

Methods for testing for texture convergence using abundance data: a randomisation test and a method for comparing the shape of distributions

J. B. Wilson¹ and B. Smith²

¹ Botany Department, University of Otago, P.O. Box 56, Dunedin, New Zealand. Fax: +64 3 479 7583,
E-mail: Bastow@otago.ac.nz.

² Climate Impacts Group, Plant Ecology, Department of Ecology, Ekologihuset, Lund University,
223 62 Lund, Sweden

Keywords: Character distributions, Community structure, Convergence, Functional characters.

Abstract: Testing the ecological communities of different areas for convergence, in the sense of remarkable similarity in the characteristics of the species present, has a long history in biology. Recently, numerical methods have been developed for comparing community-level convergence to an explicit null model.

No valid method has been known for testing the significance of texture convergence when the species are weighted by their abundance. Six combinations of method variants are tested on random datasets. A valid P value (i.e., with $P \leq 0.05$ in no more than 5% of the cases) is obtained so long as for each species the distribution of abundances across sites is retained, and only the assignment of character values is randomised. Further restriction is not necessary for obtaining a valid P value, and can lead to a test with considerably lower power to detect convergence. The power of the test with free matching of character values to species is only moderate with 10 sites, though improved with larger numbers of sites.

Previous methods for detecting texture convergence have examined convergence only in the mean value for any character. It is possible that the external environment might be reflected in the community mean of a character, leaving the imprint of convergence on the shape of the distribution, rather than the mean. A method for comparing the shape is described, and it is shown that the null model is valid also for this test statistic.

Introduction

In attempts to discover the evolutionary and ecological forces structuring plant and animal communities, a common approach is to test for texture convergence. 'Texture' in this sense describes the set of values in particular characters across the species in a community, especially characters thought to be related to the niche of the species (Barkman 1979). The term 'ecomorphological convergence' has also been used (Blondel et al. 1984). For example, we might ask whether the set of values of body sizes in the avifauna of a semi-arid shrubland in Australia is similar to the set in comparable avifauna in North America (Wiens 1991).

The evolutionary concept of adaptation and the ecological concept of competition are the theoretical bases of the texture-convergence approach, and these are the concepts that the approach tests. Antecedents of this approach are attempts to find look-alike but unrelated species on

different continents, for example the marsupial mole of Australia and the placental mole, or the shark, the ichthyosaur and the dolphin (Hildebrand 1974, Young 1981). This has been described as 'species-to-species matching' (Schluter 1990). The assumption is that similar physical environments on different continents have produced similar selection pressures, and hence similar-looking species.

This process will occur at the whole-community level as well as at the level of individual species, because the species will react not only to the physical environment but also to each other, through competitive exclusion and co-evolution. The result should in theory be texture convergence, i.e., greater similarity between the character values in different sites than expected under a null model. The competitive processes that produce co-adaptation in evolutionary time will also act, via competitive exclusion of some members of the species pool, to produce convergence between sites in ecological time. The sites can be

areas on different continents, or patches within a community, or samples at an intermediate scale, depending on the ecological question. The appropriate null model is one in which:

- the species, or rather their character values, are assigned at random to sites,
- the species richness of each site is kept as in the observed data, and
- the number of occurrences of each species is kept as in the observed data.

The opposite phenomenon, of texture divergence, is also of interest (e.g., Campbell and Werger 1988).

Like all attempts to demonstrate community structure, investigations of texture convergence have been controversial. Many early investigations had no comparison with a null model (Barbour and Minnich 1990). It has even been suggested that no null model is possible (Keeley 1992). However, Schluter (1986) was able to frame a null model in a way that allowed application of the analysis of variance, using a two-tailed *F*-test and seeking a value less than 1.0, and Wilson et al. (1994) generalised the approach by using a randomisation test, allowing a wide range of null models.

However, there has been little consideration of what methodology is appropriate and valid. In any randomisation test there are two components: null model and test statistic. We consider one outstanding problem in each component:

- the appropriate null model when there is species overlap and the species occurrence data are in terms of species abundance, and
- a test statistic that captures information additional to the mean.

There are two statistical questions to be answered, as with any test: *size* and *power* (Gonzalez and Manly 1998). Size indicates whether the test is valid, in the sense that it does not give spurious significances. We can define size by the percentage of times that the method yields a significant result when analysing random data, i.e., when the null model is true. It should equal *alpha*, e.g., in 5% of random datasets *P* should be 0.05 or lower. A test with size greater than this will give spurious significance (Wilson 1995). Power is the ability to detect real effects. It is the percentage of times that a significant result is given when the null model is false, i.e., when there is non-random structure in the data of the kind that the method is seeking. For a given

degree of departure from the null model, the greater the power the better.

