

This WACV workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Enhancing Worldwide Image Geolocation by Ensembling Satellite-Based Ground-Level Attribute Predictors

Michael J. Bianco mjbianco.01@gmail.com David Eigen deigen@clarifai.com Clarifai Inc., San Francisco Michael Gormish michael.gormish@clarifai.com

Abstract

We examine the challenge of estimating the location of a single ground-level image in the absence of GPS or other location metadata. Geolocation systems are typically evaluated by measuring the Great Circle Distance between a single predicted location and the ground truth. Because this measurement only uses a single point, it cannot assess the quality of an estimated set of regions or score heatmaps. It is critical to characterize the distribution of potential geolocation areas to help when a system is applied to less wellsampled regions, such as rural and wilderness areas where finding the exact location may not be possible. This evaluation may help characterize a system in relation to follow-on procedures that further narrow down or verify predicted locations.

This paper introduces a novel metric, Recall vs Area (RvA), which assesses distributions of location estimates. RvA treats image geolocation results in a manner similar to precision-recall in document retrieval, measuring recall as a function of area. For an ordered list of (possibly noncontiguous) predicted regions, we measure the accumulated area required for the region to contain the ground truth coordinate. This produces a curve analogous to precisionrecall, enabling evaluation for varying search budgets.

This view of the problem inspires a simple ensembling approach to global-scale image geolocation that combines multiple models, attribute predictors, and data sources. Specifically, we combine the geolocation models GeoEstimation [11] and GeoCLIP [2] with attribute predictors based on ORNL LandScan [16] and ESA Land Cover [1]. We find notable improvements in non-urban and underrepresented areas on Im2GPS3k and Street View datasets.

1. Introduction

Accurate localization of ground-level imagery is important to a variety of applications, including navigation, tourism, and security, and has been the subject of recent research activity [2, 3, 5, 11, 20]. Past works on this problem have primarily relied on publicly available geotagged image datasets to learn location estimators. However, even large geotagged datasets [8, 10, 11, 17, 18] sample only a small fraction of the Earth's surface, concentrating mostly in cities and other population centers, with the most highly sampled locations corresponding to the most highly trafficked sites, such as famous landmarks, parks, restaurants, etc. Thus, when these geolocation systems are applied globally, many areas are under-sampled, and so applications in rural and other under-sampled areas are challenged by lack of representation in the training set.

To help address the challenge of generalization to undersampled areas, we first introduce an novel image geolocation metric, *Recall vs Area (RvA)*. This metric naturally characterizes model performance when predicted regions are discontiguous, which is important when the exact location may not be possible to recover, such as in applications to rural scenes or wilderness. Even though a single location may not be possible to predict in these situations, geolocation systems can still be used to suggest areas that can be studied by human analysts or further refined by additional automated systems.

Inspired by this view of the problem, we examine a general ensembling approach to limit predicted area by incorporating ground-level attribute predictors corresponding to attributes visible by satellite. While ground level geotagged data are limited, satellite data have global coverage. Thus, accurate prediction of attributes mapped by satellite can extend to regions where there are few geotagged ground-level images. For example, different land cover types (forest, mountains, roads, etc.) are in both ground-level and satellite images. We consider two data products in our analysis: Oak Ridge National Laboratory (ORNL) LandScan Global [16] and the European Space Agency Climate Change Initiative (ESA-CCI) Land Cover [1].

We combine estimates from these attribute predictors with those from other global geolocation models, i.e. GeoEstimation [11] and GeoCLIP [2], to improve performance on rural regions under-sampled in the geotagged data. We evaluate our method using both the standard Im2GPS3k [18], as well as a Street View dataset (see Sec. 4.2) that is different in geographic distribution and scene concepts.

1.1. Contributions

- We develop a novel metric, Recall vs Area (RvA), for evaluating geolocation systems, which can account for predictions of multiple, non-contiguous areas. We use this method to assess image geolocation approaches.
- We show that, in isolation, geolocation methods trained only on geotagged imagery can be vulnerable to image domain shift with respect to training data.
- We develop a new approach to combine ground-level attribute predictors with geolocation models, leveraging satellite data products, and show that incorporating the satellite information improves generalization to areas of the Earth for which there are few ground-level picture samples.

