The *n*-dimensional hypervolume

Benjamin Blonder^{1,2,3*}, Christine Lamanna^{2,4}, Cyrille Violle⁵ and Brian J. Enquist^{1,2,6}

¹Department of Ecology and Evolutionary Biology, University of Arizona, 1041 E Lowell Street, Tucson, AZ 85721, USA, ²Rocky Mountain Biological Laboratory, PO Box 619, Crested Butte, CO 81224, USA, ³Center for Macroecology, Evolution, and Climate, University of Copenhagen, Universitetsparken 15, Copenhagen, DK-2100, Denmark, ⁴Sustainability Solutions Initiative, University of Maine, 5710 Norman Smith Hall, Orono, ME 04469, USA, ⁵Centre d'Ecologie Fonctionnelle et Evolutive-UMR 5175, CNRS, Montpellier, France, ⁶The Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

*Correspondence: Benjamin Blonder, Department of Ecology and Evolutionary Biology, University of Arizona, 1041 E Lowell Street, Tucson, AZ 85721, USA. E-mail: bblonder@email.arizona.edu

ABSTRACT

Aim The Hutchinsonian hypervolume is the conceptual foundation for many lines of ecological and evolutionary inquiry, including functional morphology, comparative biology, community ecology and niche theory. However, extant methods to sample from hypervolumes or measure their geometry perform poorly on high-dimensional or holey datasets.

Innovation We first highlight the conceptual and computational issues that have prevented a more direct approach to measuring hypervolumes. Next, we present a new multivariate kernel density estimation method that resolves many of these problems in an arbitrary number of dimensions.

Main conclusions We show that our method (implemented as the 'hypervolume' R package) can match several extant methods for hypervolume geometry and species distribution modelling. Tools to quantify high-dimensional ecological hypervolumes will enable a wide range of fundamental descriptive, inferential and comparative questions to be addressed.

Keywords

Environmental niche modelling, hole, Hutchinson, hypervolume, kernel density estimation, morphospace, niche, niche overlap, species distribution modelling.

INTRODUCTION

Hutchinson first proposed the *n*-dimensional hypervolume to quantify species niches (Hutchinson, 1957). In this approach, a set of *n* variables that represent biologically important and independent axes are identified and the hypervolume is defined by a set of points within this *n*-dimensional space that reflects suitable values of the variables (e.g. temperature or food size). The hypervolume concept of the niche is widely used in comparative biology (Pigliucci, 2007) and evolutionary biology (e.g. fitness landscapes; Gavrilets, 2004). Within ecology it can be applied beyond the quantification of species niches (Violle & Jiang, 2009), for instance to quantify the multivariate space of a community or a regional pool (Ricklefs & O'Rourke, 1975; Foote, 1997), to measure morphology (Raup & Michelson, 1965) or to test functional ecology hypotheses (Albert *et al.*, 2010; Baraloto *et al.*, 2012; Boucher *et al.*, 2013).

The use of hypervolumes in biology arises through the resolution of three related mathematical questions that are independent of scale and axis choice. The first question is about the geometry of the hypervolume. Given a set of observations, what can be inferred about the overall shape of the hypervolume, its

total volume and the presence or absence of holes? This question is relevant to topics including environmental or trait filtering in community assembly (Whittaker & Niering, 1965), forbidden trait combinations in physiological ecology and evolutionary biology (Wright, 1932; Maynard-Smith et al., 1985) and climate breadths in invasion ecology (Petitpierre et al., 2012). A second question about set operations can then be addressed for multiple hypervolumes whose geometry is known. How much do hypervolumes overlap, and what portion of each is unique? These questions are relevant to topics including competitive exclusion (May & MacArthur, 1972; Tilman, 1982; Abrams, 1983), species packing (Findley, 1973; Pacala & Roughgarden, 1982; Ricklefs & Miles, 1994; Tilman et al., 1997) and functional redundancy within communities (Petchey et al., 2007). The third question is about sampling from the n-dimensional space. Is a candidate point in or out of the hypervolume? Sampling questions are equivalent to species distribution modelling (Elith & Leathwick, 2009; Peterson et al., 2011), an approach in which a set of geographic points are projected into hyperspace, those points are determined to be in or out of the hypervolume, and are then back-projected into geographic space as range maps.

A Journal of Macroecology

While these three mathematical questions integrate a wide range of topics, they have not traditionally been considered in a unified framework. Indeed, independent methods have been developed for each of the above questions. For example, the geometry question has typically been addressed using volumeestimation methods such as a convex hull (Cornwell et al., 2006). The question on set operations has been primarily addressed using a range of overlap indices (Mouillot et al., 2005; Villéger et al., 2008; Warren et al., 2008). The sample question has mostly been addressed using predictive modelling techniques [e.g. generalized linear models (GLM), generalized boosted regression models (GBM), or MaxEnt (Elith et al., 2006; Wisz et al., 2008)]. However, methods that are successful for one question may not be directly transferrable to the other questions. For example, the sampling question can be resolved without delineating all the boundaries of a hypervolume (i.e. sampling the entire hyperspace). While resolving the geometry and set operation questions would effectively resolve the sampling question, existing approaches have been limited.

Here we argue that these three major questions can be addressed with a unified approach to infer hypervolumes from observations. Further, we highlight the key conceptual and dimensional issues that have previously limited the development of such approaches. We then propose a new method – a thresholded multivariate kernel density method – that can simultaneously address each of these questions. We show that our method matches extant methods for all three questions but can also be applied in high dimensions. We demonstrate the method with a simulation analysis and with two examples: the morphological hypervolume overlap of Darwin's finches, and climate hypervolume and geographic range projections of two *Quercus* species.

Extant methods have conceptual limitations

The general mathematical problem is how to best estimate a hypervolume from a set of observations. Ideally, an estimation procedure should: (1) directly delineate the boundaries of the hypervolume; (2) not assume a fixed distribution of observations; (3) account for disjunctions or holes; (4) not be sensitive to outlier points; and (5) produce a bounded result (i.e. not predict infinite volumes).

Using these criteria, many extant methods fall short (Fig. 1). For example, principal component analysis (e.g. Ricklefs & Travis, 1980), although intuitively appealing, assumes that the hypervolume is multivariate normal, violating procedure (2) above. Many empirical datasets indicate that even singledimensional responses often deviate strongly from normality (Austin et al., 1984) via high skewness or multiple modes that cannot be removed by transformation. Multivariate range boxes (e.g. Hutchinson, 1957) are also inappropriate because they assume that the hypervolume is multivariate uniform and with box axes aligned to coordinate axes, also violating procedure (2). Other ordination approaches (e.g. outlying means index (OMI); Doledec et al., 2000) have similar distributional limitations or may be better suited for discrimination than geometric applications (Green, 1971). While a convex hull (Cornwell et al., 2006) and other envelope methods (Nix, 1986) are distribution-free approaches that can provide a closer measurement of the hypervolume, they are sensitive to outlier points. As a result, estimates of the shape of the hypervolume using convex hulls can result in errors in measurements of volume and shape. More importantly, none of these three methods can model disjunctions or holes in the hypervolume, complicating the assessment of hypervolume overlap (see discussion below). A potentially more robust approach is to fit different functions to each hypervolume dimension (e.g. Gaussian mixture models; Laughlin et al., 2012); however, this method requires some choices to be made about the nature of the fitting function and results in an estimated hypervolume that may not include interactions or covariation between dimensions.

