

# Where in the World is this Image? Transformer-based Geo-localization in the Wild

Shraman Pramanick<sup>1</sup>, Ewa M. Nowara<sup>1</sup>, Joshua Gleason<sup>2</sup>,  
Carlos D. Castillo<sup>1</sup>, and Rama Chellappa<sup>1</sup>

<sup>1</sup> Johns Hopkins University

<sup>2</sup> University of Maryland, College Park

{spraman3, carlosdc, rchella4}@jhu.edu,  
ewa.m.nowara@gmail.com, gleason@umd.edu

**Abstract.** Predicting the geographic location (geo-localization) from a single ground-level RGB image taken anywhere in the world is a very challenging problem. The challenges include huge diversity of images due to different environmental scenarios, drastic variation in the appearance of the same location depending on the time of the day, weather, season, and more importantly, the prediction is made from a single image possibly having only a few geo-locating cues. For these reasons, most existing works are restricted to specific cities, imagery, or worldwide landmarks. In this work, we focus on developing an efficient solution to planet-scale single-image geo-localization. To this end, we propose TransLocator, a unified dual-branch transformer network that attends to tiny details over the entire image and produces robust feature representation under extreme appearance variations. TransLocator takes an RGB image and its semantic segmentation map as inputs, interacts between its two parallel branches after each transformer layer and simultaneously performs geo-localization and scene recognition in a multi-task fashion. We evaluate TransLocator on four benchmark datasets - Im2GPS, Im2GPS3k, YFCC4k, YFCC26k and obtain 5.5%, 14.1%, 4.9%, 9.9% continent-level accuracy improvement over the state-of-the-art. TransLocator is also validated on real-world test images and found to be more effective than previous methods.

**Keywords:** Geo-location estimation · Vision transformer · Multi-task learning · Semantic segmentation

## 1 Introduction

Can we determine the location of a scene given a single ground-level RGB image? For famous and characteristic scenes, such as the Eiffel Tower, it is trivial because the landmark is so distinctive of Paris. Moreover, there is a lot of images captured in such prominent locations under different viewing angles, at different times of the day, and even in different weather or lighting conditions. However, some scenes, especially places outside cities and tourist attractions, may not have characteristic landmarks and it is not so obvious to locate where they were snapped. This is the case for the vast majority of the places in the world. Moreover, such places are less popular and are less photographed. As a result, there are very few images from such locations and the

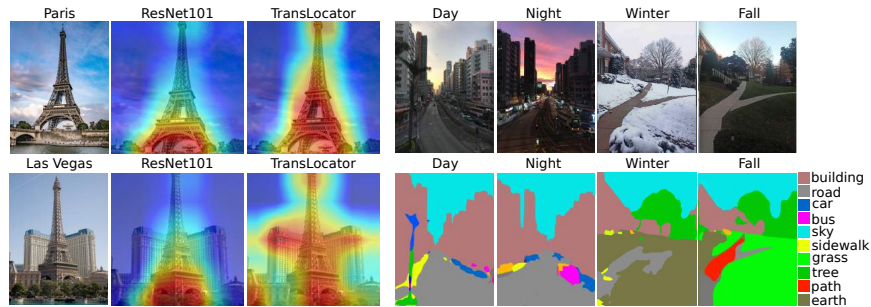


Fig. 1: *First three columns: Utilizing global context* - TransLocator focuses on the entire image to correctly geo-locate the Eiffel Tower and its replica as visualized with Grad-CAM [45] class activation maps. *Last four columns: Robustness to appearance variations* - Semantic segmentation maps are robust to extreme appearance variations.

existing ones do not capture a diversity of viewing angles, time of the day, or weather conditions, making it much harder to geo-locate. Because of the complexity of this problem, most existing geo-localization approaches have been constrained to small parts of the world [8, 16, 47, 55]. Recently, convolutional neural networks (CNNs) trained with large datasets have significantly improved the performance of geo-localization methods and enabled extending the task to the scale of the entire world [17, 18, 34, 63]. However, planet-scale unconstrained geo-localization is still a very challenging problem and existing state-of-the-art methods struggle to geo-locate images taken anywhere in the world.

In contrast to many other vision applications, single-image geo-localization often depends on fine-grained visual cues present in small regions of an image. In Figure 1, consider the photo of the Eiffel Tower in Paris and its replica in Las Vegas. Even though these two images seem to come from the same location, the buildings and vegetation in the background play a decisive role in distinguishing them. Similarly, in the case of most other images, the global context spanned over the entire image is more important than individual foreground objects in geo-localization. Recently, a few studies [12, 39] comparing vision transformer (ViT) with CNNs have revealed that the early aggregation of global information using self-attention enables transformers to build long-range dependencies within an image. Moreover, higher layers of ViT maintain spatial location information better than CNNs. Hence, we argue that transformer-based networks are more effective than CNNs for geo-localization because they focus on detailed visual cues present in the entire image.

Another challenge of single-image geo-localization is the drastic appearance variation of the exact location under different daytime or weather conditions. Semantic segmentation offers a solution to this problem by generating robust representations in such extreme variations [47]. For example, consider the drastic disparity of the RGB images of same locations in day and night or winter and fall in Figure 1. In contrast to the RGB images, the semantic segmentation maps remain almost unchanged. Furthermore, semantic segmentation provides auxiliary information about the objects present in the

image. This additional information can be a valuable pre-processing step since it enables the model to learn which objects occur more frequently in which geographic locations. For example, as soon as the semantic map detects mountains, the model immediately eliminates all flat regions, thus, reducing the complexity of the problem.

Planet-scale geo-localization deals with a diverse set of input images caused by different environmental settings (e.g., outdoors vs indoors), which entails different features to distinguish between them. For example, to geo-locate outdoor images, features such as the architecture of buildings or the type of vegetation are important. In contrast, for indoor images, the shape and style of furniture may be helpful. To address such variations, Muller et al. [34] proposed to train different networks for different environmental settings. Though such an approach produces good results, they are cost-prohibitive and are not generalizable to a higher number of environmental scenarios. In contrast, we propose a unified multi-task framework for simultaneous geo-localization and scene recognition applied to images from all environmental settings.

This work addresses the challenges of planet-scale single-image geo-localization by designing a novel dual-branch transformer architecture, TransLocator. We treat the problem as a classification task [34, 63] by subdividing the earth’s surface into a high number of geo-cells and assigning each image to one geo-cell. TransLocator takes an RGB image and its corresponding semantic segmentation map as input, divides them into non-overlapping patches, flattens the patches, and feeds them into two parallel transformer encoder modules to simultaneously predict the geo-cell and recognize the environmental scene in a multi-task framework. The two parallel transformer branches interact after every layer, ensuring an efficient fusion strategy. The resulting features learned by TransLocator are robust under appearance variation and focus on tiny details over the entire image.

