

CPlaNet: Enhancing Image Geolocalization by Combinatorial Partitioning of Maps

Paul Hongsuck Seo¹, Tobias Weyand², Jack Sim², and Bohyung Han³

¹ Dept. of CSE, POSTECH, Korea

² Google Research, USA

³ Dept. of ECE & ASRI, Seoul National University, Korea

hsseo@postech.ac.kr {weyand,jacksim}@google.com bhhan@snu.ac.kr

Abstract. Image geolocalization is the task of identifying the location depicted in a photo based only on its visual information. This task is inherently challenging since many photos have only few, possibly ambiguous cues to their geolocation. Recent work has cast this task as a classification problem by partitioning the earth into a set of discrete cells that correspond to geographic regions. The granularity of this partitioning presents a critical trade-off; using fewer but larger cells results in lower location accuracy while using more but smaller cells reduces the number of training examples per class and increases model size, making the model prone to overfitting. To tackle this issue, we propose a simple but effective algorithm, combinatorial partitioning, which generates a large number of fine-grained output classes by intersecting multiple coarse-grained partitionings of the earth. Each classifier votes for the fine-grained classes that overlap with their respective coarse-grained ones. This technique allows us to predict locations at a fine scale while maintaining sufficient training examples per class. Our algorithm achieves the state-of-the-art performance in location recognition on multiple benchmark datasets.

Keywords: Image geolocalization, combinatorial partitioning, fine-grained classification

1 Introduction

Image geolocalization is the task of predicting the geographic location of an image based only on its pixels without any meta-information. As the geolocation is an important attribute of an image by itself, it also plays as a proxy to other location attributes such as elevation, weather, and distance to a particular point of interest. However, geolocating images is a challenging task since input images often contain limited visual information representative of their locations. To handle this issue effectively, the model is required to capture and maintain visual cues of the globe comprehensively.

There exist two main streams to address this task: retrieval and classification based approaches. The former searches for nearest neighbors in a database of geotagged images by matching their feature representations [1–3]. Visual appearance of an image at a certain geolocation is estimated using the representations




Fig. 1: Visualization of combinatorial partitioning. Two coarse-grained class sets, $\mathcal{P} = \{p_1, p_2, \dots, p_5\}$ and $\mathcal{Q} = \{q_1, q_2, \dots, q_5\}$ in the map on the left, are merged to construct a fine-grained partition as shown in the map on the right by a combination of geoclasses in the two class sets. Each resulting fine-grained class is represented by a tuple (p_i, q_j) , and is constructed by identifying partially overlapping partitions in \mathcal{P} and \mathcal{Q} .

of the geotagged images in database. The latter treats the task as a classification problem by dividing the map into multiple discrete classes [3, 4]. Thanks to recent advances in deep learning, simple classification techniques based on convolutional neural networks handle such complex visual understanding problems effectively.

There are several advantages of formulating the task as classification instead of retrieval. First, classification-based approaches save memory and disk space to store information for geolocalization; they just need to store a set of model parameters learned from training images whereas all geotagged images in the database should be embedded and indexed to build retrieval-based systems. In addition to space complexity, inference of classification-based approaches is faster because a result is given by a simple forward pass computation of a deep neural network while retrieval-based methods undergo significant overhead for online search from a large index given a query image. Finally, classification-based algorithms provide multiple hypotheses of geolocation with no additional cost by presenting multi-modal answer distributions.

On the other hand, the standard classification-based approaches have a few critical limitations. They typically ignore correlation of spatially adjacent or proximate classes. For instance, assigning a photo of Bronx to Queens, which are both within New York city, is treated equally wrong as assigning it to Seoul. Another drawback comes from artificially converting continuous geographic space into discrete class representations. Such an attempt may incur various artifacts since images near class boundaries are not discriminative enough compared to data variations within classes; training converges slowly and performance is affected substantially by subtle changes in map partitioning. This limitation can be alleviated by increasing the number of classes and reducing the area of the region corresponding to each class. However, this strategy increases the number of parameters while decreasing the size of the training dataset per class.

To overcome such limitations, we propose a novel algorithm that enhances the resolution of geoclasses and avoids the training data deficiency issue. This is achieved by *combinatorial partitioning*, which is a simple technique to generate spatially fine-grained classes through a combination of the multiple configurations of classes. This idea has analogy to product quantization [5] since they both construct a lot of quantized regions using relatively few model parameters through a combination of low-bit subspace encodings or coarse spatial quantizations. Our combinatorial partitioning allows the model to be trained with more data per class by considering a relatively small number of classes at a time. Figure 1 illustrates an example of combinatorial partitioning, which enables generating more classes with minimal increase of model size and learning individual classifiers reliably without losing training data per class. Combinatorial partitioning is applied to an existing classification-based image geolocalization technique, PlaNet [4], and our algorithm is referred to as CPlaNet hereafter. Our contribution is threefold:

- We introduce a novel classification-based model for image geolocalization using combinatorial partitioning, which defines a fine-grained class configuration by combining multiple heterogeneous geoclass sets in coarse levels.
- We propose a technique that generates multiple geoclass sets by varying parameters, and design an efficient inference technique to combine prediction results from multiple classifiers with proper normalization.
- The proposed algorithm outperforms the existing techniques in multiple benchmark datasets, especially at fine scales.