While species distribution models (SDMs) provide multiple algorithms for sampling that can capture many of these nonlinearities, none of these methods permit delineation of the entire hypervolume. This is because SDMs are intended primarily for sampling points from the entire environmental space (i.e. transformed geographic coordinates), which is computationally simpler than delineating boundaries of the environmental space. Below we discuss the sampling versus delineation problem in more depth. Additionally, SDMs may generate environmental hypervolumes with unbounded volumes, because they may predict that all values along an axis greater/smaller than some threshold value are within the hypervolume (Peterson *et al.*, 2011).



Figure 1 A robust operational definition of the hypervolume is important for making correct inferences. Here we show an example of three poor definitions (b–d) and one accurate definition (e). (a) Consider a two-dimensional dataset describing a hypothetical 'Swiss cheese' hypervolume. (b) A two-dimensional range box fails to capture the rotation and holes in the hypervolume. (c) A principal components analysis has difficulties with non-normal data and does not account for the holes. (d) A convex hull also does not account for holes and is very sensitive to outlying points. (e) The best solution is to take the area enclosed by a contour of a kernel density estimate, which can account for non-normal, rotated and holey data including outliers. Some species distribution modelling approaches (e.g. generalized boosted regression models) can also approximate this shape.

Lastly, there are extant metrics and indices for different properties of the hypervolume, including breadth and overlap (Maguire, 1967; Colwell & Futuyma, 1971; Hurlbert, 1978; Abrams, 1980). However, these approaches do not give direct insight into the geometry and topology of the hypervolume that are needed for many current research questions. As a result, more powerful methods to measure and compare hypervolumes in biology are needed.

High dimensions are different (and harder)

The geometry of a high-dimensional hypervolume may differ qualitatively from a low-dimensional hypervolume in ways that have not been adequately considered, because human intuition is best suited for low-dimensional (n = 1-3) systems. Highdimensional biological hypervolumes are probably not smooth continuous shapes but rather rugged or filled with gaps or holes (Colwell & Futuyma, 1971; May & MacArthur, 1972; Hurlbert, 1978; Abrams, 1980; Jackson & Overpeck, 2000). For example, the recognition that most high-dimensional fitness landscapes are 'holey' has provided many advances in understanding evolutionary dynamics (Gavrilets, 1997; Salazar-Ciudad & Marin-Riera, 2013). While fundamental niches are often thought to have simpler and less holey geometry than realized niches (Colwell & Rangel, 2009; Araujo & Peterson, 2012), data limitations have precluded robust tests of this idea. Moreover, all distance-based SDMs (e.g. DOMAIN) are limited by the assumption of no holes. We argue that there should be no a priori reason to assume that a hypervolume (or niche) should be normally or uniformly distributed in multiple dimensions.

Many geometric questions have been pursued exclusively with low-dimensionality analyses (Broennimann *et al.*, 2012; Petitpierre *et al.*, 2012). However, some hypervolumes may be better analysed in higher dimensions. Though the number of axes necessary to describe any system can be debated, there is no reason to believe that two or three dimensions are sufficient for most systems.

In community ecology, trait axes are often implicitly used as proxies for niche axes (Westoby *et al.*, 2002). Currently a focus on measures of single traits as a metric of position along a niche axis is widespread in trait-based ecology and evolution. The goal of this work is to track assembly processes by analysing the distribution of a single trait (e.g. body size) within and between ecological communities. However, evidence suggests that community assembly is driven by integrated phenotypes rather than by single traits (Bonser, 2006), meaning that a hypervolume approach to community assembly may be more relevant.

Delineating high-dimensional hypervolumes using existing approaches is difficult. For example, quantifying a hypervolume with arbitrarily complex geometry requires that the entire hyperspace must be sampled. In low dimensions, this is simple. The geometry of a one-dimensional hypervolume can be defined by determining if each of g regularly spaced points in an interval is in or out. However, for increasingly higher-dimensional hypervolumes this procedure must be repeated independently in each dimension, requiring g^n evaluations. For

example, characterizing a hypervolume where g = 500 and n = 3 requires more than 10^8 evaluations; for n = 10 dimensions, it is more than 10^{26} evaluations. As a result, exhaustive sampling approaches are too computationally demanding to be practical. Thus, robust methods from species distribution modelling have not been useful for delineating hypervolumes. Developing new methods that can handle high-dimensional datasets will remove the limitations of these extant estimation procedures.

METHODS

Measuring the hypervolume is now possible

Direct estimation of the *n*-dimensional hypervolume from a set of observations *w* can be achieved by a multidimensional kernel density estimation (KDE) procedure. While kernel density approaches for hypervolume delineation have been successfully used in low-dimensional systems (Broennimann *et al.*, 2012; Petitpierre *et al.*, 2012), they had not previously been computationally feasible in high dimensions. Similarly, other fast kernelbased approaches such as support vector machines (Guo *et al.*, 2005) work in transformed high-dimensional spaces but had not been directly applied to geometry questions, as we show here. We now outline the KDE approach and propose a method that can resolve the sampling problem in high dimensions.

We formally define the hypervolume, z, as a set of points within an *n*-dimensional real-valued continuous space that encloses a set of *m* observations, *w*. The problem is to infer *z* from *w*. We start by assuming that *w* is a sample of some distribution *Z*, of which *z* is a uniformly random sample. Next, we compute a kernel density estimate of *Z*, \hat{Z} , using the observations *w* and bandwidth vector \bar{h} . Lastly, we choose a quantile threshold parameter $\tau \in [0,1]$. As a result, *z* can be defined as a set of points enclosed by a contour of \hat{Z} containing a fraction $1 - \tau$ of the total probability density. We illustrate this procedure graphically in Fig. 2.

We describe methods to perform this kernel density estimation of a hypervolume z for both large n and m. The computational problems associated for scaling up this method can be solved with importance-sampling Monte Carlo integration. The resulting algorithms can determine the shape, volume, intersection (overlap), union and set difference of hypervolumes. They can also perform sampling (i.e. species distribution modelling) via inclusion tests in order to determine if a given n-dimensional point is enclosed within a hypervolume or not (Fig. 3). Together, these tools make it possible to directly address our three major questions and move beyond metrics that provide incomplete descriptions of hypervolumes in high dimensions. We describe the algorithms conceptually in Box 1 and in full mathematical depth in Box 2. The software is freely available as an R package ('hypervolume'), with full documentation and several example analyses, including those presented in this paper.

Usage guidelines and caveats

To ensure that hypervolume analyses are replicable, we recommend reporting the chosen bandwidth \bar{h} (or the algorithm used

Box 1 A cartoon guide to the hypervolume algorithms

We present an example of hypervolume creation and set operations to develop the reader's conceptual understanding of how the algorithms are implemented. For clarity, the example is drawn in n = 2 dimensions but the algorithms generalize to an arbitrary number of dimensions.

