ECOGRAPHY 29: 786–792, 2006

Null models for species co-occurrence patterns: assessing bias and minimum iteration number for the sequential swap

Veiko Lehsten and Peter Harmand

Lehsten, V. and Harmand P. 2006. Null models for species co-occurrence patterns: assessing bias and minimum iteration number for the sequential swap. – Ecography 29: 786–792.

The analysis of co-occurrence matrices is a common practice to evaluate community structure. The observed data are compared with a "null model", a randomised co-occurrence matrix derived from the observation by using a statistic, e.g. the C-score, sensitive to the pattern investigated. The most frequently used algorithm, "sequential swap", has been criticised for not sampling with equal frequencies thereby calling into question the results of earlier analysis. The bias of the "sequential swap" algorithm when used with the C-score was assessed by analysing 291 published presence-absence matrices. In 152 cases, the true p-value differed by $>5\%$ from the p-value generated by an uncorrected "sequential swap". However, the absolute value of the difference was rather small. Out of the 291 matrices, there were only 5 cases in which an incorrect statistical decision would have been reached by using the uncorrected p-value (3 at the $p<0.05$ and 2 at the $p<0.01$ level), and in all 5 of these cases, the true p-value was close to the significance level. Our results confirm analytical studies of Miklos and Podani which show that the uncorrected swap gives slightly conservative results in tests for competitive segregation. However, the bias is very small and should not distort the ecological interpretation. We also estimated the number of iterations needed for the "sequential swap" to generate accurate p-values. While most authors do not exceed a number of 10^4 iterations, the suggested minimum number of swaps for 29 out of the 291 tested matrices is greater than 10^4. We recommend to use 30 000 "sequential swaps" if the required sample size is not assessed otherwise.

V. Lehsten (veiko.lehsten@uni-oldenburg.de), Landscape Ecology Group, Univ. of Oldenburg (Fak V), DE-26111 Oldenburg, Germany. – P. Harmand, Inst. of Mathematics, Univ. of Oldenburg, DE-26111 Oldenburg, Germany.

Analysing co-occurrence data has become a common practice in ecology to study the community structure within single observations (Gotelli et al. 1987) as well as to verify general ecological theories by using meta-analysis of co-occurrence matrices (Gotelli and McCabe 2002). All these analyses require a randomisation of the observed data, i.e. (0, 1)-matrices, to which the observed pattern is compared. Although a number of different null models is used to test different ecological hypotheses (Gotelli (2000) compares nine different null models), most authors use a variant of the null model proposed by Connor and Simberloff (1979). It retains row and column sums simultaneously to incorporate site effects such as island size as well as rarity of species to account for species dependent characteristics such as niche breadth. Connor and Simberloff (1979) also used a third constraint by restricting species occurrences to those islands for which the total species richness fells within the range occupied by the species. The basic assumption of the null model methodology is, that if the observed co-occurrence matrix differs by much with respect to a certain pattern from the total set of unique matrices,

Accepted 31 July 2006

ISSN 0906-7590

then there is a structure which can be ecologically interpreted. The investigated pattern is often summarised within a single score (metric or test statistic) which is extreme for structured matrices, e.g. if the matrix is structured by an ecological mechanism. If this score is not significantly different between the observed matrix and the randomised matrices, no pattern can be detected. For instance, to evaluate the co-occurrence between species, the number of perfect checkerboard pairs or the C-score (Stone and Roberts 1990) is used by several authors (Wilson 1987, Feeley 2003). Since it is only possible to calculate the total set for relatively small matrices (as will be shown below), a randomisation algorithm is applied to sample a subset of matrices, which will then be compared to the observed matrix.

A valid randomisation algorithm has to sample all matrices with fixed row and column sums at equal frequencies. The choice of the randomisation algorithm has been shown to influence the result of the study. In a re-analysis of a presence/absence matrix from the Vanuatu avian fauna, Sanderson et al. (1998) concluded that the "results from previous studies are flawed" due to an inappropriate randomisation algorithm. Gotelli and Entsminger (2001) showed by using probability calculations that the randomisation algorithm used by Sanderson et al. (1998), the "Knight's Tour", is biased towards not sampling all matrices with equal frequencies, which in turn has led to contradictory results. On the other hand, the "sequential swap" algorithm is also prone to sample matrices with unequal frequencies depending on the observed matrix (Gotelli and Entsminger 2001). This controversy about null models has lead to publications reporting results using several randomisation algorithms (Feeley 2003). Although Miklos and Podani (2004) proposed a new unbiased randomisation method, we suggest using the original "sequential swap" and performing a frequency correction afterwards as described by Zaman and Simberloff (2002). The first method converges to the unbiased uniform distribution, the latter uses a frequency correction i.e. "importance weighting" to eliminate the bias of the sequential swap algorithm.

