

# Revisiting Street-to-Aerial View Image Geo-localization and Orientation Estimation

Sijie Zhu, Taojiannan Yang, Chen Chen  
University of North Carolina at Charlotte  
{szhu3, tyang30, chen.chen}@uncc.edu

## Abstract

Street-to-aerial image geo-localization, which matches a query street-view image to the GPS-tagged aerial images in a reference set, has attracted increasing attention recently. In this paper, we revisit this problem and point out the ignored issue about image alignment information. We show that the performance of a simple Siamese network is highly dependent on the alignment setting and the comparison of previous works can be unfair if they have different assumptions. Instead of focusing on the feature extraction under the alignment assumption, we show that improvements in metric learning techniques significantly boost the performance regardless of the alignment. Without leveraging the alignment information, our pipeline outperforms previous works on both panorama and cropped datasets. Furthermore, we conduct visualization to help understand the learned model and the effect of alignment information using Grad-CAM. With our discovery on the approximate rotation-invariant activation maps, we propose a novel method to estimate the orientation/alignment between a pair of cross-view images with unknown alignment information. It achieves state-of-the-art results on the CVUSA dataset.

## 1. Introduction

Image based geo-localization aims at providing image-level GPS location by matching a query street/ground image with the GPS-tagged images in a reference dataset. Instead of relying on street-view images [25] as the reference dataset, street-to-aerial geo-localization [12] leverages GPS-tagged aerial-view images as the reference, given their more complete coverage of the Earth than street-view images. In the early work [12], street-to-aerial geo-localization is proposed to coarsely localize isolated images where no nearby ground-level image is available. With emerging deep learning techniques, recent works [8, 21] are able to achieve high geo-localization accuracy on city-scale datasets such as CVUSA [26] and Vo [21]. In scenarios

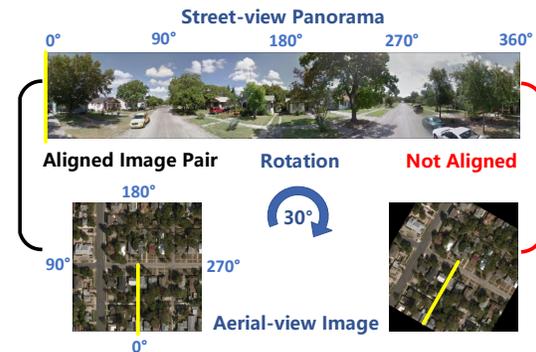


Figure 1: An example of alignment between street and aerial views. Yellow line denotes the South direction (0°).

where GPS signal is noisy [20], image geo-localization can provide additional information to achieve fine-grained localization. Street-to-aerial geo-localization is also proved effective on city-scale street navigation [10]. These practical applications make cross-view image geo-localization an important and attractive research problem in the computer vision community.

Recently, a number of works [3, 8, 13, 14, 21] are proposed to address the street-to-aerial geo-localization problem and the performance seems to be improved significantly. A key ingredient of the prior work is to learn a feature embedding, such that the distance of a matched pair of images is small whereas the distance of the unmatched pair is large in this feature space, which is also known as metric learning. However, existing works have different settings about the alignment between street and aerial views (Fig. 1), which can lead to unfair comparison. CVM-Net [8] is trained with randomly rotated aerial images without leveraging alignment information in the training set, thus is applicable for inference scenarios where alignment is not available, e.g. images on social media. [3, 21] take advantage of the alignment information as additional supervision, but they do not assume the inference image pair to be well aligned. [3, 8, 21] also apply their algorithms to cropped street view image dataset, i.e. Vo dataset [21]

(Fig. 2), where panorama street view is not available. On the contrary, [14, 17, 18] aim at image or feature transformation from one view to the other, which is *dependent on the geometric relationship between two views*. Although these techniques can boost the retrieval performance, but their presupposition does not hold when the alignment is not available for inference or only cropped street view images are provided. It can be unfair to claim state-of-the-art (SOTA) performance based on comparison with methods which do not have alignment assumption like [8]. These challenges motivate us to study three problems: 1) *how the alignment information would affect the retrieval model in terms of performance (Table 1) and beyond (Fig. 4)*; 2) *without assuming the inference image pairs are aligned, how to effectively improve the retrieval performance*; 3) *is it possible to estimate the alignment information when no explicit supervision is given?*

Validation	Training	
	Aligned	Rotate
Aligned	60.1%	43.7%
Rotate	13.5%	44.2%

Table 1: Top-1 recall accuracy of Siamese VGG with different alignment settings (Aligned or Randomly Rotated (**Rotate** for short)) for training and validation sets on CVUSA [26].

In this paper, we revisit the street-to-aerial view geo-localization problem and answer those three questions. We first shed light on the effect of alignment between two views, which is usually ignored in the discussion of previous works, by conducting ablation studies with a simple Siamese [4] network under different alignment settings. Without the alignment assumption in inference phase, improving metric learning is a promising direction, which is explored but not well-exploited in [3]. By identifying the unique challenges of cross-view image matching in the context of metric learning, we further show that two techniques (global mining strategy and a new loss function) specifically tailored to the challenges can significantly improve the retrieval performance regardless of the alignment setting. Moreover, we leverage visual explanation [16] to investigate how the image matching model works for cross-view. *Our observation reveals that the activation map of the matching model can provide geometric information which is independent of the alignment of the training data*. Inspired by this observation, we propose a novel orientation estimation method which significantly outperforms the existing approaches. The main contributions of this paper are three-fold:

- We provide an in-depth analysis on image alignment, which is ignored by previous works, for cross-view matching. Ablation study and visual explanation lead to

a key observation – the alignment has a great impact on the retrieval performance. It provides valuable information for designing robust and general frameworks, and manifesting fair comparisons with prior work.

