

Presence-Only Geographical Priors for Fine-Grained Image Classification

Oisín Mac Aodha Elijah Cole Pietro Perona
Caltech

www.vision.caltech.edu/~macaodha/projects/geopriors

Abstract

Appearance information alone is often not sufficient to accurately differentiate between fine-grained visual categories. Human experts make use of additional cues such as where, and when, a given image was taken in order to inform their final decision. This contextual information is readily available in many online image collections but has been underutilized by existing image classifiers that focus solely on making predictions based on the image contents.

We propose an efficient spatio-temporal prior, that when conditioned on a geographical location and time, estimates the probability that a given object category occurs at that location. Our prior is trained from presence-only observation data and jointly models object categories, their spatio-temporal distributions, and photographer biases. Experiments performed on multiple challenging image classification datasets show that combining our prior with the predictions from image classifiers results in a large improvement in final classification performance.

1. Introduction

Correctly classifying objects into different fine-grained visual categories is a challenging problem. In contrast to generic object recognition, it can require knowledge of subtle features that are essential for differentiating between visually similar categories. However, without having access to additional information that may not be present in an image, many categories can be visually indistinguishable. For example, the two toad species in Fig. 1 are similar in appearance but tend to be found in very different locations in Europe. Knowing *where* a given image was taken can provide a strong prior for *what* objects it may contain.

Most images that are captured and shared online today also come with additional metadata in the form of *where* they were taken, *when* they were taken, and *who* captured them. This information not only offers the possibility of helping to resolve ambiguous cases for image classification, but can also enable us to generate predictions of where, and when, different objects are likely to be observed.

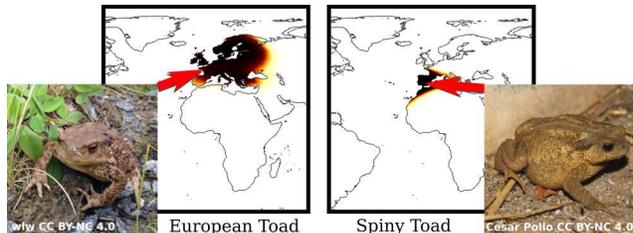


Figure 1: Differentiating between two visually similar categories such as the European (left) and Spiny (right) Toad can be challenging without additional context. To address this problem, we propose a spatio-temporal prior that encodes where, and when, a given category is likely to occur. For a known test location our prior predicts how likely it is for each category to be present. Darker colors indicate locations that are more likely to contain the object of interest.

Existing work that uses location information to improve classification performance either discretizes the input data into spatio-temporal volumes [5], store the entire training set in memory at inference time [64], or jointly train deep images classifiers along with corresponding location information [55]. Methods that discretize or store the raw training data do not scale well in terms of memory, and jointly training image classifiers with location information necessitates that location information is present at test time - which may not always be the case. We take inspiration from species distribution modeling (SDM) [20], and instead model a separate geographical prior that can be combined with the predictions of *any* image classifier. However, unlike many approaches to SDM that assume they have access to presence and absence information at training time (e.g. [56]), we make a more general assumption that only presence information is available *i.e.* we know where the categories have been observed, but have *no* explicit data regarding where they are not found.

In this work we make the following contributions: (1) An efficient spatio-temporal prior that jointly models the relationship between location, time of year, photographer, and the presence of multiple different object categories. (2) A novel presence-only training loss to capture these relation-

ships. (3) Experiments that show that combining the probabilistic predictions of image classifiers with our prior significantly improves the test time performance on challenging fine-grained image datasets.

2. Related Work

Here we discuss work related to spatio-temporal models that encode the location of a set of discrete object categories. We do not address methods that explore other uses of location information such as inferring where an image was taken given only the raw pixels [29, 60], or methods that use location to disambiguate visually similar places for image localization [59, 68].

Fine-Grained Image Classification

Correctly determining which one of multiple possible fine-grained categories is present in an image requires understanding the relationship between subtle visual features and the corresponding image-level category label *e.g.* [61, 35, 66, 58]. Existing approaches have investigated the modeling of parts [41, 71, 9, 70, 32], higher order feature interactions [40, 23], attention mechanisms [65, 72, 62], noisy web data [37], novel training losses [14], and pairwise category information [17]. Orthogonal to those works, we propose a spatio-temporal prior that can be combined with the probabilistic predictions of any image classifier to improve the final classification performance.

Location and Classification

A small number of approaches have explored the use of location information to improve image classification at test time. Berg *et al.* [5] proposed a spatio-temporal prior that when combined with the output of an image classifier increased the accuracy of bird species classification. Their approach discretized location and time into spatio-temporal cubes and used an adaptive kernel density estimator to represent the distribution of each species independently. Also in the context of predicting the presence of different biological species, Wittich *et al.* [64] evaluated different nearest neighbor based lookup strategies for retrieving the most relevant instances from a training set of geo-tagged observations. These approaches are inefficient in terms their memory requirements as they necessitate storing either the entire training set or a discretized version of it. Existing repositories of citizen science data (*e.g.* [53, 2, 1]) can contain on the order of tens of millions of observations making them prohibitively large to store and retrieve on mobile devices. Choosing the correct discretization is challenging [48], and incorrect choices can significantly affect the final performance [38, 46]. A key benefit of our approach is that discretization is not required.

Tang *et al.* [55] explored different feature encodings for incorporating location information directly into deep neural

networks at training time. This included raw location features (*i.e.* longitude and latitude), demographic information collected via a census, user provided hash-tags, and geographical map features (*e.g.* land use estimates). The disadvantage of their method is that it assumes that location information is present at test time and that all the required features can be computed for a given test location. Furthermore, they cannot use location information that does not have an associated image. They also need to retrain their entire model if new location data is collected. We instead propose an efficient spatio-temporal prior that jointly models the spatial distribution of multiple object categories that can be trained independently of the image classifier. Parallel to our work, [13] builds on [55] by exploring different ways to integrate location information into deep image classifiers.

Spatio-Temporal Distribution Modelling

Our goal is to estimate the spatio-temporal distribution of a set of object categories. Related to this, there is a rich literature exploring models for estimating the distribution of biological specimens across geographic space and time [31]. This is referred to as species distribution modelling or environmental niche modelling. Broadly, these methods can be divided into two groups, those that use *presence-absence* data and those that use *presence-only* data [28].