2. Related Work

Müller-Budack *et al.* [11] develop a hierarchical cellbased approach to geolocation. They fine-tune their model on the large MP-16 dataset [8], and quantize the Earth using S^2 cells (using the Google S2 geometry Library¹), with three resolutions {'coarse', 'medium', 'fine'} depending on data density. Their model outputs image geolocations for each level, which are combined to create a final estimate over the finest level cells. Subsequent papers have built on this work, proposing various cell division strategies, including CPlaNet [15] and Translocator [12]. These approaches rely exclusively on geotagged ground-level images as their datasource, which are assembled into cell bins of different types.

Other recent works have brought additional information to the geolocation problem. Lin et al. [9] explored using land-cover information for improving query-based geolocation, but required associated triplets of ground-level and aerial imagery and land use and was applied only to a small 40×40 km region. More recent works include relationships between the scene and variations in environment [3] or from segmentation models [12]. Luo et al. [10] and Cepeda et al. [2] exploit geographical correlations encoded in visionlanguage foundation models (using Contrastive Language-Image Pretraining [13]) to improve image geolocation performance.

There is also significant work in cross-view image localization, e.g. [6, 9, 19]. In these approaches, optical satellite imagery is matched to ground-level imagery to achieve ground-level image localization. These methods are used to pinpoint exact locations within limited-area urban scenes where dense ground-level imagery is available. This is complementary to our work, which uses ground-level attributes and satellite sources to limit search areas over a global field; indeed, these works are examples of systems that might run downstream of ours within a proposed area.

As far as we are aware, the approach we present is the first to use land cover and other attributes recognizable from ground level for global-scale geolocation. We combine predictions from multiple sources, attribute predictors associated with satellite-based maps, along with cell-based estimation systems, using a common ensembling grid. This yields an effective intersection of non-contiguous regions predicted from the available evidence to limit the geolocation area.

3. Approach

3.1. Recall vs Area Metric

Existing approaches typically measure geolocalization accuracy by calculating the distance of a single predicted location from the ground truth. Distance is calculated using great circle distance (GCD), and binned by a small set of thresholds. The fraction of images whose top prediction lies within each threshold is reported as the accuracy for that distance. For example [11] used GCD thresholds of [1, 25, 200, 750, 2500] km to evaluate model performance, even giving these distances names to represent how close the solution was, namely street, city, region, country, continent.

While this measure works well to characterize performance of a single predicted location, it does not extend to measuring performance of a larger (possibly discontiguous) predicted area. This is important for many practical applications, where the image geolocation algorithm is used to propose a set of possible locations that are used in downstream tasks. Thus, we would like to measure localization performance in terms of an area budget.

In addition, many distant regions on Earth look similar from ground-level, as illustrated by the often scattered global distribution of high confidence S^2 cell predictions by GeoEstimation (see e.g. Fig. 1(b)). In the context of a system that progressively narrows down the possible locations, it is important to maintain recall at each stage while still reducing the set of predicted regions.

We propose a new metric, Recall vs Area (RvA), to characterize image geolocation model performance. We rely on the conventions of recall and precision, and formulate a geographic analogy to document retrieval. For each image, a model predicts scores over a set of region proposal areas. To evaluate, we accumulate areas by score, up to a given search budget. Recall is determined according to whether

https://code.google.com/archive/p/s2-geometrylibrary/



Figure 1. (A) Visualization of 'low' population density mask obtained from ORNL LandScan global, see Sec. 5.1 for discussion. Masks obtained from LandScan as well as ESA-CCI Land Cover are used by our ensembling method to improve image location estimates of base geolocation models. (B) The objective of our proposed ensembling approach is, given a single ground-level RGB image, to obtain a global geolocation probability map that incorporates information from base geolocation models (in this work, GeoEstimation and GeoCLIP) and ground-level maps obtained from satellite data products. The image is passed to both the base geolocation model and the individual ground-level attribute predictors. Example output S^2 cell and ground level attribute probabilities from Algo. 1 and 2, and final ensembled probabilities are shown. The high probability area is significantly reduced in ensembling. The probability maps are thresholded for visualization.



Figure 2. Batches of images from (left) MP-16 and (right) Street View, from medium urban density regions per LandScan labeling (see Sec. 4.2, 5.1). In general, the MP-16 dataset consists of more typical urban scenes, whereas Street View contains more scenery from near roads. Further MP-16 dataset is drawn from areas much more close to urban environments.

the ground truth lies in the accumulated area. We then compute recall for different area budgets (thresholds), resulting in a ROC-like curve comparing recall on one axis and area budgets on the other.

