

On the selection and effectiveness of pseudo-absences for species distribution modeling with deep learning

Robin Zbinden^{a,*}, Nina van Tiel^a, Benjamin Kellenberger^b, Lloyd Hughes^a, Devis Tuia^a

^a Ecole Polytechnique Fédérale de Lausanne (EPFL), Environmental Computational Science and Earth Observation (ECEO) Laboratory, Switzerland

^b Yale University, Jetz Lab, USA

ARTICLE INFO

Keywords:

Species distribution modeling
Neural networks
Presence-only data
Pseudo-absences
Deep learning

ABSTRACT

Species distribution modeling is a highly versatile tool for understanding the intricate relationship between environmental conditions and species occurrences. However, the available data often lacks information on confirmed species absence and is limited to opportunistically sampled, presence-only observations. To overcome this limitation, a common approach is to employ pseudo-absences, which are specific geographic locations designated as negative samples. While pseudo-absences are well-established for single-species distribution models, their application in the context of multi-species neural networks remains underexplored. Notably, the significant class imbalance between species presences and pseudo-absences is often left unaddressed. Moreover, the existence of different types of pseudo-absences (e.g., random and target-group background points) adds complexity to the selection process. Determining the optimal combination of pseudo-absences types is difficult and depends on the characteristics of the data, particularly considering that certain types of pseudo-absences can be used to mitigate geographic biases. In this paper, we demonstrate that these challenges can be effectively tackled by integrating pseudo-absences in the training of multi-species neural networks through modifications to the loss function. This adjustment involves assigning different weights to the distinct terms of the loss function, thereby addressing both the class imbalance and the choice of pseudo-absence types. Additionally, we propose a strategy to set these loss weights using spatial block cross-validation with presence-only data. We evaluate our approach using a benchmark dataset containing independent presence-absence data from six different regions and report improved results when compared to competing approaches.

1. Introduction

In a world where climate change and human activities increasingly threaten numerous species and their habitat, there is a growing need to understand the factors that determine the presence of a species at a specific geographic location. Addressing this need, species distribution models (SDMs) seek to unveil the complex relationship between environmental conditions at a given location and the likelihood of a species occurring there (Elith and Leathwick, 2009; Franklin, 2010). In particular, SDMs are used to predict the geographic range of species, thereby playing a pivotal role in supporting conservation and restoration policies (Guisan et al., 2013; Sofaer et al., 2019). Developing reliable SDMs, however, presents several challenges, such as limited data availability

and inherent selection bias within the observations used (Beck et al., 2014; Mesaglio and Callaghan, 2021). One typical issue is the potential geographic bias arising from variations in sampling efforts across different areas (Kadmon et al., 2004). These limitations not only hinder the predictive accuracy of these models, but also impede their ability to generalize effectively to other regions (Elith et al., 2010). The scarcity of records for rare species further diminishes the significance of model predictions.

Nevertheless, recent initiatives in community science, exemplified by platforms like *iNaturalist*,¹ *eBird*,² and *Pl@ntNet*,³ have revolutionized the field and consolidated large numbers of species records contributed by enthusiasts and experts alike. The *iNaturalist* platform, for instance, comprises an impressive collection of over 150 million species

* Corresponding author.

E-mail address: robin.zbinden@epfl.ch (R. Zbinden).

¹ <https://www.inaturalist.org/>

² <https://ebird.org/home>

³ <https://plantnet.org>

observations spanning more than 400,000 species. These rich data resources present tremendous opportunities to significantly enhance the performance of SDMs and address their aforementioned associated challenges (Botella et al., 2023; Teng et al., 2023). However, they frequently consist of *presence-only* (PO) observations, providing information exclusively about the presence of species, while lacking any corresponding data regarding their absence (Elith et al., 2020; Pearce and Boyce, 2006). This disparity emerges from the difficulty of gathering data on species absence, in contrast to the more opportunistic nature of recording species presence (Franklin, 2010). From a machine learning perspective, this constraint implies that only positive samples are at our disposal (Bekker and Davis, 2020).

However, most statistical and machine learning techniques require the incorporation of negative samples to effectively discriminate between species' presence and absence. Different methods have been devised to tackle this challenge, with one common approach being the use of *pseudo-absences* sampling. This technique involves designating selected geographic locations as negative samples, even though there is no guarantee that the environmental conditions at these locations are unfavorable habitats for the target species. These pseudo-absences, also known as *background points* or *pseudo-negatives*, are usually sampled uniformly across the geographic space (*random background points*) or among the presence locations of other species (*target-group background points*), aiming to account for the sampling bias within the presence data (Ponder et al., 2001; Phillips et al., 2009; Botella et al., 2020). An illustration of these two types of pseudo-absences is provided in Fig. 1.

Incorporating pseudo-absences into model training generally involves combining them with the set of presence data (Valavi et al., 2022). However, this approach raises important questions regarding the optimal *quantity* and *type* of these pseudo-absences, potentially leading to class imbalance issues between presences and pseudo-absences and biased predictions if not carefully managed. While several studies have addressed these questions (Barbet-Massin et al., 2012; Phillips et al., 2009; Stokland et al., 2011), the answers tend to be context-dependent, making it difficult to establish general guidelines. This challenge is particularly pronounced for models that have not been the focus of these studies, as is the case with *neural networks*, the most popular family of methods of the recent wave of machine learning in ecology (Tuia et al., 2022).

We aim to bridge this gap since the use of neural networks and *deep learning* techniques stands as a promising and efficient approach for SDMs (Botella et al., 2023; Davis et al., 2023; Teng et al., 2023). These

methods have shown significant advancements when applied to large datasets, often outperforming traditional machine learning techniques in various fields (Brown et al., 2020; Christin et al., 2019; Jumper et al., 2021; Krizhevsky et al., 2012). A distinctive advantage of deep learning approaches lies in their inherent flexibility, enabling the simultaneous integration of diverse data types and the modeling of multiple species, a capability that is increasingly vital in ecological research. Recent research emphasizes the potential of deep learning methods in SDMs, showing performance that matches or exceeds traditional techniques. In particular, this improvement seems to extend to both presence-absence (Chen et al., 2016; Teng et al., 2023) and presence-only datasets (Botella et al., 2023; Deneu et al., 2021).

In this study, we investigate how to adequately employ pseudo-absences with neural networks. Our emphasis lies on multi-species models, observing that the incorporation of pseudo-absences seamlessly aligns with the multi-label prediction capabilities of neural networks. This integration is achieved through a modification of the training loss function that is used to optimize the neural network parameters. Specifically, we enhance the *full* loss function proposed by Cole et al. (2023) by assigning appropriate weights within the loss function for presences and pseudo-absences. This adjustment addresses the class imbalance between pseudo-absences and presences and recognizes that some pseudo-absences can be more informative than others in some situations (Phillips et al., 2009). Since our loss function involves weights whose optimal values may vary depending on the dataset, we show how to tune them using presences-only data with spatial block cross-validation (Roberts et al., 2017).

We assess the effectiveness of our approach using a well-established benchmark dataset in SDMs, which comprises data from six distinct regions (Elith et al., 2020). This dataset includes independent presence-absence data for model evaluation, marking our work as the first to assess different SDMs loss functions on presence-absence data. In addition, previous evaluations have predominantly focused on global scales. While crucial for establishing global biodiversity trends, finer distribution maps are essential for informing local conservation policies. The diverse regions considered in our study span various scales, from local to continental. Our approach not only exhibits superior performance compared to alternative methods for neural networks, especially when coupled with an effective cross-validation methodology, but also demonstrates adaptability to accommodate the specific characteristics of diverse datasets. This includes addressing issues such as class imbalance and sampling bias.

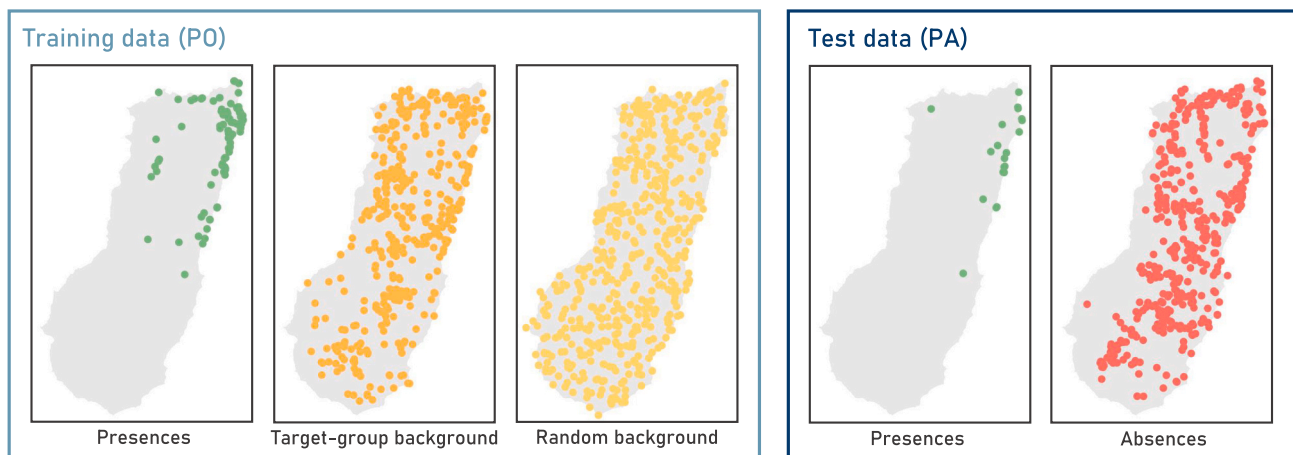


Fig. 1. Modeling species distributions often involves utilizing presence-only (PO) data, where information about species absences is unavailable. To apply machine learning techniques in such situations, pseudo-absences are used as a contrast to presence data. There are primarily two types of pseudo-absences: target-group background points and random background points. Target-group background points consist of presences of other species, sharing a similar sampling bias, while random background points are uniformly sampled within the area. In this paper, we emphasize the critical importance of how these different types are managed during training for optimal performance, especially when dealing with neural networks. We support our approach by evaluating it on an independent test set comprised of presence-absence (PA) data.

2. Background

2.1. Pseudo-absences in single-species models

Pseudo-absences were introduced as a means to train single-species models in situations where data on species absence is unavailable, as is the case with presence-only datasets (Pearce and Boyce, 2006). These pseudo-absences serve the essential purpose of creating a contrast with presence data, thus preventing the model from converging to a trivial solution where it predicts the presence of the species everywhere. Extensive research has delved into the integration of pseudo-absences in single-species modeling (Barbet-Massin et al., 2012; Jarnevich et al., 2017; Senay et al., 2013; Wisz and Guisan, 2009). These studies predominantly center on fundamental questions related to the required quantity and the appropriate type of pseudo-absences.

