ECOGRAPHY

Software note

ENMTools 1.0: an R package for comparative ecological biogeography

Dan L. Warren, Nicholas J. Matzke, Marcel Cardillo, John B. Baumgartner, Linda J. Beaumont, Michael Turelli, Richard E. Glor, Nicholas A. Huron, Marianna Simões, Teresa L. Iglesias, Julien C. Piquet and Russell Dinnage

D. L. Warren (https://orcid.org/0000-0002-8747-2451) ✉ (dan.l.warren@gmail.com), Biodiversity and Biocomplexity Unit, Okinawa Inst. of Science and Technology, Okinawa, Japan, and Senckenberg Biodiversity and Climate Research Centre, Frankfurt, Germany. – N. J. Matzke, School of Biological Sciences, Univ. of Auckland, Auckland, New Zealand. – M. Cardillo and R. Dinnage (https://orcid.org/0000-0003-0846-2819), Macroevolution & Macroecology Group, Research School of Biology, Australian National Univ., Canberra, Australia. RD also at: Inst. for Applied Ecology, Univ. of Canberra, Canberra, ACT, Australia, and Dept of Biological Sciences, Florida International Univ., Miami, FL, USA. – J. Baumgartner (https://orcid.org/0000-0002-8898-0300), Centre of Excellence for Biosecurity Risk Analysis (CEBRA), Univ. of Melbourne, Melbourne, Australia. – JB and L. J. Beaumont, Dept of Biological Sciences, Macquarie Univ., Australia. – M. Turelli, Dept of Evolution and Ecology, Univ. of California, Davis, CA, USA. – R. E. Glor (https://orcid.org/0000-0002-3359-1631), Dept of Ecology and Evolutionary Biology, Univ. of Kansas, Lawrence, KS, USA. – N. A. Huron (https://orcid.org/0000-0001-6835-1390), Integrative Ecology Lab, Center for Biodiversity, Dept of Biology, Temple Univ., Philadelphia, PA, USA. – M. Simões, Senckenberg Deutsches Entomologisches Inst. (SDEI), Müncheberg, Germany. – T. L. Iglesias, Animal Resources Section, Okinawa Inst. of Science and Technology, Okinawa, Japan. – J. C. Piquet, Island Ecology and Evolution Research Group, Inst. de Productos Naturales y Agrobiología (IPNA-CSIC), La Laguna, Tenerife, Canary Islands, Spain.

The ENMTools software package was introduced in 2008 as a platform for making measurements on environmental niche models (ENMs, frequently referred to as species distribution models or SDMs), and for using those measurements in the context of newly developed Monte Carlo tests to evaluate hypotheses regarding niche evolution. Additional functionality was later added for model selection and simulation from ENMs, and the software package has been quite widely used. ENMTools was initially implemented as a Perl script, which was also compiled into an executable file for various platforms. However, the package had a number of significant limitations; it was only designed to fit models using Maxent, it relied on a specific Perl distribution to function, and its internal structure made it difficult to maintain and expand. Subsequently, the R programming language became the platform of choice for most ENM studies, making ENMTools less usable for many practitioners. Here we introduce a new R version of ENMTools that implements much of the functionality of its predecessor as well as numerous additions that simplify the construction, comparison and evaluation of niche models. These additions include new metrics for model fit, methods of measuring ENM overlap, and methods for testing evolutionary hypotheses. The new version of ENMTools is also designed to work within the expanding universe of R tools for ecological biogeography, and as such includes greatly simplified interfaces for analyses from several other R packages.

Keywords: biogeography, ecology, ENM, ENMTools, SDM

Introduction

Here we describe the new ENMTools R package ver. 1.0, which reproduces the core functionality of the original ENMTools (Warren et al. 2010), but adds significant new functionality. It interfaces with many more modeling algorithms, performs several new statistical tests, and provides a unified interface for many different model types with automatic data processing, model construction, model evaluation and visualization. It provides new data structures designed specifically to simplify the addition of new modeling algorithms, and provides streamlined interfaces for some existing R packages including ecospat (Di Cola et al. 2017), vip (Greenwell and Boehmke 2020), CalibratR (Schwarz and Heider 2019) and phyloclim (Heibl et al. 2018).

Functionality

Species and clade objects

One of ENMTools' primary functions is to test hypotheses about similarity between environmental niche estimates for groups of organisms. To simplify this process, we have created data structures that contain the information needed to build an ENM for a species. The first of these structures is an *enmtools.species* object, which contains a species name, species occurrence data and optional range data and background points. They are created by using the *enmtools.species* function, and formatting is checked using the *check.species* function. An *enmtools.clade* object contains a list of *enmtools.species* objects and a phylogeny imported with the R package ape (Paradis and Schliep 2018). These objects are created using the *enmtools.clade* function, and can be checked for formatting by using the *check.clade* function.