Here, recall is determined by whether the geographic area returned by the model for a query image contains the image's location. More formally, for a query image $x_i \in \mathcal{X}$ with ground truth location y_i , a geolocation model will return a list of location areas (S^2 cells, points, etc.) $a_{ik}, k \in [1..K]$, along with confidence scores. We accumulate area sorted by model confidence, to create a (possibly discontiguous) region, up to an area budget (threshold). For an area threshold α , recall is then

$$recall(\alpha) = \frac{\sum_{i} \mathbb{1}[y_i \in \{a_{ik} : \sum_{k'=1}^{k} a_{ik'} \le \alpha\}]}{|\mathcal{X}|}, \quad (1)$$

i.e. the fraction of images whose ground-truth location lies within the model's top-scoring predicted areas, up to the area threshold α .

We then plot the curve $recall(\alpha)$ for different values of the area budget α , resulting in a description of model performance, akin to a precision-recall curve (see Figs. 4 and 5).² The curve is used to evaluate model performance, and can also be used when selecting appropriate trade-offs between recall and size of a proposed area.

3.2. Algorithm

Given an image x, we ensemble geolocation and groundlevel attribute predictors to find a subset of areas on Earth most likely to contain the image. Fig. 1 illustrates this process.

Because we combine the geolocation area predictions of multiple models, we map each model's output to a common grid over the surface of the Earth, with "pixel" coordinates given by latitude and longitude (WGS-84 projection). Note that we index by (lat, lon), so each "pixel" has a different area on the Earth depending on its latitude coordinate. To account for this, we also create an array A_{pix} that contains the surface area corresponding to each point in the map.

²Indeed, area itself can be thought of as a notion of inverse precision, with smaller areas being more precise and larger ones including more false positive locations and therefore less precise.

Algorithm 1 Assigning probabilities to the common grid

Given image x, model f, and mask(s) \mathcal{M}_f # compute softmax score for each mask using the model p = f(x)# combine into rasterized grid $P_f = \sum_{j=1} p_j m_j / area(m_j)$, with $m_j \in \mathcal{M}_f$ return P_f

Algorithm 2 Ensembling probabilities from ground-level attribute predictors and GeoEst./GeoCLIP

Given $x, \{f_1 \dots f_K\}, \{\mathcal{M}_1 \dots \mathcal{M}_K\}$ return $P = \prod_{k=1}^K P_{f_k}$, with P_{f_k} as in Alg.1

Algorithm 3 Model evaluation

Given dataset \mathcal{D} , model, pixel spherical areas map A_{pix} Returns accumulated areas necessary to the include ground truth location for each image. $\mathcal{A} \leftarrow \{\} \text{ # eval. areas}$ for (image x_i , g.t.location y_i) $\in \mathcal{D}$ do $P \leftarrow \text{model}(x)$ $p^* = P[y_i] \text{ # predicted prob at ground truth loc.}$ $area = \sum A_{pix} \cdot \mathbb{1}[P \ge p^*] \text{ # area of pred} \ge p^*$ $\mathcal{A} \leftarrow \mathcal{A} \cup area$ end for return \mathcal{A}

The output of each model $f(\cdot)$ is a vector of probabilities $p = [p_1 \dots p_J] \in \mathbb{R}^J$, corresponding a set of binary masks $\mathcal{M}_f = \{m_1 \dots m_J\}$, with $|\mathcal{M}| = J$, and $m \in \{0, 1\}^{H \times W}$ with H and W the vertical (latitude) and horizontal (longitude) resolution of the entire surface grid. For GeoEstimation [11], each mask is an S^2 cell and each probability is the model's softmax score for the corresponding cell. To map GeoCLIP [2] predictions to the common grid, we use the same S^2 cells as in GeoEstimation, since GeoCLIP was trained using MP-16, which is the same dataset used to distribute the S^2 cells. We obtain a score for each cell by passing the center location to GeoCLIP.

For LandScan and Land Cover attribute predictors, each mask m_j is a discontiguous region corresponding to a population density bucket, land use type, etc., and probabilities are the softmax scores for each attribute. The scores for the attribute masks and S^2 cells are assigned to the common grid, normalizing by their areas.