Another issue when applying a randomisation procedure is the necessary size of the random sample of null matrices needed for an analysis. Since the matrices sampled by the "sequential swap" are not independent of each other, this question is not straightforward. Raftery and Lewis (1996) developed a procedure which calculates the minimum number of required iterations to estimate stable state probabilities of a Markov chain which we apply to the randomisation of matrices.

The bias of the sequential swap and the necessary number of swaps calculated by the procedure of Raftery and Lewis (1996) is assessed using a large collection of published presence/absence matrices.

Material and methods

Data

The applicability of null models and especially of the "sequential swap" has been discussed using the data set of the Vanuatu avifauna (Diamond and Marshall 1979, Wilson 1987, Stone and Roberts 1990). The p-value of the C-score for this data set is calculated. To show the relevance of the approach we use 291 published matrices, collected by Patterson and Atmar (1986) and calculate the p-values of the C-scores as well as their differences obtained with and without a frequency correction of the "sequential swap".

Scores

We use the checkerboard score (C-score) to illustrate the analysis (Roberts and Stone 1990). It measures the mean number of pairs of species and islands with one species occurring on one island only and the second occurring on the second island only. The number of checkerboards involving species i and j can be calculated as follows:

$$C_{ij} = (r_i - S_{ij})(r_j - S_{ij}). \quad (1)$$

Where r_i is the sum of the ith row and S_{ij} is the number of islands that the two species share. Let $P = m(m-1)/2$ be the number of species pairs for m species, then the C-score is:

$$C = \sum_{i<j} C_{ij}/P. \quad (2)$$

Randomisation algorithms

There are two general techniques for generating random matrices with given row and column sums: filling and swapping (for a survey see Gotelli and Entsminger 2001). We will solely consider swap algorithms. The original "sequential swap" (Manly 1995) randomly selects a pair of rows and a pair of columns. If one species occurs only at the first site and the other species occurs only at the other site, their occurrences are interchanged, i.e. after the swap the first species is assigned to the second site and the other species is assigned to the first site. In this way, both row and column sums are kept constant. If swapping was not possible, a new pair of rows and columns is selected – this is not counted as a step. The algorithm starts with the observed presence/absence matrix.

The set of all matrices with given row and column sums

The best way to test the C-score of the observed occurrence matrix would be to calculate the C-score of

all elements in the set A of all possible matrices with the same row and column sums as the observed matrix. In almost all cases this is computational impossible nowadays.

Except for trivial matrices, the generation of the complete null set is computationally infeasible. Even the much simpler problem of computing |A|, the number of elements in A, i.e. the size of the null space, is often impossible to solve. A formula for |A| was found by Wang and Zhang (1998) and simplified by Perez-Salvador et al. (2002). However, this formula can only be evaluated in reasonable time if at least one dimension of the matrix is small, e.g. <11. Nevertheless, for about one third of the matrices in the Patterson and Atmar (1986) collection it was possible to find |A|. For some very small matrices this permitted enumeration of the whole set A and to compare the estimated p-value with the true global p-value. It also revealed that in some cases the size of A is really enormous: e.g. the 4×180-matrix of the data collection gives $|A|=4.7\times10^{68}$. In sharp contrast to the large number of matrices in the set A is the fact that the minimal number of swaps required to transform matrix M_1 into matrix M_2 is $<d(M_1, M_2)/2-1$, where $d(M_1, M_2)$ is the number of positions where M_1 and M_2 differ (Brualdi 1980 p. 172). By a rough estimate using only the row or column sums one obtains $d\leq282$ for the 4×180 matrix mentioned above, hence with <140 swaps one can go from any matrix to any other matrix within this set of 4.7×10^{68} matrices. Besides the formulas for the exact value of |A|, there are Markov chain Monte Carlo techniques which give approximations (chapter 4.3; Liu 2001).