Building and refining niche and distribution models

The use of *enmtools.species* objects permits ENMTools to implement a uniform interface for many of the most popular ENM algorithms, and allows for detailed informative visualizations and model validation tests with minimal additional user input. For example, the construction of a generalized additive model (GAM) for species *myspecies* using environmental rasters stored in a stack called *env* requires only one command:

```
enmtools.gam(myspecies, env)
```

This single command extracts and formats the environmental data, constructs a GAM, projects it across the study area, and evaluates model fit to training data. It returns an *enmtools.model* object that contains a suitability raster, plots of response curves for environmental variables, and objects of class *evaluate* from the R package dismo (Hijmans et al. 2017). These *evaluate* objects contain a number of widely used metrics for measuring discrimination accuracy including AUC, kappa and summary statistics from a confusion matrix that can be used to calculate many other discrimination statistics (Hijmans et al. 2017). Providing additional arguments activates extended functionality including subsampling data sets for model evaluation (random or spatially structured), Monte Carlo tests to construct null distributions for model significance as in Raes and ter Steege (2007), and manually providing model formulas for algorithms that accept them. Example outputs are shown in Fig. 1. Converting this analysis to use another algorithm is simply a matter of changing the function name; the remaining syntax is identical. Currently implemented algorithms include Bioclim, Domain, maximum entropy (using dismo's interface for the standalone Maxent software), generalized additive models, generalized linear models, Poisson point process models and random forests via both randomForest and ranger R packages (Nix and Busby 1986, Carpenter et al. 1993, Wood 2001, Liaw et al. 2002, Phillips et al. 2006, Renner and Renner 2015, Wright and Ziegler 2015, Hijmans et al. 2017). Although this implementation allows easy access to all of these algorithms using default settings, all user-facing options from the underlying modeling algorithms may also be included as arguments to the ENMTools modeling functions, allowing full control over the modeling process while minimizing many of the repetitive steps that typically accompany ENM/SDM construction in R. In addition, we note that models for multiple species can be generated by using the *lapply* function on a list of *enmtools.model* objects (for details <https://enmtools.blogspot.com/2019/01/fun-fact-you-can-run-whole-bunch-of.html>).

Measures of fit offered by ENMTools include traditional measures of discrimination accuracy on training and test occurrences in both geographic space. ENMTools also introduces methods for measuring discrimination accuracy in an N-dimensional environment space. This method evaluates model performance using latin hypercube sampling of all environmental conditions within the minima and maxima for each predictor in the training region, in a manner similar to the measures of niche breadth and overlap developed in Warren et al. (2019). ENMTools also automatically generates plots of response curves, which show the marginal estimate of habitat suitability in the context of the frequency of occurrence of different environmental conditions in the occurrence data and background points. Further visualization of model predictions can be done in 2D environmental space using any pair of layers from the predictor raster stack (Fig. 1D–E).