The rest of this paper is organized as follows. We review the related work in Section 2, and describe combinatorial partitioning for image geolocalization in Section 3. The details about training and inference procedures are discussed in Section 4. We present experimental results of our algorithm in Section 5, and conclude our work in Section 6.

2 Related Work

The most common approach of image geolocalization is based on the image retrieval pipeline. Im2GPS [1, 2] and its derivative [3] perform image retrieval in a database of geotagged images using global image descriptors. Various visual features can be applied to the image retrieval step. NetVLAD [6] is a global image descriptor trained end-to-end for place recognition on street view data using a ranking loss. Kim et al. [7] learn a weighting mask for the NetVLAD descriptor to focus on image regions containing location cues. While global features have the benefit to retrieve diverse natural scene images based on ambient information, local image features yield higher precision in retrieving structured objects such as buildings and are thus more frequently used [8–16]. DELF [17] is a deeply learned local image feature detector and descriptor with attention for image retrieval.

On the other hand, classification-based image geolocalization formulates the problem as a classification task. In [3, 4], a classifier is trained to predict the geolocation of an input image. Since the geolocation is represented in a continuous

space, classification-based approaches quantize the map of the entire earth into a set of geoclasses corresponding to partitioned regions. Note that training images are labeled into the corresponding geoclasses based on their GPS tags. At test time, the center of the geoclass with the highest score is returned as the predicted geolocation of an input image. This method is lightweight in terms of space and time complexity compared to retrieval-based methods, but its prediction accuracy highly depends on how the geoclass set is generated. Since every image that belongs to the same geoclass has an identical predicted geolocation, more fine-grained partitioning is preferable to obtain precise predictions. However, it is not always straightforward to increase the number of geoclasses as it linearly increases the number of parameters and makes the network prone to overfitting to training data.

Pose estimation approaches [9, 18–23] match query images against 3D models of an area, and employ 2D-3D feature correspondences to identify 6-DOF query poses. Instead of directly matching against a 3D model, [24, 23] first perform image retrieval to obtain coarse locations and then estimate poses using the retrieved images. PoseNet [25, 26] treats pose estimation as a regression problem based on a convolutional neural network. The accuracy of PoseNet is improved by introducing an intermediate LSTM layer for dimensionality reduction [27].

A related line of research is landmark recognition, where images are clustered by their geolocations and visual similarity to construct a database of popular landmarks. The database serves as the index of an image retrieval system [28–33] or the training data of a landmark classifier [34–36]. Cross-view geolocation recognition makes additional use of satellite or aerial imagery to determine query locations [37–40].

3 Geolocalization using Multiple Classifiers

Unlike existing classification-based methods [4], CPlaNet relies on multiple classifiers that are all trained with unique geoclass sets. The proposed model predicts more fine-grained geoclasses, which are given by combinatorial partitioning of multiple geoclass sets. Since our method requires a distinct geoclass set for each classifier, we also propose a way to generate multiple geoclass sets.

3.1 Combinatorial Partitioning

Our primary goal is to establish fine-grained geoclasses through a combination of multiple coarse geoclass sets and exploit benefits from both coarse- and fine-grained geolocalization-by-classification approaches. In our model, there are multiple unique geoclass sets represented by partitions $\mathcal{P} = \{p_1, p_2, \dots, p_5\}$ and $\mathcal{Q} = \{q_1, q_2, \dots, q_5\}$ as illustrated on the left side of Figure 1. Since the region boundaries in these geoclass sets are unique, overlapping the two maps constructs a set of fine-grained subregions. This procedure, referred to as combinatorial partitioning, is identical to the Cartesian product of the two sets, but disregards the tuples given by two spatially disjoint regions in the map. For instance,

combining two aforementioned geoclass sets in Figure 1, we obtain fine-grained partitions defined by a tuple (p_i, q_j) as depicted by the map on the right of the figure, while the tuples made by two disjoint regions, *e.g.*, (p_1, q_5) , are not considered.

While combinatorial partitioning aggregates results from multiple classifiers, it is conceptually different from ensemble models whose base classifiers predict labels in the same output space. In combinatorial partitioning, while each coarse-grained partition is scored by a corresponding classifier, fine-grained partitions are given different scores by the combinations of multiple unique geoclass sets. Also, combinatorial partitioning is closely related to product quantization [5] for approximate nearest neighbor search in the sense that they both generate a large number of quantized regions by either a Cartesian product of quantized subspaces or a combination of coarse space quantizations. Note that combinatorial partitioning is a general framework and applicable to other tasks, especially where labels have to be defined on the same embedded space as in geographical maps.

3.2 Benefits of Combinatorial Partitioning

The proposed classification model with combinatorial partitioning has the following three major benefits.

Fine-grained classes with fewer parameters Combinatorial partitioning generates fine-grained geoclasses using a smaller number of parameters because a single geoclass in a class set can be divided into many subregions by intersections with geoclasses from multiple geoclass sets. For instance in Figure 1, two sets with 5 geoclasses form 14 distinct classes by the combinatorial partitioning. If we design a single flat classifier with respect to the fine-grained classes, it requires more parameters, *i.e.*, $14 \times F > 2 \times (5 \times F)$, where F is the number of input dimensions to the classification layers.

More training data per class Training with fine-grained geoclass sets is more desirable for higher resolution of output space, but is not straightforward due to training data deficiency; the more we divide the maps, the less training images remain per geoclass. Our combinatorial partitioning technique enables us to learn models with coarsely divided geoclasses and maintain more training data in each class than a naïve classifier with the same number of classes.