The frequency correction of the swap algorithm

The generation of random matrices by the "sequential swap" can be seen as a Markov process in which each matrix is one state. There are as many ways to reach different states (possibilities to swap) as there are checkerboards within a given matrix. As an example consider the matrix published by Maly and Doolittle (1977). There are five unique matrices with the same row and column sums (M_0–M_4) representing five states of the Markov process (Box 1). The probabilities of going from one state to another are drawn in Box 1. Box 1 lists the transition probabilities, the C-score and the stable state probabilities. If a large number of swaps is performed, matrix M_0 will be sampled in 25% of the cases, while each other matrix will be sampled only in 18.75% of all cases. If this matrix would be analysed using the "sequential swap", the resulting expected C-score would be 0.2167 instead of the correct value of 0.2133. In the general case one can show that the stationary distribution and hence, in the long run, the frequencies of the possible matrices in a simulation are proportional to the C-scores. This is very plausible if the process is viewed as a random walk on a graph with the matrices as vertices and with edges joining them if there is a swap which transforms one into the other: vertices with a high number of edges, i.e. matrices with high C-score, are visited more often than others (Zaman and Simberloff 2002).

If S is any statistic and S_i denotes its value at the ith step of a simulation of length n, the usual sample mean $\overline{S}=\frac{1}{n}\Sigma_{i=1}^{n}S_i$ is biased and has to be replaced by a weighted sum $\overline{S_{corr}}=\frac{1}{n}\Sigma_{i=1}^{n}w_iS_i$. To correct the unequal frequencies observed above, the weights w_i should have the form a/C_i, where C_i is the C-score of the matrix in the ith step. The constant a is a proportional factor, determined by the condition $\Sigma_{i=1}^{n}w_i=n$ which gives $a=\overline{C_{corr}}$, with

$$\overline{C_{corr}}=\frac{n}{\sum_{i=1}^{n}\frac{1}{C_i}} \tag{3}$$

and finally

$$\overline{S_{corr}}=\frac{\overline{C_{corr}}\sum_{i=1}^{n}\frac{S_i}{C_i}}{n}. \tag{4}$$

The frequency corrected p-value can be obtained from this by choosing the statistic $S(M)=1$, if the C-score of M is greater than or equal to the C-score of the initial matrix and $S(M)=0$ otherwise. Hence, to obtain the probability of reaching a certain score, the frequency of each matrix has to be weighted by the ratio of $\overline{C_{corr}}/C_i$ (see example in Box 1). In a histogram of the C-scores generated by the original algorithm, all bars representing C-scores higher than the expected value $\overline{C_{corr}}$ would therefore become smaller and all bars of C-scores smaller than the expected value would become higher.

Minimum number of required swaps

Several authors suggest invoking a "burn in period" for the swap algorithm, i.e. the first randomised matrices are discarded (Gotelli and McCabe 2002, Zaman and Simberloff 2002), or generating the start matrix with a fill algorithm (Miklos and Podani 2004) to minimise the influence of the structure of the observed matrix. Although theoretical considerations show that, regardless of the starting matrix, the algorithm will converge to the correct stable state distribution, it may be useful not

Box 1. Example calculation of the expected frequencies and C-scores by the "sequential swap" and the frequency corrected "sequential swap".

The matrix (M_0) published by Maly and Doolittle (1977) has only 5 unique randomisations (M_0–M_4) with fixed row and column totals. Figure A shows the five matrices and the transition probabilities.

Frequency correction:

$$\overline{C_{corr}} = \frac{n}{\sum_{i=1}^{n} \frac{1}{C_i}}$$

Consider a simulation with 10 000 swaps, the expected frequencies of the C-scores generated by the "sequential swap" would be 2500 times 0.2666 and 7500 times 0.2. The mean C-score generated by the "sequential swap" is 0.2167. The correction would be calculated as:

$$\overline{C_{corr}} = \frac{10000}{2500 \times \frac{1}{0.2666} + 7500 \times \frac{1}{0.2}}$$

$$\overline{C_{corr}} = 0.2133$$

The probability of reaching a C-score as extreme as the observation by using the "sequential swap" is 0.25 while the correct probability is 0.20 (Table A).

