

Hypervolume concepts in niche- and trait-based ecology

Benjamin Blonder

B. Blonder (<http://orcid.org/0000-0002-5061-2385>) (bblonder@gmail.com), Environmental Change Inst., School of Geography and the Environment, Univ. of Oxford, Oxford, United Kingdom.

Hutchinson's n -dimensional hypervolume concept for the interpretation of niches as geometric shapes has provided a foundation for research across different fields of ecology and evolution. There is now an expanding set of applications for hypervolume concepts, as well as a growing set of statistical methods available to operationalize this concept with data. The concept has been applied to environmental, resource, functional trait, and morphometric axes and to different scales, i.e. from individuals, species, to communities and clades. Further, these shapes have been variously interpreted as niches, ecological or evolutionary strategy spaces, or proxies for community structure. This paper highlights these applications' shared mathematical framework, surveys uses of the hypervolume concept across fields, discusses key limitations and assumptions of hypervolume concepts in general, provides a critical guide to available statistical estimation methods, and delineates the situations where hypervolume concepts can be useful.

The n -dimensional hypervolume was originally proposed by Hutchinson (1957) to describe the fundamental niche of a species. These hypervolumes exist within a space defined by a set of n independent axes. The hypervolume is then defined as a subset of the space, i.e. an n -dimensional geometrical shape (Fig. 1A). Hutchinson suggested that the axes correspond to requirements of the species (e.g. food size, temperature) and that the boundaries of the shape indicate the conditions that permit the growth and reproduction of the species. The concept can be found illustrated in biology textbooks (Begon et al. 2006) and has gained wide usage as well as extensive discussion (Chase and Leibold 2003, Colwell and Rangel 2009, Holt 2009, Peterson et al. 2011).

Because this geometrical concept is apparently simple and easily explained, it has been applied to other contexts where the shape does not represent a niche and the axes do not represent limiting conditions. That is, hypervolumes can be defined using many types of axes (limiting resource, competition parameters, climate, resource, functional trait) and for many systems (individuals, populations, species, communities, clades, regions) (Table 1). The hypervolume concept has therefore come to inspire a range of other research areas throughout ecology and evolutionary biology (Peters 1991, McInerny and Etienne 2012, Garnier et al. 2016), e.g. in niche-based ecology, functional ecology, community phylogenetics, coexistence theory, and morphometrics.

This paper surveys the scope of uses of the hypervolume concept across ecology and evolution, identifies the common limitations and opportunities across all these uses, and provides guidelines on how best to use hypervolumes subject to these limitations. In this process it explores how

the hypervolume concept has been operationalized, synthesizes recent efforts for making the concept measurable, and delineates the contexts in which hypervolumes are useful, and when they are not.

A modern definition of the hypervolume

It is useful to begin with a general mathematical definition that can unify these different usages (Fig. 1B). A hypervolume function can be defined as a function $h(x)$ mapping from an n -dimensional Euclidean space to a 1-dimensional Euclidean space, as $h(x) : \mathbb{R}^n \rightarrow \mathbb{R}$. The domain of this function may be restricted to a finite region $X \subseteq \mathbb{R}^n$, and corresponds to the possible values of axis combinations. The range of this function is always limited to non-negative values, and corresponds to possible biological states (e.g. absent/present, rare/prevalent, improbable/probable). It is also assumed that the hypervolume function has a finite integral H , i.e. $\int_X h(x) dx = H < \infty$. This allows the hypervolume function to be considered as a probability density function (p.d.f.) by normalizing the hypervolume function to $h(x)/H$. Data samples can be used to estimate $h(x)$, but these data do not in themselves define the hypervolume function.

The hypervolume function can then be summarized by several metrics. First, a metric for a 'boundary' can be defined as a level set q of $h(x)$, defining points x that satisfy $h(x) = q$ (Fig. 1C, D). Second, a metric for a 'volume' can be defined as the size of the region $X_q \subset X$ enclosed by the boundary, i.e. a value $|X_q| = \int_{X_q} dx$ such that for a value q , $\forall x \in X_q, h(x) \geq q$ (Fig. 1E). The region X_q itself defines

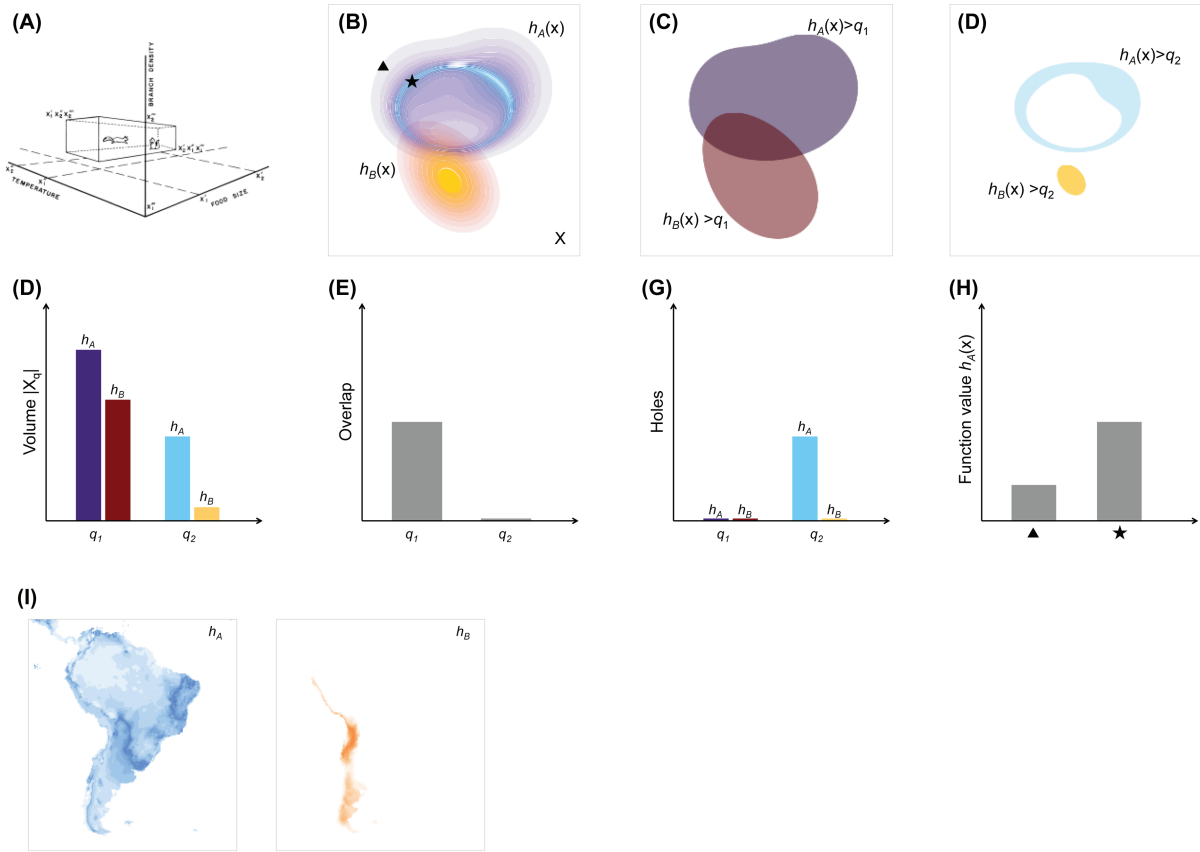


Figure 1. Comparison of original and modern hypervolume concepts. A) Hutchinson's original concept, as illustrated for the fundamental niche of a squirrel. Redrawn from Hutchinson (1978) with permission. The hypervolume is depicted as a rectangular geometrical region. B) A modern understanding of the hypervolume function – here, two example functions $h_A(x)$ and $h_B(x)$ – allows for more complex geometries and variation in function value (shown with higher probability regions in brighter colors) over a region of hyperspace X . C) A threshold value – here $q = q_1$ – can be used to delineate a contour boundary and shape for each hypervolume function. D) Changing the threshold value (here, $q = q_2$) yields a different contour boundary and shape for each hypervolume function. E) a metric of the volume can be calculated as the region enclosed by each contour boundary, $|X_q|$. F) A metric of overlap can be calculated as the volume of the region of intersection between $h(x)$ functions for a certain threshold value provides. G) The presence, size, and shape of holes varies with threshold value. F) The value of $h(x)$ can be evaluated at different points (star and triangle in panel B) and used as a metric of probability of suitability. I) Geographic maps can be obtained for niche applications by projecting discretized values of $h(x)$ from the n -dimensional space to spatial pixels.

a geometric shape that corresponds to Hutchinson's original concept of a hypervolume, and defines a uniform probability

$$\text{density function, } h_q(x) = \begin{cases} x \in X_q, 1/|X_q| \\ x \notin X_q, 0 \end{cases}.$$

There are two general cases for $h(x)$. First, if $h(x)$ is binary and obtains only the values 0 or 1 (e.g. 'absent' or 'present'), then there is only one possible non-zero value of q , and thus the shape of the hypervolume function is specified as a uniform probability density function $h_{q=1}(x)$. The region $X_{q=1}$ can be interpreted directly as an n -dimensional shape, corresponding exactly to Hutchinson's original idea. This approach effectively assumes that all points in the hypervolume are equally likely or suitable, and allows the use of geometrical concepts to describe the hypervolume.

