ECOGRAPHY

Software note

Fauxcurrence: simulating multi-species occurrences for null models in species distribution modelling and biogeography

Owen G. Osborne, Henry G. Fell, Hannah Atkins, Jan van Tol, Daniel Phillips, Leonel Herrera-Alsina, Poppy Mynard, Greta Bocedi, Cécile Gubry-Rangin, Lesley T. Lancaster, Simon Creer, Meis Nangoy, Fahri Fahri, Pungki Lupiyaningdyah, I M. Sudiana, Berry Juliandi, Justin M. J. Travis, Alexander S. T. Papadopulos and Adam C. Algar

O. G. Osborne (https://orcid.org/0000-0002-1213-1169) ✉ (owengosborne@gmail.com), D. Phillips, S. Creer and A. S. T. Papadopulos, School of Natural Sciences, Bangor Univ., Environment Centre Wales, Bangor, UK. – H. G. Fell and A. C. Algar (https://orcid.org/0000-0001-8095-0097), School of Geography, Sir Clive Granger Building, Univ. of Nottingham, University Park, Nottingham, UK. – H. Atkins, Royal Botanic Garden Edinburgh, Edinburgh, UK. – J. van Tol, Naturalis Biodiversity Center, Leiden, the Netherlands. – L. Herrera-Alsina, P. Mynard, G. Bocedi (https://orcid.org/0000-0002-9131-6670), C. Gubry-Rangin, L. T. Lancaster and J. M. J. Travis (https://orcid.org/0000-0002-5785-4272), School of Biological Sciences, Univ. of Aberdeen, Aberdeen, UK. – M. Nangoy, Faculty of Animal Husbandry, Sam Ratulangi Univ., Manado, Indonesia. – F. Fahri, Dept of Biology, Tadulako Univ., Palu, Indonesia. – P. Lupiyaningdyah, Museum Zoologicum Bogoriense, Research Center for Biology, Indonesian Inst. of Sciences, Cibinong, Indonesia. – I M. Sudiana, Microbial Ecology Research Group, Research Center for Biology, Indonesian Inst. of Sciences, Cibinong, Indonesia. – B. Juliandi, Dept of Biology, Faculty of Mathematics and Natural Sciences, IPB Univ., Bogor, Indonesia. ACA also at: Dept of Biology, Lakehead Univ., Thunder Bay, Ontario, Canada.

Defining appropriate null expectations for species distribution hypotheses is important because sampling bias and spatial autocorrelation can produce realistic, but ecologically meaningless, geographic patterns. Generating null species occurrences with similar spatial structure to observed data can help overcome these problems, but existing methods focus on single or pairs of species and do not incorporate between-species spatial structure that may occlude comparative biogeographic analyses. Here, we describe an algorithm for generating randomised species occurrence points that mimic the within- and between-species spatial structure of real datasets and implement it in a new R package – *fauxcurrence*. The algorithm can be implemented on any geographic domain for any number of species, limited only by computing power. To demonstrate its utility, we apply the algorithm to two common analysis-types: testing the fit of species distribution models (SDMs) and evaluating niche-overlap. The method works well on all tested datasets within reasonable timescales. We found that many SDMs, despite a good fit to the data, were not significantly better than null expectations and identified only two cases (out of a possible 32) of significantly higher niche divergence than expected by chance. The package is user-friendly, flexible and has many potential applications beyond those tested here, such as joint SDM evaluation and species co-occurrence analysis, spanning the areas of ecology, evolutionary biology and biogeography.

Keywords: environmental niche model, joint species distribution modelling, niche conservatism, niche divergence, niche overlap, null biogeographical model

Introduction

Eco-geographical hypothesis testing using species occurrence data can be hampered by spatial autocorrelation and the difficulty of defining appropriate null expectations (Bahn and Mcgill 2007, Beale et al. 2008, Chapman 2010, Fourcade et al. 2018, Moore et al. 2018). A major issue is that spatial clustering of conspecific, or separation of heterospecific, occurrence records can be affected by multiple factors, which are often difficult to disentangle. These include: 1) habitat suitability (Phillips et al. 2006), 2) dispersal limitation (Glor and Warren 2010), 3) interactions between individuals of the same or different species such as conspecific attraction, competitive exclusion or mutualism (Mielke et al. 2020) or 4) sampling bias, where occurrence records are more likely to be collected from more easily accessible or intensively studied areas (Phillips et al. 2009).

