vol. 187, no. 4 the american naturalist april 2016

E-Article

Do Hypervolumes Have Holes?

Benjamin Blonder*

Environmental Change Institute, School of Geography and the Environment, University of Oxford, South Parks Road, Oxford OX1 3QY, United Kingdom

Submitted May 13, 2015; Accepted October 23, 2015; Electronically published February 15, 2016

Online enhancements: video.

abstract: Hypervolumes are used widely to conceptualize niches and trait distributions for both species and communities. Some hypervolumes are expected to be convex, with boundaries defined by only upper and lower limits (e.g., fundamental niches), while others are expected to be maximal, with boundaries defined by the limits of available space (e.g., potential niches). However, observed hypervolumes (e.g., realized niches) could also have holes, defined as unoccupied hyperspace representing deviations from these expectations that may indicate unconsidered ecological or evolutionary processes. Detecting holes in more than two dimensions has to date not been possible. I develop a mathematical approach, implemented in the *hypervolume* R package, to infer holes in large and high-dimensional data sets. As a demonstration analysis, I assess evidence for vacant niches in a Galapagos finch community on Isabela Island. These mathematical concepts and software tools for detecting holes provide approaches for addressing contemporary research questions across ecology and evolutionary biology.

Keywords: niche, hypervolume, hole, geometry, Hutchinson, vacant niche, invasion.

Introduction

Hutchinson (1957) defined the niche as an *n*-dimensional hypervolume describing the set of environments that permit a species to exist. He noted that when the environmental axes "are independent in their action on the species we may regard this area as the rectangle . . . but failing such independence the area will exist whatever the shape of its sides" (p. 416). Many complex hypervolume shapes beyond rectangles are possible, especially in higher dimensions, where intuition fails. However, patterns of variation in hypervolume geometry remain understudied despite wide interest in the hypervolume concept (Leibold 1995; Colwell and Rangel 2009; Holt 2009). One key aspect of this geometry is the presence or absence of empty features within hypervolumes. Here I argue that such features (holes) may be able to provide insights into a range of contemporary research topics. I then present a mathematical approach for detecting holes and demonstrate it using both simulated and empirical data.

* E-mail: bblonder@gmail.com.

Am. Nat. 2016. Vol. 187, pp. E000–E000.
DOI: 10.1086/685444

Concepts and Terms

Hypervolumes can represent a range of ecological concepts depending on their operational definition. They can be defined for different scales (species, communities) and for different axes (climate variables, functional traits). For example, a species-scale hypervolume with climate axes could represent a Grinnellian niche (Grinnell 1924), a species-scale hypervolume with resource axes could represent an Eltonian niche (Elton 1958), and a community-scale hypervolume with trait axes could represent a set of ecological strategies viable in a particular location (McGill et al. 2006). Because the mathematical concepts I will present here are shared between these definitions, I distinguish between hypervolume types only when exploring the biological interpretation of different geometries.

Geometrical Features: Holes and Expectations

An observed hypervolume may have an arbitrarily complex geometry (Laughlin et al. 2015), for example, as often depicted for realized Grinnellian niches (Colwell and Futuyma 1971; Austin 1999). This geometry can be decomposed as an expectation minus a set of holes (fig. 1). Holes directly indicate deviations from expectations and so may reflect additional ecological or evolutionary processes beyond those structuring the expectation. The detection of holes provides an approach for generating hypotheses relevant to such processes and allows progress beyond the relatively simple baseline geometries that have commonly been used to characterize hypervolumes.

Two baseline expectations (convexity and maximality) provide useful starting points for assessing holes. One good baseline expectation is convexity. If a hypervolume is convex, then for any two points that are part of the hyper-

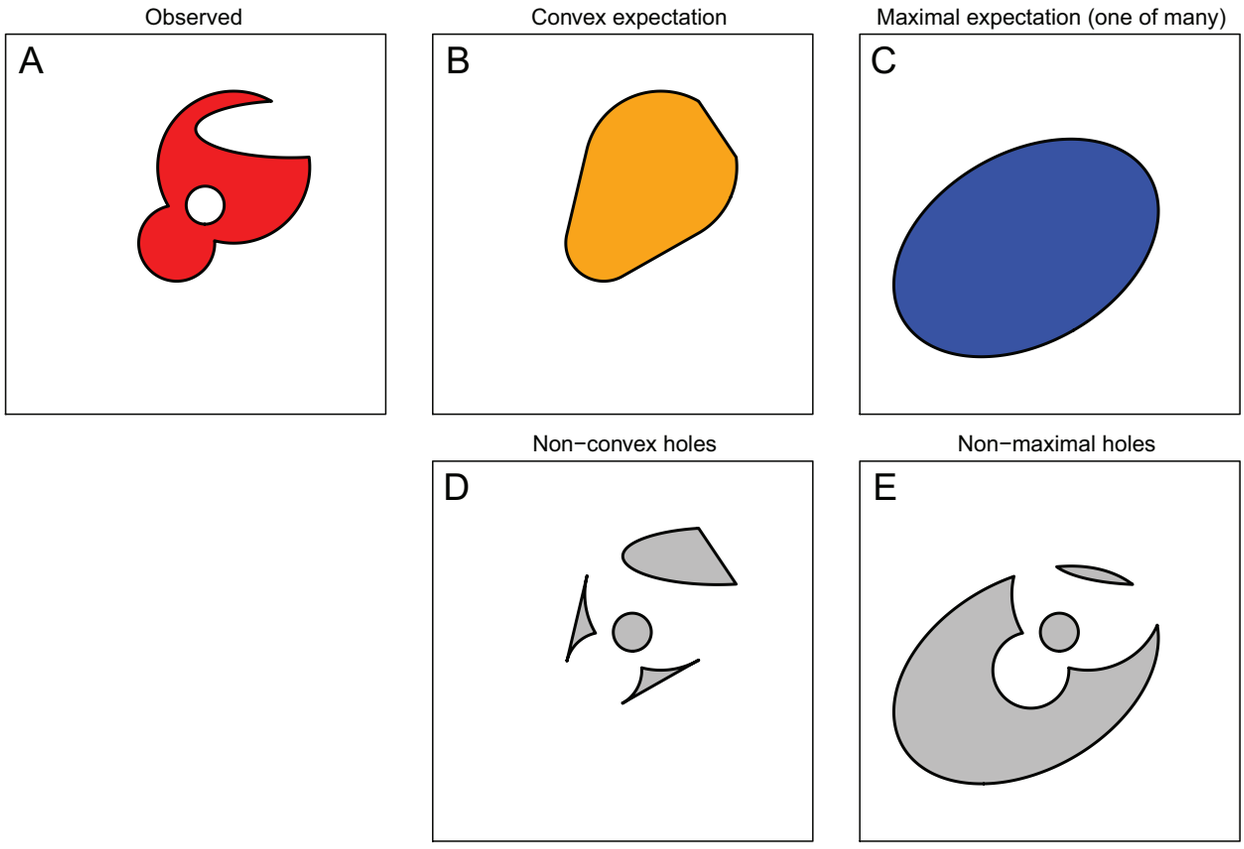

Figure 1: Typology of hypervolume geometry. *A*, An observed hypervolume (red) may have an arbitrarily complex shape. *B*, The convex expectation (orange) for this hypervolume is described by multiple linear inequality constraints. If any two points are in the convex hypervolume, then all points along the line segment connecting them are also in the convex hypervolume. *C*, A maximal expectation (blue) can have any shape, with boundaries indicating the limits of the available space. *D*, Holes (gray) can represent deviations from the convex expectation. *E*, Holes (gray) can also represent deviations from a maximal expectation.

volume, all points lying on a line segment between these points are also part of the hypervolume. An equivalent way to define a convex hypervolume is by the intersection of multiple linear inequality constraints, each defining a separating hyperplane and thus a face of the resulting convex hull (e.g., multiple upper/lower bounds). For example, an organism that can tolerate any temperature of more than 10°C (first inequality constraint) and less than 20°C (second inequality constraint) has a convex hypervolume. Rectangles, ellipses, and their multidimensional analogs (e.g., multivariate uniform and normal distributions) are examples of convex hypervolume geometries. Fundamental niches are often thought to be convex (Hutchinson 1957; Maguire 1973; Austin 1992; Soberón and Nakamura 2009) because physiological tolerances to different variables are often thought to be independent and limited only by extreme values. Species' physiological limitations (and thus environmental filters) are also widely thought to be convex (Violle et al. 2012; Lamanna et al. 2014).