ENMTools provides additional support for model refinement by implementing a simplified interface for variable selection using the vip package (Greenwell and Boehmke 2020) and visualization of correlations between predictors (Fig. 2). Two visualizations are available: a cluster plot, which does multidimensional scaling (MDS) on the correlation matrix and then plots the results in a two dimensional space so that more correlated predictors are closer to each other (Fig. 2B), and a heatmap, where colors represent the Pearson correlation coefficient between predictors (Fig. 2C). ENMTools also provides methods to measure model calibration using continuous Boyce
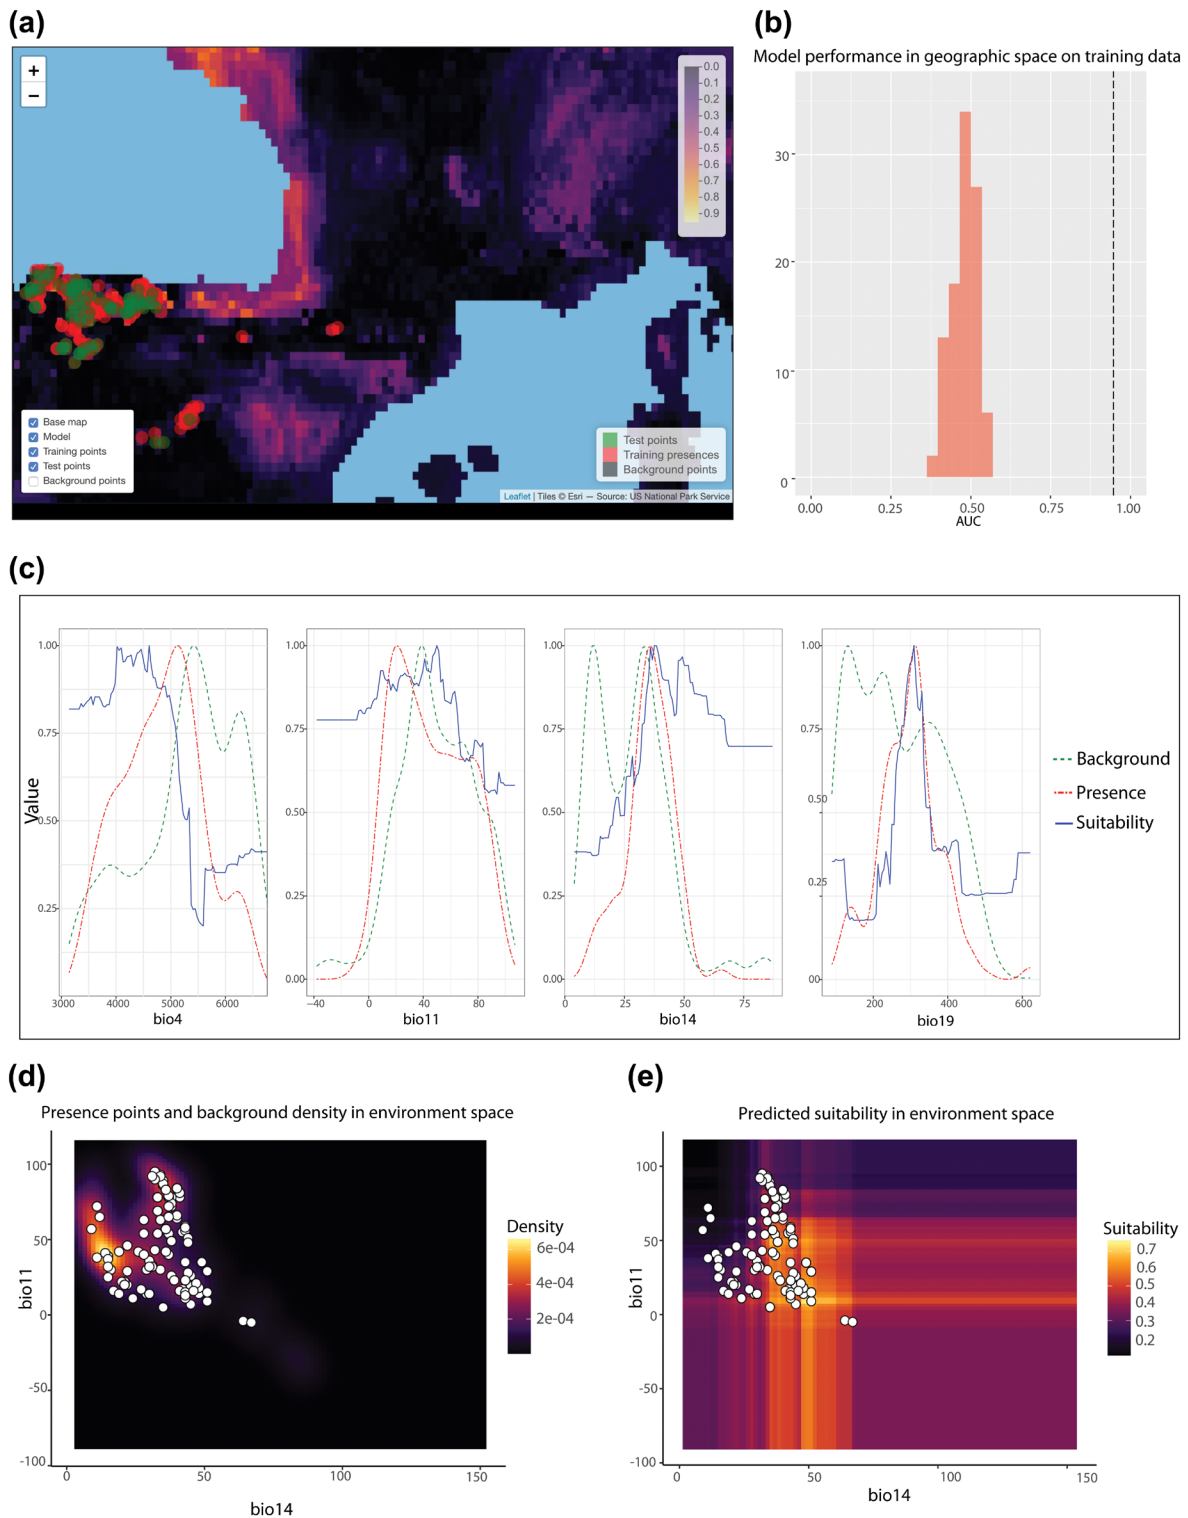
Figure 1. Sample model outputs from an *enmtools.model* object. Panel (a) shows a screen capture of an interactive plot that can be called on any *enmtools.model* object using the interactive.plot function. These visualizations can be panned and zoomed, with various layers being turned on and off to facilitate model and data exploration. Panel (b) shows the partial output of a randomization test based on Raes and ter Steege (2007), indicating that the predictive power of the model far exceeds that expected under the null hypothesis. Response curves (panel c) are presented in the context of the density distribution of each environmental variable in the species' occurrences as well as in the background data. Panels (d) and (e) show bivariate heat maps depicting the same information. Occurrence data points are represented by white points.
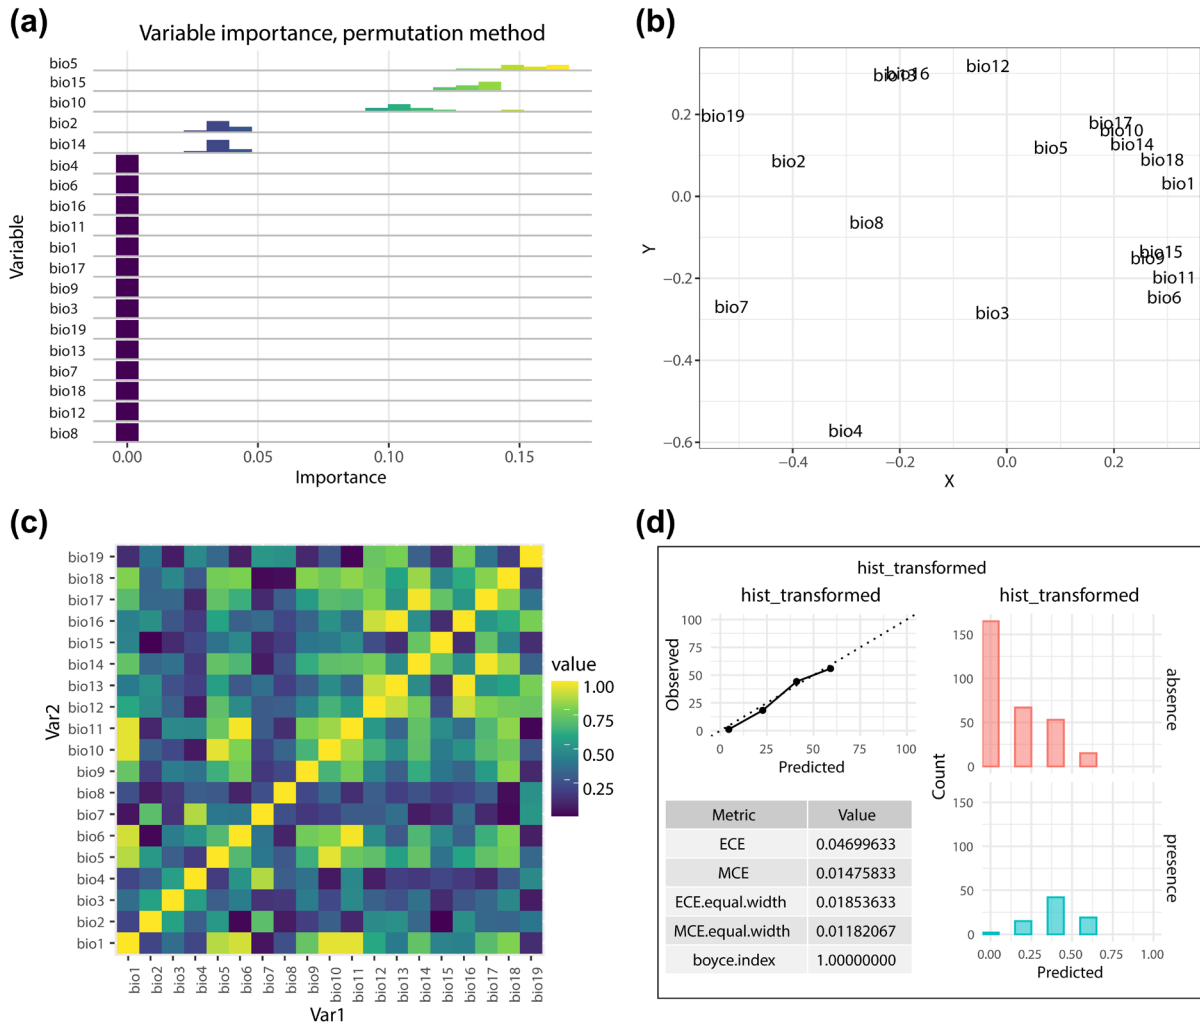
Figure 2. Variable selection and model recalibration in ENMTools. ENMTools implements a simplified interface for conducting tests of variable importance using several different methods, and plotting those results (panel a) using the enmtools.vip function. In order to explore multicollinearity of predictors, the raster.cor.plot function also provides visualizations of correlation matrices using MDS scaling to group correlated variables as well as pairwise heatmaps (panels b–c). Panel (d) shows the output of the enmtools.calibrate function, which allows users to measure model calibration and recalibrate models using a wide array of different methods.