One approach to overcome these issues has been to use null species occurrences to define the expectations if species distributions were entirely random, or if only the inherent spatial structure within species has shaped them (Raes and ter Steege 2007, Beale et al. 2008, Algar et al. 2013, Bohl et al. 2019). In general, these approaches use randomised species distributions generated without reference to variables of interest which are compared to a real species distribution. If the real species distribution shows a closer association with the focal variable than the null models, then this is taken as evidence that they are linked. This approach was pioneered for species distribution modelling by Raes and ter Steege (2007), who used randomly chosen points across the study area to produce null occurrence datasets. More recent studies (Beale et al. 2008, Algar et al. 2013) have used null models to account for the effect of inherent spatial structure within species. To define the null expectation, these approaches use an iterative procedure to produce null species distributions with similar spatial structure to observed occurrences, excluding any consideration of environment or specific geographic location. While these methods are well-suited to testing habitat-suitability hypotheses for single species, they do not take spatial structure between species into account, making them unsuitable for testing multispecies hypotheses involving, for example, niche overlap or range boundaries. Some methods designed to test niche overlap hypotheses employ null models for pairs of species, but these either do not take spatial structure into account (Warren et al. 2008), or simply translocate the entire set of occurrence points. This preserves spatial structure but limits their application to species which are range-restricted relative to the study region (Nunes and Pearson 2017).

Here, we present a method to fill this gap, which is implemented in a new *R* package – *fauxcurrence* ver. 1 (available at <https://github.com/ogosborne/fauxcurrence>). The package can produce null species occurrences which preserve the spatial structure within and between an arbitrary number of species, and provides many options to tailor these occurrences to the user's needs. We demonstrate the utility of the package using a dataset of 22 species of plants, vertebrates and arthropods. We use the resulting null occurrence points to test the significance of species distribution models (SDMs) and to test for significant deviations from null expectations of niche overlap between species.

Material and methods

Method description

Our method (Fig. 1a) has three main modes of operation, distinguished by how inter-point distances are used to define spatial structure. Within-species distances (divided into one subset per species) are always included and can also be used alone (which we refer to here as the '*Intra*' null model; Fig. 1b). The total set of between-species distances can be divided into subsets in two ways: either as sets of general between-species distances per species (i.e. the distances from a species' occurrence points to all heterospecific occurrence points; the '*Inter*' null model; Fig. 1b) or as a separate set of distances between each pair of species in the dataset (the '*Inter-sep*' model; Fig. 1b; Supporting information).

The user provides a set of species occurrence points and a raster defining the study area. The algorithm begins by randomly generating one simulated occurrence point per species. It then adds occurrence points for each species by drawing each new point D distance away from a random existing conspecific point (where D is sampled from the empirical distribution function of observed within-species distances) until each species has the same number of occurrences as in the observed dataset (Supporting information).

Once the initial set of simulated points are generated, the fit of their spatial structure to that of the observed points is iteratively improved. For each iteration, one point is replaced and the match between null and observed interpoint distances is evaluated using discrete Kullback–Leibler (KL) divergence (Kullback and Leibler 1951), where smaller values indicate a better match between the simulated and observed points. Since there are multiple interpoint distance distributions (i.e. within- and between-species distances for multiple species or pairs of species), a weighted mean of KL-divergence across all distributions is used, weighted such that within- and between-species distances contribute equally. The point replacement is only retained if it improves the match, and this procedure is repeated until either no improvement in KL divergence has been made for a set number of iterations or a maximum iteration limit has reached (Fig. 1c–e; see Supporting information for full details). The package includes full documentation and a vignette – 'Using-Fauxcurrence' – with runnable code examples of several use-cases (reproduced in the Supporting information).

Test data

We tested the method on seven species occurrence datasets from Sulawesi, Indonesia each containing between one and six species from a single genus (Supporting information).