In summary, our contributions are three-fold. (i) We propose TransLocator - a unified solution to planet-scale single-image geo-localization with a dual-branch transformer network. TransLocator is able to distinguish between similar images from different locations by precisely attending to tiny visual cues. (ii) We propose a simple yet efficient fusion of two transformer branches, which helps TransLocator to learn robust features under extreme appearance variation. (iii) We achieve state-of-the-art performance on four datasets with a significant improvement of 5.5%, 14.1%, 4.9%, 9.9% continent-level geolocational accuracy on Im2GPS [17], Im2GPS3k [18], YFCC4k [59], and YFCC26k [51], respectively. We also qualitatively evaluate the effectiveness of the proposed method on real-world images.

## 2 Related Works

### 2.1 Single-image geo-localization

**Small-scale approaches:** Planet-scale single-image geo-localization is difficult due to several challenges, including the large variety of images due to different environmental scenarios and drastic differences in the appearance of same location based on the weather, time of day, or season. For this reason, many existing approaches are limited to geo-locating images captured in very specific and constrained locations. For example, many

approaches have been restricted to geo-locating an image within a single city [3, 16, 47], such as Pittsburgh [56], San Francisco [8], or Tokyo [55]. Far fewer approaches have attempted geo-localization in natural, non-urban environments. Some have limited the problem to only very specific natural environments, such as beaches [5, 61], deserts [57], or mountains [2, 44].

**Cross-view approaches:** The challenge of large-scale image geo-localization with few landmarks and limited training data has led some researchers to propose cross-view approaches, which match a query ground-level RGB image with a large reference dataset of aerial or satellite images [21, 28, 41, 53, 54, 67, 78, 79]. However, these approaches require access to a relevant reference dataset which is not available for many locations.

**Planet-scale approaches:** Only a few approaches have attempted to geo-locate images on a scale of an entire world without any restrictions. Im2GPS [17] was the first work to geo-locate images taken anywhere on the earth using hand-crafted features and nearest neighbor search. The same authors have later improved their approach [18] by refining the search with multi-class support vector machines. PlaNet [63] was the first deep learning approach for unconstrained geo-localization which significantly outperformed the two Im2GPS approaches [17, 18]. Vo et al. [59] combined the Im2GPS and PlaNet approaches by using the deep-learning-based features to match a query image with a nearest neighbors search. The approach by Muller et al. [34] achieved the current state-of-the-art performance for unconstrained geo-localization. They fine-tuned three separate ResNet101 networks, each trained only on images of outdoor natural, outdoor urban, or indoor scenes. Detailed surveys of existing work on geo-localization can be found at [4, 32].

## 2.2 Vision Transformer

Following the success of transformers [58] in machine translation, convolution-free networks that rely only on self-attentive transformer layers have gone viral in computer vision. In particular, Vision Transformer (ViT) [13] was the first to apply a pure transformer architecture on non-overlapping image patches and surpassed CNNs for image classification. Inspired by ViT, transformers have been widely adopted for various computer vision tasks, such as object detection [6, 14, 26, 80], image segmentation [10, 48, 64, 74], video understanding [49, 62, 65, 69], low-shot learning [12, 71], image super resolution [9, 66], 3D classification [33, 73], and multimodal learning [1, 20, 23, 25, 27, 31, 68].

We are the first to address the ill-posed planet-scale single-image geo-localization problem by designing a novel dual-branch transformer architecture. Unlike [34], we use only a single network for all kinds of input images, allowing our model to take advantage of the similar features from different scenes while maintaining its ability to learn scene-specific features. Moreover, our approach is robust under extreme appearance variations and generalizes to challenging real-world images.

## 3 Proposed System - TransLocator

This section presents our proposed multimodal multi-task system, TransLocator for geo-location estimation and scene recognition. Following the literature [34, 46, 63], we

treat geo-localization as a classification problem by subdividing the earth into several geographical cells<sup>3</sup> containing a similar number of images. As the size and the number of geo-cells poses a trade-off between system performance and classification difficulty, we use multiple partitioning schemes which allow the system to learn geographical features at different scales, leading to a more discriminative classifier. Furthermore, we incorporate semantic representations of RGB images to learn robust feature representation across different daytime and weather conditions. Finally, we predict the geo-cell and environmental scenario of the input image by exploiting contextual information in a multi-task fashion.

### 3.1 Global Context with Vision Transformer

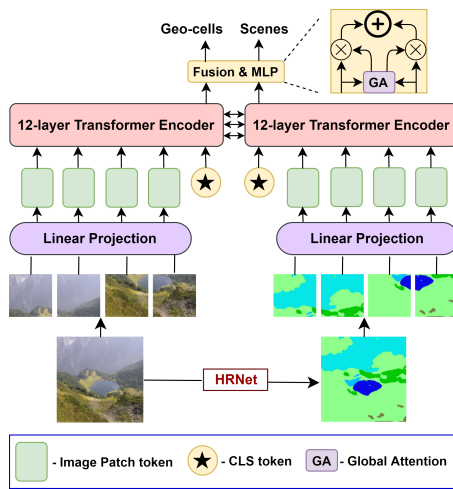


Fig. 2: **Overview of the proposed model - TransLocator.** A dual-branch transformer network is trained using RGB images and corresponding semantic segmentation maps in a multi-task learning framework to simultaneously predict the environmental scene and the geographic location (geo-cell) of the image.

We use a vision transformer (ViT) [13] model as the backbone of our architecture. Transformers have traditionally been used for sequential data, such as text or audio [11, 29, 36, 58, 70]. In order to extend transformers to vision problems, an image is split into a sequence of 2-D patches which are then flattened and fed into the stacked transformer encoders through a trainable linear projection layer. An additional classification token (CLS) is added to the sequence, as in the original BERT approach [11]. Moreover, positional embeddings are added to each token to preserve the order of the patches in the original image. A transformer encoder is made up with a sequence of blocks containing multi-headed self-attention (MSA) with a feed-forward network (FFN). FFN contains two multi-layer perceptron (MLP) layers with GELU non-linearity applied after the first layer. Layer normalization (LN) is applied before every block, and residual connections after every block. The output of  $k$ -th block  $x^{(k)}$  in the trans-

former encoder can be expressed as

$$x^{(0)} = [x_{cls}^{(0)} || x_{patch}^{(0)}] + x_{pos} \quad (1)$$

$$y^{(k)} = x^{(k-1)} + \text{MSA}(\text{LN}(x^{(k-1)})) \quad (2)$$

$$x^{(k)} = y^{(k)} + \text{FFN}(\text{LN}(y^{(k)})) \quad (3)$$

where  $x_{cls}^{(i)} \in \mathbb{R}^{1 \times C}$  is the CLS token,  $x_{patch}^{(i)} \in \mathbb{R}^{N \times C}$  is the patch tokens at  $i$ -th layer, and  $x_{pos} \in \mathbb{R}^{(1+N) \times C}$  is the positional embeddings.  $N$  and  $C$  is the patch-sequence length and embedding dimension, respectively.  $||$  denotes concatenation.