- We show that improvements on metric learning techniques can boost the retrieval performance regardless of the alignment information. By identifying the unique challenges pertaining to the cross-view problem, our specifically designed pipeline achieves the state-of-the-art results on two benchmarks when *no assumption is made on the alignment of the inference set*.
- We discover that the orientation information between cross-view images can be estimated when the alignment is unknown. The proposed novel orientation estimation method outperforms previous methods without explicit supervision.

## 2. Related Work and Motivation

### 2.1. Geo-localization

Recent works for cross-view geo-localization [3,8,13,14, 21] are all based on Siamese networks [4], while they build their pipelines on different baselines with different settings. Vo *et al.* [21] first propose a Siamese network with exhausted triplet loss [15] for city-scale geo-localization with cropped street-view images and satellite images. Their results suggest that an auxiliary orientation regression task can further improve the performance. Hu *et al.* [8] propose CVMNet which combines Siamese VGG [19] and NetVLAD [2] along with a modified version of the triplet loss. Liu *et al.* [13] leverage the alignment information between street and aerial views and improve the performance by adding orientation information in the input. Krishna *et al.* [14] utilize GANs to generate image from one view to the other and adopt feature fusion to achieve higher accuracy. Shi *et al.* [18] aim to find the optimal feature transformation between two views based on the geometric prior knowledge. Cai *et al.* [3] put more weight on the hard samples in an online negative mining manner and use a stronger backbone (ResNet [6] with attention) to achieve better result. *Although progress has been made for cross-view geo-localization, existing works with different settings on the cross-view alignment (discussed in Section 2.2) can lead to unfair comparison.*

### 2.2. Alignment Setting

CVMNet [8] is trained on randomly rotated aerial images, in which case the alignment information is not available, thus resulting in a general framework for cross-view image matching. [3, 21] leverage the alignment information in training set by adding a regression task for orientation

prediction as additional supervision, while they are applicable for unaligned inference image pairs. However, [13] takes advantage of the alignment information by adding orientation as an auxiliary input, so the orientation is also required as input for inference images. The image/feature transformation in [14, 18] also relies on the accurate geometric relationship between two views. These methods may not work well when the alignment between street and aerial views is not available or only cropped street view images are provided, which is often the case in real-world applications. It is clear that more supervision can result in better performance, but existing works [13, 14] fail to make the effect of alignment information very clear in their comparisons.

Fig. 1 shows a graphical illustration of the alignment information of cross-view images. To investigate its impact on geo-localization performance, in Table 1 we report the top-1 recall accuracy of Siamese-VGG with different alignment settings for training and validation sets on the CVUSA dataset [26]. Training with randomly rotated aerial images (the alignment information is therefore not available) yields a performance drop on the aligned validation set (from 60.1% to 43.7%) compared with training with aligned images. But this trained model is able to perform well on the randomly rotated validation set (43.7% vs 44.2%). On the other hand, the model trained with aligned images has an extremely low top-1 accuracy (13.5%) on the randomly rotated validation set. These results indicate that the alignment information actually has a great impact on accuracy. And training without alignment information makes the model generalize better. More discussion on the effect beyond the retrieval accuracy is presented in Section 4.1.

### 2.3. Metric Learning

Unlike feature transformation, metric learning techniques are *independent of the alignment assumption*. General metric learning aims to learn an embedding space where positive samples (matched pairs) are close to each other, while negative samples (unmatched pairs) have a large distance between each other. Recent works [13, 14] usually adopt the loss from [8], i.e. a modified triplet loss [15], along with the within batch negative mining [15] or assigning more weights on hard negative samples in a mini-batch [3]. Although the common techniques for metric learning, e.g. triplet loss and hard negative mining, are employed in recent geo-localization methods, the unique challenges of cross-view geo-localization are not specifically addressed.

**Challenges for cross-view geo-localization.** For street-to-aerial view matching, most of the time, there is only one matched (**positive**) aerial-view image for the query street-view image from the same location. On the contrary, all the aerial images from other locations are considered as **negative** samples. As a result, there is a significant imbalance

between positive and negative pairs. Therefore, different from the Facenet [15] dataset which contains about 20 different images for one face ID, the number of positive samples for an anchor street-view image is very limited in geo-localization, i.e. only one. The boundary between positive pairs at different locations is difficult to estimate by only one sample in the embedding space. Since most existing methods follow the form of triplet loss which gives the same weight for positive and negative samples, *the large imbalance between positive and negative samples inspires us to design a better loss function in Section 3.2 for this task.*

As the training accuracy increases, most training samples are correctly handled and have little contribution to the overall loss [15], therefore hard negative mining [15] is necessary. Several geo-localization methods [3, 8, 13, 14] use a small batch size to fit the high resolution images in memory. Although online negative mining within mini-batch is employed, it does not work well when the training accuracy is high since almost no hard pairs can be found in a mini-batch. *We solve this problem in Section 3.3 by introducing a global mining strategy.*