Figure 1. Overview of the method. A flowchart (a) shows the basic functioning of the algorithm: initial null points are generated based on observed inter-point distances, these are then iteratively improved. The algorithm finishes when the number of iterations with no improvement (*n.flat*) reaches a user-defined limit. There are three classes of inter-point distance sets in the algorithm (b), within-species (used in all models), general between-species (used in the *Inter* model) and pairwise between-species (used in the *Inter-sep* model). The circles to the right of each indicate which null models they are used in. Panels (c–e) show an example run of the *Inter model*. Kullback–Leibler divergence decreases across iterations (c) and this improvement can be clearly seen when comparing the initial and final null occurrence points (d; map panels). The left of each split violin plot (e) shows the density of observed interpoint distances for each of the distance sets in the model and the right shows those of the null, which are more similar to the observed distances in the final null (bottom) than the initial null (top).

We ran the *Intra* model on all datasets, and the *Inter* model on all datasets with over one species. Since the *Inter* and *Inter-sep* models are identical for species pairs, we only ran the *Inter-sep* model for datasets with over two species. The iterative improvement was continued until there had been no improvement for 10 000 iterations. For each dataset/model combination, we produced 1000 independent, null occurrence replicates, each with similar spatial structure, but differing in the final locations selected by the algorithm.

Comparison to a random null model

To illustrate the effect of including spatial structure in our null models, we also implemented a random null model which ignores spatial structure. Following the approach taken by Raes and ter Steege (2007), we randomly selected raster cells from the study region to produce the appropriate number of occurrences for each species. This procedure was repeated to produce 1000 random null datasets for each species.

SDM model-fitting and niche overlap

We used *Maxent* ver. 3.4.1 (Phillips et al. 2006) to build species distribution models (SDMs) for all species using the 19 *BIOCLIM* climate variables and altitude from the *WorldClim1* database (Hijmans et al. 2005). To determine if SDMs generated from the observed data performed significantly better than those generated from the null models, we calculated area under the receiver operating characteristic curve (AUC; a measure of model discrimination) and minimum training presence omission rate (OR_{MTP}; a measure of model overfitting), for SDMs built from observed data, the *fauxcurrence*-generated null model replicates, and the randomly-generated null model replicates. These statistics were then compared between observed and null SDMs, and p-values were calculated to determine significance. For datasets with more than one species, we compared null and observed niche overlap using Schoener's *D* (Schoener 1968) and Warren's *I* (Warren et al. 2008), between all pairs of congeneric species (see Supporting information for full details).

Results

Method performance

Average time per iteration ranged from 3.2 milliseconds (ms) to 21.5 ms and the mean number of iterations ranged from 48 155 to 341 610 (Supporting information). Number of occurrences was the best predictor of number of iterations to model completion, although number of species also had an effect (Supporting information). Plotting KL divergence across iterations suggested that 10 000 iterations without improvement was more than sufficient to minimise the KL divergence statistic for most datasets (Supporting information). As a result of the smaller number of distance distributions to be minimised for the *Intra* model, it was substantially faster per iteration, and required far fewer iterations to reach minimisation, than either the *Inter* or *Inter-sep* models (Supporting information).

Application to SDM model-fitting and niche overlap

The AUC values were over 0.8 for 13 of 22 species and over 0.9 for four species, although only six of these were significantly greater than null expectations according to at least one of our *fauxcurrence*-generated null model types (Fig. 2a–c; Supporting information), and where they were significantly greater than those of one null model type, they were often not significantly different to those of the others (Fig. 2a–d). In fact, SDMs for only two species, *Cyrtandra geocarpa* and *Sphenomorphus tropidonotus*, had a significantly better AUC than all applied *fauxcurrence*-generated null models (Fig. 2d). For the random null model (Raes and ter Steege 2007), AUC values were clustered around 0.5 and approximately normally distributed for all species (Supporting information), leading to all but two of the species' observed AUC values

being significantly higher than the random null (Supporting information). Only five species had OR_{MTP} significantly smaller than null expectations according to at least one of our *fauxcurrence*-generated null model types (Supporting information) and only three, *Cyrtandra geocarpa*, *Sphenomorphus tropidonotus* and *Limnonectes microtympanum*, had significantly smaller OR_{MTP} than all applied *fauxcurrence*-generated null models (Supporting information). OR_{MTP} values for the random null model were more similar to fauxcurrence-generated null values than AUC values were (Supporting information), leading to 10 species' OR_{MTP} values being significantly lower than the random null, including all species with a significant OR_{MTP} according to the *fauxcurrence*-generated null models.