<sup>3</sup> Following [34], we use the *S2 geometry library* to generate the geo-cells.

The self-attention in ViT is a mechanism to aggregate information from all spatial locations and is structurally very different from the fixed-sized receptive field of CNNs. Every layer of ViT can access the whole image and, thus, learn the global structure more effectively. This particular attribute of ViT plays a vital role in our classification system. In agreement with Raghu et al. [39], we empirically establish how the global receptive field of ViT can help to attend to small but essential visual cues, which CNNs often neglect. We provide a detailed comparative evaluation of this phenomenon in Section 5.3.

### 3.2 Semantic Segmentation for Robustness to Appearance Variation

In order to improve the network’s ability to generalize to scenes captured in different daytime and weather conditions, we train a dual-branch vision transformer on the RGB images and their corresponding semantic segmentation maps. As shown in Figure 1, compared to RGB features, the semantic layout of a scene is generally more robust to drastic appearance variations. We use HRNet [60] pre-trained on ADE20K and scene parsing datasets [76, 77] to obtain high resolution semantic maps. HRNet assigns each pixel in the RGB image to one of the pre-defined 150 object classes, such as *sky*, *tree*, *sea*, *building*, *grass*, *pier*, etc.

**Multimodal Feature Fusion (MFF):** Our proposed framework contains two parallel transformer branches, one for the RGB image and the other for the corresponding semantic map. Since the two branches carry complementary information of the same input, effective fusion is the key for learning multimodal feature representations. We propose a simple and computationally light fusion scheme, where we sum the CLS tokens of each branch after every transformer encoder layer. At each layer, the CLS token is considered as an abstract global feature representation. This strategy is as effective as concatenating all feature tokens, but avoids quadratic complexity. Once the CLS tokens are fused, the information will be passed back to patch tokens at the later transformer encoder layers. More formally,  ${}^{(i)}x^{(k)}$ , the token sequence at  $k$ -th layer of a branch  $i$  can be expressed as

$${}^{(i)}x^{(k)} = \left[ g\left( \sum_{j \in \{\text{rgb}, \text{seg}\}} f({}^{(j)}x_{\text{cls}}^{(k)}) \right) || {}^{(i)}x_{\text{patch}}^{(k)} \right] \quad (4)$$

where  $f(\cdot)$  and  $g(\cdot)$  are the projection and back-projection functions used to align the dimensions.

Since the relative importance of the two branches depends upon the structure of the input image, we attentively fuse the CLS tokens from the last layer of each branch. Motivated by [15, 37, 38], we design our attention module with two major parts – modality attention generation and weighted concatenation. In the first part, a sequence of dense layers followed by a softmax layer is used to generate the attention scores  $w_{mm} = [w_{rgb}, w_{seg}]$  for the two branches. In the second part, the CLS tokens from the last transformer layer are weighted using their respective attention scores and concatenated together as follows

$$f_{rgb} = (1 + w_{rgb})^{rgb} x_{cls}^{(last)} \quad (5)$$

$$f_{seg} = (1 + w_{seg})^{seg} x_{cls}^{(last)} \quad (6)$$

$$f_{mm} = [f_{rgb} || f_{seg}] \quad (7)$$

We use residual connections to improve the gradient flow. The final multimodal representation,  $f_{mm}$ , is fed into fully-connected classifier heads.

### 3.3 Single Model with Multi-Task Learning

Different features are essential for various environmental settings, such as indoor and outdoor urban or natural scenes. Hence, geo-localization can benefit from contextual knowledge about the surroundings, and this information can reduce the complexity of the data space, thus simplifying the classification problem. Muller et al. [34] addressed this issue and approached the problem by fine-tuning three individual networks separately on natural, urban, and indoor images. However, one immediate drawback of their approach is that using multiple separate networks is cost-prohibitive and limits the number of scene kinds. In addition, the separately trained networks can not share the learned features, which likely have semantic similarities across different scene kinds, which effectively reduces the size and potency of the training set.

We address these two drawbacks by training a single network with a multi-task learning objective for simultaneous geo-localization and scene recognition. Following [34], we use a ResNet-152 pre-trained on *Places2* dataset<sup>4</sup> [75] to label the training images with corresponding 365 different scene categories. Additionally, based on the provided scene hierarchy<sup>5</sup>, we label each image with a coarser 16 and 3 different scene labels. In multi-task learning, adding complementary tasks has been proven to improve the results of the main task [7, 40, 72]. We follow a multi-task learning approach known as *hard parameter sharing* [43] which shares the parameters of hidden layers for all tasks and uses task-specific classifier heads. This strategy reduces overfitting because learning the same weights for multiple tasks forces the model to learn a generalized representation useful for the different tasks. More specifically, we only add a classifier head on top of the fused multimodal features and train the system end-to-end for both tasks.

### 3.4 Training Objective

Since there is a trade-off between the classification difficulty and the prediction accuracy caused by the number and size of the geo-cells, we use partitioning of the earth’s surface at three different resolutions, which we call *coarse*, *middle* and *fine* geo-cells<sup>6</sup>. We feed the final multimodal feature representation  $f_{mm}$  into four parallel classifier heads for the final classification: three for geo-cell prediction and one for scene recognition. We use a cross-entropy loss for each head. Our overall loss function can be summarized as follows:

$$\mathcal{L}_{total} = (1 - \alpha - \beta)\mathcal{L}_{geo}^{coarse} + \alpha\mathcal{L}_{geo}^{middle} + \beta\mathcal{L}_{geo}^{fine} + \gamma\mathcal{L}_{scene} \quad (8)$$

<sup>4</sup> *Places2 ResNet152 model*: <https://github.com/CSAILVision/places365>

<sup>5</sup> <http://places2.csail.mit.edu/download.html>

<sup>6</sup> Details on geo-cell partitioning can be found in the supplementary material.



## 4 Experiments

We conduct extensive experiments to show the effectiveness of our proposed method. In this section, we describe the datasets, the evaluation metrics, the baseline methods, and the detailed experimental settings.

### 4.1 Datasets

We use publicly available RGB image datasets with corresponding ground truth GPS (latitude, longitude) tags for training, validation, and testing. We trained our model on the MediaEval Placing Task 2016 dataset (MP-16) [24] containing 4.72M geo-tagged images sourced from Flickr. Like Vo et al. [59], during training we excluded images taken by the same authors in our validation or test sets. We validated and tested our model on two randomly sampled subsets of images from the Yahoo Flickr Creative Commons 100 Million dataset (YFCC100M) [52], referred to as YFCC26k [51] and YFCC4k [59] containing 25,600 and 4536 images, respectively. Since the images of MP-16, YFCC26k and YFCC4k were sourced without any scene and user restrictions, these datasets contain images of landmarks and landscapes, but also ambiguous images with little to no geographical cues, such as photographs of food and portraits of people.