Second, if $h(x)$ is not binary, it may obtain a continuous range of values, there are many possible values of q , and $h(x)$ is specified as an arbitrary probability distribution function. As a result, there are multiple possible boundaries and volumes and shapes that depend on the value of q that is

chosen. This interpretation allows for variation in the likelihood of points within the hypervolume, permits the use of probabilistic theory and concepts to describe a hypervolume (Blonder et al. 2014, Carmona et al. 2016a). Geometrical concepts can also be used to describe the hypervolume once a particular value of q is specified.

Applications of the hypervolume concept

Hypervolume concepts underlie many areas of research. Hutchinson's inspiration for the geometrical hypervolume concept likely comes from the work of Grinnell, who originally described niches in terms of a set of multiple limiting factors at different points on the landscape (Grinnell 1917). The central insight of Hutchinson (1957) was to transfer these niche concepts from an explicitly geographical context to an abstract geometrical context (Colwell and Rangel 2009). This definition led to a common conceptualization of hypervolumes as geometrical

Table 1. Examples of the diversity of research questions that have been addressed using hypervolume concepts across scales and axes types, drawing from recent literature.

Axis type	Scale	Hypervolume term	Research topic	Example reference(s)
Climate	Species	Niche	How do realized niches shift under biological invasions?	Broennimann et al. (2007), Tingley et al. (2014)
Climate	Community	Niche	Do communities comprising species with narrow realized niches indicate stronger environmental filtering?	Blonder et al. (2015)
Climate	Clade	Niche	How has long-term climate change shifted the fundamental niches of related species?	Evans et al. (2009)
Climate	Region	Niche	How has short-term climate change limited the overall climate space and realized niches of species over time?	Jackson and Overpeck (2000)
Functional trait	Species	Functional diversity	When is intraspecific trait variation larger than interspecific trait variation, and why?	Hulshof and Swenson (2010)
Functional trait	Community	Functional diversity	Do communities' functional richness and overlap indicate different assembly processes?	Cornwell et al. (2006), Lamanna et al. (2014)
Functional trait	Clade	Functional diversity	What is the overall ecological strategy space of different clades?	Díaz et al. (2016)
Functional trait	Region	Functional diversity	How do climate gradients explain patterns of functional richness across landscapes?	Šímová et al. (2015), Swenson et al. (2012)
Allele frequency	Population	Fitness landscape	How difficult is evolution from one complex phenotype to another?	Gavrilets (1997)
Morphological trait	Species	Phenotype space	What are the biophysical limits to possible phenotypes?	Raup and Michelson (1965)
Morphological trait	Clade	Phenotype space	Why do some clades diversify into larger morphospaces than others?	Sidlauskas (2008)

shapes that in turn provided inspiration for many other applications.

Niche modeling

Hutchison's fundamental niche concept is relevant to the geographic distribution of species, but requires some extension. Species may occur in subsets or supersets of the fundamental niche because of variation in the conditions necessary for regeneration, establishment, or persistence niches (Grubb 1977). Moreover, dispersal limitation and biotic interactions may limit species to realized niches that are subsets of the fundamental niche (Hutchinson 1957, Pulliam 2000). Similarly, these realized niches are constrained by potential niches, which were developed to describe the set of possible conditions X available at any point in time (Jackson and Overpeck 2000). All of these processes may change the form of $b(x)$, but fundamentally remain rooted in a hypervolume context.

Hypervolume concepts have also now been used to describe potential niches and the envelopes of climate space that constrain realized niches (Jackson and Overpeck 2000, Soberón and Nakamura 2009), as well as establishment and persistence niches that take into account different demographic and dispersal processes (Araújo and Guisan 2006), with consensus converging on the utility of these approaches being for modeling the suitability rather than actual distributions of species (Holt 2009).

The modeling of realized climate niches has been an immediate area of growth for the hypervolume concept that stemmed from early observations that species distributions often occur along limited subsets of environmental or elevation gradients (von Humboldt 1807, Shreve 1911, Whittaker and Niering 1965). Early realized niche models

were probabilistic and constructed using order statistics or distribution moments (for birds, Van Valen 1965, and for protozoa; Maguire 1967). Discriminant analyses were also then used to describe hypervolumes of mollusks (Green 1971). Further developments likely were an outgrowth of growing interest in gradient analyses (Whittaker 1967) and ordination (Austin 1985, ter Braak and Prentice 1988).

Modern probabilistic realized niche and species distribution modeling applications can be re-cast in terms of these hypervolume concepts (Franklin 2010, Peterson et al. 2011). The first modern probabilistic and geometrical hypervolume model was probably developed for realized niches of *Eucalyptus* species (Austin et al. 1990). Subsequently, BIOCLIM, a geometrical approach (Booth et al. 2014, Busby 1991), was developed to model realized niches. These methods have inspired the development of a larger set of modern correlative approaches (see reviews by Elith et al. 2006, Peterson et al. 2011).

There also have been some experimental measurements of fundamental niches by exploring factorial combinations of axes values. The hypervolume then is typically considered in terms of a number of range constraints (Hutchinson 1978), e.g. air temperature above a lethal minimum and below a lethal maximum (Araújo et al. 2013). The earliest measurements actually predated Hutchinson's work and were carried out for a small set of axes for the fundamental niches of two beetle species (Birch 1953). Many other experimental efforts have followed, usually restricted to a small set of niche axes for practical reasons (Colwell and Fuentes 1975). For example, fundamental niche hypervolumes have been estimated with this approach for *Daphnia magna* (Hooper et al. 2008).

Process-based niche concepts and species distribution models have been developed that begin to move beyond hypervolumes, though some still have hypervolume interpretations

(Pulliam 2000). This includes models that incorporate spatial effects (Phillips et al. 2006, Renner and Warton 2013), and those that incorporate demographic processes (Hooten and Wikle 2008, Cabral and Schurr 2010, Soberón 2010, Nenzén et al. 2012), dispersal limitation (Hooten and Wikle 2008, Engler and Guisan 2009), and species interactions (Pollock et al. 2014, Mod et al. 2015, Anderson 2017). A recent comparison of some of these approaches is available in Zurell et al. (2016). Dynamic energy budget models for fundamental niches also move beyond the limits of the hypervolume concept, because possible density dependence among different life stages, e.g. as in models for herpetofauna (Kearney and Porter 2004, Porter and Kearney 2009) and corals (Hoogenboom and Connolly 2009). Similarly, general demographic / population modeling approaches for defining realized niches may incorporate density-dependent effects of species on themselves and on others that do not readily correspond to Euclidean niche axes (Chesson 2000). Additionally, multidimensional overlap metrics for $h(x)$ may not be predictive of interspecific competition, because competition may depend only on single niche axes or instead on biotic interactions (e.g. parasites) not described by the niche space (Holt 1987).

The zero net growth isocline approach (Tilman 1977, Chase and Leibold 2003) provides a hybrid framework that retains some of the geometrical and population dynamic aspects of niches and may continue to be valuable for bridging these divides. In this approach, species occur in environments where they draw down resources in ways that render them competitively superior or able to coexist with others, and are absent from those in which resources are insufficient (Godsoe and Harmon 2012). Nevertheless, reconciling environmental axes (which do not change with biotic context) and resource axes (which do) in a unified framework remains challenging. It is not clear that geometric or probabilistic concepts should or could map onto realized niches defined by the outcomes of species interactions. There is a clear need for more careful process-based niche models.

Functional and community ecology

Another prominent usage of geometrical hypervolume-based fundamental niches has been in for species coexistence and community assembly. MacArthur and Levins (1967) pioneered an influential model for community assembly based on the competitive exclusion principle (Gause 1934). Their model yielded a prediction of limiting similarity, in which species with fundamental niches that were too close together could not coexist. The main consequence of this model was to make prominent the idea of geometric ‘niche packing’ of species as a conceptual approach for understanding species coexistence and richness limits (May and Arthur 1972, Rappoldt and Hogeweg 1980, Abrams 1983). This in turn inspired a controversial body of work on the importance of competition and limiting similarity in structuring species assemblages (Diamond 1975, Connor and Simberloff 1979, Diamond and Gilpin 1982). This work has also resulted in various tests using character displacement as a proxy for niche differentiation (Stubbs and Bastow Wilson 2004, Huntley et al. 2008, Cornwell and Ackerly 2009, Götzenberger et al. 2012).

In functional ecology, traits are often taken proxies for niche axes (Kraft et al. 2008, Swenson et al. 2012, Lamanna et al. 2014) and used to calculate distributions and hypervolumes whose statistical properties are used to infer community assembly processes. In this context, community-scale trait hypervolumes that are loosely packed together or that have large sizes are thought to indicate competitive processes that displace species in trait/niche space (Weiher and Keddy 1995, Weiher et al. 1998, Kraft et al. 2007, Lamanna et al. 2014). Functional trait hypervolumes have also been used to predict and explain species invasions via vacant niche concepts (Herbold and Moyle 1986, Moles et al. 2008, Van Kleunen et al. 2010), on the assumption that species with more similar hypervolumes are more likely to experience competition with resident species (Weiher et al. 1998, Swenson and Enquist 2009). These perspectives have all received extensive criticism because of the weak, or generally untested, linkage between traits and niche differentiation and trait space and niche space (Kunstler et al. 2012, Kraft et al. 2015, D’Andrea and Ostling 2016, Shipley et al. 2016). There is also uncertainty regarding the multiple assembly processes that could yield the same phylogenetic or trait pattern. It was long thought that differences among species tend to promote coexistence and reduce the effect of competition. It is now clear that the reverse is often true, that some differences among species can make competition strong and coexistence difficult (Mayfield and Levine 2010). Coexistence theory that does not rely on hypervolume niche concepts (Chesson 2000, Letten et al. 2017) may ultimately be more successful.