More reasonable class sets There is no standard method to define geoclasses for image geolocalization, so that images associated with the same classes have common characteristics. An arbitrary choice of partitioning may incur undesirable artifacts due to heterogeneous images located near class territories; the features trained on loosely defined class sets tend to be insufficiently discriminative and less representative. On the other hand, our framework constructs diverse partitions based on various criteria observed in the images. We can define more

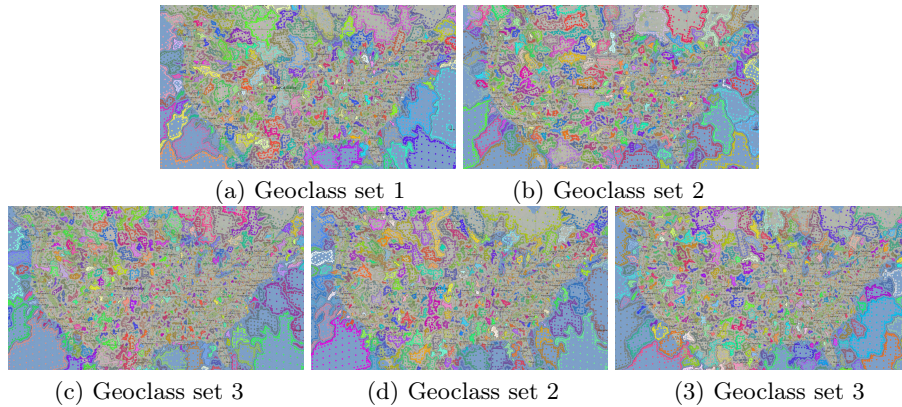


Fig. 2: Visualization of the geoclass sets on the maps of the United States generated by the parameters shown in Table 1. Each distinct region is marked by a different color. The first two sets, (a) and (b), are generated by manually designed parameters while parameters for the others are randomly sampled.

tightly-coupled classes through combinatorial partitioning by distilling noisy information from multiple sources.

3.3 Generating Multiple Geoclass Sets

The geoclass set organization is an inherently ill-posed problem as there is no consensus about ideal region boundaries for image geolocalization. Consequently, it is hard to define the optimal class configuration, which motivates the use of multiple random boundaries in our combinatorial partitioning. We therefore introduce a mutable method of generating geoclass sets, which considers both visual and geographic distances between images.

The generation method starts with an initial graph for a map, where a node represents a region in the map and an edge connects two nodes of adjacent regions. We construct the initial graph based on S2 cells⁴ at a certain level. Empty S2 cells, which contain no training image, do not construct separate nodes and are randomly merged with one of their neighboring non-empty S2 cells. This initial graph covers the entire surface of the earth. Both nodes and edges are associated with numbers—scores for nodes and weights for edges. We give a score to each node by a linear combination of three different factors: the number of images in the node and the number of empty and non-empty S2 cells. An edge weight is computed by the weighted sum of geolocational and visual distances between two nodes. The geolocational distance is given by the distance between the centers of

⁴ We use Google’s S2 library. S2 cells are given by a geographical partitioning of the earth into a hierarchy. The surface of the earth is projected onto six faces of a cube. Each face of the cube is hierarchically subdivided and forms S2 cells in a quad-tree. Refer to <https://code.google.com/archive/p/s2-geometry-library/> for more details.

Table 1: Parameters for geoclass set generation. Parameters for geoclass set 1 and 2 are manually given while the ones for rest geoclass sets are randomly sampled.

Parameter group	Parameters	1	2	3	4	5
N/A	Num. of geoclasses	9,969	9,969	12,977	12,333	11,262
	Image feature dimensions	2,048	0	1,187	1,113	14,98
Node score	Weight for num. of images (α_1)	1.000	1.000	0.501	0.953	0.713
	Weight for num. of non-empty S2 cells (α_2)	0.000	0.000	0.490	0.044	0.287
	Weight for num. of S2 cells (α_3)	0.000	0.000	0.009	0.003	0.000
Edge weight	Weight for visual distance (β_1)	1.000	0.000	0.421	0.628	0.057
	Weight for geographical distance (β_2)	0.000	1.000	0.579	0.372	0.943

two nodes while the visual distance is measured by cosine similarity based on the visual features of nodes, which are computed by averaging the associated image features extracted from the bottleneck layer of a pretrained CNN. Formally, a node score $\omega(\cdot)$ and an edge weight $\nu(\cdot, \cdot)$ are defined respectively as

$$\omega(v_i) = \alpha_1 \cdot n_{\text{img}}(v_i) + \alpha_2 \cdot n_{\text{S2+}}(v_i) + \alpha_3 \cdot n_{\text{S2}}(v_i) \quad (1)$$

$$\nu(v_i, v_j) = \beta_1 \cdot \text{dist}_{\text{vis}}(v_i, v_j) + \beta_2 \cdot \text{dist}_{\text{geo}}(v_i, v_j) \quad (2)$$

where $n_{\text{img}}(v)$, $n_{\text{S2+}}(v)$ and $n_{\text{S2}}(v)$ are functions that return the number of images, non-empty S2 cells and all S2 cells in a node v , respectively, and $\text{dist}_{\text{vis}}(\cdot, \cdot)$ and $\text{dist}_{\text{geo}}(\cdot, \cdot)$ are the visual geolocational distances between two nodes. Note that the weights $(\alpha_1, \alpha_2, \alpha_3)$ and (β_1, β_2) are free parameters in $[0, 1]$.