Observed niche divergence differed from null expectations in only two species pairs, both of which showed significant niche divergence (Supporting information). Both of these involved *Cyrtandra geocarpa*, which was also one of only two species whose SDM fit significantly better than those from all the null models in terms of both AUC and OR_{MTP} (Fig. 2; Supporting information). Both *Inter* and *Inter-sep* null models, and both niche overlap statistics found the same two species pairs to be significant (Supporting information).

Discussion

Here, we describe a new tool to help overcome the difficulties of defining null expectations when working with multi-species occurrence data. Our case study demonstrates the utility of the method. While 13 species had AUC values typically interpreted as 'excellent' or 'outstanding' discrimination (Hosmer et al. 2013), less than half of these were significantly higher than one of our null models. Encouragingly, the two species with significantly higher AUC than expected according to all null models also had significantly lower OR_{MTP} than expected according to all null models. While both AUC and OR_{MTP} values were highly correlated with their p-values, many species with very high AUC scores did not have a significantly better fit than the null models. The lowest AUC score which was significantly higher than any of the null models was 0.76, underlining that SDMs with low AUC scores (e.g. < 0.75) should be treated with great caution. Using our method to demonstrate that SDMs fit significantly better than null expectations in terms of both high discrimination and low over-fitting, will provide much greater certainty than simple inspection of AUC and OR.

Comparison between the *fauxcurrence* models and the random null model (Raes and ter Steege 2007) illustrated the importance of considering spatial structure when conducting null-model analysis. Random null distributions of AUC were very similar between species, with means of approximately 0.5 (signifying no discrimination) in all cases (Supporting information). This may be expected for a null model which consists of randomly chosen occurrences, and means that any observed AUC value substantially higher than 0.5 is likely to be significant according to this model. The contrasting results

Figure 2. The relationship between area under the receiver operating characteristic curve (AUC) of the species distribution models (SDMs) for observed occurrences, and the p-values for comparison to AUC of each null model type (a–c). Dotted lines mark 0.05 on the p-value axis, the solid line is a linear regression line and Pearson's correlation is shown in the top-right of each plot. A Venn diagram (d) shows the overlap in significance ($p < 0.05$) between the three models. Each labelled circle contains species with significantly higher observed AUC than those of simulated data of each null model type and those which were significantly higher than simulated data from multiple null models are shown in the appropriate intersection (those outside the circles were not significant: 'n.s'). Where all implemented models agree, symbols are in bold. Species symbols are shown in the legend (e).

of the *fauxcurrence*-generated null models underlines the fact that spatial autocorrelation of both occurrences and environmental variables can result in SDMs with high discrimination even where no association between species occurrence and environment exists.

The *Intra* model algorithm is largely equivalent to the approach of Beale et al. (2008) differing only in the method used to measure similarity of spatial structure between null and observed data. Comparison between the *Intra* model and the *Inter* and *Inter-sep* models showed that the main novel

feature of our approach – the inclusion of between-species spatial structure – can alter null expectations for SDMs. For example, *Draco spilonotus* had a significantly better fit than the *Intra* null model (p = 0.001), but was not significantly different from either the *Inter* or *Inter-sep* models (p = 0.18 and p = 0.27, respectively; Supporting information). While this is not evidence of it per se, such a pattern could plausibly have a biological basis. For example, if competitive exclusion between species, rather than climate suitability, was responsible for the geographical separation of species into ranges that happen to have distinct climates (Godsoe et al. 2017), such a pattern might be expected.

We also demonstrated the applicability of our method to identify significant niche divergence (or niche conservatism). Our method has advantages over existing approaches. The approach of Nunes and Pearson (2017) creates null replicates by translocating and rotating observed species occurrences. This would clearly be inappropriate for a study region such as Sulawesi, since most rotations and translocations would result in a large proportion of occurrences being translocated to the sea, leading to a high number of similar replicates, as noted by the authors (Nunes and Pearson 2017). The contorted geography of Sulawesi is not unique, and many intensely studied locations such as Isabela Island in the Galápagos archipelago and Lord Howe Island, Australia, also fall into this category.

