

Estimating environmental suitability

JOHN M. DRAKE^{1,2,3,†} AND ROBERT L. RICHARDS^{1,2}

¹*Odum School of Ecology, University of Georgia, Athens, Georgia 30602 USA*

²*Center for the Ecology of Infectious Diseases, University of Georgia, Athens, Georgia 30602 USA*

³*Department of Zoology, Oxford University, Oxford OX1 3PS UK*

Citation: Drake, J. M., and R. L. Richards. 2018. Estimating environmental suitability. *Ecosphere* 9(9):e02373. 10.1002/ecs2.2373

Abstract. Methods for modeling species' distributions in nature are typically evaluated empirically with respect to data from observations of species occurrence and, occasionally, absence at surveyed locations. Such models are relatively "theory-free." In contrast, theories for explaining species' distributions draw on concepts like *fitness*, *niche*, and *environmental suitability*. This paper proposes that environmental suitability be defined as the conditional probability of occurrence of a species given the state of the environment at a location. Any quantity that is proportional to this probability is a measure of relative suitability and the support of this probability is the niche. This formulation suggests new methods for presence-background modeling of species distributions that unify statistical methodology with the conceptual framework of niche theory. One method, the plug-and-play approach, is introduced for the first time. Variations on the plug-and-play approach were studied with respect to their numerical performance on 106 species from an exhaustively sampled presence-absence survey of vegetation in the Canton of Vaud, Switzerland. Additionally, we looked at the robustness of these methods to the presence of irrelevant information and sample size. Although irrelevant variables eroded the predictive performance of all methods, these methods were found to be both numerically and statistically robust.

Key words: ecological modeling; environmental suitability; habitat; niche; plant ecology; species distribution.

Received 20 July 2017; revised 8 March 2018; accepted 9 March 2018; final version received 15 July 2018. Corresponding Editor: Debra P. C. Peters.

Copyright: © 2018 The Authors. This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

† **E-mail:** jdrake@uga.edu

INTRODUCTION

How the occurrence of a species in nature depends on the state of its environment is one of the most fundamental problems in ecology. Conceptual frameworks for answering this question introduce such ideas as *fitness* (Fretwell and Lucas 1969, Hirzel and Le Lay 2008, Peterson et al. 2011), *niche* (Grinnell 1917, Peterson et al. 2011), and *environmental suitability* (Engler et al. 2004, Hirzel and Le Lay 2008, VanDerWal et al. 2009, Franklin 2010). Problems with operationalizing these concepts include vagueness about their meaning (Peters 1976, Hurlbert 1981, Orr 2009), failure of real systems to meet assumptions such as *distributional equilibrium*

(Elith et al. 2010) and lack of *source-sink dynamics* (Pulliam 2000), and data problems such as the commonness of *presence-only data* (Brotons et al. 2004, Pearce and Boyce 2006, Ward et al. 2009).

This paper seeks to define some of these concepts in a theoretically unifying and computationally operational way. The proposed definitions lend themselves to a new method, which we refer to as the *plug-and-play approach* to modeling environmental suitability. The plug-and-play approach is very flexible. For instance, it allows that data might come from different places at different times. Further, plug-and-play methods for estimating environmental suitability also yield an approach to ecological niche modeling. We study several instances and show that the

performance of the plug-and-play approach can be comparable to MaxEnt, a leading method for species distribution modeling (Phillips et al. 2004, Elith and Leathwick 2009, Elith et al. 2011), and superior to a recently proposed density ratio estimator (Kanamori et al. 2009, Sugiyama et al. 2013). Finally, when used in a particular way (i.e., using the regularized Gaussian estimator as its base learner), the plug-and-play method can be used for *variable identification* and fitting of *robust models* even in the presence of few records and irrelevant variables.

Environmental suitability and the ecological niche

The state of the environment at a location i may be represented by a vector of measurements $z_i = [a_i, b_i, c_i, \dots]$ where a is rainfall, b is temperature, c is vegetation, etc. We assume that z_i is constant through time. The *environmental distribution* is the joint density of environments in nature, denoted $f(z)$ (notation follows Elith et al. 2011). The *environmental distribution of species s* or *occurrence distribution* is the joint density of environments in which s is found, denoted $f_1(z)$, distinct from its *range* and *niche* (Drake 2015). The support of f_1 , that is, the values of z at which $f_1 > 0$, is the realized niche. The boundary of the support of f_1 is denoted by h_F . Sometimes one is interested in the conditional probability that a species occurs at a location given the environment there, $P(y = 1|z)$ (Keating and Cherry 2008, Ward et al. 2009, Franklin 2010, Elith et al. 2011, Royle et al. 2012, Hastie and Fithian 2013). We call this conditional probability the *suitability* of environment z , $S(z)$, for species s . By Bayes' rule,

$$S(z) = P(y = 1|z) = \frac{f_1(z)P(y = 1)}{f(z)} \quad (1)$$

where $P(y = 1)$ is the species *prevalence*. Since prevalence is a proportionality coefficient, we will sometimes wish to ignore it, in which case we have *relative suitability*

$$S_R(z) = \frac{f_1(z)}{f(z)} \propto S(z). \quad (2)$$

Typically, prevalence will be unknown, although (because it is a single quantity) it might possibly be estimated through independent surveys with less effort than is required to estimate the component densities $f_1(z)$ or $f(z)$ (Phillips et al.

2009). Assuming non-extinction, $P(y = 1) > 0$, $S_R(z) > 0$ if and only if $S(z) > 0$. That is, the support of S and S_R is identical. The *fundamental niche*, N , is defined as *the set of all environments in which the species can persist in the absence of continuous immigration from other populations*. More completely, the fundamental niche, N , is the set of all environments z such that there exists a population size n at which the probability of persisting at a location with environment z and in the absence of immigration over a large time horizon $T \gg 0$ exceeds some threshold θ , possibly close to one, in which case there is a non-zero probability that species s will be found in z ($S(z) > 0$), further implying $S_R(z) > 0$. Typically, this persistence condition will obtain when there is population size n at which average absolute individual fitness exceeds one. This is a cumbersome definition, but it points to a way that such notorious problems as source-sink dynamics (Pulliam 2000), Allee effects (Holt 2009), stochastic extinction (Hanski 1989), and niche conservatism (Wiens et al. 2010) may be conceptually incorporated and quantitatively addressed, rather than having them swept away by unrealistic model assumptions. Formally, $N = \{z: S(z) > 0\}$. The boundary of this set is designated h_N .

By providing definitions for *environmental suitability* and *niche* and relating these to observable quantities in nature (e.g., the frequency of occurrence of a species), this framework allows us to formulate and answer some fundamental ecological questions, including:

1. Is location i more or less suitable for species s than location j (for any choice of i and j)?
2. What is the dimension of the niche of species s ?
3. What are the environmental variables that comprise the niche of species s or influence its distribution?
4. How is the potential distribution of species s influenced by a given variable x ?
5. What are the boundaries of the niche of species s ? (What are its environmental *tolerances*?)
6. Is the environmental distribution of species s set by its niche (its *tolerances*) or by the set of the environments realized in nature?
7. How will the environmental distribution of species s change when the distribution of environments in nature changes?

