

Simple parametric tests for trait–environment association

Cajo J. F. ter Braak¹  | Pedro R. Peres-Neto²  | Stéphane Dray³ 

¹Biometris, Wageningen University and Research, Wageningen, the Netherlands

²Department of Biology, Concordia University, Montreal, Québec, Canada

³CNRS, UMR 5558 Laboratoire de Biométrie et Biologie Évolutive, Université Lyon I, Villeurbanne, France

Correspondence

Cajo J. F. ter Braak, Biometris, Wageningen University and Research, Wageningen, The Netherlands.

Email: cajo.terbraak@wur.nl

Co-ordinating Editor: Otto Wildi

Abstract

Question: The community-weighted mean (CWM) approach is an easy way of analysing trait–environment association by regressing (or correlating) the mean trait per plot against an environmental variable and assessing the statistical significance of the slope or the associated correlation coefficient. However, the CWM approach does not yield valid tests, as random traits (or random indicator values) are far too often judged significantly related to the environmental variable, even when the trait and environmental variable are extrinsic to (not derived from) the community data. Existing solutions are the ZS (Zelený & Schaffers) modified test and the max (or sequential) test based on the fourth-corner correlation. Both tests are based on permutations which become cumbersome when many tests need to be carried out and many permutations are required, as in methods that correct for multiple testing. The main goal of this study was to compare these existing permutation-based solutions and to develop a quick and easy parametric test that can replace them.

Methods: This study decomposes the fourth-corner correlation in two ways, which suggests a simple parametric approach consisting of assessing the significances of two linear regressions, one plot-level test as in the CWM approach and one species-level test, the reverse of the CWM approach, that regresses the environmental mean per species (i.e. the species niche centroid) on to the trait. The tests are combined by taking the maximum *p*-value. The type I error rates and power of this parametric max test are examined by simulation of one- and two-dimensional Gaussian models and log-linear models.

Results: The ZS-modified test and the fourth-corner max test are conservative in different scenarios, the ZS-modified test being even more conservative than the fourth-corner. The new parametric max test is shown to control the type I error and has equal or even higher power than permutation tests based on the fourth-corner, the ZS-modified test and variants thereof. A weighted version of the new test showed inflated type I error.

Conclusion: The combination of two simple regressions is a good alternative to the fourth-corner and the ZS-modified test. This combination is also applicable when multiple trait measurements are made per plot.

KEYWORDS

community ecology, community-level test, CWM of traits, environmental gradients, fourth-corner, functional traits, modified test, species niche centroid, species-level test, statistical ecology, trait–environment relationship

1 | INTRODUCTION

The community-weighted mean (CWM) approach is perhaps the most popular way to assess how individual traits and individual environmental variables correlate and how this link is relevant to the process of community assembly (see Kleyer et al., 2012; McGill, Enquist, Weiher, & Westoby, 2006). The CWM approach takes averages of traits of species that are present in each of n plots (communities), possibly weighted by species abundance (Figure 1), and correlates the resulting trait means to one or more environmental variables that characterize the communities (Figure 2a). However, the CWM-based correlation can be unduly variable (see Supporting information Appendix S1 for an extreme example) and presents major issues with statistical testing (Peres-Neto, Dray, & ter Braak, 2017). The problem occurs in both the parametric and permutation test versions of the CWM approach. Also, the statistic is too sensitive to only a few species contributing to the correlation between trait and environment.

For example, “yellow coloured flowers” may be more predominant in highly fertilized soils than in unfertilized soils just because a single species (e.g. dandelion) responds positively to fertilization of plots regardless of how plant species with flowers other than yellow respond to fertilization (Šmilauer & Lepš, 2014). This artificial example suggests that a species-level approach is also required as follows: it should require more than just one species to influence trait–environment correlations.

A species-level approach has also been suggested using the average environment value across all plots in which the species is present, again possibly weighted by abundance (Figure 1; Kleyer et al., 2012). The resulting means are then correlated to species traits (Figure 2b), which we refer to here as the species’ niche centroids (SNC) approach (Peres-Neto, Dray, & ter Braak, 2017). Ackerly, Knight, Weiss, Barton, and Starmer (2002) applied both the plot-level and the species-level approach and found they may generate different results and corresponding interpretations. Figure 2

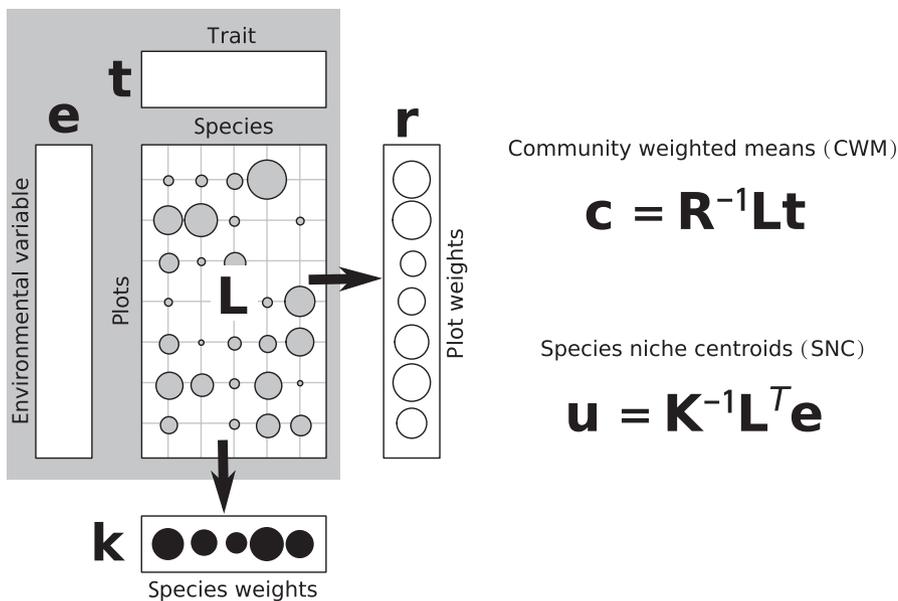


FIGURE 1 The data needed to assess the trait–environment association (abundance data table \mathbf{L} on the abundances of species across plots, values for environmental variable \mathbf{e} and functional trait \mathbf{t}) and derived statistics (CWM \mathbf{c} and species niche centroids \mathbf{u}) where $\mathbf{R} = \text{diag}(\mathbf{r})$ and $\mathbf{K} = \text{diag}(\mathbf{k})$, i.e. diagonal matrices, with the plot weights and species weights on the main diagonal, respectively. The size of circles (black for species, white for plots and grey for species–plot combination) is proportional to the abundances