### 2.4. Orientation Estimation

Several cross-view orientation estimation methods have been proposed along with geo-localization. Vo *et al.* [21] predict the rotated angle of aerial image by adding an auxiliary *supervised* regression task, with the goal of improving localization accuracy. The auxiliary regression sub-network is able to coarsely predict the orientation angle between street and aerial view images. Zhai *et al.* [26] first predict the semantic segmentation map of street-view panorama from aerial image by learning a transformation matrix between two views. Then the segmentation map is matched with the one generated directly from street-view crop in a sliding-window manner to find the best matching angle. To learn the cross-view transformation matrix, their model composing of three sub-networks has to be trained on *well-aligned* street and aerial image pairs. A pre-trained segmentation model is also required as additional supervision. In summary, these methods [21, 26] require *explicit supervision (i.e. image alignment)* on the training data to train their models for orientation estimation/prediction. *We propose an orientation estimation approach (Section 4.2) that does not rely on the alignment information for training, yet is able to achieve superior performance.*

## 3. Retrieval Framework

### 3.1. Baseline Architecture

We adopt a simple Siamese-VGG as our baseline architecture with the loss function of [8]. Given a set of training pairs including street-view images  $x_i$  and their corresponding aerial images  $y_i$ , our framework learns two mapping

functions,  $f(X, \Theta_x)$  and  $g(Y, \Theta_y)$ , which map each input pair  $(x_i, y_i)$  into a  $K$ -dimensional space. The goal is to find the best  $\Theta_x$  and  $\Theta_y$  so that  $g(y_i, \Theta_y)$  is the nearest neighbor of  $f(x_i, \Theta_x)$  in the embedding space. As shown in Fig. 2, we use the same backbone architecture (e.g. VGG) for both views, i.e.  $f = g$ . Due to the significant visual difference between two views, the convolutional layers denoted by two separate CNNs in Fig. 2 are trained without sharing weights in order to extract different low-level features from two views. On the other hand, fully connected layers of the two streams share weights, since the high-level feature (e.g. semantic information) is similar in both views.  $L_2$  normalization is used for the output feature vectors of both views.

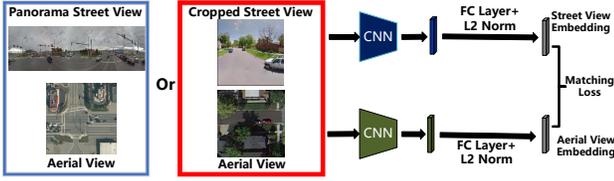


Figure 2: The overall framework of our method.

### 3.2. Binomial Loss

Triplet loss has been widely used in a variety of image matching tasks, including face clustering [15], person re-identification (Re-ID) [7, 11] and image retrieval [5, 24]. The idea is to teach the network to pull positive samples close to the anchor, while at the same time, push negative samples far away by a margin  $m$ . The triplet loss function is written as:

$$L = \frac{1}{N} \sum_i \max(0, d_i^p - d_i^n + m). \quad (1)$$

For  $N$  triplets in a mini-batch,  $d_i^p$  and  $d_i^n$  denote the distance between the  $i$ -th anchor and its corresponding positive and negative samples. Squared and non-squared Euclidean distances are commonly used in triplet loss. Different from the hard-margin function in Eq. (1), Vo *et al.* [21] propose to use a soft-margin function  $\sigma(d) = \log(1 + \exp(d))$  in the triplet network for geo-localization, where  $d = d_i^p - d_i^n$ . The soft-margin loss has a Sigmoid gradient function which is more smooth than hard margin. Moreover, Hu *et al.* [8] add a parameter  $\alpha$  to form a weighted soft-margin loss:

$$L = \frac{1}{N} \sum_i \sigma(\alpha(d_i^p - d_i^n)), \quad \alpha > 0. \quad (2)$$

The parameter  $\alpha$  should be tuned to find a suitable balance for the gradients of easy and hard samples. However,  $d_i^p$  and  $d_i^n$  always have the same weight  $\alpha$  in Eq. 2, resulting in the same magnitude of gradient. That means positive and negative samples will be pulled and pushed in the same man-

ner. For our positive and negative imbalanced case, adjusting different weights for positive and negative samples can lead to a loss function which better alleviates this problem. Yi *et al.* [23] have successfully utilized binomial deviance loss for person Re-ID. It is formulated as:

$$L = \frac{1}{N_p} \sum_i \sigma(-\alpha(s_i^p - m)) + \frac{1}{N_n} \sum_i \sigma(\alpha(s_i^n - m)). \quad (3)$$

