

Comparing isotopic niche widths among and within communities: SIBER – Stable Isotope Bayesian Ellipses in R

Andrew L. Jackson^{1*}, Richard Inger², Andrew C. Parnell³ and Stuart Bearhop²

¹Department of Zoology, School of Natural Sciences, Trinity College Dublin, Dublin 2, Ireland; ²Centre for Ecology & Conservation, School of Biosciences, University of Exeter, Cornwall Campus, Penryn, Cornwall, TR10 9EZ, UK; and

³Department of Statistics, School of Mathematical Sciences, University College Dublin, Dublin 4, Ireland

Summary

1. The use of stable isotope data to infer characteristics of community structure and niche width of community members has become increasingly common. Although these developments have provided ecologists with new perspectives, their full impact has been hampered by an inability to statistically compare individual communities using descriptive metrics.

2. We solve these issues by reformulating the metrics in a Bayesian framework. This reformulation takes account of uncertainty in the sampled data and naturally incorporates error arising from the sampling process, propagating it through to the derived metrics.

3. Furthermore, we develop novel multivariate ellipse-based metrics as an alternative to the currently employed Convex Hull methods when applied to single community members. We show that unlike Convex Hulls, the ellipses are unbiased with respect to sample size, and their estimation via Bayesian inference allows robust comparison to be made among data sets comprising different sample sizes.

4. These new metrics, which we call SIBER (Stable Isotope Bayesian Ellipses in R), open up more avenues for direct comparison of isotopic niches across communities. The computational code to calculate the new metrics is implemented in the free-to-download package Stable Isotope Analysis for the R statistical environment.

Key-words: Bayesian Inference, community ecology, R statistics program, statistical methods, trophic niche width

Introduction

The ecological niche concept has undergone a renaissance in recent years (Newsome *et al.* 2007; Soberon 2007). This is unsurprising given that it can provide insights into a wide variety of both ecological and evolutionary problems (Martínez del Río *et al.* 2009). Refinements of Hutchinson's idea that an ecological niche can be represented as an n -dimensional hypervolume have proposed that it be partitioned into scenopoetic axes, representing environmental components of niche space, and bionomic axes, which refer mostly to the trophic components of niche space (Hutchinson 1957, 1978). Recently, it has been argued that location on these axes may be quantified using stable isotopic ratios (Bolnick *et al.* 2003; Bearhop *et al.* 2004) and have been formalized in the concept of the 'isotopic niche' (Newsome *et al.* 2007). Stable isotope ratios can be used in this respect because the values measured in consumer tissues are tightly linked to those in their diet.

For example, the ratio of heavy to light stable nitrogen isotopes ($^{15}\text{N}/^{14}\text{N}$) increases in a stepwise fashion with each trophic level (Minagawa & Wada 1984; Wada *et al.* 1987; Fry 1988) and thus is analogous to a bionomic axis. Whereas stable carbon isotopes ($^{13}\text{C}/^{12}\text{C}$), although they do alter slightly with trophic level, reflect the primary carbon sources within a food web which vary markedly between aquatic and terrestrial origins (Craig 1953; Chisholm, Nelson & Schwarcz 1982), and thus can be considered as both bionomic and scenopoetic axes. Additional isotopes, for example of oxygen and hydrogen can provide more detailed information on scenopoetic as they vary over geographical scales (Newsome *et al.* 2007). Isotopic data are routinely presented as bi-plots, where the isotopic values of animal tissues may be represented in δ -space (Newsome *et al.* 2007), and this essentially delineates an animal's isotopic niche. Although the isotopic niche is likely to be tightly correlated to the trophic niche, we stress that these are not the same and should not be confused. Nevertheless, the information contained within consumer stable isotope ratios is predominantly ecological in origin

*Correspondence author. E-mail: A.Jackson@tcd.ie

and therefore can be considered descriptors of key axes in Hutchinson's hypervolume, providing ecologically relevant information about the individual, population or community they represent. Additionally, by carefully selecting tissues with appropriate turnover rates in consumers, the variance among individuals in isotope space can be linked qualitatively to among-individual variation in diet (e.g. specialist and generalist) and hence is an integrated measure of niche width (Bearhop *et al.* 2004).

These concepts have been extended by Layman *et al.* (2007a), who suggested a series of additional measurements that could be made to generate metrics of niche structure from individuals to entire communities. Whilst concerns have been raised about these approaches (Hoeinghaus *et al.* 2006), it is clear that many of the issues can be dealt with by appropriate sampling strategies or baseline corrections (Layman & Post 2008) and they have been used to great effect in a number of studies (Darimont, Paquet & Reimchen 2009; Martinez del Rio *et al.* 2009; Olsson *et al.* 2009; Quevedo, Svanbäck & Eklöv 2009). Having said this, the approach will apply better to some systems than others (as is the case with all approaches to quantifying niche), and care should be taken as always to ensure one is satisfied of their suitability to their system and hypotheses. There are however a number of ways in which the current framework proposed by Layman *et al.* (2007a) might be advanced. Here, we outline several shortcomings of the existing approach and suggest a number of ways in which these may be solved.

