GAReT: Cross-view Video <u>G</u>eolocalization with Adapters and Auto-Regressive Transformers

Manu S Pillai¹, Mamshad Nayeem Rizve², and Mubarak Shah¹

¹ Center for Research in Computer Vision, University of Central Florida, USA ² Adobe manu.pillai@ucf.edu, mrizve@adobe.com shah@crcv.ucf.edu

Abstract. Cross-view video geo-localization (CVGL) aims to derive GPS trajectories from street-view videos by aligning them with aerialview images. Despite their promising performance, current CVGL methods face significant challenges. These methods use camera and odometry data, typically absent in real-world scenarios. They utilize multiple adjacent frames and various encoders for feature extraction, resulting in high computational costs. Moreover, these approaches independently predict each street-view frame's location, resulting in temporally inconsistent GPS trajectories. To address these challenges, in this work, we propose **GAReT**, a fully transformer-based method for CVGL that does not require camera and odometry data. We introduce GeoAdapter, a transformer-adapter module designed to efficiently aggregate imagelevel representations and adapt them for video inputs. Specifically, we train a transformer encoder on video frames and aerial images, then freeze the encoder to optimize the GeoAdapter module to obtain videolevel representation. To address temporally inconsistent trajectories, we introduce TransRetriever, an encoder-decoder transformer model that predicts GPS locations of street-view frames by encoding top-k nearest neighbor predictions per frame and auto-regressively decoding the best neighbor based on the previous frame's predictions. Our method's effectiveness is validated through extensive experiments, demonstrating state-of-the-art performance on benchmark datasets. Our code is available at https://github.com/manupillai308/GAReT.

Keywords: Cross-View Video Geo-localization · Transformer Adapters · Transformers for Retrieval

1 Introduction

Recently, there has been significant interest in cross-view video geo-localization (CVGL), driven by its applications in vision-guided navigation, autonomous driving, and robot manipulation [2,4,7]. Like its image counterpart [25,40,41,44, 46], the objective of CVGL is to obtain the GPS trajectory of a video by matching every frame of a query street-view video with a reference aerial-view gallery. Prior art primarily tackled CVGL from a coarse Sequence-to-Image [28,43] matching perspective. In this setting, the objective is to obtain a *single* GPS coordinate

by matching a sequence of street-view images (typically ~ 7) [43], sampled from a street-view video with geo-tagged aerial images in the reference gallery. However, this coarse localization may not be optimal for many real-world applications [5]. Hence, [32] extends CVGL to a more fine-grained Frame-to-Frame setting, where each frame of a street-view video is matched with a corresponding reference aerial image (frame-level association). Particularly, [32] proposes a hierarchical approach and a hierarchical dataset, where they first localize a street-view video to a larger geographical area before obtaining the final frameby-frame predictions. Thus, following [23, 28, 32, 43], tackling CVGL at both coarse (sequence-to-image) and fine (frame-to-frame) levels is the primary focus of this work.

While existing works on CVGL have demonstrated promising performance, they are not without notable drawbacks. Firstly, these approaches often rely on camera intrinsic parameters and odometry data to model the temporal relation in the street-view video and to effectively match it with a reference aerial image [28]. which are not always available, especially in videos captured in uncontrolled environments. Secondly, to improve performance, prior work [32] utilizes contextual information from multiple adjacent street-view frames (0.5-second short 'clip') to obtain a feature representation for the center street-view frame. This design leads to additional computational overhead for obtaining per-frame features and, hence, is not optimal for real-time applications. Lastly, and most importantly, existing works on CVGL suffer from temporally inconsistent GPS predictions due to modeling per-frame localization as independent nearest-neighbor retrieval operations [28,32]. However, while obtaining the GPS predictions for each frame in a street-view video, the predictions should follow temporal closeness consistent with the input video. In this work, we aim to address these limitations of existing CVGL approaches.

To this end, we extend TransGeo [44], a versatile transformer-based method for cross-view image geo-localization, to the video setting and propose **GAReT**, a fully transformer-based method for cross-view video geo-localization. We hypothesize that in the context of video geo-localization, the temporal reasoning requirement, though important, is not as critical as in other video tasks (like action recognition). Following this assumption, we hypothesize that in order to geo-localize a street-view video, the representations from a model capable of geo-localizing street-view images can be efficiently aggregated to obtain videolevel representation. Towards this, we propose a transformer-adapter-based strategy (*GeoAdapter*) to aggregate image-level representations of an image geolocalization method and adapt it to video inputs. The advantage of this strategy is two-fold: (i) As the image geo-localization method is grounded on matching frame-level representations, a learnable aggregation module can effectively fuse the representations to obtain representation for video geo-localization. (ii) The efficiency of adapters allows our model to be more computationally lightweight than traditional video models [3, 18] as well as previous works [28, 32]. Additionally, our method tackles the problem of CVGL in two hierarchical stages (as opposed to four in [32]), where we first localize a street-view video to a large aerial region (using the video representation) and then individually localize each frame of the video within the larger region. As transformers can model view-specific feature attributes without explicit view transformations (e.g., polar transform or projective transform) [44], our method does not require any knowledge of camera intrinsic or odometry. Furthermore, to ensure temporal consistency in frame-to-frame predictions, we propose an encoder-decoder transformer model (*TransRetriever*) that encodes the top-k nearest neighbor predictions for each frame and auto-regressively decodes the best result based on previous frames.

We summarize our contributions as follows:

- We propose a fully transformer-based method (**GAReT**) for cross-view video geo-localization that doesn't rely on explicit viewpoint information like odometry and camera intrinsic.
- We employ a novel transformer adapter (*GeoAdapter*) for CVGL that adapts an image geo-localization encoder for video inputs, outperforming dedicated video models on benchmark datasets.
- To ensure temporally consistent frame-to-frame GPS predictions, we present the first learnable transformer-based model (*TransRetriever*) that models the independent frame retrievals through an auto-regressive decoder.