index (Hirzel et al. 2006) as well as expected and maximum calibration error, and allows recalibration via the CalibratR, caret and ecospat packages, along with automated conversion of recalibrated models to suitability rasters (Kuhn 2008, Di Cola et al. 2017, Schwarz and Heider 2019). We note here that expected and maximum calibration error are highly sensitive to prevalence, and should be used with caution.

Metrics of ENM breadth and overlap

The initial version of ENMTools (Warren et al. 2010) introduced metrics of niche breadth and niche overlap on ENMs. Breadth metrics (B1 and B2, Levins 1968) measure the uniformity of the geographic distribution of suitability scores for a model, while niche overlap metrics (*D*, Schoener 1968; *I*, Warren et al. 2008) measure the similarity of the distribution of suitable habitat for a pair of models. The new ENMTools

adds the use of the rank correlation coefficient *rho* (Spearman 1904) to measure niche overlap from ENMs. This addition is intended to supplement, rather than replace, the existing metrics and the choice of metric is intended to be driven by the goals of a given study. *D* and *I* will emphasize the potential for a pair of species (or lineages) to interact in a given geographic space, while *rho* will emphasize differences in the estimated physiological response to the predictor variables. Statistically, this is due to the fact *D* and *I* measure the similarity in the numerical value of suitability in a given set of conditions, while *rho* emphasizes the direction of response to the environmental gradient (for details <https://enmtools.blogspot.com/2018/10/why-add-correlations-for-suitability.html>). The new ENMTools further extends these metrics by allowing users to measure them in the N-dimensional environment space defined by the predictors used in model construction (Warren et al. 2019).