Nevertheless, geometrical hypervolume concepts have seen wide successful use in functional ecology for primarily descriptive applications. Much of the motivation for this work has roots in Hutchinson’s niche hypervolume-based exploration of the limits to species richness and to the idea that species need to be, to a certain degree, functionally different to coexist (Hutchinson 1959). Axes are often chosen to represent strategy axes for the taxa of interest (McGill et al. 2006, Westoby and Wright 2006), e.g. the fast-slow tradeoffs proposed for plants (Grime 1977, Westoby 1998). After an initial development of functional diversity indices applicable to single trait axes (Mason et al. 2005, Lepš et al. 2006), hypervolume approaches were pioneered to extend these concepts to multiple dimensions (Villéger et al. 2008, Mouchet et al. 2010, Carmona et al. 2016a). Such trait-based hypervolumes have been used to describe biogeographical (Swenson et al. 2012, Lamanna et al. 2014, Šimová et al. 2015), fire (Pausas and Verdú 2008), successional (Loranger et al. 2016), and land use (Laliberté et al. 2010) gradients in functional diversity. They have also been used to summarize the occupancy and overlap of ecological strategy spaces, e.g. for plant phenotypes (Díaz et al. 2016) and for domestication (Milla et al. 2015), and to characterize trait distributions for dynamic global vegetation models (van Bodegom et al. 2014).

Evolutionary applications

Separately, the geometrical hypervolume concept has been used as a tool to address evolutionary and phylogenetic

related questions. Hypervolume-based niche concepts were also translated to applications in phylogenetic community ecology. In this context, under a further assumption of phylogenetic niche conservatism (Wiens and Graham 2005), more related species should have more similar or closely packed hypervolumes, in term of fundamental niches and functional traits, and therefore be less likely to co-occur if competition is a major factor structuring communities (Webb et al. 2002). In practice, however, differences between species in terms of niches, traits and phylogeny often do not follow the same patterns (Mason and Pavoine 2013).

Similarly, hypervolume concepts have been used to describe phenotypic morphospaces (Raup and Michelson 1965, Warheit et al. 1999, Wainwright et al. 2004, Stayton 2005). They were then also used to compare phenotypic evolution within and across clades (Gould 1991, Shen et al. 2008, Díaz et al. 2016). The concept of a fitness landscape defined over multiple allele frequency or phenotypic axes, while independently developed (Wright 1932), is clearly similar to the hypervolume concept in that it defines a non-negative function $h(x)$ with geometrical interpretation, but with values corresponding to fitness rather than probability. Hutchinson's concepts also have indirectly inspired a body of work on phenotypic evolution on high-dimensional fitness landscapes, which may potentially include hypervolume-like geometrical features such as holes (Gavrilets 1997, Laughlin and Messier 2015). Nevertheless, the linkages between Hutchinsonian hypervolumes and fitness landscapes still remain incompletely explored.

Other applications

Hypervolumes also underlie work in paleoecology and climate change science. The realized niches of species are limited by the availability of climate conditions at any moment in time. As the climate changes, so do the conditions available for life, effectively leading to 'clipping' of the edges of realized niche hypervolumes (Austin et al. 1990). This perspective has led to a range of probabilistic and geometrical hypervolume models for understanding realized niche shifts in response to climate shifts (Jackson and Overpeck 2000) or invasion of novel habitats (Broennimann et al. 2007, Tingley et al. 2014).

Last, geometrical hypervolume concepts have also recently been used to understand ecosystem/community state stability (Barros et al. 2016). In this framework a system can be defined in terms of multiple state variables (e.g. the abundance of different species in a community, or the weighted means of different trait values) and variation in these state variables over time can be quantified. Temporal variation in these state variables can be used to delineate a hypervolume, and is closely linked to concepts of ecosystem stability (Beisner et al. 2003), where smaller volumes represent more stable systems and shifts in these volumes represent transitions to alternate states.

Current challenges

Despite Hutchinson's concept being developed more than half a century ago, many of these applications have only been developed over the past decade. Part of this is probably due to

shifts in opinion about the utility of the concept for realized or fundamental niche modeling (Holt 2009, McNerny and Etiene 2012), which in turn has shaped support for other applications of the concept. However statistical estimation of hypervolume functions from data requires computational resources that may not have been available until relatively recently, limiting past interest in applications. Additionally, only recent have informatics initiatives made it possible to compile and share the large datasets needed to estimate trait (Jones et al. 2009, Kattge et al. 2011) or niche (Soberón and Peterson 2004, Enquist et al. 2016) hypervolumes, though data limitations remain important especially for resource axes (Miller et al. 2005). As a result, interest in hypervolumes and applications of the concept are rapidly growing. We summarize some key recent studies across research areas in Table 1, highlighting variation in underlying research questions, axis definitions, and scale of application.

Mathematical operations using hypervolumes

The overall geometry and boundaries (limits/constraints) of $h(x)$ can be quantified using geometrical metrics (Fig. 1C, D). For example the elongation of the hypervolume can be determined via the ratio of its longest principal axis length to its shortest principal axis length, and its compactness via the ratio of its volume to the volume of the minimum convex hull enclosing the hypervolume. Such metrics can provide insight into constraints on the processes generating the hypervolume, e.g. developmental constraints on body plans in morphospaces, strength of environmental filters in climate niches, resource drawdowns for zero net growth isoclines, or correlated selection in functional spaces.

The volume can be measured by finding the area enclosed by contours of $h(x)$ (Fig. 1F). This operation allows estimation of niche size and breadth for niche axes, or of functional richness for traits axes.

The overlap can be measured between two hypervolumes $h_1(x)$ and $h_2(x)$ by defining a function $g(h_1(x), h_2(x))$ that also satisfies the definition of a hypervolume function (for example, $g(h_1(x), h_2(x)) = h_1(x) \times h_2(x)$) (Fig. 1G). Additionally, it is possible to calculate a range of symmetric or asymmetric similarity indices based on volume ratios, e.g. Jaccard, Sørensen, unique fraction indices (Blonder et al. 2014), overlap index (Junker et al. 2016)), or for niches, resampled geographic species distributions (Broennimann et al. 2012, Godsoe and Case 2015). These overlaps can be useful for studies of niche overlap and phylogenetic niche conservatism, and can be weighted for community-scale hypervolumes to account for variation in species abundance (Colwell and Fuentes 1975, Carmona et al. 2016a, Blonder et al. 2017).

Geometrical holes in $h(x)$ can be detected (Fig. 1H). These holes are defined as regions within X that have low/zero values of $h(x)$ and represent unrealized combinations of axes. In functional spaces, they can indicate a range of biological processes, e.g. forbidden strategies or invasion opportunities in functional spaces, allowing determination of whether invader species are similar to resident species. For realized niches, they can determine whether novel climate

conditions may be suitable for a species, or highlight the effect of biotic interactions, dispersal limitation, or climate availability. For fitness landscapes, they can indicate low-fitness valleys (Gavrilets 1997). Concepts surrounding holes have only recently been developed (Blonder 2016) and is unclear yet whether holes are found in empirical hypervolumes for any axis combinations or scales.

The value of $h(x)$ can be estimated at a given point in the hyperspace (Fig. 1H). For realized niche applications, this is the process underlying geographic projection in many correlative species distribution models (Peterson et al. 2011) (Fig. 1I). In this approach, a set of geographic points in \mathbb{R}^2 are transformed into a Euclidean environmental space in \mathbb{R}^n , by using georeferenced layers corresponding to different variables. A hypervolume function $h(x)$ is estimated based on the species' presence and possibly absence in this hyperspace. Suitability maps are then obtained by projecting discretized geographic pixels within the mapped region from \mathbb{R}^2 into \mathbb{R}^n using the georeferenced layers, and using these to calculate values of $h(x)$. These values can then be interpreted as suitabilities or presence probabilities depending on the data used to calibrate the model, acknowledging that non-niche factors may further limit the actual distribution of the species (Guisan and Zimmermann 2000, Godsoe 2010), and that this probability may be conditioned on unmeasured variables due to sampling bias (Soberón and Nakamura 2009).

For functional applications, estimates of $h(x)$ in a probabilistic context are also useful for understanding functional diversity patterns (Carmona et al. 2016a). Delineating values of $h(x)$ across the hyperspace X can be used to estimate metrics of functional evenness, dispersion, and redundancy are available (Villéger et al. 2008, Mouillot et al. 2014, Ricotta et al. 2016).

Mathematical assumptions of hypervolume approaches

There is a shared set of fundamental assumptions relevant to the definition of a hypervolume that arise from the mathematics, regardless of the identity of the axes or the scale of the system being considered. These assumptions highlight conceptual limitations that are important for determining the relevance and applicability of hypervolumes for any particular question.

The hypervolume assumes that the n axes defining the system are fixed: to compare hypervolumes, each must have the same number and identity of axes, for the same reason the volume of a square and a cube are not comparable, nor are a unit of time and a unit of length. The definition must be valid across all the units of interest, e.g. different spatial locations, different time points, different species. However the axes defining one species' niche may not be a good description for another species, and likewise for the traits used to characterize the functional diversity of one community or another. While in some cases this problem can be resolved by adding more axes, in other cases it may not be possible. For example, soil electrical conductivity may be an appropriate trait axis for all species of plants, but inappropriate when applied to all species of birds. This problem becomes even

less tractable if the identity of axes depends on the biotic context for the system, e.g. if temperature is a niche axis for an animal species at high population densities but not at low densities (Holt 2009). In these cases the geometric conception of a hypervolume can become less useful.