The appropriate null model when there is species overlap and abundance data

The simplest comparisons are those between areas with entirely distinct species pools. Even with inter-continental comparisons, this is rarely achievable, because of the existence of cosmopolitan species and because of the spread of species as exotics. Investigations of ecological rather than evolutionary sorting must use within-community comparisons (e.g., Smith et al. 1994), which inevitably involves overlap in species composition. When texture is measured using species presence/absence (i.e., binary) data, it is easy to ensure that species overlap does not cause spurious convergence (Wilson et al. 1994, Smith et al. 1994). It is necessary only to ensure that in each randomisation each species occurs in the same number of sites as in the observed data (and without any one species being assigned twice to the same site). So long as this is done, although there is species overlap in the observed data, there is the same overlap in the randomisations. Therefore any conclusion of convergence is safe against objections that the 'convergence' found is due to the presence of some species in common between sites.

However, it is desirable to take into account the abundance of each species. Sites are hardly convergent that have complements of species that match species-for-species in the values of a particular functional character, but in which abundances of matching pairs of species are very different. Texture calculations should therefore use abundance values (ideally of biomass or density for animals, biomass for plants) to produce a weighted mean value for each character over the species present. When abundance weightings are used, and those are randomised across all species, another problem arises. In a randomised dataset, an occurrence of a species may become associated with any abundance value from the observed site. Species that occur several times in the observed dataset, and in each occurrence with high abundance, will on average be assigned lower abundances in the randomised data. As a result, such species will influence texture more in the observed data than in the randomised data. Abundant species common to more than one site, having similar characters and being heavily weighted, would cause the sites to be similar in the weighted mean value for any character, and thus lead to a low variance in texture among the observed sites. However, their net effect on the among-site variance in the randomised datasets would be much lower. As a result, observed sites would tend to be more similar than most randomised ones, even if there were no overall convergence. In the example of Table 1,

Table 1. A demonstration of spurious evidence for convergence, due to an inappropriate null model which randomises the abundances between species. Sites '1' and '2' are more similar in weighted mean in the observed data than in the randomised, because in the latter the effect of species '3', with high abundance, is dissipated. The units are arbitrary.

	Character value (e.g. leaf width)	Observed abundance (biomass or density)		Randomised abundance	
		Site 1	Site 2	Site '1'	Site '2'
Species 1	1	20	30	2	-
Species 2	5	10	5	-	80
Species 3	100	90	80	20	1
Species 4	20	2	-	90	5
Species 5	2	-	1	10	30
Weighted mean character value over species		74.7	69.5	31.3	5.7

the observed sites are similar in weighted-mean character value (74.7 vs 69.5) because of the dominance of species '3' with a consistent leaf width, but because of the randomisation of abundances the randomised sites are not similar (31.3 vs 5.7). This will give a spurious indication of convergence when there is really none. Smith et al. (1994) saw this problem, but could see no way to overcome it.

A possible solution is to not randomise the abundance values, thus avoiding the problem of splitting up the abundances originally belonging to one species. But would this solve the problem, or would the method still give spurious significance? Would such a method be able to detect texture convergence?

There is also the issue of randomising the character values of species that occur in more than one site. Clearly, the number of occurrences of each species (i.e., each character value) should be kept in the randomisations the same as those observed. When the value of one of their characters is randomised, should the value be allowed to associate with any species, only with a species that occurs in the same number of sites, or only with a species that occurs in the same sites?

We especially bear in mind in our investigations the worst-case situation of a collection of samples from one area, with considerable species overlap, though comparable situations arise for any set of samples with appreciable species overlap.

A test statistic that captures information additional to the mean

Most convergence work has based the measurement of convergence in the mean value of each character in the different sites, calculating as test statistic something like the variance of these means between sites, or the mean absolute deviation (MAD). Convergence is then identified as greater similarity in mean than expected at random, i.e., low variance or MAD. Evolutionary and/or ecological convergence would tend to produce an identical

distribution of species values in character space (Wiens 1989), and as a test statistic the mean represents only a fraction of this information. It is possible that communities might converge in the shape of the distribution of character values without converging in mean (Fig. 1). If there is a difference in the external environment of the communities being compared, which is inevitable to some degree, it is quite likely that this would affect all the species similarly and thus affect the mean, but perhaps retain the shape of the distribution.

In the presence/absence case, the shape of the distribution could be seen simply by ranking the species in order, and basing the test statistic on the shape of the graph (Fig. 2). When the contributions of species are being weighted by their abundance (perhaps transformed), the calculation of difference in the distributions needs to take this weighting into account. We suggest this could be done by ranking the species at one site on a histogram in order of their character values (the latter adjusted for the overall mean in that site), and making the width of the histogram columns equal to the weight for that species (i.e., its abundance or transformed abundance; Figs 3, 4). Then the difference between two sites can be calculated as the area between the histograms (Figs. 3, 4). The test statistic is obtained by adding these differences over all possible pairs of sites. We examined the test statistic by testing it on the random communities.