8. How will the spatial distribution of species s change as the distribution of environments in nature changes?

The framework can also be used to empirically address questions in community ecology:

1. What is the niche similarity of species s_1 and s_2 ?
2. Do communities $\mathcal{S}_1 = \{s_1, s_2, s_3, \dots, s_n\}$ and $\mathcal{S}_2 = \{s_{n+1}, s_{n+2}, s_{n+3}, \dots, s_{n+m}\}$ exhibit more within-community niche variation or between-community niche variation?
3. What is the *saturation* or fraction of species in the species pool that could persist in z ?

An estimator for environmental suitability

In this section, we propose that this conceptual framework lends itself to modeling, that is building numerical or computer-learned models of environmental suitability or species’ niches from records of species occurrence in nature. Here is the core of the idea. First, many applications do not require an estimate of absolute environmental suitability, but (for instance) only a rank ordering. In such cases, relative suitability is adequate. Particularly, we shall argue below that *niche identification*, which is what we call the process of building a model of h_N , only requires information about relative suitability. The plug-and-play approach to modeling relative environmental suitability proposes that the ratio of two estimates, $f_1(z)$ and $f(z)$, be used as an estimator of S_R :

$$\widehat{S}_R = \widehat{f_1(z)} / \widehat{f(z)}. \tag{3}$$

Because $f_1(z)$ and $f(z)$ are just probability densities, they can be estimated using any of a number of techniques for probability density estimation, such as kernel density estimation or Parzen’s window estimator (Wand and Jones 1995). Alternatively, one might substitute a parametric expression, which is basically what is done by MaxEnt, a maximum entropy algorithm commonly used for species distribution modeling (Phillips et al. 2004). Specifically, the MaxEnt algorithm for species distribution modeling stipulates that $S_R = e^{\alpha + \beta h(z)}$, where α is a normalizing constant chosen so that $f_1(z)$ integrates to one, β is a vector of fit coefficients, and $h(z)$ are

transformations of the covariates referred to as *features*. (This function h is different from h_F and h_N as defined above, but used here for consistency with Elith et al. 2011). Thus, MaxEnt is a special case of environmental suitability modeling, but does not adopt the plug-and-play approach, which allows for the substitution of alternatives for f_1 and f , including nonparametric options, as may be suggested by theory (either statistical theory or biological theory), the objectives of a study, or the properties of a data set (such as sample size).

This approach can be extended to the problem of ecological niche identification. For a given estimate of S_R , the estimated boundary of the realized niche is just $\widehat{h}_R = \{z : S_R(z) = \epsilon\}$, where $\epsilon > 0$ is a small threshold parameter. However, Drake (2015) argued that if a species is rarely found in unsuitable habitats and the environmental distribution is “broad” with respect to the species’ niche (that is, that the range of f contains the extreme environments in N), then $h_F \approx h_N$ (Drake 2015). Substituting \widehat{h}_N for h_F , we have an estimator for the boundary of the fundamental niche, that is, $\widehat{h}_N = \{z : S_R(z) = \epsilon\}$.

An alternative to the estimator in Eq. 3 is to estimate the ratio directly, that is,

$$\widehat{S}_R = \left(\frac{\widehat{f_1(z)}}{\widehat{f(z)}} \right). \tag{4}$$

This problem has been addressed generically by Sugiyama et al. (2013) and specifically by Kanamori et al. (2009) in the context of covariate shift adaptation (adapting statistical analyses to changing distributions of independent variables) and Kanamori et al. (2009) and Hido et al. (2011) in the context of outlier detection. The technical similarity between outlier detection and ecological niche modeling has been noted before and used to motivate presence-only models for ecological niche modeling (Drake et al. 2006, Drake and Bossenbroek 2009, Drake 2015). Here, the analogy is used to motivate a presence-background approach. For this purpose, Kanamori et al. (2009) introduce unconstrained least-squares importance fitting (uLSIF), which they show to be superior to a variety of alternatives when applied to a simulated classification task. We therefore used uLSIF as a direct density ratio estimator for comparison with the plug-and-play estimator.

METHODS

Plug-and-play with simple component densities

Implementing the plug-and-play approach requires estimating the component densities f and f_1 . Candidate estimators for the component density functions include parametric multivariate density estimators (Kotz et al. 2000), robust versions of these (which may be important when data are not normally distributed or contain extreme examples; Huber and Ronchetti 2009), sparse estimators that reduce variance through “regularization” (which may be important when fitting to small samples from high-dimensional spaces; Schäfer et al. 2015), and nonparametric methods (Wand and Jones 1995).

To compare among these classes, we studied the performance of relative suitability models comprised of three kinds of Gaussian density estimates: (1) ordinary multivariate Gaussian densities, (2) robustly estimated Gaussian densities, (3) densities estimated using a shrinkage estimator proposed by Schäfer and Strimmer (Schäfer et al. 2015) in which the pairwise correlation coefficients are scaled by a shrinkage intensity parameter ($\lambda = \min(1, \max(0, 1 - \lambda^*))$), where λ^* is the analytic optimal regularization parameter of Ledoit and Wolf (2003). For comparison, we also studied the nonparametric kernel density estimator of Li and Racine (2003), which would be expected to be superior in cases where one or both of the component densities f and f_1 are strongly non-normal, for instance if they are multi-modal or highly skewed.

Plug-and-play with ensemble component densities

The plug-and-play approach is not limited to such simple component densities as elaborated in the previous section, but can be applied to basically any approach to probability density estimation that can be numerically evaluated. Recently, *ensemble methods* have been shown in many areas of statistical application to provide robust probabilistic models that exhibit both low variance and low bias. For instance, random forests are frequently used in classification and regression problems (Breiman 2001). For applications requiring species distribution modeling, the gradient boosting machine is popular (Elith et al. 2008, Elith and Leathwick 2009). To be effective, ensemble models must exhibit improvement

compared to a single model. Typically this is achieved by “voting” the predictions of a large number of minimally biased (underfit) models (Drake 2014). Ensembles of highly tuned models may even erode performance (Mainali et al. 2015). Thus, it is important when constructing ensemble learners to optimize the entire ensemble, not the base models.

Here, we explore how *bootstrap aggregation* or *bagging* can be used to improve the estimate of f or f_1 by reducing its variance and thereby improve the performance of the plug-and-play estimator. Bagging is performed by constructing bootstrap samples (with replacement) from a given data set, fitting a model to each, and then reporting the average prediction of the ensemble of models (Breiman 1996). Bagging has been found to be a very robust approach to ensemble modeling. Typically, it will be the case that there are many more records of environmental background than species occurrence, so the estimate of f will be more precise than that of f_1 (Phillips et al. 2004). Motivated by this observation, we first propose only to bootstrap the estimation of the occurrence distribution in the numerator (f_1), yielding a method we call *NumBag*. In the implementation studied here, a kernel density estimate was obtained for each of $v = 100$ bootstrap samples from the occurrence distribution and averaged before dividing by the kernel density estimate of the background points. Alternatively, we use bagged estimators for both f_1 and f , a procedure we call *DoubleBag*. Since the estimated density \hat{f} may depend on a very large number of background points, *DoubleBag* is expected to be much more computationally costly than *NumBag*. Therefore, it is of interest to determine the performance of *NumBag* relative to *DoubleBag* as well as the performance of both in absolute terms. *NumBag* and *DoubleBag* are special cases of the plug-and-play idea that offer potentially greater performance than the simple estimators described above, although at a cost of increased computation time.

