Integrating Fundamental Concepts of Ecology,

Biogeography, and Sampling into Effective Ecological

Niche Modeling and Species Distribution Modeling

by A. Townsend Peterson and Jorge Soberón

2012

Integrating Fundamental Concepts of Ecology, Biogeography, and Sampling into Effective Ecological Niche Modeling and Species Distribution Modeling

A. Townsend Peterson and Jorge Soberón

Biodiversity Institute, The University of Kansas, Lawrence, Kansas 66045 USA

Abstract

Correlative techniques for estimating environmental requirements of species—variably termed ecological niche modeling or species distribution modeling—are becoming very popular tools for ecologists and biogeographers in understanding diverse aspects of biodiversity. These tools, however, are frequently applied in ways that do not fit well into knowledge frameworks in population ecology and biogeography, or into the realities of sampling biodiversity over real-world landscapes. We offer 10 'fixes'—adjustments to typical methodologies that will take into account population ecological and biogeographic frameworks to produce better models.

Introduction

The past 15 years have seen a massive increase in the popularity of techniques that link known occurrences of species with environmental variation

across landscapes to estimate ecological niches and geographic distributions, generally termed ecological niche modeling or species distribution modeling (for further discussion regarding this latter term, see below). The literature taking advantage of this novel analytical functionality has increased massively (Figure 1), and two book-length syntheses have now appeared (Franklin 2010, Peterson et al. 2011). Two recent papers have seen massive citation in the field—Elith et al. (2006) has been cited 1050 times, and Phillips et al. (2006) has been cited 842 times (Web of Science, consulted 30 January 2012)—such massive attention in the literature indicates considerable popularity. More importantly, these coarse-resolution summaries of ecology and distribution have been incorporated into the basic 'toolkit' of the macroecologist and biogeographer, such that optimizing their use and implementation becomes critical.

Many uses of niche modeling in the literature, however, have been rather inappropriate. That is, the computational tools that have been developed for niche modeling are easily used, and frequently have been used in ways that are not in good accord either with what is known of population biology of species, ideas from modern biogeography, or the realities of sampling biodiversity phenomena across real-world landscapes. These misuses, unfortunately, detract from the genuine potential utility of the tools, and cause mistrust and misunderstanding on the part of the broader biodiversity science community.

A comparison with another field of biodiversity science is perhaps illustrative. Modern phylogenetics can arguably be stated to have begun with the publication of Willi Hennig's framework for reconstruction of evolutionary history

(Hennig 1950). This thinking framework preceded by 2-3 decades the appearance of computational tools for implementation and use of phylogenetic thinking in systematics—the first software packets for cladistic analyses did not appear until the 1980s (Wiley 1981), and were not broadly available for several years more (Felsenstein 1986, Berlocher and Swofford 1997). As such, the thinking framework for cladistics was available long before the technique was easy to implement. Niche modeling, however, has seen the opposite evolutionary trajectory: tools that, in effect, estimate niches have been around for decades (Nix 1986, Austin et al. 1990, Stockwell and Noble 1992), yet a conceptual framework for the technique has been much slower to appear (Soberón and Peterson 2005, Soberón 2007, Soberón and Nakamura 2009, Godsoe 2010, Soberón 2010, Peterson et al. 2011). We argue that this mismatch between practice and theory has handicapped the development of this emerging field, and has limited the inferences that have been possible.

In this contribution, we outline 10 critical considerations that must be taken into account in development of ecological niche models as powerful tools in ecology and biogeography. In each case, the consideration is not widely appreciated in this field; some have appeared in the literature, and others will appear soon. The point, nonetheless, is that, for lack of a solid conceptual framework for the field, key conceptual-to-empirical links have failed, and the field has been handicapped as a consequence.

BAM Scenario

The niche concept was originally presented in a verbal model (Grinnell 1917, 1924) that lacked formal structure. A more quantitative and comprehensive presentation (Hutchinson 1957), however, described biotic and abiotic dimensions of niche, but failed to acknowledge spatial considerations in structuring the environmental and geographic distributions of species. Addition of this third dimension of determinants of species' distributions did not appear broadly in the ecological literature until much later (Pulliam 2000, Soberón and Peterson 2005). This three-part concept (**B**iotic, **A**biotic, **M**obility, or **BAM**) of how species' distributions are structured now provides a useful guide in developing ecological niche models.

The **BAM** concept envisions a universe of areas under consideration **G**, within which the species encounters areas presenting appropriate abiotic conditions that can be labeled **A**, and areas presenting appropriate biotic conditions **B**. The intersection of these two areas ($\mathbf{A} \cap \mathbf{B} = \mathbf{G}_P$) is the area that is fully suitable for the species, or what can be termed the potential distribution of the species. Not all of $\mathbf{G}_P$, however, is readily available to the species, perhaps being inaccessible owing to the presence of a dispersal barrier or simply being too far away—this consideration of access is conceptualized as **M**. We can term the area that is not accessible but that is suitable as the invadable distributional area, or $\mathbf{G}_I$. As a result, the "actual" or occupied distributional area $\mathbf{G}_O = \mathbf{A} \cap \mathbf{B} \cap \mathbf{M} = \mathbf{G}_P \cap \mathbf{M}$.