Briefly, six niche/community metrics were proposed by Layman *et al.* (2007a), which together use summary information on the spread and extent of the data points. These are as follows:

1. $\delta^{15}\text{N}$ range (NR) providing information on the trophic length of the community;
2. $\delta^{13}\text{C}$ range (CR) giving an estimate of the diversity of basal resources;
3. total area of the convex hull encompassing the data points (TA) giving an indication of niche width;
4. mean distance to centroid (CD) providing additional information on niche width but also species spacing;
5. mean nearest neighbour distance (MNND) providing a measure of density and clustering of species within the community;
6. standard deviation of the nearest neighbour distance (SDNND) giving a measure of evenness of spatial density and packing.

To date, most interest appears to centre around the use of the convex hull area (TA) describing the niche width of the organism (Quevedo, Svanbäck & Eklöv 2009) or community in question (Layman *et al.* 2007a). There are two distinct situations in which the existing metrics have been applied. The first is to describe individual components of the community. Here, the metrics may be applied within a single group member within the community and the metrics calculated on a sample of n individual observations of a taxonomic or functional group of interest; e.g. to describe the isotopic niche

width of grey snappers *Lutjanus griseus* (Layman *et al.* 2007b) or perch *Perca fluviatilis* (Quevedo, Svanbäck & Eklöv 2009). The second approach is to apply the metrics to entire (or at least large parts of) communities of different species or functional groups (e.g. Layman *et al.* 2007a). In this instance, the data are distilled down to the mean $\delta^{15}\text{N}$ and $\delta^{13}\text{C}$ values for each of m taxonomic or functional groups, and the six metrics calculated on the resultant m data points.

One of the shortcomings identified herein is that application of the metrics to single community group members is sensitive to sample size, particularly across the range of sizes frequently encountered in ecological studies ($n < 50$). Such behaviour is highly undesirable in situations where sample sizes differ among samples within studies, or when comparisons across multiple studies are conducted. As an alternative we suggest using metrics based on standard ellipses (Batschelet 1981), comparable to SD in univariate cases, to draw inference on isotopic niche width instead of convex hulls and other extreme value metrics. Additionally, we present an alternative sample size correction for the standard ellipses, to our knowledge not previously published. These new methods will allow robust meta-analyses between studies that contain different sample sizes.

The second shortcoming of the current framework is that when applied across an entire community, the metrics do not incorporate any of the natural variability within the system into the subsequent summary statistics, and thus provide only a point estimate of each metric. A similar problem has recently been discussed with respect to isotopic mixing models (Jackson *et al.* 2009; Semmens, Moore & Ward 2009a; Semmens *et al.* 2009b). Ignoring the sampling and other sources of uncertainty makes any subsequent statistical comparisons among individual communities challenging at best. Here, we describe a method based on Bayesian inference techniques for propagating sampling error on the estimates of the means of community members to provide measures of uncertainty surrounding Layman *et al.*'s (2007a) metrics, which allows for robust statistical comparisons to be made among communities.

Finally, we compare our new Bayesian metrics to their existing analogues using extensive simulation techniques. The new methods are presented in detail here and made available to researchers via the package Stable Isotope Analysis in R (SIAR – Parnell *et al.* 2008, 2010) available to download from the packages section of the Comprehensive R Archive Network site (CRAN) – <http://cran.r-project.org/>. SIAR currently contains a Bayesian model for solving isotope mixing models and is now expanded to include additional functionality for analysing isotope data in the context of isotopic niche width.

Materials and methods

These methods are considered in two sections, the first dealing with the case for describing niche width on a single member of a community (e.g. a single taxonomic species), and the second when inference across multiple groups is desired (e.g. across foraging guilds or entire

communities). All analyses are conducted in the R statistical computing package (R Development Core Team, 2007). Throughout all our analyses, we make the assumption that the isotope data under consideration are adequately described by a multivariate normal distribution. Such an assumption seems appropriate for many instances based mechanistically on the nature of the data and also reference to published data (Inger *et al.* 2006, 2008). The validity of this assumption can and should be tested using one of several statistical tests (e.g. Székely & Rizzo 2005; e.g. Tan *et al.* 2005); by default, we provide output from the Shapiro–Wilk test using `mshapiro.test()` function in R (Royston 1982) as a guide, but not definitive rule, for suitability. Where this assumption is not met for a particular data set, then alternative analytical options could be employed such as transformations or nonparametric measures of dispersion. Indeed, such scenarios may represent ecologically interesting processes such as phenotypic variation or plasticity rather than simply being a nuisance.

CASE 1 – METRICS ON SINGLE MEMBERS

Standard ellipses

Central to the proposed new method is the standard ellipse, which is to bivariate data as SD is to univariate data (Batschelet 1981). The standard ellipse describing some data x and y is underpinned by its associated covariance matrix $\left(\Sigma = \begin{bmatrix} \sigma_x^2 & \text{cov}(x,y) \\ \text{cov}(y,x) & \sigma_y^2 \end{bmatrix}\right)$, which defines its shape and size, and the means of the x and y that define its location. The eigenvalues (λ) and eigenvectors (v) of Σ then give the lengths of the semi-major axis $a = \lambda_1^{-1}$ and the semi-minor axis $b = \lambda_2^{-1}$, and the angle of a with the x axis $\theta = \sin^{-1}(v_{12})$. Plotting is then straight forward (Fig. 1), as is computation of the area of the resultant ellipse (SEA, Standard Ellipse Area) which is given by $SEA = \pi ab$. Functions provided in SIAR will perform these tasks for the user.