The hypervolume axes must comprise continuous variables. This is necessary in order to work within the definition of a Euclidean hyperspace, and to be able to calculate overlaps and volume integrals. While ordered factors can be used by representing them as integer codes (e.g. 'low/medium/high temperature' becomes 1,2,3), their inclusion is not recommended and makes it difficult to make precise definitions of the size of a hypervolume. Typically in these cases, ordered factors would be ordinated to continuous via a Gower transformation (Gower 1971). Alternatively it is possible to construct a notion of volume as a fractional occupancy of each factor level (Whittaker et al. 2014). Regardless, the hypervolume concept cannot be used with unordered factors, e.g. ('site A/B/C') where distances between values are not well defined.

The hypervolume also assumes that all axes contribute independently and equally to Euclidean distances and volumes. It also means that even if one axis contributes more to a biological performance variable (e.g. an ecosystem flux or organismal fitness) it is given equal weight in the hypervolume. This issue can be partially circumvented by scaling values along individual axes.

Hypervolume also must have axes with comparable units, because distances and volumes require summation and summation can only proceed on variables with the same units. Typically axes with units are log-transformed or z -transformed to unitless coordinates before use in an analysis so that volumes or overlaps can be defined (Lamanna et al. 2014, Carmona et al. 2016b). Log-transformation is transferable to new datasets, while z -transformation depends on the data. Regardless, in species distribution models where hypervolume models are not examined for their own properties but used instead to make predictions of species presence in geographical landscapes, describing volumes or distances is not a relevant question, and variables with different units can be treated with their units implicitly considered as dimensionless.

Conceptual limitations of hypervolumes

A central issue is whether a static geometrical model is sufficient to describe the system of interest. For realized niche hypervolumes, this issue translates to a question of whether intraspecific or interspecific density dependence is important. Such biotic context dependence would mean the hypervolume would change depending on the system's population dynamics in ways not described by adding additional axes to the hypervolume, e.g. if different life stages respond differently to density (Holt 2009). For functional hypervolumes, the same issue holds: when axes identity or number is context-dependent, a geometrical model becomes insufficient. These issues suggest that concepts should be revised to consider the effect of species on their environment in addition to their responses or traits in their environment (Chase and Leibold 2003, D'Andrea and Ostling 2016). More strongly,

this perspective suggests that any dynamical questions involving hypervolumes, e.g. coexistence, invasion, or competition, may require more sophisticated models that fundamentally cannot be addressed with hypervolumes. For example, the growing body of research on demographic models based on niches, traits, and species interactions (reviewed by Soberón 2010, Zurell et al. 2016, Anderson 2017) suggests that such extensions will be important.

When hypervolumes are used to quantify functional diversity, there are two tacit and often unsupported assumptions made. The first assumption is that the measured trait axes are proxies for species fitness or niche axes (D’Andrea and Ostling 2016, Shipley et al. 2016), a link that remains little explored empirically (but see Adler et al. 2013, Kraft et al. 2015, Kunstler et al. 2016). While verbal arguments can be made that certain traits represent different tradeoffs (Westoby and Wright 2006), there is no reason to necessarily assume that the mapping between traits and niche axes is linear, or that the dimensionality of the measured trait space and the dimensionality of the true niche space are equivalent or even comparable. Indeed, recent process-based models have shown that fundamental niches often emerge as nonlinear functions of several measurable trait and environment variables (Hoogenboom and Connolly 2009, Porter and Kearney 2009). The third assumption is that the multidimensional distance between trait variables is a proxy for niche differentiation and therefore competition. Neither of these proxy relationships are necessarily valid. Empirical studies have shown that trait differences for commonly measured functional traits are more likely to be related to fitness differences than niche differences (Kraft et al. 2015, Funk and Wolf 2016, Kunstler et al. 2016), and models have demonstrated that lower niche differentiation does not always lead to higher competition (Scheffer and van Nes 2006). These findings cast doubt onto the robustness of using trait axes as proxies for niche axes in phylogenetic and functional community ecology.

The statistical estimation problem

In addition to these conceptual issues, it is also useful to explore the practical issues around statistical estimation of the hypervolume function, $h(x)$, from data. Suppose that the investigator has an $m \times n$ matrix of data, corresponding to m observations of n variables representing unique axes. These data are either implicitly or explicitly associated with a paired $m \times 1$ vector of biological state values. If all of the observations are positive (‘present’), then the state vector has one class and always takes value 1; if some of the observations are positive (‘present’) and others are negative (‘absent’), then the state vector has two classes and takes values 1 or 0; if the observations are quantitative (e.g. ‘low abundance’, ‘high abundance’), then the state vector has multiple or infinite classes (Fig. 2).

The statistical problem is then to find a method that can estimate $h(x)$ based on the observations and their state vectors. It is often assumed that the sampled data are a random sample of a true distribution over the hyperspace X , i.e. independent and identically distributed, though this may not always be realistic because of sampling bias and coverage (Soberón and Nakamura 2009). For geometrical models without conceptual interpretations as probability distributions (e.g. range box models), sampling processes have rarely been considered (Junker et al. 2016) and are discussed further below.

The properties of the state vector determine the types of methods that are available to estimate $h(x)$. If the state vector is one-class, then methods must predict presence and absence with knowledge only of presences. This can lead to large biases in probabilistic models, and/or inadvertent conditioning of probability distributions on certain sets of conditions, e.g. available environments or species presence (Soberón and Nakamura 2009). Alternatively, if the state vector is two-class or quantitative, then the model must predict values of $h(x)$ directly from observed values of $h(x)$. Two-class and

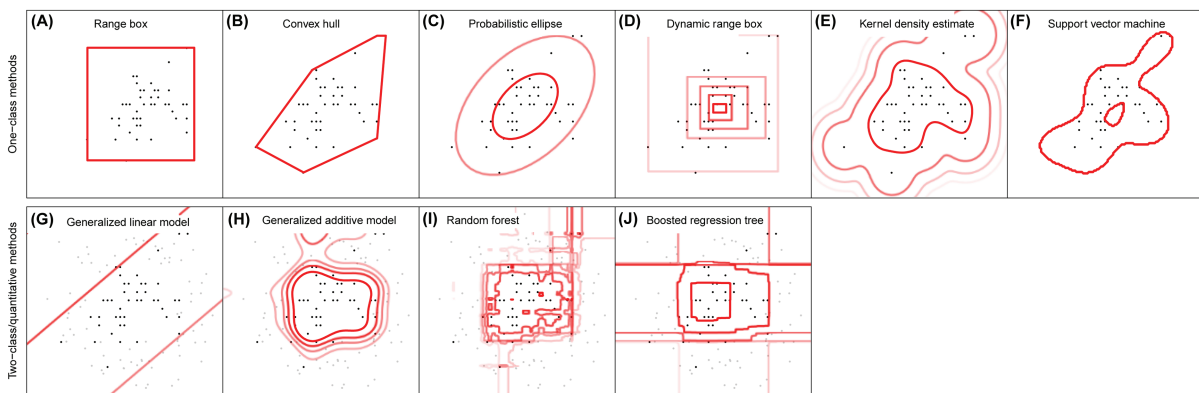


Figure 2. Comparison of hypervolume estimation methods for one-class (A–F) and two-class (G–J) data. Where models yield one-class or two-class (0–1 extremes only) values for $h(x)$, a single red line delineating class membership is drawn; where models yield quantitative (continuous) output for $h(x)$, several values are drawn at different shading levels. Presence data points are shown as closed black circles. For two-class methods, absence data are shown as open gray circles. A) Range boxes enclose the data within each axis independently. B) Convex hulls find the minimum set of linear inequality constraints that enclose the data. C) Probabilistic ellipses fit a multivariate normal distribution to the data. D) Dynamic range boxes fit a range box to each quantile of the data. E) Kernel density estimates fit a probability density function to the data. F) Support-vector machines find a high-dimensional classifier for the data. G) Generalized linear models fit a linear regression surface to the data. H) Generalized additive models use a set of splines to fit a nonlinear regression surface to the data. I) Random forests simultaneously use multiple regression trees to vote on the data. J) Boosted regression trees iteratively improve the performance of regression trees by fitting new trees to the model residuals.

quantitative methods tend to be more accurate, because they are given more complete input data. However two-class and quantitative data is not readily obtainable in all contexts: for example, in measurements of functional diversity, it is unclear which trait values should be considered truly absent from a community when the trait values measured to be present are surely only a limited sample of a true distribution. Additionally, two-class methods focused on discriminating classes rather than estimating boundaries may suffer from biases in niche modeling due to the structure of absence data (Drake and Beier 2014, Maher et al. 2014).

Historically, functional hypervolumes have been estimated with one-class methods, while realized niches have been estimated with one-class, two-class, and quantitative methods. Recognizing the common underlying statistical problem may be useful for synthesizing or transferring hypervolume estimation methods across disciplines. For example, one solution that has been used for realized niche data but not trait data is to simulate absence data by drawing at random from a background distribution, e.g. all the available environmental conditions. A similar approach could potentially be used for trait data in order to use two-class and quantitative methods where one-class methods have previously been used.