A second baseline expectation is maximality. If a hypervolume is maximal, then its geometry is equivalent to that of the available n-dimensional space. An axis defining the hypervolume may have finite bounds or support only certain sets of values, limiting the shapes of observed hypervolumes (Jackson and Overpeck 2000). In practice, this means that a hypervolume with complex geometry may purely be the outcome of a shape reaching the "edges" of the available n-dimensional space. The available climate space can also be thought of as maximal (Soberón and Nakamura 2009). Morphological and trait spaces can also be thought of as maximal when they encompass all potential points that are not mathematically or biophysically impossible.

Holes are robust geometric features. They are invariant to affine transformations—for example, translation, scaling, rotation, or shearing of axes—and therefore should be detectable across a wide range of hypervolume definitions and axis specifications.

Ideas for Applications of Holes

Thinking about hypervolumes in terms of holes provides a complementary perspective to existing metrics. Functional richness indices such as convex hull volume (Cornwell et al. 2006) or Bayesian ellipses (Jackson et al. 2011) inherently do not consider holes. Metrics of dispersion, functional divergence, or evenness (e.g., minimum spanning trees or nearest-neighbor distances; Weiher and Keddy 1995; Villéger et al. 2008) can provide only summary statistics for the overall clustering of points within a hypervolume. Metrics of distributional moments (e.g., skewness and kurtosis; Enquist et al. 2015) are applicable only to single-dimensional analyses and have the same summary statistic limitation. And covariance matrices used to describe genotypes and phenotypes (G and P matrices) describe multivariate ellipses that inherently have no holes (Lynch and Walsh 1998). In contrast, hole detection directly delineates the empty portions of an n-dimensional space.

The detection of a hole is a necessary but not sufficient condition for inference of processes structuring an observed hypervolume. However, determining the location and shape of holes can be a useful starting point for generating and testing specific mechanisms or hypotheses, assuming that potential issues of incomplete or biased sampling can be ruled out. Broadly, thinking about hypervolumes in terms of holes can provide new approaches to addressing several contemporary research questions.

Does Invasion or Succession Preferentially Occur in Holes?

Successful colonists can coexist through increased niche differences, decreased fitness differences (Chesson 2000; Adler et al. 2013), or mass effects (Leibold et al. 2004). There is evidence that in some cases niche and trait differences predict invasiveness (Van Kleunen et al. 2010) but also that low fitness differences allow invasive species to share niche and trait values with resident species (Lai et al. 2015). These analyses have been conducted only in single dimensions. A key hypothesis is that invasive species preferentially occupy multidimensional holes in a community's trait space or realized niche space.

Do Geographic Range Shifts of Species Reflect Similar Niche Space Dynamics?

Patterns of geographic range dynamics can reflect multiple processes (Channell and Lomolino 2000), which could potentially lead to parallel gain/loss of niche space. For example, the mechanisms driving the Late Pleistocene extinction of many megafaunal species (Lorenzen et al. 2011) have not been assessed in terms of niche geometry. One hypothesis is that species interactions or loss of appropriate climate conditions could lead to the formation of holes in trait space or realized niche space.

Will Contemporary Global Change Lead to No-Analog Climates beyond Species' Potential Niches?

Climate change may lead to no-analog climate spaces (Williams and Jackson 2007) that have not been experienced in recent history and include holes. A key hypothesis is that species with potential climate niches falling completely within a hole might not be able to migrate to nearby regions with more favorable climate and so become locally extinct (Burrows et al. 2014).

Do Species' Niches Evolve Holes?

Niche evolution is traditionally explored in terms of means or extremes. However, niche evolution does not always proceed through simple shifts (Pearman et al. 2008). One hypothesis is that genetic constraints, developmental constraints, or natural selection prevent the existence of certain phenotypes (Arnold 1992), leading to the existence of holes in trait or potential niche space that persist within or across lineages over time.

Are Fitness Landscapes Holey?

Advances in high-dimensional neutral networks show that adaptive landscapes are not peaked but instead have fitness contours that fold over large regions of phenotype space. Such landscapes have been inferred on the basis of the presence of a small number of intermediate genotypes connecting populations with very different phenotypes (e.g., in *Drosophila*; Weber 1996). Thus, one hypothesis is that contours of high-dimensional fitness landscapes should enclose holes corresponding to large regions of lower fitness.

Holes Are Difficult to Detect

The shape of a hypervolume is almost always inferred on the basis of a set of observations. In this context, holes resist detection in more than two dimensions. In two dimensions, when raw data are plotted, a hole appears as a low-density or vacant region that is easily detected by an investigator. This approach is viable when visualizing simple data—for example, precipitation-temperature niches of *Eucalyptus* species (Austin et al. 1990)—but fails for $n \geq 3$. In higher dimensions, data visualized with multiple bivariate projections do not indicate holes because additional data points will almost certainly lie above or below the vacant region of interest (fig. 2). "Slicing" through data could alternatively reveal holes, but it is very inefficient in high dimensions. This approach requires choosing a partic-

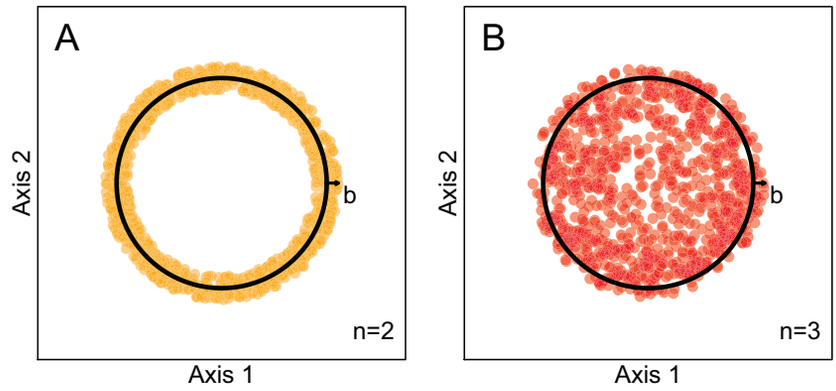

Figure 2: The detectability of holes varies with dimensionality n and bandwidth choice b. Consider a hyperannulus (black lines) data set corresponding to a ring in $n = 2$ dimensions and a hollow sphere in $n = 3$ dimensions. A hole is easy to detect in $n = 2$ dimensions (A) but disappears in $n = 3$ dimensions (B) because bivariate projections obscure interior features.

ular two-dimensional projection, then specifying constant (or near-constant) values for all combinations of $n - 2$ other axes. This approach therefore has a geometrically large parameter space and is not an effective way to detect holes. Thus, holes may exist in many sets of observations but have not been found for lack of appropriate tools.

A Probabilistic Approach to Inferring Holes

I present an approach to determine, for a given observational data set, the subregions that are empty relative to a given expectation. Briefly, the algorithm assumes that the data points are sampled from a probability density function with a certain shape **H**. The shape of **H** is inferred on the basis of padding each observation by a certain amount characterized by a kernel bandwidth vector b. The shape of **H** is then compared to the shape of a baseline expectation **B**, and the low-probability regions of **H** that are enclosed within **B** are inferred to be a set of holes $\mathbf{N} = \{\mathbf{N}_i\}$. These steps are implemented in the *hypervolume* R package (ver. 1.3.0 and newer), freely available on CRAN (http://cran.r-project.org). Specifically, the algorithm carries out the following six steps, also illustrated in figure 3.

Step 1: Obtain input data. The algorithm takes as input a $m \times n$ matrix of m data points in an n-dimensional space. The correct number and identity of axes must be determined before analysis. Axes must also be continuous. Categorical data can be converted to continuous data by ordination—for example, principal coordinates analysis after Gower transformation—although caution should be taken if the transformation is nonmetric. Axes must be on comparable scales so that Euclidean distances can be computed—for example, via rescaling and recentering—relative to a global range for the axis. Last, data must not have any missing values along any axis because each observation must be placed in an n-dimensional space; observations with missing measurements can be omitted. Gap-filling techniques (Rubin 1996; Schrodt et al. 2015) are not recommended because they will fill in the holes of interest.