Hypothesis testing for species pairs

The primary purpose of the ENMTools R package is, as with previous versions, to test hypotheses regarding the evolution of differences between species, populations and other groups of organisms using ENMs. Many of these tests are conducted by comparing the overlap between empirical species distribution models to a null distribution. These distributions are estimated using Monte Carlo methods, in which species occurrences are chosen randomly in a manner consistent with the chosen null hypothesis. While older versions only implemented these tests using Maxent models, the new R version allows users to conduct tests using any method that has an accompanying ENMTools modeling function (above). A brief description of each test follows.

Identity test
The identity (also called equivalency) test was first developed in Warren et al. (2008) to address whether a pair of closely related species are effectively identical in their realized environmental distributions. In this test, models are built using the empirical data for each species, and overlap is measured between them. The null distribution is constructed by repeatedly randomizing the group identity (e.g. species or lineage) of each data point, keeping the sample size for each group consistent with the empirical data. Models are built for each replicate, and overlap between them is measured. The distribution of overlaps constructed in this way reflects the expected similarity between groups if their environmental distributions represent repeated draws from the same underlying distribution. By comparing the empirical value to this distribution (Fig. 3), we can reject (or fail to reject) this hypothesis at a given level of confidence (e.g. $p < 0.05$).

Background test
In many cases, groups that are partially or fully allopatric will have different suites of environments available to them. As such they will often be found to have different environmental distributions for reasons that may not reflect any underlying divergence in ecological tolerances or preferences. To correct for the differential availability of habitat for allopatric groups, Warren et al. (2008) developed the 'background test', also known as the 'similarity test'. This method is broadly similar to the identity test, but the null distribution is constructed by building ENMs using randomly chosen points from the broad region where each group lives, rather than points where the species has actually been detected. The ENMs thus constructed represent the distributions of available habitat in each geographic region, rather than the environmental associations of either group. The null distribution therefore reflects the expected similarity between groups based on the available environments, and rejecting the null hypothesis therefore implies that species are more (or less) similar in their environmental distributions than expected given their respective ranges. As originally implemented, this test was conducted by randomizing occurrences for each group one at a time; the null distribution was constructed by comparing randomized

points from one species' range to the empirical points for the other species' range. The ENMTools R package allows users to choose between this method (an 'asymmetric' background test) and one in which both groups' points are chosen randomly in each replicate (a 'symmetric' background test). This latter approach mirrors one implemented for kernel density niche estimates in ecospat (Di Cola et al. 2017).

Ecospat tests
Broennimann et al. (2012) developed methods for constructing kernel density estimates of species' environmental distributions and the distributions of available environments to quantify niche overlap in a two-dimensional environment space, and used these to conduct hypothesis tests similar to those implemented in ENMTools (Warren et al. 2010). The R version of ENMTools provides a greatly simplified user interface for ecospat hypothesis tests via the use of *enmtools. species* objects. This includes automatically conducting principal components analysis when more than two predictor layers are provided, in order to reduce the predictors to a two dimensional environment space. The new ENMTools also implements identity and background tests in N-dimensional environment space without PCA using the new environment space overlap metrics introduced in Warren et al. (2019).

Rangebreak tests
Glor and Warren (2011) introduced Monte Carlo tests using ENMs to explore the environmental and historical significance of biogeographic boundaries. In these tests, the overlap between the ENMs for two allopatric groups is compared to a null distribution that is constructed by randomizing the location and shape of the biogeographic boundary that separates them. The null distribution in this case represents the expected overlap between a pair of groups (typically species) that are allopatric for reasons uncorrelated with the distribution of environmental variables in the study region. Rejection of this null supports the alternative hypothesis that the location of the boundary between the species' distributions in the empirical data partitions the two sets of occurrences into environmental distributions that are unusually disjunct.