Density ratio estimation

In contrast to plug-and-play, the uLSIF algorithm estimates the relative suitability quotient directly. Direct estimation might be desirable if, for instance, estimation of the component densities is more difficult than estimation of the ratio or

if estimation errors in the component densities are compounded when the ratio is taken. Sugiyama et al. (2013) discuss density ratio estimation in general and the uLSIF estimator in particular. One advantage of uLSIF is that it is possible to analytically compute the leave-one-out cross-validation score for a given set of model parameters, greatly decreasing the time needed for model tuning in comparison with numerical cross-validation. MaxEnt (Phillips et al. 2004) may also be viewed as a density ratio estimator and is probably the most popular approach to species distribution modeling in general.

Presence-only methods

Finally, these methods were compared to two recently introduced presence-only modeling methods, LOBAG-OC (Drake 2014) and range bagging (Drake 2015). In many cases, presence-only algorithms perform as well and in some cases perform better than related presence-absence algorithms (Maher et al. 2014). Briefly, LOBAG-OC applies bootstrap aggregation to low bias (weakly regularized) one-class support vector machines and votes the result. LOBAG-OC, therefore, follows Vapnik's principle (Vapnik 1998) in that it seeks the set of covariate values for which the probability density is non-zero (the support), rather than attempting to estimate the full probability density (Drake et al. 2006). Range bagging similarly votes a number of base learners trained on bootstrap samples of presence points. In this case, the base learner is the convex hull of a reduced number of environmental covariates randomly selected for each bootstrap sample. The current analysis provides, therefore, not only a comparison of the plug-and-play method with other presence-background estimators, but is also the first comparison of these two presence-only methods themselves.

Computation

Statistical fitting is straightforward. For this illustration, all fitting was performed in the statistical programming environment R (R Core Team 2017). In the case of the ordinary Gaussian density, the covariance was estimated from the unbiased sample covariance using the R function `cov` (R Core Team 2017). The robust version was estimated using the minimum covariance

determinant (Rousseeuw and Van Driessen 1999) (function `covRob` in the `robust` package, Wang et al. 2014). The regularized version was estimated using a regularized density estimator (function `cov.shrink` in the `corpcor` package, Schäfer et al. 2015). Since this estimator truncates the fit distribution to the first two moments, we refer to this as a "regularized Gaussian" estimator. The kernel density estimates (KDE) were using function `npudens` in package `np` (Hayfield and Racine 2008). Bandwidth of this estimator was selected automatically using the analytic rule of thumb of Li and Racine (2003). Preliminary experiments suggested that computationally intensive bandwidth selection procedures like cross-validation could improve performance, as measured by the area under the receiver-operator curve (AUC), by an average of only about 0.01 at a cost of ~6000-fold increase in computing time. MaxEnt models were fit using the function `maxent` at the default settings in the R package `dismo` (Hijmans 2012). This function makes use of the java program of Phillips et al. (2004). Density ratio estimates obtained via uLSIF were computed using the R code of Kanamori et al. (2009). Code and data to reproduce analysis reported in this paper are available in Data S1 and on the Dryad Data Repository at: <https://doi.org/doi:10.5061/dryad.g5p7d1c>.

Case study: Ecological niches of sub-alpine vegetation

Drake et al. (2006) studied the performance of support vector machines as niche models using data on 106 plant species in 550 8×8 m plots in the Swiss Alps between 400 and 3200 m in elevation. Observed species prevalence (i.e., fraction of sampling plots in which a species was found), ranged from 1.1% to 35.2%. As the above-ground species in these sampling plots were exhaustively enumerated, these data provide a rare opportunity with known absences to benchmark the performance of alternative modeling approaches. Additionally, these data have been studied by Maher et al. (2014), who established that presence-only methods could perform comparably to presence-absence methods, and Drake (2014) in an evaluation of the LOBAG-OC algorithm. Together, these papers provide a benchmark against which to compare the performance of the

plug-and-play approach to modeling environmental suitability. Locations were randomly assigned to either training (80%) or testing (20%) subsets. All predictor data were rescaled by subtracting the mean and dividing by standard deviation of samples in the training set. A model was fit for each species using the plug-and-play estimator where \hat{f}_1 was obtained from the component density fit only to occurrence records in the training set and \hat{f} was fit to all records in the training set. Model performance was evaluated by predicting the relative suitability for each record in the test data set and computing AUC, the area under the receiver operator characteristic, which is a measure of a model's discriminative ability (Phillips and Elith 2010). Models were tuned using 10-fold cross-validation on the training set with mean AUC calculated for each model on each species across the 10-folds. After fitting, tuned models were evaluated on the test data set. To investigate calibration and the propensity of these models to be overfit by cross-validation, mean AUCs from cross-validation were compared to the test AUC with Spearman rank-order correlation.

Effect of irrelevant variables

As the volume of automatically recorded and remotely sensed environmental data accelerates, an increasing problem for modeling environmental suitability or ecological niche identification is the determination of relevant variables and optimizing performance in the presence of irrelevant information. Intuitively, one expects model performance to decline with the number of irrelevant variables as the learning algorithm has to sift a smaller and smaller fraction of true correlates from among the many possibilities. Additionally, the number of parameters sharing the available degrees of freedom increases with the number of irrelevant variables, diminishing the amount of information available to estimate each and eventually resulting in ill-posedness for those methods lacking a regularization scheme to solve this problem. On the other hand, it has recently been shown that, amazingly, in some cases pattern recognition is actually improved by the presence of irrelevant variables (Helmbold and Long 2012).

We studied the performance of plug-and-play, LOBAG, range bagging, and MaxEnt methods in the presence of irrelevant variables by simulating

1, 2, 4, 8, 16, or 32 normally distributed random variates with mean zero and unit variance and combining these with the 10 genuine variables prior to model fitting. As before, model performance was evaluated by calculating AUC on the withheld test data. Results were summarized by inspecting the erosion of performance (decline in AUC) with respect to the number of irrelevant variables, a measure of statistical robustness, and by tabulating the number of models that converged, a measure of numerical robustness.

Learning rate analysis

Finally, we studied the effect of sample size on model performance. An important property of any statistical model is the way in which its performance changes with the number of instances available to learn from, the *learning rate*. This problem is particularly acute in species distribution modeling, where a species may be known from only a small number of unique localities. To investigate the learning rate of the algorithms studied here, we identified the eight plant species with the greatest number of occurrence records in the training data set (*Geranium sylvaticum*, *Anthyllis vulneraria*, *Polygonum viviparum*, *Achillea millefolium*, *Lathyrus pratensis*, *Astrantia major*, *Plantago media*, and *Pimpinella major*). Training points were randomly assigned to 10 cross-validation folds. Within folds, each method was trained on randomly selected subsets of the data ranging in size from $n = 2$ to the total number of available presence points in the smallest set of training folds. Each of these models (m_2, m_3, \dots, m_n) was then applied to the points in the test fold, and an AUC score for each was calculated as above. These results were visualized by plotting the mean test AUC for each learning method as a function of the training set size.