While we expect *fauxcurrence* to work on a wider range of datasets due to our use of 'as similar as possible' rather than identical spatial structure and the capability to include overland distances, it could have similar issues in extreme cases. In such cases, where there are a very small number of possible null model configurations which fit the observed spatial structure, the high proportion of similar null models would likely lead to underestimation of model significance. Since there are more spatial constraints on models including between-species spatial structure, this is more likely to be the case for the *Inter* and *Inter-sep* models, as can be seen by their less efficient minimisation of KL divergence (Supporting information). Thus we recommend carefully inspecting null occurrence distributions, particularly in cases where species are widely distributed relative to the study region, where summary statistics have little variance between null SDMs, or when the significance estimates from different null models differ.

Aside from the two applications shown here, there are many other potential uses for the package and possible refinements to the way biogeographic models could be assessed using *fauxcurrence*. For example, for widely distributed species, using a spatially-independent evaluation dataset to assess model performance will produce more robust results (Bohl et al. 2019). In the context of *fauxcurrence*, this would involve withholding some of the data as an evaluation dataset, using only the remaining data to build null models and construct the SDM, and assessing both null and observed SDMs using the evaluation dataset. The performance of joint species distribution models (Pollock et al. 2014), which jointly model environmental and community effects on species distributions, could be assessed with our approach in a similar way to our tests of SDM fit. While not tested here, sampling bias could be accounted for by masking the input raster to remove cells with no sampling records across a large multi-species database, an approach taken by Raes and ter Steege (2007). Our method could even be extended to apply to environmental space: while not implemented in the current release, it could be modified to generate null occurrences based on proximity in environmental space by replacing the geographical distance matrix with an environmental distance matrix. Such an approach would maintain the environmental structure in the occurrences, allowing for tests of non-randomness in geographical space. Other potential applications include identifying significant effects of biotic factors (other species) on a focal species' distribution (Algar et al. 2013, Giannini et al. 2013) where the comparison of different null models can give insight into the relevance of pairwise and complex biotic interactions, and testing for significant co-occurrence of range-boundaries across clades (Swenson and Howard 2005). More generally, *fauxcurrence*-generated occurrences could be used in any theoretical biogeographical application where realistic occurrences of species and clades are required. Overall, the method is easy to use, flexible, and can add rigour and insight into investigations of a wide range of problems in ecology, evolution and biogeography.

To cite Fauxcurrence or acknowledge its use, cite this Software note as follows, substituting the version of the application that you used for 'version 1.0':

Osborne, O. G. et al. 2022. Fauxcurrence: simulating multi-species occurrences for null models in species distribution modelling and biogeography. – Ecography 2022: 1–7 (ver. 1.0).

Author contributions

Owen G. Osborne: Conceptualization (equal); Data curation (equal); Formal analysis (lead); Investigation (lead); Methodology (lead); Software (lead); Visualization (lead); Writing – original draft (lead); Writing – review and editing (lead). **Henry G. Fell**: Conceptualization (equal); Formal analysis (supporting); Investigation (supporting); Methodology (supporting); Software (supporting); Writing – review and editing (supporting). **Hannah Atkins**: Data curation (equal); Writing – review and editing (supporting). **Jan van Tol**: Data curation (equal); Writing – review and editing (supporting). **Daniel Phillips**: Data curation (equal); Writing – review and editing (supporting). **Leonel