MOSES tests
Finally, ENMTools provides basic functionality for a new variety of test that goes by the acronym MOSES (MOdel Selection for Ecological Similarity). This method represents a parametric approach to the question of similarity between species. Using the Akaike information criterion (Burnham and Anderson 2002), MOSES treats the question of ecological similarity as a model-partitioning problem. By comparing the combined AIC of models built for two lineages separately to a model built for the combined presence/background data for both lineages, MOSES effectively determines whether the two lineages are sufficiently different in their environmental associations to justify the additional parameters needed to estimate separate models for each of them. This method is currently in development and is intended only for experimental use; simulation studies to test its efficacy are underway.
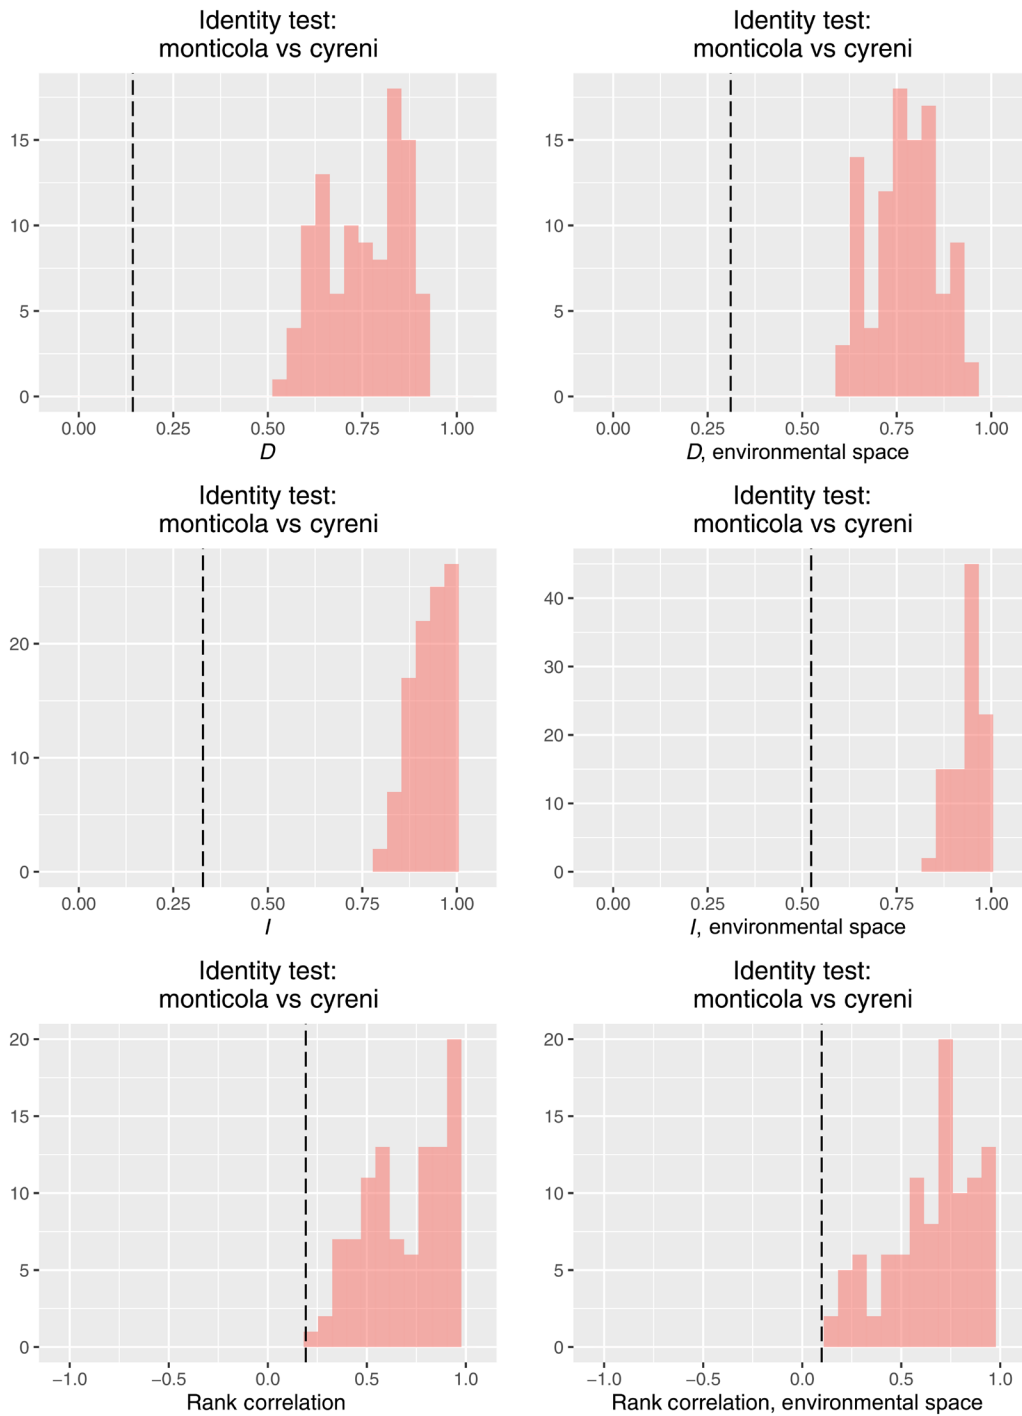
Figure 3. Sample output for identity test for all three implemented niche overlap metrics. Histograms represent the distribution of overlaps for each metric from the null distribution, while the dashed vertical line represents the overlap between the models built using the empirical data.

Hypothesis testing for clades

Measurements of similarity between species (e.g. range overlap) are difficult to analyze in the context of traditional phylogenetic comparative methods, as measurements are made across nodes of the phylogeny (i.e. comparisons between two or more species) rather than at the tips. Age-overlap correlation (AOC) is a generalization of age-range correlation methods (ARC), which allow the analysis of overlap- and distance-style metrics in a phylogenetic context (Barraclough et al. 2000, Fitzpatrick and Turelli 2006, Grossenbacher and Whittall 2011, Cardillo and Warren 2016). In an AOC analysis, the average overlap

