

Generating Binary Species Range Maps

Filip Dorm¹ Christian Lange¹ Scott Loarie² Oisín Mac Aodha¹

¹University of Edinburgh ²iNaturalist

Abstract. Accurately predicting the geographic ranges of species is crucial for assisting conservation efforts. Traditionally, range maps were manually created by experts. However, species distribution models (SDMs) and, more recently, deep learning-based variants offer a potential automated alternative. Deep learning-based SDMs generate a continuous probability representing the predicted presence of a species at a given location, which must be binarized by setting per-species thresholds to obtain binary range maps. However, selecting appropriate per-species thresholds to binarize these predictions is non-trivial as different species can require distinct thresholds. In this work, we evaluate different approaches for automatically identifying the best thresholds for binarizing range maps using presence-only data. This includes approaches that require the generation of additional pseudo-absence data, along with ones that only require presence data. We also propose an extension of an existing presence-only technique that is more robust to outliers. We perform a detailed evaluation of different thresholding techniques on the tasks of binary range estimation and large-scale fine-grained visual classification, and we demonstrate improved performance over existing pseudo-absence free approaches using our method.

Keywords: Species Range Maps · Thresholding · SINR

1 Introduction

Accurate range maps, depicting where species are present and absent, are indispensable tools for conservation. Statistics derived from their ranges are some of the key information used to estimate the threatened status of different species [39]. However, traditional expert-derived range maps are costly and time consuming to create. To alleviate this, advancements in statistics and machine learning have enabled the development of species distribution models (SDMs) that can estimate the probability of a species being present, among other quantities, at different locations [2]. Recently, deep learning-based multi-species SDMs, trained on crowd-sourced observation data, have been demonstrated to be capable of predicting the presence of thousands of different species simultaneously for an input location of interest [4, 11, 12, 24, 30, 41]. However, to obtain a binary range map from these continuous probabilities, they need to be binarized using a threshold. Binary range maps are used in visualization in addition to being necessary for computing spatial metrics for determining the threatened status of a species.

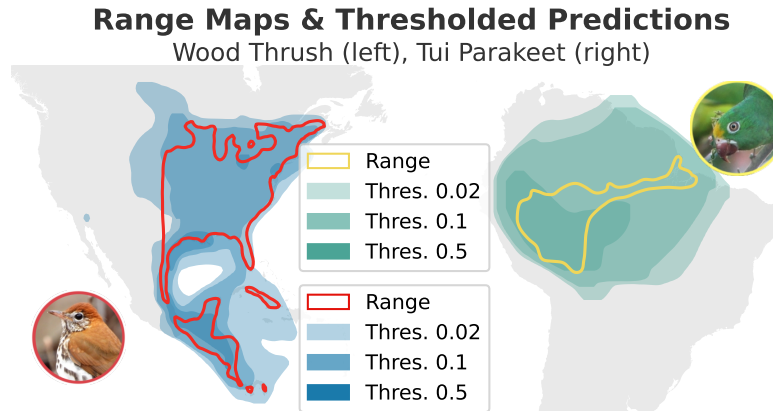


Fig. 1: Binary range maps for two different species. These ranges are generated by the SINR [11] species distribution model (SDM) for **Wood Thrush** (left) and **Tui Parakeet** (right), where the expert range map is denoted via solid outline. Converting the continuous SDM outputs to binary range maps requires setting thresholds (*e.g.* 0.02, 0.1, or 0.5) which result in very different range maps depending on the values chosen. More importantly, here the same threshold value is not the best for both species.

Finding the most effective threshold is a challenging problem, as the one that leads to the most accurate range maps will not necessarily be the same across different species (see Fig. 1). This is because the outputs of multi-species deep SDMs can be highly miscalibrated, which is likely the result of a number of factors such as imbalances in the species present in the training data, spatial biases, and issues stemming from species co-occurrences. The threshold selection task has been previously explored in the ecology literature under the setting where calibration data, which denotes a subset of locations where the species of interest is present or absent, is available [27]. However, the more challenging setting, and the one explored in this work, is where we do not have any information related to confirmed absences [29]. This presence-only data paradigm is most relevant to the opportunistic data available from large citizen science platforms such as iNaturalist [20], which contains over 200 million species observations to date.

A number of different thresholding approaches have been proposed for the case when only species presence observations are available [29]. However, many of these approaches require the generation of ‘pseudo-absences’, *i.e.* hallucinated data that simulates locations where a species is said to not be found. As a result, the precise design choices made when generating these pseudo-absences can have a large impact on the accuracy of the resulting binary range maps. In this work, we compare these approaches to alternative thresholding techniques that do not require absence data of any form. Furthermore, we outline a simple extension of these absence-free techniques that obtain superior performance. Our resulting approach is conceptually simple and more efficient than the more computationally expensive pseudo-absence-based alternatives.

We make the following contributions: (i) We perform a detailed evaluation of multiple different thresholding techniques for generating binary species range

maps. Our experiments are conducted using recent multi-species SDMs over a large set of expert-derived evaluation range data. (ii) We show that binary range maps can also be used as priors to improve large-scale fine-grained image classification and that they are superior to existing continuous priors. (iii) Finally, we introduce an extension of an existing thresholding approach that does not require any pseudo-absence data yet still outperforms existing presence-only methods on both range estimation and image classification tasks. Code to reproduce our experiments is available here: https://github.com/filipgdorm/binary_range_maps

2 Related Work

Species Distribution Models (SDMs). There is a large body of work on developing models to estimate the spatial and temporal distributions of species from raw observation data [2, 13]. These models are typically trained using one of two different types of data: presence-only (PO) or presence-absence data. PO data consists of a set of observations (*i.e.* locations in space and time where a species has been observed), without any confirmed absences. PO data is significantly easier to collect compared to presence-absence data, requiring only the recording of locations where a species has been observed, without the need for extensive, resource-intensive surveys to confirm absences. As a result, PO data is abundant and widely available through online citizen science initiatives such as iNaturalist [20].