2.1.1. Quantity of pseudo-absences

Conventional approaches sample a specific number of pseudo-absences, typically fixed (Elith et al., 2006) or relative to the number of available occurrences per species (Valavi et al., 2021). That said, the optimal number often depends on the model in use. For instance, in the case of *Maxent* (Phillips et al., 2006), incorporating a larger number of pseudo-absences is generally more advantageous, up to a saturation point (Phillips and Dudík, 2008). For classification models such as *Random Forests*, it is recommended to also sample a large number of pseudo-absences, but then to employ bootstrapping to ensure a balanced count of presences and pseudo-absences for each tree (Barbet-Massin et al., 2012; Valavi et al., 2021).

2.1.2. Type of pseudo-absences

Diverse methods exist for sampling pseudo-absences. The standard approach involves randomly selecting points within the study area, which we refer to as “random background points” throughout this paper. Some alternative approaches restrict the sampling of pseudo-absences to areas geographically distant from known presences (Barbet-Massin et al., 2012; Mateo et al., 2010). Another method focuses on the environmental space, where a first habitat suitability map is constructed using presence data only and envelope modeling (Araújo and Peterson, 2012), and pseudo-absences are drawn from areas predicted to have low suitability (Engler et al., 2004). Nonetheless, several studies have indicated that these methods do not consistently outperform the random background point approach and may even amplify biases originating from the presence data (Lyu et al., 2022; Wisz and Guisan, 2009).

Alternatively, the presences of other similar species can be used as pseudo-absences; these are referred to “target-group background points” in this context. This approach often shows significant improvement in performance compared to random background points sampling, particularly for datasets marked by substantial sampling bias (Phillips et al., 2009; Ponder et al., 2001). Several factors contribute to the often superior performance of these pseudo-absences, with a key factor being that target-group background points often incorporate a bias similar to that of the presences of the target species (Botella et al., 2020; Hertzog et al., 2014). This alignment of biases enables a finer discrimination between presence and absence.

A visual example of the difference in distribution between the target-group and random background points is represented in the left side of Fig. 1. In practice, studies rely on either random background points or target-group background points, but seldom combine both. However, combining them may potentially enhance model performance, an aspect we investigate in this work.

2.2. Pseudo-absences in multi-species neural network models

In this section, our focus shifts towards modeling the distributions of multiple species simultaneously, often referred to as multi-species distribution models (Poggiato et al., 2021). In machine learning

terminology, this corresponds to a multi-label classification task. This differs from the more conventional multi-class classification, where only one species would be assumed to be present at a given location. In the context of presence-only data, not only are we restricted to presence records, but we often have only one species occurrence record per location, as exemplified by datasets like iNaturalist. We essentially find ourselves in a situation described as *single positive multi-label learning* by Cole et al. (2021). In this setup, we have knowledge of one positive label (observed presence) and unknown labels for (usually all) other species. However, similar to single-species modeling, we still need to employ pseudo-absences to prevent the model from incorrectly predicting the presence of all species at every location. The simplest way to approach pseudo-absences in this case is to assume that the missing labels are absences. This approach, known as “assume negative”, corresponds to the definition of target-group background points. As target-group background points alone may not adequately cover the entire geographic area, it can be beneficial to include additional random background points. In the context of neural networks, an effective way to combine these two types of pseudo-absences is by adjusting the training loss function.

In the field of machine learning, training a model often involves minimizing a loss function that quantifies the model’s error on training data. Generally, deep learning-based SDMs are trained with the binary cross-entropy loss (Benkendorf and Hawkins, 2020; Deneu et al., 2021; Zbinden et al., 2023; Zhang et al., 2022), with the target-group background points often used as the de facto negative samples when absence data is unavailable. Nevertheless, the loss function can be modified to reflect the specificities of the problem. In our case, we can adapt the loss function to inform the model about how pseudo-absences are integrated. This can be achieved by adjusting or assigning weights to the different components of the loss function. Recently, Cole et al. (2023) introduced the following *full assume negative loss*, designed to account for both target-group background points and random background points:

$$\mathcal{L}_{\text{full}}(\mathbf{y}, \hat{\mathbf{y}}) = -\frac{1}{S} \sum_{s=1}^S \left[\underbrace{1_{[y_s=1]} \lambda \log(\hat{y}_s)}_{\text{presences}} + \underbrace{1_{[y_s=0]} \log(1 - \hat{y}_s)}_{\text{target-group background}} + \underbrace{\log(1 - \hat{y}_s')}_{\text{random background}} \right] \quad (1)$$

where y_s is 1 if species s is present (the species has been observed) and 0 otherwise, \hat{y}_s denotes the predicted suitability score for species s (ranging from 0 to 1), \hat{y}_s' represents the model’s prediction for species s at a random location, S is the number of species considered, and $1_{[\cdot]}$ is the indicator function, returning 1 if the condition inside the brackets is true and 0 otherwise. It is important to note that this loss function is defined for a single presence location, and therefore the overall loss for the entire training dataset is computed by averaging over all the presence locations in the training set. We observe that $\mathcal{L}_{\text{full}}$ is composed of three distinct terms, each corresponding to a different type of sample. The first term addresses presences, favoring a high score for the species that is present, and is weighed by a coefficient λ to compensate for the relatively small impact of presences in the loss function. The second term represents target-group background points and assumes that species not observed are absent. Finally, the third term corresponds to random background points, where it is assumed that all species are absent. Consequently, $\mathcal{L}_{\text{full}}$ effectively combines target-group and random background points in a straightforward, clear, and complementary manner. This combination enables informing the model about the spatial bias through the target-group background points, while retaining the advantage of using random background points to cover the entire region of interest.

2.2.1. Limitations of $\mathcal{L}_{\text{full}}$

Primarily, this loss function fails to adequately address the issue of class imbalance within datasets. Two key aspects of class imbalance need to be tackled here:

1. There is often a disparity between the quantity of pseudo-absences required and the number of presences per species in typical datasets, often with a larger number of pseudo-absences. This imbalance can impede the effective training of machine learning models (Johnson and Khoshgofaar, 2019) and has been shown to lead to suboptimal results for SDMs (Benkendorf et al., 2023). While the introduction of the weight λ , whose value is usually proportional to the number of species (Cole et al., 2023; Mac Aodha et al., 2019), aims to alleviate this issue, it does not address the second aspect described below.
2. We have to ensure that each species is equally weighted within the loss function. The uniformity of λ across all species fails to account for variations in the number of presences among species. For example, Fig. 2a shows that a significantly higher number of reported presences is available for some species, which results in a long-tailed distribution. In their work, Cole et al. (2023) address this issue by downsampling the presences for species with a large number

3. Methods

We introduce a new loss function for SDMs, described in Section 3.1. This loss function incorporates dataset-dependent weights that get tailored to the specific characteristics of the data. In Section 3.2, we advocate for the use of block cross-validation to determine the optimal values of these weights. While the loss function is applicable to any deep learning method, we employ in our experiments a particular multi-species neural network model that is detailed in Section 3.3. Subsequently, we provide an overview of the dataset in use in Section 3.4, and elaborate on our evaluation methodology in Section 3.5.

3.1. Full weighted loss function

In light of the considerations discussed above, we propose an extension of $\mathcal{L}_{\text{full}}$. We refer to this new loss function as the *full weighted loss function*, defined as follows:

$$\mathcal{L}_{\text{full-weighted}}(\mathbf{y}, \hat{\mathbf{y}}) = -\frac{1}{S} \sum_{s=1}^S \left[\underbrace{1_{[y_s=1]} \lambda_1 w_s \log(\hat{y}_s)}_{\text{weighted presences}} + \underbrace{1_{[y_s=0]} \lambda_2 \frac{1}{\left(1 - \frac{1}{w_s}\right)} \log(1 - \hat{y}_s)}_{\text{weighted target-group background}} + \underbrace{(1 - \lambda_2) \log(1 - \hat{y}_s)}_{\text{weighted random background}} \right]. \quad (2)$$

of occurrences. However, this approach underutilizes available data, potentially leading to a decline in performance.

In addition to the class imbalance issue, the target-group background points and random background points are assigned the same weight within $\mathcal{L}_{\text{full}}$. However, empirical studies suggest that target-group background points typically provide more valuable information than random ones (Botella et al., 2020; Phillips et al., 2009) when there is a significant bias in the presence data, as it is often the case with presence-only data (see Fig. 2b and Beck et al. (2014), Mesaglio and Callaghan (2021)). Ideally, the loss function should offer flexibility in determining the proportion of target-group background points relative to random background points, since this proportion is dependent on the specific dataset under consideration. Our method tackles all these aspects and proposes a new loss function, at the same time accounting for species observations imbalance, different types of pseudo-absences and their relative proportions.

The key distinction between $\mathcal{L}_{\text{full}}$ and our new formulation lies in the introduction of weighting terms, with the addition of the coefficients w_s , λ_1 , and λ_2 . We will now provide a detailed explanation of each of these weights.

3.1.1. Species weights

To address the class imbalance issues, we introduce the species weights denoted as w_s . For each species s , the weight is defined as:

$$w_s = \frac{n}{n_{p(s)}} = \frac{1}{\text{freq}(s)} \quad (3)$$

where $n_{p(s)}$ represents the number of presence records for species s and n is the total number of presence locations in the training set. We explore alternative definitions of w_s in Section 4. In this formulation, w_s corresponds to the inverse of the frequency of species s . Weighting the species' presences this way ensures that the effective contribution of the presence records to the loss is equivalent for every species. Nevertheless,

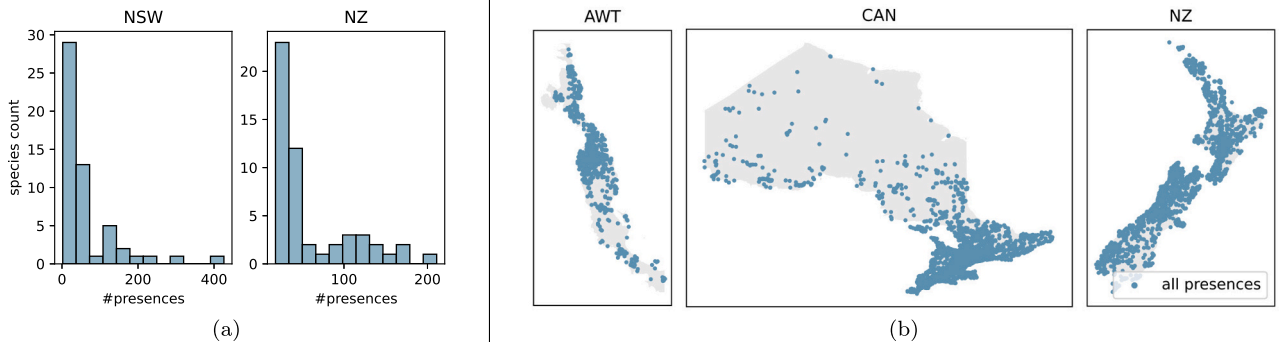


Fig. 2. Species occurrence records generally exhibit sampling biases, as depicted here by training presences in the dataset from Elith et al. (2020). (a) The number of presences per species follows a long-tailed distribution, with many species having only a limited number of available observations. To address this issue, we incorporate a species weight w_s for each species in our loss function. (b) The geographic distribution of the presence records of all species shows varying biases across regions. We introduce the pseudo-absence weight λ_2 to mitigate this bias. Additional plots for the remaining regions, not presented here, can be found in Appendix A.

solely applying this weight to presences is insufficient to guarantee the unbiased treatment of all species by the model, as the number of target-group background points may be disproportionately high for species with few observations. Hence, to ensure that their contribution is consistent for every species, we also weight target-group background points with the following term:

$$\frac{1}{\left(1 - \frac{1}{w_s}\right)} = \frac{n}{n_{\text{tgbg}(s)}} = \frac{1}{1 - \text{freq}(s)} \quad (4)$$

where $n_{\text{tgbg}(s)} = n - n_{\text{p}(s)}$ denotes the number of target-group background points. Importantly, by weighting both presences and target-group background points, we guarantee that the individual contributions to the loss function of presences, target-group background points, and random background points in the training set are equivalent. In this way, we avoid assuming that the number of presences in the training set reflects the prevalence of the species, as prevalence cannot be determined solely from presence data (Hastie and Fithian, 2013). However, if additional information about prevalence is available, it can be integrated by scaling the species prediction scores a posteriori, with the flexibility to scale differently for each species.