Making a presence-absence observation at a given location requires that every species from a predefined set of interest be confirmed as either present or absent for that sampling event. In practice, this kind of data is onerous to collect because it requires intense survey effort to confirm that a species is absent with a high degree of certainty [43]. However, once this data is collected it can be combined with standard supervised classification approaches such as logistic regression [28], probit regression [52], Gaussian processes [26], decision trees [20], and neural networks [67, 49, 45], among others [18, 47]. Presence-absence data is also compatible with traditional multi-label learning [34, 8, 69, 12, 63]. Recently deep models have been applied to this problem in order to jointly model the location preferences of different species [27, 11, 22, 56, 7] and human sampling biases [10].

In contrast, a presence-only (*i.e.* incidental) observation may be recorded wherever an object of interest is encountered - *without* requiring any absences to be verified. While presence-only data can be much easier to collect, the lack of absence information makes it more difficult to model. This limitation is typically dealt with in one of three different ways. The first approach is to generate ‘pseudo-negatives’ and then apply one of the presence-absence approaches from above. As no true negative information is available, these approaches randomly sample a set of locations and make the assumption that these locations are absences *e.g.* [19, 50, 3]. The second commonly used ap-

proach is to train a highly regularized model directly on the presence-only data *e.g.* by fitting a maximum entropy distribution [51] or a low-rank model [21], forcing the model to explain data where it has been observed and to be uncertain elsewhere. Finally, and most related to our work, there are approaches that use additional information such as the detectability of a given species and a photographer’s propensity to image them *e.g.* [44, 24].

Unlike many of the classic approaches for spatio-temporal distribution modelling, in this work we jointly learn a continuous spatio-temporal prior for each category of interest using a neural network to amortize the computation. In contrast to previous deep distribution models *e.g.* [27, 11, 56], we do not require presence-absence data or additional environmental features as input. We instead exploit the structure that exists in online image repositories, such as those collected by citizen scientists, to jointly model objects, their locations, and photographer biases.

3. Methods

Here we outline our spatio-temporal prior, which models the geographical and temporal distribution of a set of object categories and photographers. During training we assume that we have access to a set of tuples $\mathcal{D} = \{(I_i, \mathbf{x}_i, y_i, p_i) \mid i = 1, \dots, N\}$, where I_i is an image, $y_i \in \{1, \dots, C\}$ is the corresponding class label, $\mathbf{x}_i = [\text{lon}_i, \text{lat}_i, \text{time}_i]$ represents the location (longitude and latitude) and time the image was taken, and p_i is the individual, *i.e.* photographer, who captured the image. Note that the location does not need to be captured alongside the image. \mathcal{D} can be assembled from unrelated image and location datasets as long as both contain the same categories.

At test time, given an image and where and when it was taken we aim to estimate which category it contains *i.e.* $P(y|I, \mathbf{x})$. One approach is to model the joint distribution $P(I, \mathbf{x})$ as in [55], but this necessitates that the location information is *always* available at test time. Instead, inspired by [5], we can incorporate location information as a Bayesian spatio-temporal prior. If we assume that I and \mathbf{x} are conditionally independent given y , then

$$P(y|I, \mathbf{x}) = \frac{P(I, \mathbf{x}|y)P(y)}{P(I, \mathbf{x})} \quad (1)$$

$$= \frac{P(I)P(\mathbf{x})}{P(I, \mathbf{x})} \frac{P(y|I)P(y|\mathbf{x})}{P(y)} \quad (2)$$

$$\propto P(y|I)P(y|\mathbf{x}), \quad (3)$$

where we assume a uniform prior $P(y) = 1/C$ for $y \in \{1, \dots, C\}$. In reality an image may contain location information unrelated to the class label (*e.g.* the background), but we assume this factorization is valid. By factoring the distribution in this way we can represent the image classifier, $P(y|I)$, and spatio-temporal prior, $P(y|\mathbf{x})$, separately. Note

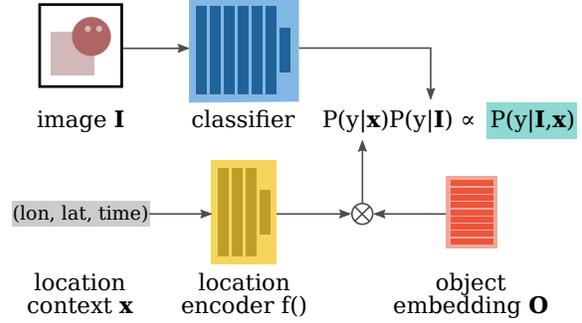


Figure 2: **Inference time.** Our goal is to estimate if an object category y is present in an input image I . At test time we make use of additional spatio-temporal information \mathbf{x} in the form of where and when the image was taken.

that at test time we do not assume that we have any knowledge of the individual p who captured the image. In this work we focus our attention on representing $P(y|\mathbf{x})$. For $P(y|I)$ we can use any discriminative model that produces a probabilistic output *e.g.* a convolutional neural network.

Presence-Absence Loss

As we are modeling the spatio-temporal prior independently from the image classifier our training data is now of the form $\mathcal{D} = \{(\mathbf{x}_i, y_i, p_i) \mid i = 1, \dots, N\}$. In the ideal case we would have complete information consisting of where and when a given category has both been observed to be present and observed to be *not* present *e.g.* as in [11, 56]. Then instead of $y_i \in 1, \dots, C$, each spatio-temporal location \mathbf{x}_i would be associated with a binary multi-label vector $\mathbf{y}_i = [y_i^1, \dots, y_i^C]$ where each entry $y_i^c \in \{0, 1\}$ indicates whether or not category c has been observed as being present at \mathbf{x}_i . This formulation results in a standard multi-label learning problem, enabling us to estimate the parameters of the spatio-temporal model by solving

$$\max_{\theta} \sum_{i=1}^N \sum_{c=1}^C y_i^c \log(\hat{y}_i^c) + (1 - y_i^c) \log(1 - \hat{y}_i^c), \quad (4)$$

where we define $\hat{y}_i^c = P(y_i^c|\mathbf{x}_i)$ and P is parameterized by θ . However, as discussed previously, presence-absence information is both difficult and time consuming to acquire in real world settings.

Presence-Only Loss

In this work we explore the more challenging presence-only setting where each spatio-temporal location \mathbf{x}_i is associated with a single label $y_i \in \{1, \dots, C\}$ indicating which category was observed. In essence, we have a label vector \mathbf{y}_i where there is only one affirmative entry, *i.e.* $y_i^c = 1$ for

some c , and the remaining entries are unknown. In this setting, Eqn. 4 can be written as

$$\max_{\theta} \sum_{i=1}^N \log(\hat{y}_i^{c_i}) + A_i, \quad (5)$$

where A_i represents a proxy absence term for the i^{th} training example and c_i is the corresponding observed category. Now the question becomes how to choose A_i .