Here  $s_i^p$  and  $s_i^n$  denote the **cosine similarity** between the  $i$ -th anchor and its positive and negative samples.  $N_p$  and  $N_n$  represent the number of positive and negative pairs, respectively. Although the original loss function uses the same parameters  $\alpha$ ,  $m$  for positive and negative samples, we make use of this formula while assigning different  $\alpha_p$ ,  $\alpha_n$  and  $m_p$ ,  $m_n$ . We use the normalized  $\sigma$  formulated as  $\hat{\sigma}(\alpha, d) = \frac{1}{\alpha} \sigma(\alpha d)$ . Then the gradient is given by  $\partial \hat{\sigma}(\alpha, d) / \partial d = 1 / (1 + \exp(-\alpha d))$ , and our loss function becomes:

$$L = \frac{\sum_i \sigma(-\alpha_p(s_i^p - m_p))}{\alpha_p N_p} + \frac{\sum_i \sigma(\alpha_n(s_i^n - m_n))}{\alpha_n N_n}. \quad (4)$$

By setting a large  $\alpha_n$ , the gradients of negative pairs drop fast as  $s^n$  decreases from  $m_n$ , which means only to push negative samples away by a small distance. However, with a small  $\alpha_p$ , the gradients of positive pairs drop slowly as  $s^p$  increases from  $m_p$ , resulting in pulling positive samples to the anchor until  $s_i^p$  is much greater than  $m_p$ . When positive samples are much fewer than negative samples, as in cross-view geo-localization with only one positive match, **it would be easier to pulling the only matched sample close to the anchor rather than pushing all negative samples away**. Therefore, we assign a much smaller value to  $\alpha_p$  than  $\alpha_n$  to validate this idea. To avoid too many hyper-parameters, we simply set  $m_p$  and  $m_n$  as the average values of  $s^n$  and  $s^p$ .

### 3.3. Global Mining Strategy

As the training accuracy increases, most of the negative samples contribute zero loss and the convergence becomes slow. In [8], mining the hardest triplet in a batch helps when there exists some hard negative samples, but it does not apply when a small batch size, e.g. 12 pairs (24 images), is used due to the high resolution images, e.g. panoramic street-view images. To find the global hard negative samples in a subset with  $N_m$  pairs, the embedding vectors of samples in the subset have to be updated every  $p$  steps ( $p$  can be adjusted to find the best trade-off between the computational cost and accuracy). To avoid large computational cost of this offline update strategy, our online mining computes the output vectors of each batch in the back-propagation step and saves them into a mining pool in an FIFO (first in first out) manner. Although the vectors in the mining

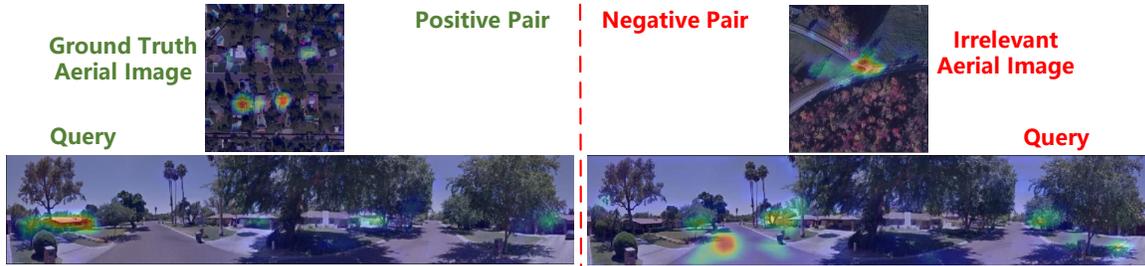


Figure 3: Grad-CAM activation map (overlaid on images) of our baseline on positive (left) and negative (right) image pairs.

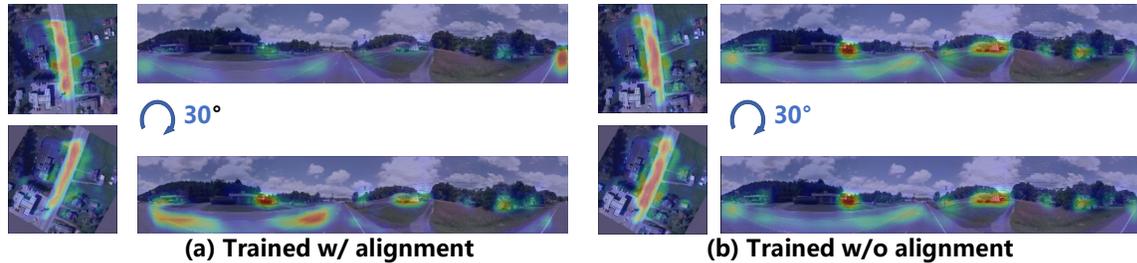


Figure 4: Grad-CAM activation maps (overlaid on images) of our Siamese VGG baseline trained on (a) aligned and (b) randomly rotated aerial images.

pool have the values of several steps before, the delay is always less than one epoch. This strategy generates approximate hard negative samples in each step for every sample in a batch with negligible additional computation. When the training set is small, we save the vectors of the whole training set such that the global hard negative samples can be found. For each positive pair, we randomly select one of the  $r$  hardest negative samples from the mining pool.

## 4. Alignment Analysis

### 4.1. Visual Explanation

It is evident that using the cross-view alignment information can improve the performance, but how the alignment information may affect the geo-localization model? In addition to the performance (Table 1), we take a closer look at the trained model with or without alignment, through the lens of Grad-CAM [16] – a widely used visualization technique, in order to understand how the model actually works. The class activation map generated by Grad-CAM highlights the important regions which contribute the most to the final similarity between two images. Specifically, the class activation map is obtained by computing the gradient till the last convolution layer from the inner product of two views’ embedding features without  $L_2$  normalization [27].