2 Related Works

Cross-view geo-localization represents an active field of research in computer vision with various applications in real-world systems. Diverse endeavors have been undertaken in the realm of cross-view *image* geo-localization with limited efforts in its video counterpart. Cross-view Image geolocalization: In cross-view image geo-localization, the aim is to obtain a GPS coordinate for a street-view image by matching it with aerial reference images. In [25], a polar transform of aerial images was used to align views, alongside the introduction of the Spatial-Aware Feature Aggregation (SAFA) module for learning embeddings with spatial correspondence between views. To align views, [17] proposed using an orientation mask that provides geometric orientation correspondence to the feature extractor. In [26, 27], the authors introduce Cross-DSM, a dynamic similarity matching (DSM) module to compute a rough estimate of the azimuth angle before performing matching. [35] showed polar transform of aerial images is not necessarily required for better performance and introduced a transformer encoder with self-cross attention layers. [45] identified key factors in cross-view image geo-localization, especially the need for orientation alignment of streetview and aerial-view image pairs. VIGOR [46] dataset was introduced to address the lack of challenging benchmark datasets. Following this, [44] proposed the first pure transformer-based architecture that performed significantly better than CNN counterparts without explicit view-specific transformation. [42] propose a novel disentangled feature learning framework that outperforms other methods on cross-area benchmarks. Recently, a new direction in cross-view image geo-localization has emerged. The goal is to accurately determine the location of the street-view camera and estimate its pose [11, 16, 24, 33, 34].

As previously mentioned, the research landscape for video geo-localization is considerably limited. Cross-view Video geolocalization: [28] developed a cross-view video-based localization approach using street-and-aerial geometric correspondences. Another work [43] proposes a novel temporal feature aggregation technique to learn better global representations of street-view video. Though promising results were obtained for both these works, they tackled the problem of CVGL from a Sequence-to-Image standpoint. [32] introduced the most recent work in Frame-to-Frame CVGL, inspired by [23]. Their contribution entailed a pioneering hierarchical approach for cross-view video geo-localization where the query street-view videos are first localized to a larger aerial region, then localizing each video frame within the larger region. They employ different models for each localization stage, resulting in increased computational complexity. Additionally, existing approaches in CVGL adhere to the nearest-neighbor paradigm for inference, where the best match from the reference set (aerial images) is taken as the nearest neighbor of the query (street-view video frame). This greedy nearestneighbor strategy yields a suboptimal local solution by independently retrieving matches for rather temporally related frames. Several works have shown improvement over this standard local optimum by posing a global constraint or gain on the retrieval task [29, 38, 39], but all of them are heuristic-based approaches. In this work, our primary focus centers on mitigating these issues. We demonstrate how a purely transformer-based approach, equipped with our proposed *GeoAdapter* module, can surpass current state-of-the-art methods by a significant margin. Furthermore, to ensure a temporally consistent solution, we propose the first transformer-based approach, TransRetriver, that models independent frame retrievals through an auto-regressive transformer decoder.

3 Methodology

We first present the problem formulation and an overview of our approach in Sec. 3.1. Then, we introduce the different components of our method in detail in Sec. 3.2 and Sec. 3.3. Finally, we present *TransRetriever*, our novel transformer-based learnable architecture for temporally consistent frame-to-frame predictions in Sec. 3.4.

3.1 Problem Formulation & Approach

We focus on CVGL in a frame-to-frame setting. Given a query street-view video, we need to geo-localize each frame individually by visually matching the frames to geotagged aerial *reference* images. Formally, given a set of query street-view videos $\{V^s\}$ and aerial-view reference images $\{I^a\}$, our objective is to learn an embedding space where every frame of query street-view videos is close to its corresponding ground-truth aerial image. Following the frame-to-frame setting proposed in [32], we assume that each street-view video V^{s_i} is associated with a *large* aerial image in addition to aerial images (*small*) corresponding to each street-view video frame. The large aerial image facilitates the hierarchical approach of localizing an entire street-view video to a larger geographical region (sequence-to-image) before individually obtaining the predictions for each frame (frame-to-frame). Precisely, each video V^{s_i} of n frames will have n ground-truth small aerial images (denoted hereon by $\{I_{s_j}^{a_i}\}_{j=1}^n$) and one ground-truth large aerial image (denoted hereon by $I_L^{a_i}$).

As shown in Fig. 1, we propose a hierarchical solution towards CVGL. During training, we optimize our feature encoder with image and video inputs to learn a unified model (\mathbf{U}) for feature extraction. Our unified model is a combination of a vision transformer (ViT) (T) as in [44] equipped with our proposed GeoAdapter $(\mathbf{G}_{\mathbf{A}})$ module, i.e., $\mathbf{U} = \{\mathbf{T}, \mathbf{G}_{\mathbf{A}}\}$. We begin by optimizing the ViT parameters with image pairs (street-view frame and small aerial image) as inputs while keeping the adapter module frozen. Then, we take video pairs (street-view video and large aerial image) as inputs and freeze the transformer encoder parameters, and optimize the adapter module. During inference, we localize a given streetview video to a larger region using our unified model and then independently obtain frame-to-frame predictions using the vision transformer (without the proposed GeoAdapter) encoder within the localized larger region. Finally, instead of independently retrieving frame-to-frame predictions through nearest-neighbor retrieval, we use our proposed TransRetriever (\mathbf{T}_{AR}) to autoregressively predict GPS coordinates for each frame by selecting matching small aerial images from a set of possible candidates (see Sec. 3.4).

3.2 Image Representations for CVGL

Following [44], we use DeiT [30] as image encoder for our unified model **U**. Note that, for simplicity, we consider one pair of street-view and aerial images (V^{s_i}, I^{a_i}) and ignore subscripts in notations.