Methods

Random datasets, to determine size

Generation of random datasets. 100000 simulated datasets were generated, each with:

- 10 or 30 sites,
- 30 species.

Each species received:

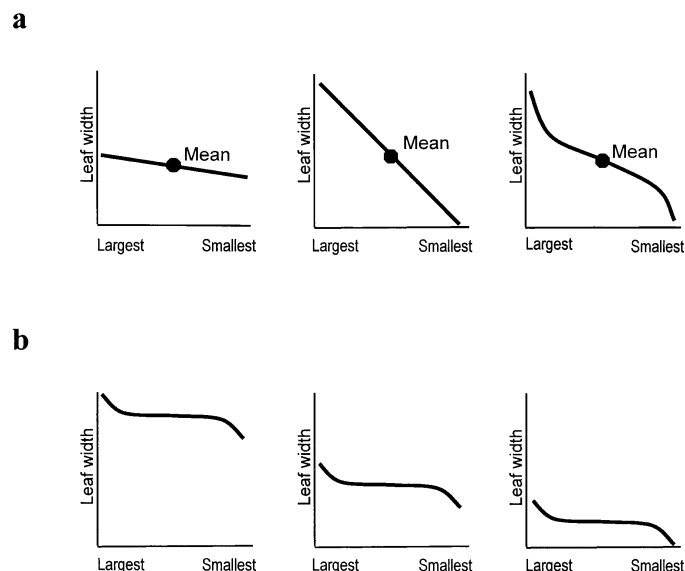


Figure 1. Comparison between means and distributions. We can look for convergence both in the mean value of, e.g., leaf width over the species, and in the shape of the distribution across species. Suppose we arrange species in order of decreasing leaf width. **a:** Three communities convergent in mean, but not in distribution. **b:** Three communities convergent in the shape of the distribution, but not in mean.

- a random character value in the range 1 to 100, with a rectangular distribution, the same value over all the sites in which the species occurred.
- a mean abundance (A) in the range 1 to 10000, with a geometric distribution (Wilson 1991), i.e., $A = \exp(\text{Random}(\ln(1), \ln(10000)))$, where ‘Random’ returns a random value with a rectangular distribution




Figure 2. The shape of the distribution of character values in two sites, expressed in terms of presence/absence of the species.

between the two limits given.

The occurrence of each species in each site was at random, with a 0.8 chance of occurrence. If it did occur, its abundance was determined at random (with a geometric distribution) in the range $A \times 0.75$ to $A \times 1.33$. That is, $\text{abundance} = \exp(\text{Random}(\ln(A \times 0.75), \ln(A \times 1.33)))$. The rather narrow range was chosen to provide a strong test of the ability of the method to cope with species of consistently low or high abundance.

Testing for convergence. Each simulated dataset was analysed by a randomisation test as if it were an ‘observed’ dataset. The analysis was based on that of Wilson et al. (1994) and Smith et al. (1994), and was performed separately for each character. The test statistic being used (see below) was calculated on the ‘observed’ data. Three hundred and fifteen randomisations were then performed (the number was selected to minimise boundary problems between P-value categories). In these randomisations, the pattern of species occurrence across sites was always held the same as in the observed. The values for the character were then randomised as specified under the null model, by replacing the value of each species with that of another species, selected at random without replacement, using one of three variants:

- *Character values matched freely:* the character value of a species in a random dataset was selected from any of the observed species,

- *Character values matched by frequency*: the character value of a species in a random dataset was selected from a species of equal frequency (i.e., occurring in the same number of sites), or
- *Character values matched by pattern*: the character value of a species in a random dataset was selected from a species with the same pattern of occurrences across sites (i.e., occurring in the same set of sites).
- *Abundance values left undisturbed*: the abundance values for each species were left as in the 'observed' data, only the character values were randomised.
- *Abundance values randomised*: the abundance values for each site were kept with that site, but were randomised among the species present at that site.

The abundance values were used in one of two ways:

The test statistic (TS) was calculated on each randomised dataset. As is usual in randomisation tests, the value of the test statistic calculated from the 'observed' dataset (observed TS) was compared with the expected TS, calculated as the mean of the TS values from the randomised datasets. The probability of the observed TS was calcu-

Species	Character value	Site A				Site B				Overall ⁴
		Abundance	Relative abundance ¹	Accumulated abundance ²	Adjusted character value ³	Abundance	Relative abundance	Accumulated abundance	Adjusted Character value	
Sp. 1	160	10	0.435	0.435	154.92	4	0.190	0.190	165.57	
Sp. 2	130	0	0.000	0.435	124.92	2	0.095	0.286	135.57	
Sp. 3	50	1	0.043	0.478	44.92	9	0.429	0.714	55.57	
Sp. 4	20	5	0.217	0.696	14.92	6	0.286	1.000	25.57	
Sp. 5	15	7	0.304	1.000	9.92	0	0.000	1.000	20.57	
Weighted mean character value ⁵		80.65				70.00				75.57