Step 2: Compute a hypervolume **H** *that encloses all data points (*hypervolume *R function).* The algorithm assumes that the data are sampled from an unknown probability distribution and then computes a kernel density estimate using a hyperbox kernel with bandwidth vector b (corresponding to a diagonal smoothing matrix with variable amounts of smoothing along each axis). Vector values for b are chosen using the *estimate_bandwidth* function, discussed in more depth below. The algorithm then "slices" through the probability density function by choosing a quantile value such that the probability density above a certain threshold value, integrated across the entire space, is approximately equal to that quantile value. The algorithm uses a hyperbox kernel that decays to zero probability density after a finite distance. The resulting subset of n-dimensional space is inferred to be the hypervolume. The algorithm's output is a stochastic geometry representation of the hypervolume $\mathbf{R_H}$, constituting a set of r_H uniformly random n-dimensional points that are within the hypervolume's boundaries; an inferred volume, V_H; and a point density, $\rho_H = r_H / V_H$. The user chooses a value r such that the output hypervolume contains approximately $r_h = mr$ points; by default, $r = 1{,}000\,n$. This algorithm has been described elsewhere (Blonder et al. 2014).

Step 3: Compute a hypervolume **B** *with stochastic geometry representation* $\mathbf{R_B}$ *for the baseline expectation (*expectation_convex *or* expectation_maximal *R functions).* In the convex case, the algorithm computes a convex polytope (hull) that minimally encloses all or a subset of the data points. A "quickhull" beneath-beyond algorithm (Barber et al. 1996) is used to infer the hull volume V_B. Enclosing all data points is most accurate but computationally expensive—therefore,

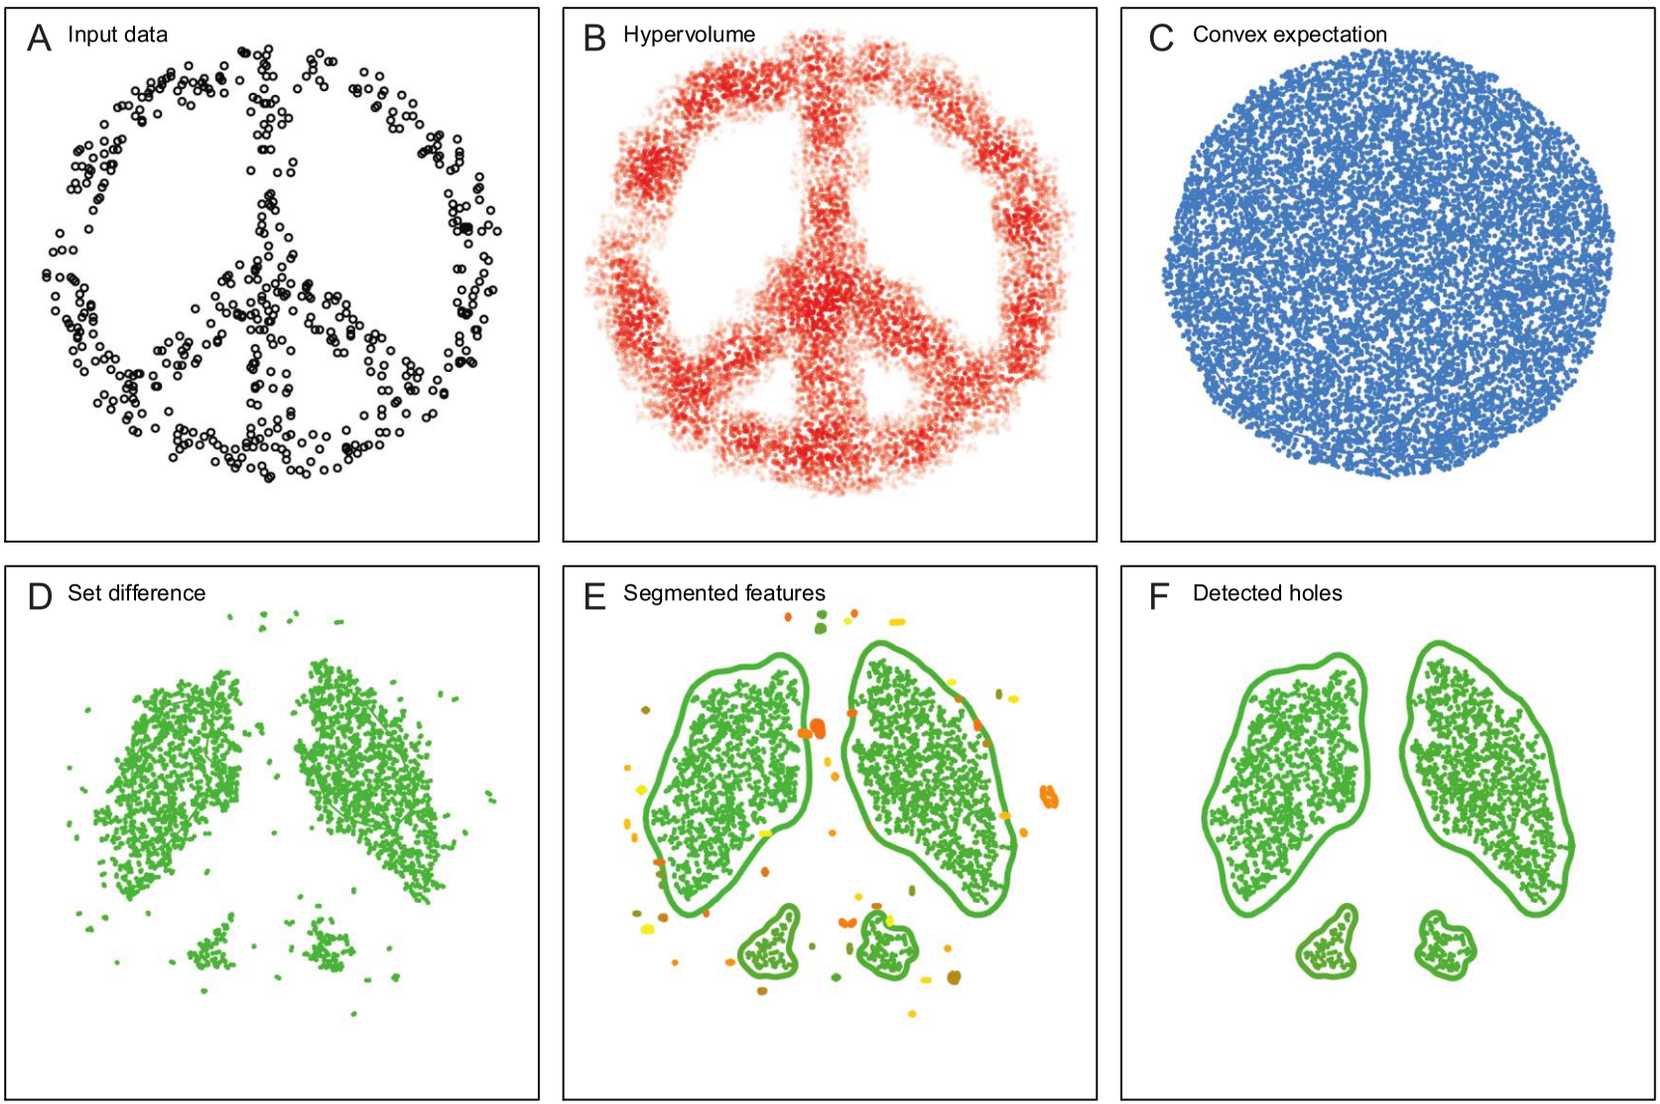

Figure 3: Major steps of the hole-detection algorithm, as illustrated in $n = 2$ dimensions. *A*, The user supplies an input data set (black circles). *B*, A hypervolume is constructed assuming that the data are sampled from an unknown probability distribution using the *hypervolume* function. The hypervolume is inferred by uniformly randomly sampling from a thresholded kernel density estimate (red points). The kernel bandwidth vector controls the amount of padding around each input data point. *C*, A convex expectation is generated by uniformly randomly sampling from a convex hull enclosing all data points (blue points) using the *expectation_convex* function. *D*, A set difference between the convex expectation and the hypervolume (green points) is computed using a stochastic approach to determine whether each point of each point cloud is enclosed by the other point cloud using the *hypervolume_holes* function. *E*, Optionally, a hierarchical clustering algorithm is used to segment the stochastic set difference into unique holes (variably colored contours) using the *hypervolume_segment* function. *F*, Optionally, spurious holes below a certain minimum volume (representing noise inherent in the stochastic methods) can be removed using the *hypervolume_prune* function, leaving behind a final set of holes (variably colored contours).

the algorithm by default estimates the convex hull on a random subset of h_B points, weighting points by their distance from their centroid. By default, $h_B = \min(m, 10^{\sqrt{n}})$ provides reasonable performance on many data sets. The user can choose a larger value to improve accuracy or a smaller value to improve runtime. The algorithm then uses rejection sampling to obtain a uniformly random set of points $\mathbf{R_B}$ from within the convex polytope **B**. Specifically, the algorithm generates a set of r_B uniformly random points from the minimal hyperbox enclosing all data points and then tests whether each of these points is "in" or "out" of the convex polytope using a dot product orientation test (D'Errico 2006). To optimize performance for the next step in the algorithm, r_B is chosen as the minimum of $10^{(3+\sqrt{n})}$ and a value estimated to match the point density of **H**, $\rho_H V_H$. The user can also choose a larger value to improve accuracy or a smaller value to improve runtime. In the maximal case, the investigator must supply this hypervolume from some external source of information. An appropriate hypervolume can be computed from a set of data points describing the available n-dimensional space. The algorithm uses random sampling to select a subset of points $\mathbf{R_B}$ from **B**. The user is recommended to choose $r_B = \rho_H V_H$ but can choose a larger value to improve accuracy or a smaller value to improve runtime.