One common approach for representing A_i is to generate ‘pseudo-negatives’ [3] by randomly sampling absence data from some parametric distribution. For instance, one might set

$$A_i = \log(1 - P(y_i | \mathbf{r}_i)). \quad (6)$$

where \mathbf{r}_i is a randomly selected spatio-temporal location with $[\text{lon}(\mathbf{r}_i), \text{lat}(\mathbf{r}_i)] \sim \text{Unif}(\mathbb{S}^2)$ and $\text{time}(\mathbf{r}_i) \sim \text{Unif}([0, 1])$. The implicit assumption is that each category (whether man-made or naturally occurring) occurs in a relatively small subset of $\mathbb{S}^2 \times [0, 1]$, so the probability of a category occurring at a randomly chosen location $\mathbf{r} \in \mathbb{S}^2 \times [0, 1]$ is small as well. To the extent that this assumption holds, these pseudo-negatives are likely to be valid.

An alternative approach is to instead sample absences over locations and times where the presence data for other categories occurs. In this case we would set A_i according to Eqn. 6 but sample negative locations from the positive occurrence locations *i.e.* $\mathbf{r}_i \sim \text{Unif}(\{\mathbf{x}_1, \dots, \mathbf{x}_N\})$. This biases the training towards regions that contain valid data.

3.1. Our Approach

In this section we outline how we model and train our spatio-temporal prior $P(y|\mathbf{x})$.

Location and Object Embedding

In many contexts, different objects do not occur independently at a given spatio-temporal location. Knowing that object A is present may provide information regarding the presence or absence of object B at the same place and time. Similarly, different spatio-temporal locations are not independent, and may share commonalities. We exploit this structure to encode low dimensional embeddings of objects and spatio-temporal locations.

Taking inspiration from [11], we model our spatio-temporal prior as $P(y|\mathbf{x}) \propto s(f(\mathbf{x})\mathbf{O})$. Here, $f : \mathcal{R}^3 \rightarrow \mathcal{R}^D$ is a multi-layered fully-connected neural network that maps a spatio-temporal location \mathbf{x} to a D dimensional embedding vector. $\mathbf{O} \in \mathcal{R}^{D \times C}$ represents an object embedding matrix, where each column is a different category. The product $f(\mathbf{x})\mathbf{O}$ results in a C dimensional vector, where each element represents the affinity that a spatio-temporal location \mathbf{x} has for category y . The intuition is that we are representing spatio-temporal locations and object categories

in a shared embedding space where the inner product between the embedding of a location \mathbf{x} and an object y is large if y is likely to occur at location \mathbf{x} . Finally, $s()$ is an entry-wise sigmoid operation to ensure that the resulting prediction are in the range $[0, 1]$.

Photographer Embedding

In online image collections we often have access to additional information at training time in the form of the photographer $p \in \mathcal{P}$ who captured the image. To see why this information is valuable, consider the following example. Suppose a photographer p visits location \mathbf{x} and does *not* report object y . If p has never taken an image of an object like y , then this non-report gives us little information. However, if p has a history of reporting categories similar to object y , then this constitutes weak evidence that y might actually be absent at that location. Thus, we can interpret the same presence-only information in different ways depending on the individual who provides it.

To capture photographer biases, we embed photographers into the same shared embedding space as the objects and locations. This is achieved by learning a photographer embedding matrix $\mathbf{P} \in \mathcal{R}^{D \times |\mathcal{P}|}$ at training time. Like different object categories, photographers may have affinities for particular locations and times, and share similarities in their spatio-temporal patterns with other photographers. This enables us to represent both a photographer’s preference for a given location $P(p|\mathbf{x}) \propto s(f(\mathbf{x})\mathbf{P})$, and a photographer’s affinity for a given object category $P(y|p) \propto s(\mathbf{O}^T \mathbf{P})$. Once trained, the photographer embeddings \mathbf{P} are not required at test time, see Fig. 2.

Joint Embedding Loss

Our goal at training time is to estimate the set of parameters $\theta = [\theta_f, \mathbf{O}, \mathbf{P}]$, where θ_f denotes the weights of the location embedding network $f()$, \mathbf{O} is the category embedding matrix, and \mathbf{P} is the photographer embedding matrix.

We start with the constraint that our model should be conservative *i.e.* if a category y has been observed at the spatio-temporal location \mathbf{x} in the training set, then $s(f(\mathbf{x})\mathbf{O}_{:,y})$ should be close to 1, otherwise it should be close to 0. Here, $\mathbf{O}_{:,y}$ indicates the y^{th} column of \mathbf{O} . We rely on the location embedding function $f()$ to interpolate between presence locations. This is conservative in the sense that it assumes that an object is absent if it has not been observed. This is a very strong assumption, but it enables the spatio-temporal prior to be aggressive in downweighting incorrect predictions from the image classifier.

Our first loss encourages the model to predict the presence of objects where they have been observed in the training set and downweight their likelihood where they have not

been observed:

$$\begin{aligned} \mathcal{L}_{o.loc}(\mathbf{x}, \mathbf{r}, \mathbf{O}, y) = & \lambda \log(s(f(\mathbf{x})\mathbf{O}_{:,y})) + \\ & \sum_{\substack{i=1 \\ i \neq y}}^C \log(1 - s(f(\mathbf{x})\mathbf{O}_{:,i})) + \\ & \sum_{i=1}^C \log(1 - s(f(\mathbf{r})\mathbf{O}_{:,i})). \end{aligned} \quad (7)$$

λ is a hyperparameter used to weight the positive observations and \mathbf{r} is a uniformly random spatio-temporal datapoint. Next, we want the affinity between a photographer p and a location \mathbf{x} be high if p was present at \mathbf{x} , and low otherwise:

$$\mathcal{L}_{p.loc}(\mathbf{x}, \mathbf{r}, \mathbf{P}, p) = \log(s(f(\mathbf{x})\mathbf{P}_{:,p})) + \log(1 - s(f(\mathbf{r})\mathbf{P}_{:,p})). \quad (8)$$

We assume that a photographer has a low affinity for a category unless they have previously observed it:

$$\begin{aligned} \mathcal{L}_{p.o}(\mathbf{O}, \mathbf{P}, y, p) = & \lambda \log(s(\mathbf{O}_{:,y}^T \mathbf{P}_{:,p})) + \\ & \sum_{\substack{i=1 \\ i \neq y}}^C \log(1 - s(f(\mathbf{O}_{:,i}^T \mathbf{P}_{:,p}))). \end{aligned} \quad (9)$$

Finally, to estimate the parameters of our prior we maximize

$$\mathcal{L} = \mathcal{L}_{o.loc} + \mathcal{L}_{p.loc} + \mathcal{L}_{p.o}, \quad (10)$$

by iterating over each of the datapoints in the training set.

