

# BIRD DISTRIBUTION MODELLING USING REMOTE SENSING AND CITIZEN SCIENCE DATA

**Mélisande Teng**

Mila, Université de Montréal  
tengmeli@mila.quebec

**Amna Elmustafa**

Mila, Stanford University  
amna97@stanford.edu

**Benjamin Akera**

Mila, McGill University  
akeraben@mila.quebec

**Hugo Larochelle**

Mila, Google Research, Brain Team  
hugolarochelle@google.com

**David Rolnick**

Mila, McGill University  
drolnick@cs.mcgill.ca

## ABSTRACT

Climate change is a major driver of biodiversity loss, changing the geographic range and abundance of many species. However, there remain significant knowledge gaps about the distribution of species, due principally to the amount of effort and expertise required for traditional field monitoring. We propose an approach leveraging computer vision to improve species distribution modelling, combining the wide availability of remote sensing data with sparse on-ground citizen science data. We introduce a novel task and dataset for mapping US bird species to their habitats by predicting species encounter rates from satellite images, along with baseline models which demonstrate the power of our approach. Our methods open up possibilities for scalably modelling ecosystems properties worldwide.

## 1 INTRODUCTION

Climate change presents significant threats to global biodiversity, both directly and by compounding the effects of other anthropogenic stressors, such as habitat loss, pollution, and introduced species. Biodiversity loss in turn also impacts ecosystem services necessary to ensure food, water, and human health and well-being (Marselle et al., 2021). It is crucial to understand the changing distributions of species globally to inform policy decisions, for example in shaping land use and land conservation choices. However, traditional methods for species distribution models (SDMs) generally focus either on narrow sets of species or narrow geographical areas. This is in part due to the methods’ computational cost, the insufficient ability of models to account for complex relationships between different variables, and the type of data used, whose availability can be limited (Phillips et al., 2005; Newbold, 2010).

Machine learning algorithms for remote sensing have increasingly seen wide applicability across sustainability-related domains (see e.g. Rolf et al. (2021)) and have been suggested as a promising tool for SDMs (Beery et al., 2021). Moreover, the surge of data collection on citizen science platforms along with improved data quality validation processes in recent years offers tremendous opportunity for scientific research, as species observation records from these sources can cover a larger temporal and geographic extent at a finer resolution and at lower cost than traditional sampling methods. Indeed, citizen science data and remotely sensed ecosystem attributes such as vegetation indices have been shown to improve the performance of models, especially for less widespread species (Arenas-Castro et al., 2022). Recently, the GeoLifeCLEF challenge (Lorieul et al., 2022) was introduced with the goal of directly predicting plant and animal abundance from aerial images at 1m resolution, using deep computer vision. However, the GeoLifeCLEF benchmark has proven extremely challenging, potentially since only one species is associated with each location.

We propose to use remote sensing to infer the joint distribution of many species for a given location, using publicly available citizen science observation records as ground truth. Our approach leverages the fact that a species’ presence or absence at a location depends on the ecosystem present there, and therefore the abundances of different species are highly correlated. Specifically, we predict the encounter rates of 684 bird species at sites across the continental USA. We focus on birds due to

the availability of high-quality bird observations and the ecological importance of birds, which play a vital role in almost every terrestrial ecosystem and are threatened significantly by climate change (Rodenhouse et al., 2008; Stephens et al., 2016). Our contributions include: (i) framing the task of predicting bird species encounter rates at a specific location using remote sensing data, (ii) building a dataset for this task, obtained from publicly available bird observation and satellite data sources, and (iii) showing the efficacy of baseline deep computer vision methods for this task.

**Problem definition.** We consider presence-absence bird sighting records from the citizen science database eBird (Kelling et al., 2013). Each location (termed a *hotspot*) in the eBird database is associated to a number of *complete checklists* containing all species a birdwatcher was able to observe at a specific date and time. If  $h$  is a hotspot, and  $s_1, \dots, s_n$  the species of interest, then our goal is to build a machine learning model that takes as input a satellite image of  $h$  (and optionally other data) and predicts the vector  $\mathbf{y}^h = (y_{s_1}^h, \dots, y_{s_n}^h)$ , where  $y_s^h$  is the number of complete checklists reporting species  $s$  at  $h$  divided by the total number of complete checklists at  $h$ . This ratio  $y_s^h$  can be understood as an *encounter rate*, the probability for a visitor to observe a species if they visit the hotspot. We aim at jointly predicting this quantity for all bird species in the region. Thus, our task can be considered as a supervised multi-output regression problem. We consider encounter rates as our target variable because they are a widely used measure in species distribution modeling, and are indeed used extensively on the eBird platform (in the form of “hotspot bar charts”) in order to summarize the species in a given location for birdwatchers and ornithologists.