We have additionally tested our model on two smaller datasets commonly used for geo-localization – Im2GPS [17] and Im2GPS3k [18]. Im2GPS contains 237 manually selected geo-localizable images. In contrast to the previous three datasets, Im2GPS is specially designed to evaluate geo-localization systems and contains images from popular landmarks and famous tourist locations. Im2GPS3k is an extended version of Im2GPS. Im2GPS3k contains 2997 images with geo-tags. Unlike Im2GPS, this dataset was not manually filtered and hence, it is a slightly more challenging test compared to Im2GPS.

### 4.2 Baselines

We compare our method to several existing geo-estimation methods, including **Im2GPS** [17], **[L]kNN** [59], **MvMF** [22], **Planet** [63], **CPlanet** [46], and **ISNs** [34]. ISNs reports the state-of-the-art results on Im2GPS and Im2GPS3k. More details about the baselines are provided in the supplementary materials. Since neither of the baselines reported their performance on all considered test sets, we re-implement the missing ones. We also provide a detailed ablation study of our method by removing one component at a time from TransLocator in Table 2.

### 4.3 Evaluation Metrics

We evaluate the performance of our approach using geolocational accuracy at multiple error levels, i.e. the percentage of images correctly localized within a predefined distance from the ground truth GPS coordinates [34, 46, 63]. Formally, if the predicted and ground truth coordinates are  $(lat_{pred}, lon_{pred})$  and  $(lat_{gt}, lon_{gt})$ , the geo-localational accuracy  $a_r$  at scale  $r$  (in km) is defined as follows for a set of  $N$  samples:

$$a_r \equiv \frac{1}{N} \sum_{i=1}^N u(GCD((lat_{pred}^{(i)}, lon_{pred}^{(i)}), (lat_{gt}^{(i)}, lon_{gt}^{(i)})) - r) \quad (9)$$



where GCD is the great circle distance and  $u(x) = \begin{cases} 1, & \text{if } x < 0 \\ 0, & \text{otherwise} \end{cases}$  is an indicator function whether the distance is smaller than the tolerated radius  $r$ . We report the results at 1 km, 25 km, 200 km, 750 km, and 2500 km ranges from the ground truth, which correspond to the scale of the same street, city, region, country, and continent, respectively [63].

Following Muller et al. [34], we assign the predicted geo-cell a GPS tag by using the mean locations of all training images in that cell. Since the models are trained in a classification framework and evaluated using a distance metric, we empirically verified a strong linear relationship between classification and geolocational accuracy. The average Pearson correlation coefficient between these two metrics on Im2GPS and Im2GPS3k test sets for TransLocator is 0.981, which implies the validity of the classification framework for geo-localization (see Figure 3). More details regarding the strong correlation between these two metrics can be found in the supplementary material.

#### 4.4 Implementation Details

We use ViT-B/16 [13] as the backbone of both RGB and segmentation channels. While training a single-channel (RGB only) multi-task transformer network, which we denote as ViT-MT, we use the pre-trained weights on the large ImageNet21K [42] dataset containing 14 million images and 21 thousand classes. We fine-tuned it for the geo-localization task. Both channels of the dual-branch TransLocator system are initialized with the weights of the ViT-MT backbone trained for 10 epochs. The weights of the classifier heads are randomly initialized with a zero-mean Gaussian distribution with a standard deviation of 0.02. Following the standard ViT literature, we linearly project the non-overlapping patches of  $16 \times 16$  pixels for both channels, and we add the CLS token and the positional embeddings.

Since the training set contains images in various resolutions and scales, we use extensive data augmentation, as detailed in the supplementary materials. We implement the methods in Pytorch [35] framework. Our ViT-MT and TransLocator took 6/10 days to train on 20 NVIDIA RTX 3090 GPUs, respectively, with 24 GB dedicated memory in each GPU. We train both systems using a AdamW [30] optimizer with an base learning




Fig. 3: **Strong linear relationship between the geolocational and classification accuracy metrics.** The positive correlation enables us to treat geo-localization as a classification problem.

Table 1: **Geolocational accuracy of the proposed systems compared to several baselines across four datasets and five scales.** The methods re-implemented by [46] are denoted by dagger( $\dagger$ ), and the methods re-implemented by us are denoted by double-dagger( $\ddagger$ ).  $\Delta$  indicates improvement over state-of-the-art achieved by TransLocator.

Dataset	Method	Distance ( $a_r$ [%] @ km)				
		Street 1 km	City 25 km	Region 200 km	Country 750 km	Continent 2500 km
Im2GPS [18]	Human [59]	–	–	3.8	13.9	39.3
	[L]kNN, $\sigma = 4$ [59]	14.4	33.3	47.7	61.6	73.4
	MvMF [22]	8.4	32.6	39.4	57.2	80.2
	PlaNet [63]	8.4	24.5	37.6	53.6	71.3
	CPlaNet [46]	16.5	37.1	46.4	62.0	78.5
	ISNs (M, f, $S_3$ ) [34]	16.5	42.2	51.9	66.2	81.0
	ISNs (M, $f^*$ , $S_3$ ) [34]	16.9	43.0	51.9	66.7	80.2
	ViT-MT	18.2	46.4	62.1	74.5	85.2
	TransLocator	<b>19.9</b>	<b>48.1</b>	<b>64.6</b>	<b>75.6</b>	<b>86.7</b>
$\Delta_{\text{Ours}} - \text{ISNs}$	3.0 $\uparrow$	5.1 $\uparrow$	12.7 $\uparrow$	8.9 $\uparrow$	5.5 $\uparrow$	
Im2GPS 3k [17]	[L]kNN, $\sigma = 4$ [59]	7.2	19.4	26.9	38.9	55.9
	PlaNet $\dagger$ [63]	8.5	24.8	34.3	48.4	64.6
	CPlaNet [46]	10.2	26.5	34.6	48.6	64.6
	ISNs (M, f, $S_3$ ) [34]	10.1	27.2	36.2	49.3	65.6
	ISNs (M, $f^*$ , $S_3$ ) [34]	10.5	28.0	36.6	49.7	66.0
	ViT-MT	11.0	29.0	42.6	54.8	71.6
	TransLocator	<b>11.8</b>	<b>31.1</b>	<b>46.7</b>	<b>58.9</b>	<b>80.1</b>
	$\Delta_{\text{Ours}} - \text{ISNs}$	1.3 $\uparrow$	3.1 $\uparrow$	6.1 $\uparrow$	9.2 $\uparrow$	14.1 $\uparrow$
YFCC 4k [51]	[L]kNN, $\sigma = 4$ [59]	2.3	5.7	11.0	23.5	42.0
	PlaNet $\dagger$ [63]	5.6	14.3	22.2	36.4	55.8
	CPlaNet [46]	7.9	14.8	21.9	36.4	55.5
	ISNs (M, f, $S_3$ ) $\ddagger$ [34]	6.5	16.2	23.8	37.4	55.0
	ISNs (M, $f^*$ , $S_3$ ) $\ddagger$ [34]	6.7	16.5	24.2	37.5	54.9
	ViT-MT	8.1	18.0	26.2	40.0	59.9
	TransLocator	<b>8.4</b>	<b>18.6</b>	<b>27.0</b>	<b>41.1</b>	<b>60.4</b>
	$\Delta_{\text{Ours}} - \text{CPlaNet}$	0.5 $\uparrow$	3.8 $\uparrow$	5.1 $\uparrow$	4.7 $\uparrow$	4.9 $\uparrow$
$\Delta_{\text{Ours}} - \text{ISNs}$	1.7 $\uparrow$	2.1 $\uparrow$	2.8 $\uparrow$	3.6 $\uparrow$	5.5 $\uparrow$	
YFCC 26k [59]	PlaNet $\ddagger$ [63]	4.4	11.0	16.9	28.5	47.7
	ISNs (M, f, $S_3$ ) $\ddagger$ [34]	5.3	12.1	18.8	31.8	50.6
	ISNs (M, $f^*$ , $S_3$ ) $\ddagger$ [34]	5.3	12.3	19.0	31.9	50.7
	ViT-MT	6.9	17.3	27.5	40.5	59.5
	TransLocator	<b>7.2</b>	<b>17.8</b>	<b>28.0</b>	<b>41.3</b>	<b>60.6</b>
	$\Delta_{\text{Ours}} - \text{ISNs}$	1.9 $\uparrow$	5.5 $\uparrow$	9.0 $\uparrow$	9.4 $\uparrow$	9.9 $\uparrow$