This **BAM** scenario (see Figure 2 for a visual presentation of these same ideas) will guide much of concept-based thinking in ecological niche modeling.

For instance, if we assume **B** not to constitute a strong constraint on species' distributions, of the resulting four "**AM**" configurations [e.g., $\mathbf{A} \subset \mathbf{M}$, $\mathbf{M} \subset \mathbf{A}$, $\mathbf{M} \approx \mathbf{A}$, $\mathbf{G}_O \subset \mathbf{G}_P = \mathbf{M} \cap \mathbf{A} \subset \mathbf{G}_A$], niche models can be developed successfully for two, but niche models are rarely if ever better than random expectations for the other two ($\mathbf{A} \subset \mathbf{M}$, $\mathbf{M} \approx \mathbf{A}$). These **BAM** configurations are not uncommon in nature (e.g., landscapes that are highly dissected, with limited area-to-area access for individuals of species), and thus **BAM**-based considerations have obvious and important implications for niche modeling (Saupe et al. In review).

Fundamental *versus* Existing Fundamental Ecological Niche

Working from the same **BAM** framework, an immediate corollary is simultaneously subtle and of critical importance. If **A** is determined by physiological limitations of the species, which can be termed the fundamental ecological niche ($\mathbf{N}_F$), but if the broader representation of those conditions on real landscapes is constrained also by geographic factors, not all of the conditions within $\mathbf{N}_F$ are necessarily available to the species for colonization (Figure 2). That is, either **G** or **M** may limit the manifestation of the species' ecological breadth to some subset of $\mathbf{N}_F$, which we term the existing fundamental ecological niche and denote as $\mathbf{N}_F^*$. The existing fundamental niche may, in turn, be reduced still further by interactions with other species (Hutchinson 1957, 1978).This concept, namely that the real-world manifestation of a fundamental ecological niche is invariably only partial (Colwell and Futuyma 1971), turns out to have very real implications for many aspects of this field (Soberón and Peterson 2011).

A particularly important implication of these ideas is the following inequality:

$$\mathbf{N}_R \subseteq \mathbf{N}_F^* = \eta(\mathbf{M}) \bigcap \mathbf{N}_F \subseteq \mathbf{N}_F$$

where $\eta(\mathbf{M})$ denotes the set of environments manifested within $\mathbf{M}$. As pointed out by Soberón and Peterson (2011), this inequality immediately handicaps any attempts to use observed environmental distributions of species as a proxy for niche estimates, particularly in testing for niche differentiation: quite simply, such tests (e.g., Broennimann et al. 2007, Fitzpatrick et al. 2007, Medley 2010) will be confused by additional factors that "distort" the estimate of the fundamental niche: those related to uneven sampling of existing environments, and those related to interactions with other species (Peterson 2011). Rather, the only tests for such differentiation that will be appropriate will be those that take $\mathbf{M}$ into account explicitly in the tests (e.g., background similarity tests in Warren et al. 2008).

This inequality also has less-obvious implications for model calibration. Specifically, because $\mathbf{M}$ affects the set of areas and environments that will be available to the species, the known set of occurrences $\mathbf{G}_+$ must come from $\mathbf{G}_O$, such that the associated environments $\eta(\mathbf{G}_+)$ are already filtered by $\mathbf{M}$. The result is that niche models that have been calibrated based on $\mathbf{M}$-limited areas will frequently under-characterize the true niches of the species in question—an excellent illustration of this point can be drawn from the criticisms of the work of

Beale et al. (2008) by two other lab groups (Araújo et al. 2009, Jiménez-Valverde et al. 2010).

Models of Niche or Models of Distribution?

The **BAM** framework also allows clear-minded reflection on the issue of what this suite of techniques should most appropriately be named. A Google Scholar search (20 Feb 2012) on "ecological niche model" yields 659 matches, while a search on "species distribution model" yields 831. Hence, the field appears more or less evenly divided between the two terminologies.

In a **BAM** framework, however, the resolution of the debate is more or less clear. These models that relate known occurrences of species to underlying environmental characteristics seek an environmental association of G_O. Nonetheless, **M** will frequently be defined in terms that are not characterizable in terms of environment—it might be a fine barrier, such as a river, or a hard barrier that is in effect the "end of the world," such as an ocean for terrestrial species. As a result, modeling algorithms will estimate G_P more directly, rather than G_O, which would require a hypothesis of the geometry of **M**. Put another way, estimating G_O requires information beyond the usual occurrence and environment characterization that are fed into niche/distribution modeling algorithms.

M Governs Everything

A recent publication (Barve et al. 2011) emphasized the crucial role of **M** in determining the outcome of many niche-model-related exercises. That is, **M** delineates the set of areas to which the species has had access over relevant time periods. As a consequence, only areas within **M** have the potential to offer presence records, and only areas within **M** offer a clear, environment-based interpretation of absence data—in these areas, the species has had the potential to visit, but has not established populations there. As a result, **M** becomes the critical arena for development of these models; Barve et al. (2011) demonstrated that **M** affects model calibration, model evaluation, and model comparisons in important ways, and therefore hypotheses of **M** in effect predetermine almost all results in niche modeling studies.