We then combine probabilities to produce a final set of scores, estimated over the common grid. The scores are obtained using a simple element-wise product of probability maps, which we found to be effective even though the masks are not independent:

$$P_{Ens.} = \prod_{k} P_{f_k}(x) \tag{2}$$

where each $P_{f_k}(x)$ is the probability map of the k^{th} predictor (GeoEst, Land Cover regions, etc.) as described in Algorithms 1 and 2.

For evaluation, locations (pixels) on the common grid are accumulated in score order for each image until the resulting region contains the ground truth location. Note that since we index the grid by (lat, lon), each element has a different area according to latitude, we accumulate area using a pre-computed matrix A_{pix} that contains the area of each pixel. This process is defined in Algorithm 3. This gives us an array with the smallest cumulative area needed to contain the ground truth for each image — i.e. the smallest area threshold α at which each image prediction is a "hit". We then sort this array to produce the RvA curve. In deployment, the area threshold can be used to calibrate between desired recall and an area budget.

4. Datasets

4.1. Ground-level attribute datasets and processing

In this work, we utilize two satellite data products: Land-Scan Global [16] from ORNL and Land Cover [1], from the ESA-CCI. Both data products provide high-resolution estimates of surface attributes: LandScan Global predicts 24-hour average population density, while ESA Land Cover classifies 38 natural and anthropogenic land cover types.

We used three image datasets. MP-16 [17] and Im2GPS3k [18] are commonly used training and evaluation benchmarks in image geolocation research. MP-16 was employed to train our attribute classifiers, and Im2GPS3k was used for testing. Additionally, we collected images from Google Street View, using training-set panorama IDs from [10], to assess how well the models generalize to image location distributions and capture pipelines that differ from MP-16 and Im2GPS3k.

4.1.1 LandScan Global

The ORNL LandScan Global dataset provides a global map of population density at a resolution of approximately 1 km^2 [16], as shown in Fig. 1(a). The data is a large TIFF image, 21600×43200 pixels (lat x lon), and contains 24-hour average population densities.

Each pixel indicates whether the corresponding region is land and if so, its value contains the population density. The image uses WGS-84 coordinates, a standard lat-lon representation, with a resolution of approximately 1 km².



Figure 3. Example overlay centered on Europe, showing S^2 cells (red boxes, fine level from GeoEstimation [11]) with LandScan [16] medium population density mask (white). The mask is much more restrictive than the S^2 cells alone, yielding improvements in RvA from model ensembling.

4.1.2 ESA Landcover

ESA-CCI Land Cover is a global land cover and land use data product [1]. The attributes are predicted using deep learning classifiers applied to Sentinel-2 satellite images. The dataset is a 64800×129600 pixel (lat x lon, WGS-84) GeoTIFF with 300 m resolution. Each pixel is labeled with the most likely land cover from 38 classes (grouped into 22 super-classes), including cropland, grassland, tree cover, and developed land use categories. Maps are available for the years 1992-2020. We used the Land Cover 2015 dataset for the development of region reduction classifiers.

4.2. Image datasets

The most comprehensive geotagged ground-level image dataset currently available is MP-16 [8], a subset of the Yahoo Flickr Creative Commons 100 Million datasets (YFCC100M) [17] that includes geotags. This dataset was used to train ground-level attribute predictors.

The size of the dataset is 4.7M images. of which we were able to obtain 4.2 million, as approximately 500 k images were no longer available at the URLs provided by [11]. The dataset includes a diverse range of image concepts and objects, featuring both indoor and outdoor scenes of natural and human-made environments.

MP-16 was originally the training set used in the GeoEstimation paper [11] and continues to be used in recent research, such as in [2,12]. In addition, we use the Im2GPS3k dataset [18] as a benchmark for testing our models.

For evaluation, we rely on Im2GPS3k [18] and a set of images from Google Street View [4]. To specifically assess the models' generalization and robustness to domain gaps, we do not train on any Street View images and use them exclusively for evaluation.

The Google Street View data [4] was obtained using panorama IDs from Luo et al. [10]. This geo-diverse dataset includes at least 426 panoramas from 90 countries. From this overall distribution, we randomly constructed train, validation, and test sets. We first sampled 50000 panorama IDs from a total of 322536 panoramas. Of the 50000 IDs chosen, 49829 IDs were still valid. We thus obtained 49829 images—each a single 90° FOV slice from a full panorama, with 15° downward pitch (toward ground to capture road markings) from the Street View API, along with corresponding image geotags.