3.1.2. Pseudo-absence weight

We introduce the pseudo-absence weight λ_2 as a way to indicate which type of pseudo-absence should have more importance in the learning process. Some studies have demonstrated that using target-group background points may lead to more accurate predictions (Botella et al., 2020; Phillips et al., 2009), but including a portion of random background points can be useful for covering the entire geographic area of interest (Cole et al., 2023; VanDerWal et al., 2009). With $\lambda_2 \in [0, 1]$, our proposed loss modulates the emphasis between these two types of pseudo-absence. When $\lambda_2 = 0$, only random background points are used, while $\lambda_2 = 1$ means that only target-group background points are employed. The full loss of Cole et al. (2023) is equivalent to $\lambda_2 = 0.5$, assuming an equal weight on target-group background points and random background points.

3.1.3. Presence weight

In addition, we introduce the presence weight λ_1 as a means to adjust the weighting of the presences compared to the pseudo-absences. When $\lambda_1 > 1$, more emphasis is placed on correctly classifying presences, whereas when $\lambda_1 < 1$, the focus shifts to pseudo-absences. Notably, considering that presence data should be more reliable than pseudo-absences, a relatively higher value of λ_1 can prove advantageous as it instructs the model to prioritize correctly classifying presences. This weight also makes $\mathcal{L}_{\text{full-weighted}}$ a generalization of $\mathcal{L}_{\text{full}}$, enabling direct comparison. $\mathcal{L}_{\text{full}}$ is recovered by setting $\lambda_1 = S/2$, $\lambda_2 = 0.5$, and removing the species weights w_s .

3.2. Tuning loss weights

Our full weighted loss function includes the weights λ_1 and λ_2 , whose optimal values are data-dependent and, as hyperparameters, require careful tuning. In machine learning, hyperparameter values are typically

determined through cross-validation, involving the partitioning of the training data and designating a portion as the validation set. The selection of hyperparameter values, or equivalently the selection of the model, is then based on the performance on this validation set. However, constructing an effective validation set with presence-only data is challenging, as evaluating model performance with such data may lead to choosing a model that makes biased predictions (El-Gabbas and Dormann, 2018). To alleviate this issue, we employ spatial block cross-validation (Roberts et al., 2017; Valavi et al., 2023), which involves spatially splitting the presence observations into training and validation sets. This approach makes it difficult for the model to perform well on the validation set, favoring models capable of generalizing to unseen areas (Smith et al., 2021). This is particularly important since presence-only data often exhibits biases towards specific areas, and we need to assess the model's ability to counter them.

Hence, we perform k -fold block cross-validation, dividing the region of interest into 5×5 geospatial blocks and assigning them to k distinct folds, such that each fold contains approximately the same number of presence observations. We choose $k = 5$, resulting in that for each partition, the model is trained on 80% of the presences and validated on the remaining 20%. This procedure is equivalent to the spatial blocks based on rows and columns of Valavi et al. (2018), and is illustrated in Fig. 3. The model is then evaluated on each fold and the results are averaged. The model with the best performance, determined here by the mean AUC (Area Under the receiver operating characteristic Curve) over all the species, is selected. The target-group background points are used as pseudo-absences to compute the AUC on the validation set, similarly to Valavi et al. (2023).

3.3. Model and training details

We adopt a configuration similar to the multi-species model presented in Zbinden et al. (2023). Instead of developing individual models for each species, we train a single *multi-layer perceptron* (MLP) model for each region (details about the six regions considered are given in Section 3.4). These models predict suitability scores for every species within the region following a multi-label classification approach. MLPs are a type of neural network consisting of multiple layers of fully-connected neurons with non-linear activation functions between layers (Gorishniy et al., 2021), enabling them to capture complex interactions among the input environmental covariates. This architecture is well-suited for the tabular format of the dataset in use (see Section 3.4). For the model and training hyperparameters, we stick to standard values (Gorishniy et al., 2021; Mac Aodha et al., 2019), while also incorporating recent advancements in deep learning to enhance performance. Specifically, our MLP architecture consists of four layers, each containing 512 neurons and connected with residual connections (He et al., 2016). We employ batch normalization (Ioffe and Szegedy, 2015) and the Rectified Linear Unit (ReLU) activation function in all layers except the final one, where instead a sigmoid function is used to enable multi-label classification. The model is trained with a batch size of 256 for 30 epochs using the AdamW optimizer (Loshchilov and Hutter, 2017). Both the weight decay and learning rate are set to 0.0001. Additionally, we employ a learning rate scheduler with exponential decay of 0.95, and introduce dropout




Fig. 3. k -fold block cross-validation (Roberts et al., 2017) is used to find the optimal value of the pseudo-absence weight λ_2 . It involves the spatial partitioning of presence observations into the training and validation sets, comprising respectively 80% and 20% of the samples. The presence records considered here pertain to the Swiss region of the dataset described in Section 3.4.

(Srivastava et al., 2014) with a probability of 0.01.

3.4. Dataset

We use the benchmark dataset from Elith et al. (2020), which comprises occurrence data for 226 anonymized species from six regions around the world: Australian Wet Tropic (AWT), Canada (CAN), New South Wales (NSW), New Zealand (NZ), South America (SA), and Switzerland (SWI). Recently made public, this dataset has already been employed to evaluate and compare different methods (Elith et al., 2006; Valavi et al., 2022), as well as for conducting analyses of target-group background points (Phillips et al., 2009). The training set consists of opportunistically sampled presence-only data, while the test set is built with full presence-absence data for each species (Fig. 1, right). As predictors, the dataset contains a varying set of 10 to 13 environmental covariates for each region, including both climatic and pedological variables (see Elith et al. (2020) for a full description).

This dataset is also representative of the biases and challenges often encountered when training SDMs. Notably, the number of presences per species follows a long-tail distribution, as illustrated in Fig. 2a. Additionally, the geographic distribution is often biased, particularly in certain areas of some regions. For instance, as observed in Fig. 2b and quantified by Phillips et al. (2009), there is a substantial bias in the distribution of presences for the AWT and CAN regions, while NZ exhibits less bias. Finally, some regions encompass species belonging to different biological groups, such as plants or birds. While previous studies used only presences of species within the same group as target-group background points (Phillips et al., 2009; Ponder et al., 2001), for simplicity we utilize all species within each region to train our models. As a result, we have a total of six models, one for each region.

3.5. Evaluation and baselines

We compare our approach with various other prediction methods and loss functions. Firstly, we examine the different loss functions proposed by Cole et al. (2023). The $\mathcal{L}_{\text{full}}$ presented in Section 2.2 is a combination of two losses: the SSDL (same species, different location) loss, which exclusively uses the random background points and presences, and the SLDS (same location, different species) loss, which only uses the target-group background points and presences. Notably, our full weighted loss generalizes these three loss functions by appropriately setting the λ_1 and λ_2 weights. Specifically, $\mathcal{L}_{\text{full-weighted}}$ is similar to the SSDL loss when $\lambda_1 = S$ and $\lambda_2 = 0$, the SLDS loss corresponds to $\lambda_1 = S$ and $\lambda_2 = 1$, and the full loss is equivalent when $\lambda_1 = S/2$ and $\lambda_2 = 0.5$. In each case, no species weights are applied. We then train the multi-species model defined in Section 3.3 on these loss functions. For completeness, in Appendix B, we also discuss and evaluate other machine learning loss functions commonly used in similar settings,

particularly the categorical cross-entropy loss, which has been employed in previous SDMs studies (Deneu et al., 2021; Zhang et al., 2022).

Additionally, we include results for the Maxent and Boosting Regression Trees models from Phillips et al. (2009), which were trained on the same dataset, but with target-group background points only. Maxent is the conventional approach to SDMs (Phillips et al., 2006), while the Boosting Regression Trees model represents one of the existing tree-based approaches that perform well in modeling species distributions (Valavi et al., 2021). However, both these approaches necessitate the creation of one model per species, resulting in the management of 226 independent models.

All the different methods are evaluated on the presence-absence test set by computing the Area Under the receiver operating characteristic Curve (AUC) for each species and then calculating the mean across all species per region. We opt for AUC due to its widespread use in SDMs and its general high agreement with independent testing data (Konovalik and Nosol, 2021), as well as to avoid binarizing predictions. For completeness, we still compute the correlation and the area under the precision-recall gain curve (Flach and Kull, 2015), as done in Valavi et al. (2022), and include these metrics in the Appendix (see Table C.7 and Table C.8). Finally, we report the average performance over ten different random seeds to add statistical significance.

4. Results and discussion

We first compare our approach to the other baselines, and then perform ablation studies to understand the role played by every weight in the loss function.

4.1. Comparison with other approaches

The full weighted loss $\mathcal{L}_{\text{full-weighted}}$ is compared to the methods and loss functions in Table 1 on the test set of the Elith et al. (2020) dataset. We evaluate the loss functions introduced in Cole et al. (2023) and observe that the use of target-group background points (SLDS loss) consistently yields superior results compared to using random background points (SSDL loss). Additionally, considering both types of pseudo-absences enhances performance in four out of six regions (Full loss). In comparison, our approach, which incorporates species weights and where the pseudo-absence weight λ_2 is set to 0.8 to put more weight on the target-group background points, achieves slightly superior results on average than the full loss. Moreover, fine-tuning the value λ_2 through block cross-validation leads to additional improvements in the CAN region. Notably, these results are achieved without the need for specific λ_1 values. This superior performance is also visible in the correlation (see Table C.7) and the area under the precision-recall gain curve (see Table C.8) metrics. Ultimately, the single-species models exhibit a performance range comparable to that of our approach but still showcase a

Table 1

Comparison of the mean AUC over the species of our approach to other works. Our loss generalizes the losses proposed by (Cole et al., 2023). The best mean AUC for each column, achieved across the different loss functions, is highlighted in bold, whereas the best mean AUC obtained from single-species models is underlined.