To give some illustrations of this point, models calibrated within areas that are overly broad (i.e., including areas that were not accessible to the species over relevant time periods) will inevitably be confused. Areas within **G**$_I$ will be tallied as absence data, even though they present conditions that are perfectly suitable for the species to establish populations. Furthermore, because most algorithms have some convergence criterion, models calibrated over too-broad areas will tend to be overly general. At the other end of the spectrum, if the definition of **M** is too narrow, models will be calibrated based on comparisons that are not particularly comprehensive, and will thus be representative of too little of environmental space (see discussion of MESS below). Barve et al. (2011) also documented significant effects of **M** definitions on model evaluations (too-broad definitions of **M** make for easy conclusion of significant predictivity of models)

and model comparisons (too-broad definitions of **M** make for easy conclusion of significant model similarity). Hence, careful choice of **M** for such studies is crucial, and biological and methodological bases for these choices must be stated clearly and explicitly as part of the Methods sections in any publications of these analyses.

S-intersect-M as Study Area

A modification of the basic schema of **M** as determining the area of analysis takes into account the fact that not all sectors of the distribution of the species may have been sampled equally thoroughly (we will refer to the area that was sampled to the point that occurrences have some probability of being detected as **S**). In particular, in many cases, some areas are very well sampled, while others remain unsampled or only lightly sampled. These imbalances can introduce biases into model calibration, and can change model results rather dramatically (Ward et al. 2009, Peterson et al. 2011).

In concept, analysis should not be limited just to **M**, as we have argued before (Barve et al. 2011), but rather to the area that is delimited by $\mathbf{M} \cap \mathbf{S}$, with implications for how one delimits the study area and for how one includes or excludes known occurrence points. The problem with this observation is that sampling 'bias' can be spatial or environmental in nature. That is, an ecological niche is in itself a sort of bias of distribution in environmental space; the sampling biases to which we are referring in this section are biases in geographic space, in

which some regions are sampled more intensively than others, regardless of whether the particular species of interest is detected.

The quick and easy approach to these sampling biases is to detect oversampled areas visually, and to reduce their sampling density via manual subsampling at random; this approach, however, can be quite subjective, and does not capture the fine details of the sampling landscape. A more rigorous approach, however, is to characterize the sampling landscape quantitatively, as was explored in presentation of a methodology for evaluating probability of true absences (Anderson 2003)—basically, under circumstances in which sampling of the species of interest was assembled by application of a standard methodology that accumulated other species as well, one can create a picture of sampling intensity by summarizing the sampling of those other species. This surface of sampling intensity can then be used to weight presence points such that each is weighted commensurate with the intensity of sampling from that region. For a more technical summary of problems involved in sampling for niche modeling, see Ward et al. (2009) and Phillips et al. (2009).

Correct Balance of + and – Occurrence Error Weightings

A probabilistic summary of the generation of presence and absence data that describe the distributions of species leads to some useful insights into the relative weighting that each should be given (Peterson et al. 2011). Specifically, presence data are relatively rarely misleading: only when errors of identification or of georeferencing create a false positive record would presence data not be

useful and informative. On the other hand, absence data are much more prone to being misleading: here, areas outside of **M** will be 'absent' regardless of whether conditions are suitable or not. What is more, any non-environmental cause of apparent absence will produce false negatives as well: extirpation of populations owing to anthropogenic pressures, lack of sampling of sites, non-detection of the species during sampling, or non-reporting of members of the species that were detected. In short, many factors will produce false-negative data, but relatively few will produce false positives—this imbalance suggests that the two types of error should be weighted differently in niche modeling applications.

These considerations of differential weighting of false-positive and false-negative errors can improve model calibration dramatically. Most modeling algorithms optimize some quantity—for example, regression-based approaches might minimize least-squares model deviations, while evolutionary-computing approaches maximize a measure of predictive accuracy (Stockwell and Noble 1992). These calibration objectives, however, do not necessarily incorporate the correct balancing of error types, such that the models that result are not optimal. Improved approaches may either re-weight error types in the actual calibration process, or may instead filter replicate models post-calibration to extract those solutions that balance error components appropriately (Anderson et al. 2003).

Error weighting also enters the picture importantly in model evaluations (Peterson et al. 2011). Common approaches such as the Kappa statistic and the area under the curve (AUC) of the receiver operating characteristic (ROC) curve, unfortunately, weight the two components of error equally (Fielding and Bell

1997, Lobo et al. 2008, Peterson et al. 2008), and thereby are not appropriate measures of model predictivity in such analyses. Indeed, given that 'absence' data are almost never genuinely at hand, these approaches must commonly rely on pseudoabsence data of some sort, which becomes an arbitrary and non-objective exercise (Peterson et al. 2008, Barve et al. 2011). Correct weighting of the two error types will depend on the specific characteristics of a particular study region, environmental dimensions, and occurrence data.