5. Experiments

5.1. Ground-level attribute predictors

We trained two ground-level attribute predictors, one for each satellite data product (LandScan and Land Cover), using MP-16 imagery relabeled with attributes from each source. In our experiments, we use the acronyms LS for LandScan and LC for Land Cover. For each image, the ground-level attribute labels were assigned by indexing the nearest pixel value in the attribute maps at the native TIFF resolution. As described in Section 3.2 the outputs of these models are projected on global masks and combined to find a global distribution of probable locations.

Since the LS population densities are continuous, we categorized these into four density buckets for use in a softmax classifier. These buckets were created based on the \log_{10} of population density, using equal bin widths over the range of LS population densities. Fig. 6 shows the LS population density range sampled by MP-16 along with the bucket boundaries. This results in four buckets of ['lowest', 'low', 'medium', 'high'] density, cover land areas in percentages [93.9, 4.93, 1.10, 0.029]%, respectively. Thus the LS classifier predicts these 4 classes. In Fig. 3 the medium population density bucket is overlaid with the S^2 cells (fine) used for base (M, f^*) . The population density mask is more restrictive than the S^2 cells alone.

For LC, we labeled the MP-16 images using the 22 top-level land cover classes and then merged these into 7 broader classes that our classifier predicts (see Sec. 5.2).

To balance the classes during training, we limited the images in the 'urban area' LC category to 12.5 % of the MP-16 dataset. This adjustment was necessary because most MP-16 images are from densely populated areas, while our goal is to locate both urban and rural locations.

For our predictors, we used a pretrained ResNet-50v2 architecture from [7], as implemented in the Pytorch Image Models (timm) library [21]. Specifically, we use the Big Transfer (BIT)-M version of the model, which is pretrained on the public ImageNet21k dataset [14].

The predictors were trained on the relabeled MP-16 imagery for each of the LS classes for a maximum of 100



Figure 4. (a) *Recall vs. Area* (*RvA*) obtained for rasterized GeoEst (Alg.1), alone and ensembled with LandScan (+LS) attribute prediction. As an additional baseline, we compare always applying the LS "urban" mask instead of predicting the bucket (+urban, dotted curve). Spherical cap areas shown in blacked dashed lines, see Sec. 5.3 for details. (b) RvA, measured on data rebalanced over urban and non-urban areas by randomly sampling equal number of images from each LS mask. Our method improves performance for less-populated areas while maintaining urban area results. (c) Relative improvement measured on the balanced datasets.



Figure 5. Rasterized GeoCLIP 1 and ensembling: same analysis as Fig. 4(a–c).

epochs using distributed training over 4 GPUs. In training, the images were augmented using random cropping, rescaling, horizontal and vertical flipping, and color jitter.

5.2. Ensembling

In our experiments, the common grid size was H = 5400(latitude) by W = 10800 (longitude). These dimensions were chosen because they are common factors of the LS and LC TIFF image sizes, and are small enough to yield reasonable time complexity and geolocalization performance. The average resolution of the common image is ~ 4 km, which corresponds to zero recall at ~ 16 km² in Figs. 4 and 5, and Tables 1 and 2.

For GeoEstimation, we used the published hierarchical model ("base (M, f^*) " from [11]). For GeoCLIP, we used

the published model architecture and weights. Because both GeoEst. and GeoCLIP are fine-tuned with MP-16, we use the S^2 cell centers, which are also distributed according to MP-16, to sample the GeoCLIP location encoder and obtain scores for each of the S2 cells.

To generate the probability map $p \in \mathbb{R}^{M \times N}$, we rasterized the S^2 cells from base (M, f^*) to the common image following Algo. 1 (see Sec. 3.2).

The LS and LC TIFFs were downsampled by factor of 4 and 12 respectively, to the common image resolution. The binary masks in Algo. 2 were obtained from the downsampled images. Due to sparse sampling, some LC classes were merged by visual similarity prior to ensembling. Specifically, LC classes including cropland were merged (values 10-40 in [1]), as were broad-leaf tree cover (50–60), and the



Figure 6. (a) Cumulative distribution of MP-16 Land-Scan (LS) population density. LS buckets shown as ['lowest', 'low', 'medium', 'high']. (b) Distribution of original (unbalanced) Im2GPS3k (top) and Street View images (bottom) in LandScan (LS) buckets.

remaining tree cover on dry land (70–100). Short vegetation, including lichen and mosses (110–150), and vegetated and flooded regions (160–190) were also merged. Urban areas were not merged, while bare areas were merged with water and permanent snow and ice (200–220). As a result, our models predict 7 classes for LC after merging.