## 2 DATASET

We introduce a dataset for the task defined in 1 with the continental USA as our region of interest. Bird distributions vary seasonally; we consider specifically the month of June, representing the breeding period for most birds in the region, and making it possible to neglect migratory effects. Our dataset will be released publicly upon acceptance.

**Species observation data.** We extracted all eBird complete checklists recorded in June from 2000 to 2021 in the continental USA. We filtered out hotspots with fewer than 5 complete checklists recorded in June to ensure more meaningful encounter rates, resulting in 8439 hotspots. We included all regularly occurring species in the region, as denoted by ABA Codes 1 and 2 (ABA, 2022), omitting species found only in Hawaii and Alaska, as well as one Code-2 seabird species with rare oceanic observations, leaving a total of 684 species. We then computed the encounter rates for each hotspot and corrected for vagrants (species seen in hotspots outside of their geographical range) by using range maps from eBird to set the target encounter rates to zero in these hotspots. We aggregate the species observations in June over twenty years, only considering seasonal change in distributions. We leave annual temporal change for future iterations of this work.

**Remote Sensing Imagery.** For each hotspot, we extracted RGB and NIR bands (reflectance measurements) from Sentinel-2 satellite tiles covering the entirety of a square of about  $5 \text{ km}^2$  centered around the hotspots. The images have a resolution of  $10 \text{ m}$  for the considered bands. We extracted images in June with cloud coverage of at most 10 %, keeping the least cloudy and most recent image of those in the period 2016-2020 if it covered the entire  $5 \text{ km}^2$  region or composing a mosaic with the extracted images otherwise, in order to minimize seams in the images of our dataset. We associate one image per hotspot for years from 2016-2020, considering that a more recent satellite image would be more representative of our species data, since there are more checklists in recent years compared to earlier years.

**Environmental data.** We extracted 19 bioclimatic and 8 pedologic variables as rasters from WorldClim 1.4 and SoilGrids in the same fashion as the GeoLifeCLEF 2020 dataset (Cole et al., 2020) for each hotspot. Additionally, we extracted land cover data patches from the ESRI Land Use Land Cover 10m-resolution maps derived from Sentinel-2 (Esri, 2020).

**Train-test split** In order to account for spatial auto-correlation and overfitting that can arise from random splits of geospatial data (Roberts et al., 2017), we first clustered hotspots such that the minimum distance between clusters is  $5 \text{ km}$  and randomly assigned them to train, validation and test splits in a 70%, 20%, 10% proportion.

### 3 METHODOLOGY

In this section, we detail baselines for this task as well as directions for further improvement.

#### 3.1 ENVIRONMENTAL BASELINE

A first naive baseline is the mean encounter rates over the training set for each species. Following the GeoLifeCLEF challenge, we also propose an environmental baseline, using Gradient Boosted Regression Trees on the bioclimatic and pedological variables extracted at each of the hotspots.

#### 3.2 CNN BASELINE

We propose a first CNN model with ResNet-18 (He et al., 2016) architecture. We initialize the network with ImageNet pretrained weights. Since ImageNet has only RGB bands, the initialization of the first layer for the NIR band is performed by sampling from a normal distribution parameterized by the mean and standard deviation of the layer’s weights for the other bands. We consider a region of interest of 640 m<sup>2</sup>, center-cropping the satellite patches to size 64 × 64 around the hotspot and normalizing the bands with our training set statistics. We use random vertical and horizontal flipping for data augmentation and train the model with cross entropy  $\mathcal{L}_{CE} = \frac{1}{N_h} \sum_h \mathcal{L}_h = \frac{1}{N_h} \sum_h \sum_{s(\text{species})} -y_h^s \log(\hat{y}_h^s) - (1 - y_h^s) \log(1 - \hat{y}_h^s)$  where  $N_h$  is the number of hotspots  $h$ ,  $y$  the predictions of the model and  $\hat{y}$  the ground truth encounter rates.

We add the bioclimatic and pedological data by normalizing it variable-wise with training set statistics, aligning it to the satellite images’ resolution and stacking the corresponding patches to the images. We add the landcover data following the same procedure.

#### 3.3 INCLUDING GEOGRAPHICAL INFORMATION

We explore different methods for making the model spatially explicit to account for the geographical ranges of some species. For example, Eastern and Western wood pewees have the same habitat but are not found in the same geographical regions, and thus do not co-occur. A model blind to geographical information might predict both species in the same place, which is undesirable when producing maps for conservation planning, and could also impede training.