Accumulation of area between the distributions for the two sites					
Site A species	Site B species	Species boundary ⁶	Distance between boundaries ⁷	Difference in adjusted character value ⁸	Area between quadrats ⁹
1	1	0.190	0.190	10.65	2.029
1	2	0.286	0.095	19.35	1.843
1	3	0.435	0.149	99.35	14.810
3	3	0.478	0.043	10.65	0.463
4	3	0.696	0.217	40.65	8.837
5	3	0.714	0.019	45.65	0.851
5	4	1.000	0.286	15.65	4.472
Total area between quadrats					33.305

1 Scaled so the total abundance in the site is 1.0

2 Accumulated in the order of calculation, starting (arbitrarily) from the largest character value

3 (Character value) - (Weighted mean character value for Site A) + (Overall weighted mean character value)

4 For all occurrences of species in both sites

5 Sum of the character values, each weighted by the abundance of the species in that site

6 The accumulated relative abundances for the two sites (omitting any absent species) arranged in order of increasing value

7 (Current boundary) - (Previous boundary)

8 (Adjusted character value for the Site A species in that range of accumulated abundance)

- (Adjusted character value for the Site B species in that range of accumulated abundance)

9 (Distance between boundaries) * abs (Difference in adjusted character value)

Figure 3. Method for comparing the shape of the distribution of the values of a character in the two sites (see Fig. 4). The overall vertical position of the two sites is adjusted additively so the weighted mean character values for the two sites are equal.

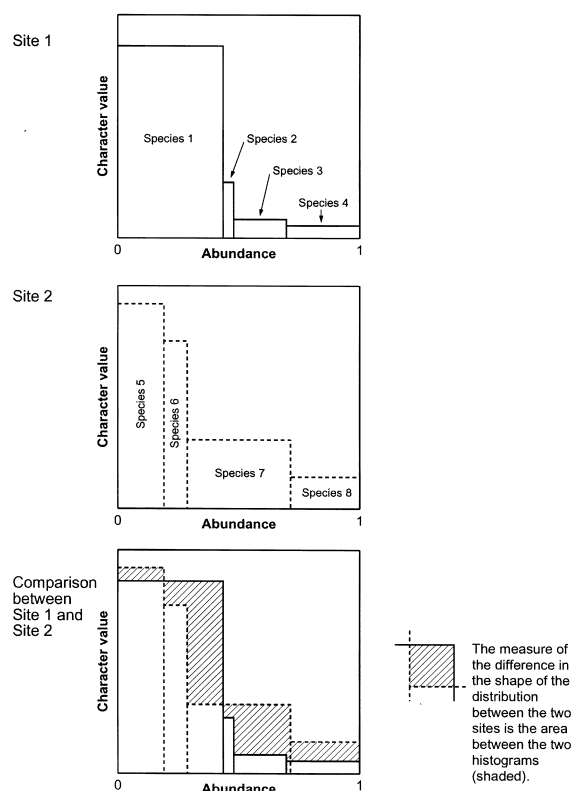


Figure 4. The shape of the distribution of character values in two sites, expressed in terms of quantitative abundance. The data are those of Fig. 3.

lated as the proportion of randomisations in which the TS value was equal to, or more extreme than, the observed TS, multiplied by 2.0 to effect a 2-tailed test. In such randomisation tests, the TS from the randomised dataset can sometimes be equal to that from the observed. Such cases must be counted in both tails. This can sometimes give a P value greater than 1.0. Such values were counted as 1.0.

Two test statistics were used: (1) the sum over the sites of the absolute difference of the site mean from the overall mean for the character (the formula given as D in Wilson et al. 1994, with $p = 1.0$), and (2) the area between the abundance-weighted distributions, after the character values in the two sites had been additively adjusted to bring them to the same weighted mean (Figs. 3, 4).

Structured datasets, to determine power

In order to test the power of the six method variants, additional datasets were created with some degree of texture convergence built in. The same process of dataset construction was used as described above. However,

whenever a species was due to be added, the method calculated:

- the overall mean value of the character over all the species,
- the mean value over the species that had already been entered into the random community, weighted by their abundance in the community,
- the mean value that the random community would have were the next species added, with its character value and its abundance determined at random, as above.

The species was rated 'unsuitable' if the deviation of the mean character value in the site (weighted by abundance) after addition of the new species would be further from the mean over all species than it was before (e.g., ecologically if a broad-leaved species was trying to enter a community that already consisted mainly of broad-leaved species, or if a large-bodied bird was trying to enter a community that already consisted mainly of large-bodied birds). When the next species to be entered was unsuitable by this criterion, the probability of entry of that species was reduced from the standard 0.8 (see above) to $0.8C$, where competitive exclusion parameter C indicates the strength of competitive exclusion (from 0.0 meaning that a species with unsuitable characters could never enter, to 1.0 meaning that there was no effect of character unsuitability).