Choosing an appropriate statistical method depends on the goals of the investigator and the limits of the available data. A comparison of leading methods is provided in the Supplementary material Appendix 1 Text A1, Table A1. Detailed information on currently available software packages is provided in Supplementary material Appendix 1 Table A2. To assist with choice of method for hypervolume estimation, Supplementary material Appendix 1 Fig. A1 also provides a decision support flowchart.

The structure of the available data will restrict the choice of method. The most critical question is whether there is sufficient data available to estimate a geometrical shape in the proposed dimensionality space. If not, none of the hypervolume methods should be used. If there is, then the investigator must choose to use either two-class or one-class methods depending on the structure of the data. Then the investigator should identify any additional assumptions about the structure of the data. For example, if it is assumed that the data represent a fundamental niche defined by simple independent limits on each axis, then fitting a range box model may be appropriate; alternatively if it is assumed that the data represent a realized functional space with complex shape, then fitting a kernel density estimate may be appropriate. This choice is beyond determination by any quantitative statistical comparison and must be determined by the investigator.

Data properties may further guide the investigator's method choice. Some methods are robust to outlying points, while others are not. If such outlying points should be given high leverage (e.g. real and extreme observations of a species' realized niche that indicate physiological tolerances) then methods that are not robust to outliers should be used; alternatively, if outlying points indicate measurement error (e.g. measurements of body size that are unrealistically high but recorded in a dataset), then methods that are robust to outliers should be used instead. Additionally, the occurrence of data with missing values can also guide choices. Most methods require complete data cases in order to locate points in the n -dimensional space. However some methods

that do not include interactions between variables (e.g. range boxes) are capable of working with incomplete cases. This is typically less of a concern with realized niche data that are obtained from complete GIS layers, but is potentially relevant to functional traits that are obtained from databases with sparse coverage.

There are sample size challenges that come from trying to estimate a hypervolume from empirical data. The estimation problem easiest is for methods that allow description of relevant statistics (e.g. size, overlap) without directly delineating a geometrical shape (Junker et al. 2016). It is also simple for methods that do delineate a geometrical shape but assume additionally that the data has a certain shape or are drawn from a certain distribution. For example, estimation is also simple when it is assumed that data are drawn from a multivariate normal distribution and have no holes (Qiao et al. 2017), or that there are no or few interactions between hypervolume axes, e.g. range boxes or generalized linear models (Busby 1991, Leathwick et al. 2006). In these cases, it becomes possible to work with low sample sizes, but at the tradeoff that the resulting hypervolumes may reflect more about the modeling approach than the data.

In methods that allow for more flexible delineation of shapes, much larger samples are required. As dimensionality increases, the number of possible shapes grows very rapidly, and so do the data required to constrain them (Blonder et al. 2014, 2017). When sufficient data are not available, options include making a stronger set of assumptions and using a simpler method, reducing the dimensionality of the analysis, or obtaining more data. In general, data with high dimensionality but low numbers of observations (e.g. gene expression data) will not be able to apply hypervolume concepts and should be handled with other approaches, while data with lower dimensionality and many observations may be able to use hypervolume concepts.

Some methods are invariant to certain data transformations. If such transformations represent uncontrollable or unwanted effects, then invariance is useful; alternatively if such transformations represent differences in biological interpretation, then invariance is not useful. Some methods (e.g. dynamic range box, support vector machine) are invariant to linear shift and scale transformations, because data are rescaled to a unit interval before analysis. Some (e.g. kernel density estimation, support vector machine, PCA, Bayesian ellipse, etc.) are invariant to rotational transformation (i.e. correlations between axes, or equivalently choice of different coordination systems). Some of the methods, such as random forests and range bagging, are invariant to general monotone transformations, e.g. logarithmic transformation. However for most other algorithms, such transforms can yield different hypervolumes after back-transformation. For example, convex hulls and kernel density estimates are not necessarily preserved by data transformation.

Most extant methods can use sampled data to calculate specific statistics of interest, but do not provide information on the accuracy of these estimates. Accomplishing the latter goal would not only require the formulation of a feasible statistical and probabilistic model, but also careful thought about whether the sampled data actually mimic the underlying population values in an unbiased, representative, and from a sampling point of view, probabilistically quantifiable

manner. This would help to avoid the danger of mistaking estimates calculated from sampled values as ‘true’ measures of size, overlap, etc. By assuming a sampling model (e.g. independent and identically distributed), it is possible to compare performance among hypervolume methods using a range of statistical approaches. These approaches either describe the quality of the model fit to the input data, or describe the predictive ability of the method on a subset of the input data, via cross-validation, or on new test data.

For some types of models it is possible to estimate a likelihood function, then use a maximum likelihood estimate (MLE) to find a best set of model parameters. Likelihood functions can also be compared to select among modeling approaches, potentially with a penalty term for model complexity (e.g. via Akaike information criterion). However this approach is based entirely on a fit to the input data and does not necessarily provide guidance on predictive ability. Additionally, some of the surveyed modeling approaches do not lend themselves to maximum likelihood analysis, e.g. range boxes, kernel density estimates, support vector machines, or convex hulls. Some of these models have no parameters, so a likelihood cannot be defined. In other cases, the MLE is not helpful. For example a kernel density estimate has one parameter – a bandwidth, which controls the padding around each data point. The MLE will always select a bandwidth of zero (Shalizi 2011) in order to yield a model most consistent with the data, but this provides no utility for predicting values of h for new data.

Cross-validation approaches provide more insight into model fit and predictive ability (Caruana and Niculescu-Mizil 2006). In these approaches the model is used to predict class membership for a new test dataset, with results summarized into a confusion matrix (number of true and false positives and negative predictions) or a divergence metric, e.g. relative entropy (Shalizi 2011). Several of these matrices can be constructed as model parameters are varied and transformed into a receiver-operator characteristic (ROC) curve. This confusion matrix can be used then to construct Type I/II error rates, sensitivities, specificities, and other metrics of performance. Importantly, the cross-validation can occur either on different folds of data resampled from the calibration or on wholly novel test data (Bahn and McGill 2013). It is also possible to use these cross-validation approaches to compare other summary statistics of hypervolumes, e.g. volume and overlap, to an a-priori expectation (Blonder 2016). For some methods, model performance has been assessed via these approaches using simulated input data with known statistical properties (Qiao et al. 2016). Among two-class methods, generalized linear model and boosted regression tree approaches tend to have better performance on simple datasets, while among one-class methods, convex hull and range box methods tend to have better performance on simple datasets, while kernel density estimates have better performance on complex datasets (Qiao et al. 2017).

Challenges and opportunities

The range of research topics across hypervolume scales and axis types and still remain largely unresolved in general cases (Table 1). Moreover, this discussion of concepts

and estimation methods has highlighted several limitations and possible synthesis areas that could be readily addressed through further research (Table 2).

Hutchinson (1957) originally noted that the hypervolume ‘will exist whatever the shape of its sides’ and discussed only a small number of the issues explored here. Carefully assessing these geometrical considerations will be critical in determining whether hypervolumes are a useful description for a given system of interest. It may be that the hypervolume concept should not be used in some research areas because of data or conceptual limitations – for example, in cases where Euclidean volumes and distances are not well defined or comparable, where data dimensionality is too high relative to the number of data points, or where multidimensional overlap metrics do not adequately capture the underlying processes in a system (Holt 1987).

Additionally, there are useful parallels between the uses of hypervolumes in modeling niches, functional diversity, fitness landscapes, and morphospaces. Understanding these commonalities and shared intellectual heritage can be useful for synthesizing ideas and methods. However in some cases these parallels have been incorrectly considered as equivalences. This issue is especially prevalent when using functional trait hypervolumes for understanding niche-based processes, when little is known about how traits and niches are actually related. In general, this suggests a need for more careful thinking about the implications of using hypervolumes and accepting their assumptions for any given analysis.

The methods described in Supplementary material Appendix 1 Table A1 have not all been compared with each other on real or simulated data, so their relative performance remains unknown (but see Elith et al. 2006, Qiao et al. 2015 for some methods). Indeed, most studies, especially in functional diversity applications, have tended to choose a single method without exploring its statistical properties or considering alternatives. There are at present no software packages available to permit statistical comparison of all these methods, especially most of the one-class methods, although some tools exist for a restricted set of methods (Thuiller et al. 2009, Naimi and Araújo 2016). In general it is unknown which types of questions and data are best suited for each modeling approach, resulting in extensive ongoing discussion (Blonder et al. 2016, 2017, Carmona et al. 2016b, Qiao et al. 2017). More interaction between different research communities that rely on hypervolume concepts would be fruitful.

Table 2. Priority questions for hypervolume research.

-
- How do hypervolumes change through time in response to different factors at different scales?
 - What biological processes can be inferred from hypervolume geometry?
 - How can mechanistic models predict hypervolume boundaries at different biological scales and axis types?
 - What shapes and dimensionalities do hypervolumes have for different applications?
 - When are static hypervolume concepts insufficient to address biological questions?
 - Which descriptive and inferential methods are best for estimating hypervolumes?
-

This performance issue underlines another fundamental knowledge gap. In general there is still not a good expectation for the dimensionality and shape complexity that would be sufficient to characterize hypervolumes in different contexts. For example, there is still active debate as to whether fundamental niches have simple shapes or not, although it is widely recognized that realized and potential niches do have complex shapes (Jackson and Overpeck 2000). Similarly, explorations of the dimensionality of trait space and ecological strategy spaces for different taxa have suggested that anywhere from 4 to 10 dimensions characterize many systems (Harpole and Tilman 2007, Eklöf et al. 2013, Laughlin 2014). If the dimensionality of the system of interest is at the high end of this range or even higher, then it may be difficult to apply the hypervolume concept because of data limitations. For example in a ten-dimensional trait space it may be difficult to apply a trait packing hypervolume approach when a community has only a few dozen species – too much of the space is empty for the concept of packing to be reasonable, or for it to be possible to have any insight into the shape of the multidimensional community trait distribution. Further meta-analysis to establish the properties of empirical hypervolumes would be useful.