Estimating the parameters of the ellipse from data

As the ellipse is likely to be estimated on a sample rather than the entire population, it is necessary to first determine the covariance matrix of a sample. In the univariate case, the sample variance (s^2) provides an unbiased estimate of the population variance (σ^2) for data x and y . The covariance between the data x and y similarly has an unbiased algebraic form. Whilst the SE of these individual compo-




Fig. 1. Example standard ellipse based on p135 (Batschelet 1981) showing the sample data (filled circles), standard ellipse, semi-major axis a and semi-minor axis b as labelled (solid lines).

nents of the covariance matrix can be calculated on their marginal distributions, to our knowledge, their joint distribution does not have a described analytical solution and variation is underestimated in small sample sizes (Ricklefs & Nealen 1998) meaning that uncertainty around these mean estimates is difficult to quantify using algebraic formulae.

Alternatively, it is straight forward to estimate the covariance matrix of the data using Bayesian inference. In this case, vague normal priors are assigned to the means, and a vague Inverse-Wishart prior is used for Σ (this is a standard conjugate prior for covariance matrices, McCarthy 2007). The posterior estimate of Σ is then simulated using Markov chain Monte Carlo (MCMC) (we employ the function `rmultireg()` in the R package `bayesm`; Rossi, Allenby & McCulloch 2005) according to the formulation:

The Priors

$$\mu_x \sim \text{dnorm}(0, \sigma^2 = 10^3) \tag{eqn 1a}$$

$$\mu_y \sim \text{dnorm}(0, \sigma^2 = 10^3) \tag{eqn 1b}$$

$$\Sigma \sim \text{wishart}^{-1}\left(\rho = 2, V = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}\right) \tag{eqn 2}$$

The Likelihood

$$Y_i \sim \text{MVN}([\mu_x, \mu_y], \Sigma) \tag{eqn 3}$$

In the Bayesian approach, the posterior estimate is calculated by combining the priors and the likelihoods, and comprises a set of k iterative draws from the MCMC simulation. For each of these k draws, posterior values of the means μ_x and μ_y and covariance matrix Σ value are returned, which can be used to construct an ellipse and derive metrics such as the area which we refer to as SEA_B . This process is repeated for all simulated values producing a range of probable values for the calculated metric reflecting the uncertainty in the estimated Σ . These estimated values can then be compared directly in a probabilistic manner in terms of how similar or not they are.

Sensitivity of ellipses to sample size

We first suppose a very simple case: a population of organisms' isotope data are described by a multivariate normal distribution with means $\mu_x = \mu_y = 0$ and covariance matrix $\Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$. That is, their data form a circle around the origin, whose standard ellipse describes a circle with radius equal to 1.

A typical sampling exercise would involve sampling from some population and using these values to draw inference on the population as a whole and make comparisons between populations. We sample from this population, drawing n samples from a multivariate normal distribution centred on the origin with $Y_{i=1:n} \sim \text{MVN}([0 \ 0], \Sigma)$ and calculate all six of Layman *et al.*'s (2007a) metrics from the resultant Y values as well as the SEA which we propose as a more robust alternative. We varied sample size from the set $n = [3, 4 \dots 29, 30, 40 \dots 100, 200 \dots 1000]$ and repeated each sampling process and associated metric calculations 10^4 times for each value of n .

Two examples are shown in Fig. 2 for $n = 10$ and $n = 200$. Note that as the convex hull is drawn around the extreme most points, the area encompassed by the hull increases with sample size, as the addition of more points can only have either no effect or make the hull larger, not smaller. The standard ellipse on the other hand contains c. 40% of the data regardless of sample size (Batschelet 1981). The

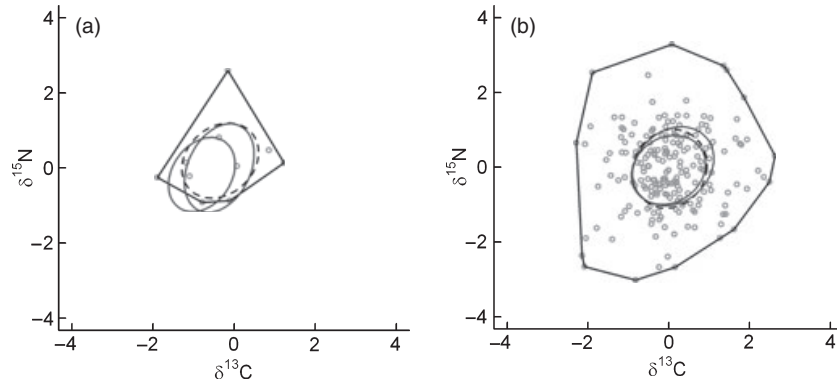


Fig. 2. Samples drawn from the same population (open circles) and their respective convex hulls (solid black lines), frequentist standard ellipses (dotted black lines) and two posterior estimates of the Bayesian standard ellipses (solid grey lines) for (a) $n = 10$ and (b) $n = 200$. The true population standard ellipse for both examples is a circle with radius = 1.

ellipse can be rescaled to include more or less of the data as desired; for instance, multiplying the semi-major and minor axes by 2.45 (the square-root of the inverse of the chi-squared distribution with 2 d.f.) will create an ellipse that has 95% probability of containing a subsequently sampled datum (Chew 1966). Rescaling clearly affects the absolute size of the resultant ellipses, but not the relative size among similarly scaled ellipses. On the other hand, Bayesian estimation of the standard ellipse is subject to fluctuations in the estimate of the location of the ellipse (arising from uncertainty in the means) and in the shape (arising from uncertainty in the covariance matrix) and thus has the benefit of producing more information than a simple point estimate. Uncertainty in the Bayesian standard ellipse is expected to be larger for smaller sample sizes and should asymptote to the population value as $n \rightarrow \infty$ (Fig. 2, grey solid lines). Indeed, the Bayesian ellipses in Fig. 2b for $n = 200$ are almost coincidental, whereas those in 2a for $n = 10$ are considerably different.