As shown in Fig. 3, the activation map of the same query image can be dramatically different when the retrieved image is different. For the positive image pair in this example, the model mainly focuses on the discriminative objects in both views, i.e. *houses*. However, for the negative pair, the

model highlights the *trees and roads* areas in the query image, as there is no house in the retrieved image. The visual explanation further demonstrates that the similarity score predicted by the metric learning model is mainly based on similar patterns in different views, and the activated regions of two views are highly relevant to each other. A reasonable hypothesis is that the activated regions of two views are likely to be the same objects. *However, what if the image pair is not aligned?* In Fig. 4, we show the changes of activation map corresponding to a rotation of  $30^\circ$  in the aerial image, using two different models: (a) baseline trained with aligned images, (b) baseline trained with unaligned images. For model (a), the activation map changes dramatically on both views when the aerial image is rotated (Fig. 4 (a)), because the model trained with aligned images relies on specific geometric relationship between two views. However, model (b) which is trained with randomly rotated aerial images only leverages similar patterns for matching (Fig. 4 (b)), thus leading to approximate *rotation-invariant* activation map, which may provide geometric information for orientation (or camera pose) estimation (Section 4.2).

### 4.2. Orientation Estimation Approach

Since the most activated regions are likely to be the same objects, the angle distributions (Fig. 5) of activated pixels from two views would be similar if the image pair is well aligned. A rotation of aerial image will cause the angle distribution to shift by a certain degree. With the observation of the *approximate rotation-invariant activation*

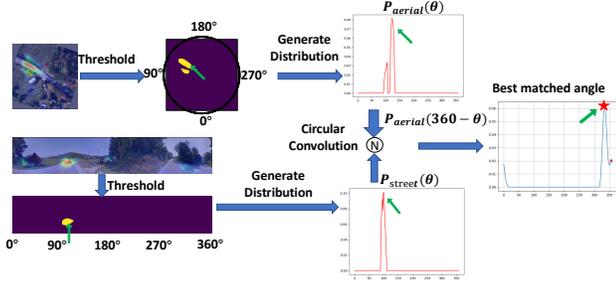


Figure 5: Framework of our orientation estimation approach. (Zoom in to view)

maps (Section 4.1), which are obtained from our matching model trained with rotated aerial images, we propose to estimate the orientation of an unaligned image pair by matching the angle distributions of activated pixels in two views. As shown in Fig. 5, we first compute the activation maps using Grad-CAM and select pixels with values higher than a threshold. Then we calculate the angle distribution of the selected pixels in each view as  $p_{street}(\theta)$  and  $p_{aerial}(\theta)$ . To find the angle  $\phi$  so that  $p_{aerial}(\theta + \phi)$  best matches  $p_{street}(\theta)$ , circular convolution of  $p_{street}(\theta)$  and  $p_{aerial}(360 - \theta)$  is computed efficiently using the Fast Fourier Transform (FFT) algorithm as follows ( $\textcircled{N}$  denotes circular convolution):

$$\begin{aligned}
 p(\theta) &= p_{street}(\theta) \textcircled{N} p_{aerial}(360 - \theta) \\
 &= \text{ifft}(\text{fft}(p_{street}(\theta))\text{fft}(p_{aerial}(360 - \theta))), \quad (5)
 \end{aligned}$$

where  $\text{fft}$  and  $\text{ifft}$  are the Discrete Fourier Transform (DFT) operation and inverse DFT operation. Finally, the highest peak is selected as the predicted angle. *Since the training image pairs are not required to be aligned, the proposed method does not leverage explicit supervision of orientation, which is a clear advantage as compared with previous supervised counterparts [21, 26].*

## 5. Experiment

**Dataset.** We conduct experiments on two popular benchmark datasets, i.e. CVUSA [26] (panorama street view) and Vo [21] (cropped street view), see Fig. 2. These two datasets are city-scale benchmarks for cross-view geo-localization. The original CVUSA (Cross-View USA) [22] dataset contains more than 1 million of ground-level and aerial images from across the US. Furthermore, Zhai *et al.* [26] make use of the camera’s extrinsic parameters to generate aligned pairs by warping the panoramas, resulting in 35,532 image pairs for training and 8,884 image pairs for testing. We use the same selected CVUSA [26] in our experiment.

Vo *et al.* [21] introduce a cross-view dataset which consists of about one million image pairs from 11 cities in the

US. They randomly query street-view panorama images and generate several crops from each panorama with overhead images from Google Map. Small crops and overlapping images make this dataset more challenging. For a fair comparison, we follow the splitting in [21] using 8 cities for training and Denver city as the test set.

**Implementation Details.** For the proposed baseline, we adopt VGG-16 [19] as the backbone architecture for both views with the loss function in Eq. 2 [8] ( $\alpha = 20$ ). Adam [9] optimizer with learning rate decay is used for training and the dimension of the embedding space is the same as CVMNet [8]. Random aerial view rotation is adopted for training without alignment. The proposed mining strategy and simple backbone facilitate the training on a single GPU (Nvidia 1080Ti) with a small batch size, i.e. 12 pairs for CVUSA and 32 pairs for Vo. All experiments are implemented based on Tensorflow [1]. The mining pool is the whole training set for CVUSA and we use a subset of 10,000 pairs for Vo with  $r = 100$ . The mining pool is updated every epoch and the binomial loss is adopted after the distribution is stable (30 epochs). We set  $\alpha_p = 5$ ,  $\alpha_n = 20$  for both CVUSA and Vo. We simply set  $m_p, m_n$  as 0, 0.7 based on the average values of the distributions of  $s^n$  and  $s^p$  as shown in Fig. 8.