	AWT	CAN	NSW	NZ	SA	SWI	avg
Cole et al. (2023) losses							
SSDL loss, i.e., $\lambda_1 = S$ and $\lambda_2 = 0$	0.673	0.542	0.691	0.731	0.800	0.801	0.706
SLDS loss, i.e., $\lambda_1 = S$ and $\lambda_2 = 1$	0.689	0.724	0.709	0.737	0.807	0.845	0.752
Full loss, i.e., $\lambda_1 = S/2$ and $\lambda_2 = 0.5$	0.698	0.673	0.723	0.742	0.814	0.836	0.748
Full weighted loss (ours)							
$\lambda_1 = 1$ and $\lambda_2 = 0.8$, with w_s	0.704	0.696	0.719	0.741	0.815	0.836	0.752
$\lambda_1 = 1$ and fine-tuned λ_2 , with w_s	0.704	0.714	0.719	0.741	0.815	0.838	0.755
Single-species (Phillips et al., 2009)							
Maxent	<u>0.732</u>	0.716	<u>0.741</u>	0.738	<u>0.798</u>	0.837	<u>0.760</u>
BRT	0.700	<u>0.728</u>	0.738	<u>0.740</u>	0.792	<u>0.842</u>	0.757

Table 2

Different values for the species weights w_s , with $\lambda_1 = 1$ and $\lambda_2 = 0.8$. w_s is required to address the strong class imbalance between presences and pseudo-absences. The clamp operation here restricts low frequencies to a minimum threshold of 0.05. The best mean AUC obtained for each column is highlighted in bold.

	AWT	CAN	NSW	NZ	SA	SWI	avg
no species weights	0.653	0.719	0.652	0.702	0.741	0.840	0.718
$w_s = 1/\text{freq}(s)$	0.704	0.696	0.719	0.741	0.815	0.836	0.752
$w_s = 1/\text{clamp}(\text{freq}(s))$	0.700	0.698	0.718	0.738	0.813	0.837	0.751
$w_s = 1/\sqrt{\text{freq}(s)}$	0.697	0.704	0.710	0.737	0.799	0.839	0.748

slight advantage.

4.2. Loss weights

We next conduct ablation studies to understand the role of each weight in our loss function. We adjust each of the three weights in question, one at a time, and keep the other weights at their base values, which are $\lambda_1 = 1$, $\lambda_2 = 0.8$, and $w_s = 1/\text{freq}(s)$.

4.2.1. Species weights

We start by assessing the impact of the species weights w_s , as presented in Table 2. We compare its influence with a scenario in which no species weights are considered, and also test different values of w_s . Our observations reveal a significant improvement when incorporating species weights, with positive effects evident in four out of six regions, resulting in a higher overall average improvement. Particularly, regions NSW and NZ show substantial increases, effectively addressing the challenge posed by the long-tail distribution of presences for these two regions (see Fig. 2a). Alternative approaches involving the square root or clamping low frequency at 0.05 produce comparable, albeit slightly less accurate results on average. Furthermore, the left panel of Fig. 4 illustrates that incorporating species weights is especially advantageous for species with fewer presence records, with diminishing benefits as the number of presence records increases.

Finally, with the inclusion of the species weights, the addition of the presence weight λ_1 to the presences, as in $\mathcal{L}_{\text{full}}$, is no longer necessary, as demonstrated by the results in Table C.5 in the Appendix. Indeed, the improvements resulting from varying the value of λ_1 become only marginal.

4.2.2. Pseudo-absence weight

In Table 3, we explore the impact of varying the ratio of target-group background points to the pseudo-absences using the weighting parameter λ_2 . As previously observed in Table 1, the incorporation of target-group background points consistently leads to improvements, extending the observation of Phillips et al. (2009) to neural networks. When examining the overall performance across datasets, employing either

Table 3

Varying the value of the pseudo-absence weight λ_2 has a notable impact on the mean AUC over the species. The results presented in the last row are obtained using block cross-validation, where λ_2 is selected based on the performance on the validation set. In each result, the species weight w_s is employed, while λ_1 is fixed to 1. The best mean AUC in each column is highlighted in bold.

	AWT	CAN	NSW	NZ	SA	SWI	avg
$\lambda_2=0$	0.651	0.547	0.685	0.731	0.804	0.794	0.702
$\lambda_2=0.2$	0.673	0.600	0.710	0.740	0.812	0.814	0.725
$\lambda_2=0.4$	0.691	0.634	0.720	0.743	0.815	0.824	0.738
$\lambda_2=0.6$	0.702	0.665	0.722	0.743	0.816	0.831	0.746
$\lambda_2=0.8$	0.704	0.696	0.719	0.741	0.815	0.836	0.752
$\lambda_2=1$	0.696	0.714	0.713	0.738	0.811	0.838	0.752
λ_2 found by cv	0.704	0.714	0.719	0.741	0.815	0.838	0.755

only target-group background points ($\lambda_2 = 1$) or a percentage of 80% ($\lambda_2 = 0.8$) tends to produce the best outcomes, although the optimal proportion varies from one region to another. Specifically, for the three regions NSW, NZ, and SA, performance improvements level off for λ_2 values exceeding 0.6. This suggests that the ideal ratio of target-group background points is dependent on the specific data characteristics of each region. In particular, this dependency is linked to the sampling bias, as quantified by Phillips et al. (2009). Regions such as AWT, CAN, and SWI, which exhibit more pronounced sampling bias, benefit more from the inclusion of target-group background points compared to other regions. Consequently, employing a block cross-validation procedure to ascertain the optimal value of λ_2 for each region proves to be beneficial, as visible in the final row of Table 3. This procedure identifies the optimal λ_2 value for three out of the six regions, with the performances of the remaining regions closely approaching the best mean AUC achievable. This leads to a higher average performance across the regions.

In the right panel of Fig. 4, we present the AUCs obtained when using only target-group background points ($\lambda_2 = 1$) versus using only random background points ($\lambda_2 = 0$) for each species. Points above the diagonal black line indicate species where the use of target-group background points leads to a higher AUC than using random background points, and

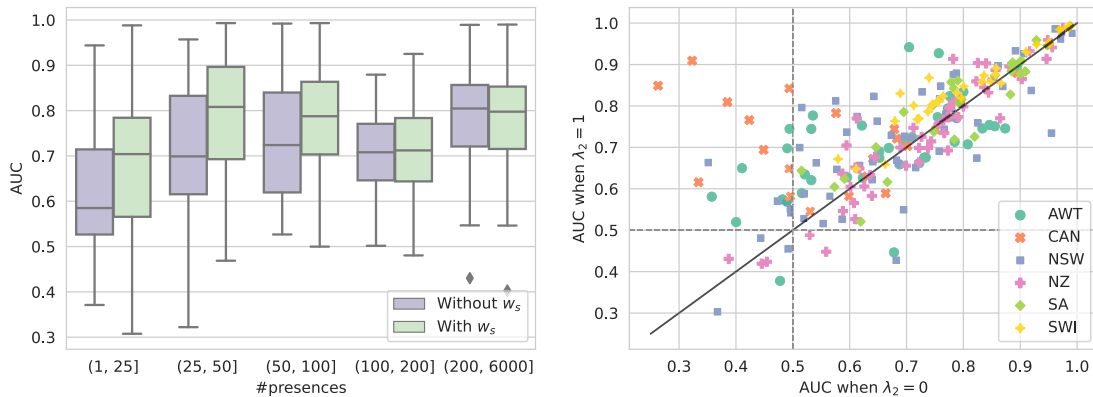


Fig. 4. Left: Impact on the AUC when using the species weight w_s in the loss function, grouped by the number of presences records in the training set. The gain of employing w_s is more pronounced for species with fewer presence records. Right: Impact on the AUC when using random ($\lambda_2 = 0$) or target-group ($\lambda_2 = 1$) background points, with every symbol representing a species. While many species benefit from using only target-group background points, not all do.