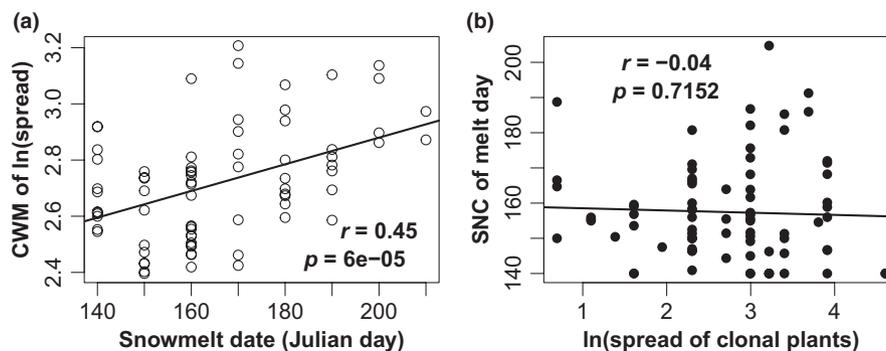


FIGURE 2 Regression analyses of the CWM of lateral spread on snowmelt day (a) and SNC of snowmelt day on lateral spread (b) in the aravo data set of the R package ade4 (black circles for species, white circles for plots), showing that the significance of the one regression does not imply significance of the other (p stands for p -value) and that even the signs of the Pearson correlations (r) may differ between analyses. The max test results in a p -value of 0.7152 (the maximum of the two p -values). The data thus provide no evidence for association between lateral spread and snowmelt day (R code of this example in Supporting information Appendix S7)

demonstrates this issue in the aravo data available in the R package *ade4* (Choler, 2005; Dray & Dufour, 2007): the mean lateral spread of plants is positively correlated with the snowmelt day along the snowmelt gradient (Figure 2a). However, at the species level, there is no such correlation: the mean melt day of plants is not related to their lateral spread; the correlation coefficient is even slightly negative (Figure 2b). This raises the question of how to combine the results of the plot-level approach and the species-level approach into one unique and logical answer.

Zelený and Schaffers (2012) also identified the issue of the CWM approach in the context of (Ellenberg) indicator values. The mean indicator value is clearly a CWM (where weights are either cover or simply 1 for presence and 0 for absence). By taking random indicator values, standard regression and ANOVA approaches resulted in increased type I error rates (i.e. rejected the null hypothesis more often than expected by the significance level when H_0 is true) as verified by Wildi (2016). Wildi (2016) pointed out that Zelený and Schaffers (2012) arguably considered an even more complex problem than we described in Figure 1; they related the mean indicator values to groups or gradients derived by cluster analysis and ordination, respectively. These groups or gradients were derived from the abundance data table (*L* in Figure 1), and therefore, it would have been surprising if they were not correlated to the mean indicator value that is also derived from *L* (Figure 1). Hawkins et al. (2017) referred to this case as relating two intrinsic variables. However, whereas the mean indicator value is “intrinsic,” the indicator value itself is not (Wildi, 2016). The null hypothesis considered by Zelený and Schaffers (2012) was simpler, namely that the indicator value is unrelated to species abundances. To test this null hypothesis, they proposed a “modified permutation test” (Zelený & Schaffers, 2012; Hawkins et al., 2017; hereafter referred as to the ZS-modified test). Inspired by a preprint (Zelený, 2016), we investigated in this paper whether their modified permutation test works in the context of trait–environment relationships (two extrinsic variables). In the ZS-modified permutation test, significance is assessed by a permutation test in which the indicator values (i.e. trait values in our case) are permuted among species and the test statistic is the CWM-based correlation (Figure 3). The test thus combines species-based permutation with a plot-level test statistic.

An alternative to these approaches is the fourth-corner correlation (Figure 3), which calculates a (weighted) Pearson's correlation between a standardized trait and environmental variable using the community table (abundance or presence–absence) as weight matrix (Dray et al., 2014; Legendre, Galzin, & Harmelin-Vivien, 1997). The fourth-corner correlation uses all the plot–species combinations (Figure 3) and so avoids the issue that the plot- and species-level correlations are unequal or even different in sign. Statistical testing of the fourth-corner correlation proceeds by permutation testing, for which five different permutation tests have been proposed and evaluated in Dray and Legendre (2008), among which permutation of the plot–species combinations. However, it turned out that a valid statistical test of trait–environment association can only be obtained by combining two tests (ter Braak, 2017; ter Braak, Cormont, & Dray,

2012): a test based on plot permutation of environmental values and a test based on species permutation of trait values. These two tests assess the relationship of the abundance table with the environmental variable and the trait, respectively, and both tests need to be significant to test for the correlation between trait and environment (Figure 3). In this case, correct type I error rates are assured by taking the maximum of the resulting two *p*-values as the final *p*-value. We call this the “max test” (ter Braak, 2017; earlier names used were sequential test (ter Braak et al., 2012) and row–column permutation test (Peres-Neto et al., 2017) or max *r/c* test (ter Braak, Peres-Neto, & Dray, 2017)). The max test controls for the type I error rate, which means that the rejection rate is at most the nominal level of the test (e.g. 0.05) when there is no trait–environment association (ter Braak et al., 2012). Peres-Neto et al. (2017) applied the max test also to CWM- and SNC-based correlations and verified, by simulation, that the resulting tests control the type I error, but this results in less powerful tests than the test based on the fourth-corner correlation. In this study, we complement their study with a max test in which the CWM correlation is used for the plot-level test and the SNC correlation is used in the species-level test.

The results of the simulation study by Peres-Neto et al. (2017) suggest an unpleasant property of the modified test of Zelený and Schaffers (2012). The ZS-modified test is identical to the species-level test using the CWM-based correlation and, from Figure 3a (column permutation, only Env. random) in Peres-Neto et al. (2017), it can be concluded that it is overly conservative in the scenario where there is an effect of the trait but no environmental effect in structuring species distributions. This likely leads to a reduction in statistical power in non-null scenarios, particularly for low correlation between traits and environment (this was not tested by Peres-Neto et al. (2017)).

The fourth-corner and ZS-modified test are both permutation tests and these become slow if large numbers of permutations are required, particularly in cases with multiple traits and multiple environmental variables (Dray et al., 2014). When many traits–environment correlations are tested in the same study, some form of correction for multiple testing is required to control the overall type I error rate, e.g. by Bonferroni correction or its improved form, Holm correction (Verhoeven, Simonsen, & McIntyre, 2005). At least one correlation is then judged significant when its original *p*-value is less than the nominal level divided by the number of correlations, e.g. with ten traits and ten environmental variables, $0.01/(10 \times 10) = 0.0001$. This presents a problem for permutation testing as it requires at least 10,000 random permutations to be generated per test, which may take non-trivial computing time. Note that the required number of permutations increases with the number of trait–environment combinations to be tested.

For simple problems, such as testing a correlation or regression coefficient, parametric tests often give similar answers to permutation tests based on the same units and test statistic, at least when the number of units and the number of random permutations are large and the data are not too far from normal (Anderson & Robinson, 2001; Anderson & ter Braak, 2003). When low *p*-values are required

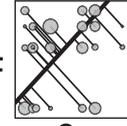
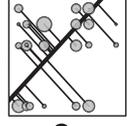
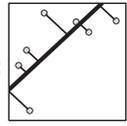
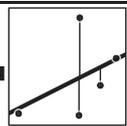
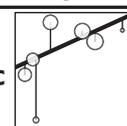
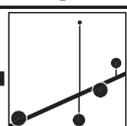
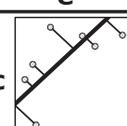
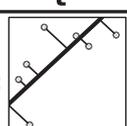
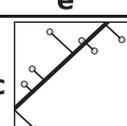
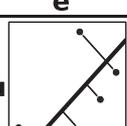
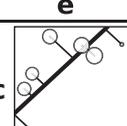
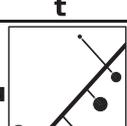
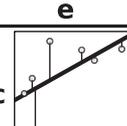
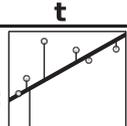
Method	Plot-level test	Species-level test
fourth-corner	 Fourth-corner statistic tested by permuting values in e	 Fourth-corner statistic tested by permuting values in t
ZS-modified test		 Correlation between c and e tested by permuting values in t
lm CWM/SNC	 Parametric test of the slope of the regression of c on e	 Parametric test of the slope of the regression of u on t
wlm CWM/SNC	 Parametric test of the slope of the weighted regression of c on e	 Parametric test of the slope of the weighted regression of u on t
cor CWM	 Correlation between c and e tested by permuting values in e	 Correlation between c and e tested by permuting values in t
cor CWM/SNC	 Correlation between c and e tested by permuting values in e	 Correlation between u and t tested by permuting values in t
wcor CWM/SNC	 Weighted correlation between c and e tested by permuting values in e	 Weighted correlation between u and t tested by permuting values in t
lm CWM	 Slope of the regression of c on e tested by permuting values in e	 Slope of the regression of c on e tested by permuting values in t