"E" in Model Calibration / Thresholding / Evaluation / Comparison

The errors mentioned above that are inherent in data documenting presences of species are dependent on numerous considerations, and must be pondered carefully for each analysis. Peterson et al. (2008) introduced the idea of a parameter E, which is an estimate of the proportion of presence data that are likely to be sufficiently erroneous as to place the species under inappropriate environmental conditions. E will be higher when data are 'found,' and not collected specifically and carefully for a given project, and also when spatial resolution is finer. E was originally proposed as part of revised methodologies for model evaluations (Peterson et al. 2008), but the concept turns out to be critical in a number of situations.

E is quite useful in model calibration and thresholding, as it offers a means to align model thresholds with performance regarding known occurrences of the species. That is, the least training presence thresholding approach (Pearson et al. 2007) offers the correct prioritization of false-negative error over false-positive

error, but does not consider that false-positive error rates may be non-zero; E offers a means of incorporating these potential errors. While the least training presence approach seeks the threshold that includes 100% of the presence data used to calibrate the model (T_{100}), a modification of this approach considers that some of those presence data may be erroneous, and seeks instead the threshold that includes (100 − E)% of the presence data (T_{100-E}). This thresholding approach both prioritizes false-negative error over false-positive error and takes into account the error inherent in the presence data available.

More briefly, in model evaluation, E also provides a useful consideration of expected error *versus* undesired error. Specifically, Peterson et al. (2008) used E to outline the area of ROC space that includes desirable model predictions, as opposed to predictions that are not useful or informative. By specifying a lower value of E, one can focus model predictions on the challenge of anticipating the entire distribution of a species, rather than on partial characterization. Finally, in model comparisons (Warren et al. 2008), E-adjusted thresholds become critical in avoiding meaningless nonsimilarity measures caused by overfitting of model predictions to well-sampled areas (Peterson et al. 2007), which can lead to spurious conclusions of niche differentiation (Maher et al. 2010, Peterson 2011).

Overfitting, Dimensionality, and Complexity of Models

Ecological niche models are frequently calibrated in highly dimensional environmental spaces, where they run considerable risk of overfitting (Peterson et al. 2011). That is, if care is not taken, models calibrated in highly dimensional

spaces will be overly specific, and will not be able to anticipate phenomena that are manifested under slightly different conditions—an excellent example is that of recent modeling efforts regarding fire ant (*Solenopsis invicta*) distributional potential that purported to document significant niche differentiation between native and invasive popultions (Fitzpatrick et al. 2007), but that turned out to be highly dependent on the environmental data employed, and particularly on their dimensionality (Peterson and Nakazawa 2008). Such overfitting has been a common cause of incorrect conclusions of niche differentiation in comparisons among species and populations (Peterson 2011).

Controlling the dimensionality of environmental spaces in niche modeling exercises is a rather complex task. One line of thinking focuses on pre-selecting a reduced set of environmental dimensions that is particularly important to species' distributional ecology (Huntley et al. 2008), but such approaches run the risk of missing critical dimensions, or under-informing models in the calibration process. More commonly, researchers attempt to reduce dimensionality of environmental spaces by removing redundant information—because intercorrelations among environmental dimensions are rampant (Jiménez-Valverde et al. 2009), it is quite feasible to reduce dimensionality either by removing members of highly correlated variable pairs manually, or via principal components analysis. These steps not only make efficient and effective calibration more feasible, but also simplify model interpretation considerably.

Spatial Autocorrelation

Another consideration is that of spatial autocorrelation—indeed, a common idea is that the 'only law' of geography is that things nearby tend to be similar, and things far away tend to be less similar. This frequent phenomenon of spatial autocorrelation can complicate niche modeling applications because occurrence points may not be independent of one another, simply because they are nearby (Diniz-Filho et al. 2003). This non-independence can cause problems in model calibration by artificially over-emphasizing certain environmental conditions, rather than allowing the algorithm to fit a model that covers the entire environmental breadth of the species; it also causes problems in model evaluation by artificially inflating sample sizes for testing model predictions (Diniz-Filho et al. 2003, Segurado et al. 2006, Dormann et al. 2007, Peterson et al. 2011).

Clearly, it is desirable to understand and incorporate spatial autocorrelation in niche modeling applications, although this factor is perhaps not as 'fatal' as has been implied by some (see discussions in Diniz-Filho et al. 2003). Several GIS and analysis programs provide the possibility of calculating spatial lag distances for environmental data sets—these distances are the distance over which the proximity effect no longer holds, such that points separated by these distances will be independent of one another. These distances can range over 10^1-10^3 km, and constitute a serious constraint on sample sizes in niche modeling: that is, the raw number of presence points that one has may not matter much, if they are separated by distances that are less than the spatial lag of the particular environmental dimensions in question.