Additionally, there is opportunity to develop better tools for statistical inference with hypervolumes. Inferential applications would include determining which of several shape models for a hypervolume is most likely, whether two hypervolumes' overlap is significantly greater than zero, whether one hypervolume is significantly larger in volume than another, and so on. Such methods are necessary when the data used to estimate $h(x)$ are not complete but rather are a random sample. Probabilistic models that are based on likelihoods can be compared amongst each other, but such comparison is not possible with models that do not have a likelihood. Confidence intervals around statistics like volume and overlap are not easy to obtain from likelihood-based methods, and are presently not available at all for non-likelihood based methods, e.g. convex hulls. One approach to resolve these questions would be to use bootstrapping methods to repeatedly resample the data, providing empirical sampling distributions and thus parametric or non-parametric tests for the statistics of interest. Bootstrapping can be implemented for any method (Parkinson et al. in review, Swanson et al. 2015), but may not perform well for estimating extrema of distributions (Drake 2015).

There are also limited mechanistic models that can predict the shape of hypervolumes. For fundamental niche hypervolumes, a few recent demographic and physiological models have provided first steps (Hoogenboom and Connolly 2009, Kearney and Porter 2009). Similar approaches exist for zero net growth isoclines, though parameterization of these models is difficult (Miller et al. 2005). The general problem of predicting boundaries under different environmental and biotic contexts remains open. For functional hypervolumes, there is still little predictive theory for e.g. community assembly or trait evolution that can predict the shape and size of trait hypervolumes, and even fewer empirical tests of theory.

The statistical methods for estimating hypervolumes are rapidly developing, with extensive progress in recent years. Nevertheless, there are still a range of opportunities

and questions present for methods development. Notably, there appears to be extensive room for sharing concepts across models traditionally used for niche modeling and those used for functional diversity modeling. For example, the models used for species distributions have a robust set of evaluation metrics developed but limited tools for summarizing the shape and overlap of the underlying hypervolume, while the converse is true of models used for functional diversity.

There is a now vibrant and active discussion of conceptual and methodological issues surrounding n -dimensional hypervolumes. This discussion has extended far beyond Hutchinson's original applications for niches to a range of other research areas, underscoring the potential importance of this geometrical concept across ecology and evolutionary biology. Determining when and where hypervolumes are useful will be useful in producing more conceptually sound science.

Acknowledgments – John Drake and William Godsoe greatly improved the manuscript through their thoughtful reviews. I am grateful for comments from Robert R. Junker, Arne Bathke, Jonas Kuppler, Marc Macias Fauria, Naia Morueta Holme, Imma Oliveras, Manuela Schreyer, Carlos Perez Carmona, and Francesco de Bello. I thank Christine Lamanna, Cyrille Violle, Brian J. Enquist, Cecina Babich Morrow, Brian Maitner, Drew Kerkhoff, and Hujie Qiao for their feedback.

Funding – This work was supported by the UK Natural Environment Research Council (NE/M019160/1).

References

- Abrams, P. 1983. The theory of limiting similarity. – *Annu. Rev. Ecol. Syst.* 14: 359–376.
- Adler, P. B. et al. 2013. Trait-based tests of coexistence mechanisms. – *Ecol. Lett.* 16: 1294–1306.
- Anderson, R. 2017. When and how should biotic interactions be considered in models of species niches and distributions? – *J. Biogeogr.* 44: 8–17.
- Araújo, M. B. and Guisan, A. 2006. Five (or so) challenges for species distribution modelling. – *J. Biogeogr.* 33: 1677–1688.
- Araújo, M. B. et al. 2013. Heat freezes niche evolution. – *Ecol. Lett.* 16: 1206–1219.
- Austin, M. P. 1985. Continuum concept, ordination methods, and niche theory. – *Annu. Rev. Ecol. Syst.* 39–61.
- Austin, M. et al. 1990. Measurement of the realized qualitative niche: environmental niches of five Eucalyptus species. – *Ecol. Monogr.* 60: 161–177.
- Bahn, V. and McGill, B. J. 2013. Testing the predictive performance of distribution models. – *Oikos* 122: 321–331.
- Barros, C. et al. 2016. N-dimensional hypervolumes to study stability of complex ecosystems. – *Ecol. Lett.* 19: 729–742.
- Begon, M. et al. 2006. *Ecology: from individuals to ecosystems*. – Blackwell.
- Beisner, B. E. et al. 2003. Alternative stable states in ecology. – *Front. Ecol. Environ.* 1: 376–382.
- Birch, L. 1953. Experimental background to the study of the distribution and abundance of insects: I. The influence of temperature, moisture and food on the innate capacity for increase of three grain beetles. – *Ecology* 34: 698–711.
- Blonder, B. 2016. Do hypervolumes have holes? – *Am. Nat.* 187: E93–E105.
- Blonder, B. et al. 2014. The n -dimensional hypervolume. – *Global Ecol. Biogeogr.* 23: 595–609.

- Blonder, B. et al. 2015. Linking environmental filtering and disequilibrium to biogeography with a community climate framework. – *Ecology* 96: 972–985.
- Blonder, B. et al. 2016. Pushing past boundaries for trait hypervolumes: a response to Carmona et al. – *Trends Ecol. Evol.* 31: 665–667.
- Blonder, B. et al. 2017. Using n-dimensional hypervolumes for species distribution modeling: a response to Qiao et al. (2016). – *Global Ecol. Biogeogr.* 26: 1071–1075.
- Booth, T. H. et al. 2014. BIOCLIM: the first species distribution modelling package, its early applications and relevance to most current MAXENT studies. – *Divers. Distrib.* 20: 1–9.
- Broennimann, O. et al. 2007. Evidence of climatic niche shift during biological invasion. – *Ecol. Lett.* 10: 701–709.
- Broennimann, O. et al. 2012. Measuring ecological niche overlap from occurrence and spatial environmental data. – *Global Ecol. Biogeogr.* 21: 481–497.
- Busby, J. 1991. BIOCLIM—a bioclimate analysis and prediction system. – *Plant Protect. Quart.* 6: 8–9.
- Cabral, J. S. and Schurr, F. M. 2010. Estimating demographic models for the range dynamics of plant species. – *Global Ecol. Biogeogr.* 19: 85–97.
- Carmona, C. P. et al. 2016a. Traits without borders: integrating functional diversity across scales. – *Trends Ecol. Evol.* 31: 382–94.
- Carmona, C. P. et al. 2016b. The density awakens: a reply to Blonder. – *Trends Ecol. Evol.* 31: 667–669.
- Caruana, R. and Niculescu-Mizil, A. 2006. An empirical comparison of supervised learning algorithms. – *Proceedings of the 23rd international conference on machine learning. ACM*, pp. 161–168.
- Chase, J. M. and Leibold, M. A. 2003. *Ecological niches: linking classical and contemporary approaches*. – Univ. of Chicago Press.
- Chesson, P. 2000. Mechanisms of maintenance of species diversity. – *Annu. Rev. Ecol. Syst.* 31: 343–366.
- Colwell, R. K. and Fuentes, E. R. 1975. Experimental studies of the niche. – *Annu. Rev. Ecol. Syst.* 6: 281–310.
- Colwell, R. K. and Rangel, T. F. 2009. Hutchinson's duality: the once and future niche. – *Proc. Natl Acad. Sci. USA* 106: 19651–19658.
- Connor, E. F. and Simberloff, D. 1979. The assembly of species communities: chance or competition? – *Ecology* 60: 1132–1140.
- Cornwell, W. and Ackerly, D. 2009. Community assembly and shifts in plant trait distributions across an environmental gradient in coastal California. – *Ecol. Monogr.* 79: 106–126.
- Cornwell, W. K. et al. 2006. A trait-based test for habitat filtering: convex hull volume. – *Ecology* 87: 1465–1471.
- D'Andrea, R. and Ostling, A. 2016. Challenges in linking trait patterns to niche differentiation. – *Oikos* 125: 1369–1385.
- Diamond, J. 1975. Assembly of species communities. – In: Cody, M. and Diamond, J. (eds), *Ecology and evolution of communities*. Harvard Univ. Press.
- Diamond, J. M. and Gilpin, M. E. 1982. Examination of the 'null' model of Connor and Simberloff for species co-occurrences on islands. – *Oecologia* 52: 64–74.
- Díaz, S. et al. 2016. The global spectrum of plant form and function. – *Nature* 529: 167–171.
- Drake, J. M. 2015. Range bagging: a new method for ecological niche modelling from presence-only data. – *J. R. Soc. Inter.* 12: 20150086.
- Drake, J. M. and Beier, J. C. 2014. Ecological niche and potential distribution of *Anopheles arabiensis* in Africa in 2050. – *Malaria J.* 13: 213.
- Eklöf, A. et al. 2013. The dimensionality of ecological networks. – *Ecol. Lett.* 16: 577–583.
- Elith, J. et al. 2006. Novel methods improve prediction of species' distributions from occurrence data. – *Ecography* 29: 129–151.
- Engler, R. and Guisan, A. 2009. MigClim: predicting plant distribution and dispersal in a changing climate. – *Divers. Distrib.* 15: 590–601.
- Enquist, B. J. et al. 2016. The botanical information and ecology network (BIEN): cyberinfrastructure for an integrated botanical information network to investigate the ecological impacts of global climate change on plant biodiversity. – <http://www.cyverse.org/sites/default/files/BIEN_White_Paper.pdf>.
- Evans, M. E. et al. 2009. Climate, niche evolution, and diversification of the 'bird-cage' evening primroses (*Oenothera*, Sections *Anogra* and *Kleinia*). – *Am. Nat.* 173: 225–240.
- Franklin, J. 2010. *Mapping species distributions: spatial inference and prediction*. – Cambridge Univ. Press.
- Funk, J. L. and Wolf, A. A. 2016. Testing the trait-based community framework: do functional traits predict competitive outcomes? – *Ecology* 97: 2206–2211.
- Garnier, E. et al. 2016. *Plant functional diversity: organism traits, community structure, and ecosystem properties*. – Oxford Univ. Press.
- Gause, G. F. 1934. *The struggle for existence*. – The Williams and Wilkins Company.
- Gavrilets, S. 1997. Evolution and speciation on holey adaptive landscapes. – *Trends Ecol. Evol.* 12: 307–312.
- Godsoe, W. 2010. I can't define the niche but I know it when I see it: a formal link between statistical theory and the ecological niche. – *Oikos* 119: 53–60.
- Godsoe, W. and Case, B. S. 2015. Accounting for shifts in the frequency of suitable environments when testing for niche overlap. – *Methods Ecol. Evol.* 6: 59–66.
- Godsoe, W. and Harmon, L. J. 2012. How do species interactions affect species distribution models? – *Ecography* 35: 811–820.
- Götzenberger, L. et al. 2012. Ecological assembly rules in plant communities – approaches, patterns and prospects. – *Biol. Rev.* 87: 111–127.
- Gould, S. J. 1991. The disparity of the Burgess Shale arthropod fauna and the limits of cladistic analysis: why we must strive to quantify morphospace. – *Paleobiol.* 17: 411–423.
- Gower, J. C. 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: 543–556.
- Grime, J. 1977. Evidence for the existence of three primary strategies in plants and its relevance to ecological and evolutionary theory. – *Am. Nat.* 111: 1169–1194.
- Grinnell, J. 1917. The niche-relationships of the California thrasher. – *Auk* 34: 427–433.
- Grubb, P. J. 1977. The maintenance of species-richness in plant communities: the importance of the regeneration niche. – *Biol. Rev.* 52: 107–145.
- Guisan, A. and Zimmermann, N. E. 2000. Predictive habitat distribution models in ecology. – *Ecol. Model.* 135: 147–186.
- Harpole, W. S. and Tilman, D. 2007. Grassland species loss resulting from reduced niche dimension. – *Nature* 446: 791–793.
- Herbold, B. and Moyle, P. B. 1986. Introduced species and vacant niches. – *Am. Nat.* 128: 751–760.
- Holt, R. D. 1987. On the relation between niche overlap and competition: the effect of incommensurable niche dimensions. – *Oikos* 48: 110–114.
- Holt, R. D. 2009. Bringing the Hutchinsonian niche into the 21st century: ecological and evolutionary perspectives. – *Proc. Natl Acad. Sci. USA* 106: 19659–19665.
- Hoogenboom, M. O. and Connolly, S. R. 2009. Defining fundamental niche dimensions of corals: synergistic effects of colony size, light, and flow. – *Ecology* 90: 767–780.