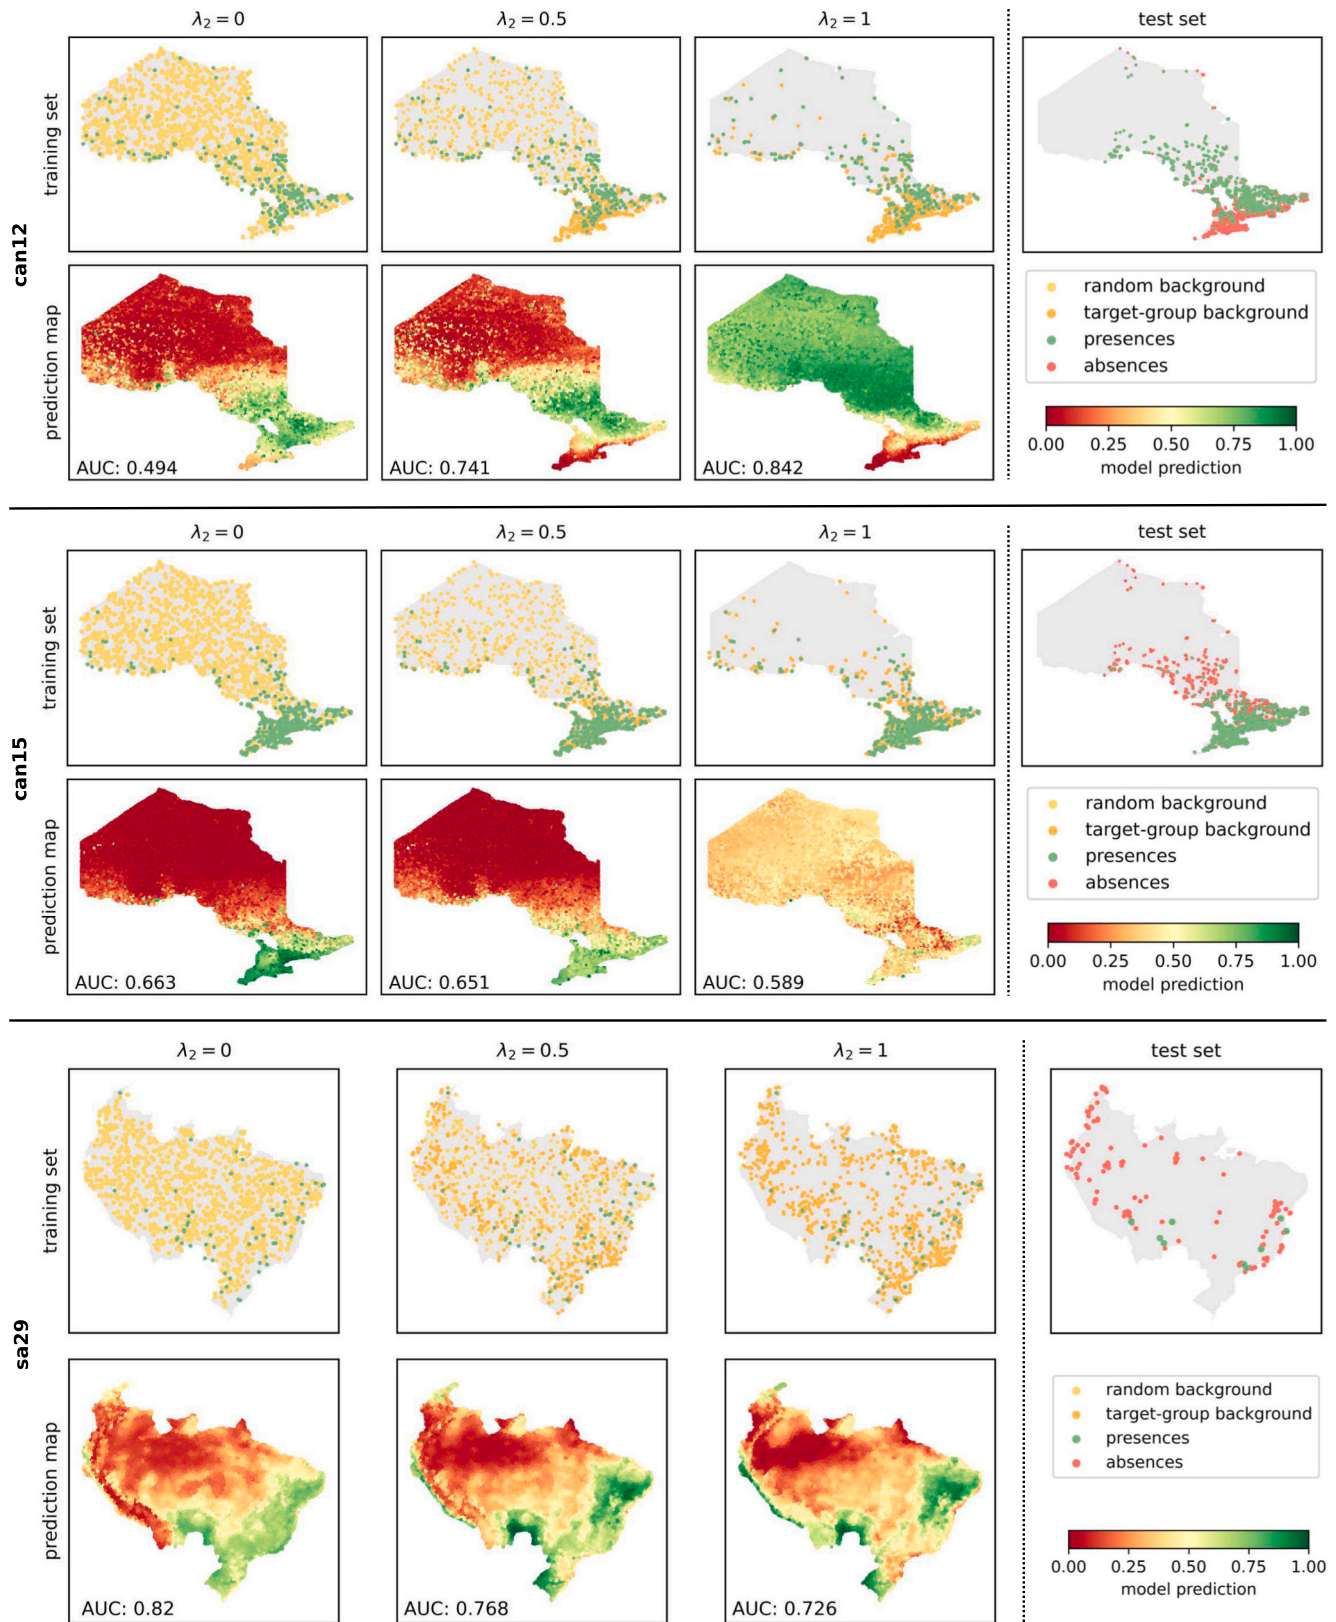


Fig. 5. Focus on three distinct (anonymized) species, labeled as can12, can15, and sa29. Each species is represented by its respective training set, test set, and prediction maps. These visualizations illustrate the role and impact of different values for the pseudo-absence weight λ_2 on the prediction maps generated by the model.

vice versa. Many species, especially in the CAN region, benefit from using target-group background points, with some achieving a boost as high as 0.5 points AUC. However, this trend is not present for all species; some obtain lower performance when using only target-group background points instead of random background points. In particular, within the same geographic region, species exhibit varied responses, highlighting that the optimal value of λ_2 can vary even within a single region, depending on the species.

To gain a concrete understanding of how the choice of λ_2 influences prediction maps, we present maps for three distinct species, each with varying values of λ_2 —ranging from 0 (no target-group background points) to 1 (only target-group background points). The selection of these three species aims at explicitly showcasing the heterogeneous response to λ_2 . Specifically, the more widely distributed species can12 from Canada significantly benefits from including target-group background points, whereas the more specialized can15 species suffers from relying on target-group background points only. This observation aligns with the findings of Ranc et al. (2017) and Botella et al. (2020), which conclude that generalist widely-distributed species gain more from target-group background points compared to specialized species. Notably, the can15 species exhibits a distribution closely mirroring the distribution of presences of all species, heavily biased towards the south (as depicted in the panel showing presences and target-group background points only). Consequently, the model struggles to distinguish between target-group background points and presences, resulting in predicted scores hovering around 0.5 for the entire map. Finally, we note that the differences in prediction maps are not always strong, as illustrated by the sa29 species. In this case, the distinction lies more in the magnitude of the predictions, although the southeast region is not accurately predicted when using only target-group background points.

5. Future work

We present several potential extensions to our approach. Firstly, our loss function offers the flexibility to include as many random background points as desired, as they are appropriately weighted. This flexibility allows one to choose the spatial extent covered by random background points that is optimal (VanDerWal et al., 2009), without being constrained by a specific number. In particular, it enables the sampling of additional random background points in an area that has been considered unsuitable, for example, based on expert knowledge or biological information, as demonstrated by Chapman et al. (2019). Random background points could also be selectively sampled from subareas within the region of interest, such as those distant from known presences (Iturbide et al., 2015; Mateo et al., 2010). Moreover, presence-only data often contains only a single species record per location. In this situation, our approach can be effectively combined with methods that aggregate nearby species presences, such as Kellenberger et al. (2022).

Furthermore, we observed in Fig. 4 (right) and Fig. 5 that different species within the same region may have different optimal values for the hyperparameter λ_2 . Consequently, setting distinct λ_2 values per species could improve model performance. In particular, higher values of λ_2 could be set for known generalist species, especially if there is a strong sampling bias. While the evaluation in this study focused on multi-species neural networks, it is straightforward to adapt our loss function to single-species neural networks by removing the summation over species and retaining only the species of interest.

We chose to use the benchmark dataset from Elith et al. (2020) due to its established reputation in SDMs and its inclusion of PA data to evaluate models trained on PO data. However, the tabular format of this dataset typically reduces the performance of neural networks (Borisov et al., 2022; Grinsztajn et al., 2022). Our results align with this limitation, as non-neural network approaches demonstrated comparable or even slightly superior performance compared to our method. Nevertheless, certain studies suggest that deep learning-based SDMs can surpass traditional machine learning methods by incorporating diverse data

types (Botella et al., 2023; Teng et al., 2023; Zhang et al., 2022). As a result, we intend to expand this research to datasets involving various types of data, such as environmental rasters, satellite images, or time series data, as exemplified by the GeoLifeCLEF 2023 dataset (Botella et al., 2023). An interesting avenue to explore is whether the behavior of the full weighted loss function remains consistent across various neural network architectures, including more complex models such as convolutional neural networks (Krizhevsky et al., 2012) or Transformers (Vaswani et al., 2017).

6. Conclusion

Employing deep learning methods for SDMs represents a promising approach to processing the substantial volume of new ecological data emerging from community science. The prevalence of presence-only observations, a common format for such datasets, poses a challenge due to the absence of information about the non-occurrence of species. To tackle this fundamental issue, pseudo-absences are frequently employed as contrasting samples to the presence observations. In this paper, we introduced a unified and flexible approach to integrating different types and quantities of pseudo-absences when using multi-species neural networks. This is achieved through the introduction of a novel loss function tailored to the specific characteristics of the datasets, effectively addressing challenges such as geographic biases and class imbalance issues. The training of multi-species neural networks with our proposed loss function yielded superior performance compared to previous approaches. Notably, our models performed similarly to Maxent on an independent test set consisting of both presence and absence data from diverse regions while maintaining the convenience of having a single model per region. Our study also sheds light on the intricate relationship between the type of pseudo-absence and the spatial bias of observations, emphasizing the importance of considering this factor in model development. By providing a comprehensive solution to the incorporation of pseudo-absences, our work opens avenues for further refinement and enhancement of multi-species neural network models, ultimately supporting more accurate and reliable predictions in ecological research and conservation efforts.

Funding

This work was supported by the Swiss National Science Foundation, under grant 200021_204057 “Learning unbiased habitat suitability at scale with AI (deepHSM)”.

CRediT authorship contribution statement

Robin Zbinden: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Nina van Tiel:** Writing – review & editing, Software, Data curation, Conceptualization. **Benjamin Kellenberger:** Writing – review & editing, Supervision, Funding acquisition, Conceptualization. **Lloyd Hughes:** Supervision, Conceptualization. **Devis Tuia:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used GPT-3.5 in order to improve the readability and language of this work. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Declaration of competing interest

The authors declare that they have no known competing financial

interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The code to reproduce the experiments is available at <https://github.com/eceo-epfl/SDM-full-weighted-loss>

Appendix A. Additional dataset details

We provide the maps of the geographic distribution of the presence records for all species (Fig. A.6), along with the distribution of the number of presence records per species (Fig. A.7) for all regions. For additional details about the dataset, refer to Elith et al. (2020).