rate of 0.1, a momentum of 0.9, and a weight decay of 0.0001. We train the network for a total of 40 epochs with a batch size of 256. During testing, we convert the *fine* geo-cells to corresponding GPS coordinates. Other necessary hyper-parameters and data augmentation details are given in the supplementary materials.

## 5 Results, Discussions and Analysis

In this section, we compare the performance of TransLocator system with different baselines, and conduct a detailed ablation study to demonstrate the importance of different components in our system. Furthermore, we visualize the interpretability of TransLocator using Grad-CAM [45] and perform an error analysis.

Table 2: **Ablation Study on the Im2GPS and Im2GPS3K datasets.** *Seg* denotes the segmentation branch of TransLocator, *MFF* represents multimodal feature fusion, and *Scene* denotes the multi-task learning framework.

Dataset	Method	Distance ( $a_r$ [%] @ km)				
		Street 1 km	City 25 km	Region 200 km	Country 750 km	Continent 2500 km
Im2GPS [18]	ResNet101	14.3	41.4	51.9	64.1	78.9
	EfficientNet-B4	15.4	42.7	52.8	64.8	79.5
	ViT base	16.9	43.4	54.5	67.8	80.7
	+ Seg	17.6	44.8	58.9	70.0	83.3
	+ Seg + MFF	19.0	47.2	62.7	73.5	85.7
	+ Seg + MFF + Scene	<b>19.9</b>	<b>48.1</b>	<b>64.6</b>	<b>75.6</b>	<b>86.7</b>
Im2GPS 3k [18]	ResNet101	9.0	25.1	32.8	46.1	63.5
	EfficientNet-B4	9.2	26.8	32.7	47.0	63.9
	ViT base	9.9	28.0	37.8	54.2	70.7
	+ Seg	10.5	29.1	42.5	55.8	73.6
	+ Seg + MFF	11.1	30.2	45.0	56.8	78.1
	+ Seg + MFF + Scene	<b>11.8</b>	<b>31.1</b>	<b>46.7</b>	<b>58.9</b>	<b>80.1</b>

## 5.1 Comparison with Baselines

Table 1 presents the performance of our proposed TransLocator system and baseline methods on all four evaluation datasets. The reported baselines have a similar number of training images and geographic classes, and hence, we can directly compare the results of our system with them. Since the Im2GPS and Im2GPS3k dataset mainly contains images of landmarks and popular tourist locations, the systems yield high accuracy on these two datasets. On Im2GPS, even the earlier methods like PlaNet and MvMF surpass human performance considerably. CPlaNet, a combinational geo-partitioning approach, brings a substantial improvement of 8.1% in street-level accuracy over PlaNet. A similar trend of results is seen in the case of Im2GPS3k. On this dataset, CPlaNet beats PlaNet by 1.7% street-level accuracy. For other distance scales, the results improve proportionally. The Individual Scene Networks (ISNs) report the state-of-the-art result on both of these datasets, achieving 16.5% and 10.1% street-level accuracy on Im2GPS and Im2GPS3k, respectively. The ensemble of hierarchical classifications (denoted by  $f^*$ ) improves their results. However, note that none of these methods use semantic maps in their framework. Thus, we first implement the single-branch multi-task ViT-MT model. Interestingly, even this model produces a significant improvement over ISNs for all scales on both datasets, which can be explained by the global context used by ViT architecture. Our final dual-branch TransLocator system improves on top of ViT-MT. Overall, we push the current state-of-the-art by 3.0% and 1.3% street-level and 5.5% and 14.1% continent-level accuracy on Im2GPS and Im2GPS3k, respectively.

The YFCC4k and YFCC26k datasets contain more challenging samples which have little to no geo-locating cues. However, these large datasets of unconstrained real-world images examine the generalizability of the systems. On YFCC4k, CPlaNet produces the best street-level accuracy among baselines, while ISNs beats CPlaNet at the continent level. Our proposed TransLocator outperforms both CPlaNet and ISNs in every distance scale, improving the state-of-the-art by 1.7% and 5.5% street-level and continent-level accuracy. The YFCC26k dataset is even more challenging than YFCC4k; the best baseline produces only 5.3% street-level accuracy. However, our proposed TransLocator

Table 3: **Effect of number of scenes on TransLocator.** Fine-grained scene information helps TransLocator to achieve superior performance.

Dataset	Method	#Scenes	Distance ( $a_r$ [%] @ km)				
			Street 1 km	City 25 km	Region 200 km	Country 750 km	Continent 2500 km
Im2GPS [17]	TransLocator	3	18.4	46.3	55.6	68.6	84.0
		16	19.0	47.1	56.5	69.7	85.4
		365	<b>19.9</b>	<b>48.1</b>	<b>57.4</b>	<b>70.9</b>	<b>86.5</b>
	$\Delta_{\text{Scenes}_{365} - \text{Scenes}_3}$	1.5 ↑	1.8 ↑	1.8 ↑	2.3 ↑	2.5 ↑	
Im2GPS 3k [18]	TransLocator	3	10.8	29.9	41.0	56.8	78.7
		16	11.6	30.5	42.1	57.6	79.4
		365	<b>11.8</b>	<b>31.1</b>	<b>42.6</b>	<b>58.9</b>	<b>80.1</b>
	$\Delta_{\text{Scenes}_{365} - \text{Scenes}_3}$	1.0 ↑	1.2 ↑	1.6 ↑	2.1 ↑	1.4 ↑	

system yields an impressive 7.2% street-level accuracy on YFCC26k, which proves the appreciable generalizability of TransLocator.