Step 4: Perform a stochastic geometry set difference $\mathbf{S} = \mathbf{R_B} \backslash \mathbf{R_H}$ *to identify holes* **N** *(*hypervolume_holes *function).* First, $\mathbf{R_B}$ and $\mathbf{R_H}$ are subsampled to the same point density ρ_S. By default, the algorithm uses $\rho_S = \min(\rho_B, \rho_H, 10^{2+n}/V_B, 10^{2+n}/V_H)$. The user can choose a larger value of ρ_S to improve accuracy or a smaller value to improve runtime. As previously noted (Blonder et al. 2014), the characteristic distance d between uniformly random points with density ρ in n dimensions is $d = \rho^{-1/n}$. Thus, a test point in $\mathbf{R_B}$ is in **S** if it is "far away" from $\mathbf{R_H}$, that is, if the distance to every point in $\mathbf{R_H}$ is at least d. The algorithm performs this test on all random points in $\mathbf{R_B}$ to compute the final set difference, which includes r_S random points with a density of ρ_S and volume $V_S = \rho_S r_S$. There may be several isolated points in **S** that reflect random points that erroneously were included in the set difference because there were no random points in $\mathbf{R_H}$ that were within a distance d of these points. Such an effect can occur on the basis of small-number sampling effects. To eliminate these spurious points, the algorithm computes the minimum pairwise distance between all points in $\mathbf{R_H}$. The set of r_C points with a minimum distance greater than d is identified as **C**. The algorithm then returns the remaining holes as hypervolume **N**, defined as $\mathbf{N} = \mathbf{S} - \mathbf{C}$, with point density $\rho_N = \rho_S$ and volume $V_N = \rho_N(rS - rC)$.

Step 5 (optional): Segment unique holes $\{\mathbf{N}_i\}$ *from the set difference (*hypervolume_segment *function).* **N** contains a uniformly random set of points sampled from the true distribution of holes, but disjunct holes are not distinguished. To segment unique holes, the algorithm computes the pairwise distance between all points in **N** and then performs a hierarchical clustering analysis on this distance matrix. Clusters that are separated by a distance of at least d^* are inferred to be a set of I unique holes $\{\mathbf{N}_i\}$. By default, $d^* = nd$.

*Step 6 (optional): Remove small holes (*hypervolume_prune *function).* $\{\mathbf{N}_i\}$ may contain some holes that do not exist and have been falsely detected. These holes arise because of small-number sampling in the various stochastic geometry algorithms. Any unique holes with volumes beneath a certain investigator-specified threshold κ are considered to be spurious and removed, resulting in a final set of $J \leq I$ holes, $\{\mathbf{N}_j\}$. Note, however, that this function also will remove true holes that are small and should be used with caution.

Bandwidth Choice

The kernel bandwidth vector b can be varied to reflect the investigator's belief about the sparseness of the data set. If the bandwidth vector elements are small relative to the number of data points, then each data point will become disjunct from other points in the kernel density estimate, and many holes will be inferred. On the other hand, if the bandwidth vector elements are large relative to the number of data points, then small holes will not be inferred.

The investigator can manually choose a bandwidth vector that reflects their belief about the balance between these two types of uncertainty. For example, if variation below the scale of measurement exists (e.g., intraspecific trait variation for a community trait hypervolume), then the bandwidth vector elements can be chosen to match this range of variation. Alternatively, the bandwidth vector can be automatically estimated by minimizing asymptotic mean integrated squared error of the kernel density estimate relative to the data. The *estimate_bandwidth* function provides wrapper functions for a simple Silverman estimator for axis-wise optimization assuming normally distributed data (Silverman 1992) and two options for multidimensional optimization: a diagonal plug-in estimator with two-stage pilot estimation and a prescaling transformation (Wand and Jones 1994), and a diagonal cross-validation estimator with a two-stage pilot estimation and a prescaling transformation (Duong and Hazelton 2005).

The plug-in and cross-validation bandwidth estimators are computationally expensive. They scale exponentially with n, becoming impractical for more than $m = 10{,}000$ data points and $n = 6$ dimensions. Runtimes exceed 2 hours on a 2.7-GHz processor for the demonstration analysis of $m \approx 100$ points in $n = 4$ dimensions presented below.

Thus, the trade-off for increased accuracy in delineating holes is computational time.

Other Computational Considerations

The algorithm's performance improves with higher values of r_B (number of random points in baseline expectation), h_B (number of points used to define convex polytope), and ρ_S (density of points used for set operations). The default parameters are known to give reasonable performance for the simulation data sets described below. Higher values, subject to available computational resources, will always produce more accurate results. However, very high values of these parameters are not sufficient to yield perfect performance. The hyperbox kernels used in this analysis have "jagged edges" that yield approximation errors. Such errors do not tend to zero unless the number of data points m also becomes very large.

The computational complexity of these algorithms fundamentally limits the systems for which holes can be found. In general, stochastic geometry analyses are subject to the "curse of dimensionality": a linear increase in the number of dimensions requires an exponential increase in the number of random points used internally because higher-dimensional spaces require more points to fill them. The most limiting step is the calculation of the convex baseline expectation. The number of inequality constraints defining the convex polytope can scale exponentially with n, leading to very high memory costs and an exponential increase in the time cost of each rejection sample. Additionally, as n increases, the probability that a rejection sample is "in" also decreases with n, because increasingly less volume is occupied, leading to an exponential increase in the time cost of accepting a fixed number of random points. Set difference calculations can also be limiting when a nonconvex baseline is used. This calculation scales quadratically with the number of random points contained in the hypervolume and its baseline expectation. Because the number of random points should scale exponentially with n to maintain a fixed point density, the set difference operations can also scale exponentially in time and memory cost.

Sources of Error

Type I errors where holes are falsely inferred can be inadvertently caused by high dimensionality. High-dimensional spaces become increasingly sparse, such that it is difficult to "fill out" a data set with enough points that the entire object does not appear full of holes, whether or not these holes are real. For data sets that are described by a small number of points (e.g., a complete census of plants in a plot) or that cannot be realistically sampled to a higher number of points, it is not reasonable to examine hypervolume holes. Thus, holes are best detected in hypervolumes with intermediate dimensionality ($n \leq 6$ approximately, a range that seems to match empirical dimensionalities, e.g., for plant ecological strategies; Laughlin 2014).

Bias in data set sampling can also result in type I errors regardless of dimensionality. If some portions of the n-dimensional space are undersampled relative to others, then they will more likely be incorrectly inferred to be empty. For example, a climate niche may have a hole where environmental conditions correspond to countries in which fieldwork was not possible. These "unknown unknowns" (Jackson and Overpeck 2000; Jackson 2012) in a data set represent a fundamental source of uncertainty that affects all inferential statistical procedures.

Simulation Analyses

I explored a set of simulated data sets to determine whether the performance of this algorithm is acceptable in real-world scenarios, asking whether type I and type II error rates would be sufficiently low for variable combinations of data set size, dimensionality, and bandwidth choice.