FIGURE 3 Methods examined and, when in grey, presented in the main text. All methods (except the ZS-modified test) consist of two tests (plot-level and species-level) that are combined by taken their maximum p -value (max test). All methods are based on a test statistic that is either a (weighted or unweighted) correlation (*cor*) or the slope of a (weighted or unweighted) linear regression model (*lm*). The size of circles (black for species, white for plots and grey for plot-species combination) is proportional to the abundances (equal sizes for unweighted statistics)

for individual tests, it may thus be an advantage to replace a permutation test by a parametric test.

The test of the fourth-corner correlation requires the combination of two permutation tests. In this study, we investigate the possibility of replacing these permutation tests with fast and straightforward parametric tests, namely one test based on the plot-level correlation as in the CWM approach and one test of the species-level correlation as in the SNC approach. This is achieved by two simple linear regressions: the regression of trait means (CWMs) on the environmental variable and the regression of species niche centroids (SNCs) on the trait (either using weighted or unweighted regression; Figure 3). We also show, through theory and simulation, how the ZS-modified test and the combination of CWM- and SNC-based tests (by permutation or parametric) relate to the fourth-corner approach and whether these tests control the type I error

also when trait and environmental variables are not necessarily normal. The parametric max test based on two unweighted simple regressions is shown to be about equally or even slightly more powerful than the fourth-corner. This max test is immensely simpler than the GLM bootstrap approach of Warton, Shipley, and Hastie (2015), equally powerful (ter Braak, 2017; ter Braak et al., 2017) and is also applicable when multiple trait measurements are taken per plot.

2 | THEORY AND METHODS

2.1 | Combining CWM- and SNC-based regressions

For simplicity of notation, we consider one environmental variable e with values $[e_i]$ ($i = 1, \dots, n$) for n plots and one trait t with values

$[t_j]$ ($j = 1, \dots, S$) of S species and let $\mathbf{L} = [l_{ij}]$ denote the plots \times species table containing abundances or presences-absences (1/0) values (Figure 1). In the CWM approach, the average trait in plot i , weighted by the abundance values, is calculated first as:

$$c_i = \sum_{j=1}^S l_{ij} t_j / l_{i+} \quad \text{or, in matrix notation, } \mathbf{c} = \mathbf{R}^{-1} \mathbf{L} \mathbf{t},$$

where l_{i+} is the sum of all abundances for plot i and \mathbf{R} is an $n \times n$ diagonal matrix with the plot total abundances $l_{1+}, l_{2+}, \dots, l_{n+}$ in the main diagonal (Figure 1). If l_{ij} is either 0 or 1, then c_i is simply the average of the trait values of the species present in the plot. Second, the averaged trait values are plotted against the environmental variable, a regression line is fitted through the points (in formula notation, $\mathbf{c} \sim \mathbf{e}$) and the corresponding correlation coefficient is calculated (see Figures 2a and 3; Supporting information Appendix S1).

A strong and significant correlation suggests that the trait means are associated in a non-random way with the environmental variable in question, but this does not yet demonstrate trait-environment association (Šmilauer & Lepš, 2014). Even if the trait values are completely random, the CWM-based regression is often significant, particularly if species abundances are strongly related to the environmental variable ($\mathbf{L} \leftrightarrow \mathbf{e}$) (Peres-Neto et al., 2017). We also verified this result here through the simulation. Moreover, with a random trait, the trait means may show little variation across plots, resulting in unduly variable correlations (Peres-Neto et al., 2017). Indeed, in the artificial data in Supporting information Appendix S1, a tiny change in two trait values for two (of five) species causes the correlation to flip from +1 to -1. As such, the CWM-based correlation is not suited to measure the strength of the trait-environment association. Moreover, the CWM approach alone is not sufficient to generate valid tests on trait-environment associations. It must be extended with a species-level approach, for example, the SNC approach.

In the SNC-approach, the average environment in which species occur, weighted by their abundance values, is calculated first (Ackerly et al., 2002; ter Braak & Looman, 1986):

$$u_j = \sum_{i=1}^n l_{ij} e_i / l_{+j} \quad \text{or, in matrix notation, } \mathbf{u} = \mathbf{K}^{-1} \mathbf{L}^T \mathbf{e},$$

where l_{+j} is the sum of all abundances for species j and \mathbf{K} is an $S \times S$ diagonal matrix with the species total abundances $l_{+1}, l_{+2}, \dots, l_{+S}$ on the main diagonal (Figure 1). Second, the averaged environment values are plotted against the trait, a regression line is fitted through the points (in formula notation, $\mathbf{u} \sim \mathbf{t}$) and the corresponding correlation coefficient is calculated (Figures 2b, 3 and Supporting information Appendix S1). A strong and significant correlation suggests that the environmental means are associated in a non-random way with the trait, but this does not yet demonstrate the trait-environment association (Šmilauer & Lepš, 2014). Even if the environmental values are completely random, the SNC-based regression is often judged significant, particularly if species abundances are strongly related to the trait ($\mathbf{L} \leftrightarrow \mathbf{t}$) (Peres-Neto et al., 2017). The SNC approach alone is

thus not sufficient to generate valid tests on trait-environment associations and must be combined with the CWM approach.

Dray and Legendre (2008) noted that trait-environment association requires two links to exist: both the link $\mathbf{L} \leftrightarrow \mathbf{e}$ and the link $\mathbf{L} \leftrightarrow \mathbf{t}$. To statistically demonstrate the existence of the link $\mathbf{L} \leftrightarrow \mathbf{e}$, we must test the null hypothesis $\mathbf{L} \nleftrightarrow \mathbf{e}$. But if \mathbf{L} is unrelated to \mathbf{e} ($\mathbf{L} \nleftrightarrow \mathbf{e}$), then any combination of \mathbf{L} , such as \mathbf{c} , is unrelated to \mathbf{e} ($\mathbf{c} \nleftrightarrow \mathbf{e}$). Through logic reasoning, we thus have that if $\mathbf{c} \leftrightarrow \mathbf{e}$, then $\mathbf{L} \leftrightarrow \mathbf{e}$. Using the CWM approach, we can statistically demonstrate the link $\mathbf{c} \leftrightarrow \mathbf{e}$, and thus in consequence the link $\mathbf{L} \leftrightarrow \mathbf{e}$. The first link can be thus tested with the CWM approach. Analogously, the second link can be tested with the SNC approach. Both tests must be significant to conclude that the trait-environment association exists (ter Braak et al., 2012). An alternative way of phrasing this result is that the final p -value of the test on the trait-environment association is the maximum of the p -values of the two tests. This combined CWM/SNC test (denoted as *lm* CWM/SNC in Figure 3, where *lm* stands for linear model as in the R computing environment) is a max test in the sense of ter Braak et al. (2012) and controls the type I error rates, provided each individual test controls the type I error. Note that "control" means here that the rejection rate of the null hypothesis of no association using this test is less than or equal to the nominal level of the test when there is in fact no association (ter Braak et al., 2012).