Herrera-Alsina: Writing – review and editing (supporting). **Poppy Mynard**: Writing – review and editing (supporting). **Greta Bocedi**: Funding acquisition (supporting); Supervision (supporting); Writing – review and editing (supporting). **Cécile Gubry-Rangin**: Funding acquisition (supporting); Supervision (supporting); Writing – review and editing (supporting). **Lesley T. Lancaster**: Funding acquisition (supporting); Supervision (supporting); Writing – review and editing (supporting). **Simon Creer**: Writing – review and editing (supporting). **Meis Nangoy**: Writing – review and editing (supporting). **Fahri Fahri**: Writing – review and editing (supporting). **Pungki Lupiyaningdyah**: Writing – review and editing (supporting). **I M. Sudiana**: Writing – review and editing (supporting). **Berry Juliandi**: Funding acquisition (equal); Supervision (supporting); Writing – review and editing (supporting). **Justin Travis**: Funding acquisition (equal); Supervision (supporting); Writing – review and editing (supporting). **Alexander S. T. Papadopulos**: Conceptualization (equal); Funding acquisition (equal); Supervision (equal); Writing – review and editing (supporting). **Adam C. Algar**: Conceptualization (equal); Formal analysis (supporting); Funding acquisition (supporting); Investigation (supporting); Methodology (supporting); Software (supporting); Supervision (equal); Writing – review and editing (supporting).

Transparent peer review

The peer review history for this article is available at <https://publons.com/publon/10.1111/ecog.05880>.

Data availability statement

Data are available from the Dryad Digital Repository: <https://doi.org/10.5061/dryad.gtht76hp8> (Osborne et al. 2022).

Supporting information

The supporting information associated with this article is available from the online version.

References

Algar, A. C. et al. 2013. Niche incumbency, dispersal limitation and climate shape geographical distributions in a species-rich island adaptive radiation. – Global Ecol. Biogeogr. 22: 391–402.

Bahn, V. and Mcgill, B. J. 2007. Can niche-based distribution models outperform spatial interpolation? – Global Ecol. Biogeogr. 16: 733–742.

Beale, C. M. et al. 2008. Opening the climate envelope reveals no macroscale associations with climate in European birds. – Proc. Natl Acad. Sci. USA 105: 14908–14912.

Bohl, C. L. et al. 2019. A new null model approach to quantify performance and significance for ecological niche models of species distributions. – J. Biogeogr. 46: 1101–1111.

Chapman, D. S. 2010. Weak climatic associations among British plant distributions. – Global Ecol. Biogeogr. 19: 831–841.

Fourcade, Y. et al. 2018. Paintings predict the distribution of species, or the challenge of selecting environmental predictors and evaluation statistics. – Global Ecol. Biogeogr. 27: 245–256.

Giannini, T. C. et al. 2013. Improving species distribution models using biotic interactions: a case study of parasites, pollinators and plants. – Ecography 36: 649–656.

Glor, R. E. and Warren, D. 2010. Testing ecological explanations for biogeographic boundaries. – Evolution 65: 673–683.

Godsoe, W. et al. 2017. Effects of biotic interactions on modeled species' distribution can be masked by environmental gradients. – Ecol. Evol. 7: 654–664.

Hijmans, R. J. et al. 2005. Very high resolution interpolated climate surfaces for global land areas. – Int. J. Climatol. 25: 1965–1978.

Hosmer, D. W. et al. 2013. Applied Logistic Regression. John Wiley and Sons, 3rd ed. pp. 173–182.

Kullback, S. and Leibler, R. 1951. On information and sufficiency. – Ann. Math. Stat. 22: 79–86.

Mielke, K. P. et al. 2020. Disentangling drivers of spatial autocorrelation in species distribution models. – Ecography 43: 1741–1751.

Moore, T. E. et al. 2018. Spatial autocorrelation inflates niche breadth–range size relationships. – Global Ecol. Biogeogr. 27: 1426–1436.

Nunes, L. A. and Pearson, R. G. 2017. A null biogeographical test for assessing ecological niche evolution. – J. Biogeogr. 44: 1331–1343.

Osborne, O. G. et al. 2022. Data from: Fauxcurrence: simulating multi-species occurrences for null models in species distribution modelling and biogeography. – Dryad Digital Repository, <https://doi.org/10.5061/dryad.gtht76hp8>.

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

Phillips, S. J. et al. 2009. Sample selection bias and presence-only distribution models: implications for background and pseudo-absence data. – Ecol. Appl. 19: 181–197.

Pollock, L. J. et al. 2014. Understanding co-occurrence by modelling species simultaneously with a joint species distribution model (JSDM). – Methods Ecol. Evol. 5: 397–406.

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

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

Swenson, N. G. and Howard, D. J. 2005. Clustering of contact zones, hybrid zones and phylogeographic breaks in North America. – Am. Nat. 166: 581–591.

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