I examined n-dimensional data sets described by m points sampled from either a solid or hollow hypersphere, with fixed inner and outer radii $r_{\min}$ and $r_{\max}$ (as in fig. 2). I ran the algorithm multiple times, comparing observed hypervolumes to a convex expectation. All other parameters were set to default values. I thus constructed a hypervolume with bandwidth b that includes a region A representing padding around the original data set:

$$A(n, r_{\min}, r_{\max}, b) = \{\vec{x}\} \subset \Re^n : \max(0, r_{\min} - b) \leq \|\vec{x}\| \leq r_{\max} + b.$$

In general, a hypersphere with radius r in n dimensions has volume Ω:

$$\Omega(n, r) = \frac{\pi^{n/2} r^n}{\Gamma[(n/2) + 1]},$$

where $\Gamma(x) = (x - 1)!$ is the gamma (factorial) function. Then A has a hole with volume V_{neg} that can be analytically calculated as

$$V_{\text{neg}}(n, r_{\min}, b) = \Omega(n, \max(0, r_{\min} - b)).$$

The overall volume of the hyperannular hypervolume V_{ann} is

$$V_{\text{ann}}(n, r_{\min}, r_{\max}, b) = \Omega(n, r_{\max} + b) - V_{\text{neg}}(n, r_{\min}, b).$$

Suppose now that the algorithm is run on a data set uniformly sampled from $A(n, r_{\min}, r_{\max}, b)$ with m points and

n dimensions with a bandwidth of b, relative to a convex expectation. Repeat this process t times. Let $V_i^*(m, n, r_{\min}, r_{\max}, b)$ be the inferred volume of holes for each algorithm run $1 \leq i \leq t$.

A type I (false positive) error rate, α, for this procedure can be defined in terms of the volume incorrectly inferred to be a hole when there are none, relative to the sum of the incorrectly inferred volume and the true volume of the hyperannulus:

$$\alpha = \frac{V_i^*(m, n, 0, r_{\max}, b)}{V_i^*(m, n, 0, r_{\max}, b) + V_{\text{ann}}(n, 0, r_{\max}, b)} = \frac{1}{1 + \dfrac{V_{\text{ann}}(n, 0, r_{\max}, b)}{V_i^*(m, n, 0, r_{\max}, b)}}.$$

Thus, when $V_i^*/V_{\text{ann}} \to 0$, $\alpha \to 0$ (i.e., 0% type I error when the inferred hole volume is very small relative to the true no-hole volume), while when $V_i^*/V_{\text{ann}} \to \infty$, $\alpha \to 1$ (i.e., 100% type I error rate when the inferred hole volume is very large relative to the true no-hole volume).

An approximate type II (false negative) error rate, $1 - \beta$, for this procedure is 1 minus a normalized difference between the average volume correctly inferred to be a hole when there is one and the true volume of the hole:

$$1 - \beta = \left| \frac{V_{\text{neg}}(n, r_{\min}, b) - V_i^*(m, n, r_{\min}, r_{\max}, b)}{V_{\text{neg}}(n, r_{\min}, b) + V_i^*(m, n, r_{\min}, r_{\max}, b)} \right| = \left| \frac{1 - \dfrac{V_i^*(m, n, r_{\min}, r_{\max}, b)}{V_{\text{neg}}(n, r_{\min}, b)}}{1 + \dfrac{V_i^*(m, n, r_{\min}, r_{\max}, b)}{V_{\text{neg}}(n, r_{\min}, b)}} \right|.$$

Thus, when $V_i^*/V_{\text{neg}} \to 0$, $1 - \beta \to 0$ (i.e., 0% type II error when the inferred and true hole volume are equal), while when $V_i^*/V_{\text{neg}} \to 0$ or ∞, $1 - \beta \to 1$ (i.e., 100% type II error rate when the inferred and true hole volume are very different).

I explored all combinations of $m = 100$, 200, 400, or 800 data points; $n = 2$, 3, 4, 5, or 6 dimensions; and $b =$ 0.1, a value automatically selected using a Silverman estimator ($b = 0.31 \pm 0.08$ SD), or 0.7. In all analyses I chose $r_{\min} = 0.8$ and repeated analyses for each parameter combination $t = 10$ times. I did not perform the optional hole segmentation and pruning (i.e., steps 5–6 of the algorithm). I ran each analysis on a 2.7-GHz machine with 8 GB of RAM using R version 3.1 on Linux.

Results for auto-selected bandwidths are most likely to reflect default usage cases. Performance in this scenario was reasonable (fig. 4*A*–4*C*). Type I error rates were always lower than 2% (mean: 0.3%), indicating that the algorithm rarely detects holes when they do not truly exist. Type II error rates averaged 39% but were much more variable. Error rates decreased to values closer to 30% for larger data sets and dimensionalities. Thus, the algorithm finds many but not all holes when they do truly exist.

In the case of a very small fixed bandwidth (0.1; fig. 4*D*–4*F*), type I error rates obtained a mean value of 12% and were always lower than 26%. Type II error rates averaged 17% and were never worse than 32%. The higher type I error rate likely arises because the hyperannulus is not completely resolved by the low number of data points—that is, it inadvertently contains small holes outside of the central hole that are detected by the algorithm. The lower type II error rate likely arises because the inner boundary of the hyperannulus becomes less blocky and jagged when estimated with a lower kernel bandwidth, so that the algorithm is better able to match the true shape of the inner hole.

In the case of a very high fixed bandwidth (0.7; fig. 4*G*–4*I*), type I error rates obtained a mean value of 0.1% and were always lower than 1.2%. Type II error rates averaged 81% and often reached 100%. This situation likely arises because large bandwidth values mean that the hyperannulus is delineated at a very blocky resolution; thus, the hole cannot be accidentally detected in the wrong location, but neither can it be properly detected in the correct location.

These contrasting results at different bandwidth values demonstrate that the qualitative conclusions taken from this type of analysis depend sensitively on the bandwidth used by the investigator. The automatic bandwidth selector provides reasonable performance, but lower bandwidth values can be used to minimize type II error rates, and higher bandwidth values can be used to minimize type I error rates.

Across usage cases, the algorithm's performance was only weakly sensitive to the number of data points or the particular random replicate. Thus, analyses do not need to be run multiple times to obtain robust results and are possible even on data sets of approximately $m = 100$ points.

Some of the variation in the scaling of error rates with dimensionality is due to the heuristics used to select the internal computational parameters r_B, h_B, and ρ_S. Overriding these heuristics and setting higher than default values will improve both error rates at the cost of longer computational times. Using default internal parameters, runtimes (fig. 4*C*, 4*F*, 4*I*) were under 2 minutes for all analyses of four dimensions or fewer and were under 2 hours for all five-dimensional analyses. However, runtimes were approximately 45 hours for some six-dimensional analyses, suggesting that computational limitations do provide a relevant practical boundary on when the algorithms can be applied. These long runtimes are driven primarily by the exponential scaling with dimensionality of rejection sampling a fixed