4. Experiments

We evaluate the effectiveness of our spatio-temporal prior by performing experiments on several image classification datasets that have location and time information. We choose image classification because for other domains (*e.g.* species distribution modeling) it is challenging to obtain accurate ground truth information regarding the true spatio-temporal distributions of the categories of interest.

4.1. Datasets

While location metadata is readily available for online image collections, many popular image classification datasets do not contain this information *e.g.* [61, 57, 16, 39]. Some datasets exist with location information, but for only a subset of the images *e.g.* [25]. However, datasets containing images of different species of plants and animals are available with location, time, and photographer information. To this end, we perform experiments on the iNaturalist 2017 and 2018 (iNat2017 and iNat2018) species classification datasets which contain images collected and annotated by citizen scientists [58]. They have 5,089 and 8,142 categories respectively. While [5] evaluated their location

prior on the BirdSnap dataset, the images and location metadata used are not provided by the authors. We recollect the images and location data from the web using the original image URLs. Despite the dataset consisting of images of species commonly found in North America, when we recollect the images and locations we found that the original images are from all over the world and 40% were missing location. Like [5], we also simulate location metadata for BirdSnap [5] and another fine-grained dataset of birds, NABirds [57], by associating each image with a species observation from eBird [53]. Our train locations and photographers are sampled from eBird 2015, and the test set is from 2016. BirdSnap and NABirds contain images from 500 and 555 different species of North America birds. Finally, we also perform experiments on YFCC100M-GEO100 [55] (YFCC). YFCC contains 100 everyday object categories with associated locations, but no date or photographer information is provided. The train and test split used in [55] is not available and so we created a new one. Unlike the other datasets, many of the object categories in YFCC are not geographically distinct *e.g.* ‘band’, ‘ford’, or ‘ipod’.

4.2. Implementation Details

Our location encoder $f()$ is a fully-connected neural network consisting of an input layer, followed by multiple residual layers [30], and a final output embedding layer. We jointly train the location encoder, along with the photographer and object embeddings using Adam [36] for 30 epochs with a batch size of 1024, using dropout to prevent overfitting. The dimensionality of the shared embedding space is set to $D = 256$. When weighting the positive instances during training we set λ to the number of categories. To counteract the heavily imbalanced nature of many of the datasets, we limit the maximum number of datapoints for each category per epoch. We set the maximum number of datapoints to 100, and for each epoch we randomly select a different subset for each category. The only exception is for YFCC, where capping the data hurt performance. Details of our network architecture are in the supplementary material.

Except where noted, at test time, our model takes three inputs – longitude, latitude, and day of the year, specifying where and when the image of interest was captured. For these three input features \mathbf{x} we explored different methods for ‘wrapping’ the coordinates *i.e.* an observation taken on December 31st should result in a similar embedding to one captured on January 1st. Similarly, we want geographical coordinates to wrap around the earth. To achieve this, for each input dimension l of \mathbf{x} we perform the mapping $[\sin(\pi x^l), \cos(\pi x^l)]$, resulting in two numbers for each dimension. Here, we assume that each dimension of the input has been normalized to the range $x^l \in [-1, 1]$.

For the image classifiers $P(y|I)$ we fine-tune a separate InceptionV3 [54] network for each of the datasets beginning with ImageNet initialized weights [16] with an image

$P(y \mathbf{x})$ - Prior Type	YFCC	BirdSnap	BirdSnap†	NABirds†	iNat2017			iNat2018		
	Test	Test	Test	Test	Val	Test Pu	Test Pr	Val	Test Pu	Test Pr
No Prior (<i>i.e.</i> uniform)	50.15	70.07	70.07	76.08	63.27	64.16	63.63	60.20	50.17	50.33
Nearest Neighbor (num)	51.78	70.82	77.76	79.99	65.34	66.04	65.61	68.70	54.54	54.58
Nearest Neighbor (spatial)	51.21	71.57	77.98	80.79	65.85	67.02	66.41	67.55	53.67	53.81
Discretized Grid	51.06	71.09	77.19	79.58	65.49	66.62	66.07	67.27	53.13	53.16
Adaptive Kernel [5]	51.47	71.57	78.65	81.11	64.86	65.83	65.59	65.23	53.17	53.21
Tang <i>et al.</i> [55]	50.43	70.16	72.33	77.34	66.15	67.08	66.53	65.61	54.12	54.25
Ours no date	50.70	71.66	78.65	81.15	69.34	70.62	70.18	72.41	57.68	57.84
Ours full	-	71.84	79.58	81.50	69.60	70.83	70.51	72.68	58.44	58.59

Table 1: **Classification accuracy.** Results after combining image classification predictions $P(y|I)$ with different spatio-temporal priors $P(y|\mathbf{x})$. All results are top 1 accuracy with classifier predictions extracted from an InceptionV3 [54] network fine-tuned on each of the respective datasets. † indicates that simulated locations, dates, and photographers from the eBird dataset [53] are used. The baseline algorithms do not use date information.

	Top1	Top3	Top5
iNat2017 - InceptionV3 299 × 299			
No Prior (<i>i.e.</i> uniform)	63.27	79.82	84.51
Ours no wrap encode	69.48	84.43	88.15
Ours no photographer	69.39	83.97	87.71
Ours no date	69.34	84.16	87.89
Ours full	69.60	84.41	88.07
iNat2018 - InceptionV3 299 × 299			
No Prior (<i>i.e.</i> uniform)	60.20	77.90	83.29
Ours no wrap encode	72.12	87.00	90.52
Ours no photographer	72.84	87.30	90.75
Ours no date	72.41	87.19	90.60
Ours full	72.68	87.26	90.79
iNat2018 - InceptionV3 520 × 520			
No Prior (<i>i.e.</i> uniform)	66.18	83.32	88.04
Ours no wrap encode	77.09	90.68	93.54
Ours no photographer	77.64	90.82	93.52
Ours no date	77.41	90.80	93.58
Ours full	77.49	90.85	93.57

Table 2: **Ablation.** Classification accuracy for different variants of our prior on the iNat2017 and iNat2018 [58] validation sets. In the case of iNat2018, we still observe improvements when combining our prior with a more powerful image classifier - see rows ‘InceptionV3 520 × 520’.

resolution of 299 × 299 (unless otherwise noted).