**Location encoder (LE).** We train a location encoder taking as input latitude and longitude information jointly with the image encoder. We use the same architecture as the location encoder of Mac Aodha et al. (2019) (cf. A.2), removing the final classification layer, and instead concatenating the obtained features to those obtained from the ResNet image encoder, before passing them to a fully-connected layer to obtain the predicted encounter rates.

**Hard-masking with range maps (RM).** Some species have known geographical range. In this approach, we use range maps available via eBird which are updated regularly as binary masks on the predictions to zero out the encounter rates of species in regions where they cannot be found. If the range map is not available for a species, the predictions are left untouched.

**Soft-masking (SM).** Here, we compute proxies for range maps from the training set data, since range maps (a) are available only for certain species and (b) do not capture finer gradations of abundance across regions. We define a regional “correction” factor  $c_R$ , which reflects the relative prevalence of species in region  $R$  compared to the prevalence in the whole geographical area  $W$ . In our case,  $R$  varies across US states, while  $W$  is the continental US. For each species  $s$ , we set:

$$c_s^R = \frac{\sum_{l \in C_R} \mathbb{1}_{s \in l}}{|C_R|} / \frac{\sum_{l \in C_W} \mathbb{1}_{s \in l}}{|C_W|} \tag{1}$$

where  $C_R$  and  $C_W$  are the set of observations in  $R$  and  $W$  resp., and  $s \in l$  denotes that  $s$  has been reported in  $l$ . For a hotspot  $h$  in  $R$ , the final prediction is  $\text{clip}_{[0,1]}(c_R \cdot y_h)$ . See A.3 for derivation.

#### 3.4 INCLUDING HOTSPOT METADATA

Additionally, we consider a weighted loss (WL) with weights as a function of the number of checklists  $n_h$  at each hotspot, and give more importance to more visited hotspots since the data at such


hotspots is likely more reliable by virtue of more checklists. The contribution of each hotspot to the loss is  $f(n_h)\mathcal{L}_h$ . For  $f$ , we consider the identity, log and square root functions.

### 3.5 EVALUATION

Beyond optimization for common regression metrics (MSE, MAE), it is desirable for the predicted most likely species in a given hotspot to coincide with those that have most frequently been observed. We therefore report top-10 and top-30 accuracies, representing for  $k = 10$  and  $k = 30$  the number of species present in both the top  $k$  predicted species and the top  $k$  observed species, divided by  $k$ . However, the number of species observed varies considerably across hotspots (cf. A.1), unlike in e.g. the GeoLifeCLEF challenge for which there is only one ground truth species for each hotspot. We therefore also define an *adaptive top-k accuracy* as the top- $k$  value where  $k$  is the total number of species observed at that hotspot. We also analyse the predictions of the model by looking at precision and recall per species across hotspots and provide further details in A.6.

## 4 EXPERIMENTS

Numerical results on the test set for the models presented in Section 3 are provided in Table 1 and validation performance is reported in A.4. All models were trained with batch size 64, Adam optimizer, initial learning rate 0.0003, with a scheduler that decreases the rate based on validation loss. While the CNN baseline with RGB and NIR bands does not outperform the environmental baseline on all metrics, combining satellite and environmental data outperforms the latter, particularly on the top-k metric. This highlights the importance of remote sensing data for our task. Adding landcover data did not improve the model, likely due to the coarseness of the 10 classes, so it was not used for further experiments. While adding geographical information through RM masking improves the performance of the model, the LE and the SM methods did not, perhaps because of the simplicity of the architecture of the former and the coarse region unit (US states) considered for the latter. We plan to explore other location encoding methods, such as space2vec (Mai et al., 2020) and consider smaller units for the soft-masking. Fig. 1 shows the geographical distribution of performance on the test set of the hard-masking with RM model and we provide further examples of predictions in A.5. Species-wise precision and recall are reported in A.6. We find that range-restricted species, along with very common species, are the best performing species according to recall. Interestingly, the weighted loss yields performance similar to that of regular cross-entropy loss. The best performance was achieved with the log function. This suggests that checklists from less frequented hotspots (e.g. with fewer than 5 complete checklists) could be included in future iterations without compromising performance.



**Figure 1:** Test set hotspots colored by adaptive top-k performance for the RGBNIR + Env + range maps masking model.

**Table 1:** Results for the proposed models on the test set. “Env” indicates the use of environmental variables. All MSE and MAE scores are reported  $\times 10^3$  and  $\times 10^2$  respectively. “Top- $k$ ” denotes the adaptive top- $k$  accuracy.