Age–overlap correlation from Monte Carlo test
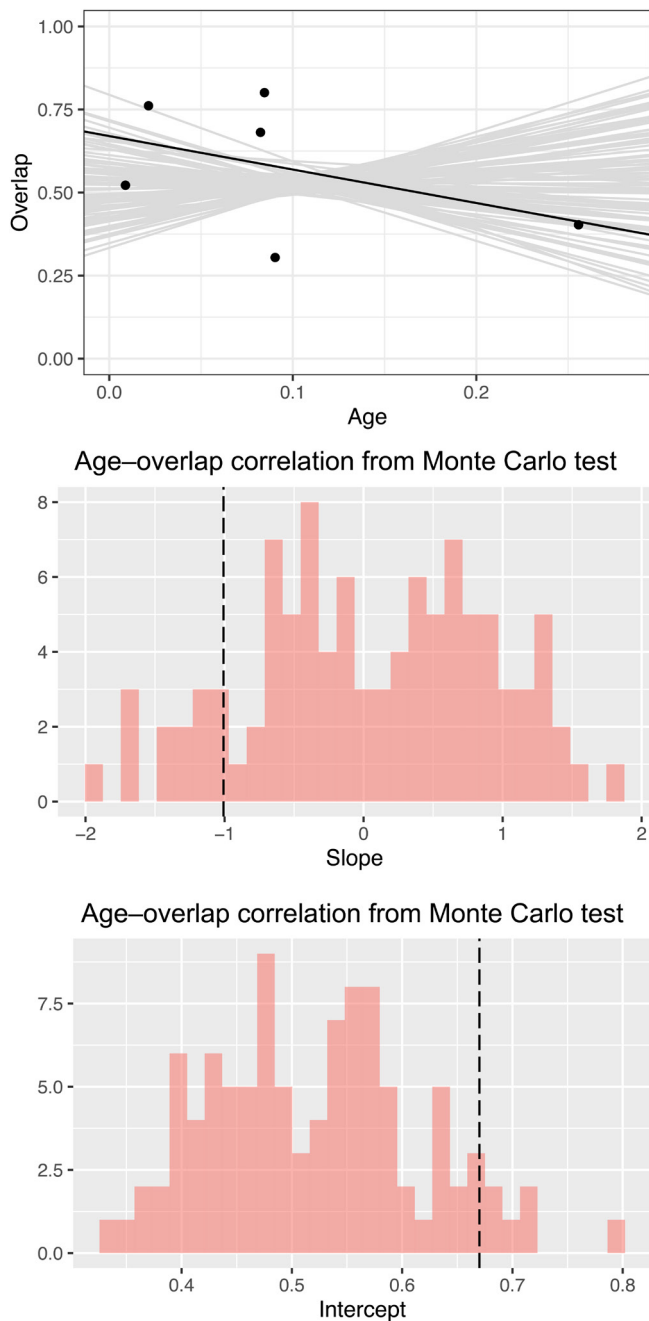

Age–overlap correlation from Monte Carlo test

Figure 4. Sample output for age-overlap correlation test. The top panel shows the distribution of topologically averaged overlaps in the empirical data (black line and points), as well as the distribution of linear regressions from the Monte Carlo replicates (grey lines). The middle and bottom panels show the distribution of slopes and intercepts separately, with pink histograms representing the values from Monte Carlo replicates and vertical dashed lines representing the values from the empirical data.

across each internal node is calculated using an algorithm that accounts for tree topology, and this average overlap is regressed against node age. As overlap metrics are often bounded [0,1] and any given species' range will be used to calculate average overlaps across multiple internal nodes, residuals from these

regressions are not expected to be independent or normally distributed. Statistical significance is therefore estimated by using a Monte Carlo test. Each replicate randomizes the identity of the tips of the phylogeny, recalculates overlap at each internal node, and performs a linear regression. The slope and intercept from the regression on empirical data are compared to the null distribution of slopes and intercepts from the Monte Carlo replicates, allowing users to test 1) whether there is more or less overlap between recently diverged species than expected and 2) whether overlap between tips increases or decreases over evolutionary time.

ENMTools streamlines age-overlap correlation analyses, allowing users to simply use the function *enmtools.aoc* on an *enmtools.clade* object (Fig. 4). For range overlap, point pattern similarity and ENM overlap in both geographic and environment space, ENMTools can automatically calculate overlaps, including automated model construction for measurements that require ENMs. Alternatively, users can supply their own overlap matrix containing any measurement of ecological similarity.

Training resources

ENMTools contains an extensive vignette introducing much of the core functionality of the package. Additionally, a set of video introductions to many of the concepts and functions implemented in ENMTools are available at <https://www.youtube.com/playlist?list=PLw-jIvFxpgIBgNzIbfh52-TQyTEbcSvSH> and in-depth training scripts are available at <https://github.com/danlwarren/ENMTools-Tutorials>. These materials are continuously being updated and new content added as new functionality is implemented.

Conclusions

This new version of ENMTools provides a toolkit of both novel and now-standard methods that will be accessible to end users at all skill levels, thereby contributing to further development of comparative ecological biogeography. This release also signals the coming end of support for older versions of ENMTools (Warren et al. 2010), which have become increasingly difficult to maintain and troubleshoot. Further, we hope that the deployment of the ENMTools R package on GitHub helps to encourage community contributions to this toolkit, and that the species- and clade-oriented approach developed here serves to encourage further efforts to make sophisticated biogeographic analyses more widely accessible.

To cite ENMTools 1.0 or acknowledge its use, cite this Software note as follows, substituting the version of the application that you used for 'version 0': Warren, D. L. et al. 2020. ENMTools 1.0: an R package for comparative ecological biogeography. – Ecography 43: 000–000 (ver. 0).