However, a challenge in training SDMs with PO data is the need for negative signal (*i.e.* species’ absences) during training. Without having absences in the training data, SDMs can naively learn to simply predict presence, for every species, at every location. One common approach to overcome this issue is to create *pseudo*-absences (PAs), *i.e.* artificially generated data points representing species’ absences [1]. Various approaches for generating PAs have been proposed, with popular techniques including selecting locations uniformly at random (*i.e.* ‘random background points’) or locations where other species have been observed (*i.e.* ‘target background points’) [5, 37, 38]. The effectiveness of different methods for generating PAs for training multi-species deep SDMs has recently been explored in [43].

Recently, [11] introduced a deep learning-based approach for modeling species’ distributions, termed Spatial Implicit Neural Representations (SINRs). SINRs are parameterized as fully connected neural networks and can be trained on noisy crowd-sourced PO data. The authors demonstrated the ability to jointly estimate the ranges of thousands of species within a single model using PO training data. They introduced new benchmarks for evaluating the accuracy of predicted range maps by comparing model outputs to expert-derived ranges. They also quantified the improvement in fine-grained visual classification for different species when using SDMs as priors to modify image classifier predictions based on the location where the image was captured [30]. However, one of the limitations of [11]’s range estimation evaluation is that it uses a continuous, *i.e.* area under the curve, metric. As a result, the task of *binary* range estimation is not directly evaluated as the evaluation protocol is agnostic to the choice of

threshold used to produce a binary range map. We use the models and data from [11] for our evaluation and demonstrate that the threshold required can vary significantly across species and how to select the best one is not obvious.

Binarizing Range Maps. The output of SDMs such as [11] are a set of continuous predictions representing species’ presences at the evaluated locations. Binarizing these predictions to create binary range maps requires setting thresholds, which can result in very different maps depending on the thresholds selected. Furthermore, as seen in Fig. 1, the choice of best threshold can be species dependent. Without true absence data for calibration [15, 27], the problem of converting continuous outputs to binary ranges, given only PO data, is challenging.

The problem of selecting appropriate thresholds for SDMs has been previously explored in the ecology literature [27–29]. However, these existing works tend to focus on an older generation of single species machine learning methods such as Maxent [36], Boosted Regression Trees [14], and Random Forests [6]. Additionally, these experiments are often limited to a small number of species, sometimes using only simulated data [29], and constrained to smaller geographic regions. Previous work has shown that fixed constant thresholds (*e.g.* 0.5) do not work well when the training data is imbalanced [17, 22]. In the context of PO data, the impact of different thresholding techniques has been evaluated in [28, 29]. Existing techniques can be broadly categorized based on the data they require, *e.g.* no data, presences, or presences and PAs. When no data is available, approaches such as choosing a fixed threshold or simply setting the threshold to be the average of the predictions can be used. Alternatively, if presence data is available, one approach is to choose the lowest/minimum presence threshold based on the SDM predictions associated with the presence observations in the training data [34, 36]. In essence, this approach attempts to maximize recall by minimizing the number of false negatives. Another approach is to generate continuous predictions for all test locations (*i.e.* not just training presences) and choose the threshold that results in a fixed percentage of the lowest predictions being set as absences [34]. However, the selected percentage is heavily influenced by the ratio between presence and absence data. If one generates PAs, then it is possible to select the threshold that directly maximizes the downstream evaluation metric of interest (*e.g.* F1 score) [29] or to achieve a desired target value (*e.g.* 95% sensitivity) [33]. However, these approaches are sensitive to mismatches between the distribution of true absences and generated PAs.

We build on the above existing work by extending the evaluation of thresholding techniques beyond simulation data, and instead use a recent dataset of globally distributed expert-derived range maps of 3,000 different species from [11]. Furthermore, our evaluation is based on recent multi-species neural-network-based SDMs that are jointly trained on PO data from thousands of species.

SDMs as Geo Priors for Vision. In addition to directly generating range maps, SDMs have also been shown to be helpful for assisting image classification. Several works have demonstrated that combining the probabilistic outputs of image classifiers with an estimate of the spatial distribution, *i.e.* as a ‘geo prior’, of the categories of interest can improve classification accuracy. Additional

image metadata, which encodes the geographic location where each image was captured, can be used to estimate the spatial distributions of different classes. Once trained, the geo prior downweights the classifier predictions for categories that are unlikely to be present, according to the prior, at the location where the test image was taken. Different approaches have been explored in the literature, ranging from methods that spatially bin observations [3], ones that jointly train classifiers with additional location inputs [9, 40], or methods that combine classifier predictions with separate spatial distribution models [30]. In our later experiments we show that binary range maps provide a small, but non-trivial, improvement in classification accuracy compared to existing approaches.

3 Method

Our goal is to generate binary species range maps. We begin with an overview of how SDMs can be trained to generate continuous range maps, and then introduce different approaches for converting these continuous predictions to binary ones.

3.1 Predicting Species’ Presence

Recently [11] introduced a neural network-based approach, SINR, for predicting continuous range maps. Their model is trained simultaneously on thousands of different species from crowd-sourced presence-only data. At test time, the input to the model is a longitude-latitude tuple denoting a location of interest, and the output is a multi-label classification indicating the probability that each of several thousand different species is present there. Below we provide an explanation of how their model is parameterized and subsequently trained.

SINR denotes an input location by its longitude and latitude as $\mathbf{x}_i = [lon, lat]$. As in [30], these inputs are wrapped using a sinusoidal encoding to limit the impact of boundary effects. Other more complex input transformations are possible, but this type of wrapped input is still competitive on related spatial representation learning tasks [42]. While it is possible to use other input features such as ones representing local environmental conditions, such as temperature and rainfall, in this work we mainly focus on location only models. The true absence (0) or presence (1) of S different species at \mathbf{x}_i is denoted by the multi-label vector $\mathbf{y}_i \in \{0, 1\}^S$. In practice, in the presence-only setting we do not have any training data that indicates if a species is absent. Instead, similar to [10], the observed training data is represented as $\mathbf{z}_i \in \{1, \emptyset\}^S$, where $z_{ij} = 1$ if a species is present and $z_{ij} = \emptyset$ if the presence or absence of the species is unknown. SINR models estimate \mathbf{y}_i at any location \mathbf{x}_i within a spatial domain \mathcal{X} , given observational training data $\{(\mathbf{x}_i, \mathbf{z}_i)\}_{i=1}^N$.