RESULTS

Using the classical Gaussian for both component distributions, the plug-and-play method performed moderately well when it could be fit (mean AUC: 0.761), which was approximately 85% of cases (Fig. 1). Models using the robust Gaussian for both components typically yielded a poorer fit (mean AUC: 0.717) and could be fit in a similar fraction of cases (83%; Fig. 1), whereas models using the regularized Gaussian




Fig. 1. Performance of six plug-and-play species distribution models compared with two presence-only species distribution models (LOBAG-OC and RangeBag), MaxEnt, and a density ratio estimator (uLSIF). The plug-and-play methods were classical multivariate Gaussian densities (Gauss), robust Gaussian (Rob. Gauss), regularized Gaussian (Reg. Gauss), and kernel density estimator (KDE) applied to both f and f_1 ; numerator bagged plug-and-play (NumBag); and double bagged plug-and-play (DoubleBag). Range in AUC was smallest for the regularized Gaussian, KDE, NumBag, and DoubleBag methods (top panel). Methods with the highest average AUC were the regularized Gaussian, KDE, NumBag, and DoubleBag plug-and-play approaches and MaxEnt (bottom panel). The classical Gaussian and robust Gaussian plug-and-play methods and presence-only methods LOBAG-OC and RangeBag performed considerably more poorly; the density ratio estimate obtained using uLSIF was intermediate. Error bars are mean \pm standard error. Both the classical Gaussian and robust Gaussian were unable to fit all models due to numerical instability or ill-posedness (89 and 87 out of 106 fit, respectively).

routinely performed very well (mean AUC: 0.829) and could always be fit. The model using KDE for the components was similarly robust in the sense that a reasonable model could always be obtained. Initially, we thought this improvement might come at considerable cost in terms of computational complexity to allow for empirically optimizing bandwidth parameters, but ultimately we found the model could be estimated using the analytic rule of thumb of Li and Racine (2003) with very little loss of performance (mean AUC: 0.840). The ensemble approaches NumBag and DoubleBag were among the best performing methods we studied, with high average AUC

values (0.836 and 0.837, respectively) and the large majority of models able to be fit (98% and 97%). LOBAG-OC (mean AUC: 0.756, ~93% fit), range bagging (mean AUC: 0.777), and uLSIF (mean AUC: 0.789) all performed less well. Thus, both the plug-and-play method (when using either KDE or regularized Gaussian component distributions) and the ensemble approaches NumBag and DoubleBag were found to perform comparably to MaxEnt (mean AUC: 0.841; Fig. 1) and better than two other methods we have recently introduced.

To assess the tunability of each method and vulnerability to overfitting, we compared AUC

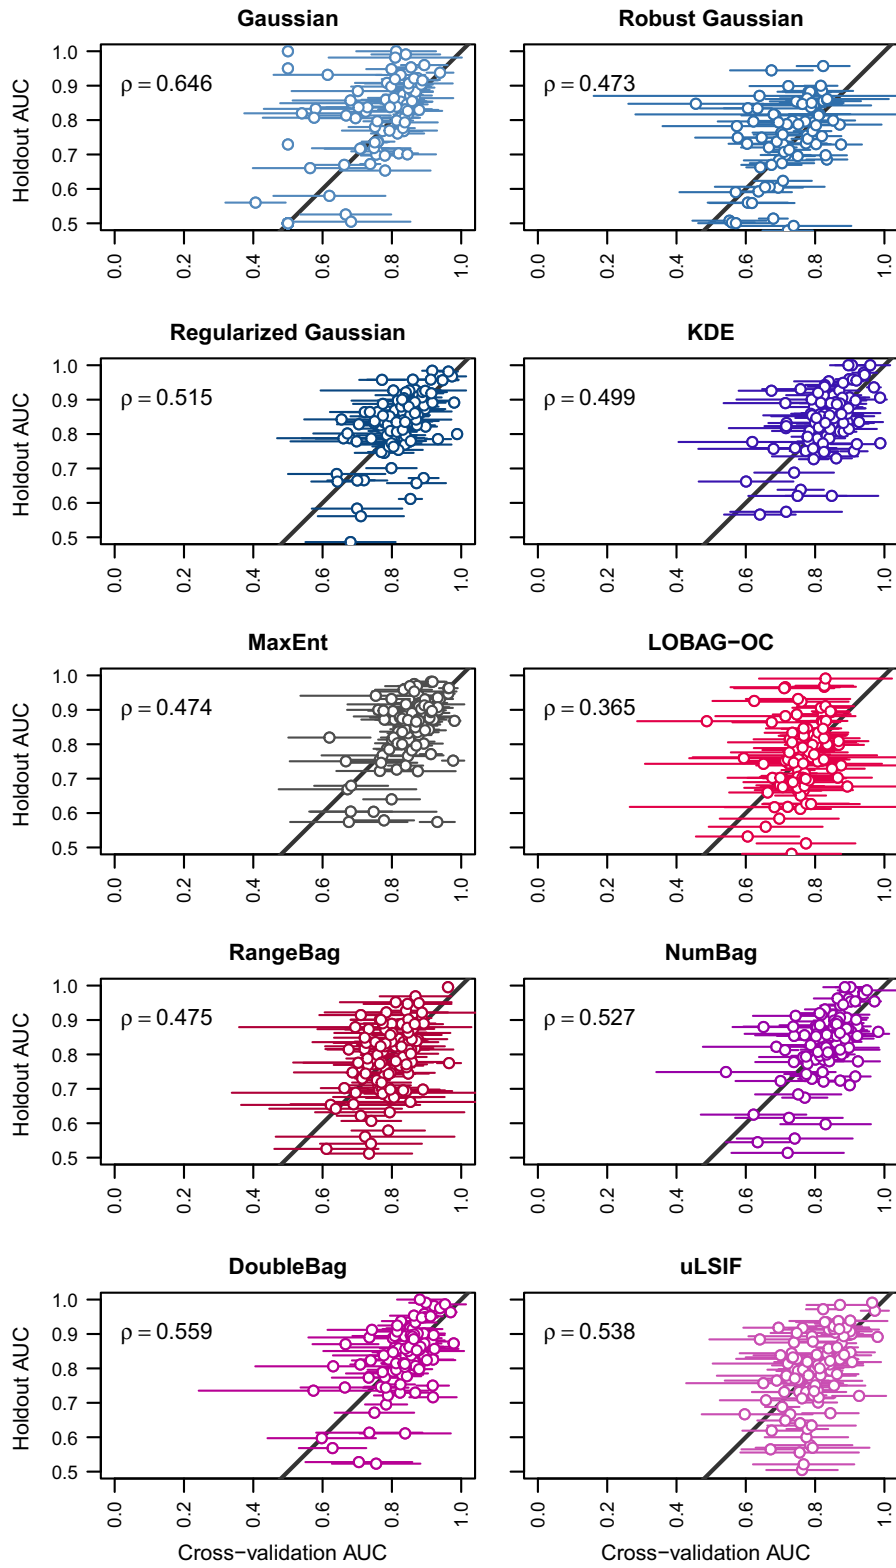


Fig. 2. Mean performance (AUC) across 10 cross-validation folds is compared to performance (AUC) of each