2.2 | Relation to the fourth-corner correlation and the ZS-modified test

The CWM- and SNC-based correlations in the previous section are unweighted Pearson's correlations and do not always have the same sign as in the example shown in Figure 2. In contrast, the weighted CWM- and SNC-Pearson's correlations, with the abundance totals for plots and for species, respectively, as weights, always have the same signs and can be directly linked to the fourth-corner correlation. These weighted correlations are indeed proportional to the fourth-corner correlation (Peres-Neto et al., 2017).

For the CWM-based weighted correlation, the proportionality described above with the fourth-corner correlation does not depend on the environmental variable \mathbf{e} . In fact, a series of permutations of the environmental values generates a set of CWM-based weighted correlations and fourth-corner correlations that are proportional. Consequently, the plot-based permutation tests based on either test statistic (CWM-based weighted correlation or fourth-corner correlation) yields the same p -value. The same principle applies to the SNC-based weighted correlation. Its proportionality does not depend on the trait \mathbf{t} , so that the species-based permutation tests on the SNC-based weighted correlation generates the same p -value as the one based on the fourth-corner correlation. We verified these mathematical equivalences in the simulation study. Additional details are provided in Supporting information Appendix S2.

In this study, we investigate whether these permutation tests can be replaced by parametric tests based on weighted linear regression without loss of type I error control and power. In Figure 3,

these approaches are denoted by *wcor* CWM/SNC and *wlm* CWM/SNC, respectively.

In the ZS-modified test, the test statistic is the CWM-based correlation and its significance is assessed by a permutation test in which the indicator values (i.e. trait values) are permuted among species (Figure 3). Supporting information Appendix S3 provides additional details on the relationships between the CWM-based and the fourth-corner correlations, providing some related test statistics and testing methods.

2.3 | Set-up of the simulation study

We considered a variety of models that link species abundances to traits and environmental variables to assess and contrast the different permutation and parametric tests. Abundance data were generated by three simulation models: one-dimensional and two-dimensional Gaussian response models and a log-linear model, in which an interaction term between trait (**t**) and environmental variable (**e**) represents the trait–environment association. The two-dimensional Gaussian response model (with ellipsoidal contours with axes parallel to the gradients) is detailed in Supporting information Appendix S4; the other two models in Supporting information Appendix S1 of ter Braak et al. (2017) with one difference, namely that the observed trait and environmental values were drawn from a 50–50% mixture of two normal distributions with means at -1 and 1 and common variance 1 and then rescaled to mean 0 and variance 1 . The purpose of this was to evaluate whether the parametric approach also works in the case of non-normal variables.

In the Gaussian models, the strength of the trait–environment association is determined by the correlations of the observed trait and environmental variables with the true single gradient (ρ_t and ρ_e , respectively) or, in the two-dimensional case, the two gradients (with the first gradient: (ρ_{t1}, ρ_{e1}) and with the second gradient: (ρ_{t2}, ρ_{e2})). In addition, the tolerance (niche width) along each gradient and the response distribution for the data play a role, expressed as in the simulation model as the maximum of the true species tolerances (Peres-Neto et al., 2017). In Gaussian models, there is a true association between the observed trait and environmental variable only when both ρ_t and ρ_e are non-zero in the one-dimensional model and, in the two-dimensional model, when both ρ_{t1} and ρ_{e1} and/or both ρ_{t2} and ρ_{e2} are non-zero. All parameters (ρ_{t1}, ρ_{e1}) and the tolerances are varied using a grid of values so as to obtain a wide range of scenarios and associated powers.

In the log-linear model, the strength of the trait–environment association is determined by the regression coefficient associated with the interaction between the trait and environmental variable (b_{te}). The log-linear model contained additional interaction terms, notably the interaction between the observed trait with a latent (unobserved) environmental variable (**x**) and the interaction between the observed environmental variable with a latent trait (**z**), with regression coefficients b_{tx} and b_{ze} , respectively. These latent variables and the true gradients in the Gaussian models play a similar role of providing structured community data with and without true

links between the observed trait and environmental variable. All latent variables were independent standard normals. In the log-linear model, there is true association between the observed trait and environmental variable only when b_{te} is non-zero. All parameters (b_{te} , b_{tx} and b_{ze}) are varied using a grid of values so as to obtain a wide range of scenarios and associated estimations of statistical power.

Note that Gaussian models with equal niche breadths among species can be expressed as log-linear models, but the Gaussian models with unequal niche breadths, as in this simulation study, cannot. Also, not all log-linear models can be easily expressed as equitolerance Gaussian models and the parameters, such as tolerance in the Gaussian models and regression coefficients in the log-linear models, require different ranges to generate realistic abundance values. The models therefore complement each other and, without further information, none is necessarily more realistic.

2.4 | Statistical analysis

Each simulated data set was analysed using a number of methods to detect associations between **t** and **e**, without using any true gradient values or latent variable values. All methods consisted of a plot-based and a species-based significance test which were combined in a max test by taking the maximum of their *p*-values, with the exception of the ZS-modified test method, which uses a species-based significance test only. Significance of traits–environment associations was assessed by either permutation or parametric tests. The plot-based permutation test randomly reshuffled the environmental values in **e**. The species-based permutation test randomly reshuffled the trait values in **t**.

In the main article, two permutation methods and two parametric methods are compared (Figure 3). The two permutation tests are (starting with abbreviations between parentheses):

- (*fourth-corner*). The fourth-corner correlation (Dray & Legendre, 2008).
- (*ZS-modified test*). The CWM-based correlation (Pearson's correlation between CWMs and **e**) with species-based permutation only (Hawkins et al., 2017; Zelený & Schaffers, 2012).

The two parametric methods are:

- (*lm* CWM/SNC). The plot-level test used the significance level of the slope of the regression of **c**~**e** and the species-level test used the significance level of the slope of the regression of **u**~**t**.
- (*wlm* CWM/SNC). As in linear regression, but using weighted regression. The weights are the plot and species total abundance (l_{jt} and l_{je}) in the plot-level and species-level test, respectively.

In the methods based on simple linear regression, the significance level of the slope was obtained by the usual Student's *t*-test, with $n - 2$ and $S - 2$ *df* in the plot- and species-level tests, respectively. Equivalently, the significance levels of the (weighted) Pearson correlations between **c** and **e** and between **u** and **t** could have been used.

These approaches are thus the parametric versions of the *cor* CWM/SNC method and the *wcor* CWM/SNC method, respectively (Figure 3).

Type I error rates at the 5% level (false-positive rate in null models, i.e. models without true trait–environment association) of methods were determined on the basis of 10,000 independent replicates for the simulated data and 199 random permutations in each permutation test, unless stated otherwise. The power at this level (true-positive rate in models with true trait–environment association) was based on 5,000 replicates and 499 random permutations, except in the two-dimensional Gaussian model, where 56% of the runs was with 99 permutations to reduce computation time. More permutations slightly increase statistical power; the power difference between 99 and 499 permutations was at most 12% (median 6%) for powers >0.25. All permutation methods used identical random permutations of species and/or plots to reduce the effect of the number of permutations on power comparisons. Power comparisons were based on the odds ratio with respect to a reference method, as the odds ratio highlights power differences when the power is low but also when the power is high (low type II error rate). The *cor* CWM/SNC method was chosen as reference as it is a permutation method allowing more precise comparisons with the other permutation methods and most clearly showed our main findings. Four additional permutation tests (Figure 3) were also considered (Supporting information Appendix S5).