The range estimation model we wish to learn is denoted as $\hat{\mathbf{y}}_i = h_\phi(f_\theta(\mathbf{x}_i))$, where $f_\theta : \mathcal{X} \rightarrow \mathbb{R}^k$ is a location encoder parameterized by θ , and $h_\phi : \mathbb{R}^k \rightarrow [0, 1]^S$ is a multi-label classifier parameterized by ϕ . The location encoder f_θ is a fully connected neural network with residual connections and the species classifiers in h_ϕ are simple per-species linear projections followed by sigmoid nonlinearities. The prediction $\hat{\mathbf{y}}_i \in [0, 1]^S$ represents a prediction of the probability

of each species being present at location \mathbf{x}_i . The parameters of both components of the model can be estimated using stochastic gradient descent:

$$\theta^*, \phi^* = \operatorname{argmin}_{\theta, \phi} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(\hat{\mathbf{y}}_i, \mathbf{z}_i), \quad (1)$$

where \mathcal{L} is a suitable loss function and $\hat{\mathbf{y}}_i = h_\phi(f_\theta(\mathbf{x}_i))$ is the model’s prediction. The output of this model is thus a probability distribution indicating species’ presence, or absence, over a location of interest.

Training Loss. [11] evaluate different loss functions \mathcal{L} for training SINR models. At their core, these losses amount to different ways of addressing the lack of true absence data. To address this, different mechanisms for generating pseudo-absences (PAs) are explored. For example, taking inspiration from the task of single positive multi-label learning [10], the heuristic that all unobserved labels are negative is used. This assumption also holds for species range estimation because, given an arbitrary location, most species will not actually be present there. Here we focus on their best performing approach, ‘Full Assume Negative’ ($\mathcal{L}_{\text{AN-full}}$), which is inspired by losses introduced in [30]. We describe the other losses in Appendix C.

For the $\mathcal{L}_{\text{AN-full}}$ loss function, PAs are selected in two different ways. First, each species presence is paired with a PA at a location \mathbf{r} that is sampled uniformly from the surface of the earth, *i.e.* $\mathbf{r} \sim \text{Uniform}(\mathcal{X})$. Second, each species observation is also paired with a PA at the same location for a *different* species, which generates PAs that align with the distribution of the data. The final loss is thus represented as:

$$\mathcal{L}_{\text{AN-full}}(\hat{\mathbf{y}}_i, \mathbf{z}_i) = -\frac{1}{S} \sum_{j=1}^S [\mathbb{1}_{[z_{ij}=1]} \lambda \log(\hat{y}_{ij}) + \mathbb{1}_{[z_{ij} \neq 1]} \log(1 - \hat{y}_{ij}) + \log(1 - \hat{y}'_j)], \quad (2)$$

where \hat{y}_{ij} is the predicted probability of species j being present, z_{ij} indicates the presence or absence of species j , $\lambda > 0$ is a hyperparameter used to prevent PAs from dominating the loss, and $\hat{\mathbf{y}}' = h_\phi(f_\theta(\mathbf{r}))$ are the predicted probabilities of the species of interest being present at a randomly chosen location \mathbf{r} . This loss is efficient in that it provides a training signal for all entries in $\hat{\mathbf{y}}$ for each training example, where a positive signal is provided for the observed present species and a negative signal from PAs is provided for all other species.

3.2 Generating Binary Range Maps via Thresholding

Thresholding techniques aim to estimate species-specific thresholds that lead to the most accurate binary range maps. Thresholds are applied after the SDM has been trained and after the predictions for each location have been generated. Once a threshold is identified, all predictions above or equal to it are marked as presences, and all below as absences. Next we outline different threshold generation techniques.

Single Fixed Threshold. This uses a single constant threshold (*e.g.* 0.5) across all species [26,31,32]. If the model is well calibrated, this should be sufficient.

Single Best Threshold. This method also uses a single threshold, but unlike the previous approach here we use the presence and absence data from the expert ranges in the evaluation data to select the threshold that leads to the highest mean F1 score. This threshold is selected by generating 20 linearly spaced candidate thresholds in the range 0-1 and choosing the best one.

Random Sampling. When true absences are not available, one common choice is to generate pseudo-absences (PAs) [11,43]. This method takes presence observations from the *training* data and generates randomly distributed PAs for each species where it has not been observed and selects per-species thresholds that maximize the mean F1 score on this data. This is achieved by binning the observations into H3 [19] cells of resolution four. The cells chosen as presences are those where the species has been observed, the absences are randomly sampled from the rest of the cells. In our later evaluation, we explore the impact of varying the ratio between the true presence and generated PA data.

Target Sampling. To match the distribution of the training data, this approach generates PAs in locations where other species have been observed, but the species of interest is absent [11,43]. We use H3 [19] cells of resolution four to bin the presence data, where cells with at least one observation of a species are counted as presences. All cells where there are at least N_t observations across all species, but the species of interest is not present, are counted as PAs. The value of N_t is calculated by first identifying all the cells where the species of interest has been observed at least once. These cells are ranked based on the ratio of the total number of observations in each cell to the number of observations specifically for the species of interest within that cell. N_t represents the 95th percentile value of this ratio, indicating the minimum number of total observations in a cell needed to be reasonably confident that the species of interest is not present in that cell. Given the resulting presence and PA data, we can select per-species thresholds that maximize the F1 score.