(Fig. 2. *Continued*)

model on the testing subset. A one-to-one line is plotted on each graph, and 95% confidence intervals for mean cross-validation AUCs are shown. Points falling on the one-to-one line are models where the estimated performance in cross-validation was equal to performance on the test set. Points below the line are overfit (higher performance in cross-validation than with withheld data) while points above the line are underfit (higher performance in withheld data than in cross-validation). AUC values were significantly and positively correlated for all models (Spearman's ρ shown).

in the cross-validation folds with the AUC calculated on the test data. An overfit model will have higher AUC on training data than on test data. Additionally, the correlation between the mean AUC in cross-validation folds and test AUC indicates the extent to which the observed AUC in training is predictive of the AUC that can be expected with respect to unseen examples. In this study, all cross-validation AUC values were significantly ($P < 0.01$) and positively correlated with test values (Fig. 2). The strongest correlations were shown by the plug-and-play method with classical Gaussian components ($\rho = 0.646$), DoubleBag ($\rho = 0.559$), and uLSIF ($\rho = 0.538$) algorithms, suggesting that measured performance will be most indicative of future performance in novel analyses performed with these three models.

Performance of the studied methods varied greatly in the presence of irrelevant variables (Fig. 3). The plug-and-play approach with KDE and regularized Gaussian component densities, NumBag, DoubleBag, and MaxEnt models clustered together as the best performing overall (Fig. 3). Specifically, although performance declined as the number of irrelevant variables increased, it declined very slowly so that the reduction in AUC was only about 6% in the presence of 32 irrelevant variables, at which point random variables outnumbered real variables by more than three to one. Thus, in the sense that introducing irrelevant information had little effect on the final output, these three models were all statistically very robust. In contrast, the performance of the plug-and-play approach with classical Gaussian and robust Gaussian component densities, range bagging, LOBAG-OC, and density ratio estimation with uLSIF declined much more rapidly (Fig. 3). Additionally, although the plug-and-play approach with KDE and regularized Gaussian components and MaxEnt models converged in all cases, the plug-and-play approach

with classical Gaussian and robust Gaussian components increasingly failed to converge as the dimension of the environmental data increased.

Analyses of learning rate showed similar performance for some methods across the eight most abundant species in our data and variable performance for others (Fig. 4). Additionally, the niches of some species (e.g., *Plantago media*) seemed inherently more learnable. The group of high-performing methods (i.e., DoubleBag, NumBag, KDE, regularized Gaussian, and MaxEnt) tended to perform relatively poorly with fewer than around 15 training points, but improved quickly approaching their maximum performance at around 20 or 30 observations. At the other end of the spectrum (i.e., with around 100 observations in *Achillea millefolium* and *Pl. media*), MaxEnt tended to perform slightly better than DoubleBag and other high-performing methods, although this was not always true (compare *Pimpinella major*). Interestingly, the performance of the density ratio estimated with uLSIF varied substantially across species. In some cases (e.g., *Pi. major* and *Pl. media*) uLSIF performed consistently less well than the group of highest fit models, while in other cases (e.g., *Geranium sylvaticum* and *Lathyrus pratensis*) uLSIF overtook the performance of DoubleBag with increasing numbers of training points. The robust Gaussian method was a clear outlier overall, routinely exhibiting lower performance at all sample sizes.

DISCUSSION

This study proposes numerical definitions for the ecological concepts of *environmental suitability*, *relative suitability*, and *niche* that we believe are consistent with common usage. A virtue of these definitions is the potential to unify ecological niche theory with dynamical theory, as alluded to by Peterson et al. (2011), particularly

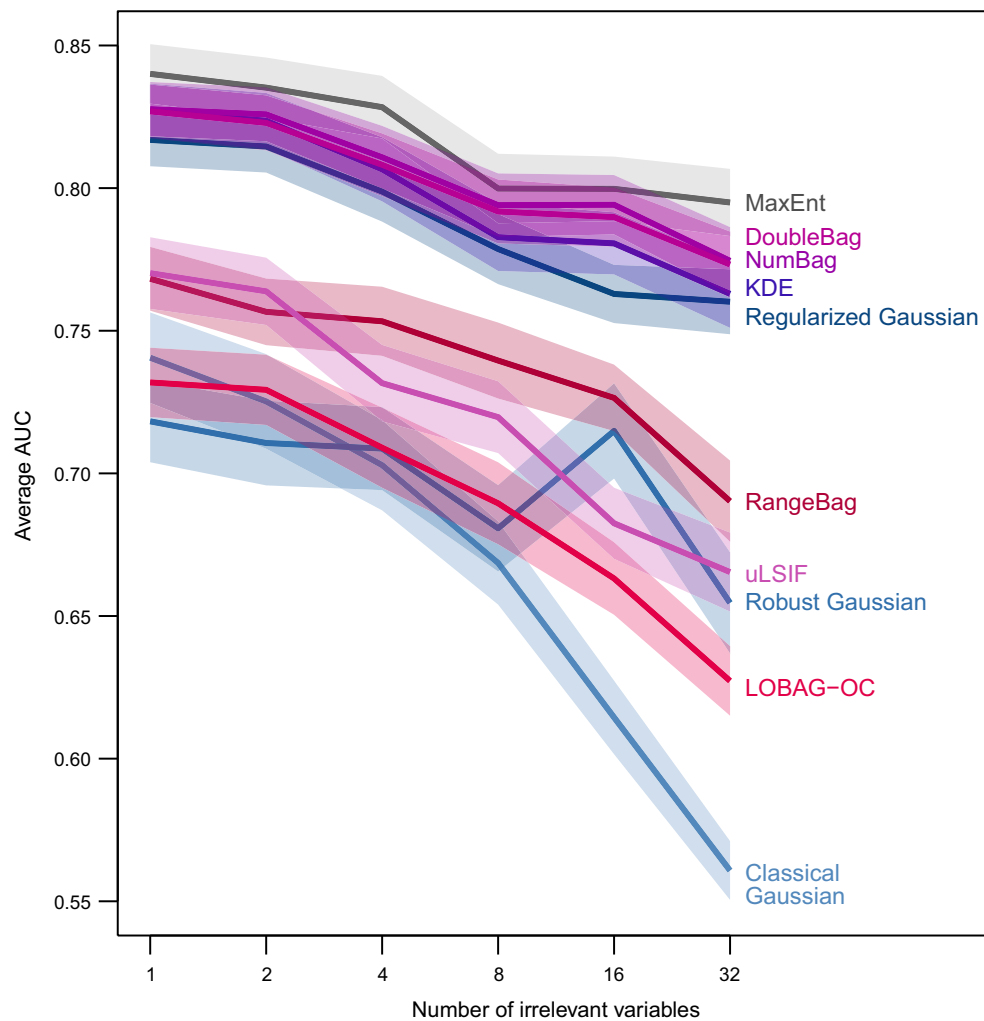


Fig. 3. Performance of six plug-and-play species distribution models compared with two presence-only species distribution models (LOBAG-OC and RangeBag), MaxEnt, and a density ratio estimator in the presence of irrelevant variables. Performance was degraded the least by irrelevant variables in the regularized Gaussian, kernel density estimates (KDE), DoubleBag, NumBag, and MaxEnt models (top panel; error bars are mean \pm standard error).

where those theories make probabilistic statements about occurrence within a landscape (i.e., metapopulation theories) or persistence over long times (i.e., stochastic population theory). A practically useful feature of this framework is that it suggests a new approach—the plug-and-play method—for estimating relative suitability.