3 | RESULTS

Figure 4 shows the type I error rates of methods for the two-dimensional Gaussian model across four scenarios in which the trait and the environmental variable were not related to the same dimension (gradient), so that there is no true trait–environment association.

When neither trait nor environmental variable was related to the true gradients of the model, $(\rho_{e1}, \rho_{t2}) = (0,0)$, the plot- and species-level test methods all had about correct type I error rates (the

nominal level), except that the rate in the *wlm* CWM/SNC method is slightly above ($p \approx 0.061$). The rate of the max test is less than the nominal level for all methods, particularly in the *lm* CWM/SNC method. In this scenario, both the max test and the ZS-modified test controlled the type I error. When the environmental variable corresponded to the first gradient, but the trait was unrelated to both gradients, $(\rho_{e1}, \rho_{t2}) = (1,0)$, all plot-level tests had highly inflated error rates (> 0.30), whereas the other tests controlled the type I error rates, although the rate with the *wlm* CWM/SNC method was again slightly inflated. In the reverse case, when the trait corresponded to the second gradient, but the environmental variable was unrelated to both gradients, $(\rho_{e1}, \rho_{t2}) = (0,1)$, all species-level tests had highly inflated error rates (> 0.30), except for the ZS-modified test. All the other tests (plot-level and max) controlled for inflated type I error rates. When the environmental variable corresponded to the first gradient and the trait to the second gradient, $(\rho_{e1}, \rho_{t2}) = (1,1)$, all plot-level tests had highly inflated error rates (> 0.20), and the species-level tests had somewhat inflated error rates, except for the ZS-modified test, which was grossly conservative in this scenario. The max test using weighted regression had a slightly inflated p -value in this scenario ($p \approx 0.066$). In summary, the max tests and the ZS-modified test controlled the type I error (Figure 4), except when using weighted linear regression. With 100 instead of 50 species in the $(\rho_{e1}, \rho_{t2}) = (1,1)$ scenario, the max test gave an inflated type I error with the fourth-corner correlation ($p = 0.077$) and weighted linear regression ($p = 0.093$), but not with the other tests.

Figure 5 shows type I error rates of methods for the log-linear model without interaction between the observed trait and environmental variable. The four scenarios in Figure 5 were similar to the ones in Figure 4, but were parametrized differently. The most striking differences in relation to the results shown in Figure 4 are the gross type I error inflation in the weighted linear regression method. The fourth-corner and unweighted linear regression method had mildly inflated error of their max test ($p \approx 0.08$) in the (1,1) scenario

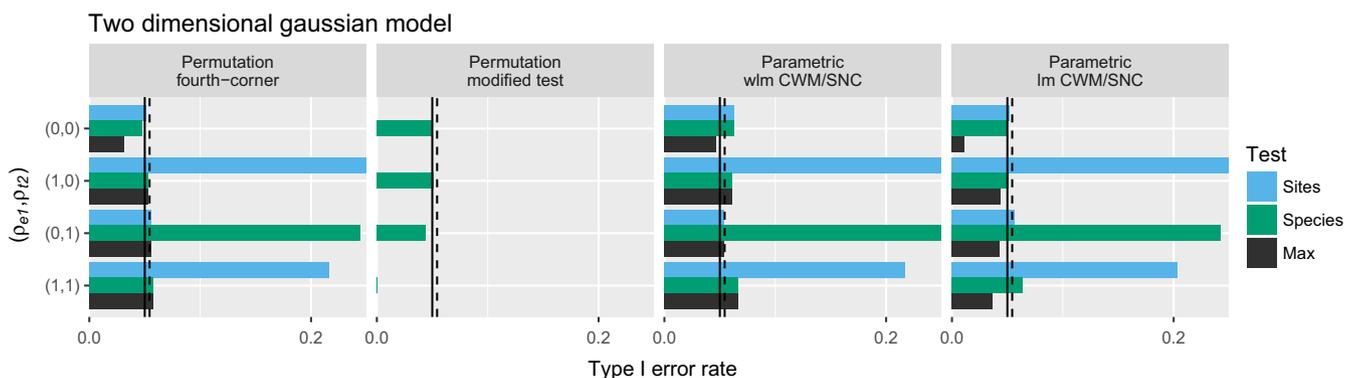


FIGURE 4 Type I error rates of permutation and t -tests on trait–environment association in simulated data from the two-dimensional Gaussian model without true trait–environment association (100 plots, 50 species and a maximum tolerance of two for both gradients) with $(\rho_{e2}, \rho_{t1}) = (0,0)$. The combinations for (ρ_{e1}, ρ_{t2}) are (0,0): both e and t random; (1,0): e corresponds to the first gradient, t random; (0,1): e random, t corresponds to the second gradient; (1,1): e corresponds to the first gradient, t corresponds to the second gradient. The max test uses the maximum of the p -values of the plots and species test in each data set. For details, see text. The solid vertical lines are at the nominal significance level ($p = 0.05$); error rates above the dotted line ($p = 0.054$) are significantly >0.05. Type I error rates >0.3 shown as 0.3 [Colour figure can be viewed at wileyonlinelibrary.com]

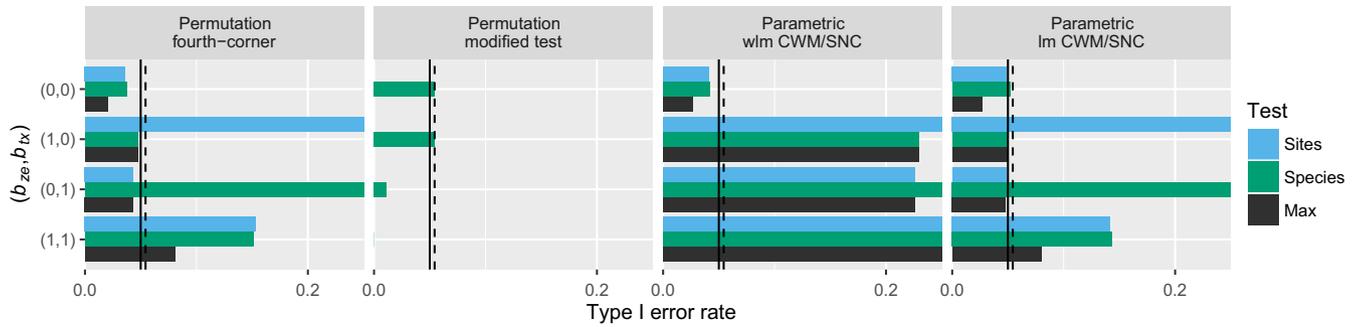
log-linear model with $b_{te} = 0$ 

FIGURE 5 Type I error rates of permutation and t -tests on trait-environment association in simulated data from the log-linear model without true trait-environment association (30 plots, 30 species, $b_{te} = 0$). The combinations for (b_{ze}, b_{tx}) are (0,0): no interaction of \mathbf{e} and \mathbf{t} with latent variables; (1,0): interaction of \mathbf{e} with a latent trait, no interactions involving \mathbf{t} ; (0,1): no interactions involving \mathbf{e} , interaction of \mathbf{t} with a latent environmental variable; (1,1): interaction of both \mathbf{e} and \mathbf{t} with a latent variable. The max test uses the maximum of the p -values of the plots and species test in each data set. For details, see text. The solid vertical lines are at the nominal significance level ($p = 0.05$); error rates above the dotted line ($p = 0.054$) are significantly >0.05 . Type I error rates >0.3 shown as 0.3 [Colour figure can be viewed at wileyonlinelibrary.com]

in which both trait and environmental variable interact with a latent variable but not with one another. In this scenario, the ZS-modified test was again grossly conservative.