- Hooper, H. L. et al. 2008. The ecological niche of *Daphnia magna* characterized using population growth rate. – *Ecology* 89: 1015–1022.
- Hooten, M. B. and Wikle, C. K. 2008. A hierarchical Bayesian non-linear spatio-temporal model for the spread of invasive species with application to the Eurasian collared-dove. – *Environ. Ecol. Stat.* 15: 59–70.
- Hulshof, C. M. and Swenson, N. G. 2010. Variation in leaf functional trait values within and across individuals and species: an example from a Costa Rican dry forest. – *Funct. Ecol.* 24: 217–223.
- Huntley, J. W. et al. 2008. Testing limiting similarity in Quaternary terrestrial gastropods. – *Paleobiology* 34: 378–388.
- Hutchinson, G. 1957. Concluding remarks. – *Cold Spring Harb. Symp. Quant. Biol.* 22: 415–427.
- Hutchinson, G. 1959. Homage to Santa Rosalia or why are there so many kinds of animals? – *Am. Nat.* 93: 145–159.
- Hutchinson, G. 1978. An introduction to population ecology. – Yale Univ. Press.
- Jackson, S. T. and Overpeck, J. T. 2000. Responses of plant populations and communities to environmental changes of the late Quaternary. – *Paleobiol.* 26: 194–220.
- Jones, K. E. et al. 2009. PanTHERIA: a species-level database of life history, ecology, and geography of extant and recently extinct mammals. – *Ecology* 90: 2648–2648.
- Junker, R. R. et al. 2016. Dynamic range boxes – a robust nonparametric approach to quantify size and overlap of n-dimensional hypervolumes. – *Methods Ecol. Evol.* 7: 1503–1513.
- Kattge, J. et al. 2011. TRY – a global database of plant traits. – *Global Change Biol.* 17: 2905–2935.
- Kearney, M. and Porter, W. P. 2004. Mapping the fundamental niche: physiology, climate, and the distribution of a nocturnal lizard. – *Ecology* 85: 3119–3131.
- Kearney, M. and Porter, W. 2009. Mechanistic niche modelling: combining physiological and spatial data to predict species' ranges. – *Ecol. Lett.* 12: 334–350.
- Kraft, N. et al. 2007. Trait evolution, community assembly and the phylogenetic structure of ecological communities. – *Am. Nat.* 170: 271–283.
- Kraft, N. J. et al. 2008. Functional traits and niche-based tree community assembly in an Amazonian forest. – *Science* 322: 580–582.
- Kraft, N. J. et al. 2015. Plant functional traits and the multidimensional nature of species coexistence. – *Proc. Natl Acad. Sci. USA* 112: 797–802.
- Kunstler, G. et al. 2012. Competitive interactions between forest trees are driven by species' trait hierarchy, not phylogenetic or functional similarity: implications for forest community assembly. – *Ecol. Lett.* 15: 831–840.
- Kunstler, G. et al. 2016. Plant functional traits have globally consistent effects on competition. – *Nature* 529: 204–207.
- Laliberte, E. et al. 2010. Land-use intensification reduces functional redundancy and response diversity in plant communities. – *Ecol. Lett.* 13: 76–86.
- Lamanna, C. et al. 2014. Functional trait space and the latitudinal diversity gradient. – *Proc. Natl Acad. Sci. USA* 111: 13745–13750.
- Laughlin, D. C. 2014. The intrinsic dimensionality of plant traits and its relevance to community assembly. – *J. Ecol.* 102: 186–193.
- Laughlin, D. C. and Messier, J. 2015. Fitness of multidimensional phenotypes in dynamic adaptive landscapes. – *Trends. Ecol. Evol.* 30: 487–496.
- Leathwick, J. et al. 2006. Comparative performance of generalized additive models and multivariate adaptive regression splines for statistical modelling of species distributions. – *Ecol. Model.* 199: 188–196.
- Lepš, J. et al. 2006. Quantifying and interpreting functional diversity of natural communities: practical considerations matter. – *Preslia* 78: 481–501.
- Letten, A. D. et al. 2017. Linking modern coexistence theory and contemporary niche theory. – *Ecol. Monogr.* 87: 161–177.
- Loranger, J. et al. 2016. Occupancy and overlap in trait space along a successional gradient in Mediterranean old fields. – *Am. J. Botany* 103: 1050–1060.
- MacArthur, R. and Levins, R. 1967. The limiting similarity, convergence, and divergence of coexisting species. – *Am. Nat.* 101: 377–385.
- Maguire, B. 1967. A partial analysis of the niche. – *Am. Nat.* 101: 515–526.
- Maher, S. P. et al. 2014. Pattern-recognition ecological niche models fit to presence-only and presence-absence data. – *Methods Ecol. Evol.* 5: 761–770.
- Mason, N. W. et al. 2005. Functional richness, functional evenness and functional divergence: the primary components of functional diversity. – *Oikos* 111: 112–118.
- Mason, N. W. H. and Pavoine, S. 2013. Does trait conservatism guarantee that indicators of phylogenetic community structure will reveal niche-based assembly processes along stress gradients? – *J. Veg. Sci.* 24: 820–833.
- May, R. M. and Arthur, R. H. M. 1972. Niche overlap as a function of environmental variability. – *Proc. Natl Acad. Sci. USA* 69: 1109–1113.
- Mayfield, M. M. and Levine, J. M. 2010. Opposing effects of competitive exclusion on the phylogenetic structure of communities. – *Ecol. Lett.* 13: 1085–1093.
- McGill, B. J. et al. 2006. Rebuilding community ecology from functional traits. – *Trends Ecol. Evol.* 21: 178–185.
- McInerny, G. J. and Etienne, R. S. 2012. Ditch the niche – is the niche a useful concept in ecology or species distribution modelling? – *J. Biogeogr.* 39: 2096–2102.
- Milla, R. et al. 2015. Plant domestication through an ecological lens. – *Trends Ecol. Evol.* 30: 463–469.
- Miller, T. E. et al. 2005. A critical review of twenty years' use of the resource-ratio theory. – *Am. Nat.* 165: 439–48.
- Mod, H. K. et al. 2015. Biotic interactions boost spatial models of species richness. – *Ecography* 38: 913–921.
- Moles, A. T. et al. 2008. A new framework for predicting invasive plant species. – *J. Ecol.* 96: 13–17.
- Mouchet, M. A. et al. 2010. Functional diversity measures: an overview of their redundancy and their ability to discriminate community assembly rules. – *Funct. Ecol.* 24: 867–876.
- Mouillot, D. et al. 2014. Functional over-redundancy and high functional vulnerability in global fish faunas on tropical reefs. – *Proc. Natl Acad. Sci. USA* 111: 13757–13762.
- Naimi, B. and Araújo, M. B. 2016. sdm: a reproducible and extensible R platform for species distribution modelling. – *Ecography* 39: 368–375.
- Nenzén, H. K. et al. 2012. demoniche – an R-package for simulating spatially-explicit population dynamics. – *Ecography* 35: 577–580.
- Parkinson, J. H. et al. in review. A fast and robust way to estimate overlap of d-dimensional niches, and draw inference. – *Int. J. Biostat.*
- Pausas, J. G. and Verdú, M. 2008. Fire reduces morphospace occupation in plant communities. – *Ecology* 89: 2181–2186.
- Peters, R. H. 1991. A critique for ecology. – Cambridge Univ. Press.
- Peterson, A. T. et al. 2011. Ecological niches and geographic distributions (MPB-49). – Princeton Univ. Press.
- Phillips, S. J. et al. 2006. Maximum entropy modeling of species geographic distributions. – *Ecol. Model.* 190: 231–259.
- 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.