Threshold Classifier. For this method we introduce a novel supervised machine learning approach to predict the best threshold per-species. For the input features for each species we use the final learned species embedding vector from [11], where the assumption is that species with similar embeddings will require similar thresholds. For the output of the learning algorithm, we bin the thresholds into 20 possible values, evenly distributed between 0-1, and treat it as a classification problem. 75% of the 2,418 IUCN species are used to train the classifier and the remaining 25% are used for evaluation. As a result, the final scores for this approach are not directly comparable to the other techniques. We explore two different variants of this approach: a Random Forest (RF) classifier which uses the default values from scikit-learn [35] or a Multilayer Perceptron (MLP) with two hidden layers, containing 100 neurons each.

Mean Predicted Threshold. Following the approach outlined in [29], we select the threshold based on the average of the SDM predictions. Specifically, for each species, we generate predictions across all locations. Subsequently, we determine

the threshold by calculating the mean predicted value from a randomly selected subset of these locations.

Lowest Presence Threshold (LPT). This approach extracts the continuous predictions for a species of interest associated with each presence observation from the training data and selects the threshold as the lowest value from this set [34]. This then ensures that the resulting binary range is guaranteed to encompass each of the presence observations.

Lowest Presence Threshold - Robust (LPT-R). This method builds on the previous LPT approach [34], but we extend it to make it more robust to outliers. The motivation is that LPT is highly susceptible to outliers as inaccurate observations or ‘vagrants’ (individuals lost outside their ranges [25]) can massively bias the threshold selection. So instead of choosing a threshold that includes all training presences, we instead select one that encompasses the vast majority of them. Specifically, we perform the same steps as LPT but set the threshold by sorting all the predicted continuous scores and choosing it to be equal to the 5th percentile. Fig. A5 in the appendix illustrates a visual example of this method.

4 Implementation Details

4.1 Datasets

Presence-Only Training Data. Our training data is publicly available crowd-sourced occurrence data from iNaturalist [20]. The original raw data contains annotated images of species with geographic locations and timestamps. We use the same subset of 35.5 million total observations containing 47,375 species from [11]. Here, observations are of ‘research grade’ status, meaning that there is a higher confidence that the classifications are correct. While iNaturalist offers numerous advantages, it also has limitations. For instance, the data contained within is biased towards the Global North.

Presence-Absence Evaluation Data. For evaluation, we require data that contains both confirmed presence *and* absence observations distributed across the globe. The S&T and IUCN datasets from [11] satisfy these requirements. The *eBird Status and Trends* (S&T) dataset covers 535 bird species biased to North America [16]. This dataset contains estimated relative abundance maps which have been converted into binary range maps. The second dataset consists of ranges of 2,418 species from the *International Union for Conservation of Nature* (IUCN) [21]. In contrast to the S&T dataset, this dataset is more taxonomically diverse and more globally distributed. The two evaluation datasets are represented as H3 [19] cells at resolution five, which corresponds to 2,016,842 evaluation locations distributed across the entire globe.

Image Classification Evaluation Data. In addition to range estimation, we also evaluate different models as geographic priors on the task of assisting image classification predictions as in [3, 11, 30]. We again use the same dataset from [11], which contains 282,974 images from iNaturalist, which represents a subset of 39,444 species from our training set. The evaluation metric for this task is the

top-1 image classification accuracy resulting from combining, in our case binary, range predictions with the probabilistic classification outputs of an Xception classifier [8] trained on images from iNaturalist.

4.2 Evaluation Metrics

Continuous range predictions can be evaluated using area under the curve-based metrics such as average precision. These scores can then be averaged across different species to come up with a single final score as in [11]. However, our goal is to evaluate binary range maps, not continuous predictions. Thus after thresholding a continuous probability of presence output from a model, we consider all locations above the threshold as representing predicted presences and all below as predicted absences. Evaluation is performed using the F1 score, which combines precision and recall into a single value and is particularly useful in binary classification problems where there is an imbalance between the classes, which is the case here as the presences are typically outnumbered by absences. The F1 score is computed per-species after binarization:

$$F_1 = 2 \cdot \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}} = \frac{\text{TP}}{\text{TP} + \frac{1}{2}(\text{FP} + \text{FN})}, \quad (3)$$

where TP denotes True Positives, FP False Positives, and FN False Negatives. The final reported mean F1 score is the average per-species F1 score across the species in each of the respective evaluation sets. Unless stated otherwise, the held-out evaluation datasets are completely detached from any threshold identification methods to ensure that there is no data leakage in the process.

5 Experiments

In this section we evaluate different techniques for generating binary range maps. Unless specified otherwise, experiments are carried out with the $\mathcal{L}_{\text{AN-full}}$ SDM from [11] capped at 1000 training samples per-species without environmental input features. The model’s predictions are binarized using different thresholding techniques and then evaluated by computing the mean F1 score against the held-out test set. This score is presented together with information about how far the performance is from the upper bound. We also provide additional quantitative results, including different input features, and analysis in Appendix A.

5.1 Binarizing Range Maps

Evaluating Different Thresholding Techniques. Tab. 1 presents the mean F1 score for a wide set of binarization techniques when evaluated on the IUCN dataset from [11]. The robust Lowest Presence Threshold (LPT-R) method is the highest-performing technique, closely followed by the **Target Sampling** approach. Both the LPT and LPT-R methods have the advantage that they do not require any pseudo-absences to be defined. Neither **Threshold Classifier** techniques work particularly well here when predicting thresholds, indicating that

Table 1: Binary range estimation performance of different thresholding techniques. Here we report the average mean F1 score for five different random initializations of the $\mathcal{L}_{\text{AN-full}}$ SDM on the IUCN evaluation set, where the upper bound is 67.2%. (†) denotes threshold classifier scores which are computed on a 25% subset of the full evaluation set, as they are trained on the rest. (✓) indicates whether a thresholding technique uses true absences (TA), pseudo-absences (PA), or one single overall threshold (ST). Bold entries indicate best methods, and underline are second best.

