowledgements

This work is part of the research program Efficient Deep Learning (EDL) with project number P16-25, which is (partly) financed by the Dutch Research Council (NWO).

References

- [1] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Paszler, and Josef Sivic. NetVLAD: CNN architecture for weakly supervised place recognition. In *CVPR*, pages 5297–5307, 2016. 2
- [2] Boaz Ben-Moshe, Elazar Elkin, Harel Levi, and Ayal Weissman. Improving accuracy of GNSS devices in urban canyons. In *CCCG*, pages 511–515, 2011. 1
- [3] Hao Cai, Zhaozheng Hu, Gang Huang, Dunyao Zhu, and Xiaocong Su. Integration of GPS, monocular vision, and high definition (HD) map for accurate vehicle localization. *Sensors*, 18(10):3270, 2018. 1
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 6
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 3
- [6] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *IJRR*, 32(11):1231–1237, 2013. 6, 8
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 3
- [8] Henry Howard-Jenkins and Victor Adrian Prisacariu. LaLaLoc++: Global floor plan comprehension for layout localisation in unvisited environments. In *ECCV*, pages 693–709, 2022. 3
- [9] Henry Howard-Jenkins, Jose-Raul Ruiz-Sarmiento, and Victor Adrian Prisacariu. LaLaLoc: Latent layout localisation in dynamic, unvisited environments. In *CVPR*, pages 10107–10116, 2021. 3
- [10] Sixing Hu, Mengdan Feng, Rang MH Nguyen, and Gim Hee Lee. CVM-Net: Cross-view matching network for image-based ground-to-aerial geo-localization. In *CVPR*, pages 7258–7267, 2018. 2
- [11] Wenmiao Hu, Yichen Zhang, Yuxuan Liang, Yifang Yin, Andrei Georgescu, An Tran, Hannes Kruppa, See-Kiong Ng, and Roger Zimmermann. Beyond geo-localization: Fine-grained orientation of street-view images by cross-view matching with satellite imagery. In *ACM Multimedia*, pages 6155–6164, 2022. 3
- [12] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. In *NeurIPS*, pages 18661–18673, 2020. 4
- [13] Thomas King, Holger Füllner, Matthias Transier, and Wolfgang Effelsberg. Dead-reckoning for position-based forwarding on highways. In *WIT*, pages 199–204, 2006. 1
- [14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [15] Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of applied mathematics*, 2(2):164–168, 1944. 3
- [16] Songlian Li, Zhigang Tu, Yujin Chen, and Tan Yu. Multi-scale attention encoder for street-to-aerial image geo-localization. *CAAI TIT*, 2022. 2
- [17] Tsung-Yi Lin, Yin Cui, Serge Belongie, and James Hays. Learning deep representations for ground-to-aerial geolocalization. In *CVPR*, pages 5007–5015, 2015. 2
- [18] Liu Liu and Hongdong Li. Lending orientation to neural networks for cross-view geo-localization. In *CVPR*, pages 5624–5633, 2019. 2
- [19] Stephanie Lowry, Niko Sünderhauf, Paul Newman, John J Leonard, David Cox, Peter Corke, and Michael J Milford. Visual place recognition: A survey. *T-RO*, 32(1):1–19, 2015. 1
- [20] Xiaohu Lu, Zuoyue Li, Zhaopeng Cui, Martin R Oswald, Marc Pollefeys, and Rongjun Qin. Geometry-aware satellite-to-ground image synthesis for urban areas. In *CVPR*, pages 859–867, 2020. 3
- [21] Xiufan Lu, Siqi Luo, and Yingying Zhu. It’s okay to be wrong: Cross-view geo-localization with step-adaptive iterative refinement. *TGRS*, 60:1–13, 2022. 2
- [22] Donald W Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *SIAM*, 11(2):431–441, 1963. 3
- [23] Zhixiang Min, Naji Khosravan, Zachary Bessinger, Manjunath Narayana, Sing Bing Kang, Enrique Dunn, and Ivaylo Boyadzhiev. LASER: Latent space rendering for 2d visual localization. In *CVPR*, pages 11122–11131, 2022. 3
- [24] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. In *arXiv preprint arXiv:1807.03748*, 2018. 4, 7
- [25] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *ECCV*, pages 194–210, 2020. 3
- [26] Krishna Regmi and Mubarak Shah. Bridging the domain gap for ground-to-aerial image matching. In *ICCV*, pages 470–479, 2019. 2
- [27] Tyler GR Reid, Sarah E Houts, Robert Cammarata, Graham Mills, Siddharth Agarwal, Ankit Vora, and Gaurav Pandey. Localization requirements for autonomous vehicles. *SAE IJ-CAV*, 2(12-02-03-0012):173–190, 2019. 3
- [28] Thomas Roddick and Roberto Cipolla. Predicting semantic map representations from images using pyramid occupancy networks. In *CVPR*, pages 11138–11147, 2020. 3
- [29] Royston Rodrigues and Masahiro Tani. Global assists local: Effective aerial representations for field of view constrained image geo-localization. In *WACV*, pages 3871–3879, 2022. 2
- [30] Avishkar Saha, Oscar Mendez Maldonado, Chris Russell, and Richard Bowden. Translating images into maps. In *ICRA*, 2021. 3
- [31] Torsten Sattler, Bastian Leibe, and Leif Kobbelt. Fast image-based localization using direct 2d-to-3d matching. In *ICCV*, pages 667–674, 2011. 1
- [32] Yujiao Shi, Dylan Campbell, Xin Yu, and Hongdong Li. Geometry-guided street-view panorama synthesis from satellite imagery. *T-PAMI*, 44(12):10009–10022, 2022. 3
- [33] Yujiao Shi and Hongdong Li. Beyond cross-view image retrieval: Highly accurate vehicle localization using satellite image. In *CVPR*, pages 17010–17020, 2022. 1, 2, 3, 6, 7, 8