## 5.1. Retrieval Performance

**Evaluation Metrics.** For geo-localization, we use the top- $n$  recall accuracy as the evaluation metric on both datasets. For each query image, the retrieval is considered successful if the ground-truth reference is ranked within the top  $n$  retrieved images. Instead of only providing the top-1% accuracy, we also report the top-1 accuracy to better demonstrate the performance of our approach. The top-1% accuracy is useful when the performance is poor, but it is not discriminative anymore if the top-1% is already higher than 95% as in this paper. Since a high top-1 accuracy is the ultimate goal of geo-localization for practical applications, we highly recommend including the top-1 accuracy in future research.

We compare our method with existing methods on two datasets in Table 2. We use the published results of [3, 8, 13, 14, 21, 22] for top-1% accuracy comparison. Not all the methods have the top-1 accuracy. For CVM-Net [8], we reproduce their test result on CVUSA with their pre-trained model for top-1 and top-1% accuracy<sup>1</sup>, while the result of CVM-Net on Vo is from their paper since the pre-trained model for this dataset is not provided.

**CVUSA.** As shown in Table 2, the overall proposed pipeline significantly surpasses existing methods without the alignment assumption on inference set. Note that we do not adopt any recent strong network backbones as in [3]

<sup>1</sup>The top-1% accuracy we generated using their model and code is 2% higher than their published result in [8].

Method	CVUSA		Vo	
	Top-1%	Top-1	Top-1%	Top-1
Scott [22] <sub>(ICCV'15)</sub>	34.3%	-	15.4%	-
Zhai [26] <sub>(CVPR'17)</sub>	43.2%	-	-	-
Vo [21] <sub>(ECCV'16)</sub>	63.7%	-	59.9%	-
CVMNet [8] <sub>(CVPR'18)</sub>	93.6%	22.5%	67.9%	-
Lending [13] <sub>(CVPR'19)</sub>	93.19%	31.71%	-	-
Reweight [3] <sub>(ICCV'19)</sub>	98.3%	46.0%	78.3%	-
GAN [14] <sub>(ICCV'19)</sub>	95.98%	48.75%	-	-
<b>Ours</b>	<b>97.7%</b>	<b>54.5%</b>	<b>88.3%</b>	<b>11.8%</b>

Table 2: Top-1 and top-1% recall accuracy comparison on CVUSA and Vo.

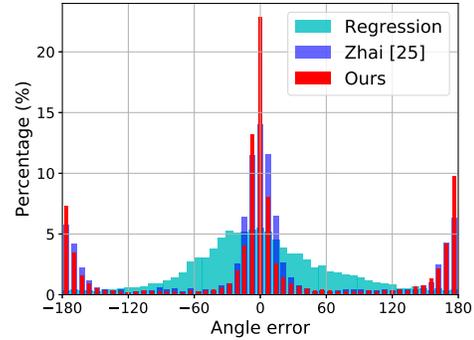
(ResNet+attention), we believe our approach would gain extra improvement from using stronger network backbones. Our method also enjoys simple architecture and therefore is easy to reproduce and requires less computation compared with CVMNet. Moreover, our global mining strategy brings superiority on convergence speed for the training process (details in Section 5.3 and Fig. 7).

**Vo.** Vo [21] is a much more challenging dataset compared with CVUSA, as it contains around 10 times of images for validation and the cropped street-view images provide less information than panorama. Furthermore, the image/feature transformation based on geometric prior knowledge is infeasible for this case. Therefore, as shown in Table 2, every method yields a much lower accuracy on Vo than that on CVUSA. Also, none of the existing methods reports top-1 accuracy. Again, the proposed framework outperforms all existing methods by a large margin (**10% on top-1% accuracy**). Note that [3, 21] utilize the alignment information of the training data by adding a rotation angle classification task (an auxiliary task) to their matching frameworks to boost the geo-localization performance, but we do not use this technique in our method.

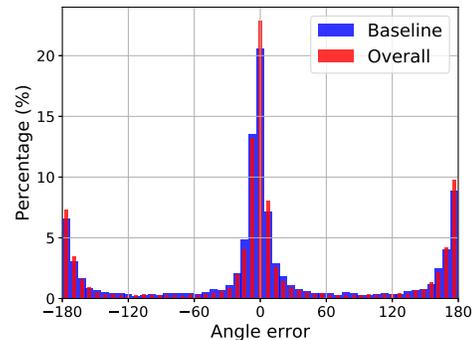
## 5.2. Orientation Estimation Performance

As described in Section 4.2, our trained matching model is able to roughly predict the angle between paired aerial and street view images without any alignment information in the training set. We conduct experiments on CVUSA with randomly rotated aerial images and evaluate the distribution of angle prediction errors. For comparison, we train our baseline with an auxiliary orientation regression task as in [21]. We also compare our result with another supervised method [26] which learns a cross-view transformation matrix and predict the orientation in a sliding window manner.