The key results of the larger simulation are summarized in Fig. 3 for the area of the convex hull and the area of the standard ellipse (SEA). Note that the results for all six of Layman *et al.*'s (2007a) metrics are included in Fig. S1. It is clear that even for this simple case where the data describe a circle, the layman metric of convex hull area (TA) is highly sensitive to sample size and increases with sample size. In reality, biological or chemical processes will constrain this to a maximum, however, the sensitivity across ecologically relevant sam-

ple sizes is what is important to consider here. In contrast, the SEA asymptotes quickly to a constant value (in this case π , as both a and $b = 1$), at around $n = 30$. The underestimation of the population SEA for small sample sizes is an undesirable property, as it is likely to confound comparisons between studies with unequal sample sizes. In the calculations of the parameters a and b described earlier, which determine the size, shape and location of the ellipse, the standard Bessel's correction of $(n-1)$ was applied to all estimations of variance and covariance. Given that this correction was derived for 1-dimensional data and corresponds to the loss of a single degree of freedom, we argue that $(n-2)$ would be a more suitable correction for 2-dimensional data considered here. We therefore propose the calculation of a corrected standard ellipse based on this formulation, whose associated statistics we denote with the subscript c . The application of this correction leads to a slightly larger ellipse for small sample sizes, but one that has the same geometrical shape (i.e. $a/b = a_c/b_c$). The relationship between SEA and SEA_c can be understood by $SEA_c = -SEA(n-1)(n-2)^{-1}$. This correction has the desirable property of approaching 1 as $n \rightarrow \infty$, is equal to 2 when n takes the minimum suitable value 3 and provides a highly satisfactory correction for all sample sizes (dashed line, Fig. 3b). However, to be convinced that this correction is applicable in situations other than that of a specific covariance structure Σ , the simulation process was extended to a more general situation.

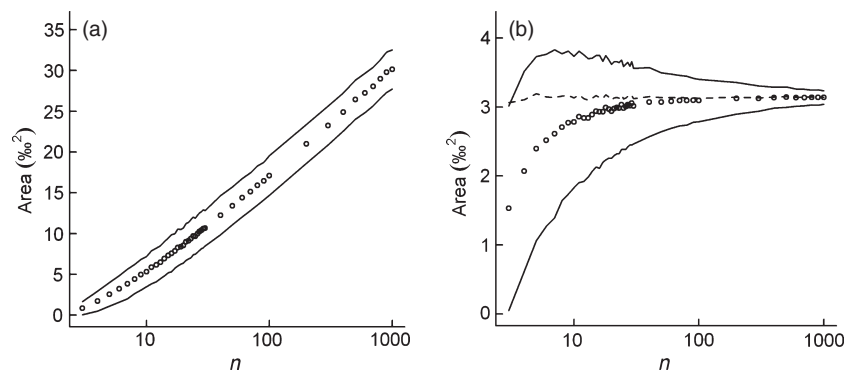


Fig. 3. The effect of sample size (n , on log scale) on (a) the area of the convex hull (TA) and (b) the standard ellipse area (SEA) for a fixed population with $\Sigma = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$, which has area equal to π . The results depict the mean (open circles) \pm the SD (solid lines) of 1000 replicates for each sample size. The dotted line in (b) is the corresponding mean SEA_c for the data.

Comprehensive simulation of SEA, SEA_c and SEA_B

The behaviour of SEA and its relationship to sample size was explored more fully in a simulation that varied both sample size and the population covariance matrix Σ governing the size and shape of the data.

A sampling process on a population of related (e.g. taxonomically or functionally) individual organisms was simulated. The observations were assumed to come from a bivariate multivariate normal distribution describing their $\delta^{13}\text{C}$ and $\delta^{15}\text{N}$ values. The simulations were centred on the origin for simplicity, with no loss of generality as it is the area and shape of the data rather than its location that is of interest here.

A single sample is generated as follows:

1. A sample size is drawn randomly from the set:

$$n \sim N[3, 4, \dots, 29, 30, 40, \dots, 100, 200, \dots, 1000]$$

This ensures that small sample sizes are well represented in the simulation study, and only sample sizes > 3 are included as a convex hull or ellipse cannot be generated on fewer values.

2. A population is defined by creating a covariance matrix (Σ) describing the relationship and structure of the $\delta^{13}\text{C}$ and $\delta^{15}\text{N}$ values. This is achieved by drawing from a Wishart distribution.

$$\Sigma \sim \text{Wishart}\left(\rho = 2, V = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}\right) \quad \text{eqn 4}$$

Where ρ is the degrees of freedom of the Wishart distribution and V is the scale matrix. The resultant distribution of the values for Σ is included in the Fig. S2.

3. $n \times 2$ bivariate samples are then drawn from a multivariate normal distribution centred on the origin with:

$$Y_{i=1:n} \sim \text{MVN}([0 \ 0], \Sigma) \quad \text{eqn 5}$$

4. Based on these samples Y , the estimated SEA $\hat{S}EA$ and the sample size corrected version $\hat{S}EA_c$ were calculated. Additionally, the SEA was estimated using the Bayesian methods outlined earlier, to yield $\hat{S}EA_B$. The bias, or difference between the estimated values of $\hat{S}EA$ and $\hat{S}EA_c$, and the population value SEA (determined directly from the known population covariance matrix Σ) were also calculated.