Method	MSE [ $10^{-3}$ ]	MAE [ $10^{-2}$ ]	Top- $k$	Top-30	Top-10
Mean encounter rate	7.18	2.91	51.46	43.91	26.45
Env baseline	4.83±0.00	2.05 ± 0.00	68.86 ± 0.01	62.3 ± 0.03	43.11 ± 0.06
RGBNIR	5.39±0.00	1.98 ± 0.00	67.0 ± 0.02	60.1 ± 0.17	41.83 ± 0.30
RGBNIR + Env	4.57±0.02	1.78 ± 0.01	72.91 ± 0.19	67.0 ± 0.00	48.0 ± 0.00
RGBNIR + Env + LE	4.49 ± 0.05	1.73± 0.01	72.01 ± 0.08	65.00 ± 0.00	44.50 ± 0.70
<b>RGBNIR + Env + RM</b>	<b>4.45 ± 0.05</b>	1.75± 0.01	<b>73.38 ± 0.08</b>	<b>67.68 ± 0.11</b>	<b>48.15 ± 0.20</b>
RGBNIR + Env + SM	13.3±0.39	2.74 ± 0.08	65.0 ± 1.39	54.7 ± 2.3	31.7 ± 1.04
RGBNIR + Env + RM + WL	4.76 ± 0.00	<b>1.60 ± 0.00</b>	72.45±0.02	65.90 ± 0.02	46.30 ± 0.006

## 5 CONCLUSION

We introduce a new dataset built from publicly available data to leverage the combination of satellite imagery and presence-absence citizen science data for joint species distribution modelling. We are currently extending our approach and dataset to other taxonomic groups and regions. In East Africa, for example, sparse data poses additional challenges but also underscores the importance of scalable remote sensing approaches to supplement scarce on-the-ground data. We hope to integrate our algorithm into eBird’s existing tool that lists the “likely species” in a given area. This tool currently relies on past checklists and is therefore only available in well-monitored locations. By approximating species distributions at poorly monitored locations using remote sensing to evaluate changing land use patterns and habitat, we anticipate significantly expanding the ability of ornithologists to rapidly estimate biodiversity across the world.

## REFERENCES

- ABA. *American Birding Association checklist: Birds of the continental United States and Canada*. 2022.
- Salvador Arenas-Castro, Adrián Regos, Ivone Martins, João Honrado, and Joaquim Alonso. Effects of input data sources on species distribution model predictions across species with different distributional ranges. *Journal of Biogeography*, 49(7):1299–1312, 2022.
- Sara Beery, Elijah Cole, Joseph Parker, Pietro Perona, and Kevin Winner. Species distribution modeling for machine learning practitioners: A review. In *ACM SIGCAS Conference on Computing and Sustainable Societies*, pp. 329–348, 2021.
- Elijah Cole, Benjamin Deneu, Titouan Lorieul, Maximilien Servajean, Christophe Botella, Dan Morris, Nebojsa Jojic, Pierre Bonnet, and Alexis Joly. The GeoLifeCLEF 2020 dataset. *Preprint arXiv:2004.04192*, 2020.
- Esri. Esri land cover 2020. <https://livingatlas.arcgis.com/landcover/>, 2020. Accessed: 2022-12-06.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Steve Kelling, Jeff Gerbracht, Daniel Fink, Carl Lagoze, Weng-Keen Wong, Jun Yu, Theo Damoulas, and Carla Gomes. eBird: A human/computer learning network for biodiversity conservation and research. *AI Magazine*, 34, 03 2013.
- Titouan Lorieul, Elijah Cole, Benjamin Deneu, Maximilien Servajean, Pierre Bonnet, and Alexis Joly. Overview of GeoLifeCLEF 2022: Predicting species presence from multi-modal remote sensing, bioclimatic and pedologic data. In *CLEF 2022-Conference and Labs of the Evaluation Forum*, volume 3180, pp. 1940–1956, 2022.
- Oisín Mac Aodha, Elijah Cole, and Pietro Perona. Presence-only geographical priors for fine-grained image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9596–9606, 2019.
- Gengchen Mai, Krzysztof Janowicz, Bo Yan, Rui Zhu, Ling Cai, and Ni Lao. Multi-scale representation learning for spatial feature distributions using grid cells. In *International Conference on Learning Representations*, 2020.
- Melissa R Marselle, Terry Hartig, Daniel TC Cox, Siân De Bell, Sonja Knapp, Sarah Lindley, Margarita Triguero-Mas, Katrin Böhning-Gaese, Matthias Braubach, Penny A Cook, et al. Pathways linking biodiversity to human health: A conceptual framework. *Environment International*, 150: 106420, 2021.
- Tim Newbold. Applications and limitations of museum data for conservation and ecology, with particular attention to species distribution models. *Progress in physical geography*, 34(1):3–22, 2010.

Steven J Phillips et al. A brief tutorial on Maxent. *AT&T Research*, 190(4):231–259, 2005.