The frequency correction would give a corrected p-value as follows:

$$\overline{P_{corr}} = \frac{\sum_{i=1}^{n} \frac{C_{corr}}{C_i}}{n} \quad \text{for } C_i \geq C_{obs}$$

$$\overline{P_{corr}} = \frac{\frac{0.2133}{0.266} \times 2500}{10000}$$

$$\overline{P_{corr}} = 0.2$$

The frequency corrected "sequential swap" gives a p-value of 0.200 which is equal to the theoretical expectation while the uncorrected version results in a p-value of 0.25.

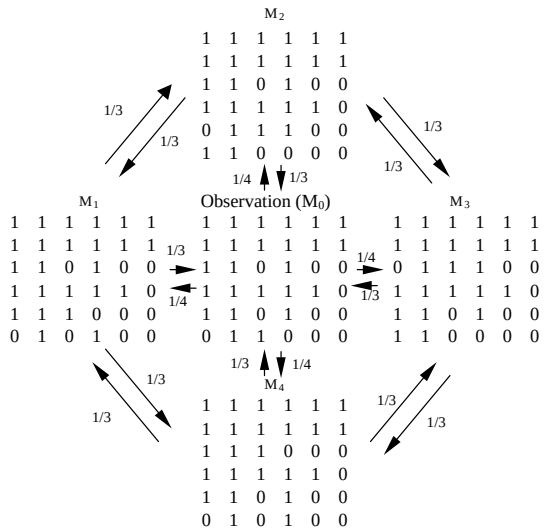

Fig. A. Unique matrices and transition probabilities of the Maly and Doolittle (1977) data set

Table A. Lists of the transition probabilities, the C-score and the stable state probabilities of each matrix using the "sequential swap". Note that the stable state probabilities are proportional to the C-score.

Matrix nr	0	1	2	3	4	C-score	Stable state prop.
0	0	1/4	1/4	1/4	1/4	0.2666	0.25
1	1/3	0	1/3	0	1/3	0.2000	0.1875
2	1/3	1/3	0	1/3	0	0.2000	0.1875
3	1/3	0	1/3	0	1/3	0.2000	0.1875
4	1/3	1/3	0	1/3	0	0.2000	0.1875
					Mean	0.2133	

The C-score derived by the "sequential swap" is: $4 \times 0.20 \times 0.1875 + 1 \times 0.2666 \times 0.25 = 0.2167$, while the correct C-score using equal frequencies for each state is: $4 \times 0.20 \times 0.20 + 1 \times 0.2666 \times 0.20 = 0.2133$.

only to discard the first set of matrices, but also to use only every kth which is referred to as "thinning" the chain. While this will remove some of the correlation between consecutive iterations, it will not yield independence.

Raftery and Lewis (1996) developed a procedure to estimate the number of iterations required to calculate the stable state distribution of a Markov chain with a given precision. Using their procedure, we calculated the required number of swaps to detect a significant deviation at the $p = 0.025$ level (as a compromise between the 0.05 and the 0.01 level) with a precision of ± 0.0125 for each matrix at a 95% confidence level. A trial run of the swap algorithm is performed and the results are used to calculate the required minimum number of iterations. If the suggested number was greater than the trial run, the size of the pilot sample was increased until it was greater than the suggested minimum number of iterations.

Results

The Vanuatu data set

The C-score of the Vanuatu data set (Diamond and Marshall 1979) is 9.5299. Performing 10^6

swaps gives a mean value $\overline{C}=9.1299$ (min = 8.8279; max = 9.7227; SD = 0.0886) and a corrected mean value $\overline{C_{corr}}=9.1290$. The p-value generated by the "sequential swap" is 9.4×10^{-5} and the frequency corrected p-value is 8.969×10^{-5}. Both are highly significant. The minimum number of required iterations is computed to 10751 ± 3603 (using 50 independent pilot samples) by the algorithm of Raftery and Lewis (1996) using the above mentioned parameters.

A rough estimate of the minimum number of swaps to be performed to go from any one of the matrices to any other yields a number <336. Hence we have no indication, that the performed 10^6 iterations did not sample the full variance of the null space.