Steps 1–4 are then repeated for 10^5 replicates (10^4 times for Bayesian estimate as it is much more computationally intensive) and the results summarized in the following section.

The estimated SEA provides a good estimate of the population SEA although the persistent tendency for underestimation at small sample sizes remains (Fig. 4a). Calculating SEA_c instead appears to nullify this bias entirely across the full range of sample sizes and data structures as determined by their randomly varied covariance matrix Σ (Fig. 4b). The Bayesian estimate, SEA_B (Fig. 4c), captures all the same properties as SEA_c, being unbiased with respect to sample size and exhibiting more uncertainty with smaller sample size.

CASE 2 – COMMUNITY METRICS

In contrast to Case 1, we now are concerned with calculating metrics to describe the arrangement of taxonomic/functional groups as part of a larger community. The current method employed here is to take the mean of the samples for each group and calculate the six metrics based on these means (Layman *et al.* 2007a). This process can ignore important ecological information by neglecting to incorporate any uncertainty in the derived means and therefore this error is not included in the presentation of the metrics such as convex hull area (TA). Here, we develop Bayesian methods for propagating uncertainty in the mean arising from a natural sampling process into the resultant metrics.

Data for this situation were simulated by drawing each group that makes up a member of the community from its own independent multivariate normal distribution. We created $j = 1, 2, \dots, m$ groups which each comprise $i = 1, 2, \dots, n$ samples. The population mean $\delta^{13}\text{C}$ and $\delta^{15}\text{N}$ values for each of the m groups were drawn from uniform distributions between -20 and 0 , and -5 and 10 , respectively. The i th observation comprising two dependent random numbers, which belongs to the j th group is then simulated by:

$$\Sigma_j \sim \text{wishart}\left(\rho = 2, V = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}\right) \quad \text{eqn 6}$$

$$Y_{ij} \sim \text{MVN}([\delta^{13}\bar{C}, \delta^{15}\bar{N}], \Sigma_j) \quad \text{eqn 7}$$

An example simulated dataset is shown in Fig. 5, created with $m = 8$ groups and $n = 5$ samples per group. Here, the sampled data are presented in Fig. 5a, and the estimated convex hull is based on a mean of these samples as the solid black line. For comparison, the ‘true’ population convex hull is shown as a dashed line.




Fig. 4. The bias in the estimation of the population standard ellipse area (SEA) as a function of sample size (n , on log scale) based on (a) $\hat{S}EA$, (b) after small sample size correction $\hat{S}EA_c$ and (c) the Bayesian estimation $\hat{S}EA_B$. Points represent the mean of the posterior estimate. Note that the y -axes have been restricted for clarity leaving some extreme values lie outside their depicted bounds. Grey points are the results of simulations, and the heavy solid line represents their mean value for a given n . The thin black line shows perfect estimate at $y = 0$ (only visible in a for $n < 30$).

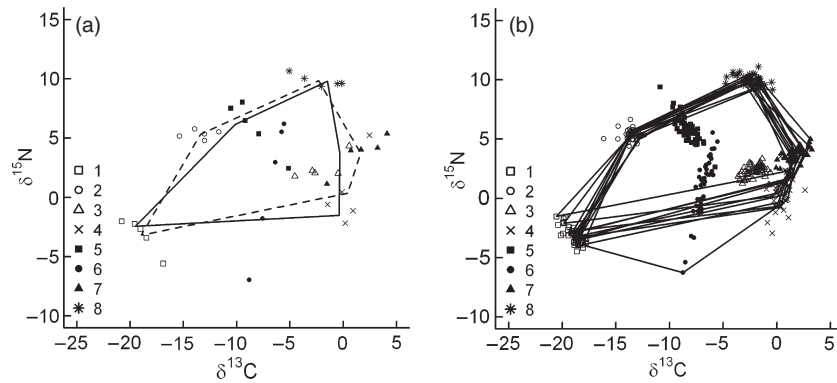


Fig. 5. (a) Data sampled according to the methodologies in the text, with $m = 8$ groups comprising $n = 5$ samples each. The convex hull based on the true population means is shown as the dashed line, while that based on the sample means is shown as the solid line. (b) Depicts 50 posterior draws for the sample means fitted to the sample data in (a), and convex hulls based on the first 10 of these posterior draws.

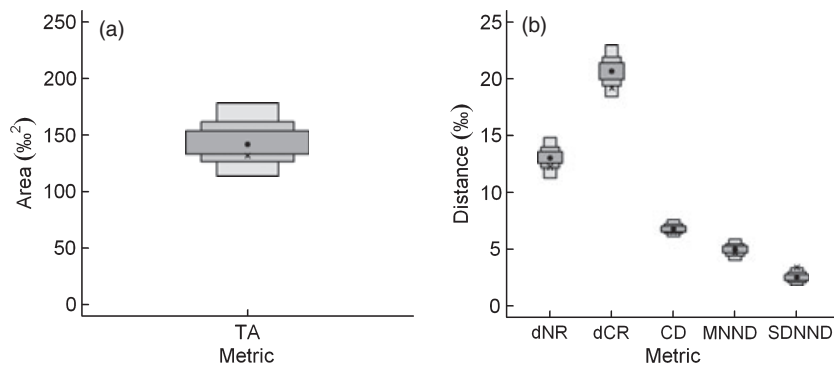


Fig. 6. Resultant uncertainty in the six Layman metrics arising from the simulated data depicted in Fig. 5a. As the scales differ greatly, the area of the convex hull (TA) is shown in (a), whereas the other five metrics (NR = dNR range; CR = dCR range) are shown in (b). Black dots represent their mode, true population values as crosses and the shaded boxes representing the 50%, 75% and 95% credible intervals from dark to light grey.