Given a pair (V^s, I^a) , we sample every k^{th} second frame from V^s and obtain corresponding matching $I_{s_k}^a \in \mathbb{R}^{256 \times 256}$. For each view-specific input, similar to [25,40,41,44,46], we train two separate image encoders, i.e., $\mathbf{T_a}, \mathbf{T_s}$ to generate embedding features. Specifically, we divide an input image $(V_k^s \text{ or } I_{s_k}^a)$ of size H, W into P non-overlapping patches of size $p \times p$. Each patch is flattened and fed into a trainable linear projection layer to generate token representations. The tokens are prepended with a learnable token representation called [CLS] token to learn global features from the image. Each token is then added with positional embeddings [31] and goes through multiple transformer blocks. The encoded [CLS] token representation after the last transformer block is passed through a linear projection layer to obtain the final embedding of the input image (additional details are provided in the supplementary material).

The encoders are optimized through contrastive learning [22] to facilitate the model in learning embeddings where matching pairs are encoded closer together while non-matching pairs are encoded farther apart. Formally, we optimize the encoders on soft-margin triplet loss [12] as given below:

$$\mathcal{L}_{1} = \sum_{V^{s}} \sum_{k} \log \left(1 + e^{\alpha(||\mathbf{T}_{s}(V_{k}^{s}) - \mathbf{T}_{a}(I_{s_{k}}^{a})||_{2} - ||\mathbf{T}_{s}(V_{k}^{s}) - \mathbf{T}_{a}(I_{s_{i\neq k}}^{a})||_{2})} \right)$$
(1)



Fig. 1: Overview of our proposed approach GAReT. (A) We begin by optimizing our image transformer encoders $\mathbf{T}_{\mathbf{a},\mathbf{s}}$ (B) with street-view frame V_k^s and matching small aerial image $I_{s_k}^a$ pair. (C) Then, for adapting our image encoder to video inputs, we add our GeoAdapter G_A module and only optimize the adapter parameters with video pairs as inputs, i.e., a street-view video V^s and corresponding large aerial image I_L^a . For training, we sample every k^{th} frame from the street-view video and partition the large aerial image into non-overlapping patches. (D) In G_A , we apply temporal selfattention (TSA) computation only on the CLS tokens. For TSA computation, we reuse the spatial self-attention weights. (E) During inference, we first perform a Sequence-to-Image inference procedure, where given a query street-view video, our unified module $\mathbf{U} = \{\mathbf{T}, \mathbf{G}_{\mathbf{A}}\}$ produces feature embeddings for both the V^s and I_L^a . Then, using the embeddings, we retrieve the t nearest neighbor large aerial images (here we show t =1) and construct a small aerial image gallery \mathcal{G} . (F) Finally, $\mathbf{G}_{\mathbf{A}}$ is removed, and feature embeddings for $I_{s_k}^a$ and V_k^s are obtained. These features are then passed to our TransRetriever T_{AR} model to obtain final frame-by-frame GPS predictions to construct a GPS trajectory.

where, $|| \cdot ||_2$ denotes l^2 norm and $\mathbf{T}_{\mathbf{a},\mathbf{s}} : \mathbb{R}^{H \times W \times 3} \to \mathbb{R}^d$.

3.3 GeoAdapter: Adapting to Video for CVGL

In order to localize a street-view video to a larger geographical region, we extend our image geolocalization model $\mathbf{T}_{\mathbf{a},\mathbf{s}}$ to video inputs. We propose GeoAdapter ($\mathbf{G}_{\mathbf{A}}$), the second and final component of our unified model \mathbf{U} to efficiently aggregate the representations of $\mathbf{T}_{\mathbf{a},\mathbf{s}}$ to obtain video embeddings. $\mathbf{G}_{\mathbf{A}}$ is a transformer adapter module inspired by [20, 36] which consists of two 2-layer MLPs with GeLU activation function, separated by a temporal self-attention (TSA) layer. The TSA layer models the temporal relation between frames of the input video by leveraging self-attention across the temporal dimension of the input. Furthermore, to induce temporal ordering in tokens, we add temporal positional embeddings to each of the tokens before computing TSA. To effectively learn rich contextual information from the video in different levels, we add our proposed module to every block of the transformer encoder $\mathbf{T}_{a,s}$. Formally, in block l of our unified model $\mathbf{U} = \{\mathbf{T}, \mathbf{G}_A\}$, we have,

$$\hat{h}^{l} = h^{l-1} + \text{SA}(\text{LN}(h^{l-1}))$$

$$h^{l} = \hat{h}^{l} + \mathbf{G}_{\mathbf{A}}(\hat{h}^{l}, \text{TE})$$

$$h^{l} = h^{l} + \text{MLP}(\text{LN}(h^{l}))$$
(2)

where $\mathbf{G}_{\mathbf{A}}(\hat{h}^{l}, \mathrm{TE}) = \mathcal{F}_{2}(\mathrm{TSA}(\mathrm{LN}(\mathcal{F}_{1}(\hat{h}^{l} + \mathrm{TE})))), h^{l-1}$ is output of block l-1, SA, LN and TE is self-attention, layer-norm and temporal positional embedding respectively. $\mathcal{F}_{1,2}$ is 2-layer MLPs with GeLU activation function. Not that for TSA computation, we reuse the spatial self-attention weights similar to [20, 36]. **Discussion.** As compared to [20, 36], our adapter module has two major differences. (1) As the image transformer encoder $\mathbf{T}_{\mathbf{a},\mathbf{s}}$ is trained using frames of the video, we don't need to perform any spatial adaptation. (2) As discussed in Sec. 1, the temporal modeling for CVGL can be greatly relaxed due to the simplicity of the task. Particularly, to model temporal relations, we do not need to compute temporal attention with every patch of every input frame. Instead, we can use the global representation of every frame ([CLS] token) and attend them across the temporal dimension. To facilitate this, we compute the spatial self-attention prior to computing TSA in every block, which is in reverse order to [36]. We experimentally demonstrate that, for CVGL, computing TSA on [CLS] tokens outperforms TSA on all patches in terms of retrieval performance and computational efficiency (see Sec. 4.6).