Performance of six versions of the plug-and-play method was studied with respect to 106 plant species that had been exhaustively sampled for species presence and absence. Mean AUC of these models ranged from 0.717 to 0.841. Except for the plug-and-play approach with classical and

robust Gaussian components, the current models are all superior to those of Drake et al. (2006), which used support vector machines in three different ways and achieved average AUC of <0.8 in all cases. On the other hand, Drake et al. (2006) fit models only to data on species occurrences (i.e., presence-only models) whereas plug-and-play is a presence-background approach. Similarly, Maher et al. (2014) studied nine presence-only and seven presence-absence niche modeling methods using these data. All 16 of those methods returned average AUC of <0.8 , with the exception of a k -nearest neighbor method configured to fit

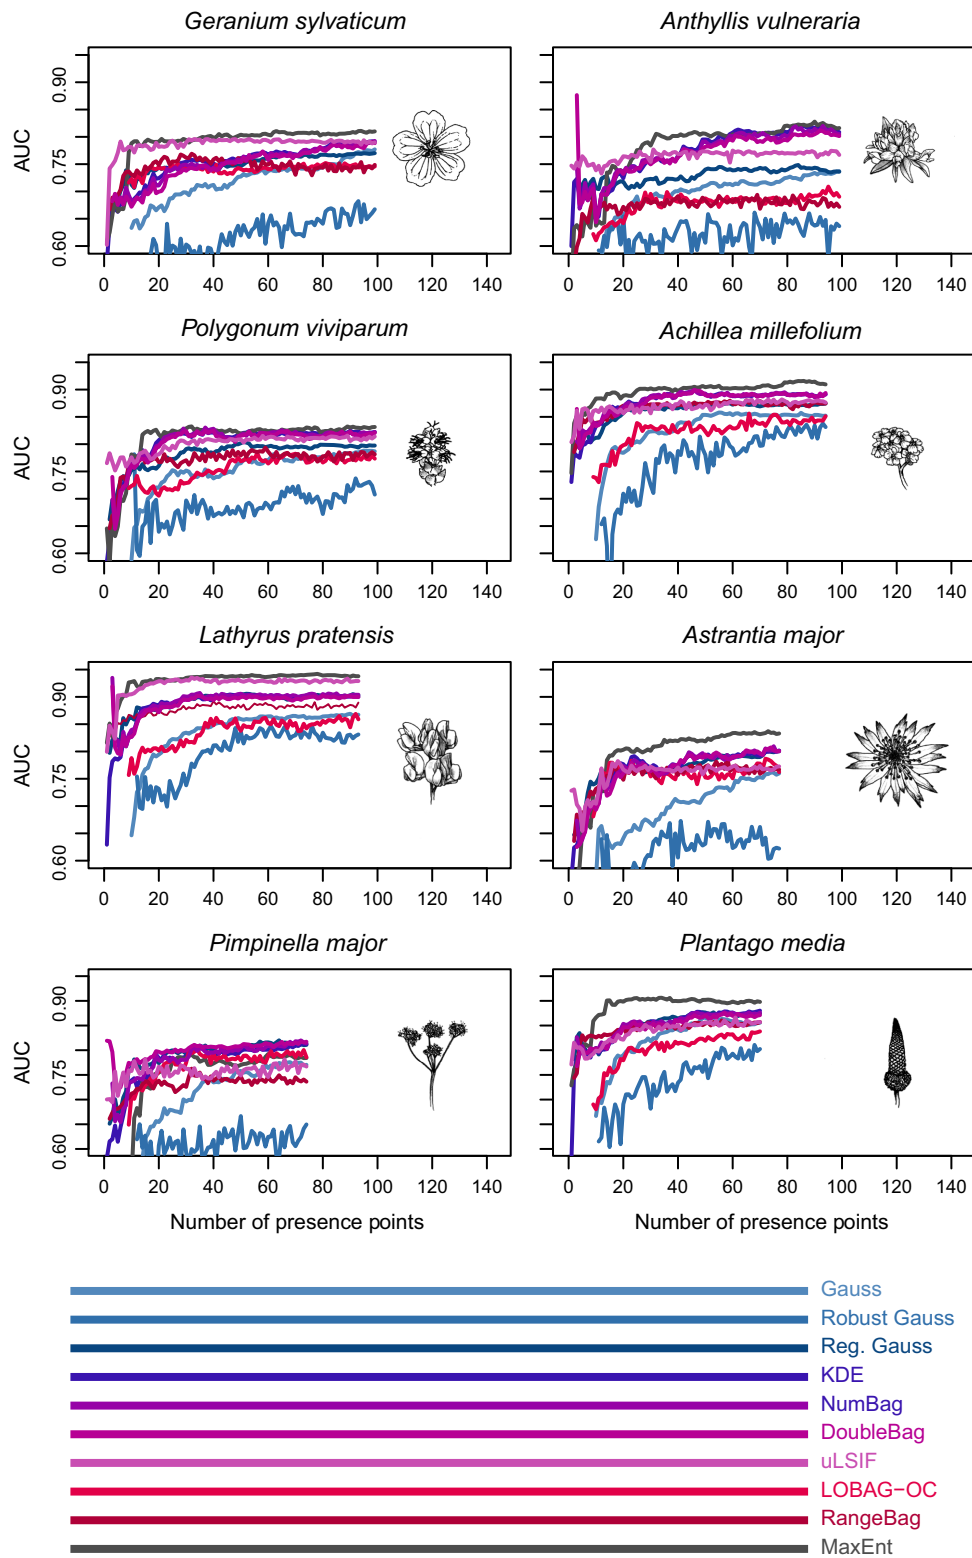


Fig. 4. Learning rates of six plug-and-play models, two presence-only models, MaxEnt, and a density ratio

(Fig. 4. *Continued*)

estimator on data for eight abundant species. Models generally perform more poorly at low numbers of training points and gain performance as more training points are offered. Kernel density estimates (KDE), NumBag, DoubleBag, and regularized Gaussian tended to approach a performance plateau at large samples sizes while uLSIF shows inconsistent behavior, occasionally outperforming all other models at all sample sizes (e.g., *Geranium sylvaticum*), but more commonly showing superior performance only at small sample sizes (e.g., *Anthyllis vulneraria*).

both presence and absence data where the average AUC was 0.804. Thus, the presence-background approaches studied here (with the exception of classical and robust Gaussian components) are superior to all the presence-only and presence-absence methods studied by Maher et al. (2014) with AUC improvements on the order of 4–6%. One conclusion from this study together with earlier analyses (Drake et al. 2006, Drake 2014, Maher et al. 2014) is that there may yet be scope for improvement, but also that many methods may be contrived to yield very similar results. Particularly, a number of methods perform similarly to MaxEnt, which is currently a widely preferred method. Ultimately, if ecologists wish to choose a method on the basis of empirical superiority (compare Brotons et al. 2004, Elith et al. 2006), many more data sets will be needed.