The Patterson and Atmar data set

Using 291 matrices collected by Patterson and Atmar (1986), the p-value of the C-score for each matrix was calculated using the "sequential swap" with and without frequency correction (Fig. 1). Compared to the original algorithm, the frequency corrected "sequential swap" identifies three more matrices as statistically significant ($p<0.05$) and two more matrices as strongly significant ($p<0.01$). These matrices have uncorrected p-values of 0.0576, 0.0613, 0.0648, 0.014 and 0.0106. The corrected p-values are 0.0485, 0.049, 0.0347, 0.0086 and 0.0096 respectively. The total and relative differences of the p-values derived by the two algorithms are displayed in Fig. 2. Though the total p-values are relatively similar (maximum total difference is 0.062), they differ substantially in relation to each other. In 152 out of 291 matrices the differences are over 5% and in 85 matrices they are $>10\%$.

The "burn in period" as suggested by the procedure of Raftery and Lewis (1996) is in all cases negligible, compared to the total of required iterations with a maximum ratio of 0.0063 to the suggested total number of iterations. For some matrices it was found, that invoking a "burn in" using the algorithm of Raftery and Lewis (1996) decreases the suggested number of iterations. The procedure may hence not always deliver the optimal combination of "burn in period" and minimum of iterations, but still a valid one. Figure 3 plots the minimum number of required swaps. For 29 matrices this number is above 10^4. While the bias of the C-score is especially pronounced for small matrices, we found no relation between the required number of iterations and the size of the matrix.

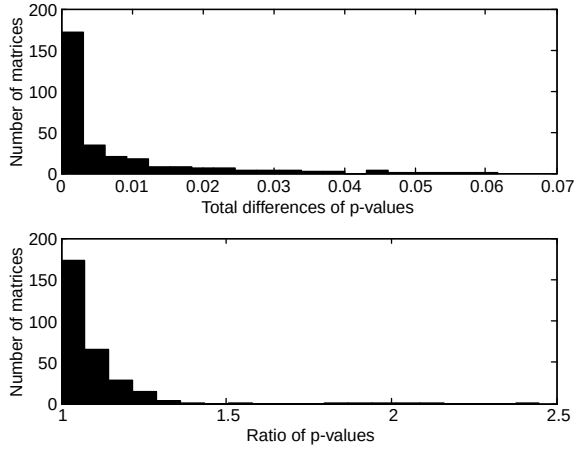

Fig. 2. Histograms of the total and relative differences of the p-values for the C-score generated by the "sequential swap" and the frequency corrected "sequential swap", using 10^5 swaps and 291 published data sets.

Discussion

The search for structure in presence/absence matrices has a long history in ecology, as in many cases these are the only available data. Since Diamond (1975) published his assembly rules, there has been an ongoing debate on the methods for detecting structure and how to interpret them (for a review see Gotelli and Graves 1996). One of the questions remaining open is the choice of the best algorithm for generating random matrices with fixed row and column totals. Although the "sequential swap" has been shown to oversample certain matrices depending on the observed data set, Gotelli and Entsminger (2001)

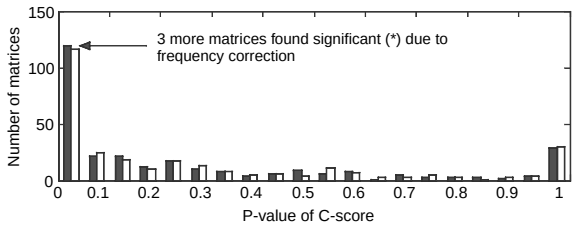

Fig. 1. Histogram of p-values of the deviation of the observed C-score from the expected value using the "sequential swap" (empty bars) and the frequency corrected "sequential swap" (solid bars). 291 published sets of data were analysed using 10^5 iterations (swaps).

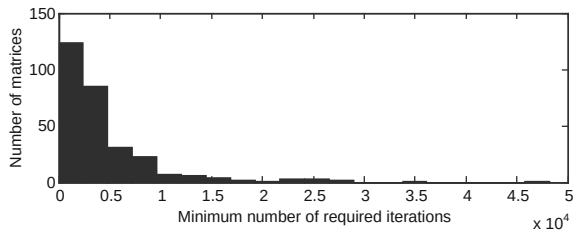

Fig. 3. Histogram of the minimum number of required iterations to estimate a p-value of 0.025 with a precision of ±0.0125 at a 95% confidence level, calculated by the algorithm of Raftery and Lewis (1996) for 291 published data sets. Note that 10 percent of all matrices require more than 10^4 iterations.

suggest using the "sequential swap" since the possible bias seems to be low and is out-weighted by the low computational demand. Zaman and Simberloff (2002) investigated the statistical properties of the "sequential swap" and suggest weighting the calculated statistic by the number of neighbouring matrices (which is related to the C-score). Miklos and Podani (2004) suggested a "trial swap" in which also the unsuccessful attempts count when calculating the statistic. The number of attempts is an unbiased estimator for the "importance weights" precisely calculated by Zaman and Simberloff (2002). Theoretical considerations suggest that estimating these values – instead of calculating the precise value – introduces more variance in the sampling distribution and hence increases the number of successful swaps necessary to calculate the p-value. Despite the expected higher variance of the Mikos and Podani method, it is mathematically proven that both methods give valid results.