On the basis of these results, the weighted linear regression method (*wlm CWM/SNC*) is not further investigated. Of the permutation methods, only the max test and the ZS-modified test are considered further.

Figure 6 compares the power of methods in the two-dimensional Gaussian model by way of the odds ratio with respect to the power of the *cor CWM/SNC* method, which is the unweighted method that uses permutation whereas (unweighted) linear regression uses parametric distributions. The linear regression method was indeed almost identical to the *cor CWM/SNC* method in Figure 3, confirming that the parametric test is a valid and efficient replacement for the corresponding permutation test. The power of the fourth-corner method and the ZS-modified test was higher for the lower powers ($p < 0.25$), but generally lower at higher powers where it matters ($p > 0.25$), particularly when the niche width was small (tolerances 0.5 and 2.0). The pattern for the log-linear model (Figure 7) shows large losses in statistical power for the ZS-modified test, and about equal power for the other methods. The power loss increased as a function of the noise due to additional latent variables governed by the parameters b_{ze} and b_{tx} in the model.

Similar figures with four more permutation methods are shown in Supporting information Appendix S5, where the results for the one-dimensional Gaussian model are also shown. The comparison of the panels for fourth-corner and the *wcor CWM/SNC* method (Figure 3) in these figures confirms that these permutation methods are exactly the same.

4 | DISCUSSION

This study highlights, once again, but with more elaborate models than previous studies (ter Braak, 2017; ter Braak et al., 2017;

Two dimensional gaussian model

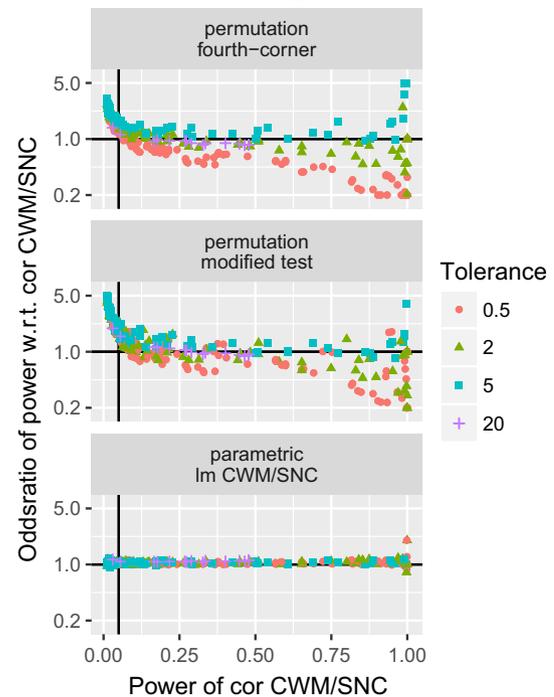


FIGURE 6 Odds ratio of the power of methods with respect to method *cor CWM/SNC* in simulated data against the power of *cor CWM/SNC* in the two-dimensional Gaussian model with true trait-environment association along the first gradient (100 plots, 50 species). Tolerance is the maximum tolerance along the first gradient. The solid vertical lines are at the nominal significance level ($p = 0.05$) [Colour figure can be viewed at wileyonlinelibrary.com]

Peres-Neto et al., 2017), that plot-based statistical tests (parametric or permutation-based) alone cannot validly detect trait-environment associations. A species-level test is also needed, which can then be combined into a max test by taking the maximum p -value of the plot-level test and the species-level test (ter Braak et al., 2012). In this study, the tests were based on the

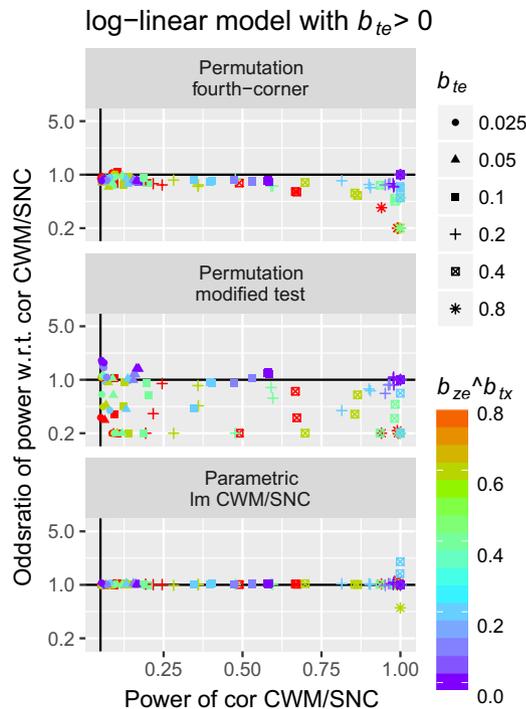


FIGURE 7 Odds ratio of the power of methods with respect to method *cor CWM/SNC* in simulated data against the power of *cor CWM/SNC* in the log-linear model with true trait–environment association (30 plots, 30 species, $b_{te} > 0$). $b_{ze} \wedge b_{tx}$ is the maximum of b_{ze} and b_{tx} . The solid vertical lines are at the nominal significance level ($p = 0.05$) [Colour figure can be viewed at wileyonlinelibrary.com]

fourth-corner correlation and a combination of CWM-based and SNC-based correlations. We showed, by theory and simulation (Supporting information Appendix S5), that the permutation test based on the weighted version of this combination is identical with that based on the fourth-corner correlation. In simulations, the unweighted version was shown to be slightly more powerful when the niche widths of species were small compared to the range of the environmental gradient. We tried to replace the permutation tests in both versions with parametric tests based on linear regressions (or the corresponding correlation coefficients) and found, rather unexpectedly, that the weighted version gave inflated type I error rates, whereas the unweighted version did control the type I error rates and performed very similarly to the unweighted version that used permutation (the *cor CWM/SNC* method). Two parametric tests, one of the CWM-based correlation and one of the SNC-based correlation (or of the corresponding linear regressions), can thus be combined into a max test to give a simple, easy and quick test of trait–environment association.

This result may be useful in itself when many tests need to be performed, but can also be the basis of extending the test to: (a) detect non-linear relations (replace the linear regressions by non-linear ones, e.g. spline regressions) and (b) allow for phylogeny and/or spatial autocorrelation (by including a covariance matrix in the linear regression, i.e. generalized least squares). The extensions deserve further investigation. It is, for example, an open question whether

the test in such an extension can still be parametric or whether it needs some kind of permutation or resampling (Lapointe & Garland, 2001; Wagner & Dray, 2015). The unexpected result of our simulations was that the weighted linear regression approach mimicking the fourth-corner generated inflated type I error rates. This result makes it more likely that any standard weighted least squares approach that attempts to correct for spatial correlation and/or phylogeny will require resampling rather than permutation. In Supporting information Appendix S6, we introduce a very simple model in which weighted linear regression gives inflated type I error rates, and also propose a heuristic modification, based on Hill's effective number of occurrences (Hill, 1973), to account for this. A linear mixed model for the trait–environment association that accounts for phylogeny was examined by Li and Ives (2017) and a related Bayesian model by Ovaskainen et al. (2017).