## 5.2 Ablation Study

**Role of Vision Transformer:** We conduct a detailed ablation study to understand the contributions of different components proposed in TransLocator architecture. We start by comparing a base ViT-B/16 encoder with two conv networks - ResNet101 [19] and EfficientNet-B4 [50]. We re-train these three systems using only RGB images of the MP-16 dataset. As shown in Table 2, the base ViT architecture produces consistent improvements over the conv models on Im2GPS and Im2GPS3k, which confirms the effectiveness of the larger receptive field of ViT for geo-localization.

**Role of Segmentation Maps:** We then add the segmentation branch to base-ViT by concatenating the CLS tokens from the last layers. This method yields an improvement of 0.7% over base-ViT on Im2GPS. However, this kind of fusion is not optimal because the two channels do not interact in-between. Hence, we then incorporate our proposed multimodal feature fusion (MFF), which sums CLS tokens after each transformer layer. MFF improves on the base-ViT by 2.1% and 1.2% street-level accuracy on Im2GPS and Im2GPS3k, suggesting that the two branches learn complementary and robust features for different images.

**Role of Multi-task Learning:** Next, we incorporate an additional classifier head for scene recognition, which significantly improves the performance. Moreover, we evaluate the effect of coarse- and fine-grained contextual knowledge by varying the number of scene categories. As shown in Table 3, a higher number of scenes improves street-level accuracy by 1 – 1.5% across the two datasets.

## 5.3 Interpretability of TransLocator

In this section, we comprehend the interpretability of TransLocator by generating visual explanations using Grad-CAM [45]. First, we focus on similar-looking images coming from different portions of the world. As shown in Figure 4, these images<sup>7</sup> can be discriminated only by close attention to the global context.

<sup>7</sup> Collected from the Internet under creative commons license.

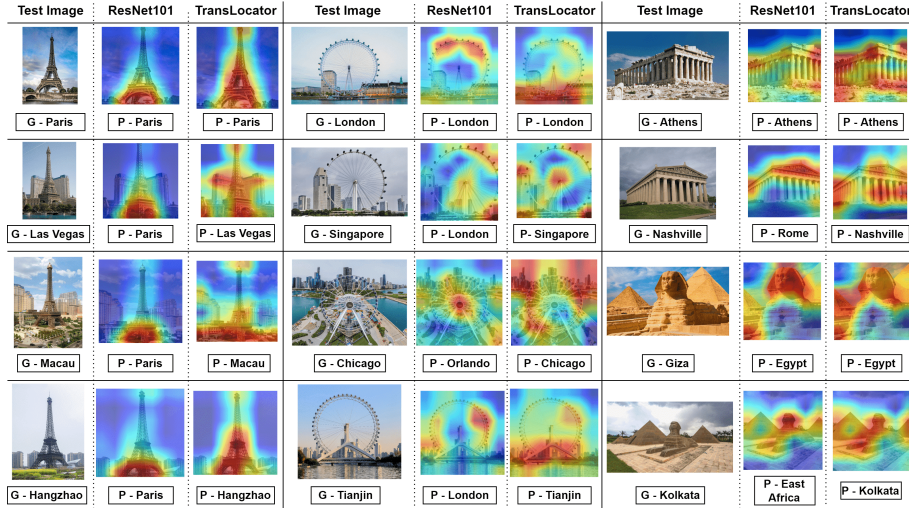


Fig. 4: **Qualitative comparison of TransLocator and ResNet101 on images with similar landmarks but from different geographic locations.** Unlike ResNet101, TransLocator focuses both on the foreground and background and is able to correctly geo-locate the very similar looking images.  $G$  and  $P$  denotes ground truth and predicted location.

For example, the famous Eiffel Tower in Paris closely resembles the towers in Las Vegas, Macau, and Hangzhou. However, there are differences in the background pixels between these images, such as different characteristic buildings and vegetation. While the ResNet101 architecture fails to locate the image from Las Vegas, our proposed TransLocator network correctly discriminates it from the Eiffel Tower in Paris by adequately focusing on the background structure. The similar superior performance of TransLocator is also shown for other images containing Ferris Wheel, Parthenon-like, and pyramid-like constructions in different geographic locations.

Next, we investigate the case of drastic appearance variation in the same location. Figure 5 shows three images<sup>8</sup> taken in Richmond, Virginia by a highway surveillance camera in the morning, dusk, and

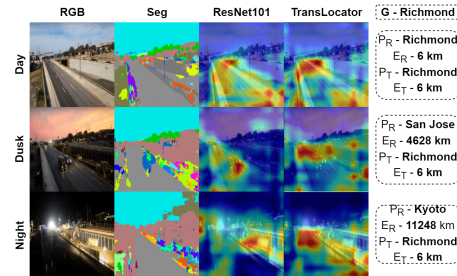


Fig. 5: **Qualitative comparison of TransLocator and ResNet101 on images with same location but under challenging appearance variations.** Unlike ResNet101, TransLocator attends to similar regions in each image and locates all three images correctly.  $G$  denotes ground truth,  $P_R$ ,  $E_R$ ,  $P_T$  and  $E_T$  denotes predicted location and prediction error by ResNet101 and TransLocator, respectively.

<sup>8</sup> Collected from <https://www.virginiadot.org> under creative commons license.

night. Though the RGB image varies with the change in the daytime, the corresponding semantic segmentation maps remain similar, and thus TransLocator can learn robust multimodal features. In contrast to ResNet101, which suffers from such appearance variation, TransLocator attends to similar regions in each image and locates all three images correctly.

#### 5.4 Error Analysis

Although TransLocator achieves better quantitative and qualitative results than baselines, there are still some open geo-localization problems that TransLocator does not solve. TransLocator can not locate images without geo-locating cues. For example, a photo of a cherry blossom tree can come anywhere from Tokyo, Paris, Washington DC, Dublin, etc. Hence TransLocator can not locate such an image. Similarly, pictures of sea beaches, deserts, or pictures of foods with no background cues can never precisely be geo-located. A few samples of such incorrectly located images are illustrated in Figure 6.