4.3. Quantitative Evaluation

In Table 1 we evaluate how much our spatio-temporal prior improves image classification performance by comparing it to several baselines. We found that adding a uniform prior to the outputs of the nearest neighbor based baselines increases their performance. This adds robustness in cases where there are no objects from the training set present near the test locations. The lack of this uniform prior explains the poor results for nearest neighbor based approaches in [55]. For the comparison to Tang *et al.* [55], we jointly train a linear layer to embed the raw location information along with an output layer to combine the location embedding with the features from the last linear layer of the image classifier. The rest of the weights of the image classifier are not updated. For each of the baseline al-

gorithms we select their hyperparameters (*e.g.* the number of neighbors) on a held out validation set for each dataset. When location information is not available at test time, we assume a uniform prior over the categories.

Our model performs on par, or better, than the baselines across all datasets. The advantage of our approach is that it is computationally efficient at test time and does not require features from the image classifier during training. Compared to nearest neighbor based methods, it only requires a forward pass through a compact fully-connected neural network. In addition, it also captures structural information such as object and photographer biases. One failure case that is worth noting are the results on YFCC [55]. We observe that all methods perform similar to using no location information (No Prior). This can be explained by the relative lack of spatio-temporal structure in the object categories present in the dataset. Again, this is consistent with the findings in [55], where the authors had to use additional features to increase the performance.

4.3.1 Ablation Study

In Table 2 we compare the performance of different variants of our model on iNat2017 and iNat2018 [58]. Again, across all metrics there is a large increase in performance compared to the baseline uniform prior. In some cases, we even observe that there is an additional boost in performance when we explicitly model photographer biases.

Training fine-grained image classifiers with larger input images can significantly increase classification performance [15]. We observe that the benefit of our spatio-temporal prior is still apparent even when we use a more powerful classifier that has been training for longer with larger images. This increase in accuracy is also present when we evaluate performance using more lenient evaluation metrics *i.e.* top 5 vs. top 1 accuracy. This is significant because it highlights that for some datasets the performance boost provided by the spatio-temporal prior is orthogonal to improvements in the underlying image classifier.

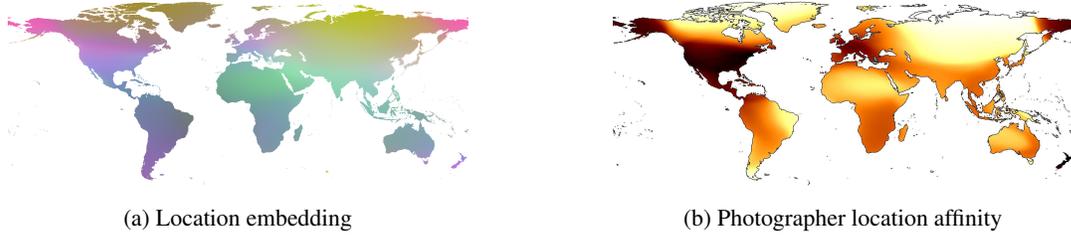


Figure 3: **Spatial predictions.** (a) Embeddings for each location on the earth for a model trained on iNat2018 [58]. We observe that the embeddings appears to capture information related to climate zones, despite not being trained on any climate data. (b) Log plot of estimated photographer location preferences. Darker colors indicate that more photographers have captured images in those locations. We can see that there is a large bias towards North America, Europe, and New Zealand.

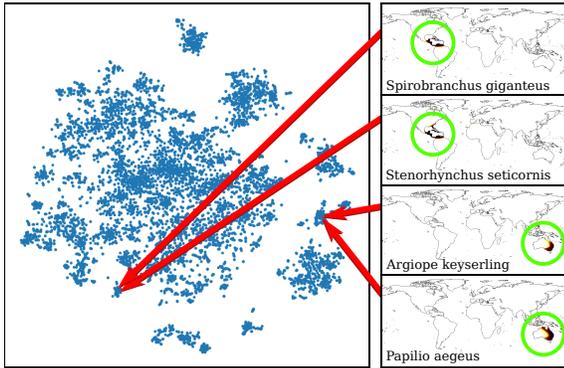


Figure 4: **Object embedding.** t-SNE [42] plot of the learned embedding \mathbf{O} for all 8,142 categories from iNat2018 [58]. The location in the object embedding space encodes a category’s preferences for a particular geographical region. We observe that categories that have similar spatio-temporal distributions tend to be close.

4.4. Qualitative Evaluation

Our model captures the relationship between objects, locations, and photographers. In Fig. 3 (a) we can see the resulting embeddings for each input location from our model trained on iNat2018 [58]. By applying the embedding function $f(\cdot)$ to each location we can generate its D dimensional embedding vector. We then use ICA [33] to project the embedded features to a three dimensional space and mask out the ocean for visualization. Perhaps as expected, there is low frequency structure in the resulting image *i.e.* nearby locations tend to support similar objects. One advantage of our approach is that we are not restricted to a fixed discretization. As a result we can generate embeddings for any location and time. In Fig. 4 we visualize our learned object embedding \mathbf{O} . Objects that have similar spatio-temporal distributions tend to result in similar embedding vectors.

Distinct from other work, our prior also models the relationship between photographers and locations, and photographers and object categories. In Fig. 3 (b) we plot the

estimated affinity for each input location across all photographers *i.e.* $\sum_p s(f(\mathbf{x})\mathbf{P}_{:,p})$. We only show results for photographers who provided at least 100 observations in the iNat2018 [58] training set, resulting in 634 individuals. In Fig. 5 we display the estimated affinity for each object category for a set of photographers *i.e.* $P(y|p) \propto s(\mathbf{O}^T\mathbf{P})$. We observe that the embedding captures the similarity in object affinity held by different photographers.

Finally, in Fig. 6 we use our prior to generate spatio-temporal predictions for several different species from iNat2018 [58]. Each image is generated by querying every location on the surface of the earth, on a specified day of the year, to generate $P(y = y^*|\mathbf{x})$ for the category of interest. In practice, we evaluate 1000×2000 spatial locations for each time point (*e.g.* first day of the month). This step is very efficient as we can pre-compute $f(\mathbf{x})$ for every location, independent of the category of interest. Again, for visualization we mask out the predictions over the ocean.