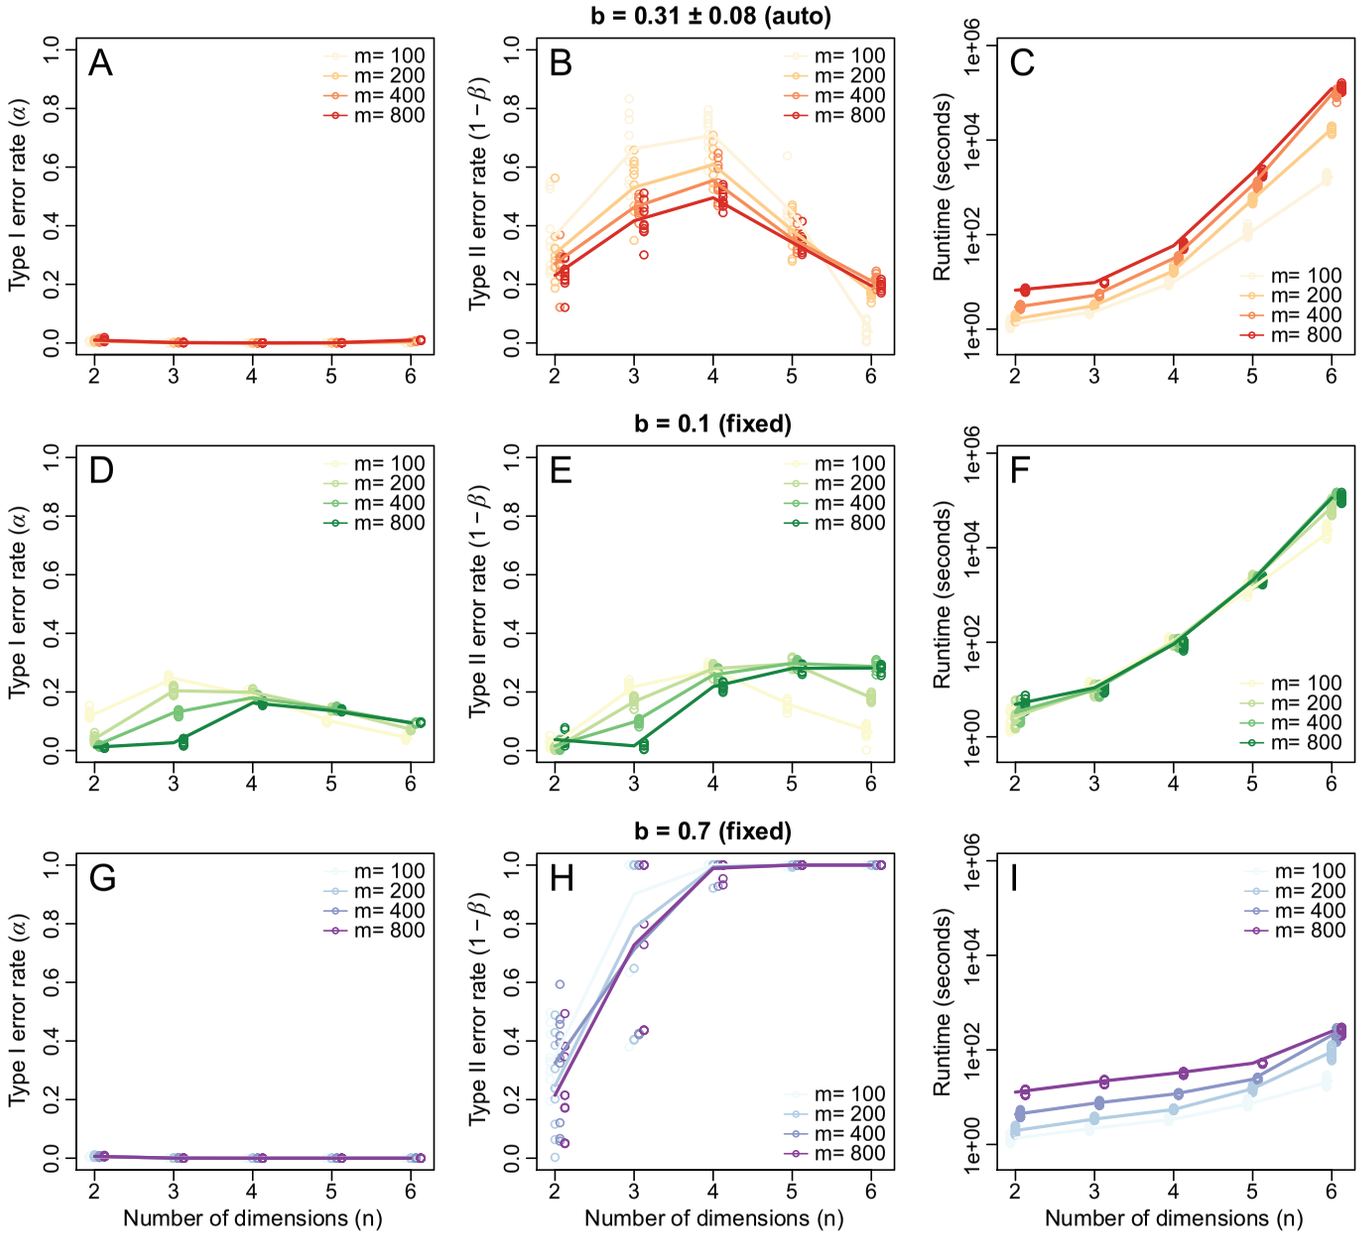

Figure 4: Summary of simulation analysis for holes inferred in a hollow or solid hypersphere data set sampled with m points in n dimensions using kernel bandwidth b. *A–C*, Kernel bandwidth b automatically selected using a Silverman estimator; *D–F*, b fixed to a very low value; *G–I*, b fixed to a very high value. *A*, *D*, *G*, Type I error rate inferred from the holes incorrectly found in the solid hypersphere; *B*, *E*, *H*, Type II error rate inferred from the holes correctly detected in the hollow hypersphere; *C*, *F*, *I*, approximate runtimes (note logarithmic *Y*-axis) using a 2.7-GHz machine with 8 GB of RAM.

number of uniformly random points from a convex hull and are independent of this software implementation.

Demonstration Analysis with Empirical Data

As proof of concept, I examined evidence for vacant niches in the context of invasion ecology. Vacant niches exist in a community if a region of a community's trait or morphological space is empty relative to a convex expectation (fig. 5). These vacant niches could have several causes, for example, (1) no colonization yet by an appropriate species in the regional pool, (2) absence of an appropriate species in the regional pool, or (3) prevention of occupation of this region by species interactions.

To determine whether any of these processes could be operating, I examined morphological hypervolume of a community of finches co-occurring on Isabela Island (Galapagos) on the basis of collections from the Snodgrass and Heller (1904) expedition. Code and data to replicate this analysis are available by running *demo(holes_finch)*.

The data set includes $m = 166$ observations for male and female individuals for eight species along $n = 4$ axes: body length, wing length, tail length, and beak width. Data were standardized before analysis by subtracting axis means and dividing by axis standard deviations.

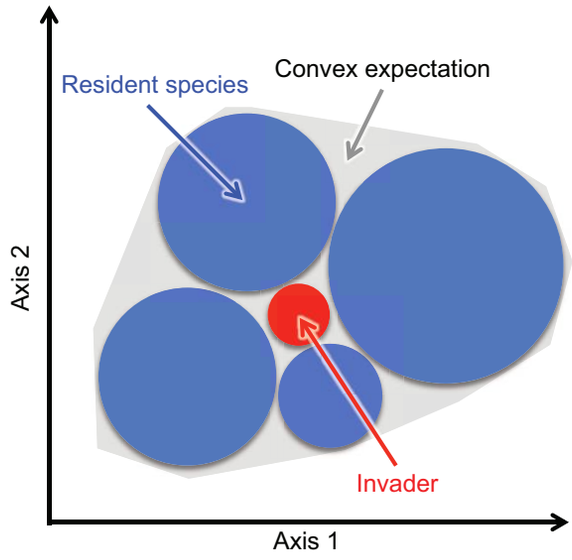

Figure 5: Cartoon example of holes mediating invasion. Consider a hyperspace describing the functional hypervolume occupied by different species in a community. Here four species (blue circles) are shown relative to the convex expectation (gray polygon). Although the convex expectation determines the limits to the ecological strategies that are viable in this community, there may still be unoccupied regions of ecological strategy space—for example, the center gray region. The existence of this central hole suggests that a species with an appropriate ecological strategy (small red circle) might be able to successfully invade this community.

I computed a community hypervolume and convex expectation using default R package parameters and bandwidth values that were the default Silverman estimate multiplied by factors of 0.1 to 2.0, yielding actual values of $b = 0.076$ to $b = 1.5$. I repeated all analyses 10 times for each value of b to assess uncertainty arising from the stochastic geometry algorithms. I also calculated estimates of the "correct" value for b using mean values reported by the Silverman, plug-in, and cross-validation estimators. I did not perform the optional hole segmentation and pruning.

The ratio between the volume of the detected holes and the volume of the convex expectation ($V_{\text{holes}}/V_{\text{convex}}$) decreased as the bandwidth b increased (fig. 6*A*). Different bandwidth estimators all led to detection of holes, with very limited variation across replicate analyses. The Silverman estimator (mean $b = 0.76$) led to $V_{\text{holes}}/V_{\text{convex}} = 0.011 \pm 0.0027$ SD, averaged across all replicate analyses. This number, while apparently very small, actually corresponds to approximately $0.011^{1/4} = 32\%$ of each axis being unoccupied. The cross-validation estimator (mean $b = 0.60$) led to $V_{\text{holes}}/V_{\text{convex}} = 0.083 \pm 0.011$ SD, corresponding to approximately 53% of each axis being unoccupied. Last, the plug-in estimator (mean $b = 0.587$) led to $V_{\text{holes}}/V_{\text{convex}} = 0.011 \pm 0.0051$ SD, corresponding to approximately 57% of each axis being unoccupied. Thus, the qualitative inference (presence of holes) did not depend on the bandwidth estimator or the replicate analysis.