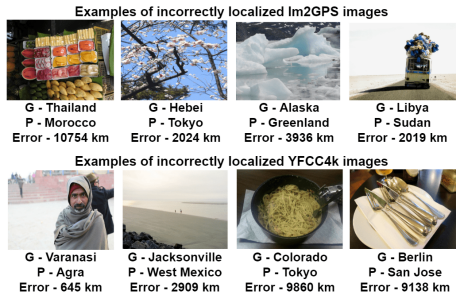


Fig. 6: **Limitations of our method:** images without salient geo-locating cues cannot be geo-located correctly. *G* and *P* denotes ground truth and predicted location.

## 6 Conclusion

Planet-scale single-image geo-localization is a highly challenging problem. These challenges include images with a large diversity in various environmental scenarios and appearance variation due to daytime, season, or weather changes. Hence, most existing approaches limit geo-localization in the scale of landmarks, a specific area, or an environmental scenario. Some approaches propose to use separate systems for different environments. In this paper, we address this challenging problem by proposing TransLocator, a unified dual-branch transformer network that attends to tiny details over the entire image and produces robust feature representation under extreme appearance variations. TransLocator takes an RGB image with its semantic segmentation map as input, interacts between its two parallel channels after each transformer layer, and concatenates the learned RGB and semantic representations using global attention. We train TransLocator in a unified multi-task framework for simultaneous geo-localization and scene recognition, and thus, our system can be applied to images from all environmental settings. Extensive experiments with TransLocator on four benchmark datasets - Im2GPS [17], Im2GPS3k [18], YFCC4k [59] and YFCC26k [51] shows a significant improvement of 5.5%, 14.1%, 4.9%, 9.9% continent-level accuracy over current state-of-the-art. We also obtain better qualitative results when we test TransLocator on challenging real-world images.

**Acknowledgement:** This research is partially supported by an ARO MURI Grant No. W911NF-17-1-0304.

## References

1. Akbari, H., Yuan, L., Qian, R., Chuang, W.H., Chang, S.F., Cui, Y., Gong, B.: Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text. *Advances in Neural Information Processing Systems* **34**, 24206–24221 (2021)
2. Baatz, G., Saurer, O., Köser, K., Pollefeys, M.: Large scale visual geo-localization of images in mountainous terrain. In: *European Conference on Computer Vision*. pp. 517–530. Springer (2012)
3. Berton, G., Masone, C., Caputo, B.: Rethinking visual geo-localization for large-scale applications. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4878–4888 (2022)
4. Brejcha, J., Čadík, M.: State-of-the-art in visual geo-localization. *Pattern Analysis and Applications* **20**(3), 613–637 (2017)
5. Cao, L., Smith, J.R., Wen, Z., Yin, Z., Jin, X., Han, J.: Bluefinder: Estimate where a beach photo was taken. In: *Proceedings of the 21st International Conference on World Wide Web*. pp. 469–470 (2012)
6. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: *European Conference on Computer Vision*. pp. 213–229. Springer (2020)
7. Caruana, R.: Multitask learning. *Machine learning* **28**(1), 41–75 (1997)
8. Chen, D.M., Baatz, G., Köser, K., Tsai, S.S., Vedantham, R., Pylvänäinen, T., Roimela, K., Chen, X., Bach, J., Pollefeys, M., et al.: City-scale landmark identification on mobile devices. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 737–744 (2011)
9. Chen, H., Wang, Y., Guo, T., Xu, C., Deng, Y., Liu, Z., Ma, S., Xu, C., Xu, C., Gao, W.: Pre-trained image processing transformer. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12299–12310 (2021)
10. Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention mask transformer for universal image segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1290–1299 (2022)
11. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. pp. 4171–4186. Association for Computational Linguistics, Minneapolis, Minnesota (Jun 2019). <https://doi.org/10.18653/v1/N19-1423>, <https://aclanthology.org/N19-1423>
12. Doersch, C., Gupta, A., Zisserman, A.: Crosstransformers: spatially-aware few-shot transfer. *Advances in Neural Information Processing Systems* **33**, 21981–21993 (2020)
13. 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. In: *International Conference on Learning Representations* (2020)
14. Fang, Y., Liao, B., Wang, X., Fang, J., Qi, J., Wu, R., Niu, J., Liu, W.: You only look at one sequence: Rethinking transformer in vision through object detection. *Advances in Neural Information Processing Systems* **34** (2021)
15. Gu, Y., Yang, K., Fu, S., Chen, S., Li, X., Marsic, I.: Hybrid attention based multimodal network for spoken language classification. In: *ACL*. vol. 2018, p. 2379 (2018)
16. Hausler, S., Garg, S., Xu, M., Milford, M., Fischer, T.: Patch-netvlad: Multi-scale fusion of locally-global descriptors for place recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14141–14152 (2021)



17. Hays, J., Efros, A.A.: Im2gps: estimating geographic information from a single image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1–8. IEEE (2008)
18. Hays, J., Efros, A.A.: Large-scale image geolocation. In: Multimodal Location Estimation of Videos and Images, pp. 41–62. Springer (2015)
19. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)
20. Hu, R., Singh, A.: Unit: Multimodal multitask learning with a unified transformer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1439–1449 (2021)
21. 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/CVF Conference on Computer Vision and Pattern Recognition. pp. 7258–7267 (2018)
22. Izbicki, M., Papalexakis, E.E., Tsotras, V.J.: Exploiting the earth’s spherical geometry to geolocate images. In: Joint European Conference on Machine Learning and Knowledge Discovery in Databases. pp. 3–19. Springer (2019)
23. Kant, Y., Batra, D., Anderson, P., Schwing, A., Parikh, D., Lu, J., Agrawal, H.: Spatially aware multimodal transformers for textvqa. In: European Conference on Computer Vision. pp. 715–732. Springer (2020)
24. Larson, M., Soleymani, M., Gravier, G., Ionescu, B., Jones, G.J.: The benchmarking initiative for multimedia evaluation: Mediaeval 2016. *IEEE MultiMedia* **24**(1), 93–96 (2017)
25. Lei, J., Li, L., Zhou, L., Gan, Z., Berg, T.L., Bansal, M., Liu, J.: Less is more: Clipbert for video-and-language learning via sparse sampling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7331–7341 (2021)
26. Li, L.H., Zhang, P., Zhang, H., Yang, J., Li, C., Zhong, Y., Wang, L., Yuan, L., Zhang, L., Hwang, J.N., et al.: Grounded language-image pre-training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10965–10975 (2022)
27. Li, X., Yin, X., Li, C., Zhang, P., Hu, X., Zhang, L., Wang, L., Hu, H., Dong, L., Wei, F., et al.: Oscar: Object-semantics aligned pre-training for vision-language tasks. In: European Conference on Computer Vision. pp. 121–137. Springer (2020)
28. Lin, T.Y., Belongie, S., Hays, J.: Cross-view image geolocation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 891–898 (2013)
29. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019)
30. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representations (2018)
31. Lu, J., Batra, D., Parikh, D., Lee, S.: Vlbnet: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in Neural Information Processing Systems* **32** (2019)
32. Masone, C., Caputo, B.: A survey on deep visual place recognition. *IEEE Access* **9**, 19516–19547 (2021)
33. Misra, I., Girdhar, R., Joulin, A.: An end-to-end transformer model for 3d object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2906–2917 (2021)
34. Muller-Budack, E., Pustu-Iren, K., Ewerth, R.: Geolocation estimation of photos using a hierarchical model and scene classification. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 563–579 (2018)