Adapting the aerial encoder $\mathbf{T_a}$: In order to extend the aerial encoder $\mathbf{T_a}$ to obtain embeddings for larger geographical regions, we perform a pre-processing step before adding $\mathbf{G_A}$. Since each *large* aerial image covers approximately 49 times (for GAMa dataset [32]) more area than a *small* aerial image, to ensure that there is no distribution shift in the encoder $\mathbf{T_a}$, which has been trained on small aerial images, we divide the large aerial image into 49 non-overlapping patches. Each patch is of the same size as the small aerial image. Formally, given $I_L^a \in \mathbb{R}^{kn \times kn}$ with $I_s^a \in \mathbb{R}^{n \times n}$, we break I_L^a into k^2 non-overlapping patches of size $n \times n$, i.e., $I_L^a = \{I_{p_i}^a\}_{i=1}^{k^2}$. These patches of images can be considered analogous to the street-view frames but for the aerial encoder. Note that there is no one-to-one correspondence between the aerial patch $I_{p_i}^a$ and street-view frame V_k^s . We only know matching pairs of street-view videos and large aerial images.

Similar to Sec. 3.2, the unified model **U** is trained using contrastive learning with soft-margin triplet loss. We freeze the encoders $\mathbf{T}_{\mathbf{a},\mathbf{s}}$ and optimize only the added GeoAdapter modules $\mathbf{G}_{\mathbf{A}}$. Given street-view videos and large aerial images, our unified model learns to embed matching pairs closer to each other while pushing non-matching pairs farther apart. Formally, we optimize the following

objective:

$$\mathcal{L}_{2} = \sum_{i} \log \left(1 + e^{\alpha(||\mathbf{U}_{\mathbf{s}}(V^{s_{i}}) - \mathbf{U}_{\mathbf{a}}(I_{L}^{a_{i}})||_{2} - ||\mathbf{U}_{\mathbf{s}}(V^{s_{i}}) - \mathbf{U}_{\mathbf{a}}(I_{L}^{a_{k}\neq i})||_{2})} \right)$$
(3)

where, $|| \cdot ||_2$ denotes l^2 norm and $\mathbf{U}_{\mathbf{a},\mathbf{s}} : \mathbb{R}^{N \times H \times W \times 3} \to \mathbb{R}^d$.

3.4 TransRetriever: Temporally Consistent Retrieval

Given V_s , we obtain the k nearest neighbor large aerial images in the embedding space of **U**. Using the obtained k large images, we create a small aerial reference gallery \mathcal{G} , similar to [32]. For each frame V_i^s of the street-view video, we obtain t nearest neighbor small aerial images in the embedding space of **T** from \mathcal{G} . In order to obtain final temporally consistent frame-to-frame GPS predictions, we formulate the retrieval task as the Generalized Traveling Salesman problem (GTSP) [19]. Precisely, given $V^s = \{V_1^s, V_2^s, ..., V_n^s\}$ with $N = \{N_1, N_2, ..., N_n\}$ where $N_i = \{I_{s_1}^a, I_{s_2}^a, ..., I_{s_t}^a\}$ is the set of t small aerial nearest neighbor of V_i^s , we want to predict one small aerial image from each set N_i such that, the prediction from N_i and N_{i+1} is in closer proximity to each other.

Overview: Similar to [6], we cast our GTSP problem as a 'translation' problem where the source 'language' is the set N (analogous to clusters of cities), and the target 'language' is n small aerial images, one from each N_i (analogous to one city from each cluster). Analogous to translation architectures in NLP [8], our proposed TransRetriever has an encoder-decoder architecture, i.e., $\mathbf{T}_{\mathbf{AR}} = \{\mathbf{E}, \mathbf{D}\}$. The source 'language', i.e., the set N, is encoded by the encoder, and the decoder autoregressively predicts the target small aerial images using the encoded representations. Since the elements of the set N are ordered (as the frames of V^s are ordered), we use positional encoding to induce the ordering before encoding. To obtain a global context before decoding, we append a learnable token [START] (analogous to [CLS] token in VIT [10]) to the set N. The decoder takes the encoder representation of the [START] token to begin the decoding phase, where at every iteration i, we obtain a small aerial image that is the final prediction for street-view frame V_i^s . A comprehensive architecture diagram for $\mathbf{T}_{\mathbf{AR}}$ is provided in the supplementary material.

Methodology: We begin by obtaining a 3-dimensional representation for each small aerial in N. Precisely, for all $N_i \in N$, for all $I_{s_j}^a \in N_i$, we convert its GPS coordinate to UTM coordinates and obtain its similarity score $S_{ij} = C_S(\mathbf{T_a}(I_{s_j}^a), \mathbf{T_s}(V_i^s))$, where C_S is cosine similarity. Each small aerial image $I_{s_j}^a \in N_i$ is then represented by a 3-D vector $r_{ij} = [\text{UTM}_x, \text{UTM}_y, S_{ij}]$. We will refer to r_{ij} as tokens from hereon for simplicity. Since each N_i represents a set of neighbors for street-view frame V_i^s , an inherent ordering exists for all N_i and N_{i+1} . To induce such an ordering in the encoding phase, we add positional

9

embedding to each r_{ij} based on the N_i it belongs to. Specifically, we add a positional embedding PE_i to every r_{ij} in N_i . Note that all the tokens belonging to one neighbor set N_i will be added with the same positional embedding PE_i, e.g., r_{12} and r_{13} will be added with PE₁ while r_{23} and r_{21} will be added with PE₂ and so on. The resulting tokens $x_{ij} = r_{ij} + \text{PE}_i$ along with the learnable token [START] is encoded through the encoder $\mathbf{E}(\cdot)$ to obtain $[h_{[\text{START}]}, h_{ij}] = \mathbf{E}([[\text{START}], x_{ij}])$.