- Porter, W. P. and Kearney, M. 2009. Size, shape and the thermal niche of endotherms. – *Proc. Natl Acad. Sci. USA* 106: 19666–19672.
- Pulliam, H. R. 2000. On the relationship between niche and distribution. – *Ecol. Lett.* 3: 349–361.
- Qiao, H. et al. 2015. Marble algorithm: a solution to estimating ecological niches from presence-only records. – *Sci. Rep.* 5: 14232.
- Qiao, H. et al. 2016. NicheA: creating virtual species and ecological niches in multivariate environmental scenarios. – *Ecography* 39: 805–813.
- Qiao, H. et al. 2017. A cautionary note on the use of hypervolume kernel density estimators in ecological niche modelling. – *Global Ecol. Biogeogr.* 26: 1066–1070.
- Rappoldt, C. and Hogeweg, P. 1980. Niche packing and number of species. – *Am. Nat.* 116: 480–492.
- Raup, D. M. and Michelson, A. 1965. Theoretical morphology of the coiled shell. – *Science* 147: 1294–1295.
- Renner, I. W. and Warton, D. I. 2013. Equivalence of MAXENT and Poisson point process models for species distribution modeling in ecology. – *Biometrics* 69: 274–281.
- Ricotta, C. et al. 2016. Measuring the functional redundancy of biological communities: a quantitative guide. – *Methods Ecol. Evol.* 7: 1386–1395.
- Scheffer, M. and van Nes, E. H. 2006. Self-organized similarity, the evolutionary emergence of groups of similar species. – *Proc. Natl Acad. Sci. USA* 103: 6230–6235.
- Shalizi, C. 2011. Estimating distributions and densities. – <<https://www.stat.cmu.edu/~cshalizi/402/lectures/06-density/lecture-06.pdf>>.
- Shen, B. et al. 2008. The Avalon explosion: evolution of Ediacara morphospace. – *Science* 319: 81–84.
- Shipley, B. et al. 2016. Reinforcing loose foundation stones in trait-based plant ecology. – *Oecologia* 180: 923–931.
- Shreve, F. 1911. The influence of low temperatures on the distribution of the giant cactus. – *Plant World* 14: 136–146.
- Sidlauskas, B. 2008. Continuous and arrested morphological diversification in sister clades of characiform fishes: a phylomorphospace approach. – *Evolution* 62: 3135–3156.
- Šimová, I. et al. 2015. Shifts in trait means and variances in North American tree assemblages: species richness patterns are loosely related to the functional space. – *Ecography* 38: 649–658.
- Soberón, J. and Peterson, T. 2004. Biodiversity informatics: managing and applying primary biodiversity data. – *Phil. Trans. R. Soc. B* 359: 689–698.
- Soberón, J. and Nakamura, M. 2009. Niches and distributional areas: concepts, methods and assumptions. – *Proc. Natl Acad. Sci. USA* 106 Suppl 2: 19644–50.
- Soberón, J. 2010. Niche and area of distribution modeling: a population ecology perspective. – *Ecography* 33: 159–167.
- Stayton, C. T. 2005. Morphological evolution of the lizard skull: a geometric morphometrics survey. – *J. Morph.* 263: 47–59.
- Stubbs, W. J. and Bastow Wilson, J. 2004. Evidence for limiting similarity in a sand dune community. – *J. Ecol.* 92: 557–567.
- Swanson, H. K. et al. 2015. A new probabilistic method for quantifying n-dimensional ecological niches and niche overlap. – *Ecology* 96: 318–324.
- Swenson, N. G. and Enquist, B. J. 2009. Opposing assembly mechanisms in a Neotropical dry forest: implications for phylogenetic and functional community ecology. – *Ecology* 90: 2161–2170.
- Swenson, N. G. et al. 2012. The biogeography and filtering of woody plant functional diversity in North and South America. – *Global Ecol. Biogeogr.* 21: 798–808.
- ter Braak, C. J. and Prentice, I. C. 1988. A theory of gradient analysis. – *Adv. Ecol. Res.* 18: 271–317.
- Thuiller, W. et al. 2009. BIOMOD – a platform for ensemble forecasting of species distributions. – *Ecography* 32: 369–373.
- Tilman, D. 1977. Resource competition between plankton algae: an experimental and theoretical approach. – *Ecology* 58: 338–348.
- Tingley, R. et al. 2014. Realized niche shift during a global biological invasion. – *Proc. Natl Acad. Sci. USA* 111: 10233–10238.
- van Bodegom, P. M. et al. 2014. A fully traits-based approach to modeling global vegetation distribution. – *Proc. Natl Acad. Sci. USA* 111: 13733–13738.
- Van Kleunen, M. et al. 2010. A meta-analysis of trait differences between invasive and non-invasive plant species. – *Ecol. Lett.* 13: 235–245.
- Van Valen, L. 1965. Morphological variation and width of ecological niche. – *Am. Nat.* 99: 377–390.
- Villéger, S. et al. 2008. New multidimensional functional diversity indices for a multifaceted framework in functional ecology. – *Ecology* 89: 2290–2301.
- von Humboldt, A. 1807. *Essai sur la géographie des plantes: accompagné d'un tableau physique des régions équinoxiales.* – Chez Levbault, Schoell, et Compagnie.
- Wainwright, P. C. et al. 2004. A functional morphospace for the skull of labrid fishes: patterns of diversity in a complex biomechanical system. – *Biol. J. Linn. Soc.* 82: 1–25.
- Warheit, K. I. et al. 1999. Morphological diversification and adaptive radiation: a comparison of two diverse lizard clades. – *Evolution* 1226–1234.
- Webb, C. O. et al. 2002. Phylogenies and community ecology. – *Annu. Rev. Ecol. Syst.* 33: 475–505.
- Weiher, E. and Keddy, P. A. 1995. Assembly rules, null models, and trait dispersion: new questions from old patterns. – *Oikos* 74: 159–164.
- Weiher, E. et al. 1998. Community assembly rules, morphological dispersion, and the coexistence of plant species. – *Oikos* 81: 309–322.
- Westoby, M. 1998. A leaf-height-seed (LHS) plant ecology strategy scheme. – *Plant Soil* 199: 213–227.
- Westoby, M. and Wright, I. J. 2006. Land-plant ecology on the basis of functional traits. – *Trends Ecol. Evol.* 21: 261–268.
- Whittaker, R. H. 1967. Gradient analysis of vegetation. – *Biol. Rev.* 42: 207–264.
- Whittaker, R. H. and Niering, W. A. 1965. Vegetation of the Santa Catalina Mountains, Arizona: a gradient analysis of the south slope. – *Ecology* 46: 429–452.
- Whittaker, R. J. et al. 2014. Functional biogeography of oceanic islands and the scaling of functional diversity in the Azores. – *Proc. Natl Acad. Sci. USA* 111: 13709–13714.
- Wiens, J. J. and Graham, C. H. 2005. Niche conservatism: integrating evolution, ecology, and conservation biology. – *Annu. Rev. Ecol. Syst.* 36: 519–539.
- Wright, S. 1932. The roles of mutation, inbreeding, crossbreeding, and selection in evolution. – In: Jones, D. (ed), *Proceedings of the Sixth International Congress of Genetics.*
- Zurell, D. et al. 2016. Benchmarking novel approaches for modelling species range dynamics. – *Global Change Biol.* 22: 2651–2664.

Supplementary material (Appendix ECOG-03187 at <www.ecography.org/appendix/ecog-03187>). Appendix 1.