Creation

The algorithm proceeds by (a) computing an *n*-dimensional kernel density estimate by overlaying hyperbox kernels (gray boxes) around each observation (black dots), (b) sampling from these kernels randomly (gray dots), (c) importance-sampling the space using these boxes and performing range tests on random points using a recursive partitioning tree (rainbow colours are proportional to kernel density) then (d) applying a threshold that encloses a specified quantile of the total probability mass, retaining only points within the resulting volume and then using combined properties of the kernel and importance-sample to subsample the random points to a uniform point density (purple dots). These uniform-density points, along with the known point density and volume, constitute the full stochastic description of the hypervolume. The key advance of the algorithm is to develop efficient approaches for importance-sampling high-dimensional spaces using box kernels and recursive partitioning trees, as described in depth in Box 2.

Set operations

Uniformly random points in an *n*-dimensional space are likely to be separated by a characteristic distance. The algorithm uses a *n*-ball test with this distance on the candidate point relative to the hypervolume's random points. If at least one random point in the hypervolume is within the characteristic distance of the candidate point, then the point satisfies the inclusion test. An example is shown here as a zoom from the full hypervolume intersection. The algorithm uses this inclusion test to determine which random points in the first hypervolume are and are not enclosed within the second hypervolume, and vice versa. The intersection is inferred to include the points that satisfy both inclusion tests. The unique component of the first hypervolume is inferred to include the points that satisfy the first inclusion test, and vice versa for the unique component of the second hypervolume. The union is inferred to include the unique components of both hypervolumes and the intersection (as calculated above). In all cases the resulting random points are resampled to uniform density and used to infer a new point density and volume. (e) An example is shown of overlap between two hypervolumes. Each hypervolume's random uniformly sampled points are coloured as green or purple, and a ball of the appropriate radius is drawn around each point. Points that have overlapping balls (coloured blue) are inferred to constitute the intersection. (f) A zoom of the boxed region in (e).



to choose it) and the quantile τ obtained by the algorithm (which may differ slightly from the specified quantile due to some discrete approximations in the algorithm; see Box 2).

There are several issues that should be considered before an investigator applies this hypervolume method. First, our approach is best suited for continuous variables. Categorical variables are problematic because a volume is not well defined when the same distance function cannot be defined for all axes. If it is necessary to use categorical variables, the data can first be ordinated into fewer dimensions using other approaches (e.g. the Gower distance; Gower, 1971). We acknowledge that categorical variables are often biologically relevant, and wish to

Global Ecology and Biogeography, © 2014 John Wiley & Sons Ltd

Box 2 Mathematical description of the hypervolume algorithms

Hypervolume construction

For a hypervolume *z* and measurements *w*, we perform the kernel density estimation and volume measurement using a Monte Carlo importance sampling approach. Suppose we are given a set of *m* points $w = \{w_1, \ldots, w_m\} \in \Re^n$ drawn from an unknown probability distribution *Z*, a kernel *k*, a threshold τ , and an *r*-fold replication parameter corresponding to the number of Monte Carlo samples. We wish to find the volume of *z* and return a set of uniformly random points II with point density π from within *z*.

The idea is to choose a space *P* (such that $Z \subset P \subset \Re^n$) and randomly sample points II = { p_1, \ldots, p_r } $\subset P$. At each point $\overline{p_i}$, evaluate the kernel density estimate

$$\hat{Z}(\overline{p_i}) = \alpha \sum_{j=1}^m k \left(\frac{\overline{p_i} - \overline{w_j}}{\overline{h}} \right)$$

where vector division indicates division in each dimension independently and α is a normalization constant. Now flag the indices $(i_1, \ldots, i_d): \hat{Z}(\overline{p_i}) > q$ for the *q* that satisfies

$$\sum_{\{i\}} \hat{Z}(\overline{p_i}) \ge \tau \sum_{\{i\}} \hat{Z}(\overline{p_i}).$$

If the volume of z is |z| and the volume of P is |P|, then

$$\lim_{x\to\infty}\frac{1}{s}\sum_{\{i\}}\hat{Z}(\overrightarrow{p_i})=\frac{|z|}{|P|}.$$

That is, the ratio of the sum of the kernel density estimates to the number of sampled points converges to the ratio of the true volume to the volume of the sampled space.

We assume an axis-aligned box kernel,

$$k(\vec{x}) = \begin{cases} 1 & : |\vec{x} \cdot \vec{e}_i| < 1/2, \forall_i \in \{1, \dots, n\} \\ 0 & : \text{otherwise} \end{cases}$$

where \vec{e}_i is the *i*th Euclidean unit vector. The proportionality constant is now

$$\alpha = \frac{1}{m \prod_{i=1}^{n} \vec{h} \cdot \vec{e}_{i}}$$

The kernel bandwidth vector \vec{h} can be specified by the investigator. Alternatively, it can be chosen quasi-optimally using a Silverman bandwidth estimator for one-dimensional normal data as

$$\vec{h} = \sum_{i=1}^{n} 2\sigma_i (4/3n)^{1/5} \vec{e}_i$$

where σ is the standard deviation of points in *w* in the *i*th dimension:

$$\sigma_{i} = \sqrt{\frac{1}{m} \sum_{k=1}^{m} (\vec{w}_{k} \cdot \vec{e}_{i} - \mu_{i})^{2}}, \quad \mu_{i} = \frac{1}{m} \sum_{j=1}^{m} \vec{w}_{j} \cdot \vec{e}_{i}$$

We choose this kernel representation because (1) it reaches zero in a finite distance and (2) has constant non-zero value, enabling the evaluation of the kernel density estimate to be reduced to a counting problem.

In practice, random sampling of *P* is impractical because most regions of a high-dimensional space are likely to be empty. Instead, we proceed by importance sampling. Because of the choice of *k*, we know that *Z* has non-zero probability density only within regions that are within an axis-aligned box (with widths given by h_i) surrounding each point w_j , i.e.

$$P = \left\{ p_k : \forall i \in \{1, \dots, n\}, \bigcup_{j=1}^m \left| \left(\overline{p_k} - \overline{w_j} \right) \cdot \overline{e_i} \right| < h_i \right\}.$$

We therefore generate a uniformly random set of points drawn only from axis-aligned boxes centred around each w_j , each of which has point density

$$\pi = \frac{r}{m \prod_{i=1}^{n} h_i}.$$

This process yields

$$\Pi = \left\{ \bigcup_{i=1}^{m} \bigcup_{j=1}^{r/m} U\left(\overline{w_i} - \vec{h}, \overline{w_i} + \vec{h}\right) \right\}$$

where U(a, b) represents a single draw from the uniform distribution scaled to the interval (a, b).