After constructing the initial graph, we merge two nodes hierarchically in a greedy manner until the number of remaining nodes becomes the desired number of geoclasses. To make each geoclass roughly balanced, we select the node with the lowest score first and merge it with its nearest neighbor in terms of edge weight. A new node is created by the merge process and corresponds to the region given by the union of two merged regions. The score of the new node is set to the sum of the scores of the two merged nodes.

The generated geoclass sets are diversified by the following free parameters: 1) the desired number of final geoclasses, 2) the weights of the factors in the node scores, 3) the weights of the two distances in computing edge weights and 4) the image feature extractor. Each parameter setting constructs a unique geoclass set. Note that multiple geoclass set generation is motivated by the fact that geoclasses are often ill-defined and the perturbation of class boundaries is a natural way to address the ill-posed problem. Figure 2 illustrates generated geoclass sets using different parameters described in Table 1.

4 Learning and Inference

This section describes more details about CPlaNet including network architecture, and training and testing procedure. We also discuss data structures and the detailed inference algorithm.




Fig. 3: Network architecture of our model. A single Inception v3 architecture is used as our feature extractor after removing the final classification layer. An image feature is fed to multiple classification branches and classification scores are predicted over multiple geoclass sets.

4.1 Network Architecture

Following [4], we construct our network based on the Inception architecture [41] with batch normalization [42]. Inception v3 without the final classification layer (fc with `softmax`) is used as our feature extractor, and multiple branches of classification layers are attached on top of the feature extractor as illustrated in Figure 3. We train the multiple classifiers independently while keeping the weights of the Inception module fixed. Note that, since all classifiers share the feature extractor, our model requires marginal increase of memory to maintain multiple classifiers.

4.2 Inference with Multiple Classifiers

Once the predicted scores in each class set are assigned to the corresponding regions, the subregions overlapped by multiple class sets are given cumulated scores from multiple classifiers. A simple strategy to accumulate geoclass scores is to add the scores to individual S2 cells within the geoclass. Such a simple strategy is inappropriate since it gives favor to classifiers that have geoclasses corresponding to large regions covering more S2 cells. To make each classifier contribute equally to the final prediction regardless of its class configuration, we normalize the scores from individual classifiers with consideration of the number of S2 cells per class before adding them to the current S2 cell scores. Formally, given a geoclass score distributed to S2 cell g_k within a class in a geoclass set \mathcal{C}^i , denoted by $\text{geoscore}(g_k; \mathcal{C}^i)$, an S2 cell is given a score $s(\cdot)$ by

$$s(g_k) = \sum_{i=1}^N \frac{\text{geoscore}(g_k; \mathcal{C}^i)}{\sum_{t=1}^K \text{geoscore}(g_t; \mathcal{C}^i)}, \quad (3)$$

where K is the total number of S2 cells and N is the number of geoclass sets. Note that this process implicitly creates fine-grained partitions because the regions defined by different geoclass combinations are given different scores.

After this procedure, we select the S2 cells with the highest scores and compute their center for the final prediction of geolocation by averaging locations of images in the S2 cells. That is, the predicted geolocation l_{pred} is given by

$$l_{\text{pred}} = \frac{\sum_{k \in \mathcal{G}} \sum_{e \in g_k} \text{geolocation}(e)}{\sum_{k \in \mathcal{G}} |g_k|}, \quad (4)$$

where $\mathcal{G} = \text{argmax}_k s(g_k)$ is an index set of the S2 cells with the highest scores and $\text{geolocation}(\cdot)$ is a function to return the ground-truth GPS coordinates of a training image e . Note that an S2 cell g_k may contain a number of training examples.

In our implementation, all fine-grained partitions are precomputed offline by generating all existing combinations of the multiple geoclass sets, and an index mapping from each geoclass to its corresponding partitions is also constructed offline to accelerate inference. Moreover, we precompute the center of images in each partition. To compute the center of a partition, we convert the latitude and longitude values of GPS tags into 3D Cartesian coordinates. This is because a naïve average of latitude and longitude representations introduces significant errors as the target locations become distant from the equator.

5 Experiments

5.1 Datasets

We train our network using a private dataset collected from Flickr, which has 30.3M geotagged images for training. We have sanitized the dataset by removing noisy examples to weed out unsuitable photos. For example, we disregard unnatural images (e.g., clipart images, product photos, etc.) and accept photos with a minimum size of 0.1 megapixels.

For evaluation, we mainly employ two public benchmark datasets—Im2GPS3k and YFCC4k [3]. The former contains 3,000 images from the Im2GPS dataset whereas the latter has 4,000 random images from the YFCC100m dataset. In addition, we also evaluate on Im2GPS test set [1] to compare with previous work. Note that Im2GPS3k is a different test benchmark from the Im2GPS test set.

5.2 Parameters and Training Networks

We generate three geoclass sets using randomly generated parameters, which are summarized in Table 1. The number of geoclasses for each set is approximately between 10K and 13K, and the generation parameters for edge weights and node scores are randomly sampled. Specifically, we select random axis-aligned subspaces out of the full 2,048 dimensions for image representations to diversify dissimilarity metrics between image representations. Note that the image representations are extracted by a reproduced PlaNet [4] after removing the final classification layer. In addition to these geoclass sets, we generate two more sets with manually designed parameters; the edge weights in these two cases are given by either

Table 2: Geolocational accuracies [%] of models at different scales on Im2GPS3k.