The max test and the ZS-modified test were shown to be conservative in different scenarios. The max test is conservative (by design) in the scenario where neither the trait nor the environmental variable was important, whereas the ZS-modified test is conservative when the environmental variable is random, and even more conservative when trait and environment variable were related to different gradients (Figures 4 and 5 and Supporting information Appendix S5). As a consequence, the ZS-modified test (Zelený & Schaffers, 2012) did not show type I error inflation in the latter scenario, whereas the max test showed moderate type I error inflation. The ZS-modified test lacks theoretical underpinning, but this empirical finding is a strong point in its favour. It comes at the cost of moderate-to-high loss in power (Figures 6 and 7, respectively), however.

ter Braak et al. (2017) detected moderate type I error rate inflation in the max test using GLM for testing the trait–environment association. In this study, we showed similar problems using the fourth-corner and CWM- and SNC-based methods in both the log-linear model and the two-dimensional Gaussian model. In the two-dimensional Gaussian model, the problematic scenario is when the trait is strongly related to one gradient and the environmental variable to a second gradient that is orthogonal to the first. In such a case, both trait and the environmental variable influence abundance but not via a common gradient. ter Braak et al. (2017) argued that trait–environment interaction is scale- and thus model-dependent. Arguably, there is a trait–environment interaction in this scenario when viewed on any other scale than the log-linear scale. Our simulation results showed that the ZS-modified test is not vulnerable to these problems and, of the other tests, that the unweighted tests were the least affected.

In this study, the trait and the environmental variable are extrinsic to (not at all derived from) the vegetation data. Without measurements on the environment, it is tempting to derive environmental variables by calculating mean Ellenberg indicator values or by more advanced procedure (cluster analysis, ordination). But note that it does not make sense to statistically test the relationship of such derived variables with community-weighted trait means. If the vegetation data are related to the trait, such a test very often leads to a correlation between the trait mean and the derived environmental

variable that is statistically significant, even when starting from random Ellenberg values, as Zelený and Schaffers (2012) and Wildi (2016) confirmed by simulation. The patterns of derived variables can be very intriguing (Hawkins et al., 2017) but are non-random by construction, so that statistical testing of the null hypothesis is not of interest (Wildi, 2016). When traits and environmental variables are measured, instead of being constructed from the abundance data, the null hypothesis of trait and environment being unrelated is of interest. This study showed once again that valid statistical testing of this null hypothesis is non-trivial but can be achieved using the max test based on both a plot-level and a species-level test. Two tests are needed because there are two sources of randomness as both plots and species are statistical units (ter Braak et al., 2017).

The fact that plots and species are two sources of randomness led Pollock, Morris, and Vesik (2012) and Jamil, Ozinga, Kleyer, and ter Braak (2013) to cast the problem of trait–environment analysis in terms of GLMM, in which both plots and species are treated as random factors. These advanced approaches can be used for pair-wise testing of single traits with single environmental variables, as considered in this study, but also allow multiple traits and environmental variables and selection of relevant traits and environmental variables. One question is whether the simple approaches of this study can be extended similarly. Bernhardt-Römermann et al. (2008) do so in the context of RLQ (Dolédec, Chessel, ter Braak, & Champely, 1996) and Peres-Neto et al. (2017) hint at extending single fourth-corner correlations to a multivariate regression approach. The logical endpoint of the latter is double-constrained correspondence analysis (Kleyer et al., 2012), which is the natural regression extension of the fourth-corner correlation (ter Braak, 2017; ter Braak, Šmilauer, & Dray, 2018) and which is implemented in Canoco 5.1 (ter Braak & Šmilauer, 2018).

So far in the literature, trait means per community or plot (i.e. CWMs) were considered as the sole summary of the trait variation across plot. In addition, the within-plot trait variance can be studied, e.g. by relating it to the environment. Similarly, the environmental means per species (i.e. SNCs) were considered as the only summary of the environmental variation across species, but in addition the within-species environmental variance (niche breadth) can be studied further, e.g. by relating it to traits (ter Braak et al., 2018; Supporting information Appendix S6). Further work is required to evaluate the performance of the procedures to test the link between trait and environment when variances are considered instead of means.

When multiple trait measurements are made per plot (instead of taking trait data from an existing database), there is also intra-specific trait variation in the data, which deserves further examination beyond the scope of this study. But for the question of trait–environment association, the methods of this study are still applicable and efficient. The reason for this is that, also with multiple trait measurements, the trait mean and trait variance values per plot are still natural numerical summaries that can be regressed on to the environmental variable in the plot-based analysis. In the species-based analysis, the environmental mean and variance per species

(i.e. the SNC and niche breadth) can then be regressed on the trait means of species. The mathematical underpinning of these claims is analogous to that in Supporting information Appendix S6 of ter Braak et al. (2018). Note that, in the case of multiple measurement, variance can be split into intra- and inter-species trait variance, a situation not covered by ter Braak et al. (2018) and requiring further study.

The conclusion of this study is that the combination of two simple regressions is a good alternative for the fourth-corner and the ZS-modified test when numerous tests need to be performed; their combination into a max test is powerful and fast; the resulting *p*-value is not bounded from below by the number of permutations, so that the *p*-value can be usefully adjusted for the number of tests when correcting for multiple testing.

ACKNOWLEDGEMENTS

We thank the Editor, David Zelený and another reviewer for useful comments.

ORCID

Cajo J. F. ter Braak  <http://orcid.org/0000-0002-0414-8745>

Stéphane Dray  <http://orcid.org/0000-0003-0153-1105>

Pedro R. Peres-Neto  <http://orcid.org/0000-0002-5629-8067>

REFERENCES

- Ackerly, D., Knight, C., Weiss, S., Barton, K., & Starmer, K. (2002). Leaf size, specific leaf area and microhabitat distribution of chaparral woody plants: Contrasting patterns in species level and community level analyses. *Oecologia*, 130, 449–457. <https://doi.org/10.1007/s004420100805>
- Anderson, M. J., & Robinson, J. (2001). Permutation tests for linear models. *Australian & New Zealand Journal of Statistics*, 43, 75–88. <https://doi.org/10.1111/1467-842X.00156>
- Anderson, M. J., & ter Braak, C. J. F. (2003). Permutation tests for multi-factorial analysis of variance. *Journal of Statistical Computation and Simulation*, 73, 85–113. <https://doi.org/10.1080/00949650215733>
- Bernhardt-Römermann, M., Römermann, C., Nuske, R., Parth, A., Klotz, S., Schmidt, W., & Stadler, J. (2008). On the identification of the most suitable traits for plant functional trait analyses. *Oikos*, 117, 1533–1541. <https://doi.org/10.1111/j.0030-1299.2008.16776.x>
- Choler, P. (2005). Consistent shifts in alpine plant traits along a mesotopographical gradient. *Arctic, Antarctic, and Alpine Research*, 37, 444–453. [https://doi.org/10.1657/1523-0430\(2005\)037\[0444:CSIAPT\]2.0.CO;2](https://doi.org/10.1657/1523-0430(2005)037[0444:CSIAPT]2.0.CO;2)
- Dolédec, S., Chessel, D., ter Braak, C. J. F., & Champely, S. (1996). Matching species traits to environmental variables: A new three-table ordination method. *Environmental and Ecological Statistics*, 3, 143–166. <https://doi.org/10.1007/BF02427859>
- Dray, S., Choler, P., Dolédec, S., Peres-Neto, P. R., Thuiller, W., Pavoine, S., & ter Braak, C. J. F. (2014). Combining the fourth-corner and the RLQ methods for assessing trait responses to environmental variation. *Ecology*, 95, 14–21. <https://doi.org/10.1890/13-0196.1>
- Dray, S., & Dufour, A. B. (2007). The ade4 package: Implementing the duality diagram for ecologists. *Journal of Statistical Software*, 22, 1–20. <https://doi.org/10.18637/jss.v022.i04>