However each axis-aligned box may intersect with multiple other axis-aligned boxes, so \prod is not uniformly random. Our next step is therefore to determine which random points $\overline{p_i} \in \Pi$ fall within regions of \hat{Z} with higher probability density and correct for their oversampling. Because of the choice of k, we know that each $\overline{p_i}$ has a kernel density $\hat{Z}(\overline{p_i})$ estimate proportional to the number of data points whose kernels (i.e. axis-aligned boxes) intersect this $\overline{p_i}$. We build an *n*-dimensional recursive partitioning tree *T* from the data points *w*. Then for each $\overline{P_i}$, we perform a range-box query on *T*, where the range is chosen to be \overline{h} , and count the number of non-zero returns, which is proportional to $\hat{Z}(\overline{p_i})$.

Each point is now over-sampled by a factor proportional to $1/\hat{Z}(\vec{p_i})$, yielding an effective number of sampled points given by

$$\rho = \sum_{i=1}^{r} \frac{1}{\hat{Z}(\overline{p_i})}.$$

The total volume of z is therefore the original point density divided by the effective number of points, or

$$|z| = \frac{\pi}{\rho}.$$

Finally we obtain a uniformly random sample of points from z, Π^* , by sampling σ points from Π , weighting each point by $1/\hat{Z}(\overline{p_i})$ and retaining only the ρ^* unique points $\overline{r^*}$, where $\sigma = \pi |z|$ reflects the original uniformly sampled point density. We then calculate the final point density π^* as $\pi^* = \rho^*/|z|$.

Inclusion test

Consider a set of *n*-dimensional points $p = \{\overline{p_i}\} \in \Pi_a^*$ with point density π_a^* . We wish to determine if a point $\overline{X_j}$ is within the volume sampled by Π_a^* . The characteristic distance between uniformly random points is $d = (\pi_a^*)^{-1/n}$; this means that $\overline{X_j}$ is likely to be within Π_a^* if $\exists i : |\overline{p_i} - \overline{x_j}| < d$. This is implemented by a ball test using an *n*-dimensional recursive partitioning tree built from points in Π_a^* .

Hypervolume set operations (intersection, union, unique subset)

Consider two hypervolumes z_a and z_b described by volumes $|z_a|$ and $|z_b|$, uniformly random point samples Π_a^* and Π_b^* , and point densities π_a^* and π_b^* . We wish to find $z_c = z_a \cap z_b$ as described by $|z_c|$, Π_c^* , and π_c^* . First, both z_a and z_b are uniformly randomly sampled to a point density of $\pi_c = \min(\rho, \pi_a^*, \pi_b^*)$, where ρ is a user-specified value (using high point densities can be computationally costly but not necessarily significantly more accurate), yielding m_a and m_b random points respectively in each hypervolume. Then we use the inclusion test (described above) to find the set of points contained in both z_a and z_b by identifying the r_{ab} , random points in z_a enclosed by $z_b(\Pi_{ab})$, and the r_{ba} random points in z_b , enclosed by z_a (Π_{ba}). We calculate the final volume conservatively as the number of points divided by the point density,

$$|z_c|=\frac{1}{\pi_c}\min(r_{ab},r_{ba}).$$

The uniformly random sample of points in z_c is then $\prod_{c}^{\star} = \prod_{ab} \bigcup \prod_{ba}$ and the final point density is $\pi_c^{\star} = 2\pi_c$.

We also wish to characterize the union and unique components of these hypervolumes, z_{union} and $z_{un} a$ and $z_{un} b$. We apply the intersection algorithm (described above) to subsample Π_a^* and Π_b^* to the same point density π_{int}^* , then to find the intersection hypervolume z_{int} , and also to flag the points not in each hypervolume, $c_{ab} = \Pi_a^* \notin z_b$, and $c_{ba} = \Pi_b^* \notin z_a$. We then determine the final volume as $|z_{union}| = |z_a| + |z_b| - |z_{int}|$. The random sample of points is $\Pi_{union}^* = c_{ab} \cup c_{ba} \cup \Pi_{int}^*$ and the final point density is $\pi_{union}^* = \pi_{int}^*$. We take a similar approach using the flagged points in one hypervolume and not the other to determine the unique components of each hypervolume.



Figure 2 Illustration of the hypervolume estimation procedure. Consider a one-dimensional set of observations, assumed to be a sample from a probability distribution. Estimate this distribution with a kernel density estimate. Slice (subset) the distribution at different probability levels until at least chosen fraction $1 - \tau$ of the probability density is enclosed by the distribution. The estimated hypervolume is then defined by the axis values of this subset (e.g. Fig. 1e; here shown in black for several values of τ where hypervolumes for each probability fraction are colour-coded where hotter colours include cooler ones). The kernel density estimation and slicing process can naturally be extended to multiple dimensions using importance-sampling Monte Carlo methods described conceptually in Box 1 and in detail in Box 2.

highlight this issue as a major unavoidable limitation of the hypervolume concept that also applies to the other methods discussed.

Missing data may restrict the dimensionality of the analysis. Any observation with at least one missing variable cannot be used for hypervolume estimation because an n-dimensional object is not well defined in fewer than n dimensions. In these situations it will be necessary to remove observations with missing values, reduce the dimensionality of the analysis or fill in missing data values via some other approach, for example multiple imputation (Rubin, 1996) or hierarchical probabilistic matrix factorization (Shan *et al.*, 2012).

Choosing comparable units for each axis is critical. Because volume scales with the product of the units of all dimensions it can be difficult to ascribe a change in volume to one axis if units are not comparable. Similarly, non-comparable dimensions mean that results would not be invariant to changes of units and or scale (e.g. redefining an axis in millimetres instead of metres). Thus the observational data must be normalized (e.g. using *z*-scores or log-transformation) before the hypervolume method can be applied. The units of the output hypervolume will therefore be the product of axis units [e.g. in powers of standard deviations (SDs) or logarithmic units].

It is important to clearly identify appropriate and biologically relevant axes. For example, it may be unclear what variables should and should not be included in an analysis, and how sensitive a given result may be to this choice (Petchey & Gaston, 2006; Bernhardt-Römermann *et al.*, 2008). Inclusion of dimen-



Figure 3 Hypervolume geometry operations. From two sets of observations (red and blue) (a), hypervolumes can be created (b), enabling measurement of shape and volume. (c) The total volume occupied by two hypervolumes can be determined as the union of both hypervolumes. (d) Overlap can be measured after finding the intersection between two hypervolumes. (e) The components of each hypervolume that are unique can be identified by set difference operations. (f) An inclusion test can determine if a given point is found within a hypervolume, and enable sampling applications such as species distribution modelling.

sions with limited or highly correlated variation will produce degenerate results, such that the hypervolume is effectively constrained to a hyperplane (Van Valen, 1974). These problems can be identified by high variance in the SDs of each dimension or by high Pearson correlation coefficients between each pair of dimensions. Additionally, choice of the number of dimensions to include may also influence results. For example, hypervolumes that appear to overlap in low dimensions may not overlap if more dimensions are added, and conversely, with the addition of extra redundant dimensions, estimates of overlap may be falsely inflated. Nonetheless, we do expect that hypervolume metrics should be comparable across datasets with identical axes.