Models	1 km	5 km	10 km	25 km	50 km	100 km	200 km	750 km	2500 km
ImageNetFeat	3.0	5.5	6.4	6.9	7.7	9.0	10.8	18.5	37.5
Deep-Ret [3]	3.7	–	–	19.4	–	–	26.9	38.9	55.9
PlaNet (reprod) [4]	8.5	18.1	21.4	24.8	27.7	30.0	34.3	48.4	64.6
ClassSet 1	8.4	18.3	21.7	24.7	27.4	29.8	34.1	47.9	64.5
ClassSet 2	8.0	17.6	20.6	23.8	26.2	29.2	32.7	46.6	63.9
ClassSet 3	8.8	18.9	22.4	25.7	27.9	29.8	33.5	47.8	64.1
ClassSet 4	8.7	18.5	21.4	24.6	26.8	29.6	33.0	47.6	64.4
ClassSet 5	8.8	18.7	21.7	24.7	27.3	29.3	32.9	47.1	64.5
Average[1-2]	8.2	18.0	21.1	24.2	26.8	29.5	33.4	47.3	64.2
Average[1-5]	8.5	18.4	21.5	24.7	27.1	29.5	33.2	47.4	64.3
CPlaNet[1-2]	9.3	19.3	22.7	25.7	27.7	30.1	34.4	47.8	64.5
CPlaNet[1-5]	9.9	20.2	23.3	26.3	28.5	30.4	34.5	48.8	64.6
CPlaNet[1-5,PlaNet]	10.2	20.8	23.7	26.5	28.6	30.6	34.6	48.6	64.6

visual or geolocational distance exclusively, and their node scores are based on the number of images to mimic the setting of PlaNet. Figure 2 visualizes five geoclass sets generated by the parameters presented in Table 1.

We use S2 cells at level 14 to construct the initial graph, where a total of ~ 2.8 M nodes are obtained after merging empty cells to their non-empty neighbors. To train the proposed model, we employ the pretrained model of the reproduced PlaNet with its parameters fixed while the multiple classification branches are randomly initialized and fine-tuned using our training dataset. The network is trained by RMSprop with a learning rate of 0.005.

5.3 Evaluation Metrics

Following [3, 4], we evaluate the models using geolocational accuracies at multiple scales by varying the allowed errors in terms of distances from ground-truth locations as follows: 1 km, 5 km, 10 km, 25 km, 50 km, 100 km, 200 km, 750 km and 2500 km. Our evaluation focuses more on high accuracy range compared to the previous papers as we believe that fine-grained geolocalization is more important in practice. A geolocational accuracy a_r at a scale is given by the fraction of images in the test set localized within radius r from ground-truths, which is given by

$$a_r \equiv \frac{1}{M} \sum_{i=1}^M u[\text{geodist}(l_{\text{gt}}^i, l_{\text{pred}}^i) < r], \quad (5)$$

where M is the number of examples in the test set, $u[\cdot]$ is an indicator function and $\text{geodist}(l_{\text{gt}}^i, l_{\text{pred}}^i)$ is the geolocational distance between the true image location l_{gt}^i and the predicted location l_{pred}^i of the i -th example.

Table 3: Geolocational accuracies [%] on YFCC4k.

Models	1 km	5 km	10 km	25 km	50 km	100 km	200 km	750 km	2500 km
Deep-Ret [3]	2.3	-	-	5.7	-	-	11.0	23.5	42.0
PlaNet (reprod) [4]	5.6	10.1	12.2	14.3	16.6	18.7	22.2	36.4	55.8
CPlaNet[1-5]	7.3	11.7	13.1	14.7	16.1	18.2	21.7	36.2	55.6
CPlaNet[1-5,PlaNet]	7.9	12.1	13.5	14.8	16.3	18.5	21.9	36.4	55.5

5.4 Results

Benefits of Combinatorial Partitioning Table 2 presents the geolocational accuracies of the proposed model on the Im2GPS3k dataset. The proposed models outperform the baselines and the existing methods at almost all scales on this dataset. ClassSet 1 through 5 in Table 2 are the models trained with the geoclass sets generated from the parameters presented in Table 1. Using the learned models as the base classifiers, we construct two variants of the proposed method—CPlaNet[1-2] using the first two base classifiers with manual parameter selection and CPlaNet[1-5] using all the base classifiers.

Table 2 presents that both options of our models outperform all the underlying classifiers at every scale. Compared to naïve average of the underlying classifiers denoted by Average[1-5] and Average[1-2], CPlaNet[1-5] and CPlaNet[1-2] have $\sim 16\%$ and $\sim 13\%$ of accuracy gains at street level, respectively, compared to their counterparts. We emphasize that CPlaNet achieves substantial improvements by a simple combination of the existing base classifiers and a generation of fine-grained partitions without extra training procedure. The larger performance improvement in CPlaNet[1-5] compared to CPlaNet[1-2] makes sense as using more classifiers constructs more fine-grained geoclasses via combinatorial partitioning and increases prediction resolution. Note that the number of distinct partitions formed by CPlaNet[1-2] is 46,294 while it is 107,593 in CPlaNet[1-5].