David R Roberts, Volker Bahn, Simone Ciuti, Mark S Boyce, Jane Elith, Gurutzeta Guillera-Arroita, Severin Hauenstein, José J Lahoz-Monfort, Boris Schröder, Wilfried Thuiller, et al. Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*, 40(8):913–929, 2017.

Nicholas L Rodenhouse, SN Matthews, KP McFarland, JD Lambert, LR Iverson, A Prasad, T Scott Sillett, and Richard T Holmes. Potential effects of climate change on birds of the Northeast. *Mitigation and adaptation strategies for global change*, 13:517–540, 2008.


Esther Rolf, Jonathan Proctor, Tamma Carleton, Ian Bolliger, Vaishaal Shankar, Miyabi Ishihara, Benjamin Recht, and Solomon Hsiang. A generalizable and accessible approach to machine learning with global satellite imagery. *Nature communications*, 12(1):1–11, 2021.

Philip A Stephens, Lucy R Mason, Rhys E Green, Richard D Gregory, John R Sauer, Jamie Allison, Ainars Aunins, Lluís Brotons, Stuart HM Butchart, Tommaso Campedelli, et al. Consistent response of bird populations to climate change on two continents. *Science*, 352(6281):84–87, 2016.

## A APPENDIX

### A.1 TRAINING SET DISTRIBUTION

A characteristic of the presence-absence data considered for this task is the zero-inflated nature of the targets. Indeed, while 684 species are considered, hotspots in our training set have 57 species with non-zero encounter rate on average. Fig. 2 shows the distribution of number of species encountered.




**Figure 2:** Distribution of the number of species encountered in training set hotspots.

### A.2 LOCATION ENCODER

We include latitude-longitude information through a separate encoder as shown in Fig. 3. We follow the implementation for Mac Aodha et al. (2019) for the location encoder. We first map each spatial coordinate  $x^l$  (latitude and longitude) to  $[\cos(\pi x^l), \sin(\pi x^l)]$ . We end up with an input vector of size 4, which is passed through an initial fully connected layer, followed by 4 residual blocks, each consisting of two fully connected layers with a dropout layer in between.

### A.3 DERIVATION OF THE SOFT MASKING FACTOR

In this section we explain further our choice of the soft masking factor which draws inspiration from the work of Mac Aodha et al. (2019). The underlying assumption is that while satellite images are



**Figure 3:** Schematic view of the model including location information through a separate encoder.

informative about habitat, location also encodes geography, and both are important for designing a spatially-explicit model in our setting. To encode location here, we take  $y$  as the target,  $x$  as the input image and  $l$  as the location in a region  $R$ . Predictions conditioned on both  $x$  and  $l$  can be expressed using Bayes rule:

$$p(y|x, l) = \frac{p(y|x) * p(x) * p(l|y, x)}{p(x, l)} \tag{2}$$

Under the assumption that  $l$  and  $x$  are independent and also independent w.r.t  $y$ , we can approximate the formula as follows, where  $p(y|x)$  is the output of the CNN and  $c^R$  is the correction factor. In fact, we account for geography at the level of region  $R$ , meaning that predictions in all location within  $R$  will be corrected with the same correction factor.

$$p(y|x, l) = \frac{p(y|x) * p(l|y)}{p(l)} \tag{3}$$

$$c^R = \frac{p(l|y)}{p(l)} \tag{4}$$

which is interpreted as the proportion of checklists in  $R$  reporting  $y$  over the proportion of checklists anywhere reporting  $y$ . This can be expressed in terms of counts for each species  $s$  as :

$$c_s^R = \frac{\sum_{l \in C_R} \mathbb{1}_{s \in l}}{|C_R|} / \frac{\sum_{l \in C_W} \mathbb{1}_{s \in l}}{|C_W|} \tag{5}$$

Note that although this factor is easy to calculate from the training data, and allow scalability to sparse data regimes, this factor can results in predictions values greater than 1, we propose to handle this by clipping the values to 1. However, we would like to think in the future about better options. Moreover, the assumption of the independence of  $l$  and  $x$  is very strong, because satellite images are inherently tied to the location they correspond to. We suspect this may be a reason for the degradation in performance in table 1 when using this factor. We aim to explore more about cases when this assumption doesn't hold.