Sampling issues also deserve careful consideration. The hypervolume approach assumes that the input set of observations is an unbiased sample of the actual distribution. Meeting this assumption may be difficult, depending on the methodology used for data collection. Unavoidable spatial and taxonomic biases can conflate the occurrence and observation processes in real-world datasets. For example, realized climate niches of species may be oversampled in easily accessed regions and undersampled regions that are more difficult to access. This will lead to incorrect inference of holes. While these biases are common to all niche modelling algorithms (Phillips *et al.*, 2009), the kernel density approach used in our method may be more prone to overfitting the data.

Our method can be used with any number of observations regardless of dimensionality. However, analyses with few observations (m/n < 10), as a rough guideline) will be very sensitive to the choice of bandwidth and are not recommended. For example, the volume inferred for a single point is necessarily equal to the product of the kernel bandwidth along each axis and is not biologically relevant. In general, choosing a smaller bandwidth (or large threshold) will lead to a smaller hypervolume, with each observation appearing disjoint from others, while choosing a larger bandwidth (or small threshold) for the same dataset will lead to a larger volume with more observations appearing to be connected. The investigator must thus carefully consider and potentially standardize the choice of bandwidth and threshold for the hypervolume construction process. A bandwidth can be estimated from the data using a quasi-optimal approach (e.g. a Silverman estimator; Silverman, 1992) that pads each observation by an amount that depends on

the number of available observations and the total range of variation between observations (reflecting an increasing level of confidence that the observations have sampled the extreme boundaries of the hypervolume).

RESULTS

Application to simulated data

Dataset choice

We next compared our approach with other extant methods using simulated data of a variety of complexities, dimensionalities (*n*), and number of unique observations (*m*). We constructed two test datasets of varying complexity. The first dataset, T_C , is defined by *m* samples from a single *n* dimensional hypercube (Fig. 4a):

$$T_{\rm C}(m,n) = \{ \vec{x}_i : \forall i \in (1,...,m), H(\vec{x}_i,n) = 1 \}$$

where H is the hypercube function

$$H(\bar{x},n) = \begin{cases} 1 & |x_i| < 0.5, \forall i \in (1, \dots, n) \\ 0 & \text{otherwise} \end{cases}$$

The second dataset, T_{DC} , is defined by *m* samples from double *n*-dimensional hypercubes, each offset from the origin by two units (Fig. 4b):



Figure 4 Hypervolume geometric analysis of simulated data. (a) Data sampled from a hypothetical single hypercube (T_c) dataset. (b) Data sampled from a hypothetical double hypercube (T_{DC}) dataset, with each hypercube offset by two units from the origin in all axes. In (a) and (b), data are shown at the same scale in two dimensions for clarity but were simulated in up to eight dimensions for the analyses. (c) Comparison of volumes estimated by different methods (hypervolume, minimum volume ellipsoid, range box, convex hull) for the $T_{\rm C}$ dataset. Each boxplot represents the distribution of volumes inferred from 10 independent samples of *m* points from a *n*-dimensional dataset. Boxes that are closer to the black line (the true volume) indicate better methods. The y-axis is log-transformed and normalized by n to reflect the geometric scaling of volume with dimension. (d) Comparison of volumes for the $T_{\rm DC}$ dataset.

$$T_{\rm DC} = \{ \bar{x}_j : \forall i \in (1, \dots, m), (H(\bar{x}_j - \bar{2}) = 1) \lor (H(\bar{x}_j + \bar{2}) = 1) \}$$

In the first example, $T_{\rm C}$ or volume 1 is intended to represent a simple hypervolume that other methods should easily estimate, while in the second example, $T_{\rm DC}$ or volume 2 is intended to represent a complex disjoint hypervolume that may challenge extant methods. For each example, we generated 10 independently sampled test datasets for each parameter combination of m = 10, 100 and 1000 observations and n = 2, 4, 6 and 8 dimensions.

Geometric application

We estimated the volume of each dataset using our method and a range of alternatives: range boxes, minimum volume ellipsoids (similar to a principal components analysis) and a convex hull, and compared results with the known volume of each dataset. Hypervolumes were inferred using a Silverman bandwidth estimator and a quantile threshold of 0.5.

We found that, for the $T_{\rm C}$ dataset, the range box, convex hull and hypervolume methods consistently performed well, but the minimum volume ellipsoid method consistently overestimated volumes (Fig. 4c). This result indicates that the hypervolume method performs well in comparison with extant volume estimation methods for simple datasets. However, for the disjoint $T_{\rm DC}$ dataset we found that the minimum volume ellipsoid and range box consistently overestimated volumes. The convex hull performed best and the hypervolume method performed second best when the sampling effort was high (large m) (Fig. 4d). Nevertheless, unlike the convex hull and minimum volume ellipsoid methods, the results of our hypervolume method were consistent regardless of dimension. This result indicates that our approach provides a viable consistent tool for estimating the volume of complex hypervolumes. The overestimation of volume is not necessarily a problem, and arises because (as previously discussed) the hypervolume method estimates volumes that are due to the choice of kernel bandwidth specified by the researcher.

Sampling application

We next tested the ability of multiple methods to predict presences and absences, using the same test datasets and combinations of dimensions and observations as described above. We compared the hypervolume method with two common and high-performing species distribution modelling algorithms: GLM and GBM (Wisz *et al.*, 2008). We did not evaluate MaxEnt because it is formally equivalent to a GLM (Renner & Warton, 2013). For both GBM and GLM we used fixed thresholds of 0.5 to convert predictions to binary presence/absences, mirroring the threshold used for the hypervolume method. Although more robust approaches are available for determining threshold values (Peterson *et al.*, 2011), this simple approach enables us to facilitate comparisons between all of the methods under equally challenging conditions. The hypervolume method works using only presence data. For the GLM/GBM approaches, we built models using pseudoabsences obtained by sampling from a hyperspace consisting of the region $\{\bar{x}_k : H(\bar{x}_k/6, n)=1\}$, i.e. the hypercube spanning (-3,3) in each axis. We then generated a set of *n*-dimensional test points, half of which were known to be in the analytically defined hypervolume and half of which were known to be outside it. Next, we used each model to make predictions for these points, and then computed two metrics to compare the performance of each model: sensitivity, which measures the true positive rate for predictions, and specificity, which measures the true negative rate for predictions. Better-performing methods have sensitivities and specificities closer to one.

We found that for the simpler $T_{\rm C}$ dataset, the hypervolume and GBM methods had equivalent perfect sensitivity, with the GLM method showing lower sensitivity for large values of m (Fig. 5a). The hypervolume and GBM methods showed higher specificities than the GLM method, with the hypervolume method performing best for smaller values of m (Fig. 5b). For the more challenging $T_{\rm DC}$ dataset, the hypervolume method had similar sensitivity to the GBM model regardless of m (Fig. 5c) but clearly outperformed the GBM model in high dimensions. Additionally, the hypervolume method had consistently higher specificity than the other approaches regardless of m or n (Fig. 5d).