35. 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, high-performance deep learning library. *Advances in Neural Information Processing Systems* **32** (2019)
36. Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., Zettlemoyer, L.: Deep contextualized word representations. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. pp. 2227–2237 (2018)
37. Pramanick, S., Roy, A., Patel, V.M.: Multimodal learning using optimal transport for sarcasm and humor detection. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 3930–3940 (2022)
38. Pramanick, S., Sharma, S., Dimitrov, D., Akhtar, M.S., Nakov, P., Chakraborty, T.: Momenta: A multimodal framework for detecting harmful memes and their targets. In: *Findings of the Association for Computational Linguistics: EMNLP 2021*. pp. 4439–4455 (2021)
39. Raghu, M., Unterthiner, T., Kornblith, S., Zhang, C., Dosovitskiy, A.: Do vision transformers see like convolutional neural networks? *Advances in Neural Information Processing Systems* **34** (2021)
40. Ranjan, R., Patel, V.M., Chellappa, R.: Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **41**(1), 121–135 (2017)
41. Regmi, K., Shah, M.: Bridging the domain gap for ground-to-aerial image matching. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 470–479 (2019)
42. Ridnik, T., Ben-Baruch, E., Noy, A., Zelnik-Manor, L.: Imagenet-21k pretraining for the masses. *arXiv preprint arXiv:2104.10972* (2021)
43. Ruder, S.: An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098* (2017)
44. Saurer, O., Baatz, G., Köser, K., Pollefeys, M., et al.: Image based geo-localization in the alps. *International Journal of Computer Vision* **116**(3), 213–225 (2016)
45. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 618–626 (2017)
46. Seo, P.H., Weyand, T., Sim, J., Han, B.: Cplanet: Enhancing image geolocation by combinatorial partitioning of maps. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 536–551 (2018)
47. Seymour, Z., Sikka, K., Chiu, H.P., Samarasekera, S., Kumar, R.: Semantically-aware attentive neural embeddings for image-based visual localization. *arXiv preprint arXiv:1812.03402* (2018)
48. Strudel, R., Garcia, R., Laptev, I., Schmid, C.: Segmenter: Transformer for semantic segmentation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 7262–7272 (2021)
49. Sun, C., Myers, A., Vondrick, C., Murphy, K., Schmid, C.: Videobert: A joint model for video and language representation learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 7464–7473 (2019)
50. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: *International Conference on Machine Learning*. pp. 6105–6114. PMLR (2019)
51. Theiner, J., Müller-Budack, E., Ewerth, R.: Interpretable semantic photo geolocation. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 750–760 (2022)
52. Thomee, B., Shamma, D.A., Friedland, G., Elizalde, B., Ni, K., Poland, D., Borth, D., Li, L.J.: Yfcc100m: The new data in multimedia research. *Communications of the ACM* **59**(2), 64–73 (2016)

53. Tian, Y., Chen, C., Shah, M.: Cross-view image matching for geo-localization in urban environments. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3608–3616 (2017)
54. Toker, A., Zhou, Q., Maximov, M., Leal-Taixé, L.: Coming down to earth: Satellite-to-street view synthesis for geo-localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6488–6497 (2021)
55. Torii, A., Arandjelovic, R., Sivic, J., Okutomi, M., Pajdla, T.: 24/7 place recognition by view synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1808–1817 (2015)
56. Torii, A., Sivic, J., Pajdla, T., Okutomi, M.: Visual place recognition with repetitive structures. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 883–890 (2013)
57. Tzeng, E., Zhai, A., Clements, M., Townshend, R., Zakhori, A.: User-driven geolocation of untagged desert imagery using digital elevation models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 237–244 (2013)
58. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems. pp. 5998–6008 (2017)
59. Vo, N., Jacobs, N., Hays, J.: Revisiting im2gps in the deep learning era. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2621–2630 (2017)
60. Wang, J., Sun, K., Cheng, T., Jiang, B., Deng, C., Zhao, Y., Liu, D., Mu, Y., Tan, M., Wang, X., et al.: Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **43**(10), 3349–3364 (2020)
61. Wang, Y., Cao, L.: Discovering latent clusters from geotagged beach images. In: Advances in Multimedia Modeling, pp. 133–142. Springer (2013)
62. Wang, Y., Xu, Z., Wang, X., Shen, C., Cheng, B., Shen, H., Xia, H.: End-to-end video instance segmentation with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8741–8750 (2021)
63. Weyand, T., Kostrikov, I., Philbin, J.: Planet-photo geolocation with convolutional neural networks. In: European Conference on Computer Vision. pp. 37–55. Springer (2016)
64. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems* **34** (2021)
65. Xu, H., Ghosh, G., Huang, P.Y., Okhonko, D., Aghajanyan, A., Metzger, F., Zettlemoyer, L., Feichtenhofer, C.: Videoclip: Contrastive pre-training for zero-shot video-text understanding. In: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing. pp. 6787–6800 (2021)
66. Yang, F., Yang, H., Fu, J., Lu, H., Guo, B.: Learning texture transformer network for image super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5791–5800 (2020)
67. 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)
68. Yang, J., Li, C., Zhang, P., Xiao, B., Liu, C., Yuan, L., Gao, J.: Unified contrastive learning in image-text-label space. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19163–19173 (2022)
69. Yang, L., Fan, Y., Xu, N.: Video instance segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5188–5197 (2019)
70. Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R.R., Le, Q.V.: Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in Neural Information Processing Systems* **32** (2019)

71. Ye, H.J., Hu, H., Zhan, D.C., Sha, F.: Few-shot learning via embedding adaptation with set-to-set functions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8808–8817 (2020)
72. Zhang, Y., Yang, Q.: A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering* (2021)
73. Zhao, H., Jiang, L., Jia, J., Torr, P.H., Koltun, V.: Point transformer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16259–16268 (2021)
74. Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., Fu, Y., Feng, J., Xiang, T., Torr, P.H., et al.: Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6881–6890 (2021)
75. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **40**(6), 1452–1464 (2017)
76. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 633–641 (2017)
77. Zhou, B., Zhao, H., Puig, X., Xiao, T., Fidler, S., Barriuso, A., Torralba, A.: Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision* **127**(3), 302–321 (2019)
78. 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)
79. Zhu, S., Yang, T., Chen, C.: Vigor: Cross-view image geo-localization beyond one-to-one retrieval. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3640–3649 (June 2021)
80. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. In: International Conference on Learning Representations (2020)