5.3. Results

We evaluate the geolocation methods using Algo. 3 for both Im2GPS3k and the Street View dataset, based on pano-IDs from [10]. Results are presented in Figs. 4 and 5, and Tables 1 and 2.

Both Im2GPS3k and Street View datasets have uneven population density distributions (Im2GPS3k is predominantly urban, while Street View has relatively more nonurban). To reduce this bias and better assess the effects on each population density type, we created rebalanced versions of the datasets by randomly sampling an equal number of images from each LandScan mask bucket. We evaluate on both the original and rebalanced distributions.

The results are summarized in terms of recall vs. area (RvA) in Figs. 4 and 5, both with and without balancing, and Tables 1 and 2 present results for the balanced imagery. We show results for the baseline rasterized algorithm for GeoEstimation or GeoCLIP predictions (Algo. 1); ensembling (Algo. 2) with GeoEstimation or GeoCLIP combined with LS (+LS); GeoEstimation or GeoCLIP combined with LC (+LC); and GeoEstimation or GeoCLIP combined with both LS and LC (+LS+LC). Black dashed lines in the figures indicate areas of spherical caps with radii corresponding to the GCD thresholds typically used in literature, e.g., [2, 11]. It should be noted that the resolution of our common image prevents the model from returning areas corresponding to the 1 km GCD threshold spherical cap (see Sec. 5.2). The recall of the remaining ensembles is summarized as a function of global spherical cap areas corresponding to typical GCD thresholds used in literature with radius of the thresholds $(3.1, 2.0 \times 10^3, 1.3 \times 10^5, 1.8 \times 10^6, \text{ and } 2.0 \times 10^7 \text{km}^2)$. Since in our proposed approach we accumulate all proposed areas in the final probability map (per Eq. 1), recall is considered as a function of the total area of these often discontinuous regions.

We also quantify performance using conventional GCD. For direct comparison, the top-1 accuracy is computed using the common grid we used to implement RvA. The prediction is chosen as the highest-probability location from the common grid; note that due to the ~ 4 km grid resolution and the fact that pixels in non-contiguous areas can share values, there is no clear analog to a "cell center" in our approach, so this selection method results in somewhat worse GCD relative to published results.

Because GCD only evaluates the top-1 location, it is not able to capture the gains our method provides in limiting region extent: indeed, our method shows little improvement in the GCD for the single most confident point, even though it produces tangible gains in performance, as shown by the RvA metric.

For Im2GPS3k, the image counts for each LandScan mask (from lowest to highest population density) are [540, 548, 1159, 715]. To rebalance, we randomly sample 540 images from each bucket, resulting in a total of 2160 images. Similarly, for Street View, the images in each mask bucket are [11846, 17623, 18668, 1692], with the smallest number of images in the highest population density mask. Rebalancing results in 6768 images.

In Figs. 4, we observe larger gains for Street View, with the most significant gains occurring at the larger area thresholds, which generally correspond to rural areas. Note that there is also a performance drop for both base models on the out-of-domain Street View dataset compared to the typical Im2GPS3k evaluation set; however, we are able to recover a portion of this difference. Figs. 4-5(c) display the absolute improvement in recall over the base models.

We also include a constant "urban prior" model as an additional baseline, labeled "+LS, urban", which always applies a mask restricting the data to the two highest-density LS buckets. Since the original imagery is primarily drawn from urban regions, such a constant prior might be effective. However, we find that the +LS attribute predictors outperform the constant prior, as expected, particularly in larger area regions. The urban prior model discards these regions wholesale, whereas predicted population density improves performance for these areas. Thus, the gains are not simply due to applying an urban region mask, but rather predicting different density attribute values is necessary to achieve improvements.