The presence of holes in the Galapagos finch data set (with b determined by the plug-in estimator) is readily visible in a pairs plot (fig. 6*B*) or rotating three-dimensional projection (video 1, available online). The regions of morphospace that are empty correspond to a mean centroid hole location of (body length, wing length, tail length, beak width) in terms of transformed coordinates (standard deviation above the mean) of (1.56, 1.65, 1.39, 0.87) or in untransformed coordinates (mm) of (143, 75, 47, 9). These values are each occupied along single dimensions, suggesting that single-dimensional analyses would not have been able to identify this empty region. Additionally, these coordinates can be used to identify candidate species that could potentially invade this island. I calculated species-mean morphology axis values for all individuals collected by the Snodgrass and Heller (1904) expedition with the smallest rescaled differences relative to the Isabela Island hole centroid. Of these, the smallest difference was for *Geospiza scandens* subsp. *scandens*, with species-mean morphological coordinates of (137, 70, 45, 8). This species was recorded during the expedition as being collected only on the nearby Santiago Island. Thus, one hypothesis is that an experimental introduction of this species to Isabela Island might be more successful than that of other finch

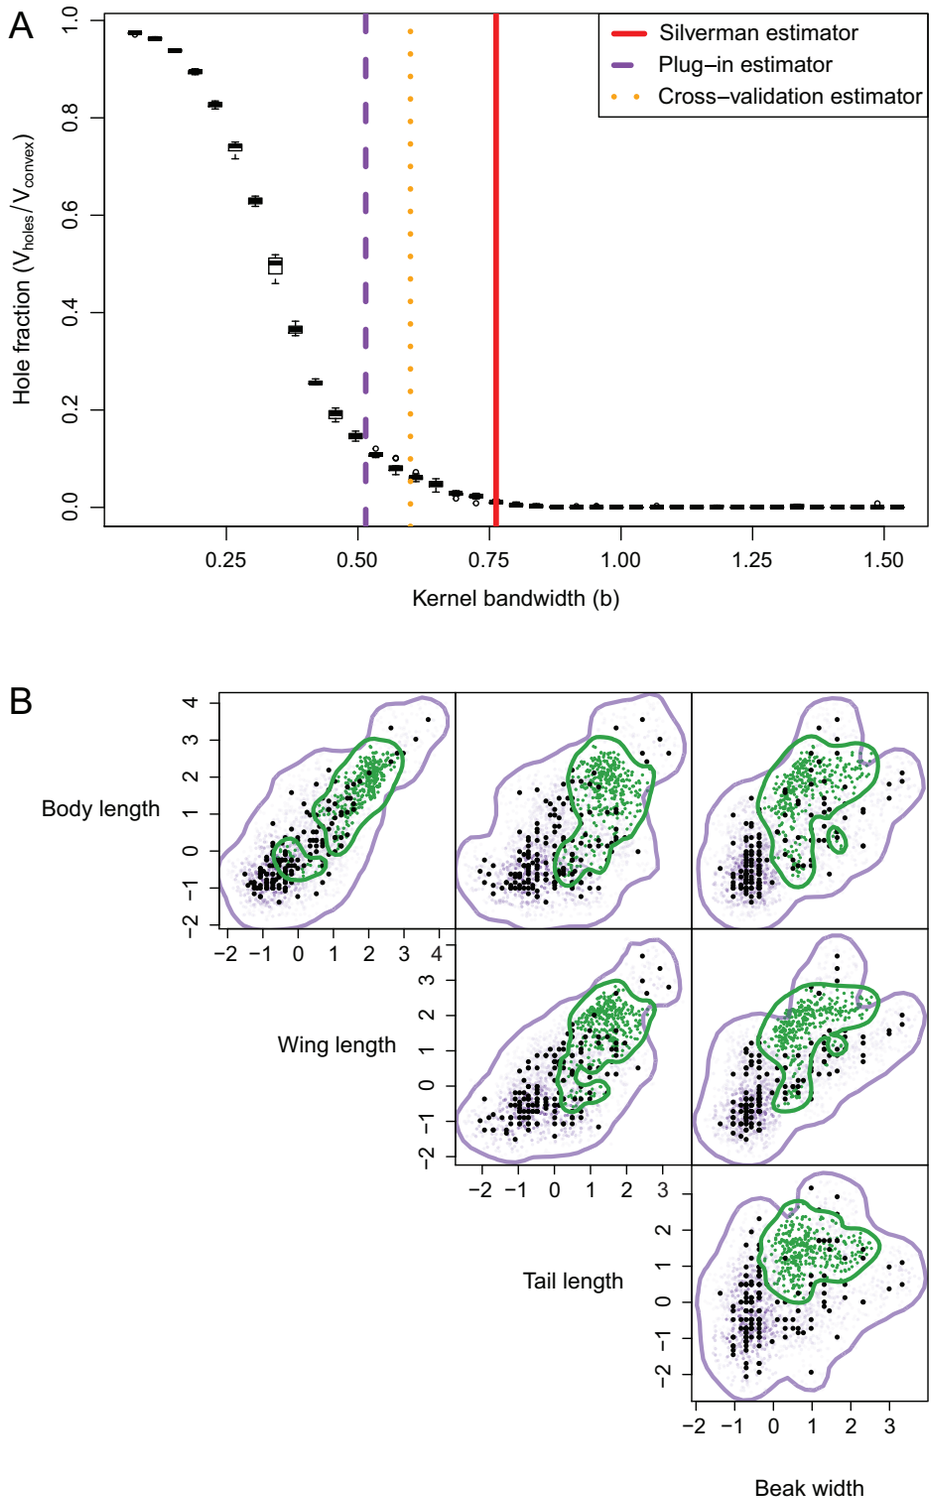

Figure 6: Vacant niche analysis for five species of finch co-occurring on Isabela Island using $n = 4$ morphological axes. *A*, Hole detection depends on the bandwidth value b used in the hypervolume analysis. The ratio between the volume of a detected hole and the volume of the convex expectation (V_{holes}/V_{convex}) is shown for a range of values of b. Boxes represent distribution of 10 replicate analyses of the data set plotted in *B*. The mean values of b proposed by different kernel bandwidth estimators (Silverman, plug-in, cross validation) are indicated as vertical lines. *B*, Pairs plot for hypervolumes computed using the plug-in estimator for b. Each panel is a two-dimensional projection of the four-dimensional space. Points represent observed data (black), uniformly random points sampled from the hypervolume (purple), and uniformly random points sampled from a detected hole (green). Lines represent approximate contour boundaries for each two-dimensional projection of the four-dimensional space. A three-dimensional projection of this hypervolume is animated as video 1 (available online).

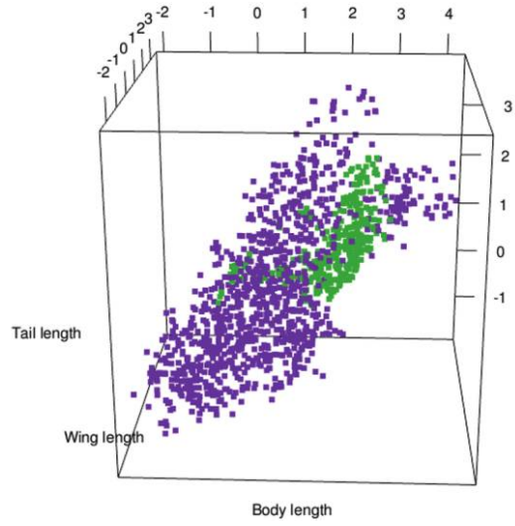

Video 1: Still photograph from a video (video 1, available online) showing morphological community hypervolume for five Galapagos finches co-occurring on Isabela Island, as seen in an animated three-dimensional project of tail length, wing length, and body length. Dots represent uniformly random points sampled from the community hypervolume (purple) or an inferred four-dimensional hole (green).

species; alternatively, the hole centroid may indicate an ecological strategy that is not viable on Isabela Island for other reasons, for example, lack of resources available to an organism of such morphology.

The point of this demonstration analysis is not to propose that specific experiments should be carried out in the absence of more detailed knowledge of these finches' biology. The point is also not to ascribe particular causes to observed patterns. Rather, this analysis highlights the utility of the hole-detection approach for generating novel types of hypotheses. It also shows that well-known data sets may be described in terms of hypervolumes with holey geometries that are potentially consistent with unexplored mechanisms.

Guidance on Usage

The algorithm can be used on data sets of at least $m = 100$ observations in $n \leq 6$ continuous dimensions. A rule of thumb is that holes can be detected successfully when $\log_e(m) > n$. In these cases, for a given choice of bandwidth vector b, detecting a hole very likely indicates that an empty region exists in the hypervolume, while detecting the absence of a hole probably (but not certainly) indicates that no empty region exists in the hypervolume. A good value of b can be automatically chosen using a Silverman estimator at very low computational cost. Better estimates of b can be obtained using a plug-in or cross-validation estimator with the trade-off of very high computational cost. In cases of small data sets, the investigator must be content with a fairly high type II error rate (low power). Future mathematical insights will hopefully improve the situation further.

Looking Forward

Holes may be present but not yet detected in many data sets. To address this problem, I have provided a first set of concepts and algorithms that can be applied across biological scales and axes. The open challenges are now to begin detecting the presence and absence of holes and to develop models of ecological and evolutionary dynamics that make specific predictions for the conditions under which such holes are expected. Holes may not often exist—but we will not know until we begin to look.