One of the key assumptions of the approaches introduced here is that environments are constant and species distributions do not primarily reflect past environmental conditions. Modeling environmental suitability in the presence of *dynamic environments* is an important problem for further research. Extensions of the current approach to dynamic environments are easy to imagine. Suppose we supplement our spatial indexing (i) with a temporal indexing (t). Occurrence records also must be indexed by time. Recognizing that a location's past environment affects its present species composition, one would model the probability of occurrence as a function of both present and past environments, perhaps downweighting the effect of conditions in the distant past or considering only those environmental conditions within a window of time. Now, the probability of occurrence is given by a function of both present and past environments or a weighted mixture of densities. The definition of environmental suitability with respect to the present time is retained, however. A simpler way to implement this idea is to introduce time-lagged covariates into the feature set, that is, $z_{i,t} = [a_{i,(t)}, a_{i,(t-1)}, a_{i,(t-2)}, \dots, b_{i,(t)}, b_{i,$

$(t-1)}, b_{i,(t-2)}, \dots]$. Joint densities f and f_1 could be fit as done here and present suitability returned either by marginalizing over historical environments or by evaluating the full model (including time-lagged variables) for all locations' actual environmental histories. As the current study of irrelevant variables shows, if the lagged environmental variables are not grossly irrelevant or too numerous, the effects they have on a suitably chosen model (i.e., plug-and-play approach with KDE or regularized Gaussian components or MaxEnt) are minimal, and we may expect the resulting model to be fairly robust.

In principle, it might seem that this approach also makes the assumption that the species distribution is at equilibrium (since it is assumed that f and f_1 are stationary). Applied naively to all the data that one might acquire, this is true. However, because the problem does not explicitly include spatial dependence, the modeler can easily restrict the data to a subset prior to estimating f and f_1 . Thus, for instance, the set of records for estimating f and f_1 could be just the locations that were searched for the species. (In fact, this is what has been done in any study where the background data reflect only study locations, including the current study.) Alternatively, the data could be selected to be drawn only from regions considered to be "accessible" to the species (Peterson et al. 2011). This is very similar to the problem of study extent addressed by Barve et al. (2011).

This paper advocates estimating relative suitability or niche boundary rather than absolute suitability in applications where species' prevalence is irrelevant (for instance to rank sites by their value to species conservation). The plug-and-play approach nevertheless solves an unnecessary intermediate problem: estimating the densities f and f_1 (or, in the case of sample selection bias, f_2). It would seem that estimating the ratio f_1/f directly might be a more efficient approach, as has been advocated by Sugiyama et al. (2013). In comparative studies, Sugiyama et al. found

models based on density ratio estimation to be superior to models based on density estimation itself (Sugiyama et al. 2013). The empirical analysis reported here shows this not to be generally true, however, as uLSIF exhibited performance inferior to plug-and-play methods. Except for MaxEnt (Elith et al. 2011), to our knowledge, density ratio estimation has not previously been used for species distribution modeling. What, then, explains the superior performance of MaxEnt compared with uLSIF? Possibly it is due to the fitting criterion, which MaxEnt stipulates to be the Kullback-Leibler divergence between \hat{f}_1 and \hat{f} . In contrast, plug-and-play methods propose no such constraint. Thus, for instance, in the current study the alternative component distributions are chosen to meet different criteria applying to different plausible conditions that might restrict the performance of a simple estimator (like the classical Gaussian). As it turns out, some of these options degrade performance (i.e., robust Gaussian components), while others improve it (i.e., KDE and regularized Gaussian components). The set of possible estimators for the component densities f_1 and f and the rationale for choosing among them is an important problem for further study. For now, we recommend the plug-and-play approach with KDE or regularized Gaussian components, or MaxEnt when computing resources are limited. In contrast, when the size of the data is small or when computing resources are not a concern, NumBag and DoubleBag are our methods of choice.

ACKNOWLEDGMENTS

This work was begun in 2012 while JMD was a Leverhulme Visiting Professor at Oxford University. Additional support was provided by the National Oceanic and Atmospheric Administration Agreement No. NA10NOS4780218. Annakate Schatz assisted with the preparation of the manuscript. JD proposed the model, JD and RR wrote the code and performed the analysis, JD wrote the first draft of the manuscript, and all authors contributed substantially to revisions.

LITERATURE CITED

- Barve, N., V. Barve, A. Jiménez-Valverde, A. Lira-Noriega, S. P. Maher, A. T. Peterson, J. Soberón, and F. Villalobos. 2011. The crucial role of the accessible area in ecological niche modeling and species distribution modeling. *Ecological Modelling* 222: 1810–1819.
- Breiman, L. 1996. Bagging predictors. *Machine Learning* 24:123–140.
- Breiman, L. 2001. Random forests. *Machine Learning* 45:5–32.
- Brottons, L., W. Thuiller, M. B. Araújo, and A. H. Hirzel. 2004. Presence-absence versus presence-only modelling methods for predicting bird habitat suitability. *Ecography* 27:437–448.
- Drake, J. M. 2014. Ensemble algorithms for ecological niche modeling from presence-background and presence-only data. *Ecosphere* 5:art76.
- Drake, J. M. 2015. Range bagging: a new method for ecological niche modelling from presence-only data. *Journal of the Royal Society Interface* 12:20150086.
- Drake, J. M., and J. M. Bossenbroek. 2009. Profiling ecosystem vulnerability to invasion by zebra mussels with support vector machines. *Theoretical Ecology* 2:189–198.
- Drake, J. M., C. Randin, and A. Guisan. 2006. Modelling ecological niches with support vector machines. *Journal of Applied Ecology* 43:424–432.
- Elith, J., C. H. Graham, R. P. Anderson, M. Dudík, S. Ferrier, A. Guisan, R. J. Hijmans, F. Huettmann, J. R. Leathwick, A. Lehmann, and J. Li. 2006. Novel methods improve prediction of species' distributions from occurrence data. *Ecography* 29:129–151.
- Elith, J., M. Kearney, and S. J. Phillips. 2010. The art of modelling range-shifting species. *Methods in Ecology and Evolution* 1:330–342.
- Elith, J., and J. R. Leathwick. 2009. Species distribution models: ecological explanation and prediction across space and time. *Annual Review of Ecology, Evolution, and Systematics* 40:677–697.
- Elith, J., J. R. Leathwick, and T. Hastie. 2008. A working guide to boosted regression trees. *Journal of Animal Ecology* 77:802–813.
- Elith, J., S. J. Phillips, T. Hastie, M. Dudík, Y. E. Chee, and C. J. Yates. 2011. A statistical explanation of MaxEnt for ecologists. *Diversity and Distributions* 17:43–57.
- Engler, R., A. Guisan, and L. Rechsteiner. 2004. An improved approach for predicting the distribution of rare and endangered species from occurrence and pseudo-absence data. *Journal of Applied Ecology* 41:263–274.
- Franklin, J. 2010. Mapping species distributions: spatial inference and prediction. Cambridge University Press, Cambridge, UK.
- Fretwell, S. D., and H. L. Lucas. 1969. On territorial behavior and other factors influencing habitat distribution in birds – I. Theoretical development. *Acta Biotheoretica* 19:16–36.