Next, we utilize the encoded representation $h_{[START]}$ as a starting token for the decoder to predict the small aerial image for the first street-view frame V_1^s . Similar to the encoding phase, to induce ordering as well as to enforce positionspecific decoding, we add the encoded representation with its corresponding positional embedding before passing it to the decoder. Formally, in the decoder $\mathbf{D}(\cdot)$, to predict the small aerial image for the *i*th street-view frame V_i^s , we add the encoded token $h_{(i-1)j}$ (the prediction of the previous frame V_{i-1}^s) with positional embedding PE_i and pass it to the decoder and compute causal selfattention (masked self-attention) [6, 31] to induce autoregressive conditioning followed by a cross-attention with only the encoded representations $r_{ij} \in N_i$. This enforces that the prediction of the i^{th} street-view frame comes only from set N_i . The decoding starts with $h_{0j} = h_{[START]}$. For each input token h_{ij} , the decoder predicts a probability distribution p_i that represents the probability of each encoded token in N_i to be selected. During inference, we take the token with the highest probability as prediction and pass it to the next iteration of the decoder. Formally, we have, $p_i = \mathbf{D}(h_{(i-1)j} + \mathrm{PE}_i, [h_{i1}, h_{i2}, ..., h_{it}])$ with $\hat{j}_i = \operatorname{argmax}(p_i)$ and $h_{(i+1)j} = h_{i\hat{j}_i}$. Evidently, the prediction for the *i*th streetview frame V_i^s will be $r_{i\hat{j}_i}$ which we convert back to GPS coordinates to obtain the final GPS predictions.

4 Experiments and Results

4.1 Training and Implementation details

We use DeiT-m [30] as our image transformer architecture and use PyTorch [21] to implement our framework. We train the models \mathbf{T} and $\mathbf{G}_{\mathbf{A}}$ for 100 epochs with a batch size of 64 and 8, respectively. For both models, we use Adam [15] optimizer with a learning rate of 1e-4 and default settings for other parameters. We use pre-trained weights from ImageNet to initialize our image encoder \mathbf{T} and initialize all the parameters in the $\mathbf{G}_{\mathbf{A}}$ module with 0. For our proposed TransRetriever architecture $\mathbf{T}_{\mathbf{A}\mathbf{R}}$, we have a 6-layer encoder with a 2-layer decoder with embedding dimension 128 and feed-forward dimension 512. Following [6], we use Batch Normalization [13] instead of Layer Normalization [1]. We use sinusoidal positional embedding [31] and train the model for 500 epochs.

4.2 Datasets

GAMa: GAMa [32] dataset is a cross-view video geo-localization (CVGL) dataset that aims to provide a comprehensive benchmark for evaluating CVGL methods in a frame-to-frame setting. The datasets provide ~ 40 seconds long street-view videos from the BDD100K [37] dataset with matching aerial images

corresponding to each geo-tagged street-view frame of the video. Additionally, each video is accompanied by a large aerial image that helps to localize the entire video to a larger geographical region before obtaining frame-by-frame predictions. The dataset provides 45029 street-view videos and large aerial image pairs with 1.68 million frame-level matching aerial images.

SeqGeo: SeqGeo [43] dataset is an image sequence geolocalization dataset consisting of high-resolution frontal camera street-view images of 1920×1080 resolution with a field-of-view of approximately 120° . Each street-view image is a part of a video shot at 1 FPS, with the distance between each frame (or capture point) at approximately 8 meters. Additionally, each street-view image accompanies a GPS location and camera heading (compass direction) angle. The total number of street-view images is 118,549, grouped into sequences of ~ 7 images covering 50 meters of ground distance per sequence. Each sequence includes an aerial image of the region covered by the street-view images. In total, there are 38,863 pairs of aerial images and street-view image sequences covering approximately 500 kilometers of roads in Vermont.

4.3 Sequence-to-Image Inference

To assess the effectiveness of our approach, we initially showcase the performance of our method in retrieving large aerial images given a street-view video. We use our trained unified model **U** to encode the given street-view video and find the best matching large aerial image by calculating the nearest neighbor in the feature space. Following previous works [28, 32, 43], we report top-k retrieval recall at k = 1, 5, 10, and 1% to compare our model with other state-of-the-art methods. A match is correct if the retrieved k nearest neighbours consists the correct large aerial image. We compare our proposed approach on the GAMa [32] dataset against Screening Network from [32], CVLNet [32] (since CVLNet requires ground-truth camera parameters, we compare our method with the "Ours w/o GVP (Unet)" and "Ours w/o GVP" variants for fairness.), and generalpurpose video models VideoSWIN [18] (initialized with pretrained weights of Kinetics400 [14]), TimeSformer [3](initialized with pretrained weights of ImageNet [9]). To further establish the effectiveness of our GeoAdapter module $\mathbf{G}_{\mathbf{A}}$, we compare our method with two additional baselines that are extensions of TransGeo [44] to video inputs.

Baseline 1: As the first baseline, we train the image encoder $\mathbf{T}_{\mathbf{a},\mathbf{s}}$ with video inputs without our proposed adapter $\mathbf{G}_{\mathbf{A}}$. Specifically, we perform the preprocessing step on the large aerial image I_L^a (see Sec. 3.3) and encode each patch individually using the encoder $\mathbf{T}_{\mathbf{a}}$. The embedding of the large aerial image is thus the average of the individual embeddings of each patch, $I_{p_i}^a$. Precisely, $f_L = \frac{1}{k^2} \sum_{i=1}^{k^2} \mathbf{T}_{\mathbf{a}}(I_{p_i}^a)$, where k^2 is the total number of non-overlapping patches of the large aerial image I_L^a . Similarly, the embedding for a street-view video V^s is obtained by computing the average of the individual embeddings of each frame V_i^s , i.e., $f_V = \frac{1}{n} \sum_{i=1}^{\mathbf{T}_s} (V_i^s)$, where n is the total number of frames in V^s . We train the encoders $\mathbf{T}_{\mathbf{a},\mathbf{s}}$ with Eq. (3) with f_V and f_L as video and large

aerial embeddings. Note that, for a fair comparison, we initialize the encoders $T_{a,s}$ with weights obtained after pretraining with the images following Sec. 3.2. **Baseline 2:** To establish the effectiveness of our preprocessing step during adaptation of T_a with our GeoAdapter module (see Sec. 3.3), we compare our approach with another baseline. In this, we perform the adaptation of only the street-view encoder T_s and train T_a (without the adapter) by resizing the large aerial image to size 512×512 . Additional details about these experiments can be found in supplementary material.