For the rebalanced datasets, we observe overall reductions in recall for the Im2GPS3k set and improvements for Street View. This is because the original Im2GPS3k dataset

	RvA					GCD					
Method	Street 1 km	City 25 km	Region 200 km	Country 750 km	Continent 2500 km	Street 1 km	City 25 km	Region 200 km	Country 750 km	Continent 2500 km	
GeoEst.		39.42	65.98	85.37	97.02	4.21	23.24	31.25	43.89	60.97	
+LS	_	35.08	65.58	85.89	97.45	3.33	20.79	28.94	41.02	58.94	
+LS, urban prior	—	31.60	43.20	47.45	47.50	3.61	21.25	28.94	42.22	59.77	
+LC	_	36.91	65.90	86.78	97.23	3.43	22.55	31.20	43.94	61.02	
+LS+LC		32.75	61.81	84.20	97.07	3.01	20.46	28.70	41.16	58.94	
GeoCLIP		41.19	68.34	88.42	97.40	4.91	30.37	43.84	62.22	79.81	
+LS	_	40.96	70.97	90.57	97.84	4.07	28.61	42.59	61.39	78.98	
+LS, urban prior		32.96	44.18	47.22	47.22	4.35	27.96	41.34	61.34	78.94	
+LC	_	40.29	71.29	90.24	97.86	4.40	29.12	43.01	62.82	79.68	
+LS+LC	-	37.48	69.14	89.88	97.86	3.80	27.82	42.31	61.30	78.75	

Table 1. Results on Im2GPS3k [18] (balanced)

	RvA				GCD					
Method	Street 1 km	City 25 km	Region 200 km	Country 750 km	Continent 2500 km	Street 1 km	City 25 km	Region 200 km	Country 750 km	Continent 2500 km
GeoEst.	—	7.96	34.73	68.21	92.45	0.01	4.79	12.37	32.68	57.24
+LS	_	9.85	42.14	74.27	94.39	0.07	5.27	13.39	33.85	59.84
+LS, urban prior	—	9.36	33.63	41.06	41.21	0.01	4.68	12.60	32.95	57.24
+LC	_	8.29	40.25	75.43	94.37	0.01	4.82	12.75	33.73	57.74
+LS+LC		9.39	39.83	73.95	94.67	0.10	5.29	13.45	34.10	59.65
GeoCLIP		13.28	44.70	77.51	93.97	0.15	10.96	29.03	61.07	81.47
+LS	_	16.56	54.42	83.37	94.62	0.19	11.14	29.95	62.38	82.77
+LS, urban prior		15.76	38.06	42.03	42.04	0.18	10.83	29.18	61.23	81.56
+LC	_	15.69	51.33	82.69	94.97	0.13	10.92	29.40	61.39	81.32
+LS+LC		16.09	53.25	83.21	95.03	0.24	11.18	30.04	62.53	82.71

Table 2. Results on Street View [10] (balanced)

is biased toward regions with the highest population density, whereas Street View samples from the highest density areas considerably less. Using our ensembling approach, recall on Im2GPS3k is reduced at smaller area scales for both the original and rebalanced datasets, but improvements are seen at larger scales (see Fig. 4–5, Tables 1 and 2). Peak improvements in absolute recall are 0.21 for GeoEst and 2.96 for GeoCLIP base models.

At smaller scales, e.g. spherical cap areas corresponding to 25 and 200 km, the +LS ensemble gave the largest improvements. The +LC and full +LS+LC ensembles provide benefits at larger scales.

The benefits of ensembling the ground-level attribute predictors at a global scale with the baseline image geolocation predictors are clear. The attribute predictors bring additional ground-truth information to the solution and improve performance relative to baseline.

6. Conclusions and Future Work

We have introduced a new image geolocation metric, *Recall vs Area* (RvA), along with an ensembling approach to global-scale image geolocation that incorporates ground-level attribute predictors based on satellite data products. Our metric addresses the challenge of accounting for ground-level image similarity across different geographic regions and can measure discontiguous predicted areas. Since practical image geolocation systems often prioritize search area, RvA naturally aligns with these needs. By ensembling our ground-level attribute predictors, we improved upon the state-of-the-art (SOTA) image geolocation method, GeoCLIP [2], in terms of RvA.

Our ensembling approach enhances image geolocation recall as a function of search area by combining groundlevel attribute predictors with geolocation models. These attribute predictors leverage satellite-based attribute masks, adding valuable information to the solution and improving performance over strong baselines, particularly in regions under-sampled in the original dataset.