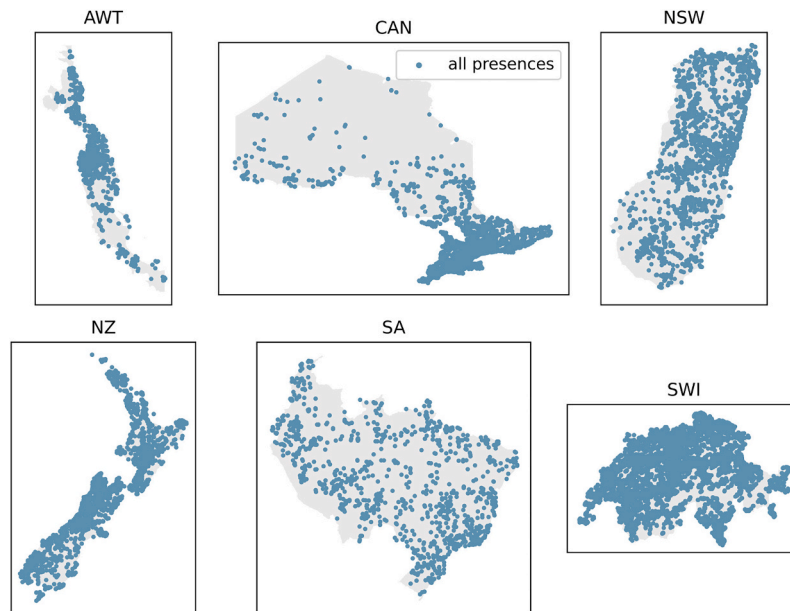


Fig. A.6. The geographic distribution of the presence records of all species within each region of the dataset in use (Elith et al., 2020). The intensity of the sampling bias varies significantly from one region to another.




Fig. A.7. Distribution of the number of presence records per species in each region of the dataset in use (Elith et al., 2020). The distributions typically exhibit a long-tailed pattern, with many species observed only a few times, while a few others are observed more frequently.

Appendix B. Comparison with other ML losses

In this section, we discuss and compare the full weighted loss to other losses commonly used by the machine learning community to address similar setups.

Cross-entropy loss. Previous studies applying deep learning approaches to SDMs have frequently employed the categorical cross-entropy (CE) loss function (Deneu et al., 2021; Zhang et al., 2022). However, conceptually, we argue that the CE loss may not be well-suited for the specific multi-species setup we are dealing with. This is because a single location can accommodate an arbitrary number of species, whereas the CE loss assumes the presence

of only one species per location, thereby introducing score dependencies between species. Furthermore, it may exacerbate issues related to data bias, as it tends to favor species that are more prevalent in the training dataset and penalize rare species. It is also noteworthy that the CE loss implicitly considers target-group background points (and no random background points), as it pits species classes against each other via the softmax operation, preventing an all-one prediction. Nevertheless, given the common usage of the CE loss, we include in Table B.4 the performance results obtained using it. We note that while the CE loss yields good and comparable performance in certain regions, namely CAN, NSW, and NZ, its performance is considerably worse in others and on average. Additionally, weighting the CE by the inverse frequency improves performance in the AWT region, but decreases performance in other regions.

Focal loss. The focal loss function was originally introduced to mitigate class imbalance in object detection (Lin et al., 2017), and works by down-weighting already well-classified samples. Although initially designed for multi-class classification, it can be adapted to multi-label classification. Therefore, we evaluate this loss function in two settings: one combined with the cross-entropy loss and the other combined with our full weighted loss function. However, we observe in Table B.4 that neither of these settings leads to an improvement in performance.

LDAM loss. The Label-Distribution-Aware Margin (LDAM) loss function was introduced to mitigate class imbalance in multi-class classification by regularizing the minority classes more strongly, encouraging larger margins for these classes (Cao et al., 2019). It has been employed by Estopinan et al. (2022) to train SDMs with deep learning. As shown in Table B.4, the performance of this loss function is similar to the CE loss, since it is based on and therefore inherits the same issues described above.

DB loss. Unlike the focal and LDAM loss functions, the Distribution-Balanced (DB) loss function was specifically designed for the multi-label classification setting (Wu et al., 2020). It incorporates considerations of label co-occurrences and aims to mitigate the over-suppression of negative labels to address class imbalance effectively. Based on the hyperparameters values used in Wu et al. (2020), we vary λ while setting $\alpha = 0.1, \beta = 10, \mu = 0.2$, and $\kappa = 0$. However, the performance is lower than that of other baselines on average (Table B.4).

Entmax loss. Zhou et al. (2022) introduced the Entropy-Maximization (Entmax) loss function to tackle the particular setting of single positive multi-label learning encountered with presence-only observations (see Section 2.2). It has previously been tested by Cole et al. (2023), and similarly, we observe lower performance than the baselines, as shown in Table B.4.

Table B.4

Mean AUC performance of other alternative ML loss functions. The performance is averaged over 10 random seeds.

	AWT	CAN	NSW	NZ	SA	SWI	avg
CE	0.663	0.718	0.716	0.741	0.796	0.827	0.744
weighted CE	0.670	0.706	0.714	0.736	0.795	0.816	0.739
Focal CE ($\gamma = 0.5$)	0.663	0.718	0.716	0.741	0.796	0.827	0.744
Focal CE ($\gamma = 1$)	0.663	0.718	0.716	0.741	0.796	0.827	0.744
Focal CE ($\gamma = 2$)	0.662	0.718	0.716	0.741	0.795	0.826	0.743
Focal CE ($\gamma = 5$)	0.660	0.719	0.715	0.741	0.794	0.824	0.742
Focal full weighted ($\gamma = 0.5$)	0.680	0.684	0.682	0.729	0.773	0.836	0.731
Focal full weighted ($\gamma = 1$)	0.686	0.682	0.696	0.734	0.787	0.835	0.737
Focal full weighted ($\gamma = 2$)	0.689	0.678	0.705	0.735	0.800	0.832	0.740
Focal full weighted ($\gamma = 5$)	0.682	0.665	0.699	0.723	0.805	0.825	0.733
LDAM ($C = 0.1$)	0.663	0.718	0.716	0.741	0.796	0.827	0.744
LDAM ($C = 1$)	0.662	0.718	0.716	0.741	0.797	0.828	0.744
LDAM ($C = 10$)	0.648	0.719	0.709	0.738	0.799	0.831	0.741
DB loss ($\lambda = 3$)	0.669	0.704	0.679	0.716	0.769	0.835	0.729
DB loss ($\lambda = 5$)	0.668	0.698	0.679	0.710	0.776	0.829	0.727
Entmax ($\alpha = 0.01$)	0.665	0.687	0.710	0.740	0.796	0.817	0.736
Entmax ($\alpha = 0.05$)	0.680	0.697	0.699	0.719	0.810	0.823	0.738
Entmax ($\alpha = 0.1$)	0.676	0.683	0.688	0.703	0.808	0.819	0.729
Full weighted ($\lambda_2 = 0.5$)	0.698	0.650	0.721	0.744	0.816	0.828	0.743
Full weighted ($\lambda_2 = 0.8$)	0.704	0.696	0.719	0.741	0.815	0.836	0.752
Full weighted ($\lambda_2 = 1$)	0.696	0.714	0.713	0.738	0.811	0.838	0.752

Appendix C. Additional results

Presence weight. In Table C.5, we evaluate the impact of the presence weight λ_1 on performance. Notably, the performance changes from variations in the value of λ_1 are minor, suggesting that the species weights adequately address the class imbalance between presences and pseudo-absences.

Table C.5

Adjusting the value of the presence weight λ_1 has a minimal impact on the mean AUC over the species, indicating that the species weights w_s effectively address class imbalance issues. λ_2 is fixed at 0.8 and w_s is used. The best mean AUC in each column is highlighted in bold.

	AWT	CAN	NSW	NZ	SA	SWI	avg
$\lambda_1=0.1$	0.709	0.701	0.701	0.734	0.784	0.837	0.744
$\lambda_1=0.25$	0.705	0.696	0.714	0.739	0.805	0.836	0.749
$\lambda_1=0.5$	0.702	0.695	0.718	0.740	0.812	0.836	0.751
$\lambda_1=1$	0.704	0.696	0.719	0.741	0.815	0.836	0.752
$\lambda_1=2$	0.704	0.698	0.719	0.743	0.814	0.837	0.752
$\lambda_1=4$	0.703	0.699	0.718	0.743	0.812	0.838	0.752
$\lambda_1=10$	0.695	0.696	0.712	0.737	0.806	0.840	0.748

Block cross-validation. We explore the effectiveness of employing spatial block cross-validation in determining the optimal value for the pseudo-

absence weight λ_2 . We compare this approach with traditional cross-validation methods, including plain cross-validation (randomly splitting presence data) and species-stratified cross-validation, where the proportion of presences is maintained for each species in each subset (Sechidis et al., 2011). During this process, we include the target-group background points associated with the presences in the validation set, along with (optionally) an equivalent number of random background points. The objective is then to identify the λ_2 value that maximizes the average validation AUC across folds. This is achieved through a grid search for λ_2 , exploring its values within the set $\{0, 0.2, 0.4, 0.6, 0.8, 1\}$.

Results in Table C.6 indicate that block cross-validation outperforms other approaches when both random and target-group background points are utilized. This is particularly effective in regions with significant sampling bias, such as AWT, CAN, and SWI. However, superior average performance is achieved by restricting the validation set to include only target-group background points, regardless of the validation method used. Notably, all methods converge on the same selection of λ_2 values.

Table C.6

Mean AUC over the species for different methods employed to construct the validation set used to determine the optimal value of the pseudo-absence weight λ_2 . λ_1 is fixed at 1, and w_s is used. The best mean AUC in each column is underlined when both target-group and random background points are included in the validation set. When exclusively using target-group background points in the validation set, all methods converge on the selection of identical λ_2 values.

	AWT	CAN	NSW	NZ	SA	SWI	avg
target-group + random							
plain	0.673	0.634	<u>0.720</u>	0.740	0.812	0.824	0.734
species-stratified	0.691	0.634	<u>0.720</u>	0.740	<u>0.815</u>	0.824	0.737
block	<u>0.704</u>	<u>0.696</u>	<u>0.720</u>	<u>0.741</u>	0.811	<u>0.831</u>	<u>0.751</u>
target-group							
plain	0.704	0.714	0.719	0.741	0.815	0.838	0.755
species-stratified	0.704	0.714	0.719	0.741	0.815	0.838	0.755
block	0.704	0.714	0.719	0.741	0.815	0.838	0.755

Additional metrics. In Tables C.7 and C.8, we present the mean correlation and the mean area under the precision-recall gain curve across species, respectively. The results are consistent across the different metrics. Remarkably, employing our full weighted loss with the fine-tuned value of λ_2 consistently yields the best average performances over the regions.

Table C.7

Mean Pearson correlation coefficient over the species. The best mean correlation in each column is highlighted in bold, while the second-best mean correlation is underlined.