Thresholding Method	ST PA TA	† Mean F1	↓ Upper Bound Δ
Threshold Classifier - RF †	✓	55.4	-11.8
Threshold Classifier - MLP †	✓	56.3	-10.9
Single Fixed Threshold - 0.5	✓	40.2	-27.0
Single Best Threshold - 0.1	✓ ✓	55.2	-12.0
Random Sampling - #Absences=#Presences	✓	57.6	-9.6
Random Sampling - #Absences=5#Presences	✓	58.1	-9.1
Random Sampling - #Absences=10#Presences	✓	56.9	-10.3
Random Sampling - 100 Absences	✓	57.3	-9.9
Random Sampling - 1000 Absences	✓	57.6	-9.6
Random Sampling - 10000 Absences	✓	43.9	-23.3
Target Sampling	✓	<u>59.7</u>	<u>-7.5</u>
Mean Predicted Threshold		37.7	-29.5
Lowest Presence Threshold (LPT)		54.3	-12.9
Lowest Presence Threshold - Robust (LPT-R)		60.8	-6.4

the relationship between SDM species’ weights and final thresholds is not very effective for this type of model. Still, these methods, along with the vast majority of techniques, work better than using one **Single Fixed Threshold** of 0.5 for binarization. The main exception to this is the **Mean Predicted Threshold** technique which performs worse than this baseline. Unlike [29], which highlights this approach as viable in the PO setting, our evaluation on a larger set of globally distributed species indicates that it is not very effective. Generally the SDM outputs are very small because most species are absent from most places. Thus, selecting the threshold by averaging predictions results in a small threshold.

Qualitative examples of the resulting binary ranges for several species are depicted in Fig. 2. This figure displays a representative set of species and highlights how the two best-performing methods compare to expert-derived range maps. In many cases, such as the *Wood Thrush* and *Tui Parakeet*, the methods perform similarly. The former case is an example of a North American bird, a species that has abundant PO training data. The latter is a South American species with much less training data, but where the two methods yield almost identical threshold and range maps. The *Tolai Hare* and *Eastern Green Snake* are both examples where the LPT-R approach generates much more accurate ranges. Both of these species are in regions with less training data and have large ranges, as seen in expert ranges on the left. In these examples, the **Target Sampling** greatly underestimates the ranges. This seems to often be the case for species with large ranges and few training samples. However, for species like the *Yosemite Toad* with few training samples and a small range, the **Target Sampling** approach’s more conservative range estimation appears to work better.

Table 2: Binary range estimation across different input SDMs. Here we report the mean F1 score for three different input SDMs, where each is trained with a different loss function. The reported scores are the result of five repeated runs, with different initializations, for each SDM. We also report results for an ensemble of the five different $\mathcal{L}_{\text{AN-full}}$ SDMs, where we average the model predictions before thresholding.

Model	Method	S&T		IUCN	
		↑ Mean F1 ↓ Upper Bound Δ	↑ Mean F1 ↓ Upper Bound Δ	↑ Mean F1 ↓ Upper Bound Δ	↑ Mean F1 ↓ Upper Bound Δ
$\mathcal{L}_{\text{AN-full}}$	Random Sampling	67.4	-8.7	57.6	-9.6
	Target Sampling	70.5	-5.6	<u>59.7</u>	<u>-7.5</u>
	LPT	60.0	-16.1	54.3	-12.9
	LPT-R	<u>69.1</u>	<u>-7.0</u>	60.8	-6.4
$\mathcal{L}_{\text{AN-SSDL}}$	Random Sampling	57.8	-11.9	50.5	-9.6
	Target Sampling	64.1	-5.6	<u>53.1</u>	<u>-7.0</u>
	LPT	58.2	-11.5	49.9	-10.2
	LPT-R	<u>61.3</u>	<u>-8.4</u>	53.5	-6.6
$\mathcal{L}_{\text{AN-SLDS}}$	Random Sampling	71.1	-5.1	<u>38.9</u>	<u>-8.5</u>
	Target Sampling	71.0	-5.2	37.4	-10.0
	LPT	59.5	-16.7	30.0	-17.4
	LPT-R	71.1	-5.1	39.4	-8.0
Ensemble $\mathcal{L}_{\text{AN-full}}$	Random Sampling	67.5	-9.4	58.3	-10.0
	Target Sampling	71.0	-5.9	<u>60.5</u>	<u>-7.8</u>
	LPT	62.8	-14.1	56.2	-12.1
	LPT-R	<u>70.0</u>	<u>-6.9</u>	61.7	-6.6

Evaluating Different SDMs. Next, in Tab. 2 we evaluate the best-performing techniques more using different input SDMs and additional datasets to quantify their robustness. Here, the different SDMs correspond to training SINR models with one of three different loss functions. We include results on the smaller S&T dataset from [11], which is biased toward birds from North America with larger ranges. As can be seen in the results, for the more challenging IUCN dataset, the LPT-R approach outperforms the other methods for all three input SDMs. The **Target Sampling** approach is the second best-performing method for $\mathcal{L}_{\text{AN-full}}$ and $\mathcal{L}_{\text{AN-SSDL}}$, while being the second worst method for $\mathcal{L}_{\text{AN-SLDS}}$. The performance gap is at most a few percentage points, but considering that the LPT-R approach foregoes several computation steps associated with PA generation used by the other methods, this highlights its utility. As LPT-R performs well on the IUCN dataset, it perhaps implies that it may work better for a wider set of species, even those that are less prevalent on iNaturalist. This is important, given that these types of species are often the ones that require monitoring.

For the S&T evaluation dataset, the **Target Sampling** approach generally performs best. However, for the $\mathcal{L}_{\text{AN-SLDS}}$ SDM, the performance of the different top techniques is almost identical. As noted earlier, the S&T dataset has a North American bias, which is also where most of the iNaturalist training data is from. Similar to the findings in [11], the $\mathcal{L}_{\text{AN-SLDS}}$ SDM performs well on the S&T dataset, which is likely due to the fact that this loss function uses a version of target background sampling when generating its pseudo-absences which will also result in a bias towards North America in the context of this dataset.