#### A.4 VALIDATION RESULTS

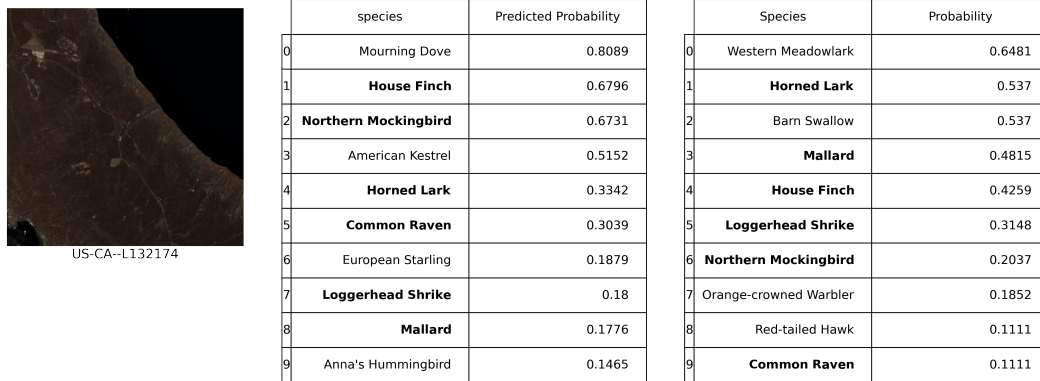
We provide results on the validation set of the proposed baseline models in Table 2.

**Table 2:** Results for the proposed models on the validation set. The ranking of models is the same as on the test set. MSE and MAE are reported  $\times 10^3$  and  $\times 10^2$  respectively.

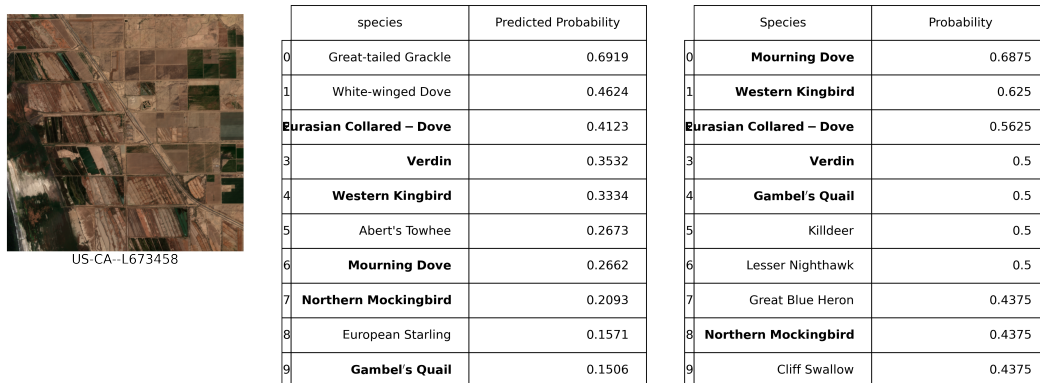
Method	MSE [ $10^{-3}$ ]	MAE [ $10^{-2}$ ]	Top- $k$	Top-30	Top-10
Mean ER	7.393	2.944	50.75	43.91	26.45
Env baseline	$4.93 \pm 0.0$	$2.08 \pm 0.0$	$68.60 \pm 0.00$	$61.64 \pm 0.01$	$42.41 \pm 0.01$
RGBNIR	$5.5 \pm 0.05$	$1.99 \pm 0.00$	$66.9 \pm 0.30$	$59.7 \pm 0.30$	$41.83 \pm 0.20$
RGBNIR + Env	$4.56 \pm 0.02$	$1.7 \pm 0.0$	$73.0 \pm 0.00$	$66.0 \pm 0.00$	$48.0 \pm 0.00$
RGBNIR + Env + LE	$4.9 \pm 0.04$	$1.7 \pm 0.01$	$72.5 \pm 0.00$	$64.5 \pm 0.70$	$44.0 \pm 0.00$
RGBNIR + Env + RM	$4.6 \pm 0.02$	$1.8 \pm 0.00$	$73.3 \pm 0.00$	$66.6 \pm 0.50$	$48.2 \pm 0.30$
RGBNIR + Env + SM	$13.6 \pm 0.45$	$2.78 \pm 0.11$	$64.5 \pm 1.56$	$54.2 \pm 1.75$	$31.0 \pm 1.41$
RGBNIR + Env + RM + WL	$4.76 \pm 0.00$	$1.77 \pm 0.00$	$72.4 \pm 0.01$	$65.4 \pm 0.02$	$46.01 \pm 0.02$

### A.5 TOP 10 SPECIES PREDICTIONS AT DIFFERENT HOTSPOTS

All predictions reported at those obtained with the RGBNIR + Env + RM + WL model. We sampled hotspots from our test set and evaluated the model’s performance in predicting the top 10 bird species at each location. We selected hotspots from geographically distinct regions, featuring different land cover types, including the Pacific Coast (California), Northeast (Maine) and Central US (Ohio). The following figures illustrate the results from these regions.




**Figure 4:** Hotspot and species at San Clemente Island, California: Our model predicts the presence of various species in this region, including the house finch and loggerhead shrike, which are consistent with the ground truth and highly reported on eBird.