Intuitively, the sample-derived and population hulls are not identical but appear reassuringly close. The question now is: how does uncertainty in the estimation of the population means affect the convex hulls? To address this question, we use Bayesian inference to fit a multivariate normal distribution to each of the m groups independently (i.e. with separate means and covariance matrices) using the R package *bayesm* (Rossi, Allenby & McCulloch 2005), just as we did for Case 1 in the calculation of SEA_B except that here we monitor the posterior means rather than the covariance matrices Σ to locate the values that define the convex hull extremities. We use an uninformative prior for the means describing the likely range of the $\delta^{13}C$ and $\delta^{15}N$ values taking a normal distribution with means = 0 and variances = 10^6 . Additionally, a Wishart prior is used for the covariance matrix with $\rho = 2$ and $V = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$. We then draw 10^4 posterior samples using the R function *rmultreg* (Rossi, Allenby & McCulloch 2005) consistent with our data and our specified priors. Together, posterior draws represent an estimate of the probability range of the population means of each taxonomic/functional group based on the samples and the priors. For illustrative purposes, the first 100 posterior draws are shown (Fig. 5b) for a model fitted to the sample data depicted in Fig. 5a. Also shown are the convex hulls calculated on the first 10 posterior draws. All six of the existing metrics (Layman *et al.* 2007a) are similarly calculated according to these posterior estimates of the sample mean (all 10^4 of them), and their uncer-

tainty is naturally propagated through to the metrics (Fig. 6). With the corresponding probability distributions for the metrics, it is now possible to make statistical comparisons between communities.

Sensitivity to sample size

Simulation tests were run on 10^4 randomly created communities. Communities comprised randomly between 3 and 10 members (m), each of which had an independent location (randomly between -20 and 0 on the $\delta^{13}C$ axis, and -5 and 10 on the $\delta^{15}N$ axis) and shape in isotope space defined by a random covariance matrix. Additionally, each community consisted of randomly between 2 and 50 isotope samples (n) per member. The true population Layman metrics were calculated on the known population means of each of the m community members. The estimated values were calculated based on the Bayesian methodology described earlier. The results of these simulations suggest that the model's estimate of the 95% credible interval for a given metric contains the true population value on 94–95% of occasions. Furthermore, there appears to be some sensitivity to sample size (Fig. 7) for $n < 10$ such that small sample sizes typically result in underestimation of the population TA. In contrast, variation in the number of group members (m) in the community does not generally result in systematic bias in the estimates of the metrics. Similar trends are observed in all calculated metrics (see Fig. S3).




Fig. 7. Bias in the estimate of convex hull total area compared to the true population value as a function of samples (n) per group. The solid black line represents the mean of the individual simulations across a given sample size (grey circles).

Discussion

The use of stable isotopes in ecological studies has increased in recent years, and it is clear that some of the more complex applications (Bearhop *et al.* 2004; Layman *et al.* 2007a; Newsome *et al.* 2007) have the potential to revolutionize the way in which we investigate questions relating to community structure and niche occupancy. However, it is now clear that there are several caveats associated with applying the point-estimate versions of community and niche metrics that overlook the inherent variability present in ecological systems. Here, we have developed the original metrics proposed by Layman *et al.* (2007a) using standard ellipses and Bayesian methodology and demonstrate that these new approaches overcome some of these caveats providing more robust and honest descriptors of community structure.

In particular, the application of convex hull area to describe isotopic niche width of community members should be employed with due caution given to sample sizes, as eluded by Podani (2009). Comparisons between community members comprising different sample sizes will be prone to bias in the direction of larger convex hull area (TA) for the larger sample sizes. This means that meta-analyses among studies or systems are likely to be rendered statistically invalid. Indeed, as they stand, the metrics of δC_{range} , δN_{range} and TA described by Layman *et al.* (2007a) will all increase towards some biologically and ecologically constrained maximum with sample size, that of CD will increase asymptotically and MNND and SDNND will both decrease with sample size (Fig. S1). This trend is to be expected given that estimating maximum and minimum values from a sample is impossible, as the best estimate you have is the largest one you have collected to date, and can only be improved by obtaining a more extreme value in future samples. CD is a measure of deviance from the centre of gravity of the data

cloud which will stabilize to the population value with increasing sample size (Fig. S1). The metrics of MNND and the associated standard deviation of nearest neighbour distance will fall to zero at an infinite density of sample data points (Fig. S1). These general trends are not restricted to the multivariate normal distribution we consider here, but will be general features of any continuous random variates. Furthermore, the uncertainties in these metrics are inherently larger for small sample sizes, and hence caution is to be urged when looking for trends in data that is prone to unquantified error.

We propose the application of a measure of isotopic niche width based on the SEA which, after application of a small sample size correction, is shown to be insensitive to bias associated with sample size. As far as we are aware, the use of $(n-2)$ on the denominator for bivariate data instead of the standard $(n-1)$ Bessel's correction has not been suggested before (Rocchi *et al.* 2005) in the context of estimating the covariance matrix, and our simulations suggest that it produces an unbiased estimate of the standard ellipse and its associated measures such as SEA.