Ensembling SDMs. Finally, in the bottom rows of Tab. 2 we study the impact of ensembling five different random initializations of the $\mathcal{L}_{\text{AN-full}}$ SDM. This ensemble was used for both threshold identification and, subsequently, evaluation on both S&T and IUCN datasets. The ensemble results in more accurate range maps compared to a single model, as illustrated by the higher mean F1 score for both datasets. We also observe a reduction in the difference between LPT-R and Target Sampling on S&T as a result of ensembling.

5.2 Geo Priors for Image Classification

SDMs can also be used as ‘geo priors’ to assist image classifiers by combining their outputs with a classifier’s probabilistic predictions. Here, we conduct an evaluation similar to the one carried out in [11], with the main difference being that with binary range maps, locations outside a species range would be downweighted to zero, *i.e.* the SDM would prevent the classifier from predicting that species. We use the image classification dataset from [11], which consists of 282,974 images of different species collected from iNaturalist, covering 39,444 species from the PO training set. Each image is accompanied by the latitude and longitude at which it was taken. In Tab. 3, the results of using different binary range maps are presented. Using a geo prior to weight predictions increases the Top-1 accuracy in all experiments compared to the vision-only prediction model (75.4%), except when binarizing with a constant threshold of 0.5 for all species. As previously established, different species require different thresholds, and a threshold of 0.5 is generally too high. This results in most species being incorrectly marked as absent from nearly all locations, which degrades model performance. However, setting a lower threshold of 0.1 for binarizing all species significantly improves accuracy, and outperforms the vision-only model. This shows that the vision model benefits from a geo prior, even if the range maps are not very accurate when making predictions.

Switching to more refined methods of setting species-specific thresholds improves the performance even further. Out of all methods, the Target Sampling and LPT-R approaches work best. However, Tab. 3 shows that this performance is lower than when using the raw prediction scores without binarizing them, *i.e.* the original continuous outputs of the SDM (81.6%). We hypothesize that the continuous scores allow the model to express uncertainty about a species being at a given location rather than choosing between the two extremes of classifying as either present or absent. To address this, we add a small constant $+\delta$ to the predicted absences after thresholding when using them as priors, *e.g.* 0.01. In all cases, adding this constant improves the performance of all binarization techniques as it prevents them from saying that a species is never present at a location. As a result, the binarized range maps generated through both Target Sampling and LPT-R outperform the continuous baseline.

5.3 Limitations

Despite the promise of our work, there are still some important limitations to note. First, the training data contains biases, both spatial biases at the national (*e.g.* Europe and North America are disproportionately represented) and local

Thresholding Method	Top-1 Acc.	
		+ δ
No Prior (<i>i.e.</i> vision-only)	75.4	-
Continuous - No Threshold	81.6	81.4
Single Fixed Threshold - 0.5	52.7	75.7
Single Fixed Threshold - 0.1	77.6	81.2
Single Fixed Threshold - 0.01	80.9	81.3
Random Sampling	78.5	81.5
Target Sampling	80.1	81.7
LPT	80.4	80.8
LPT-R	80.0	81.8

Table 3: Image classification with geo priors using different thresholding techniques. Here we use a $\mathcal{L}_{AN-full}$ SDM as a geo prior to assist image classification. The results are the average of five different random initializations of the SDM. ‘+ δ ’ indicates that a small constant is added after thresholding to ensure that it is not possible for a thresholding technique to predict that a species is never present.

scale (*e.g.* it is biased towards locations that are easier for people to get to), in addition to some species being much more prevalent than others. Spatial biases in citizen science data have been well characterized in the literature [18] and various approaches have been proposed to mitigate it [7, 23]. We leave the investigation of such methods for future work. Another data related issue is that the expert-derived range maps in the test set can contain errors or simply be outdated [11]. However, to the best of our knowledge the test set used represents the best source of large-scale global evaluation data available. The main results in this work use species distribution models that only use coordinate features as input. Results can change if additional environmental covariates are used as input, see Appendix A. We also do not perform any regularization on the binary maps to make them more spatially contiguous. This could be investigated in future work, although we note that our predictions are typically already contiguous. Finally, caution should be exercised when using the resulting binary range maps in conservation or assessment applications as the results demonstrate that performance is still likely insufficient for many species.

6 Conclusion

We explored the problem of automatically converting continuous species distribution model outputs into binary range maps via thresholding. Through detailed evaluation, we compared the performance of multiple different thresholding techniques, using different underlying models, across a range of different species. We also proposed an extension of an existing method, which we call LPT-R, and demonstrate that leads to more accurate binary range maps in a number of cases. Our approach circumvents the need to create pseudo-absences, making it computationally cheaper while also yielding more accurate ranges. Additionally, we showed that binary range maps can also be used as geographic priors to improve fine-grained image classification accuracy. A possible future extension to our work would be to leverage temporal data to generate time-conditioned range maps, which could significantly enhance our understanding of distribution changes over time.

Acknowledgments. We wish to thank the iNaturalist community for sharing their observation data. Thanks also to Alison Johnston for helpful comments. FD and OMA were in part supported by a Royal Society Research Grant.