4.5. Limitations

We are limited by the quality of the provided location data *e.g.* it can be inaccurate or intentionally obfuscated. We also make strong assumptions about a photographer’s affinity for an individual object category. In reality, these interactions may be complex *i.e.* once a photographer captures an image of a particular category they may be less likely to take an image of the same object in the near future. There are also known spatial biases in the types of citizen science data we use [4, 10]. However, this may not be a major issue as we can assume that the distribution of test locations and dates is similarly biased. We currently only use location, time, and photographer ID during training. In practice, additional data such as environmental variables may be a valuable signal for specific object categories [6].

5. Conclusion

We introduce a spatio-temporal prior to help disambiguate fine-grained categories resulting in improved test time image classification performance. In addition to helping image classification, our model also naturally captures

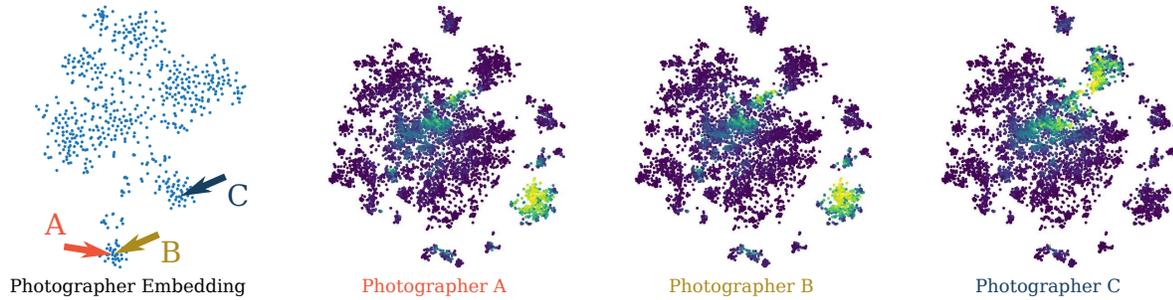


Figure 5: **Photographer object affinity.** On the left we see a t-SNE [42] plot of the photographer embeddings \mathbf{P} for iNat2018 [58]. The three plots on the right depict the predicted affinities for three different photographers (A, B, and C) visualized on the category embedding from Fig. 4. Brighter colors indicate a higher affinity for a given category. We observe that individuals that are close in the photographer embedding space \mathbf{P} (e.g. A and B) have similar category affinities, compared to those that are far away (e.g. C).

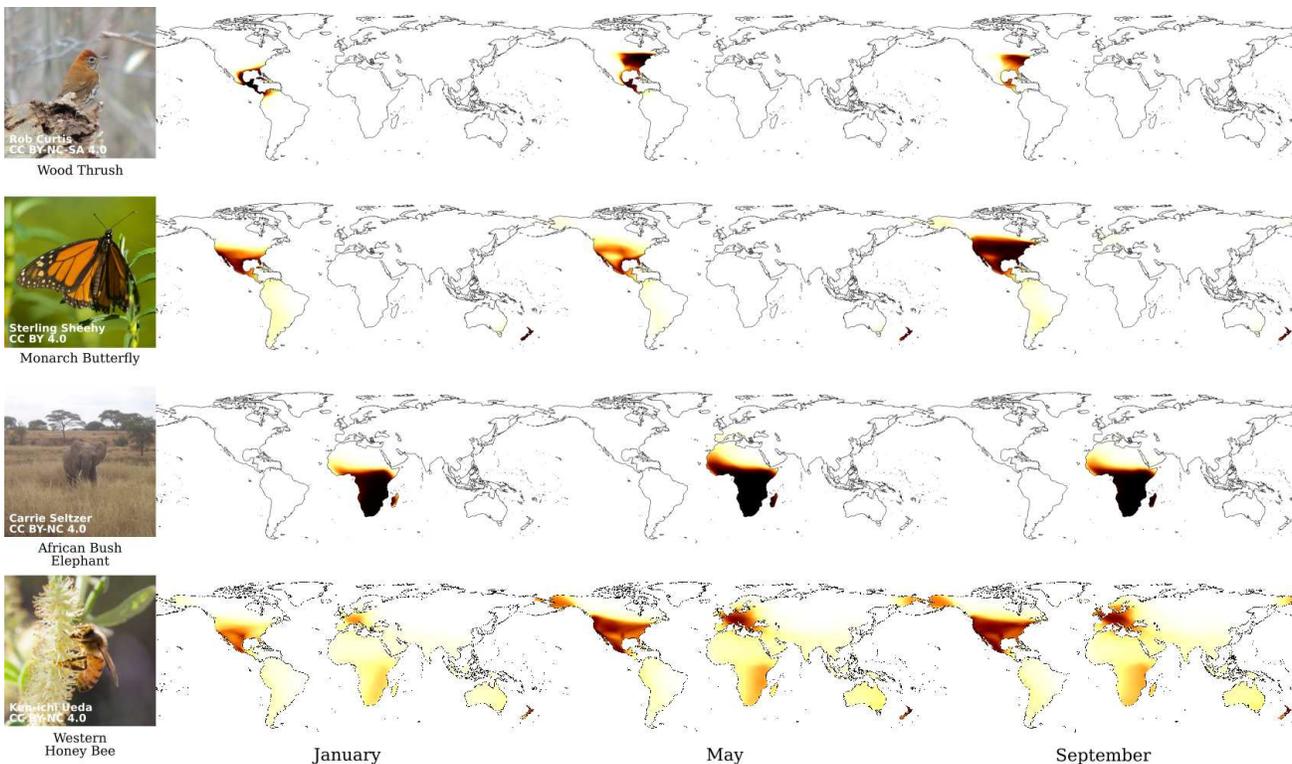


Figure 6: **Spatio-temporal predictions.** Predicted distributions for several object categories for three different time points using our full model trained on iNat2018 [58]. Darker colors indicate locations where the categories are predicted to be found. In the first two rows we observe that our model captures seasonal migratory behaviors. On the bottom row, our model correctly predicts that the Western Honey Bee can be found on several different continents. It is worth noting that the results are affected by geographical sampling biases in the iNat2018 dataset.

the relationships between locations and objects, objects and objects, photographers and objects, and photographers and locations in an interpretable manner. Importantly, our prior is efficient at test time, both in terms of model size and inference speed, and scales to large numbers of categories.