	AWT	CAN	NSW	NZ	SA	SWI	avg
Cole et al. (2023) losses							
SSDL loss, i.e., $\lambda_1 = S$ and $\lambda_2 = 0$	0.253	0.037	0.156	0.168	0.285	0.208	0.184
SLDS loss, i.e., $\lambda_1 = S$ and $\lambda_2 = 1$	0.279	0.181	0.171	0.171	0.299	0.264	0.228
Full loss, i.e., $\lambda_1 = S/2$ and $\lambda_2 = 0.5$	<u>0.296</u>	0.142	0.191	<u>0.179</u>	<u>0.319</u>	0.267	0.232
Full weighted loss (ours)							
$\lambda_1 = 1$ and $\lambda_2 = 0.8$, no w_s	0.226	<u>0.178</u>	0.119	0.144	0.238	0.310	0.203
$\lambda_1 = 1$ and $\lambda_2 = 0$, with w_s	0.223	0.042	0.153	0.171	0.296	0.222	0.184
$\lambda_1 = 1$ and $\lambda_2 = 0.2$, with w_s	0.254	0.091	0.179	<u>0.179</u>	0.311	0.249	0.210
$\lambda_1 = 1$ and $\lambda_2 = 0.4$, with w_s	0.279	0.117	<u>0.189</u>	0.182	0.318	0.264	0.225
$\lambda_1 = 1$ and $\lambda_2 = 0.6$, with w_s	<u>0.296</u>	0.137	0.191	0.182	0.320	0.274	0.233
$\lambda_1 = 1$ and $\lambda_2 = 0.8$, with w_s	0.300	0.158	0.187	<u>0.179</u>	0.318	0.280	<u>0.237</u>
$\lambda_1 = 1$ and $\lambda_2 = 1$, with w_s	0.290	0.172	0.180	0.176	0.314	<u>0.283</u>	0.236
$\lambda_1 = 1$ and fine-tuned λ_2 , with w_s	0.300	0.172	0.187	<u>0.179</u>	0.318	<u>0.283</u>	0.240

Table C.8

Mean area under the precision-recall gain curve (AUPRG) over the species. The best mean AUPRG in each column is highlighted in bold, while the second-best mean AUPRG is underlined.

	AWT	CAN	NSW	NZ	SA	SWI	avg
Cole et al. (2023) losses							
SSDL loss, i.e., $\lambda_1 = S$ and $\lambda_2 = 0$	0.346	-1.606	-0.175	0.303	0.704	0.773	0.057
SLDS loss, i.e., $\lambda_1 = S$ and $\lambda_2 = 1$	0.411	<u>-0.333</u>	0.006	0.321	0.700	0.854	0.327
Full loss, i.e., $\lambda_1 = S/2$ and $\lambda_2 = 0.5$	0.410	-0.552	0.203	0.404	0.721	0.833	0.337
Full weighted loss (ours)							
$\lambda_1 = 1$ and $\lambda_2 = 0.8$, without w_s	0.339	-0.036	-0.300	-0.451	0.602	0.846	0.167
$\lambda_1 = 1$ and $\lambda_2 = 0$, with w_s	0.291	-1.267	-0.107	0.320	0.690	0.759	0.114
$\lambda_1 = 1$ and $\lambda_2 = 0.2$, with w_s	0.340	-1.057	0.123	0.355	0.710	0.800	0.212
$\lambda_1 = 1$ and $\lambda_2 = 0.4$, with w_s	0.385	-0.782	0.186	0.374	0.722	0.816	0.284
$\lambda_1 = 1$ and $\lambda_2 = 0.6$, with w_s	0.427	-0.629	<u>0.202</u>	0.384	<u>0.728</u>	0.829	0.324
$\lambda_1 = 1$ and $\lambda_2 = 0.8$, with w_s	<u>0.456</u>	-0.430	0.166	<u>0.388</u>	0.730	0.839	<u>0.358</u>

(continued on next page)

Table C.8 (continued)

	AWT	CAN	NSW	NZ	SA	SWI	avg
$\lambda_1 = 1$ and $\lambda_2 = 1$, with w_s	0.461	−0.427	0.076	<u>0.388</u>	0.730	<u>0.847</u>	0.346
$\lambda_1 = 1$ and fine-tuned λ_2 , with w_s	<u>0.456</u>	−0.427	0.166	<u>0.388</u>	0.730	<u>0.847</u>	0.360

Optimizing the full loss. The sole adjustable parameter of the losses proposed by Cole et al. (2023) is the λ value of the full loss function, which they set to 2048. In Table C.9, we show the performance variability of the full loss across different values of λ . Notably, simply using $\lambda = 2048$ results in relatively low performance in our experiments. Additionally, other values of λ also demonstrate inferior performance compared to our full weighted loss on average. Ultimately, setting λ to be proportionate to the number of species S yields satisfactory performance on average (last row of Table C.9).

Table C.9

Mean AUC performance of the $\mathcal{J}_{\text{full}}$ loss with varying values of λ . S is the number of species, therefore the last line corresponds to our automatic tuning of λ . Note that the $\mathcal{J}_{\text{full}}$ loss is equivalent to the $\mathcal{J}_{\text{full-weighted}}$ loss when $\lambda_1 = \frac{\lambda}{2}$ and $\lambda_2 = 0.5$. The performance is averaged over 10 random seeds.

λ	AWT	CAN	NSW	NZ	SA	SWI	avg
1	0.682	0.682	0.689	0.730	0.784	0.837	0.734
2	0.690	0.677	0.707	0.737	0.803	0.836	0.742
4	0.695	0.673	0.716	0.740	0.811	0.836	0.745
8	0.697	0.675	0.721	0.741	0.814	0.836	0.747
16	0.699	0.680	0.724	0.742	0.814	0.837	0.749
32	0.695	0.684	0.725	0.745	0.811	0.837	0.749
64	0.684	0.687	0.724	0.746	0.806	0.836	0.747
128	0.672	0.688	0.723	0.744	0.799	0.834	0.743
256	0.662	0.684	0.718	0.740	0.793	0.831	0.738
512	0.654	0.675	0.711	0.733	0.789	0.825	0.731
1024	0.649	0.660	0.702	0.727	0.786	0.817	0.723
2048	0.647	0.640	0.695	0.721	0.785	0.807	0.716
S	0.698	0.673	0.723	0.742	0.814	0.836	0.748

References

Araújo, M.B., Peterson, A.T., 2012. Uses and misuses of bioclimatic envelope modeling. *Ecology* 93 (7), 1527–1539.

Barbet-Massin, M., Jiguet, F., Albert, C.H., Thuiller, W., 2012. Selecting pseudo-absences for species distribution models: how, where and how many? *Methods Ecol. Evol.* 3 (2), 327–338.

Beck, J., Böller, M., Erhardt, A., Schwanghart, W., 2014. Spatial bias in the gbif database and its effect on modeling species' geographic distributions. *Eco. Inform.* 19, 10–15.

Bekker, J., Davis, J., 2020. Learning from positive and unlabeled data: a survey. *Mach. Learn.* 109, 719–760.

Benkendorf, D.J., Hawkins, C.P., 2020. Effects of sample size and network depth on a deep learning approach to species distribution modeling. *Eco. Inform.* 60, 101137.

Benkendorf, D.J., Schwartz, S.D., Cutler, D.R., Hawkins, C.P., 2023. Correcting for the effects of class imbalance improves the performance of machine-learning based species distribution models. *Ecol. Model.* 483, 110414.

Borisov, V., Leemann, T., Seßler, K., Haug, J., Pawelczyk, M., Kasneci, G., 2022. Deep neural networks and tabular data: a survey. *IEEE Trans. Neural Networks Learn. Syst.* 1–21.

Botella, C., Joly, A., Monestiez, P., Bonnet, P., Munoz, F., 2020. Bias in presence-only niche models related to sampling effort and species niches: lessons for background point selection. *PLoS One* 15 (5), e0232078.

Botella, C., Deneu, B., Gonzalez, D.M., Servajean, M., Larcher, T., Leblanc, C., Estopinan, J., Bonnet, P., Joly, A., 2023. Overview of Geolifeclef 2023: Species Composition Prediction with High Spatial Resolution at Continental Scale Using Remote Sensing. Working Notes of CLEF.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D., 2020. Language models are few-shot learners. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (Eds.), *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc, pp. 1877–1901.

Cao, K., Wei, C., Gaidon, A., Arechiga, N., Ma, T., 2019. Learning imbalanced datasets with label-distribution-aware margin loss. *Adv. Neural Inf. Proces. Syst.* 32.

Chapman, D., Pescott, O.L., Roy, H.E., Tanner, R., 2019. Improving species distribution models for invasive non-native species with biologically informed pseudo-absence selection. *J. Biogeogr.* 46 (5), 1029–1040.

Chen, D., Xue, Y., Chen, S., Fink, D., Gomes, C., 2016. Deep multi-species embedding. *arXiv preprint. arXiv:1609.09353*.

Christin, S., Hervet, É., Lecomte, N., 2019. Applications for deep learning in ecology. *Methods Ecol. Evol.* 10 (10), 1632–1644.

Cole, E., Mac Aodha, O., Lorieul, T., Perona, P., Morris, D., Jojic, N., 2021. Multi-label learning from single positive labels. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 933–942.

Cole, E., Van Horn, G., Lange, C., Shepard, A., Leary, P., Perona, P., Loarie, S., Aodha, O. Mac, 2023. Spatial implicit neural representations for global-scale species mapping. *arXiv preprint. arXiv:2306.02564*.

Davis, C.L., Bai, Y., Chen, D., Robinson, O., Ruiz-Gutierrez, V., Gomes, C.P., Fink, D., 2023. Deep learning with citizen science data enables estimation of species diversity and composition at continental extents. *Ecology* e4175.

Deneu, B., Servajean, M., Bonnet, P., Botella, C., Munoz, F., Joly, A., 2021. Convolutional neural networks improve species distribution modelling by capturing the spatial structure of the environment. *PLoS Comput. Biol.* 17 (4), e1008856.

El-Gabbas, A., Dormann, C.F., 2018. Improved species-occurrence predictions in data-poor regions: using large-scale data and bias correction with down-weighted poisson regression and maxent. *Ecography* 41 (7), 1161–1172.

Elith, J., Leathwick, J.R., 2009. Species distribution models: ecological explanation and prediction across space and time. *Annu. Rev. Ecol. Syst.* 40 (1), 677–697.

Elith, J., Graham, C.H., Anderson, R.P., Dudík, M., Ferrier, S., Guisan, A., Hijmans, R.J., Huettmann, F., Leathwick, J.R., Lehmann, A., et al., 2006. Novel methods improve prediction of species' distributions from occurrence data. *Ecography* 29 (2), 129–151.

Elith, J., Kearney, M., Phillips, S., 2010. The art of modelling range-shifting species. *Methods Ecol. Evol.* 1 (4), 330–342.