Our results demonstrate the importance of applying a correction because there are published data sets for which the original "sequential swap" indicates no significant difference in the C-score while the frequency corrected "sequential swap" do find such a difference (Feeley 2003). The bias depends on the relative differences of the C-score (or number of checkerboards), which are relatively low for large matrices like the Vanuatu data set, with a C-score ranging from 9.7227 to 8.8279 (10%, see Results) and higher (25%) for small matrices like the one in Box 1. However, in only 5 out of 291 matrices the bias had an influence on the significance at the levels 0.01 and 0.05. Hence in most cases the influence of other factors (e.g. overlooking a species at a certain site, incorrect species identity or species status determination) can be expected to influence the result more than the bias of the "sequential swap" algorithm.

In most studies a histogram of the C-score is plotted and the individual C-scores are stored for the significance analysis. In this case the correction can be performed afterwards without repeating the randomisation. The potential bias of other scores is unknown, it may be smaller, but it can also be much higher without a frequency correction. We therefore suggest using the frequency correction for any metric.

So far the number of generated random matrices has been chosen by experience, educated guess, or was limited by the computing power. The estimation procedure by Raftery and Lewis (1996) suggested using $>10^4$ iterations for $>10\%$ of the matrices, where 10^4 is the number of iterations mostly used in the cited literature. Based on our investigation we suggest to use 3×10^4 iterations as a rule of thumb in case no estimation of the required value can be performed. This value can be expected to be sufficient for the vast majority of matrices and can be realised with the EcoSim software (Gotelli and Entsminger 1999).

Manly (1995, p. 1111) pointed out the difference between global and conditional p-values and addresses the problem that the swap algorithm without "burn in" incorporates the structure of the original matrix and is hence biased towards the observation. On the other hand, classical results from Markov chain theory show that p(n), the estimated p-value after n steps, converge to the correct global p-value independent of the starting matrix (Zaman and Simberloff 2002). However, quantitative versions of this convergence or error estimates remain open in general. To our knowledge the procedure of Raftery and Lewis (1996) is the best of what is available to determine the minimum number of required iterations. Calculating the minimum number of iterations for any null model algorithm which can be regarded as a Markov chain ("sequential swap") is computational easy because the authors implemented their procedure in a program available for download (see cited paper for details).

This study analyses the bias of the sequential swap using the C-score. However, we are aware that there are other null models and other metrics. The C-score indicates only a deviation from randomness as it would be expected in the presence of competition, it does not indicate which factors structure the competitive interactions. Using other statistics than the C-score, null models are also applicable to answer more specific ecological questions like "Which factors structure competitive interactions (e.g. environmental conditions, species traits)?". This more rewarding task has been approached by the "fourth corner method" of Legendre et al. (1997), a null model technique relating habitat to species traits. In an analysis of a greenhouse experiment (Lehsten 2005), we applied the "sequential swap" to highlight the relationship between plant traits and the treatment. To assure the validity of our analysis we applied the frequency correction of Zaman and Simberloff (2002).

Though the null model of retaining species diversity and number of occurrences has been frequently used before, we suggest to design a null model specific to the ecological mechanism of interest and the available data. In the optimal case, the design of the null model is done in conjunction with the design of the experiment and the sampling. This allows to make the best use of the advantages of this technique, for instance, its ability to deal with dependent data, or the lower amount of data needed. The latter is a result of the fact that no standard distribution has to be assumed (which can in itself be problematic), instead the test distribution is calculated from the recorded data.

Though we showed that in practice the bias of the "sequential swap" on the C-score may be low, as long as the matrix is relatively large, we suggest using the frequency correction by Zaman and Simberloff (2002) and calculating the number of required swaps using the

procedure of Raftery and Lewis (1996) to assure the validity of the calculated deviation regardless of the calculated metric.