**Figure 5:** Hotspot and species at Salton Sea, California: Our models predict the presence of various species in this region, including the western kingbird, verdin, and Eurasian collared dove. These species are consistent with the ground truth and are highly reported on eBird for this location.

### A.6 ANALYZING PREDICTIONS

In addition to adaptive Top- $k$ , Top-10 and Top-30 accuracy, which evaluate the model globally across species, we consider precision and recall per species to identify which are the species for which our models perform the best or worse. Our targets are continuous but to compute the precision and recall, we treat them in the same way as for our accuracy metrics. In Table 4, we report the top 50 species according to recall metric for the RGBNIR + Env + RM model. These are the species whose presence the model does not fail to predict. We also list some species which have 0 recall in Table 3 but note that all species which have low recall are seen in very few hotspots.



US-ME--L350189

	species	Predicted Probability
0	<b>Song Sparrow</b>	0.6435
1	Red-winged Blackbird	0.6036
2	<b>Yellow Warbler</b>	0.4973
3	Mallard	0.4899
4	American Crow	0.4446
5	Herring Gull	0.4227
6	American Robin	0.4221
7	Mourning Dove	0.3833
8	American Goldfinch	0.3614
9	Common Grackle	0.3471

	Species	Probability
0	Great Black-backed Gull	0.75
1	<b>Song Sparrow</b>	0.625
2	Double-crested Cormorant	0.625
3	Piping Plover	0.5
4	Gray Catbird	0.5
5	Snowy Egret	0.5
6	<b>Yellow Warbler</b>	0.5
7	Tree Swallow	0.5
8	Canada Goose	0.375
9	Black-bellied Plover	0.375

**Figure 6:** Hotspot and species at Rachel Carson National Wildlife Refuge, Maine: Our models correctly identified the song sparrow and the yellow warbler in this region, which exist in the top-10 species in both the ground truth and predicted list. However, it is worth noting that shoreline species are absent from the predicted top-10 list, representing a failure of our model. This may be a consequence of the partial cloud cover in the satellite image.




US-ME--L7072367

	species	Predicted Probability
0	Red-winged Blackbird	0.6798
1	<b>Song Sparrow</b>	0.6249
2	<b>Common Yellowthroat</b>	0.6018
3	<b>Yellow Warbler</b>	0.5948
4	Tree Swallow	0.5402
5	<b>American Robin</b>	0.5232
6	Common Grackle	0.4937
7	Eastern Kingbird	0.4921
8	Cedar Waxwing	0.491
9	<b>American Goldfinch</b>	0.4907

	Species	Probability
0	Veery	0.6667
1	Gray Catbird	0.6111
2	Alder Flycatcher	0.6111
3	<b>Common Yellowthroat</b>	0.5556
4	<b>Song Sparrow</b>	0.5556
5	<b>American Goldfinch</b>	0.5556
6	<b>American Robin</b>	0.5556
7	Black-capped Chickadee	0.5
8	<b>Yellow Warbler</b>	0.5
9	Red-eyed Vireo	0.5

**Figure 7:** At Valentine Farm Conservation Center in Maine, our models correctly identified the song sparrow, common yellowthroat, yellow warbler, and American goldfinch. These species exist in both the ground truth and are consistent with Maine’s ecology.



US-OH--L1188602

	species	Predicted Probability
0	<b>American Robin</b>	0.8373
1	Song Sparrow	0.8058
2	<b>Northern Cardinal</b>	0.7935
3	<b>Gray Catbird</b>	0.7639
4	American Goldfinch	0.7492
5	<b>European Starling</b>	0.7173
6	Brown-headed Cowbird	0.6497
7	Indigo Bunting	0.6149
8	<b>Red - winged Blackbird</b>	0.6051
9	<b>Cedar Waxwing</b>	0.6037

	Species	Probability
0	Blue Jay	0.8
1	House Sparrow	0.8
2	Chipping Sparrow	0.8
3	<b>European Starling</b>	0.8
4	<b>American Robin</b>	0.8
5	<b>Red - winged Blackbird</b>	0.6
6	<b>Cedar Waxwing</b>	0.6
7	<b>Northern Cardinal</b>	0.6
8	<b>Gray Catbird</b>	0.6
9	Barn Swallow	0.6

**Figure 8:** At West Park Cemetery in Ohio, our models accurately identified the American robin, northern cardinal, gray catbird, and six of the top ten species in this hotspot. This is consistent with the species reported on eBird.