Together, these results indicate that the hypervolume method not only compares favourably with other species distribution modelling methods for simple geometries, but can outperform other methods when the dataset being measured has a complex geometry (e.g. high specificity). While these preliminary results are limited in their scope, they do suggest that the hypervolume method also should be considered as a viable candidate for predicting species distributions.

Application to real-world data

We next show two applications of our approach using real data. Code and data to duplicate these analyses are included as demonstrations within the R package.

First, we present a demonstration analysis of the ninedimensional morphological hypervolumes of two species of Darwin's finches (Box 3). A prominent hypothesis for these birds, stemming from Darwin's original observations, is that species co-occurring on the same islands would have experienced strong resource competition and character displacement, and therefore should have evolved to occupy non-overlapping regions of morphospace (Brown & Wilson, 1956). We tested this hypothesis on Isabela Island with data from the Snodgrass-Heller expedition (Snodgrass & Heller, 1904). We used log10transformed nine-dimensional data to construct hypervolumes for the five species with at least 10 complete observations. Hypervolumes were constructed using a Silverman bandwidth estimator and a quantile threshold of 0%. We found that of the possible $\binom{5}{2} = 10$ overlaps, only two species pairs had non-zero fractional overlaps: less than 1% between Geospiza fuliginosa parvula and Geospiza prosthemelas prosthemelas, and 11%



Figure 5 Hypervolume sampling analysis of simulated data, reflecting a species distribution modelling application. We assessed the ability of multiple methods - hypervolume, generalized linear model (GLM) and generalized boosted regression model (GBM) - to correctly predict sampled points as being in or out of the single $(T_{\rm C};$ Fig. 4a) or double $(T_{\rm DC};$ Fig. 4b) hypercubes. Each boxplot represents the prediction statistic calculated from 10 independent samples of *m* points from each *n*-dimensional dataset. (a) Sensitivity (true positive rate) and (b) specificity (true negative rate) statistics for the single hypercube dataset. (c) Sensitivity and (d) specificity statistics for the double hypercube dataset. For all panels, boxes that are closer to one indicate better methods.

between *Geospiza fortis fortis* and *Geospiza fortis platyrhyncha*. Thus, the original hypothesis of character displacement cannot be rejected except perhaps in two cases of weak overlap, of which one case applied to two very closely related subspecies. In contrast, a single-trait analysis would reject the hypothesis, leading to very different biological inferences. Note also that this hypothesis could be tested at the community level, comparing the union of all morphological hypervolumes of all species on each island.

Second, we present a four-dimensional analysis of the climate hypervolumes of two closely related oak species that are common in the eastern United States, Quercus alba and Quercus rubra (Box 4). We tested the hypothesis that the species occupying a larger region of climate space would also have a larger geographic range. We obtained occurrence data for each species from the BIEN database (http://bien.nceas.ucsb.edu/bien/) and transformed these to climate values using four main WorldClim layers: BIO1, mean annual temperature; BIO4, temperature seasonality; BIO12, mean annual precipitation; and BIO15, precipitation seasonality (Hijmans et al., 2005). Each climate layer was transformed relative to its global mean and SD before analysis; hypervolumes were then constructed using a Silverman bandwidth estimator and a 0% quantile threshold. We found that Q. alba had a smaller volume [0.13 standard deviations (SD⁴)] compared with Q. rubra (0.15 SD⁴). We then used the hypervolume method for a sampling application, projecting each of these climate hypervolumes into geographic space. We found that Q. alba also had a smaller range (9097 vs. 11,343 10-arcmin

pixels), supporting the original hypothesis and consistent with expert-drawn range maps (Little, 1977). Returning to the climate hyperspace, we then identified the unique (and sometimes disjoint) regions of the *Q. rubra* climate hypervolume that contributed to this larger volume. In fact, *Q. rubra* contributed more than twice as much unique volume as *Q. alba* (0.04 vs 0.02 SD⁴) to the combined hypervolumes of both species.

DISCUSSION

The future of the hypervolume

Our approach can unify several previously separate lines of ecological inquiry through direct measurement of hypervolumes. Our demonstration analyses provide preliminary evidence that this new approach can perform as well as several existing approaches, and can also enable new types of analyses.

The application of this method to species distribution modelling, while exciting, is preliminary. An advantage of our hypervolume approach is that it is conceptually simple, does not require absence or pseudo-absence data and enables hypervolume geometry to be simultaneously measured. Moreover, the method performed well in our initial tests. We therefore suggest that the method warrants further comparison with other approaches (e.g. Elith *et al.*, 2006).

The development of our method also highlights several issues that are relevant to hypervolume-related inquiry. First, there has so far been limited understanding of the properties of real

Box 3 Morphological hypervolumes of five species of Darwin's finches co-occurring on Isabela Island

Estimated nine-dimensional hypervolumes for the five species listed in the bottom left inset are shown as pair plots. Variables have original units of mm but have been log10-transformed: BodyL, body length; WingL, wing length; TailL, tail length; BeakW, beak width; BeakH, beak height; LBeakL, lower beak length; UBeakL, upper beak length; N-UBkL, nostril–upper beak length; TarsusL, tarsus length. The coloured points for each species reflect the stochastic description of each hypervolume, i.e. random points sampled from the inferred hypervolume rather than original observations. The inset shows all possible pairwise overlaps between species pairs (2 × shared volume/summed volume). Only two species pairs had non-zero overlap despite apparent overlap in each pair plot. These analyses can be replicated by running the demo 'finch' within the R package.



high-dimensional biological hypervolumes. The generality and prevalence of holes in genotypic, phenotypic and climatic hypervolumes is under-studied (Austin *et al.*, 1990; Jackson & Overpeck, 2000; Soberón & Nakamura, 2009). Our method can detect holes by calculating the difference between hypervolumes constructed with larger and smaller quantile thresholds. Additionally, little is known about how hypervolumes change over time, for example for climate niche evolution (Peterson, 1999;

Box 4 Climate hypervolumes of two oak species

Occurrence data (shown in the inset) were mapped to climate data and used to infer hypervolumes. Hypervolumes are shown as pair plots (*Quercus rubra*, red; *Quercus alba*, blue; four-dimensional climate-space overlap, purple). The coloured points for each species reflect the stochastic description of each hypervolume, i.e. random points sampled from the inferred hypervolume rather than original observations. *Quercus rubra* had a larger volume than *Q. alba*. We also projected these hypervolumes into geographic space using the inclusion test for sampling (geographic ranges are shown in inset using the same colour scheme; purple indicates geographic range overlap) and showed that *Q. rubra* also had a larger range. These analyses can be replicated by running the demo 'quercus' within the R package.



Jackson & Overpeck, 2000). With appropriate palaeo-data, our method can detect changes in hypervolume and overlap across time periods. Finally, our method is relevant to studies of community assembly, where multivariable analyses of trait hypervolume may produce fundamentally different (and more realistic) insights than single-variable analyses. Such approaches will require robust null modelling approaches (Lessard *et al.*,

2012) to compare observed hypervolume geometry and overlap with expectations under different null hypotheses. The low runtime of our methods now makes this application computationally tractable.