Acknowledgements – We want to thank N. Gotelli for providing the Vanuatu data set and helpful comments on the manuscript. We also thank J. Sanderson for his comments which improved the paper.

References

Brualdi, R. A. 1980. Matrices of zeros and ones with fixed row and column sum vectors. – Linear Algebra Appl. 33: 159–231.

Connor, E. F. and Simberloff, D. 1979. The assembly of species communities: chance or competition? – Ecology 60: 1132–1140.

Diamond, J. M. 1975. Assembly of species communites. – In: Cody, M. L. and Diamond, J. M. (eds), Ecology and evolution of communities. Harvard Univ. Press, pp. 342–444.

Diamond, J. M. and Marshall, A. G. 1979. Origin of the Hebridean avifauna. – Emu 76: 187–200.

Feeley, K. 2003. Analysis of avian communities in Lake Guri, Venezuela, using multiple assembly rule models. – Oecologia 137: 104–113.

Gotelli, N. J. 2000. Null model analysis of species co-occurrence patterns. – Ecology 81: 2606–2621.

Gotelli, N. J. and Graves, G. R. 1996. Null models in ecology. – Smisonian Inst. Press.

Gotelli, N. J. and Entsminger, G. L. 1999 EcoSim: null models software for ecology. – Aquired Intelligence and Kesey-Bear, <http://together.net/~gentsmin/ecosim.htm>.

Gotelli, N. J. and Entsminger, G. L. 2001. Swap and fill algorithms in null model analysis: rethinking the knight's tour. – Oecologia 129: 281–291.

Gotelli, N. J. and McCabe, D. J. 2002. Species co-occurrence: a meta-analysis of J. M. Diamond's assembly rules model. – Ecology 83: 2091–2096.

Gotelli, N. J. et al. 1987. Body-size differences in a colonizing amphipod-mollusk assemblage. – Oecologia 72: 104–108.

Legendre, P. et al. 1997. Relating behavior to habitat: solutions to the fourth-corner problem. – Ecology 78: 547–562.

Lehsten, V. 2005. Functional analysis and modelling of vegetation. – Ph.D. thesis, Carl von Ossietzky Univ., Oldenburg, <http://docserver.bis.uni-oldenburg.de/publikationen/dissertation/2005/lehfun05/lehfun05.html>.

Liu, J. S. 2001. Monte Carlo strategies in scientific computing. – Springer.

Maly, E. J. and Doolittle, W. L. 1977. Effects of island area and habitat on Bahamian land and freshwater snail distribution. – Am. Midl. Nat. 97: 59–67.

Manly, B. F. J. 1995. A note on the analysis of species co-occurrences. – Ecology 76: 1109–1115.

Miklos, I. and Podani, J. 2004. Randomization of presence-absence matrices: comments and new algorithms. – Ecology 85: 86–92.

Patterson, B. D. and Atmar, W. 1986. Nested subsets and the structure of insular mammalian faunas and archipelagoes. – Biol. J. Linn. Soc. 28: 65–82.

Perez-Salvador, B. R. et al. 2002. A reduced formula for the precise number of (0,1)-matrices in A(R,S). – Discret. Math. 256: 361–372.

Raftery, A. E. and Lewis, S. M. 1996. Implementing MCMC. – In: Gilks, W. R. et al. (eds), Markov chain Monte Carlo in practice. Chapman and Hall, pp. 115–130.

Roberts, A. and Stone, L. 1990. Island-sharing by archipelago species. – Oecologia 83: 560–567.

Sanderson, J. G. et al. 1998. Null matrices and the analysis of species co-occurrences. – Oecologia 116: 275–283.

Stone, L. and Roberts, A. 1990. The checkerboard score and species distributions. – Oecologia 85: 74–79.

Wang, B. Y. and Zhang, F. Z. 1998. On the precise number of (0,1)-matrices in U(R,S). – Discret. Math. 187: 211–220.

Wilson, J. B. 1987. Methods for detecting nonrandomness in species co-occurrences – a contribution. – Oecologia 73: 579–582.

Zaman, A. and Simberloff, D. 2002. Random binary matrices in biogeographical ecology – instituting a good neighbor policy. – Environ. Ecol. Stat. 9: 405–421

Subject Editor: Andrew Liebhold.