The combinatorial partitioning of the proposed model is not limited to geoclass sets from our generation methods, but is generally applicable to any geoclass sets. Therefore, we construct an additional instance of the proposed method, CPlaNet[1-5,PlaNet], which also incorporates PlaNet (reprod), reproduced version of PlaNet model [4] with our training data, additionally. CPlaNet[1-5,PlaNet] shows extra performance gains over CPlaNet[1-5] and achieves the state-of-the-art performance at all scales. These experiments show that our combinatorial partitioning is a useful framework for image geolocalization through ensemble classification, where multiple classifiers with heterogeneous geoclass sets complement each other.

We also present results on YFCC4k [3] dataset in Table 3. The overall tendency is similar to the one in Im2GPS3k. Our full model outperforms Deep-Ret [3] consistently and significantly. The proposed algorithm also shows substantially better performance compared to PlaNet (reprod) in the low threshold range while two methods have almost identical accuracy at coarse-level evaluation.

On the Im2GPS dataset, our model outperforms other classification-based approaches—Deep-Cls and PlaNet, which are single-classifier models with a different geoclass schema—significantly at every scale, as shown in Table 4. The

Table 4: Geolocational accuracies [%] on Im2GPS.

Models		1 km	5 km	10 km	25 km	50 km	100 km	200 km	750 km	2500 km
Retrieval	Im2GPS [1]	-	-	-	12.0	-	-	15.0	23.0	47.0
	Im2GPS [2]	2.5	12.2	16.9	21.9	25.3	28.7	32.1	35.4	51.9
	Deep-Ret [3]	12.2	-	-	33.3	-	-	44.3	57.4	71.3
	Deep-Ret+ [3]	14.4	-	-	33.3	-	-	47.7	61.6	73.4
Classifier	Deep-Cls [3]	6.8	-	-	21.9	-	-	34.6	49.4	63.7
	PlaNet [4]	8.4	19.0	21.5	24.5	27.8	30.4	37.6	53.6	71.3
	PlaNet (reprod) [4]	11.0	23.6	26.6	31.2	35.4	30.5	37.6	64.6	81.9
	CPlaNet[1-2]	14.8	28.7	31.6	35.4	37.6	40.9	43.9	60.8	80.2
	CPlaNet[1-5]	16.0	29.1	33.3	36.7	39.7	42.2	46.4	62.4	78.5
	CPlaNet[1-5,PlaNet]	16.5	29.1	33.8	37.1	40.5	42.6	46.4	62.0	78.5

performance of our models is also better than the retrieval-based models at most scales. Moreover, our model, like other classification-based approaches, requires much less space than the retrieval-based models for inference. Although Deep-Ret+ improves Deep-Ret by increasing the size of the database, it even worsens space and time complexity. In contrast, the classification-based approaches including ours do not require extra space when we have more training images.

Figure 4 presents qualitative results of CPlaNet[1-5] on Im2GPS. It shows how the combinatorial partitioning process improves the geolocalization quality. Given an input image, each map shows an intermediate prediction as we accumulate the scores on different geoclass sets one by one. The region with the highest score is progressively sharded into a smaller region with fewer S2 cells, and the center of the region gradually approaches to the ground-truth location as we integrate more classifiers for inference.

Computational Complexity Although CPlaNet achieves competitive performance through combinatorial partitioning, one may be concerned about potential increase of time complexity for its inference due to additional classification layers and overhead in combinatorial partitioning process. However, it turns out that the extra computational cost is negligible since adding few more classification layers on top of the shared feature extractor does not increase inference time substantially and the required information for combinatorial partitioning is pre-computed as described in Section 4.2. Specifically, when we use 5 classification branches with combinatorial partitioning, theoretical computational costs for multi-head classification and combinatorial partitioning are only 2% and 0.004% of that of feature extraction process. In terms of space complexity, classification based methods definitely have great advantages over retrieval based ones, which need to maintain the entire image database. Compared to a single-head classifier, our model with five base classifiers requires just four additional classification layers, which incurs moderate increase of memory usage.

Importance of Visual Features For geoclass set generation, all the parameters of ClassSet 1 and 2 are set to the same values except for the relative importance of