Results

Simulated datasets, to determine size

When the abundance values were randomised, significant 'convergence' in the mean was seen much more often than the 5% of datasets in which it should (i.e., the size of the test was too high; Fig. 5b, d, f). When abundance values were left undisturbed, and species were either Matched-freely or Matched-by-frequency, the dis-

Table 2. Size of randomisation tests using the shape of the distribution (Fig. 4) as the test statistic, i.e., the proportion of random datasets that give results significant at $P = 0.05$. An ideal test would give a size of 0.05.

Character values	Abundances	Size
Matched-freely	Undisturbed	0.051
	Randomised	0.677
Matched-by-frequency	Undisturbed	0.050
	Randomised	0.830
Matched-by-pattern	Undisturbed	0.032
	Randomised	0.827

tribution of P values was very close to rectangular (Fig. 5a, c), with only an excess of P values close to 1.0. However, with abundances left undisturbed and species Matched-by-pattern, the percentage of P values < 0.05 was somewhat lower than 5%, with a considerable excess of P values close to 1.0 (Fig. 5e).

Results in the same direction were found using the shape of the distribution as the test statistic (Table 2).

Structured datasets, to determine power

When convergence was built into the simulated datasets (Fig. 6), the Matched-freely and Matched-by-frequency methods gave a considerable excess of values

< 0.05 , as they should do (Fig. 6a,c). Power was only about 0.35 with 10 sites, and the power of Matched-by-frequency was slightly greater than that of Matched-freely. When there were more sites, the power of the Matched-freely method increased, and that of Matched-by-frequency increased, becoming equal to Matched-freely (Fig. 6b, d). The Matched-by-pattern variant gave a high P value (> 0.95) in more than half of the analyses, and practically never gave the a significant result (i.e., $P < 0.05$).

Discussion

The results make it clear that the appropriate test for convergence in texture/ecomorphy with abundance