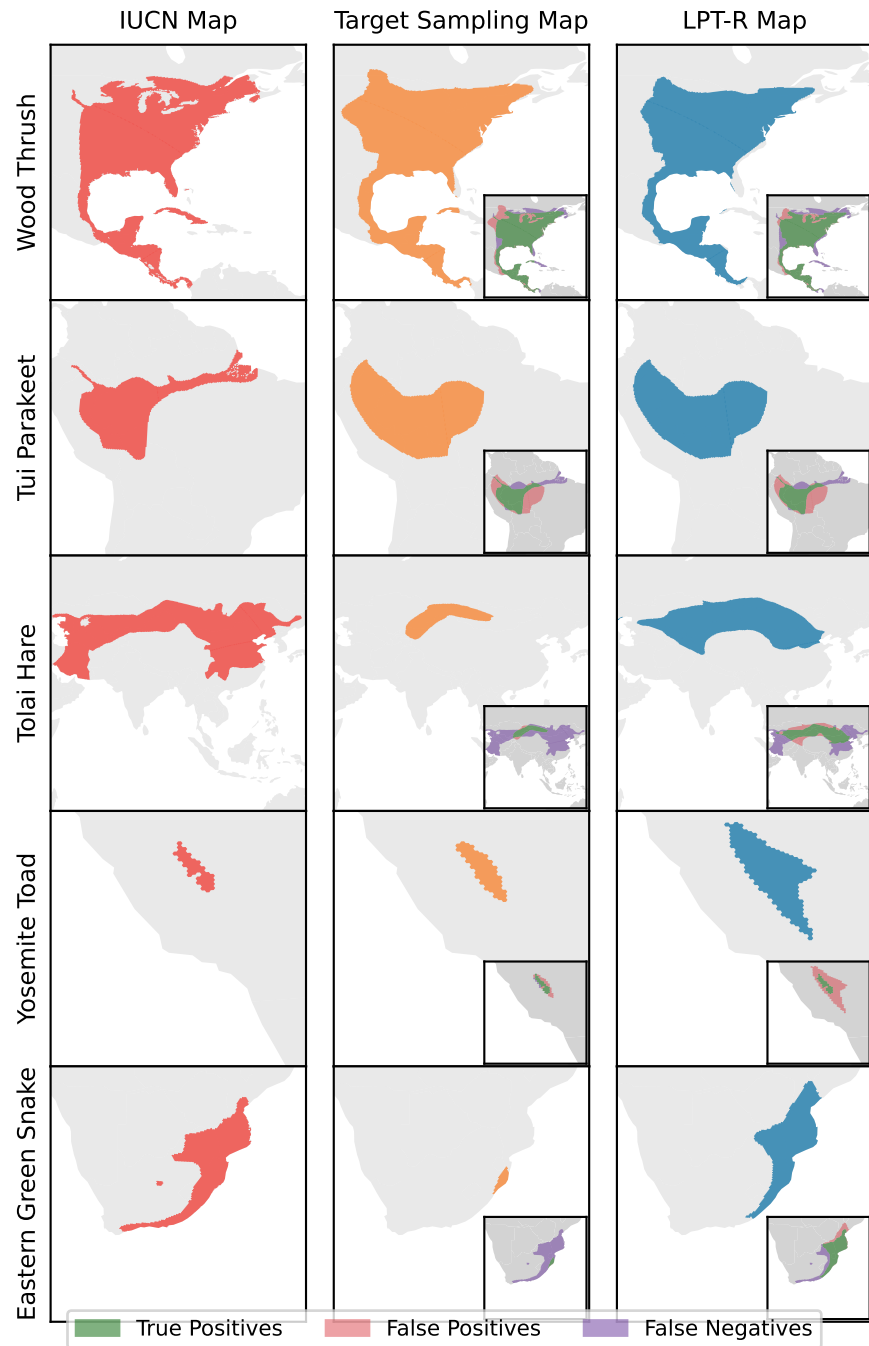


Fig. 2: Qualitative examples of estimated binary ranges. Each row depicts a different species, and the columns show the expert-derived ranges and the outputs from the **Target Sampling** and **LPT-R** approaches, respectively. Inset, we also display the different types of errors. We use an ocean mask for visualization purposes.

References

1. Barbet-Massin, M., Jiguet, F., Albert, C.H., Thuiller, W.: Selecting pseudo-absences for species distribution models: How, where and how many? *Methods in Ecology and Evolution* (2012) [3](#)
2. Beery, S., Cole, E., Parker, J., Perona, P., Winner, K.: Species distribution modeling for machine learning practitioners: A review. In: *Conference on Computing and Sustainable Societies* (2021) [1](#), [3](#)
3. Berg, T., Liu, J., Woo Lee, S., Alexander, M.L., Jacobs, D.W., Belhumeur, P.N.: Birdsnap: Large-scale fine-grained visual categorization of birds. In: *CVPR* (2014) [5](#), [8](#)
4. Botella, C., Joly, A., Bonnet, P., Monestiez, P., Munoz, F.: A deep learning approach to species distribution modelling. *Multimedia Tools and Applications for Environmental & Biodiversity Informatics* (2018) [1](#)
5. Botella, C., Joly, A., Monestiez, P., Bonnet, P., Munoz, F.: Bias in presence-only niche models related to sampling effort and species niches: Lessons for background point selection. *PLoS One* (2020) [3](#)
6. Breiman, L.: Random forests. *Machine learning* (2001) [4](#)
7. Chen, D., Gomes, C.P.: Bias reduction via end-to-end shift learning: Application to citizen science. In: *AAAI* (2019) [13](#)
8. Chollet, F.: Xception: Deep learning with depthwise separable convolutions. In: *CVPR* (2017) [9](#)
9. Chu, G., Potetz, B., Wang, W., Howard, A., Song, Y., Brucher, F., Leung, T., Adam, H.: Geo-aware networks for fine-grained recognition. In: *ICCV Workshops* (2019) [5](#)
10. Cole, E., Mac Aodha, O., Lorieul, T., Perona, P., Morris, D., Jovic, N.: Multi-label learning from single positive labels. In: *CVPR* (2021) [5](#), [6](#)
11. Cole, E., Van Horn, G., Lange, C., Shepard, A., Leary, P., Perona, P., Loarie, S., Mac Aodha, O.: Spatial implicit neural representations for global-scale species mapping. In: *ICML* (2023) [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [9](#), [11](#), [12](#), [13](#)
12. Davis, C.L., Bai, Y., Chen, D., Robinson, O., Ruiz-Gutierrez, V., Gomes, C.P., Fink, D.: Deep learning with citizen science data enables estimation of species diversity and composition at continental extents. *Ecology* (2023) [1](#)
13. Elith, J., Leathwick, J.R.: Species distribution models: ecological explanation and prediction across space and time. *Annual review of ecology, evolution, and systematics* (2009) [3](#)
14. Elith, J., Leathwick, J.R., Hastie, T.: A working guide to boosted regression trees. *Journal of animal ecology* (2008) [4](#)
15. Fielding, A.H., Bell, J.F.: A review of methods for the assessment of prediction errors in conservation presence/absence models. *Environmental Conservation* (1997) [4](#)
16. Fink, D., Auer, T., Johnston, A., Strimas-Mackey, M., Robinson, O., Ligocki, S., Hochachka, W., Jaromczyk, J., Wood, C., Davies, I., Iliff, M., Seitz, L.: eBird status and trends, data version: 2020; released: 2021 (10 2020) [8](#)
17. Freeman, E.A., Moisen, G.G.: A comparison of the performance of threshold criteria for binary classification in terms of predicted prevalence and kappa. *Ecological Modelling* (2008) [4](#)
18. Geldmann, J., Heilmann-Clausen, J., Holm, T.E., Levinsky, I., Markussen, B., Olsen, K., Rahbek, C., Tøttrup, A.P.: What determines spatial bias in citizen science? exploring four recording schemes with different proficiency requirements. *Diversity and Distributions* (2016) [13](#)