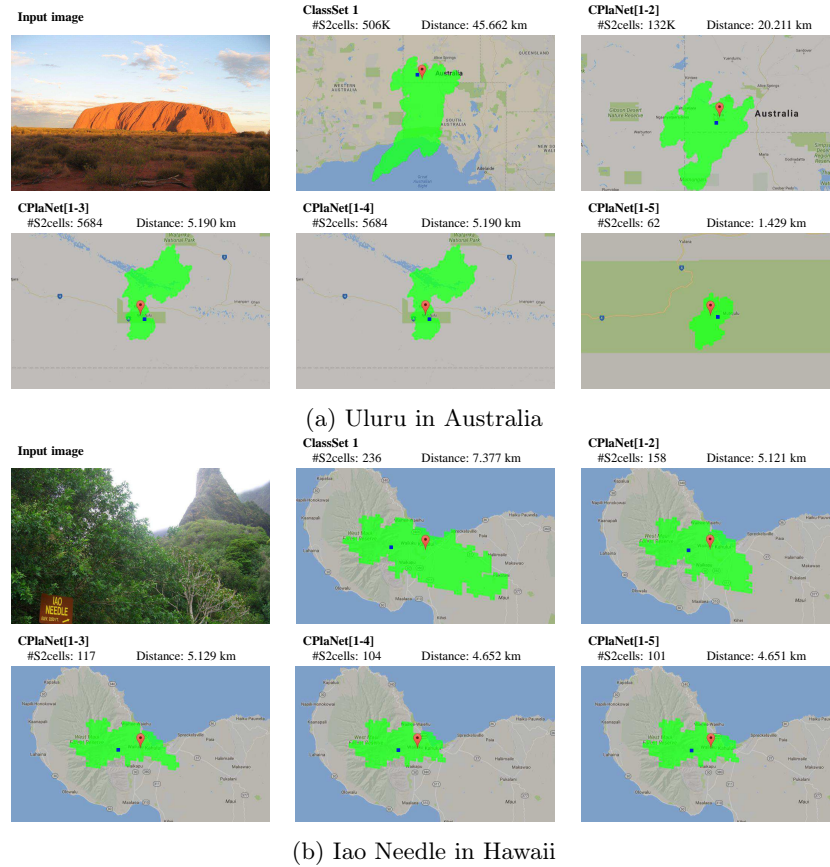


Fig. 4: Qualitative results of CPlaNet[1-5] on Im2GPS. Each map illustrates the progressive results of combinatorial partitioning by adding classifiers one by one. S2 cells with the highest score and their centers are marked by green area and red pins respectively while the ground-truth location is denoted by the blue dots. We also present the number of S2 cells in the highlighted region and distance between the ground-truth location and the center of the region in each map.

two factors for edge weight definition; edge weights for ClassSet 1 are determined by visual distances only whereas those for ClassSet 2 are based on geolocalational distances between the cells without any visual information of images. ClassSet 1 presents better accuracies at almost all scales as in Table 2. This result shows how important visual information of images is when defining geoclass sets.

Moreover, we build another model (ImageNetFeat) learned with the same geoclass set with ClassSet 1 but using a different feature extractor pretrained on ImageNet [43]. The large margin between ImageNetFeat and ClassSet 1 indicates importance of feature representation methods, and implies unique characteristics of visual cues required for image geolocalization compared to image classification.

Table 5: Comparisons between the models with and without normalization for combinatorial partitioning on Im2GPS3k. Each number in parentheses denotes the geoclass set size, which varies largely to highlight the effect of normalization for this experiment.

Models	1 km	5 km	10 km	25 km	50 km	100 km	200 km	750 km	2500 km
ClassSet 1 (9969)	8.4	18.3	21.7	24.7	27.4	29.8	34.1	47.9	64.5
ClassSet 2 (9969)	8.0	17.6	20.6	23.8	26.2	29.2	32.7	46.6	63.9
ClassSet 3 (3416)	4.2	15.9	19.1	22.8	24.9	28.0	31.4	46.1	63.5
ClassSet 4 (1444)	1.8	9.5	13.2	16.8	21.2	24.5	29.5	44.4	61.8
ClassSet 5 (10600)	8.2	19.1	22.3	25.2	27.3	29.9	33.6	47.3	65.5
SimpleSum	9.7	19.4	23.1	26.6	28.1	30.6	33.8	47.7	64.0
NormalizedSum	9.8	19.8	23.6	26.8	28.8	31.1	34.9	48.3	65.0

Balancing Classifiers We normalize the scores assigned to individual S2 cells as discussed in Section 4.2, which is denoted by NormalizedSum, to address the artifact that sums of all S2 cell scores are substantially different across classifiers. To highlight the contribution of NormalizedSum, we conduct an additional experiment with classsets that have large variations in number of classes. Table 5 presents that NormalizedSum clearly outperforms the combinatorial partitioning without normalization (SimpleSum) while SimpleSum still illustrates competitive accuracy compared to the base classifiers.

6 Conclusion

We proposed a novel classification-based approach for image geolocation, referred to as CPlaNet. Our model obtains the final geolocation of an image using a large number of fine-grained regions given by combinatorial partitioning of multiple classifiers. We also introduced an inference procedure appropriate for classification-based image geolocation. The proposed technique improves image geolocation accuracy with respect to other methods in multiple benchmark datasets especially at fine scales, and also outperforms the individual coarse-grained classifiers.

Acknowledgment

The part of this work was performed while the first and last authors were with Google, Venice, CA. This research is partly supported by the IITP grant [2017-0-01778], and the Technology Innovation Program [10073166] funded by the Korea government MSIT and MOTIE, respectively.