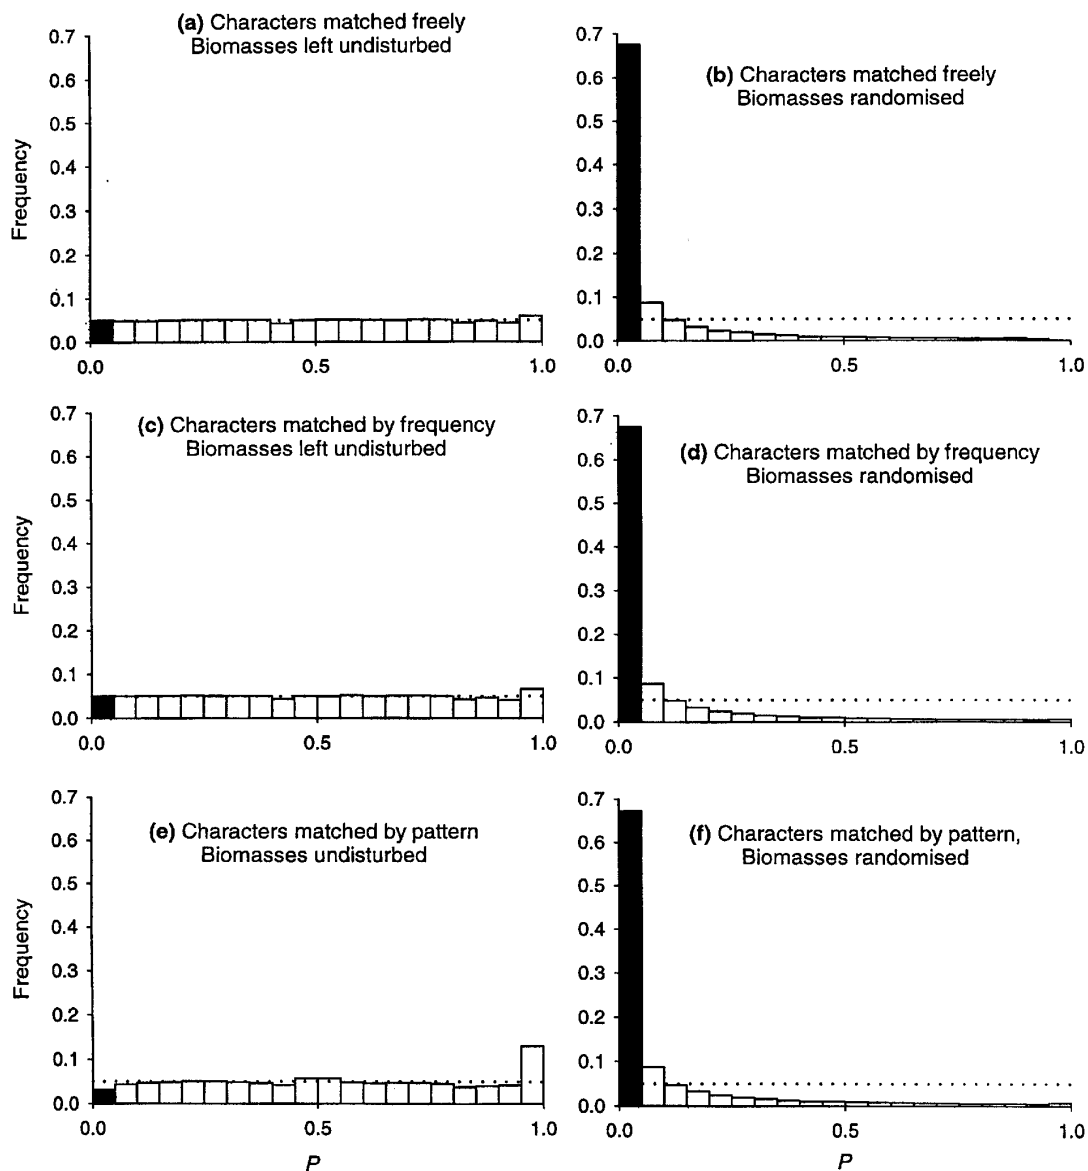


Figure 5. The size of the tests for convergence of the mean: frequency distributions of P values obtained analysing random datasets by six methods. For a valid test, the distribution should be approximately rectangular, and especially the proportion of significances (black column: ■) should be 0.05 (as indicated by).