Acknowledgments

I thank R. Colwell, W. Driscoll, N. Gotelli, C. Lamanna, S. Michaletz, L. Sloat, J. Soberón, J.-C. Svenning, C. Violle, and the Norwegian University of Science and Technology (NTNU) Beehive group for their thoughtful comments.

Literature Cited

Adler, P. B., A. Fajardo, A. R. Kleinhesselink, and N. J. B. Kraft. 2013. Trait-based tests of coexistence mechanisms. Ecology Letters 16: 1294–1306.

Arnold, S. J. 1992. Constraints on phenotypic evolution. American Naturalist 140:S85–S107.

Austin, M. P. 1992. Modelling the environmental niche of plants: implications for plant community response to elevated CO_2 levels. Australian Journal of Botany 40:615–630.

———. 1999. A silent clash of paradigms: some inconsistencies in community ecology. Oikos 1999:170–178.

Austin, M. P., A. O. Nicholls, and C. R. Margules. 1990. Measurement of the realized qualitative niche: environmental niches of five *Eucalyptus* species. Ecological Monographs 60:161–177.

Barber, C. B., D. P. Dobkin, and H. Huhdanpaa. 1996. The quickhull algorithm for convex hulls. ACM Transactions on Mathematical Software 22:469–483.

Blonder, B., C. Lamanna, C. Violle, and B. J. Enquist. 2014. The *n*-dimensional hypervolume. Global Ecology and Biogeography 23: 595–609.

Burrows, M. T., D. S. Schoeman, A. J. Richardson, J. G. Molinos, A. Hoffmann, L. B. Buckley, P. J. Moore, et al. 2014. Geographical limits to species-range shifts are suggested by climate velocity. Nature 507:492–495.

Channell, R., and M. V. Lomolino. 2000. Trajectories to extinction: spatial dynamics of the contraction of geographical ranges. Journal of Biogeography 27:169–179.

Chesson, P. 2000. Mechanisms of maintenance of species diversity. Annual Review of Ecology and Systematics 31:343–366.

Colwell, R. K., and D. J. Futuyma. 1971. On the measurement of niche breadth and overlap. Ecology 52:567–576.

Colwell, R. K., and T. F. Rangel. 2009. Hutchinson's duality: the once and future niche. Proceedings of the National Academy of Sciences of the USA 106:19651–19658.

Cornwell, W. K., D. W. Schwilk, and D. D. Ackerly. 2006. A trait-based test for habitat filtering: convex hull volume. Ecology 87: 1465–1471.

D'Errico, J. 2006. Efficient test for points inside a convex hull in *n* dimensions. http://www.mathworks.com/matlabcentral/fileexchange/10226-inhull/content/inhull.m.

Duong, T., and M. L. Hazelton. 2005. Cross-validation bandwidth matrices for multivariate kernel density estimation. Scandinavian Journal of Statistics 32:485–506.

Elton, C. S. 1958. The ecology of invasions by plants and animals. Methuen, London.

Enquist, B. J., J. Norberg, S. P. Bonsor, C. Violle, C. T. Webb, A. Henderson, L. L. Sloat, and V. M. Savage. 2015. Scaling from traits to ecosystems: developing a general trait driver theory via integrating trait-based and metabolic scaling theories. Advances in Ecological Research 52:249–318.

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

Holt, R. D. 2009. Bringing the Hutchinsonian niche into the 21st century: ecological and evolutionary perspectives. Proceedings of the National Academy of Sciences of the USA 106:19659–19665.

Hutchinson, G. 1957. Concluding remarks. Cold Spring Harbor Symposium on Quantitative Biology 22:415–427.

Jackson, A. L., R. Inger, A. C. Parnell, and S. Bearhop. 2011. Comparing isotopic niche widths among and within communities: SIBER—stable isotope Bayesian ellipses in R. Journal of Animal Ecology 80:595–602.

Jackson, S. T. 2012. Representation of flora and vegetation in Quaternary fossil assemblages: known and unknown knowns and unknowns. Quaternary Science Reviews 49:1–15.

Jackson, S. T., and J. T. Overpeck. 2000. Responses of plant populations and communities to environmental changes of the late Quaternary. Paleobiology 26:194–220.

Lai, H. R., M. M. Mayfield, J. M. Gay-des-combes, T. Spiegelberger, and J. M. Dwyer. 2015. Distinct invasion strategies operating within a natural annual plant system. Ecology Letters 18:336–346.

Lamanna, C., B. Blonder, C. Violle, N. J. B. Kraft, B. Sandel, I. Šímová, J. C. Donoghue, et al. 2014. Functional trait space and the latitudinal diversity gradient. Proceedings of the National Academy of Sciences of the USA 111:13745–13750.

Laughlin, D. C. 2014. The intrinsic dimensionality of plant traits and its relevance to community assembly. Journal of Ecology 102:186–193.

Laughlin, D. C., C. Joshi, S. J. Richardson, D. A. Peltzer, N. W. H. Mason, and D. A. Wardle. 2015. Quantifying multimodal trait distributions improves trait-based predictions of species abundances and functional diversity. Journal of Vegetation Science 26:46–57.

Leibold, M. A. 1995. The niche concept revisited: mechanistic models and community context. Ecology 76:1371–1382.

Leibold, M. A., M. Holyoak, N. Mouquet, P. Amarasekare, J. M. Chase, M. F. Hoopes, R. D. Holt, et al. 2004. The metacommunity concept: a framework for multi-scale community ecology. Ecology Letters 7:601–613.

Lorenzen, E. D., D. Nogues-Bravo, L. Orlando, J. Weinstock, J. Binladen, K. A. Marske, A. Ugan, et al. 2011. Species-specific responses of Late Quaternary megafauna to climate and humans. Nature 479:359–364.

Lynch, M., and B. Walsh. 1998. Genetics and analysis of quantitative traits. Vol. 1. Sinauer, Sunderland, MA.

Maguire, B., Jr. 1973. Niche response structure and the analytical potentials of its relationship to the habitat. American Naturalist 107:213–246.

McGill, B. J., B. J. Enquist, E. Weiher, and M. Westoby. 2006. Rebuilding community ecology from functional traits. Trends in Ecology and Evolution 21:178–185.

Pearman, P. B., A. Guisan, O. Broennimann, and C. F. Randin. 2008. Niche dynamics in space and time. Trends in Ecology and Evolution 23:149–158.

Rubin, D. B. 1996. Multiple imputation after 18+ years. Journal of the American Statistical Association 91:473–489.

Schrodt, F., J. Kattge, H. Shan, F. Fazayeli, J. Joswig, A. Banerjee, M. Reichstein, et al. 2015. BHPMF—a hierarchical Bayesian approach to gap-filling and trait prediction for macroecology and functional biogeography. Global Ecology and Biogeography 24:1510–1521. doi: 10.1111/geb.12335.

Silverman, B. W. 1992. Density estimation for statistics and data analysis. Chapman & Hall, London.

Snodgrass, R., and E. Heller. 1904. Papers from the Hopkins-Stanford Galapagos Expedition, 1898–99 XVI. Birds. Proceedings of the Washington Academy of Science 5:231–372.

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

Van Kleunen, M., E. Weber, and M. Fischer. 2010. A meta-analysis of trait differences between invasive and non-invasive plant species. Ecology Letters 13:235–245.

Villéger, S., N. W. H. Mason, and D. Mouillot. 2008. New multidimensional functional diversity indices for a multifaceted framework in functional ecology. Ecology 89:2290–2301.

Violle, C., B. J. Enquist, B. J. McGill, L. Jiang, C. H. Albert, C. Hulshof, V. Jung, et al. 2012. The return of the variance: intraspecific variability in community ecology. Trends in Ecology and Evolution 27:244–252.

Wand, M., and M. Jones. 1994. Multivariate plug-in bandwidth selection. Computational Statistics 9:97–116.

Weber, K. E. 1996. Large genetic change at small fitness cost in large populations of *Drosophila melanogaster* selected for wind tunnel flight: rethinking fitness surfaces. Genetics 144:205–213.

Weiher, E., and P. A. Keddy. 1995. Assembly rules, null models, and trait dispersion: new questions from old patterns. Oikos 74:159–164.

Williams, J. W., and S. T. Jackson. 2007. Novel climates, no-analog communities, and ecological surprises. Frontiers in Ecology and the Environment 5:475–482.

Associate Editor: Axel G. Rossberg
Editor: Yannis Michalakis