- Dray, S., & Legendre, P. (2008). Testing the species traits–environment relationships: The fourth-corner problem revisited. *Ecology*, 89, 3400–3412. <https://doi.org/10.1890/08-0349.1>
- Hawkins, B. A., Leroy, B., Rodríguez, M. Á., Singer, A., Vilela, B., Villalobos, F., ... Zelený, D. (2017). Structural bias in aggregated species-level variables driven by repeated species co-occurrences: A pervasive problem in community and assemblage data. *Journal of Biogeography*, 44, 1199–1211. <https://doi.org/10.1111/jbi.12953>
- Hill, M. O. (1973). Diversity and evenness: A unifying notation and its consequences. *Ecology*, 54, 427–432. <https://doi.org/10.2307/1934352>
- Jamil, T., Ozinga, W. A., Kleyer, M., & ter Braak, C. J. F. (2013). Selecting traits that explain species–environment relationships: A generalized linear mixed model approach. *Journal of Vegetation Science*, 24, 988–1000. <https://doi.org/10.1111/j.1654-1103.2012.12036.x>
- Kleyer, M., Dray, S., Bello, F., Lepš, J., Pakeman, R. J., Strauss, B., ... Lavorel, S. (2012). Assessing species and community functional responses to environmental gradients: Which multivariate methods? *Journal of Vegetation Science*, 23, 805–821. <https://doi.org/10.1111/j.1654-1103.2012.01402.x>
- Lapointe, F.-J., & Garland, J. T. (2001). A generalized permutation model for the analysis of cross-species data. *Journal of Classification*, 18, 109–127. <https://doi.org/10.1007/s00357-001-0007-0>
- Legendre, P., Galzin, R. G., & Harmelin-Vivien, M. L. (1997). Relating behavior to habitat: Solutions to the fourth-corner problem. *Ecology*, 78, 547–562. <https://doi.org/10.2307/2266029>
- Li, D., & Ives, A. R. (2017). The statistical need to include phylogeny in trait-based analyses of community composition. *Methods in Ecology and Evolution*, 8, 1192–1199. <https://doi.org/10.1111/2041-210X.12767>
- McGill, B. J., Enquist, B. J., Weiher, E., & Westoby, M. (2006). Rebuilding community ecology from functional traits. *Trends in Ecology & Evolution*, 21, 178–185. <https://doi.org/10.1016/j.tree.2006.02.002>
- Ovaskainen, O., Tikhonov, G., Norberg, A., Guillaume Blanchet, F., Duan, L., Dunson, D., ... Abrego, N. (2017). How to make more out of community data? A conceptual framework and its implementation as models and software. *Ecology Letters*, 20, 561–576. <https://doi.org/10.1111/ele.12757>
- Peres-Neto, P. R., Dray, S., & ter Braak, C. J. F. (2017). Linking trait variation to the environment: Critical issues with community-weighted mean correlation resolved by the fourth-corner approach. *Ecography*, 40, 806–816. <https://doi.org/10.1111/ecog.02302>
- Pollock, L. J., Morris, W. K., & Vesik, P. A. (2012). The role of functional traits in species distributions revealed through a hierarchical model. *Ecography*, 35, 716–725. <https://doi.org/10.1111/j.1600-0587.2011.07085.x>
- Šmilauer, P., & Lepš, J. (2014). *Multivariate analysis of ecological data using CANOCO 5* (2nd ed.). Cambridge, UK: Cambridge University Press. <https://doi.org/10.1017/cbo9781139627061>
- ter Braak, C. J. F. (2017). Fourth-corner correlation is a score test statistic in a log-linear trait–environment model that is useful in permutation testing. *Environmental and Ecological Statistics*, 24, 219–242. <https://doi.org/10.1007/s10651-017-0368-0>
- ter Braak, C. J. F., Cormont, A., & Dray, S. (2012). Improved testing of species traits–environment relationships in the fourth-corner problem. *Ecology*, 93, 1525–1526. <https://doi.org/10.1890/12-0126.1>
- ter Braak, C. J. F., & Looman, C. W. N. (1986). Weighted averaging, logistic regression and the Gaussian response model. *Vegetatio*, 65, 3–11. <https://doi.org/10.1007/BF00032121>
- ter Braak, C. J. F., Peres-Neto, P., & Dray, S. (2017). A critical issue in model-based inference for studying trait-based community assembly and a solution. *PeerJ*, 5, e2885. <https://doi.org/10.7717/peerj.2885>
- ter Braak, C. J. F., & Šmilauer, P. (2018). *Canoco reference manual and user's guide: Software for ordination (version 5.1)*. Ithaca, NY: Microcomputer Power.
- ter Braak, C. J. F., Šmilauer, P., & Dray, S. (2018). Algorithms and biplots for double constrained correspondence analysis. *Environmental and Ecological Statistics*, 25, 171–197. <https://doi.org/10.1007/s10651-017-0395-x>
- Verhoeven, K. J. F., Simonsen, K. L., & McIntyre, L. (2005). Implementing false discovery rate control: Increasing your power (vol 108, pg 643, 2005). *Oikos*, 109, 208. <https://doi.org/10.1111/j.0030-1299.2005.13727.x>
- Wagner, H. H., & Dray, S. (2015). Generating spatially constrained null models for irregularly spaced data using Moran spectral randomization methods. *Methods in Ecology and Evolution*, 6, 1169–1178. <https://doi.org/10.1111/2041-210X.12407>
- Warton, D. I., Shipley, B., & Hastie, T. (2015). CATS regression – a model-based approach to studying trait-based community assembly. *Methods in Ecology and Evolution*, 6, 389–398. <https://doi.org/10.1111/2041-210X.12280>
- Wildi, O. (2016). Why mean indicator values are not biased. *Journal of Vegetation Science*, 27, 40–49. <https://doi.org/10.1111/jvs.12336>
- Zelený, D. (2016). Bias in community-weighted mean analysis of plant functional traits and species indicator values. *bioRxiv*. <https://doi.org/10.1101/046946>
- Zelený, D., & Schaffers, A. P. (2012). Too good to be true: Pitfalls of using mean Ellenberg indicator values in vegetation analyses. *Journal of Vegetation Science*, 23, 419–431. <https://doi.org/10.1111/j.1654-1103.2011.01366.x>

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Appendix S1 Instability of CWM-based correlation. **Table S1** Spreadsheet example of the CWM and SNC approach applied to the artificial data of Appendix S1

Appendix S2 Relation of fourth-corner and CWM/SNC-based correlations

Appendix S3 The ZS-modified test compared with CWM-based correlation and regression

Appendix S4 The two-dimensional Gaussian response model

Appendix S5 Additional results on type I error and power

Appendix S6 Weighted linear regression and type I error control

Appendix S7 Zip file with the R code of the simulations and the snowmelt example.

How to cite this article: ter Braak CJF, Peres-Neto PR, Dray S. Simple parametric tests for trait–environment association. *J Veg Sci*. 2018;29:801–811. <https://doi.org/10.1111/jvs.12666>