One approach to quantifying the isotopic d -space (Newsome *et al.* 2007) inhabited by data would be to estimate the SEA using Bayesian inference as we describe (we provide functions in R for this purpose). Qualitatively similar to bootstrapping techniques, the Bayesian approach returns a distribution (the posterior) representing estimates of SEA that reflect uncertainty arising from the sampling process, with larger uncertainty associated with smaller sample sizes. In this manner, the resultant ellipses can be compared in a quantitative manner and differences among samples stated with reference to a robust probability associated with the difference. Additionally, other factors such as differences in the distribution of the resource base in the lower trophic levels with respect to a consumer are also likely to influence dispersion in d -space and hence potentially complicate comparisons among communities (Hoeninghaus & Zeug 2008; Layman & Post 2008), and more work is required to refine methods for normalising across such communities such as implemented by Araújo *et al.* (2007).

The issue of comparability also extends to the situation where these metrics are applied to entire communities by taking the means of the community members. Uncertainty in the means with small sample size creates similar issues of bias and uncertainty that are currently ignored. Our Bayesian implementation of these metrics results in improved estimates of the community metrics, including their uncertainty (Fig. 6). We recommend a minimum of 10 samples per group member when describing the community in isotope-space to mitigate against bias towards underestimating the population value (Fig. 7).

In conclusion, the practice of using convex hull area to describe the niche width of a single community member is inappropriate, particularly should comparisons be drawn between groups of different sample sizes, thereby hampering meta-analyses and comparative studies. By assuming that the data represent a sample from some underlying distribution (currently assumed to be multivariate normal), we provide a

Bayesian method for passing the uncertainty associated with sampling on to generate robust measures of isotopic niche width of both community members and entire communities. We believe that these developments will allow ecologists to ask increasingly complex and searching questions about key factors driving community structure.

Acknowledgements

We are grateful to the associate editor, one anonymous reviewer and Carlos Martínez del Río for detailed, thoughtful and constructive comments which have improved this paper.