- Grinnell, J. 1917. The niche-relationships of the California thrasher. *Auk* 34:427–433.
- Hanski, I. 1989. Metapopulation dynamics: Does it help to have more of the same? *Trends in Ecology and Evolution* 4:113–114.
- Hastie, T., and T. Fithian. 2013. Inference from presence-only data: the ongoing controversy. *Ecography* 36:864–867.
- Hayfield, T., and J. S. Racine. 2008. Nonparametric econometrics: the np package. *Journal of Statistical Software* 27:1–32.
- Helmbold, D. P., and P. M. Long. 2012. On the necessity of irrelevant variables. *Journal of Machine Learning Research* 13:2145–2170.
- Hido, S., Y. Tsuboi, H. Kashima, M. Sugiyama, and T. Kanamori. 2011. Statistical outlier detection using direct density ratio estimation. *Knowledge and Information Systems* 26:309–336.
- Hijmans, R. J. 2012. Cross-validation of species distribution models: removing spatial sorting bias and calibration with a null model. *Ecology* 93:679–688.
- Hirzel, A. H., and G. Le Lay. 2008. Habitat suitability modelling and niche theory. *Journal of Applied Ecology* 45:1372–1381.
- Holt, R. D. 2009. Bringing the Hutchinsonian niche into the 21st century: ecological and evolutionary perspectives. *Proceedings of the National Academy of Sciences of the USA* 106:19659–19665.
- Huber, P. J., and E. M. Ronchetti. 2009. *Robust statistics*. Wiley, Hoboken, New Jersey, USA.
- Hurlbert, S. H. 1981. A gentle depilation of the niche: Dicean resource sets in resource hyperspace. *Evolution Theory* 5:177–184.
- Kanamori, T., S. Hido, and M. Sugiyama. 2009. A least-squares approach to direct importance estimation. *Journal of Machine Learning Research* 10:1391–1445.
- Keating, K. A., and S. Cherry. 2008. Use and interpretation of logistic regression in habitat selection studies. *Journal of Wildlife Management* 68:774–789.
- Kotz, S., N. Balakrishnan, and N. Johnson. 2000. *Continuous multivariate distributions, volume 1: models and applications of series in probability and statistics*. John Wiley & Sons, New York, New York, USA.
- Ledoit, O., and M. Wolf. 2003. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance* 10:603–621.
- Li, Q., and J. S. Racine. 2003. Nonparametric estimation of distributions with categorical and continuous data. *Journal of Multivariate Analysis* 86: 266–292.
- Maher, S. P., C. F. Randin, A. Guisan, and J. M. Drake. 2014. Pattern-recognition ecological niche models fit to presence-only and presence-absence data. *Methods in Ecology and Evolution* 5:761–770.
- Mainali, K. P., D. L. Warren, K. Dhileepan, A. McConachie, L. Strathie, G. Hassan, D. Karki, B. B. Shrestha, and C. Parmesan. 2015. Projecting future expansion of invasive species: comparing and improving methodologies for species distribution modeling. *Global Change Biology* 21:4464–4480.
- Orr, H. A. 2009. Fitness and its role in evolutionary genetics. *Nature Reviews Genetics* 10:531–539.
- Pearce, J. L., and M. S. Boyce. 2006. Modelling distribution and abundance with presence-only data. *Journal of Applied Ecology* 43:405–412.
- Peters, R. H. 1976. Tautology in evolution and ecology. *American Naturalist* 110:1.
- Peterson, A. T., J. Soberón, R. G. Pearson, R. P. Anderson, E. Martínez-Meyer, M. Nakamura, and M. B. Araújo. 2011. *Ecological niches and geographic distributions*. Monographs in population biology. Princeton University Press, Princeton, New Jersey, USA.
- Phillips, S. J., M. Dudík, J. Elith, C. H. Graham, A. Lehmann, J. R. Leathwick, and S. Ferrier. 2009. Sample selection bias and presence-only distribution models: implications for background and pseudo-absence data. *Ecological Applications* 19:181–197.
- Phillips, S. J., M. Dudík, and R. E. Schapire. 2004. A maximum entropy approach to species distribution modeling. *Proceedings of the Twenty-First International Conference on Machine Learning* 69:655–662.
- Phillips, S. J., and J. Elith. 2010. POC plots: calibrating species distribution models with presence-only data. *Ecology* 91:2476–2484.
- Pulliam, H. R. 2000. On the relationship between niche and distribution. *Ecology Letters* 3:349–361.
- R Core Team. 2017. *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Rousseeuw, P. J., and K. Van Driessen. 1999. A fast algorithm for the minimum covariance determinant estimator. *Technometrics* 41:212–223.
- Royle, J. A., R. B. Chandler, C. Yackulic, and J. D. Nichols. 2012. Likelihood analysis of species occurrence probability from presence-only data for modelling species distributions. *Methods in Ecology and Evolution* 3:545–554.
- Schäfer, J., R. Opgen-Rhein, V. Zuber, M. Ahdesmäki, A. P. D. Silva, and K. Strimmer. 2015. corpcor: efficient estimation of covariance and (partial) correlation. <http://cran.r-project.org/package=corpcor>
- Sugiyama, M., T. Kanamori, and T. Suzuki. 2013. *Density ratio estimation in machine learning*. Cambridge University Press, New York, New York, USA.
- VanDerWal, J., L. P. Shoo, C. N. Johnson, and S. E. Williams. 2009. Abundance and the environmental

- niche: Environmental suitability estimated from niche models predicts the upper limit of local abundance. *American Naturalist* 174:282–291.
- Vapnik, V. 1998. *Statistical learning theory*. Wiley, New York, New York, USA.
- Wand, M. P., and M. C. Jones. 1995. *Kernel smoothing*. Chapman & Hall/CRC, Boca Raton, Florida.
- Wang, J., et al. 2014. *robust: robust library*. <http://cran.r-project.org/package=robust>
- Ward, G., T. Hastie, S. Barry, J. Elith, and J. R. Leathwick. 2009. Presence-only data and the EM algorithm. *Biometrics* 65:554–563.
- Wiens, J. J., D. D. Ackerly, A. P. Allen, B. L. Anacker, L. B. Buckley, H. V. Cornell, E. I. Damschen, T. J. Davies, J. A. Grytnes, S. P. Harrison, and B. A. Hawkins. 2010. Niche conservatism as an emerging principle in ecology and conservation biology. *Ecology Letters* 13:1310–1324.

DATA AVAILABILITY

Data and code are archived at Dryad (<https://doi.org/10.5061/dryad.g5p7d1c>).

SUPPORTING INFORMATION

Additional Supporting Information may be found online at: <http://onlinelibrary.wiley.com/doi/10.1002/ecs2.2373/full>