References

1. Hays, J., Efros, A.A.: Im2GPS: Estimating geographic information from a single image. In: CVPR. (2008)
2. Hays, J., Efros, A.A.: Large-scale image geolocalization. In: Multimodal Location Estimation of Videos and Images. Springer (2015) 41–62
3. Vo, N., Jacobs, N., Hays, J.: Revisiting IM2GPS in the deep learning era. In: ICCV. (2017)
4. Weyand, T., Kostrikov, I., Philbin, J.: Planet-photo geolocation with convolutional neural networks. In: ECCV. (2016)
5. Jegou, H., Douze, M., Schmid, C.: Product quantization for nearest neighbor search. TPAMI **33**(1) (2011) 117–128
6. Arandjelovic, R., Gronat, P., Torii, A., Pajdla, T., Sivic, J.: NetVLAD: CNN architecture for weakly supervised place recognition. In: CVPR. (2016)
7. Kim, H.J., Dunn, E., Frahm, J.M.: Learned contextual feature reweighting for image geo-localization. In: CVPR. (2017)
8. Baatz, G., Koeser, K., Chen, D., Grzeszczuk, R., Pollefeys, M.: Handling urban location recognition as a 2d homothetic problem. In: ECCV. (2010)
9. Cao, S., Snavely, N.: Graph-based discriminative learning for location recognition. IJCV **112**(2) (2015) 239–254
10. Chen, D., Baatz, G., Köser, K., Tsai, S., Vedantham, R., Pylvänäinen, T., Roimela, K., Chen, X., Bach, J., Pollefeys, M., Girod, B., Grzeszczuk, R.: City-scale landmark identification on mobile devices. In: CVPR. (2011)
11. Kim, H.J., Dunn, E., Frahm, J.M.: Predicting good features for image geolocalization using per-bundle vlad. In: ICCV. (2015)
12. Knopp, J., Sivic, J., Pajdla, T.: Avoiding confusing features in place recognition. In: ECCV. (2010)
13. Philbin, J., Chum, O., Isard, M., Sivic, J., Zisserman, A.: Object retrieval with large vocabularies and fast spatial matching. In: CVPR. (2007)
14. Schindler, G., Brown, M., Szeliski, R.: City-scale location recognition. In: CVPR. (2007)
15. Zamir, A.R., Shah, M.: Accurate image localization based on google maps street view. In: ECCV. (2010)
16. Zamir, A.R., Shah, M.: Image geo-localization based on multiple nearest neighbor feature matching using generalized graphs. PAMI **36**(8) (2014)
17. Noh, H., Araujo, A., Sim, J., Weyand, T., Han, B.: Large-scale image retrieval with attentive deep local features. In: ICCV. (2017)
18. Irschara, A., Zach, C., Frahm, J.M., Bischof, H.: From structure-from-motion point clouds to fast location recognition. In: CVPR. (2009)
19. Li, Y., Snavely, N., Huttenlocher, D., Fua, P.: Location recognition using prioritized feature matching. In: ECCV. (2010)
20. Li, Y., Snavely, N., Huttenlocher, D., Fua, P.: Worldwide pose estimation using 3d point clouds. In: ECCV. (2012)
21. Liu, L., Li, H., Dai, Y.: Efficient global 2d-3d matching for camera localization in a large-scale 3d map. In: ICCV. (2017)
22. Sattler, T., Leibe, B., Kobbelt, L.: Fast image-based localization using direct 2d-to-3d matching. In: ICCV. (2011)
23. Sattler, T., Weyand, T., Leibe, B., Kobbelt, L.: Image retrieval for image-based localization revisited. In: BMVC. (2012)

24. Sattler, T., Torii, A., Sivic, J., Pollefeys, M., Taira, H., Okutomi, M., Pajdla, T.: Are large-scale 3d models really necessary for accurate visual localization? In: CVPR. (2017)
25. Kendall, A., Cipolla, R.: Geometric loss functions for camera pose regression with deep learning. CVPR (2017)
26. Kendall, A., Grimes, M., Cipolla, R.: PoseNet: A convolutional network for real-time 6-dof camera relocalization. In: ICCV. (2015)
27. Walch, F., Hazirbas, C., Leal-Taixé, L., Sattler, T., Hilsenbeck, S., Cremers, D.: Image-based localization using lstms for structured feature correlation. In: ICCV. (2017)
28. Avrithis, Y., Kalantidis, Y., Toliás, G., Spyrou, E.: Retrieving landmark and non-landmark images from community photo collections. In: MM. (2010)
29. Gammeter, S., Quack, T., Van Gool, L.: I know what you did last summer: Object-level auto-annotation of holiday snaps. In: ICCV. (2009)
30. Johns, E., Yang, G.Z.: From images to scenes: Compressing an image cluster into a single scene model for place recognition. In: ICCV. (2011)
31. Quack, T., Leibe, B., Van Gool, L.: World-scale mining of objects and events from community photo collections. In: CIVR. (2008) 47–56
32. Zheng, Y.T., Zhao, M., Song, Y., Adam, H., Buddemeier, U., Bissacco, A., Brucher, F., Chua, T.S., Neven, H.: Tour the world: Building a web-scale landmark recognition engine. In: CVPR. (2009)
33. Weyand, T., Leibe, B.: Visual landmark recognition from internet photo collections: A large-scale evaluation. CVIU **135** (2015) 1–15
34. Bergamo, A., Sinha, S.N., Torresani, L.: Leveraging structure from motion to learn discriminative codebooks for scalable landmark classification. In: CVPR. (2013)
35. Li, Y., Crandall, D.J., Huttenlocher, D.P.: Landmark classification in large-scale image collections. In: ICCV. (2009)
36. Gronat, P., Obozinski, G., Sivic, J., Pajdla, T.: Learning per-location classifiers for visual place recognition. In: CVPR. (2013)
37. Workman, S., Souvenir, R., Jacobs, N.: Wide-area image geolocalization with aerial reference imagery. In: ICCV. (2015)
38. Lin, T.Y., Belongie, S., Hays, J.: Cross-view image geolocalization. In: CVPR. (2013)
39. Lin, T.Y., Cui, Y., Belongie, S., Hays, J.: Learning deep representations for ground-to-aerial geolocalization. In: CVPR. (2015)
40. Tian, Y., Chen, C., Shah, M.: Cross-view image matching for geo-localization in urban environments. In: CVPR. (2017)
41. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: CVPR. (2015)
42. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: CVPR. (2016)
43. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: ICCV. (2009)