Tab. 1 shows the quantitative result of our approach on large aerial inference. The baseline 1 performs better than both general-purpose video models, establishing the advantage of image-level pretraining in CVGL. Also, our method with both **T** and **G**_A performs superior to all other methods, achieving a top-1 recall rate of 50.69. To further demonstrate the performance, we compare our method against the state-of-the-art cross-view sequence geo-localization benchmark SeqGeo [43]. Note that even though SeqGeo is an image sequence dataset (sequence-to-image), our method can be applied here without loss of generality. Specifically, instead of ~ 40 street-view images (as in GAMa), SeqGeo consists of ~ 7 street-view images per aerial image. As it can be inferred from the results that, our method also performs significantly better on the SeqGeo benchmark nearly doubling the state-of-the-art achieving a top-1 recall rate of 3.34.

Discussion. It is important to note that while Baseline 1 excels in large aerial inference, the encoders used to achieve this are entirely trained on video data. Therefore, once the larger aerial region is localized, the same encoder cannot be used for frame-by-frame matching. Whereas, with our proposed approach, once the large aerial region is obtained, we can disengage the adapter $\mathbf{G}_{\mathbf{A}}$ in the encoder and use it for frame-by-frame matching. Further details on this behavior can be found in the supplementary material.



Fig. 2: Examples of trajectories obtained using NN (A) based retrieval and our proposed TransRetriever (B). NN-based retrieval heavenly suffers from temporally inconsistent predictions depicted by the jumps in the trajectory while TransRetriever predictions are globally consistent which preserves the temporal coherence of the predictions.

4.4 Frame-to-frame inference

After localizing a given street-view video to a large aerial region, similar to [32], we create a small aerial image reference gallery \mathcal{G} . To obtain predictions for

Model	R@1 (↑)	R@5 (\uparrow)	R@10 (^)	R@1% (†)				
GAMa dataset								
TimeSformer [3]	20.10	44.47	55.58	83.46				
VideoSWIN [18]	20.37	45.86	59.94	88.02				
CVL-Net [28]	0.36	1.33	2.66	15.35				
GAMa(Screening-Network) [32]	12.2	-	35.3	49.3				
Baseline 1	40.07	70.76	79.72	94.04				
Baseline 2	15.39	38.48	51.36	83.16				
$\mathbf{Ours}~(\mathbf{T}+\mathbf{G}_{\mathbf{A}})$	50.69	81.77	88.71	98.26				
SeqGeo dataset								
VIGOR [46]	0.54	2.52	4.48	18.55				
SAFA [25]	0.63	2.83	5.03	21.51				
SeqGeo [43]	1.80	6.45	10.36	34.38				
$\mathbf{Ours} \ (\mathbf{T} + \mathbf{G}_{\mathbf{A}})$	3.34	11.19	17.18	44.39				

Table 1: Top-k retrieval recall score for large aerial image localization. Our approach $(\mathbf{T} + \mathbf{G}_{\mathbf{A}})$, achieves superior top-1 recall performance of 50.69 on the GAMa dataset, surpassing both baseline models and state-of-the-art methods. Additionally, our method significantly improves upon the state-of-the-art SeqGeo benchmark, achieving a top-1 recall rate of 3.34. (Arrows represent more (\uparrow) or less (\downarrow) is better, '-' denotes scores not provided in respective works)

each street-view frame of the input video, we use our image encoder model $\mathbf{T}_{\mathbf{a},\mathbf{s}}$ to retrieve nearest neighbor small aerial images from the gallery \mathcal{G} . Similar to Sec. 4.3, we report top-k retrieval recall at k = 1, 5, 10 and 1%. Due to high degree of overlap in aerial images of subsequent video frames, following the previous work [32], we consider a match correct if the predicted GPS is within the range of 0.05 miles of the ground-truth location. We compare our framework against GAMa-Net [32], L2LTR [35], VIGOR [46] and [25]. Tab. 4 shows the quantitative results of our method in comparison to other methods. Note that the recall rate for GAMa-Net uses top-1% large aerial images for creating the small aerial reference gallery, while our method uses top-10 large aerial images. Our proposed framework achieves significant improvement over state-of-art scoring 54.64% top-1 recall rate on the validation split of the GAMa dataset.

4.5 Retrieval with TransRetriever

Following Sec. 3.4, we compare our proposed TransRetriever with nearest neighbor (NN) retrieval and Dominant sets (DS) strategy from [cite]. For predicting using NN, we obtain the nearest small aerial image to a given street-view frame individually and predict its GPS coordinates as the final prediction. For DS and TransRetriever, we obtain the k nearest small aerial image candidates for each frame and predict the GPS coordinates by selecting one small aerial image from each k neighbor (see Sec. 3.4). Additional information on how we implement the DS strategy for our use case can be found in the supplementary material. In Tab. 2, we show the top-1 recall rate of our method compared to NN and DS for k = 10 and k = 20. Fig. 2 shows examples of the GPS trajectory obtained for a street-view video using NN and TransRetriever. The red point represents the

GPS prediction for a frame in the video, and lines connect adjacent frame predictions. NN-based retrieval heavenly suffers from temporally inconsistent predictions depicted by the jumps in the trajectory. At the same time, TransRetriever predictions are globally consistent, reducing inconsistencies and preserving the temporal coherence in the predictions.