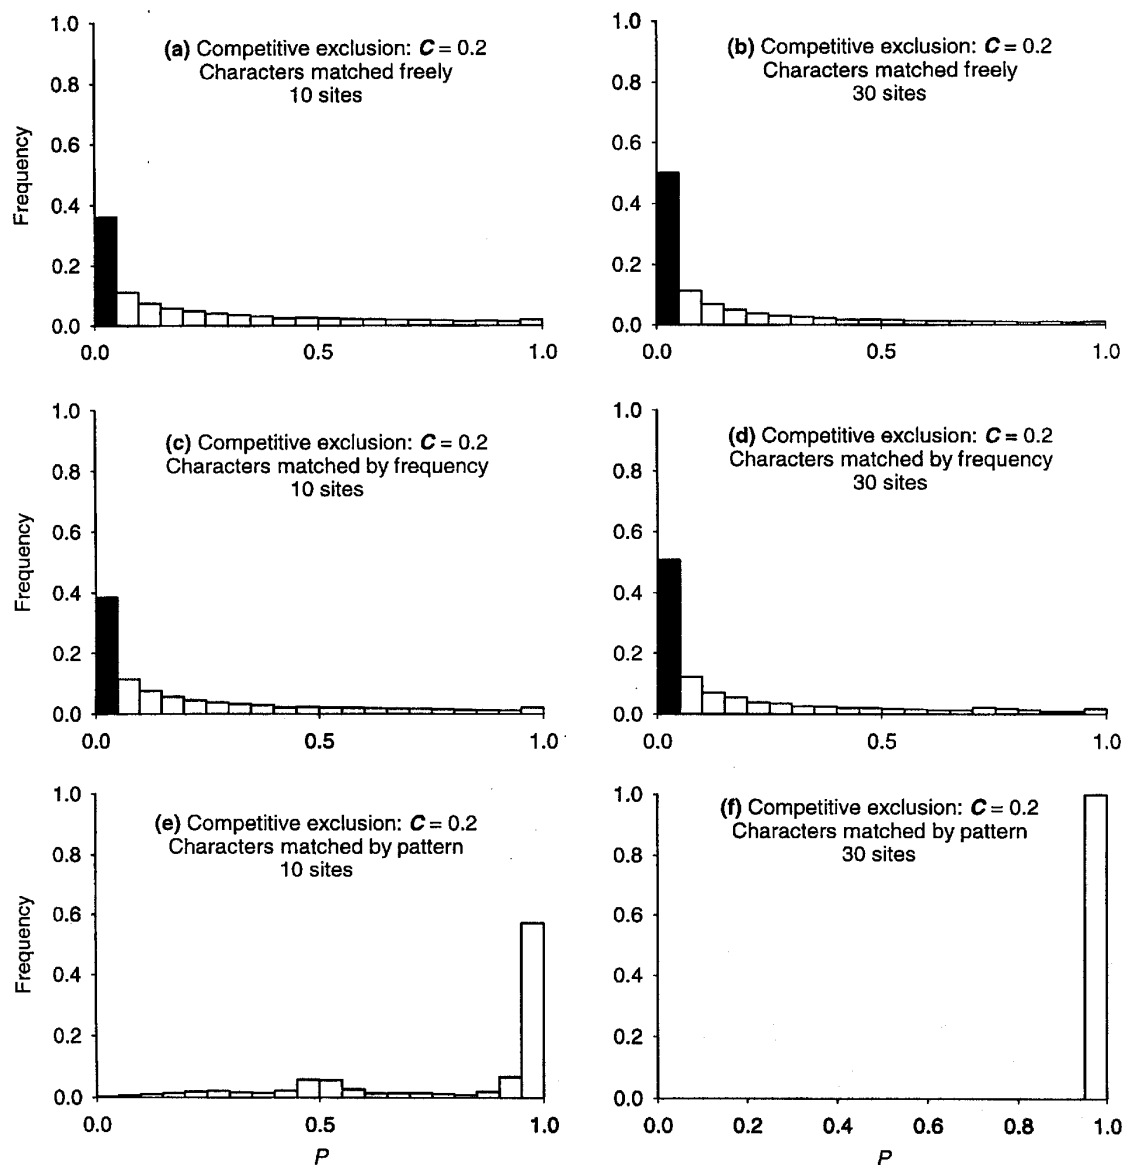


Figure 6. The power of the tests for convergence of the mean: frequency distributions of P values obtained analysing, by three methods, datasets with 10 and 30 sites, structured to contain texture convergence. For a powerful test, the proportion of significances (black column: ■) should be as high as possible.

data involves leaving the abundance values as in the observed data. This does not comprise building texture convergence from the observed data into the null model, because texture is defined in terms of the species' characters, and those are still being randomised. What it does is to retain the rank consistency (Watkins and Wilson 1994), i.e., if there is in an observed dataset one species that is dominant in all sites, that feature will be retained in the randomisations (be it associated with a different value of the characters), as will more subtle effects of this type. Since the method is not intending to test for rank consistency, this amounts to following the dictum of Tokeshi (1986) that the null model should include

every feature of the observed data except the one it is desired to test.

Not to retain the abundance structure amounts to an error of the Jack Horner type (Wilson 1995), i.e., demonstrating an obvious feature by failing to build it into the null model, and thus failing to investigate the feature of interest. Here the obvious feature is that some species are consistently more abundant and others less abundant. This becomes especially obvious in a forest, when no filmy fern could possibly have the biomass typical of forest trees. If the character of a species is more or less consistent across sites (as it must be), this will give less similarity between sites in the randomisations than in the

observed data. All these intuitive arguments are confirmed by the close-to-rectangular results of Figs. 5a, c and the far-from-rectangular results of Figs. 5b, d.

No validity is lost in the Matched-by-frequency and Matched-by-pattern methods (Figs. 5c, e). However, there is a dramatic loss of power using Matched-by-pattern (Fig. 6e). As the number of sites increases, the number of possible patterns increases greatly, so when trying to swap the characters of a species it is difficult to find any species with an identical pattern, and thus eligible for swapping. The problem is even worse with larger numbers (Fig. 6f). In the worst case, there might be no pairs of species with identical patterns, and thus no randomisation possibilities. When the number of species is large relative to the number of sites, there is little such constraint in Matched-by-frequency. However, the number of possible frequencies is the number of sites, so as the number of sites increases, opportunity for species matching decreases.

The recommended method of randomisation is similar to, and thus confirms the validity of, the test used by Pillar (1996). It has been used here with a univariate quantitative measure, but it could be used with other types of data and other test statistics. Although the example of species-level convergence has been used here, the method could be used with any taxa above or below this level which can be described by characters.

The simulations have involved quantitative data, but presence/absence data are just a special case, with all abundances 0 or 1. Therefore, all the methods with appropriate size in these tests will also have appropriate size with presence/absence data. Our simulations were focused on the worst-case situation of a set of samples from one community, with consequentially very considerable overlap in species occurrences between samples. Because this is the worst case, demonstrations of validity will apply also to cases with less overlap. However, in between-continent comparisons an additional problem arises that it is not possible to find 10 or 30 samples, as used in the simulations here, because there are too few continents. Nested sampling of sites in continents would introduce further power questions.

The method for comparing distributions allows us to use more of the information in the data. An additional, potential advantage of analysing the distribution is coping with environmental differences. All texture work seeking texture convergence is bedevilled by the problem that it is impossible to exactly match environments between continents or other areas. Since texture also responds to environmental gradients (e.g., Campbell and Werger 1988),

environmental differences can lead to divergence in texture between the areas being compared, obscuring any texture convergence that may be occurring. If samples along environmental gradients are available, matching in one environmental factor would be possible by some kind of interpolation even if no actual samples match (Pillar 1999), but with multiple environmental factors more often the same combination cannot be found anywhere in another area. The hope in our adjustment to equal overall means is that the physical environment will affect all species, leading to consistently higher or lower values of a character across the species of a community, so the adjustment will remove the effect. However, if preferred, our method of comparing distributions can be used without the adjustment to equal overall means. The allowance will not be perfect, but it is another step on the way to removing environmental effects when seeking assembly rules (Wilson 1999).