Acknowledgements This work was supported by a Google Focused Research Award and an NSF Graduate Research Fellowship (Grant No. DGE1745301). We thank Grant Van Horn and Serge Belongie for helpful discussions, along with NVIDIA and AWS for their kind donations.

References

- [1] GBIF - www.gbif.org. 2019.
- [2] iNaturalist - www.inaturalist.org. 2019.
- [3] Morgane Barbet-Massin, Frédéric Jiguet, Cecile Helene Albert, and Wilfried Thuiller. Selecting pseudo-absences for species distribution models: how, where and how many? *Methods in Ecology and Evolution*, 2012.
- [4] Jan Beck, Marianne Böller, Andreas Erhardt, and Wolfgang Schwanghart. Spatial bias in the GBIF database and its effect on modeling species' geographic distributions. *Ecological Informatics*, 2014.
- [5] Thomas Berg, Jiongxin Liu, Seung Woo Lee, Michelle L Alexander, David W Jacobs, and Peter N Belhumeur. Birdsnap: Large-scale fine-grained visual categorization of birds. In *CVPR*, 2014.
- [6] Christophe Botella, Pierre Bonnet, François Munoz, Pascal Monestiez, and Alexis Joly. Overview of GeoLifeCLEF 2018: location-based species recommendation. 2018.
- [7] Christophe Botella, Alexis Joly, Pierre Bonnet, Pascal Monestiez, and François Munoz. A deep learning approach to species distribution modelling. In *Multimedia Tools and Applications for Environmental & Biodiversity Informatics*. 2018.
- [8] Matthew R. Boutell, Jiebo Luo, Xipeng Shen, and Christopher M. Brown. Learning multi-label scene classification. *Pattern Recognition*, 2004.
- [9] Steve Branson, Grant Van Horn, Serge Belongie, and Pietro Perona. Bird species categorization using pose normalized deep convolutional nets. In *BMVC*, 2014.
- [10] Di Chen and Carla P Gomes. Bias reduction via end-to-end shift learning: Application to citizen science. In *AAAI*, 2019.
- [11] Di Chen, Yexiang Xue, Shuo Chen, Daniel Fink, and Carla Gomes. Deep multi-species embedding. In *IJCAI*, 2017.
- [12] Yao-nan Chen and Hsuan-tien Lin. Feature-aware label space dimension reduction for multi-label classification. In *NeurIPS*, 2012.
- [13] Grace Chu, Brian Potetz, Weijun Wang, Andrew Howard, Yang Song, Fernando Brucher, Thomas Leung, and Hartwig Adam. Geo-aware networks for fine grained recognition. *arXiv*, 2019.
- [14] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *CVPR*, 2019.
- [15] Yin Cui, Yang Song, Chen Sun, Andrew Howard, and Serge Belongie. Large scale fine-grained categorization and domain-specific transfer learning. In *CVPR*, 2018.
- [16] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [17] Abhimanyu Dubey, Otkrist Gupta, Pei Guo, Ramesh Raskar, Ryan Farrell, and Nikhil Naik. Pairwise confusion for fine-grained visual classification. In *ECCV*, 2018.
- [18] Jane Elith and Catherine H. Graham. Do they? how do they? why do they differ? on finding reasons for differing performances of species distribution models. *Ecography*, 2009.
- [19] Robin Engler, Antoine Guisan, and Luca Rechsteiner. An improved approach for predicting the distribution of rare and endangered species from occurrence and pseudo-absence data. *Journal of Applied Ecology*, 2004.
- [20] Daniel Fink, Wesley M Hochachka, Benjamin Zuckerberg, David W Winkler, Ben Shaby, M Arthur Munson, Giles Hooker, Mirek Riedewald, Daniel Sheldon, and Steve Kelling. Spatiotemporal exploratory models for broad-scale survey data. *Ecological Applications*, 2010.
- [21] William Fithian and Rahul Mazumder. Flexible low-rank statistical modeling with missing data and side information. *Statistical Science*, 2018.
- [22] Simone Franceschini, Emanuele Gandola, Marco Martinoli, Lorenzo Tancioni, and Michele Scardi. Cascaded neural networks improving fish species prediction accuracy: the role of the biotic information. *Scientific Reports*, 2018.
- [23] Yang Gao, Oscar Beijbom, Ning Zhang, and Trevor Darrell. Compact bilinear pooling. In *CVPR*, 2016.
- [24] Georgia E Garrard, Michael A McCarthy, Nicholas SG Williams, Sarah A Bekessy, and Brendan A Wintle. A general model of detectability using species traits. *Methods in Ecology and Evolution*, 2013.
- [25] Hervé Goëau, Pierre Bonnet, and Alexis Joly. Plant identification in an open-world (LifeCLEF 2016). In *CLEF: Conference and Labs of the Evaluation Forum*, 2016.
- [26] Nick Golding and Bethan V. Purse. Fast and flexible bayesian species distribution modelling using gaussian processes. *Methods in Ecology and Evolution*, 2016.
- [27] David J. Harris. Generating realistic assemblages with a joint species distribution model. *Methods in Ecology and Evolution*, 2015.
- [28] Trevor Hastie and Will Fithian. Inference from presence-only data; the ongoing controversy. *Ecography*, 2013.
- [29] James Hays and Alexei A Efros. IM2GPS: estimating geographic information from a single image. In *CVPR*, 2008.
- [30] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [31] Troy M. Hegel, Samuel A. Cushman, Jeffrey Evans, and Falk Huettmann. Current state of the art for statistical modelling of species distributions. In *Spatial Complexity, Informatics, and Wildlife Conservation*. 2010.
- [32] Shaoli Huang, Zhe Xu, Dacheng Tao, and Ya Zhang. Part-stacked CNN for fine-grained visual categorization. In *CVPR*, 2016.
- [33] Aapo Hyvärinen and Erkki Oja. Independent component analysis: algorithms and applications. *Neural networks*, 2000.
- [34] Julia Jones, Jeffrey Miller, and Matt White. Multi-label classification for multi-species distribution modeling. In *ICML*, 2011.
- [35] Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Fei-Fei Li. Novel dataset for fine-grained image categorization: Stanford dogs. In *CVPR Workshop on Fine-Grained Visual Categorization*, 2011.
- [36] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv*, 2014.