Interestingly, the relative gain from ensembling predicted attribute masks was greater for GeoCLIP than for GeoEst, even though GeoCLIP began with a higher performance level. This suggests that the embeddings used in GeoCLIP may be complementary to our ground-level attributes, an area we hope to explore in future work.

By linking satellite-recognizable attributes to groundlevel images, we can extend location information to areas of the world not represented in geotagged image data but included in satellite coverage. We have validated this approach using population density and land cover, and aim to extend it to additional attributes.

References

- [1] European Space Agengcy. Esa: Land Cover CCI product user guide version 2.0, 2023. 1, 4, 5, 6
- [2] Vicente Vivanco Cepeda, Gaurav Kumar Nayak, and Mubarak Shah. GeoCLIP: Clip-inspired alignment between locations and images for effective worldwide geolocalization. arXiv preprint arXiv:2309.16020, 2023. 1, 2, 4, 5, 7, 8
- [3] Brandon Clark, Alec Kerrigan, Parth Parag Kulkarni, Vicente Vivanco Cepeda, and Mubarak Shah. Where we are and what we're looking at: Query based worldwide image geo-localization using hierarchies and scenes. *arXiv preprint arXiv:2303.04249*, 2023. 1, 2
- [4] Google. Streetview API, 2023. 5
- [5] James Hays and Alexei A Efros. Im2GPS: estimating geographic information from a single image. In 2008 ieee conference on computer vision and pattern recognition, pages 1–8. IEEE, 2008. 1
- [6] Sixing Hu, Mengdan Feng, Rang MH Nguyen, and Gim Hee Lee. CVM-Net: Cross-view matching network for imagebased ground-to-aerial geo-localization. In *Proceedings* of the IEEE Conference on Computer Vision and Pattern Recognition, pages 7258–7267, 2018. 2
- [7] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big Transfer (BiT): General visual representation learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 491–507. Springer, 2020. 5
- [8] Martha Larson, Mohammad Soleymani, Guillaume Gravier, Bogdan Ionescu, and Gareth JF Jones. The benchmarking initiative for multimedia evaluation: MediaEval 2016. *IEEE MultiMedia*, 24(1):93–96, 2017. 1, 2, 5
- [9] Tsung-Yi Lin, Serge Belongie, and James Hays. Cross-view image geolocalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 891–898, 2013. 2
- [10] Grace Luo, Giscard Biamby, Trevor Darrell, Daniel Fried, and Anna Rohrbach. G[^] 3: Geolocation via guidebook grounding. arXiv preprint arXiv:2211.15521, 2022. 1, 2, 4, 5, 7, 8
- [11] Eric Muller-Budack, Kader Pustu-Iren, and Ralph Ewerth. Geolocation estimation of photos using a hierarchical model and scene classification. In *Proceedings of the European conference on computer vision (ECCV)*, pages 563–579, 2018. 1, 2, 4, 5, 6, 7
- [12] Shraman Pramanick, Ewa M Nowara, Joshua Gleason, Carlos D Castillo, and Rama Chellappa. Where in the world is this image? Transformer-based geo-localization in the wild. In Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVIII, pages 196–215. Springer, 2022. 2, 5
- [13] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervi-

sion. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2

- [14] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21k pretraining for the masses, 2021. 5
- [15] Paul Hongsuck Seo, Tobias Weyand, Jack Sim, and Bohyung Han. CPlaNet: Enhancing image geolocalization by combinatorial partitioning of maps. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 536– 551, 2018. 2
- [16] K. Sims, A. Reith, E. Brigh, J. McKee, and A. Rose. Land-Scan Global 2021 [dataset]. 2022. 1, 4, 5
- [17] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. YFCC100M: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016. 1, 4, 5
- [18] Nam Vo, Nathan Jacobs, and James Hays. Revisiting Im2GPS in the deep learning era. In *Proceedings of the IEEE international conference on computer vision*, pages 2621– 2630, 2017. 1, 2, 4, 5, 8
- [19] Xiaolong Wang, Runsen Xu, Zhuofan Cui, Zeyu Wan, and Yu Zhang. Fine-grained cross-view geo-localization using a correlation-aware homography estimator. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [20] Tobias Weyand, Ilya Kostrikov, and James Philbin. Planetphoto geolocation with convolutional neural networks. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VIII 14*, pages 37–55. Springer, 2016. 1
- [21] Ross Wightman. Pytorch image models. https: //github.com/rwightman/pytorch-imagemodels, 2019. 5