Data are rapidly becoming available to extend hypervolume analyses to higher dimensions. For example, global climate layers (e.g. WorldClim; Hijmans *et al.*, 2005) provide data for measure climate hypervolumes while trait databases: TRY for plants (Kattge *et al.*, 2011) or MammalDB for birds, mammals and reptiles (Baldridge *et al.*, 2012), are examples which provide data to measure functional hypervolumes, both along tens of axes. While obtaining data in other contexts can still be difficult, our tools will permit high-dimensional comparative analyses.

In sum, hypervolumes are relevant to individuals, genotypes, communities, biomes and clades, and can be constructed from a wide range of variables including climate, edaphic variables, functional traits and morphology. Although hypervolumes are a central though controversial concept in biology, they have not been adequately measured in enough real systems in high dimensions. We have provided the computational tools to make this concept operational, usable and tractable.

ACKNOWLEDGEMENTS

We thank Robert Colwell, Miguel Araújo, David Nogués-Bravo, John Harte, Simon Stump and Stuart Evans for their thoughts. B.B. was supported by a NSF predoctoral fellowship and a NSF Nordic Research Opportunity award. C.L. was supported by a University of Maine's Sustainability Solutions Initiative. C.V. was supported by a Marie Curie International Outgoing Fellowship within the 7th European Community Framework Program. B.J.E. was funded by a National Science Foundation Macrosystems award.

REFERENCES

- Abrams, P. (1980) Some comments on measuring niche overlap. *Ecology*, **61**, 44–49.
- Abrams, P. (1983) The theory of limiting similarity. *Annual Review of Ecology, Evolution, and Systematics*, **14**, 359–376.
- Albert, C.H., Thuiller, W., Yoccoz, N.G., Douzet, R., Aubert, S. & Lavorel, S. (2010) A multi-trait approach reveals the structure and the relative importance of intra- vs. interspecific variability in plant traits. *Functional Ecology*, **24**, 1192–1201.
- Araujo, M.B. & Peterson, A.T. (2012) Uses and misuses of bioclimatic envelope modeling. *Ecology*, 93, 1527–1539.
- Austin, M., Cunningham, R. & Fleming, P. (1984) New approaches to direct gradient analysis using environmental scalars and statistical curve-fitting procedures. *Vegetatio*, 55, 11–27.
- Austin, M., Nicholls, A. & Margules, C. (1990) Measurement of the realized qualitative niche: environmental niches of five *Eucalyptus* species. *Ecological Monographs*, **60**, 161–177.
- Baldridge, E., Myrhvold, N. & Morgan Ernest, S.K. (2012) Macroecological life-history trait database for birds, reptiles, and mammals. *97th ESA Annual Meeting*, Aug 2012, Portland, OR.
- Baraloto, C., Hardy, O.J., Paine, C.E.T., Dexter, K.G., Cruaud, C., Dunning, L.T., Gonzalez, M.-A., Molino, J.-F., Sabatier, D., Savolainen, V. & Chave, J. (2012) Using functional traits and phylogenetic trees to examine the assembly of tropical tree communities. *Journal of Ecology*, **100**, 690–701.

- Bonser, S.P. (2006) Form defining function: interpreting leaf functional variability in integrated plant phenotypes. *Oikos*, 114, 187–190.
- Boucher, F.C., Thuiller, W., Arnoldi, C., Albert, C.H. & Lavergne, S. (2013) Unravelling the architecture of functional variability in wild populations of *Polygonum viviparum* L. *Functional Ecology*, 27, 382–391.
- Broennimann, O., Fitzpatrick, M.C., Pearman, P.B., Petitpierre, B., Pellissier, L., Yoccoz, N.G., Thuiller, W., Fortin, M.-J., Randin, C., Zimmermann, N.E., Graham, C.H. & Guisan, A. (2012) Measuring ecological niche overlap from occurrence and spatial environmental data. *Global Ecology and Biogeography*, **21**, 481–497.
- Brown, W. & Wilson, E. (1956) Character displacement. Systematic Zoology, 5, 49–64.
- Colwell, R.K. & Futuyma, D.J. (1971) On the measurement of niche breadth and overlap. *Ecology*, **52**, 567–576.
- Colwell, R.K. & Rangel, T.F. (2009) Hutchinson's duality: the once and future niche. *Proceedings of the National Academy of Sciences USA*, **106**, 19651–19658.
- Cornwell, W.K., Schwilk, D.W. & Ackerly, D.D. (2006) A traitbased test for habitat filtering: convex hull volume. *Ecology*, 87, 1465–1471.
- Doledec, S., Chessel, D. & Gimaret-Carpentier, C. (2000) Niche separation in community analysis: a new method. *Ecology*, **81**, 2914–2927.
- Elith, J. & Leathwick, J.R. (2009) Species distribution models: ecological explanation and prediction across space and time. *Annual Review of Ecology, Evolution, and Systematics*, **40**, 677– 697.
- Elith, J., Graham, C., Anderson, R. *et al.* (2006) Novel methods improve prediction of species' distributions from occurrence data. *Ecography*, **29**, 129–151.
- Findley, J. (1973) Phenetic packing as a measure of faunal diversity. *The American Naturalist*, **107**, 580–584.
- Foote, M. (1997) The evolution of morphological diversity. Annual Review of Ecology and Systematics, 28, 129–152.
- Gavrilets, S. (1997) Evolution and speciation on holey adaptive landscapes. *Trends in Ecology and Evolution*, **12**, 307–312.
- Gavrilets, S. (2004) *Fitness landscapes and the origin of species*. Princeton University Press, Princeton, NJ.
- Gower, J. (1971) A general coefficient of similarity and some of its properties. *Biometrics*, 27, 857–871.
- Green, R.H. (1971) A multivariate statistical approach to the Hutchinsonian niche: bivalve molluscs of central Canada. *Ecology*, **52**, 544–556.
- Guo, Q., Kelly, M. & Graham, C.H. (2005) Support vector machines for predicting distribution of sudden oak death in California. *Ecological Modelling*, **182**, 75–90.
- Hijmans, R.J., Cameron, S., Parra, J., Jones, P. & Jarvis, A. (2005) Very high resolution interpolated climate surfaces for global

land areas. International Journal of Climatology, 25, 1965–1978.