In practice, once the lag distance has been calculated for each environmental dimension, occurrence data would ideally be filtered to be separated by at least the minimum of the lag distances across all of the environmental dimensions under consideration. In many cases, however, this step will reduce the number of presence records so much as to leave too few for model calibration; an alternative approach is to seek new environmental data sets that have a spatial lag more amenable to the presence records that are available, although such choices must be appropriate also to the temporal and spatial characteristics of the occurrence data as well (Peterson et al. 2011). By means of these steps, any similarity in environmental dimensions among presence records will be a consequence of niche preferences, rather than an artifactual result of spatial proximity.

Spatial autocorrelation enters the picture much more powerfully in model evaluation (Peterson et al. 2011). Here, the common practice of subsetting presence data at random into calibration and evaluation data sets (e.g., Manel et al. 1999) will frequently fall into two traps: (1) points that are closely positioned (and therefore non-independent) may fall into both calibration and evaluation subsets, and yet are not independent of one another; and (2) evaluation data may include many points that are close to one another (and therefore non-independent), thereby inflating sample sizes artificially. Indeed, at least the second of these problems may plague even the spatial subsetting exercises that have been purported to be superior (e.g., Peterson et al. 2007).

A solution is to subset the evaluation data at random, but subject to the constraint that they are separated by at least the lag distance characteristic of the environmental data set being used. In this way, once again, any similarity among evaluation points, or any shared tendency toward coincidence of the evaluation data with the model prediction, will be the consequence of biological factors, and not simply spatial proximity and consequent non-independence. As mentioned above for model calibration, however, these steps are rather severe, and will frequently cause serious problems of minuscule sample sizes for model evaluation.

Transferring and Extrapolating: MESS and Clamping

A final point that we will explore is that of the perils of transferring models that have been trained within $\mathbf{M} \cap \mathbf{S}$ to broader areas; we can distinguish between model transfer, which is to conditions over which the model was calibrated, and model extrapolation, which is to conditions outside of the range over which the model was calibrated—model extrapolation is, quite generally, perilous and should be avoided if possible. Elith et al. (2011) outlined a means of visualizing areas on maps that are highly different and "out of range" as referred to a reference set of areas (called "MESS"). The implementation that Elith et al. (2011) provided as an option in Maxent calculates MESS surfaces that refer to the occurrence points; we suggest that the more relevant areas for MESS calculations are instead those within $\mathbf{M}$—i.e., the areas that the species has explored and found either suitable or not. Otherwise, MESS calculations will

confound niche differences with environmental extrapolation, and will underestimate environmental similarity. A more detailed treatment of these points is currently in preparation (KU Niche Modeling Group, in prep.).

Discussion

In this paper, we have provided a broad overview of an important set of transitions that we see as necessary and important in ecological niche modeling. That is, with the broad availability of occurrence data, environmental data, and modeling algorithms, we see considerable potential for many users simply to "push the button," and interpret what comes out. This simple approach, unfortunately, will frequently lead users to inaccurate, inappropriate, and incorrect conclusions.

In this contribution, we outline 10 adjustments that must be made to the simple button-pushing. In each case, the change is intended to place the particular analysis in a context of the ecology and biogeography of species' distributions in environmental and geographic spaces, as well as in the context of the realities of sampling of species' occurrences across real-world landscapes. This shift towards development of niche models in appropriate conceptual frameworks has much to offer to the future development of this approach in ecology, biogeography, evolution, and conservation biology.

Beginning with a recent synthesis (Peterson et al. 2011), and continuing with a series of insights both from our own lab group and other lab groups, we are eager to see this emerging set of tools mature into a more synthetic science.

The 10 points treated in this paper are not the *only* 10, but rather are a first cut; returning to our earlier analogy to phylogenetics, the six decades since Hennig's initial insights (Hennig 1950) have not been static in any way … rather, new insights and methodological improvements have continued right up to the present (e.g., Alfaro and Holder 2006). This paper is our attempt to provide a compendium of recent advances and insights in niche modeling that will open doors to a firmer conceptual foundation for these new approaches.

Acknowledgments

Literature Cited

Alfaro ME, Holder MT. 2006. The posterior and the prior in Bayesian phylogenetics. Annual Review of Ecology, Evolution, and Systematics 37:19-42.

Anderson RP. 2003. Real vs. artefactual absences in species distributions: Tests for *Oryzomys albigularis* (Rodentia: Muridae) in Venezuela. Journal of Biogeography 30:591-605.

Anderson RP, Lew D, Peterson AT. 2003. Evaluating predictive models of species' distributions: Criteria for selecting optimal models. Ecological Modelling 162:211-232.

Araújo MB, Thuiller W, Yoccoz NG. 2009. Reopening the climate envelope reveals macroscale associations with climate in European birds. Proceedings of the National Academy of Sciences USA 106:E41-E43.

Austin MP, Nicholls AO, Margules CR. 1990. Measurement of the realized qualitative niche: Environmental niches of five *Eucalyptus* species. Ecological Monographs 60:161-177.

Barve N, Barve V, Jiménez-Valverde A, Lira-Noriega A, Maher SP, Peterson AT, Soberón J, Villalobos F. 2011. The crucial role of the accessible area in ecological niche modeling and species distribution modeling. Ecological Modelling 222:1810-1819.