- [37] Jonathan Krause, Benjamin Sapp, Andrew Howard, Howard Zhou, Alexander Toshev, Tom Duerig, James Philbin, and Li Fei-Fei. The unreasonable effectiveness of noisy data for fine-grained recognition. In *ECCV*, 2016.
- [38] Alex M. Lechner, William T. Langford, Simon D. Jones, Sarah A. Bekessy, and Ascelin Gordon. Investigating species-environment relationships at multiple scales: Differentiating between intrinsic scale and the modifiable areal unit problem. *Ecological Complexity*, 2012.
- [39] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014.
- [40] Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. Bilinear CNN models for fine-grained visual recognition. In *ICCV*, 2015.
- [41] Jiongxin Liu, Angjoo Kanazawa, David Jacobs, and Peter Belhumeur. Dog breed classification using part localization. In *ECCV*, 2012.
- [42] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-SNE. *JMLR*, 2008.
- [43] Darryl I. MacKenzie. What are the issues with presence-absence data for wildlife managers? *The Journal of Wildlife Management*, 2005.
- [44] Darryl I MacKenzie, James D Nichols, Gideon B Lachman, Sam Droege, J Andrew Royle, and Catherine A Langtimm. Estimating site occupancy rates when detection probabilities are less than one. *Ecology*, 2002.
- [45] Sylvain Mastrorillo, Sovan Lek, Francis Dauba, and Alain Belaud. The use of artificial neural networks to predict the presence of small-bodied fish in a river. *Freshwater Biology*, 2003.
- [46] Justin Moat, Steven P. Bachman, Richard Field, and Doreen S. Boyd. Refining area of occupancy to address the modifiable areal unit problem in ecology and conservation. *Conservation Biology*, 2018.
- [47] Anna Norberg, Nerea Abrego, F Guillaume Blanchet, Frederick R Adler, Barbara J Anderson, Jani Anttila, Miguel B Araújo, Tad Dallas, David Dunson, Jane Elith, et al. A comprehensive evaluation of predictive performance of 33 species distribution models at species and community levels. *Ecological Monographs*, 2019.
- [48] Stan Openshaw. The modifiable areal unit problem. Geobooks, Norwich, England, 1983.
- [49] Stacy L. Özemesi and Uygur Özemesi. An artificial neural network approach to spatial habitat modelling with interspecific interaction. *Ecological Modelling*, 1999.
- [50] Steven J Phillips, Miroslav Dudík, Jane Elith, Catherine H Graham, Anthony Lehmann, John Leathwick, and Simon Ferrier. Sample selection bias and presence-only distribution models: implications for background and pseudo-absence data. *Ecological Applications*, 2009.
- [51] Steven J Phillips, Miroslav Dudík, and Robert E Schapire. A maximum entropy approach to species distribution modeling. In *ICML*, 2004.
- [52] Laura J. Pollock, Reid Tingley, William K. Morris, Nick Golding, Robert B. O'Hara, Kirsten M. Parris, Peter A. Veski, and Michael A. McCarthy. Understanding co-occurrence by modelling species simultaneously with a joint species distribution model. *Methods in Ecology and Evolution*, 2014.
- [53] Brian L Sullivan, Christopher L Wood, Marshall J Iliff, Rick E Bonney, Daniel Fink, and Steve Kelling. ebird: A citizen-based bird observation network in the biological sciences. *Biological Conservation*, 2009.
- [54] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, 2016.
- [55] Kevin Tang, Manohar Paluri, Li Fei-Fei, Rob Fergus, and Lubomir Bourdev. Improving image classification with location context. In *ICCV*, 2015.
- [56] Luming Tang, Yexiang Xue, Di Chen, and Carla P Gomes. Multi-entity dependence learning with rich context via conditional variational auto-encoder. In *AAAI*, 2018.
- [57] Grant Van Horn, Steve Branson, Ryan Farrell, Scott Haber, Jessie Barry, Panos Ipeirotis, Pietro Perona, and Serge Belongie. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *CVPR*, 2015.
- [58] Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The iNaturalist species classification and detection dataset. In *CVPR*, 2018.
- [59] Kumar Vishal, CV Jawahar, and Visesh Chari. Accurate localization by fusing images and GPS signals. In *CVPR Workshop*, 2015.
- [60] Nam Vo, Nathan Jacobs, and James Hays. Revisiting IM2GPS in the deep learning era. In *ICCV*, 2017.
- [61] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The Caltech-UCSD Birds-200-2011 Dataset. 2011.
- [62] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *CVPR*, 2017.
- [63] Jiang Wang, Yi Yang, Junhua Mao, Zhiheng Huang, Chang Huang, and Wei Xu. CNN-RNN: A unified framework for multi-label image classification. In *CVPR*, 2016.
- [64] Hans Christian Wittich, Marco Seeland, Jana Wäldchen, Michael Rzanny, and Patrick Mäder. Recommending plant taxa for supporting on-site species identification. *BMC Bioinformatics*, 2018.
- [65] Tianjun Xiao, Yichong Xu, Kuiyuan Yang, Jiaying Zhang, Yuxin Peng, and Zheng Zhang. The application of two-level attention models in deep convolutional neural network for fine-grained image classification. In *CVPR*, 2015.
- [66] Linjie Yang, Ping Luo, Chen Change Loy, and Xiaoou Tang. A large-scale car dataset for fine-grained categorization and verification. In *CVPR*, 2015.
- [67] Peggy PW Yen, Falk Huettmann, and Fred Cooke. A large-scale model for the at-sea distribution and abundance of Marbled Murrelets (*Brachyramphus marmoratus*) during the breeding season in coastal British Columbia, Canada. *Ecological Modelling*, 2004.

- [68] Menghua Zhai, Tawfiq Salem, Connor Greenwell, Scott Workman, Robert Pless, and Nathan Jacobs. Learning geotemporal image features. In *BMVC*, 2018.
- [69] Min-Ling Zhang and Zhi-Hua Zhou. Multi-label neural networks with applications to functional genomics and text categorization. *IEEE Transactions on Knowledge and Data Engineering*, 2006.
- [70] Ning Zhang, Jeff Donahue, Ross Girshick, and Trevor Darrell. Part-based R-CNNs for fine-grained category detection. In *ECCV*, 2014.
- [71] Ning Zhang, Ryan Farrell, Forrest Iandola, and Trevor Darrell. Deformable part descriptors for fine-grained recognition and attribute prediction. In *ICCV*, 2013.
- [72] Bo Zhao, Xiao Wu, Jiashi Feng, Qiang Peng, and Shuicheng Yan. Diversified visual attention networks for fine-grained object classification. *IEEE Transactions on Multimedia*, 2017.