- [34] Yujiao Shi, Liu Liu, Xin Yu, and Hongdong Li. Spatial-aware feature aggregation for cross-view image based geo-localization. In *NeurIPS*, pages 10090–10100, 2019. 2, 4
- [35] Yujiao Shi, Xin Yu, Dylan Campbell, and Hongdong Li. Where am I looking at? Joint location and orientation estimation by cross-view matching. In *CVPR*, pages 4064–4072, 2020. 2, 6, 7, 8
- [36] Yujiao Shi, Xin Yu, Liu Liu, Dylan Campbell, Piotr Koniusz, and Hongdong Li. Accurate 3-dof camera geo-localization via ground-to-satellite image matching. *T-PAMI*, 2022. 1, 2, 3
- [37] Yujiao Shi, Xin Yu, Liu Liu, Tong Zhang, and Hongdong Li. Optimal feature transport for cross-view image geo-localization. In *AAAI*, pages 11990–11997, 2020. 2
- [38] Yujiao Shi, Xin Yu, Shan Wang, and Hongdong Li. CVLNet: Cross-view semantic correspondence learning for video-based camera localization. In *ACCV*, 2022. 2
- [39] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 3, 6
- [40] Aysim Toker, Qunjie Zhou, Maxim Maximov, and Laura Leal-Taixé. Coming down to earth: Satellite-to-street view synthesis for geo-localization. In *CVPR*, pages 6488–6497, 2021. 2
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 2
- [42] Sebastiano Verde, Thiago Resek, Simone Milani, and Anderson Rocha. Ground-to-aerial viewpoint localization via landmark graphs matching. *SPL*, 27:1490–1494, 2020. 2
- [43] Huayou Wang, Changliang Xue, Yanxing Zhou, Feng Wen, and Hongbo Zhang. Visual semantic localization based on hd map for autonomous vehicles in urban scenarios. In *ICRA*, pages 11255–11261, 2021. 1
- [44] Shan Wang, Yanhao Zhang, and Hongdong Li. Satellite image based cross-view localization for autonomous vehicle. In *arXiv preprint arXiv:2207.13506*, 2022. 1, 2, 3
- [45] Tingyu Wang, Zhedong Zheng, Chenggang Yan, Jiyong Zhang, Yaoqi Sun, Bolun Zheng, and Yi Yang. Each part matters: Local patterns facilitate cross-view geo-localization. *TCSVT*, 32(2):867–879, 2021. 2
- [46] Scott Workman and Nathan Jacobs. On the location dependence of convolutional neural network features. In *CVPR Workshops*, pages 70–78, 2015. 2
- [47] Scott Workman, Richard Souvenir, and Nathan Jacobs. Wide-area image geolocalization with aerial reference imagery. In *ICCV*, pages 3961–3969, 2015. 2
- [48] Zimin Xia, Olaf Booij, Marco Manfredi, and Julian FP Kooij. Geographically local representation learning with a spatial prior for visual localization. In *ECCV Workshops*, pages 557–573, 2020. 1
- [49] Zimin Xia, Olaf Booij, Marco Manfredi, and Julian FP Kooij. Cross-view matching for vehicle localization by learning geographically local representations. *RA-L*, 6(3):5921–5928, 2021. 1
- [50] Zimin Xia, Olaf Booij, Marco Manfredi, and Julian FP Kooij. Visual cross-view metric localization with dense uncertainty estimates. In *ECCV*, pages 90–106, 2022. 1, 2, 3, 4, 5, 6, 7, 8
- [51] Hongji Yang, Xiufan Lu, and Yingying Zhu. Cross-view geo-localization with layer-to-layer transformer. In *NeurIPS*, 2021. 2
- [52] Menghua Zhai, Zachary Bessinger, Scott Workman, and Nathan Jacobs. Predicting ground-level scene layout from aerial imagery. In *CVPR*, pages 867–875, 2017. 2
- [53] Sijie Zhu, Mubarak Shah, and Chen Chen. TransGeo: Transformer is all you need for cross-view image geo-localization. In *CVPR*, pages 1162–1171, 2022. 2
- [54] Sijie Zhu, Taojiannan Yang, and Chen Chen. Revisiting street-to-aerial view image geo-localization and orientation estimation. In *WACV*, pages 756–765, 2021. 2
- [55] Sijie Zhu, Taojiannan Yang, and Chen Chen. VIGOR: Cross-view image geo-localization beyond one-to-one retrieval. In *CVPR*, pages 3640–3649, 2021. 1, 2, 3, 4, 5, 6, 7, 8