Without explicit alignment information from the training set, our method achieves better result than regression and [26] as shown in Fig. 6 (a). The percentage of samples with error in  $[-3.5^\circ, 3.5^\circ]$  is around 24% of the whole test set. It is worth noticing that most failure predictions of our method have an error around  $180^\circ$ , because the ac-



(a) Our method vs. previous methods.



(b) Overall model vs. baseline model.

Figure 6: Error distribution of orientation estimation on CVUSA in percentage (best viewed in color with zoom in).

tivation map usually focuses on the road which is symmetrical in aerial view. The same pattern is found in the result of [26], while the result of regression has a very different pattern. Fewer samples of regression have an error around  $180^\circ$ , while accurate predictions (the central bar on  $0^\circ$ ) are also much fewer than the other two methods.

We further provide the comparison between our baseline and overall framework for orientation estimation in Fig. 6 (b). Our overall framework leverages superior metric learning techniques, thus achieving better performance than our simple baseline. For example, the percentages of samples with error in  $[-3.5^\circ, 3.5^\circ]$  are 24% and 21% for our overall and baseline models, respectively. The result shows that better retrieval model does improve the orientation estimation performance.

## 5.3. Ablation Study

**Effect of global mining.** To demonstrate the effectiveness of the proposed global mining strategy, we conduct experiments with three settings on our baseline, i.e. no mining, within-batch mining, and our global mining scheme. As discussed in Section 3.3, the within-batch mining [8] or the loss in [3] is not able to select hard negative samples, because almost no hard negative samples exists within a mini-batch when the top-1% accuracy is higher than 80%. We report the result of our baseline with within-batch min-

ing (“Batch mining” in Table 3) to show the superiority of global mining. As shown in Table 3, the global mining significantly improves the performance on both datasets, while the within-batch mining makes little difference on performance. In Fig. 7, we also present the convergence speeds of our baseline and baseline with global mining on CVUSA. The “Baseline + global mining” converges much faster than the “Baseline”, especially in terms of the top-1 accuracy. Since our mining strategy aims to find the hardest  $r$  samples, it brings relative larger gain on top-1 accuracy than top-1% accuracy.

Mining Strategy	CVUSA		Vo	
	Top-1%	Top-1	Top-1%	Top-1
No mining	96.9%	43.7%	84.3%	7.9%
Batch mining	96.7%	43.0%	84.6%	8.0%
Global mining	<b>97.0%</b>	<b>52.1%</b>	<b>85.8%</b>	<b>11.1%</b>

Table 3: Comparison between different mining strategies with our baseline on CVUSA and Vo.

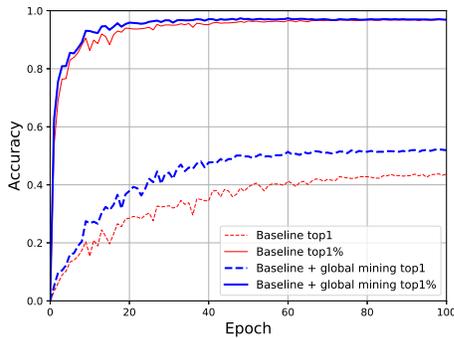


Figure 7: Top-1 and top-1% recall accuracy vs. epochs of our baseline w/o and w/ global mining on CVUSA.

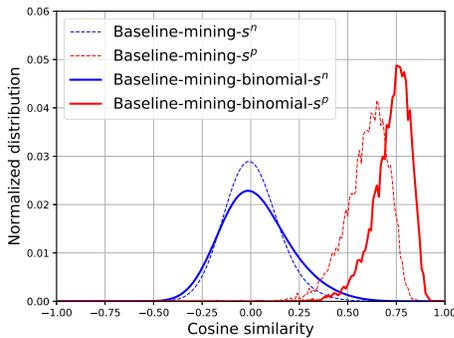


Figure 8: Cosine similarity distributions with loss in CVM-Net [8] (Eq. 2) and binomial loss (Eq. 4) on CVUSA.

**Effect of binomial loss.** As shown in Fig. 8, the average values of  $s^n$  and  $s^p$  (similarity between negative pairs and positive pairs) lie at around 0 and 0.7 for “Baseline + global mining” with the loss in Eq. 2 [8]. An intuitive idea to further improve the performance is breaking the restraint

of triplet-like loss which assigns the same weight for positive and negative pairs. As explained in Section 3.2, the binomial loss (Eq. 4) puts a stronger constraint on positive samples (larger mean value and smaller variance in Fig. 8), while the distribution of negative pairs is more scattered (larger variance). This characteristic is beneficial to the performance on both datasets as shown in Table 4.

Loss Function	CVUSA		Vo	
	Top-1%	Top-1	Top-1%	Top-1
CVMNet [8]	97.0%	52.1%	85.8%	11.1%
Binomial (Eq. 4)	<b>97.7%</b>	<b>54.5%</b>	<b>88.3%</b>	<b>11.8%</b>

Table 4: Comparison between different loss functions with our baseline and global mining on CVUSA and Vo.