19. H3: (2024), <https://h3geo.org> 7, 8
20. iNaturalist: (2024), <https://www.inaturalist.org> 2, 3, 8
21. IUCN: The IUCN red list of threatened species (2022), <https://www.iucnredlist.org>, version 2022-2 8
22. Jiménez-Valverde, A., Lobo, J.M.: Threshold criteria for conversion of probability of species presence to either–or presence–absence. *Acta Oecologica* (2007) 4
23. Johnston, A., Moran, N., Musgrove, A., Fink, D., Baillie, S.R.: Estimating species distributions from spatially biased citizen science data. *Ecological Modelling* (2020) 13
24. Lange, C., Cole, E., Horn, G., Mac Aodha, O.: Active learning-based species range estimation. *NeurIPS* (2023) 1
25. Lees, A., Gilroy, J.: *Vagrancy in birds*. Bloomsbury Publishing (2022) 8
26. Li, W., Wang, Z., Ma, Z., Tang, H.: A regression model for the spatial distribution of red-crown crane in yancheng biosphere reserve, china. *Ecological Modelling* (1997) 7
27. Liu, C., Berry, P.M., Dawson, T.P., Pearson, R.G.: Selecting thresholds of occurrence in the prediction of species distributions. *Ecography* (2005) 2, 4
28. Liu, C., Newell, G., White, M.: On the selection of thresholds for predicting species occurrence with presence-only data. *Ecology and evolution* (2016) 4
29. Liu, C., White, M., Newell, G.: Selecting thresholds for the prediction of species occurrence with presence-only data. *Journal of Biogeography* (2013) 2, 4, 7, 10
30. Mac Aodha, O., Cole, E., Perona, P.: Presence-only geographical priors for fine-grained image classification. In: *ICCV* (2019) 1, 3, 5, 6, 8
31. Manel, S., Dias, J.M., Ormerod, S.J.: Comparing discriminant analysis, neural networks and logistic regression for predicting species distributions: a case study with a himalayan river bird. *Ecological modelling* (1999) 7
32. Manel, S., Williams, H.C., Ormerod, S.J.: Evaluating presence–absence models in ecology: the need to account for prevalence. *Journal of applied Ecology* (2001) 7
33. Pearson, R.G., Dawson, T.P., Liu, C.: Modelling species distributions in britain: a hierarchical integration of climate and land-cover data. *Ecography* (2004) 4
34. Pearson, R.G., Raxworthy, C.J., Nakamura, M., Townsend Peterson, A.: Predicting species distributions from small numbers of occurrence records: a test case using cryptic geckos in madagascar. *Journal of Biogeography* (2007) 4, 8
35. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al.: Scikit-learn: Machine learning in python. *JMLR* (2011) 7
36. Phillips, S.J., Anderson, R.P., Schapire, R.E.: Maximum entropy modeling of species geographic distributions. *Ecological Modelling* (2006) 4
37. Phillips, S.J., Dudík, M., Elith, J., Graham, C.H., Lehmann, A., Leathwick, J., Ferrier, S.: Sample selection bias and presence-only distribution models: implications for background and pseudo-absence data. *Ecological applications* (2009) 3
38. Ponder, W.F., Carter, G., Flemons, P., Chapman, R.: Evaluation of museum collection data for use in biodiversity assessment. *Conservation Biology* (2001) 3
39. Standards, Group, P.W., et al.: Guidelines for using the iucn red list categories and criteria v16. International Union for Conservation of Nature, Gland, Switzerland (2024) 1
40. Tang, K., Paluri, M., Fei-Fei, L., Fergus, R., Bourdev, L.: Improving image classification with location context. In: *ICCV* (2015) 5
41. Teng, M., Elmustafa, A., Akera, B., Larochele, H., Rolnick, D.: Bird distribution modelling using remote sensing and citizen science data. *Tackling Climate Change with Machine Learning Workshop, ICLR* (2023) 1

42. Wu, N., Cao, Q., Wang, Z., Liu, Z., Qi, Y., Zhang, J., Ni, J., Yao, X., Ma, H., Mu, L., et al.: Torchspatial: A location encoding framework and benchmark for spatial representation learning. arXiv:2406.15658 (2024) 5
43. Zbinden, R., Van Tiel, N., Kellenberger, B., Hughes, L., Tuia, D.: On the selection and effectiveness of pseudo-absences for species distribution modeling with deep learning. Ecological Informatics (2024) 3, 7