Conclusions

Randomisation methods should always be validated by running them on random data, to check both the concept and the algorithm. This is not always done. Some dramatic exposures have been made in ecology by taking methods that had been used to reach ecological conclusions, and showing that they gave significant results with random data (e.g., Simberloff 1976, Connor and Simberloff 1979, Wilson 1987, 1995, Manly 1996).

A valid test for texture convergence can be obtained using either presence/absence or quantitative data, so long as the abundance values are not randomised, i.e., the values for one species are kept together. Species' character values can be randomised freely, without loss of validity. Restricting random swapping to species with the same pattern of occurrences, and to a much lesser extent restricting it to species with equal frequency, is not necessary, and can lead to loss of power in the test.

References

- Barbour, M.G. and R.A. Minnich. 1990. The myth of chaparral convergence. *Israel J. Bot.* 39: 453-463.
- Barkman, J.J. 1979. The investigation of vegetation texture and structure. In: M.J.A. Werger (ed.), *The Study of Vegetation*. Junk, The Hague, The Netherlands, pp. 123-160.
- Blondel, J., F. Vuilleumier, L.F. Marcus and E. Terouanne. 1984. Is there ecomorphological convergence among mediterranean bird communities of Chile, California, and France? *Evol. Biol.* 18: 141-213.
- Campbell, B.M. and M.J.A. Werger. 1988. Plant form in the mountains of the Cape, South Africa. *J. Ecol.* 76: 637-653.
- Connor, E.F. and D. Simberloff. 1979. The assembly of species communities: chance or competition? *Ecology* 60: 1132-1140.

- Gonzalez, L. and B.F.J. Manly. 1998. Analysis of variance by randomization with small data sets. *Environmetrics* 9: 53-65.
- Hildebrand, M. 1974. *Analysis of Vertebrate Structure*. Wiley, New York, USA.
- Keeley, J.E. 1992. A Californian's view of fynbos. In: R. M. Cowling (ed.), *The Ecology of Fynbos*, Oxford University Press, Cape Town, South Africa, pp. 372-388.
- Manly, B.F.J. 1996. Are there clumps in body-size distributions? *Ecology* 77: 81-86.
- Pillar, V.D. 1996. A randomization-based solution for vegetation classification and homogeneity testing. *Coenoses* 11: 29-36.
- Pillar, V.D. 1999. On the identification of optimal plant functional types. *J. Veg. Sci.* 10: 631-640.
- Schluter, D. 1986. Tests for similarity and convergence of finch communities. *Ecology* 67: 1073-1085.
- Schluter, D. 1990. Species-for-species matching. *Am. Nat.* 136: 560-568.
- Simberloff, D. 1976. Trophic structure determination and equilibrium in an arthropod community. *Ecology* 57: 395-398.
- Smith, B., S.H. Moore, P.B. Grove, N.S. Harris, S. Mann and J.B. Wilson. 1994. Vegetation texture as an approach to community structure: community-level convergence in a New Zealand temperate rainforest. *N.Z. J. Ecol.* 18: 41-50.
- Tokeshi, M. 1986. Resource utilization, overlap and temporal community dynamics: a null model analysis of an epiphytic chironomid community. *J. Anim. Ecol.* 55: 491-506.
- Watkins, A.J. and J.B. Wilson. 1994. Plant community structure, and its relation to the vertical complexity of communities: dominance/diversity and spatial rank consistency. *Oikos* 70: 91-98.
- Wiens, J.A. 1989. *The Ecology of Bird Communities. Vol. 1. Foundations and Patterns*. Cambridge University Press, Cambridge, UK.
- Wiens, J.A. 1991. Ecomorphological comparisons of the shrub-desert avifaunas of Australia and North America. *Oikos* 60: 55-63.
- Wilson, J.B. 1987. Methods for detecting non-randomness in species co-occurrences: a contribution. *Oecologia* 73: 579-582.
- Wilson, J.B. 1991. Methods for fitting dominance/diversity curves. *J. Veg. Sci.* 2: 35-46.
- Wilson, J.B. 1995. Null models for assembly rules: the Jack Horner effect is more insidious than the *Narcissus* effect. *Oikos* 72: 139-143.
- Wilson, J.B. 1999. Assembly rules in plant communities. In: E. Weiher and P.A. Keddy (eds.), *The Search for Assembly Rules in Ecological Communities*, Cambridge University Press, Cambridge, UK, pp. 130-164.
- Wilson, J.B., Agnew, A.D.Q. and Partridge, T.R. 1994. Carr texture in Britain and New Zealand: Community convergence compared with a null model. *J. Veg. Sci.* 5: 109-116.
- Young, Z. 1981. *The Life of Vertebrates*. Oxford University Press, Oxford, UK.