Author contributions

Julien Piquet: Conceptualization (supporting); Software (supporting). **Dan Warren**: Conceptualization (lead); Formal analysis (equal); Methodology (equal); Project administration (equal); Software (lead); Writing – original draft (lead). **Nicholas Matzke**: Conceptualization (supporting); Software (supporting). **Marcel Cardillo**: Conceptualization (equal); Writing – review and editing (equal). **John Baumgartner**: Conceptualization (equal); Methodology (equal); Software (supporting). **Linda Beaumont**: Conceptualization (equal); Methodology (equal). **Michael Turelli**: Conceptualization (equal); Methodology (equal); Writing – original draft (equal); Writing – review and editing (equal). **Richard Glor**: Conceptualization (equal); Methodology (equal). **Nicholas Huron**: Conceptualization (supporting); Software (supporting). **Marianna Simoes**: Conceptualization (supporting); Methodology (supporting); Resources (supporting). **Teresa Iglesias**: Conceptualization (supporting); Methodology (supporting); Writing – original draft (supporting). **Russell Dinnage**: Conceptualization (equal); Methodology (equal); Project administration (equal); Software (equal); Writing – original draft (supporting).

References

Barraclough, T. G. et al. 2000. Detecting the geographical pattern of speciation from species-level phylogenies. – Am. Nat. 155: 419.

Broennimann, O. et al. 2012. Measuring ecological niche overlap from occurrence and spatial environmental data. – Global Ecol. Biogeogr. 21: 481–497.

Burnham, K. P. and Anderson, D. R. 2002. Model selection and multimodel inference: a practical information–theoretic approach. – Springer Science & Business Media.

Cardillo, M. and Warren, D. L. 2016. Analysing patterns of spatial and niche overlap among species at multiple resolutions. – Global Ecol. Biogeogr. 25: 951–963.

Carpenter, G. et al. 1993. DOMAIN: a flexible modelling procedure for mapping potential distributions of plants and animals. – Biodivers. Conserv. 2: 667–680.

Di Cola, V. et al. 2017. ecospat: an R package to support spatial analyses and modeling of species niches and distributions. – Ecography 40: 774–787.

Fitzpatrick, B. M. and Turelli, M. 2006. The geography of mammalian speciation: mixed signals from phylogenies and range maps. – Evolution 60: 601–615.

Glor, R. E. and Warren, D. 2011. Testing ecological explanations for biogeographic boundaries. – Evolution 65: 673–683.

Greenwell, B. and Boehmke, B. C. 2020. vip: variable importance plots – an introduction to the vip package. – R J. 12: 343–366

Grossenbacher, D. L. and Whittall, J. B. 2011. Increased floral divergence in sympatric monkeyflowers. – Evolution 65: 2712–2718.

Heibl, C. et al. 2018. Phyloclim: integrating phylogenetics and climatic niche modelling. – R package ver. 0.9-5, <https://CRAN.R-project.org/package=phyloclim>.

Hijmans, R. J. et al. 2017. Package 'dismo'. – R package ver. 1.1-4, <https://CRAN.R-project.org/package=dismo>.

Hirzel, A. H. et al. 2006. Evaluating the ability of habitat suitability models to predict species presences. – Ecol. Model. 199: 142–152.

Kuhn, M. 2008. Building predictive models in R using the caret package. – J. Stat. Softw. 28: 1–26.

Levins, R. 1968. Evolution in changing environments: some theoretical explorations. – Princeton Univ. Press.

Liaw, A. et al. 2002. Classification and regression by randomForest. – R News 2: 18–22.

Nix, H. A. and Busby, J. 1986. BIOCLIM, a bioclimatic analysis and prediction system. – Annual Report of CSIRO Division of Water and Land Resources, Canberra.

Paradis, E. and Schliep, K. 2018. ape 5.0: an environment for modern phylogenetics and evolutionary analyses in R. – Bioinformatics 35: 526–528.

Phillips, S. J. et al. 2006. Maximum entropy modeling of species geographic distributions. – Ecol. Model. 190: 231–259.

Raes, N. and ter Steege, H. 2007. A null-model for significance testing of presence-only species distribution models. – Ecography 30: 727–736.

Renner, I. and Renner, M. I. 2015. Package 'ppmlasso'. – R package ver. 1.1, <https://CRAN.R-project.org/package=ppmlasso>.

Schoener, T. W. 1968. The *Anolis* lizards of Bimini: resource partitioning in a complex fauna. – Ecology 49: 704–726.

Schwarz, J. and Heider, D. 2019. GUESS: projecting machine learning scores to well-calibrated probability estimates for clinical decision making. – Bioinformatics 35: 2458–2465.

Spearman, C. 1904. The proof and measurement of association between two things. – Am. J. Psychol. 15: 72.

Warren, D. L. et al. 2008. Environmental niche equivalency versus conservatism: quantitative approaches to niche evolution. – Evolution 62: 2868–2883.

Warren, D. L. et al. 2010. ENMTools: a toolbox for comparative studies of environmental niche models. – Ecography 33: 607–611.

Warren, D. L. et al. 2019. New methods for measuring ENM breadth and overlap in environmental space. – Ecography 42: 444–446.

Wood, S. N. 2001. mgcv: GAMs and generalized ridge regression for R. – R News 1: 20–25.

Wright, M. N. and Ziegler, A. 2015. ranger: a fast implementation of random forests for high dimensional data in C++ and R. – arXiv: 1508.04409 [stat.ML].