Mathad	D@1 (本)	Adapter	$R@1 (\uparrow) R@5 (\uparrow) R@10 (\uparrow)$		
Method	$\frac{\mathbf{R} \otimes \mathbf{I}(1)}{\mathbf{h} - 10 \mathbf{h} - 20}$	AIM [36]	48.06	80.50	88.25
Nearest Neighbor	$\frac{\kappa - 10 \kappa - 20}{54.64}$	ST-Adapter [20]	34.29	69.98	80.75
Dominant Sets [29]	56.09 60.67	Ours	50.69	81.77	88.71
TransRetriever (TAP) 67.66 70.70				

 Table 2: The top-1 frame-by-frame recall
 bols follow the definition in Tab. 1)

Table 3: Top-k retrieval recall score of Sequence-to-Image inference with rate of our method with nearest neighbor, AIM [36], ST-Adapter [20], and our dominant sets, and our proposed Tran- proposed GeoAdapter (Ours) as adapter sRetriever with k = 10 and k = 20. (Sym-modules. (Symbols follow the definition in Tab. 1)

Model	R@1 (†)	R@5 (\uparrow)	R@10 (†)	R@1% (†)
Shi et al. [26]	9.6	18.1	26.6	71.9
L2LTR [35]	11.7	20.8	28.2	87.1
GAMa-Net [32]	15.2	27.2	33.8	91.9
GAMa-Net (Hierarchical) [32]	18.3	27.6	32.7	-
Ours (\mathbf{T})	54.64	70.45	76.36	91.92

Table 4: Top-k retrieval recall score for frame-by-frame geo-localization on the GAMa dataset. For our method, we construct the gallery \mathcal{G} using the top 10 large aerial images, while other methods use the top 1% images. Compared to the state-of-the-art, our method achieves significant improvement achieving a final score of 54.65 on frameto-frame cross-view geo-localization. (Symbols follow the definition in Tab. 1)

4.6 Analysis

In this section, we present additional analysis experiments to provide further support for our proposed architecture.

Variants of G_A : Following the discussion in Sec. 3.3, we compare our proposed GeoAdapter architecture with two other variants. Precisely, in $\mathbf{G}_{\mathbf{A}}^{\mathrm{ALL}}$, we compute temporal attention across every patch of every input frame. While in $\mathbf{G}_{\mathtt{A}}^{\mathrm{ASYM}},$ we compute temporal attention across <code>[CLS]</code> tokens in the aerial branch but compute temporal attention across all the patches of every input frame in the street-view branch. In Fig. 3, we compare both these variants with our proposed G_A where we compute temporal attention across the [CLS] tokens in both branches. It can be inferred from the figure that out of all three variants, the [CLS] token-based temporal attention is best suited for CVGL.

Effect of Top-k large aerial images: To further understand the effect of different numbers of large aerial images used to create a small aerial gallery \mathcal{G} . Fig. 4 shows the top-k recall rate of frame-by-frame inference using our method when multiple large aerial images are taken to create the gallery.

Comparison with other Adapters: In order to showcase the efficacy of our proposed GeoAdapter module G_A , we compare the Sequence-to-Image inference

performance of our methodology using adapters from [36] and [20]. Specifically, we replace our adapter module with AIM and ST-Adapter without modifying any other component. Tab. 3 shows the top-k retrieval recall score for Sequence-to-Image inference with different adapters. It should be noted that the AIM adapter performs comparably to $\mathbf{G}_{\mathbf{A}}^{\mathrm{ALL}}$ variant due to its similarity in temporal attention computation. In addition to better recall rate, the temporal attention computation in our method can be computationally lighter than in AIM or other variants of $\mathbf{G}_{\mathbf{A}}$ because we only compute attention across [CLS] token of each input frame instead of all input tokens. Additional experiments showing our method's recall rate with varying distance thresholds and runtime analysis are presented in the supplementary material.



Fig. 3: Comparison of our proposed GeoAdapter module with different variants of architectural design. With a top-1 recall rate of 50.7 (50.69), our proposed architecture best suits CVGL.



Fig. 4: Top-k retrieval recall score of frame-by-frame inference using our method when multiple large aerial images are taken to create the gallery.

5 Conclusion

In conclusion, this work introduces **GAReT**, a fully transformer-based method for cross-view video geo-localization (CVGL) that, unlike prior works, doesn't rely on camera parameters or odometry. We proposed *GeoAdapter*, a transformeradapter to aggregate image-level representations of an image geo-localization method and adapt it to video inputs, making our method computationally lightweight than existing works. Additionally, to ensure temporally consistent GPS predictions, we proposed *TransRetriever*, an encoder-decoder transformer model that encodes the top-k nearest neighbor predictions for each frame and auto-regressively decodes the best result based on the previous frame's predictions. Extensive experimental evaluations demonstrate the effectiveness of our approach, showcasing state-of-the-art performance on benchmark datasets.

Acknowledgments

This research is funded by an academic grant from the National Geospatial-Intelligence Agency (NGA) (Award # HM0476-20-1-0001, Project Title: Estimating Geospatial Trajectories of Videos Using Cross-View Image Matching) We thank all anonymous reviewers and ACs for their constructive suggestions.