Beale CM, Lennon JJ, Gimona A. 2008. Opening the climate envelope reveals no macroscale associations with climate in European birds. Proceedings of the National Academy of Sciences USA 105:14908-14912.

Berlocher SH, Swofford DL. 1997. Searching for phylogenetic trees under the frequency parsimony criterion: An approximation using generalized parsimony. Systematic Biology 46:211-215.

Broennimann O, Treier UA, Müller-Schärer H, Thuiller W, Peterson AT, Guisan A. 2007. Evidence of climatic niche shift during biological invasion. Ecology Letters 10:701-709.

Colwell RK, Futuyma D. 1971. On the measurment of niche breadth and overlap. Ecology 52:567-576.

Diniz-Filho JAF, Bini LM, Hawkins BA. 2003. Spatial autocorrelation and red herrings in geographical ecology. Global Ecology and Biogeography 12:53-64.

Dormann CF, McPherson JM, Araújo MB, Bivand R, Bolliger J, Carl G, Davies RG, Hirzel A, Jetz W, Kissling WD, Kühn I, Ohlemüller R, Peres-Neto PR, Reineking B, Schröder B, Schurr FM, Wilson R. 2007. Methods to account for spatial autocorrelation in the analysis of species distributional data: A review. Ecography 30:609-628.

Elith J, Graham C, Anderson R, Dudik M, Ferrier S, Guisan A, Hijmans R, Huettmann F, Leathwick J, Lehmann A, Li J, Lohmann L, Loisell B, Manion G, Moritz C, Nakamura M, Nakazawa Y, Overton J, Peterson A, Phillips S, Richardson K, Schachetti-Pereira R, Schapire E, Soberon J, Williams S, Wisz M, Zimmerman N. 2006. Novel methods improve prediction of species' distributions from occurrence data. Ecography 29:129-151.

Elith J, Phillips SJ, Hastie T, Dudík M, Chee YE, Yates CJ. 2011. A statistical explanation of MaxEnt for ecologists. Diversity and Distributions 17:43-57.

Felsenstein J. 1986. PHYLIP (Phylogeny Inference Package) Version 3.57c. Seattle: University of Washington.

Fielding AH, Bell JF. 1997. A review of methods for the assessment of prediction errors in conservation presence/absence models. Environmental Conservation 24:38-49.

Fitzpatrick MC, Weltzin JF, Sanders NJ, Dunn RR. 2007. The biogeography of prediction error: Why does the introduced range of the fire ant over-predict its native range? Global Ecology and Biogeography 16:24-33.

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

Godsoe W. 2010. I can't define the niche but I know it when I see it: A formal link between statistical theory and the ecological niche. Oikos 119:53-60.

Grinnell J. 1917. Field tests of theories concerning distributional control. American Naturalist 51:115-128.

Grinnell J. 1924. Geography and evolution. Ecology 5:225-229.

Hennig W. 1950. Grundzüge einer Theorie der Phylogenetischen Systematik. Berlin: Deutscher Zentralverlag.

Huntley B, Collingham YC, Willis SG, Green RE. 2008. Potential impacts of climatic change on European breeding birds. PLoS ONE 3:e1439.

Hutchinson GE. 1957. Concluding remarks. Cold Spring Harbor Symposia on
    Quantitative Biology 22:415-427.

Hutchinson GE. 1978. An Introduction to Population Ecology. New Haven: Yale
    University Press.

Jiménez-Valverde A, Barve N, Lira-Noriega A, Maher SP, Nakazawa Y, Papeş M,
    Soberón J, Sukumaran J, Peterson AT. 2010. Dominant climate influences on
    North American bird distributions. Global Ecology and Biogeography 20:114-
    118.

Jiménez-Valverde A, Nakazawa Y, Lira-Noriega A, Peterson AT. 2009. Environmental
    correlation structure and ecological niche model projections. Biodiversity
    Informatics 6:28-35.

Lobo JM, Jiménez-Valverde A, Real R. 2008. AUC: A misleading measure of the
    performance of predictive distribution models. Global Ecology and
    Biogeography 17:145-151.

Maher SP, Ellis C, Gage KL, Enscore RE, Peterson AT. 2010. Range-wide
    determinants of plague distribution in North America. American Journal of
    Tropical Medicine and Hygiene 83:736-742.

Manel S, Dias JM, Buckton ST, Ormerod SJ. 1999. Alternative methods for predicting
    species distribution: An illustration with Himalayan river birds. Journal of
    Applied Ecology 36:734-747.

Medley KA. 2010. Niche shifts during the global invasion of the Asian tiger mosquito,
    Aedes albopictus Skuse (Culicidae), revealed by reciprocal distribution
    models. Global Ecology and Biogeography 19:122-133.

Nix HA. 1986. A biogeographic analysis of Australian elapid snakes. In: Longmore R,
    editor. Atlas of elapid snakes of Australia. Australian Government Publishing
    Service, Canberra. Pages 4-15

Pearson R, Raxworthy C, Nakamura M, Peterson A. 2007. Predicting species
    distributions from small numbers of occurrence records: A test case using
    cryptic geckos in Madagascar. Journal of Biogeography 34:102 - 117.