- Hurlbert, S.H. (1978) The measurement of niche overlap and some relatives. *Ecology*, **59**, 67–77.
- Hutchinson, G. (1957) Concluding remarks. *Cold Spring Harbor Symposia on Quantitative Biology*, **22**, 415–427.
- Jackson, S.T. & Overpeck, J.T. (2000) Responses of plant populations and communities to environmental changes of the late Quaternary. *Paleobiology*, **26**, 194–220.
- Kattge, J., Díaz, S., Lavorel, S. *et al.* (2011) TRY: a global database of plant traits. *Global Change Biology*, **17**, 2905–2935.
- Laughlin, D.C., Joshi, C., Van Bodegom, P.M., Bastow, Z.A. & Fulé, P.Z. (2012) A predictive model of community assembly that incorporates intraspecific trait variation. *Ecology Letters*, **15**, 1291–1299.
- Lessard, J.P., Belmaker, J., Myers, J.A., Chase, J.M. & Rahbek, C. (2012) Inferring local ecological processes amid species pool influences. *Trends in Ecology and Evolution*, **27**, 600–607.
- Little, E.L., Jr (1977) *Atlas of United States trees*. US Department of Agriculture, Forest Service, Washington, DC.
- Maguire, B. (1967) A partial analysis of the niche. *The American Naturalist*, **101**, 515–526.
- May, R.M. & MacArthur, R.H. (1972) Niche overlap as a function of environmental variability. *Proceedings of the National Academy of Sciences USA*, **69**, 1109–1113.
- Maynard-Smith, J., Burian, R. & Kauffman, S. (1985) Developmental constraints and evolution. *Quarterly Review of Biology*, **60**, 265–287.
- Mouillot, D., Stubbs, W., Faure, M., Dumay, O., Tomasini, J.A., Wilson, J.B. & Chi, T.D. (2005) Niche overlap estimates based on quantitative functional traits: a new family of non-parametric indices. *Oecologia*, **145**, 345–353.
- Nix, H. (1986) A biogeographic analysis of Australian elapid snakes. *Atlas of elapid snakes of Australia* (ed. R. Longmore), pp. 4–15. Australian Government Publishing Service, Canberra.
- Pacala, S. & Roughgarden, J. (1982) The evolution of resource partitioning in a multidimensional resource space. *Theoretical Population Biology*, **22**, 127–145.
- Petchey, O.L. & Gaston, K.J. (2006) Functional diversity: back to basics and looking forward. *Ecology Letters*, **9**, 741–758.
- Petchey, O.L., Evans, K.L., Fishburn, I.S. & Gaston, K.J. (2007) Low functional diversity and no redundancy in British avian assemblages. *Journal of Animal Ecology*, **76**, 977–985.
- Peterson, A.T. (1999) Conservatism of ecological niches in evolutionary time. *Science*, **285**, 1265–1267.
- Peterson, A., Soberón, J., Pearson, R., Anderson, R., Martínez-Meyer, E., Nakamura, M. & Araújo, M.B. (2011) *Ecological niches and geographic distributions*. Princeton University Press, Princeton, NJ.
- Petitpierre, B., Kueffer, C., Broennimann, O., Randin, C., Daehler, C. & Guisan, A. (2012) Climatic niche shifts are rare among terrestrial plant invaders. *Science*, **335**, 1344–1348.
- Phillips, S.J., Dudík, M., Elith, J., Graham, C.H., Lehmann, A., Leathwick, J. & Ferrier, S. (2009) Sample selection bias and

presence-only distribution models: implications for background and pseudo-absence data. *Ecological Applications*, **19**, 181–197.

- Pigliucci, M. (2007) Finding the way in phenotypic space: the origin and maintenance of constraints on organismal form. *Annals of Botany*, **100**, 433–438.
- Raup, D. & Michelson, A. (1965) Theoretical morphology of the coiled shell. *Science*, 147, 1294–1295.
- Renner, I.W. & Warton, D.I. (2013) Equivalence of MAXENT and Poisson point process models for species distribution modeling in ecology. *Biometrics*, **69**, 274–281.
- Ricklefs, R. & Miles, D. (1994) Ecological and evolutionary inferences from morphology: an ecological perspective. *Ecological morphology: integrative organismal biology* (ed. by P.C. Wainwright and S. M. Reilly), pp. 13–41. University of Chicago Press, Chicago, IL.
- Ricklefs, R.E. & O'Rourke, K. (1975) Aspect diversity in moths: a temperate–tropical comparison. *Evolution*, **29**, 313–324.
- Ricklefs, R.E. & Travis, J. (1980) A morphological approach to the study of avian community organization. *The Auk*, **97**, 321–338.
- Rubin, D. (1996) Multiple imputation after 18+ years. *Journal of the American Statistical Association*, **91**, 473–489.
- Salazar-Ciudad, I. & Marin-Riera, M. (2013) Adaptive dynamics under development-based genotype–phenotype maps. *Nature*, **497**, 361–364.
- Shan, H., Kattge, J., Reich, P., Banerjee, A., Schrodt, F. & Reichstein, M. (2012) Gap filling in the plant kingdom – trait prediction using hierarchical probabilistic matrix factorization. *Proceedings of the 29th International Conference on Machine Learning* (ed. by J. Langford and J. Pineau), pp. 1303– 1310. Omnipress, Madison, WI.
- Silverman, B.W. (1992) *Density estimation for statistics and data analysis*. Chapman & Hall, London.
- Snodgrass, R. & Heller, E. (1904) Papers from the Hopkins–Stanford Galapagos expedition, 1898–99 XVI. Birds. *Proceedings of the Washington Academy of Science*, 5, 231–372.
- Soberón, J. & Nakamura, M. (2009) Niches and distributional areas: concepts, methods, and assumptions. *Proceedings of the National Academy of Sciences USA*, **106**, 19644–19650.
- Tilman, D. (1982) *Resource competition and community structure.* Princeton University Press, Princeton.
- Tilman, D., Knops, J., Wedin, D., Reich, P., Ritchie, M. & Siemann, E. (1997) The influence of functional diversity and composition on ecosystem processes. *Science*, **277**, 1300–1302.
- Van Valen, L. (1974) Multivariate structural statistics in natural history. *Journal of Theoretical Biology*, **45**, 235–247.
- Villéger, S., Mason, N.W.H. & Mouillot, D. (2008) New multidimensional functional diversity indices for a multi-faceted framework in functional ecology. *Ecology*, **89**, 2290–2301.
- Violle, C. & Jiang, L. (2009) Towards a trait-based quantification of species niche. *Journal of Plant Ecology*, **2**, 87–93.

- Warren, D.L., Glor, R.E. & Turelli, M. (2008) Environmental niche equivalency versus conservatism: quantitative approaches to niche evolution. *Evolution*, **62**, 2868–2883.
- Westoby, M., Falster, D.S., Moles, A.T., Vesk, P.A. & Wright, I.J. (2002) Plant ecological strategies: some leading dimensions of variation between species. *Annual Review of Ecology and Systematics*, **33**, 125–159.
- Whittaker, R. & Niering, W. (1965) Vegetation of the Santa Catalina Mountains, Arizona: a gradient analysis of the south slope. *Ecology*, **46**, 429–452.
- Wisz, M.S., Hijmans, R.J., Li, J., Peterson, A.T., Graham, C.H. & Guisan, A. (2008) Effects of sample size on the performance of species distribution models. *Diversity and Distributions*, 14, 763–773.
- Wright, S. (1932) The roles of mutation, inbreeding, crossbreeding, and selection in evolution. *Proceedings of the Sixth International Congress on Genetics* (ed. by D.F. Jones), pp. 355–366. Brooklyn Botanic Garden, Menasha, WI.

BIOSKETCH

Benjamin Blonder is a PhD student. He is interested in science education and the response of communities and ecosystems to changing climate.

Editor: José Alexandre Diniz-Filho