**Effect of alignment.** As discussed in Section 4.1, the alignment setting has a large impact on the performance, we report the ablation study results of CVUSA on both settings (w/ or w/o alignment) in Table 5. The “Overall” denotes our baseline with global mining and binomial loss in Eq. 4. As expected, the improvements of the proposed techniques are consistent under both settings. Apparently, training with alignment can improve the retrieval performance, because it assumes the inference images to be aligned, but this may be infeasible for challenging real-world applications.

Method	w/ alignment		w/o alignment	
	Top-1%	Top-1	Top-1%	Top-1
Baseline	98.8%	60.1%	96.9%	43.7%
Baseline+global mining	98.8%	67.0%	97.0%	52.1%
Overall	<b>99.1%</b>	<b>70.4%</b>	<b>97.7%</b>	<b>54.5%</b>

Table 5: Comparison between our methods w/ and w/o alignment (random rotated aerial images) on CVUSA. “Overall” = Baseline+global mining+binomial loss (Eq. 4).

## 6. Conclusion

In this paper, we revisit cross-view image geo-localization and orientation estimation and highlight the effect of image alignment information which is usually ignored by previous works. Our analysis indicates the alignment has a great impact on the retrieval performance. Furthermore, we identify the unique challenges of geo-localization and propose a global mining strategy along with the binomial loss to tackle them. Extensive experiments on two widely used benchmark datasets show the superiority of the proposed method on both alignment settings (w or w/o). Moreover, our model trained without alignment is able to predict the orientation (angle) of a cross-view image pair without any alignment information supervision during the training. The proposed orientation estimation method achieves state-of-the-art result on CVUSA.

**Acknowledgement.** This work is partially supported by the start-up grants of UNC-Charlotte.

## References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)*, pages 265–283, 2016.
- [2] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Paszke, and Josef Sivic. Netvlad: Cnn architecture for weakly supervised place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5297–5307, 2016.
- [3] Sudong Cai, Yulan Guo, Salman Khan, Jiwei Hu, and Gongjian Wen. Ground-to-aerial image geo-localization with a hard exemplar reweighting triplet loss. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8391–8400, 2019.
- [4] Sumit Chopra, Raia Hadsell, and Yann LeCun. Learning a similarity metric discriminatively, with application to face verification. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 1, pages 539–546. IEEE, 2005.
- [5] Albert Gordo, Jon Almazan, Jerome Revaud, and Diane Larlus. End-to-end learning of deep visual representations for image retrieval. *International Journal of Computer Vision*, 124(2):237–254, 2017.
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [7] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017.
- [8] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7258–7267, 2018.
- [9] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [10] Ang Li, Huiyi Hu, Piotr Mirowski, and Mehrdad Farajtabar. Cross-view policy learning for street navigation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8100–8109, 2019.
- [11] Wentong Liao, Michael Ying Yang, Ni Zhan, and Bodo Rosenhahn. Triplet-based deep similarity learning for person re-identification. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 385–393, 2017.
- [12] Tsung-Yi Lin, Serge Belongie, and James Hays. Cross-view image geolocation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 891–898, 2013.
- [13] Liu Liu and Hongdong Li. Lending orientation to neural networks for cross-view geo-localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5624–5633, 2019.
- [14] Krishna Regmi and Mubarak Shah. Bridging the domain gap for ground-to-aerial image matching. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 470–479, 2019.
- [15] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015.
- [16] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.
- [17] Yujiao Shi, Liu Liu, Xin Yu, and Hongdong Li. Spatial-aware feature aggregation for image based cross-view geo-localization. In *Advances in Neural Information Processing Systems*, pages 10090–10100, 2019.
- [18] Yujiao Shi, Xin Yu, Liu Liu, Tong Zhang, and Hongdong Li. Optimal feature transport for cross-view image geo-localization. *arXiv preprint arXiv:1907.05021*, 2019.
- [19] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [20] Shaohui Sun, Ramesh Sarukkai, Jack Kwok, and Vinay Shet. Accurate deep direct geo-localization from ground imagery and phone-grade gps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1016–1023, 2018.
- [21] Nam N Vo and James Hays. Localizing and orienting street views using overhead imagery. In *European conference on computer vision*, pages 494–509. Springer, 2016.
- [22] Scott Workman, Richard Souvenir, and Nathan Jacobs. Wide-area image geolocation with aerial reference imagery. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3961–3969, 2015.
- [23] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Deep metric learning for person re-identification. In *2014 22nd International Conference on Pattern Recognition*, pages 34–39. IEEE, 2014.
- [24] Baosheng Yu, Tongliang Liu, Mingming Gong, Changxing Ding, and Dacheng Tao. Correcting the triplet selection bias for triplet loss. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 71–87, 2018.
- [25] Amir Roshan Zamir and Mubarak Shah. Image geolocation based on multiplenearest neighbor feature matching using generalized graphs. *IEEE transactions on pattern analysis and machine intelligence*, 36(8):1546–1558, 2014.
- [26] Menghua Zhai, Zachary Bessinger, Scott Workman, and Nathan Jacobs. Predicting ground-level scene layout from aerial imagery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 867–875, 2017.

- [27] Sijie Zhu, Taojiannan Yang, and Chen Chen. Visual explanation for deep metric learning. *arXiv preprint arXiv:1909.12977*, 2019.