Peterson AT. 2011. Ecological niche conservatism: A time-structured review of
    evidence. Journal of Biogeography 38:817-827.

Peterson AT, Nakazawa Y. 2008. Environmental data sets matter in ecological niche
    modeling: An example with *Solenopsis invicta* and *Solenopsis richteri*. Global
    Ecology and Biogeography 17:135-144.

Peterson AT, Papeş M, Eaton M. 2007. Transferability and model evaluation in
    ecological niche modeling: A comparison of GARP and Maxent. Ecography
    30:550-560.

Peterson AT, Papeş M, Soberón J. 2008. Rethinking receiver operating characteristic
    analysis applications in ecological niche modelling. Ecological Modelling
    213:63-72.

Peterson AT, Soberón J, Pearson RG, Anderson RP, Martínez-Meyer E, Nakamura M,
    Araújo MB. 2011. Ecological Niches and Geographic Distributions. Princeton:
    Princeton University Press.

Phillips S, Anderson R, Schapire R. 2006. Maximum entropy modeling of species
    geographic distributions. Ecological Modelling 190:231-259.

Phillips S, Dudik M, Elith J, Graham C, Lehman A, Leathwick J, Ferrier S. 2009. Sample selection bias and presence-only distribution models: implications for background and pseudo-absence data. Ecological Applications 19:181-197.

Pulliam HR. 2000. On the relationship between niche and distribution. Ecology Letters 3:349-361.

Saupe E, Barve V, Myers C, Soberón J, Barve N, Hensz C, Peterson AT, Owens HL, Lira-Noriega A. In review. Variation in niche and distribution model performance: The need for *a priori* assessment of key causal factors. Ecological Modelling.

Segurado P, Araújo MB, Kunin WE. 2006. Consequences of spatial autocorrelation for niche-based models. Journal of Applied Ecology 43:433-444.

Soberón J. 2007. Grinnellian and Eltonian niches and geographic distributions of species. Ecology Letters 10:1115-1123.

Soberón J. 2010. Niche and area of distribution modeling: A population ecology perspective. Ecography 33:159-167.

Soberón J, Nakamura M. 2009. Niches and distributional areas: Concepts, methods, and assumptions. Proceedings of the National Academy of Sciences USA 106:19644-19650.

Soberón J, Peterson AT. 2005. Interpretation of models of fundamental ecological niches and species' distributional areas. Biodiversity Informatics 2:1-10.

Soberón J, Peterson AT. 2011. Ecological niche shifts and environmental space anisotropy: A cautionary note. Revista Mexicana de Biodiversidad 82:1348-1353.

Stockwell DRB, Noble IR. 1992. Induction of sets of rules from animal distribution data: A robust and informative method of analysis. Mathematics and Computers in Simulation 33:385-390.

Ward G, Hastie T, Barry S, Elith J, Leathwick JR. 2009. Presence-only data and the EM algorithm. Biometrics 65:554-563.

Warren DL, Glor RE, Turelli M. 2008. Environmental niche equivalency versus conservatism: Quantitative approaches to niche evolution. Evolution 62:2868-2883.