References

- 1. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)
- Berrabah, S.A., Sahli, H., Baudoin, Y.: Visual-based simultaneous localization and mapping and global positioning system correction for geo-localization of a mobile robot. Measurement Science and Technology 22(12), 124003 (2011)
- 3. Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? In: ICML. vol. 2, p. 4 (2021)
- Berton, G., Masone, C., Caputo, B.: Rethinking visual geo-localization for largescale applications. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4878–4888 (2022)
- Brejcha, J., Čadík, M.: State-of-the-art in visual geo-localization. Pattern Analysis and Applications 20, 613–637 (2017)
- Bresson, X., Laurent, T.: The transformer network for the traveling salesman problem. arXiv preprint arXiv:2103.03012 (2021)
- Brosh, E., Friedmann, M., Kadar, I., Yitzhak Lavy, L., Levi, E., Rippa, S., Lempert, Y., Fernandez-Ruiz, B., Herzig, R., Darrell, T.: Accurate visual localization for automotive applications. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019)
- Cho, K., Van Merriënboer, B., Bahdanau, D., Bengio, Y.: On the properties of neural machine translation: Encoder-decoder approaches. arXiv preprint arXiv:1409.1259 (2014)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A largescale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
- Fervers, F., Bullinger, S., Bodensteiner, C., Arens, M., Stiefelhagen, R.: Uncertainty-aware vision-based metric cross-view geolocalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21621–21631 (2023)
- Hu, S., Feng, M., Nguyen, R.M., Lee, G.H.: 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. pp. 7258–7267 (2018)
- Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: International conference on machine learning. pp. 448–456. pmlr (2015)
- Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017)

- 16 M. Pillai et al.
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
- Lentsch, T., Xia, Z., Caesar, H., Kooij, J.F.: Slicematch: Geometry-guided aggregation for cross-view pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17225–17234 (2023)
- Liu, L., Li, H.: Lending orientation to neural networks for cross-view geolocalization. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5624–5633 (2019)
- Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H.: Video swin transformer. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3202–3211 (2022)
- Noon, C., Bean, J.: An efficient transformation of the generalized traveling salesman problem. INFOR. Information Systems and Operational Research 31 (02 1993). https://doi.org/10.1080/03155986.1993.11732212
- Pan, J., Lin, Z., Zhu, X., Shao, J., Li, H.: St-adapter: Parameter-efficient imageto-video transfer learning. Advances in Neural Information Processing Systems 35, 26462–26477 (2022)
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, highperformance deep learning library. Advances in neural information processing systems **32** (2019)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
- Regmi, K., Shah, M.: Video geo-localization employing geo-temporal feature learning and gps trajectory smoothing. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12126–12135 (2021)
- Shi, Y., Li, H.: Beyond cross-view image retrieval: Highly accurate vehicle localization using satellite image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17010–17020 (2022)
- Shi, Y., Liu, L., Yu, X., Li, H.: Spatial-aware feature aggregation for image based cross-view geo-localization. Advances in Neural Information Processing Systems 32 (2019)
- 26. Shi, Y., Yu, X., Campbell, D., Li, H.: Where am i looking at? joint location and orientation estimation by cross-view matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4064–4072 (2020)
- 27. Shi, Y., Yu, X., Liu, L., Campbell, D., Koniusz, P., Li, H.: Accurate 3-dof camera geo-localization via ground-to-satellite image matching. IEEE Transactions on Pattern Analysis and Machine Intelligence 45(3), 2682–2697 (2022)
- Shi, Y., Yu, X., Wang, S., Li, H.: Cvlnet: Cross-view semantic correspondence learning for video-based camera localization. In: Asian Conference on Computer Vision. pp. 123–141. Springer (2022)
- Tian, Y., Chen, C., Shah, M.: Cross-view image matching for geo-localization in urban environments. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3608–3616 (2017)
- Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jegou, H.: Training data-efficient image transformers & amp; distillation through attention. In: Meila, M., Zhang, T. (eds.) Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 139, pp. 10347–10357.

 $\label{eq:PMLR} PMLR \ (18-24 \ Jul \ 2021), \ \texttt{https://proceedings.mlr.press/v139/touvron21a.html}$

- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)
- Vyas, S., Chen, C., Shah, M.: Gama: Cross-view video geo-localization. In: European Conference on Computer Vision. pp. 440–456. Springer (2022)
- Xia, Z., Booij, O., Kooij, J.F.: Convolutional cross-view pose estimation. arXiv preprint arXiv:2303.05915 (2023)
- Xia, Z., Booij, O., Manfredi, M., Kooij, J.F.: Visual cross-view metric localization with dense uncertainty estimates. In: European Conference on Computer Vision. pp. 90–106. Springer (2022)
- Yang, H., Lu, X., Zhu, Y.: Cross-view geo-localization with layer-to-layer transformer. Advances in Neural Information Processing Systems 34, 29009–29020 (2021)
- Yang, T., Zhu, Y., Xie, Y., Zhang, A., Chen, C., Li, M.: Aim: Adapting image models for efficient video action recognition. In: The Eleventh International Conference on Learning Representations (2022)
- 37. Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., Darrell, T.: Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2636–2645 (2020)
- Zamir, A.R.: Visual geo-localization and location-aware image understanding (2014)
- Zamir, A.R., Shah, M.: Image geo-localization based on multiplenearest neighbor feature matching usinggeneralized graphs. IEEE transactions on pattern analysis and machine intelligence 36(8), 1546–1558 (2014)
- Zhang, X., Li, X., Sultani, W., Chen, C., Wshah, S.: Geodtr+: Toward generic cross-view geolocalization via geometric disentanglement. arXiv preprint arXiv:2308.09624 (2023)
- Zhang, X., Li, X., Sultani, W., Zhou, Y., Wshah, S.: Cross-view geo-localization via learning disentangled geometric layout correspondence. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 3480–3488 (2023)
- 42. Zhang, X., Li, X., Sultani, W., Zhou, Y., Wshah, S.: Cross-view geo-localization via learning disentangled geometric layout correspondence. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 3480–3488 (2023)
- Zhang, X., Sultani, W., Wshah, S.: Cross-view image sequence geo-localization. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2914–2923 (2023)
- 44. Zhu, S., Shah, M., Chen, C.: Transgeo: Transformer is all you need for cross-view image geo-localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1162–1171 (2022)
- 45. Zhu, S., Yang, T., Chen, C.: Revisiting street-to-aerial view image geo-localization and orientation estimation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 756–765 (2021)
- Zhu, S., Yang, T., Chen, C.: Vigor: Cross-view image geo-localization beyond oneto-one retrieval. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3640–3649 (2021)