Species	Recall	Occurrences
Long-eared Owl	0.0	4.0
Pacific Golden-Plover	0.0	1.0
American Golden-Plover	0.0	2.0
Snow Bunting	0.0	1.0
Tufted Puffin	0.0	1.0
Greater Scaup	0.0	9.0
Black-footed Albatross	0.0	1.0
Red Phalarope	0.0	1.0
Northern Fulmar	0.0	2.0
Harris’s Sparrow	0.0	1.0

**Table 3:** Example species with 0 recall on the test set, which means the model misses them but we also note that they are encountered in very few hotspots.

## B ACKNOWLEDGEMENTS

We are grateful to Srishti Yadav for initial dataset exploration, Sal Elkafrawy for investigating the initialization of our models, and Dan Morris and Nebojsa Jojic for insights in remote sensing. This project was supported by a Microsoft-Mila collaboration grant and a Microsoft AI for Earth cloud compute grant. Mélisande Teng received the support of the NSERC-CREATE LEADS program and Benjamin Akera received support from IVADO. David Rolnick was supported in part by the Canada CIFAR AI Chairs Program. The authors also acknowledge material support from NVIDIA in the form of computational resources, and are grateful for technical support from the Mila IDT team in maintaining the Mila Compute Cluster.

## C CONTRIBUTIONS

We are happy to say that this work was a truly collaborative effort, and that this work would not have been possible without any one of us. We will just highlight a few of the many contributions each one of us has made to this project.

In particular, **Mélisande Teng** played a major role in the problem and task formulation through multiple iterations since the beginning of the project and led the dataset building efforts. **Amna Elmustafa** played an instrumental role in incorporating geographical information in our models. **Benjamin Akera** contributed to the analysis of the predictions of our models with thorough inspections of the hotspots. **Hugo Larochelle** provided valuable guidance, not only with insightful discussions about methodology but also with much appreciated moral support. **David Rolnick** was the ideal supervisor for this project and originally came up with the idea of leveraging eBird and remote sensing data for understanding ecosystems better.

Species	Recall	Precision	Occurrences
Varied Bunting	1.000	1.000	4
Botteri's Sparrow	1.000	1.000	5
Olive Sparrow	1.000	1.000	4
Pink-footed Shearwater	1.000	1.000	1
Sprague's Pipit	1.000	1.000	1
Cassia Crossbill	1.000	1.000	1
Red-faced Warbler	1.000	0.875	7
Black-crested Titmouse	1.000	0.857	6
Yellow-eyed Junco	1.000	0.833	5
Abert's Towhee	1.000	0.813	13
Bicknell's Thrush	1.000	0.800	4
Rufous-winged Sparrow	1.000	0.800	4
Great Kiskadee	1.000	0.800	4
Blue-throated Mountain-gem	1.000	0.800	4
Gambel's Quail	1.000	0.741	20
Verdin	1.000	0.714	25
Reddish Egret	1.000	0.700	7
Arizona Woodpecker	1.000	0.667	4
Mexican Chickadee	1.000	0.667	2
Bridled Titmouse	1.000	0.636	7
Gray Hawk	1.000	0.571	4
Olive Warbler	1.000	0.500	3
Baird's Sparrow	1.000	0.500	1
Mexican Duck	1.000	0.500	2
Mountain Plover	1.000	0.500	1
Buff-breasted Flycatcher	1.000	0.500	2
Snail Kite	1.000	0.500	1
Scaled Quail	1.000	0.500	5
Wilson's Plover	1.000	0.444	4
Nanday Parakeet	1.000	0.429	3
California Gnatcatcher	1.000	0.400	2
Whiskered Screech-Owl	1.000	0.400	2
Green Kingfisher	1.000	0.333	1
Golden-cheeked Warbler	1.000	0.333	1
Gray-crowned Rosy-Finch	1.000	0.333	1
Thick-billed Longspur	1.000	0.250	1
Kirtland's Warbler	1.000	0.250	1
Brown-capped Rosy-Finch	1.000	0.250	1
Red-cockaded Woodpecker	1.000	0.200	1
Northern Cardinal	0.985	0.889	480
Mourning Dove	0.982	0.897	720
Common Grackle	0.978	0.882	504
Western Gull	0.978	0.786	45
Black-billed Magpie	0.972	0.812	71
American Robin	0.969	0.914	678
Ladder-backed Woodpecker	0.968	0.682	31
American Crow	0.966	0.871	651
Carolina Chickadee	0.966	0.778	174
Blue Jay	0.961	0.897	490
Carolina Wren	0.956	0.801	294

**Table 4:** Top 50 species with highest recall on the test set, recall and precision per species scores and number of test set hotspots in which species have non-zero encounter rates, for predictions of the RGBNIR + Env + RM model. Note that only species which were encountered in the test set are considered (566/684). Rows are sorted by recall performance.