Wiley EO. 1981. The Theory and Practice of Phylogenetic Systematics. 1st edition. New York: John Wiley & Sons.
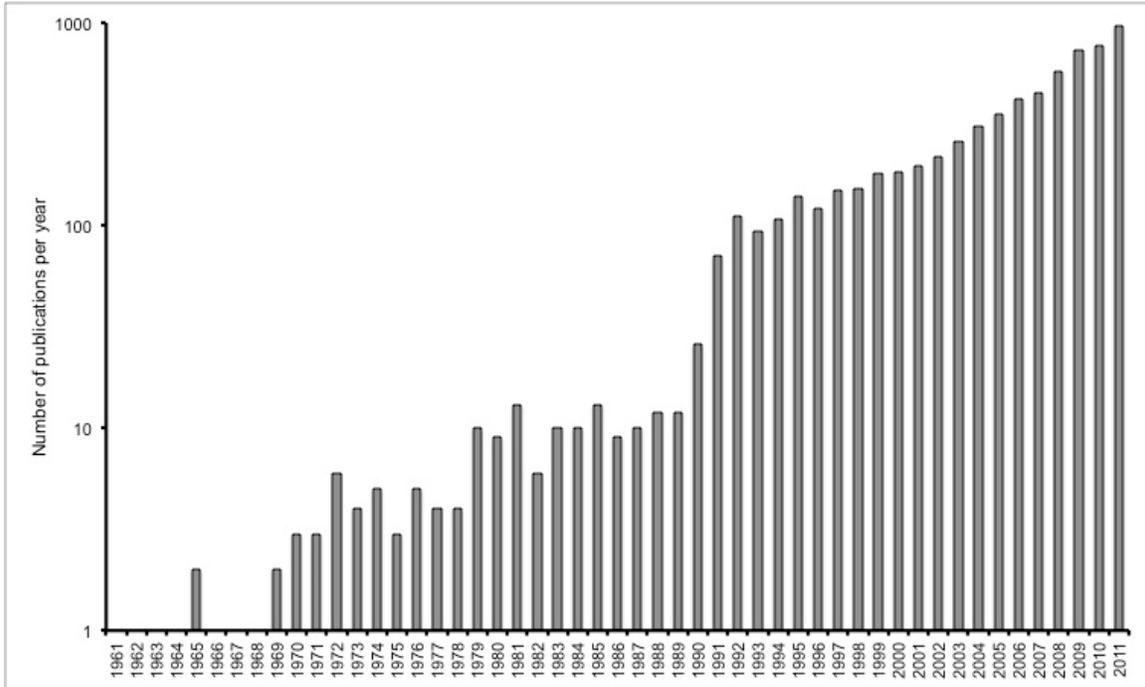
Figure 1. Summary of representation in the literature of concepts related to ecological niche modeling and species distribution modeling, based on a Web of Science search for "ecological niche" or "species distribution" under title or topic, over the period 1950-2011. Note that the number of publications is shown on a log$_{10}$ scale.
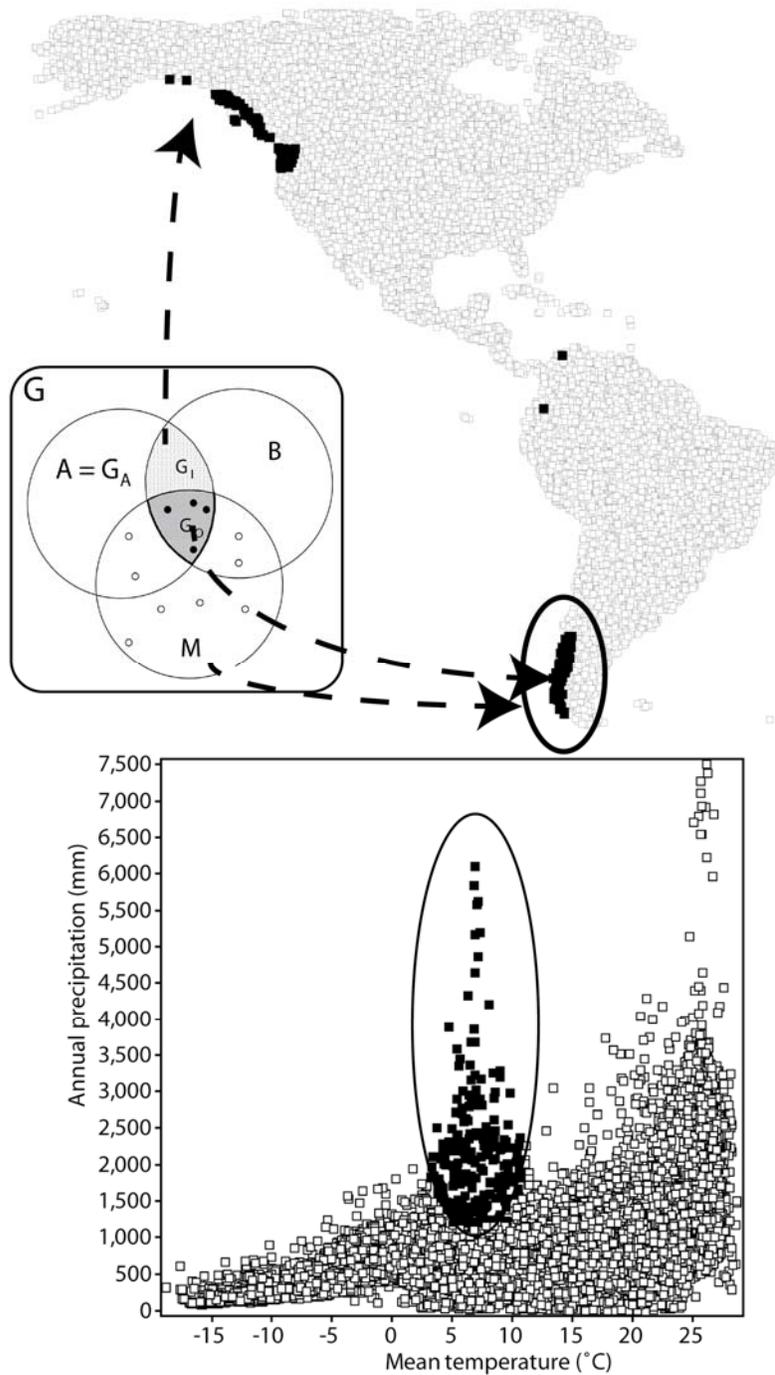
Figure 2. Summary of geographic (**G**: Venn diagram and map) and environmental spaces (**E**: scatterplot at bottom). Shown are all combinations of annual precipitation and temperature across the Americas, in which a fundamental niche has been identified (ellipse in E-space); note that the full extent of this fundamental niche is not represented anywhere in the Americas. In the map, the spatial footprint of this niche can be seen, but the species is limited to a particular **M** (see ellipse in map), and thus does not actually inhabit the remainder of its invadable area, **G**$_I$.