References

- Araújo, M., Bolnick, D., Machado, G., Giaretta, A. & dos Reis, S. (2007) Using $\delta^{13}\text{C}$ stable isotopes to quantify individual-level diet variation. *Oecologia*, **152**, 643–654.
- Batschelet, E. (1981) *Circular Statistics in Biology*. Academic Press, London.
- Bearhop, S., Adams, C.E., Waldron, S., Fuller, R.A. & Macleod, H. (2004) Determining trophic niche width: a novel approach using stable isotope analysis. *Journal of Animal Ecology*, **73**, 1007–1012.
- Bolnick, D.I., Svanbäck, R., Fordyce, J.A., Yang, L.H., Davis, J.M., Hulsey, C.D. & Forister, M.L. (2003) The ecology of individuals: incidence and implications of individual specialization. *American Naturalist*, **161**, 1–28.
- Chew, V. (1966) Confidence, prediction and tolerance regions for multivariate normal distribution. *Journal of the American Statistical Association*, **61**, 605–617.
- Chisholm, B.S., Nelson, D.E. & Schwarcz, H.P. (1982) Stable-carbon isotope ratios as a measure of marine versus terrestrial protein in ancient diets. *Science*, **216**, 1131–1132.
- Craig, H. (1953) The geochemistry of stable-carbon isotopes. *Geochimica Et Cosmochimica Acta*, **3**, 53–92.
- Darimont, C.T., Paquet, P.C. & Reimchen, T.E. (2009) Landscape heterogeneity and marine subsidy generate extensive intrapopulation niche diversity in a large terrestrial vertebrate. *Journal of Animal Ecology*, **78**, 126–133.
- Fry, B. (1988) Food web structure on Georges Bank from stable C, N, and S isotopic compositions. *Limnology and Oceanography*, **33**, 1182–1190.
- Hoeinghaus, D.J. & Zeug, S.C. (2008) Can stable isotope ratios provide for community-wide measures of trophic structure? Comment. *Ecology*, **89**, 2353–2357.
- Hoeinghaus, D.J., Winemiller, K.O., Layman, C.A., Arrington, D.A. & Jepsen, D.B. (2006) Effects of seasonality and migratory prey on body condition of *Cichla* species in a tropical floodplain river. *Ecology of Freshwater Fish*, **15**, 398–407.
- Hutchinson, G.E. (1957) Concluding remarks. In Cold Spring Harbour symposium (ed Q. Biology).
- Hutchinson, G.E. (1978) *An Introduction to Population Biology*. Yale University Press, New Haven.
- Inger, R., Ruxton, G.D., Newton, J., Colhoun, K., Robinson, J.A., Jackson, A.L. & Bearhop, S. (2006) Temporal and intrapopulation variation in prey choice of wintering geese determined by stable isotope analysis. *Journal of Animal Ecology*, **75**, 1190–1200.
- Inger, R., Gudmundsson, G.A., Ruxton, G.D., Newton, J., Colhoun, K., Auhage, S. & Bearhop, S. (2008) Habitat utilisation during staging affects body condition in a long distance migrant, *Branta bernicla hrota*: potential impacts on fitness? *Journal of Avian Biology*, **39**, 704–708.
- Jackson, A.L., Inger, R., Bearhop, S. & Parnell, A. (2009) Erroneous behaviour of MixSIR, a recently published Bayesian isotope mixing model: a discussion of Moore & Semmens (2008). *Ecology Letters*, **12**, E1–E5.
- Layman, C.A. & Post, D.M. (2008) Can stable isotope ratios provide for community-wide measures of trophic structure? Reply. *Ecology*, **89**, 2358–2359.
- Layman, C.A., Arrington, D.A., Montana, C.G. & Post, D.M. (2007a) Can stable isotope ratios provide for community-wide measures of trophic structure? *Ecology*, **88**, 42–48.
- Layman, C.A., Quattrochi, J.P., Peyer, C.M. & Allgeier, J.E. (2007b) Niche width collapse in a resilient top predator following ecosystem fragmentation. *Ecology Letters*, **10**, 937–944.
- Martínez del Río, C., Sabat, P., Anderson-Sprecher, R. & Gonzalez, S.P. (2009) Dietary and isotopic specialization: the isotopic niche of three *Cinclodes* ovenbirds. *Oecologia*, **161**, 149–159.
- McCarthy, M.A. (2007) *Bayesian Methods for Ecology*. Cambridge University Press, Cambridge.
- Minagawa, M. & Wada, E. (1984) Stepwise enrichment of ^{15}N along food-chains – further evidence and the relation between $\delta^{15}\text{N}$ and animal age. *Geochimica Et Cosmochimica Acta*, **48**, 1135–1140.
- Newsome, S.D., Martínez del Río, C., Bearhop, S. & Phillips, D.L. (2007) A niche for isotopic ecology. *Frontiers in Ecology and the Environment*, **5**, 429–436.
- Olsson, K., Stenroth, P., Nystrom, P. & Graneli, W. (2009) Invasions and niche width: does niche width of an introduced crayfish differ from a native crayfish? *Freshwater Biology*, **54**, 1731–1740.
- Parnell, A., Inger, R., Bearhop, S. & Jackson, A.L. (2008) SIAR: Stable isotope analysis in R.
- Parnell, A.C., Inger, R., Bearhop, S. & Jackson, A.L. (2010) Source partitioning using stable isotopes: coping with too much variation. *PLoS ONE*, **5**, e9672.
- Podani, J. (2009) Convex hulls, habitat filtering, and functional diversity: mathematical elegance versus ecological interpretability. *Community Ecology*, **10**, 244–250.
- Quevedo, M., Svanbäck, R. & Eklöv, P. (2009) Intrapopulation niche partitioning in a generalist predator limits food web connectivity. *Ecology*, **90**, 2263–2274.
- R Development Core Team (2007) *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Ricklefs, R.E. & Nealen, P. (1998) Lineage-dependent rates of evolutionary diversification: analysis of bivariate ellipses. *Functional Ecology*, **12**, 871–885.
- Rocchi, M.B.L., Sisti, D., Ditroilo, M., Calavalle, A. & Panebianco, R. (2005) The misuse of the confidence ellipse in evaluating statokinesigram. *Italian Journal of Sports Sciences*, **12**, 169–171.
- Rossi, P.E., Allenby, G.M. & McCulloch, R. (2005) *Bayesian Statistics and Marketing*. Wiley, Chichester.
- Royston, J.P. (1982) The W test for normality. *Applied Statistics-Journal of the Royal Statistical Society Series C*, **31**, 176–180.
- Semmens, B.X., Moore, J.W. & Ward, E.J. (2009a) Improving Bayesian isotope mixing models: a response to Jackson *et al.* (2009). *Ecology Letters*, **12**, E6–E8.
- Semmens, B.X., Ward, E.J., Moore, J.W. & Darimont, C.T. (2009b) Quantifying inter- and intra-population niche variability using hierarchical Bayesian stable isotope mixing models. *PLoS ONE*, **4**, e6187.
- Soberon, J. (2007) Grinnellian and Eltonian niches and geographic distributions of species. *Ecology Letters*, **10**, 1115–1123.
- Székely, G.J. & Rizzo, M.L. (2005) A new test for multivariate normality. *Journal of Multivariate Analysis*, **93**, 58–80.
- Tan, M., Fang, H.-B., Tian, G.-L. & Wei, G. (2005) Testing multivariate normality in incomplete data of small sample size. *Journal of Multivariate Analysis*, **93**, 164–179.
- Wada, E., Terazaki, M., Kabaya, Y. & Nemoto, T. (1987) ^{15}N and ^{13}C abundances in the Antarctic Ocean with emphasis on the biogeochemical structure of the food web. *Deep-Sea Research Part a-Oceanographic Research Papers*, **34**, 829–841.

Received 14 April 2010; accepted 6 January 2011

Handling Editor: Ryan Norris

Supporting Information

Additional Supporting Information may be found in the online version of this article.

Fig. S1. The effect of sample size on the 6 Layman metrics (Layman *et al.* 2007a) when applied to a single population of bivariate isotope data (case 1).

Fig. S2. Values drawn from the Wishart distribution with degrees of freedom $\rho = 2$ and scale matrix $V = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}$.

Fig. S3. Bias in the Bayesian estimate of each of the six layman metrics when applied to an ecological community (Case 2 in the main text) as a function of sample size (n).

As a service to our authors and readers, this journal provides supporting information supplied by the authors. Such materials may be re-organized for online delivery, but are not copy-edited or typeset. Technical support issues arising from supporting information (other than missing files) should be addressed to the authors.