Elith, J., Graham, C., Valavi, R., Abegg, M., Bruce, C., Ferrier, S., Ford, A., Guisan, A., Hijmans, R.J., Huettmann, F., et al., 2020. Presence-only and presence-absence data for comparing species distribution modeling methods. *Biodivers. Inform.* 15 (2), 69–80.

Engler, R., Guisan, A., Rechsteiner, L., 2004. An improved approach for predicting the distribution of rare and endangered species from occurrence and pseudo-absence data. *J. Appl. Ecol.* 41 (2), 263–274.

Estopinan, J., Servajean, M., Bonnet, P., Munoz, F., Joly, A., 2022. Deep species distribution modeling from sentinel-2 image time-series: a global scale analysis on the orchid family. *Front. Plant Sci.* 13, 839327.

Flach, P., Kull, M., 2015. Precision-recall-gain curves: Pr analysis done right. *Adv. Neural Inf. Proces. Syst.* 28.

Franklin, J., 2010. *Mapping Species Distributions: Spatial Inference and Prediction*. Cambridge University Press.

Gorishniy, Y., Rubachev, I., Khrulkov, V., Babenko, A., 2021. Revisiting deep learning models for tabular data. *Adv. Neural Inf. Proces. Syst.* 34, 18932–18943.

Grinsztajn, L., Oyallon, E., Varoquaux, G., 2022. Why do tree-based models still outperform deep learning on typical tabular data? In: *In Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Guisan, A., Tingley, R., Baumgartner, J.B., Naujokaitis-Lewis, I., Sutcliffe, P.R., Tulloch, A.I., Regan, T.J., Brotons, L., McDonald-Madden, E., Mantyka-Pringle, C., et al., 2013. Predicting species distributions for conservation decisions. *Ecol. Lett.* 16 (12), 1424–1435.

- Hastie, T., Fithian, W., 2013. Inference from presence-only data; the ongoing controversy. *Ecography* 36 (8), 864–867.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.
- Hertzog, L.R., Besnard, A., Jay-Robert, P., 2014. Field validation shows bias-corrected pseudo-absence selection is the best method for predictive species-distribution modelling. *Divers. Distrib.* 20 (12), 1403–1413.
- Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: *International Conference on Machine Learning*. pmlr, pp. 448–456.
- Iturbide, M., Bedia, J., Herrera, S., del Hierro, O., Pinto, M., Gutiérrez, J.M., 2015. A framework for species distribution modelling with improved pseudo-absence generation. *Ecol. Model.* 312, 166–174.
- Jarnevich, C.S., Talbert, M., Morissette, J., Aldridge, C., Brown, C.S., Kumar, S., Manier, D., Talbert, C., Holcombe, T., 2017. Minimizing effects of methodological decisions on interpretation and prediction in species distribution studies: an example with background selection. *Ecol. Model.* 363, 48–56.
- Johnson, J.M., Khoshgoftaar, T.M., 2019. Survey on deep learning with class imbalance. *J. Big Data* 6 (1), 1–54.
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Zidek, A., Potapenko, A., et al., 2021. Highly accurate protein structure prediction with alphafold. *Nature* 596 (7873), 583–589.
- Kadmon, R., Farber, O., Danin, A., 2004. Effect of roadside bias on the accuracy of predictive maps produced by bioclimatic models. *Ecol. Appl.* 14 (2), 401–413.
- Kellenberger, B., Cole, E., Marcos, D., Tuia, D., 2022. Training techniques for presence-only habitat suitability mapping with deep learning. In: *IGARSS 2022–2022 IEEE International Geoscience and Remote Sensing Symposium*, pp. 5085–5088. <https://doi.org/10.1109/IGARSS46834.2022.9883627>.
- Konowalik, K., Nosol, A., 2021. Evaluation metrics and validation of presence-only species distribution models based on distributional maps with varying coverage. *Sci. Rep.* 11 (1), 1482.
- Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. Imagenet classification with deep convolutional neural networks. *Adv. Neural Inf. Proces. Syst.* 25.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., Dollár, P., 2017. Focal loss for dense object detection. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2980–2988.
- Loshchilov, I., Hutter, F., 2017. Decoupled Weight Decay Regularization. *arXiv preprint. arXiv:1711.05101*.
- Lyu, L., Leugger, F., Hagen, O., Fopp, F., Boschman, L.M., Strijk, J.S., Albouy, C., Karger, D.N., Brun, P., Wang, Z., et al., 2022. An integrated high-resolution mapping shows congruent biodiversity patterns of fagales and pinales. *New Phytol.* 235 (2), 759–772.
- Mac Aodha, O., Cole, E., Perona, P., 2019. Presence-only geographical priors for fine-grained image classification. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9596–9606.
- Mateo, R.G., Croat, T.B., Felicísimo, Á.M., Muñoz, J., 2010. Profile or group discriminative techniques? Generating reliable species distribution models using pseudo-absences and target-group absences from natural history collections. *Divers. Distrib.* 16 (1), 84–94.
- Mesaglio, T., Callaghan, C.T., 2021. An overview of the history, current contributions and future outlook of inaturalist in Australia. *Wildl. Res.* 48 (4), 289–303.
- Pearce, J.L., Boyce, M.S., 2006. Modelling distribution and abundance with presence-only data. *J. Appl. Ecol.* 43 (3), 405–412.
- Phillips, S.J., Dudík, M., 2008. Modeling of species distributions with maxent: new extensions and a comprehensive evaluation. *Ecography* 31 (2), 161–175.
- Phillips, S.J., Anderson, R.P., Schapire, R.E., 2006. Maximum entropy modeling of species geographic distributions. *Ecol. Model.* 190 (3–4), 231–259.
- Phillips, S.J., Dudík, M., Elith, J., Graham, C.H., Lehmann, A., Leathwick, J., Ferrier, S., 2009. Sample selection bias and presence-only distribution models: implications for background and pseudo-absence data. *Ecol. Appl.* 19 (1), 181–197.
- Poggiato, G., Münkemüller, T., Bystrova, D., Arbel, J., Clark, J.S., Thuiller, W., 2021. On the interpretations of joint modeling in community ecology. *Trends Ecol. Evol.* 36 (5), 391–401.
- Ponder, W.F., Carter, G., Flemons, P., Chapman, R., 2001. Evaluation of museum collection data for use in biodiversity assessment. *Conserv. Biol.* 15 (3), 648–657.
- Ranc, N., Santini, L., Rondinini, C., Boitani, L., Poitevin, F., Angerbjörn, A., Maiorano, L., 2017. Performance tradeoffs in target-group bias correction for species distribution models. *Ecography* 40 (9), 1076–1087.
- Roberts, D.R., Bahn, V., Ciuti, S., Boyce, M.S., Elith, J., Guisera-Aroita, G., Hauenstein, S., Lahoz-Monfort, J.J., Schröder, B., Thuiller, W., et al., 2017. Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography* 40 (8), 913–929.
- Sechidis, K., Tsoumakas, G., Vlahavas, I., 2011. On the stratification of multi-label data. In: *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2011, Athens, Greece, September 5–9, 2011, Proceedings, Part III* 22. Springer, pp. 145–158.
- Senay, S.D., Worner, S.P., Ikeda, T., 2013. Novel three-step pseudo-absence selection technique for improved species distribution modelling. *PLoS One* 8 (8), e71218.
- Smith, J.N., Kelly, N., Renner, I.W., 2021. Validation of presence-only models for conservation planning and the application to whales in a multiple-use marine park. *Ecol. Appl.* 31 (1), e02214.
- Sofaer, H.R., Jarnevich, C.S., Pearse, I.S., Smyth, R.L., Auer, S., Cook, G.L., Edwards Jr., T.C., Guala, G.F., Howard, T.G., Morissette, J.T., et al., 2019. Development and delivery of species distribution models to inform decision-making. *BioScience* 69 (7), 544–557.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: a simple way to prevent neural networks from overfitting. *J. Machine Learn. Res.* 15 (1), 1929–1958.
- Stokland, J.N., Halvorsen, R., Støa, B., 2011. Species distribution modelling—effect of design and sample size of pseudo-absence observations. *Ecol. Model.* 222 (11), 1800–1809.
- Teng, M., Elmustafa, A., Akera, B., Bengio, Y., Radi, H., Laroche, H., Rolnick, D., 2023. Satbird: a dataset for bird species distribution modeling using remote sensing and citizen science data. In: *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Tuia, D., Kellenberger, B., Beery, S., Costelloe, B.R., Zuffi, S., Risse, B., Mathis, A., Mathis, M.W., van Langevelde, F., Burghardt, T., et al., 2022. Perspectives in machine learning for wildlife conservation. *Nat. Commun.* 13 (1), 792.
- Valavi, R., Elith, J., Lahoz-Monfort, J.J., Guisera-Aroita, G., 2018. blockcv: an r package for generating spatially or environmentally separated folds for k-fold cross-validation of species distribution models. *Biorxiv* 357798.
- Valavi, R., Elith, J., Lahoz-Monfort, J.J., Guisera-Aroita, G., 2021. Modelling species presence-only data with random forests. *Ecography* 44 (12), 1731–1742.
- Valavi, R., Guisera-Aroita, G., Lahoz-Monfort, J.J., Elith, J., 2022. Predictive performance of presence-only species distribution models: a benchmark study with reproducible code. *Ecol. Monogr.* 92 (1), e01486.
- Valavi, R., Elith, J., Lahoz-Monfort, J.J., Guisera-Aroita, G., 2023. Flexible species distribution modelling methods perform well on spatially separated testing data. *Glob. Ecol. Biogeogr.* 32 (3), 369–383.
- VanDerWal, J., Shoo, L.P., Graham, C., Williams, S.E., 2009. Selecting pseudo-absence data for presence-only distribution modeling: how far should you stray from what you know? *Ecol. Model.* 220 (4), 589–594.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need. *Adv. Neural Inf. Proces. Syst.* 30.
- Wis, M., Guisan, A., 2009. Do pseudo-absence selection strategies influence species distribution models and their predictions? An information-theoretic approach based on simulated data. *BMC Ecol.* 9 (8), 05. <https://doi.org/10.1186/1472-6785-9-8>.
- Wu, T., Huang, Q., Liu, Z., Wang, Y., Lin, D., 2020. Distribution-balanced loss for multi-label classification in long-tailed datasets. In: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV* 16. Springer, pp. 162–178.
- Zbinden, R., van Tiel, N., Kellenberger, B., Hughes, L., Tuia, D., 2023. Exploring the potential of neural networks for species distribution modeling. In: *ICLR Climate Change AI Workshop*.
- Zhang, X., Zhou, Y., Peng, P., Wang, G., 2022. A novel multimodal species distribution model fusing remote sensing images and environmental features. *Sustainability* 14 (21), 14034.
- Zhou, D., Chen, P., Wang, Q., Chen, G., Heng, P.-A., 2022. Acknowledging the unknown for multi-label learning with single positive labels. In: *European Conference on Computer Vision*. Springer, pp. 423–440.