APPLICATION

PIECEWISESEM: Piecewise structural equation modelling in R for ecology, evolution, and systematics

Jonathan S. Lefcheck*

Department of Biological Sciences, Virginia Institute of Marine Science, The College of William & Mary, Gloucester Point, VA 23062-1346, USA

Summary

1. Ecologists and evolutionary biologists rely on an increasingly sophisticated set of statistical tools to describe complex natural systems. One such tool that has gained significant traction in the biological sciences is structural equation models (SEM), a form of path analysis that resolves complex multivariate relationships among a suite of interrelated variables.

2. Evaluation of SEMs has historically relied on covariances among variables, rather than the values of the data points themselves. While this approach permits a wide variety of model forms, it limits the incorporation of detailed specifications. Recent developments have allowed for the simultaneous implementation of non-normal distributions, random effects and different correlation structures using local estimation, but this process is not yet automated and consequently, evaluation can be prohibitive with complex models.

3. Here, I present a fully documented, open-source package PIECEWISESEM, a practical implementation of confirmatory path analysis for the R programming language. The package extends this method to all current (generalized) linear, (phylogenetic) least-square, and mixed effects models, relying on familiar R syntax. I also provide two worked examples: one involving random effects and temporal autocorrelation, and a second involving phylogenetically independent contrasts.

4. My goal is to provide a user-friendly and tractable implementation of SEM that also reflects the ecological and methodological processes generating data.

Key-words: confirmatory path analysis, graph theory, mixed models, networks

"No aphorism is more frequently repeated in connection with field trials, than that we must ask Nature few questions, or, ideally, one question, at a time. The writer is convinced that this view is wholly mistaken. Nature, he suggests, will best respond to a logical and carefully thought out questionnaire; indeed, if we ask her a single question, she will often refuse to answer until some other topic has been discussed."

—Sir Ronald Fisher (1926)

Introduction

The desire to understand the intricate complexity of nature is arguably the single driving force behind all of science. Yet, for the last century or so, ecologists and evolutionary biologists have closely examined the impact of one or few factors on a single response. This practice was, and sometimes still is, a consequence of limited computational power, and the necessity of simplification in rigorous experimentation. However, with the advent of modern computing and the tractability of large-scale observation, there is an increasing recognition that multifaceted data sets representing complex natural systems require an equally sophisticated toolbox. Structural equation models (SEM) provide one such tool.

Structural equation models are probabilistic models that unite multiple predictor and response variables in a single causal network. They are often represented using path diagrams, where arrows indicate directional relationships between observed variables (Figs 1 and 2). These relationships can be captured in a series of structured equations that correspond to the pathways in the model. Two primary characteristics of SEMs separate them from more traditional modelling approaches:

1. *Paths represent hypothesized causal relationships*. This is a departure from the phrase, 'correlation does not imply causation.' In fact, correlation does imply causation, but the direction of causality is unresolved, since one cannot know whether, for instance, A causes B, B causes A, or both A and B are a consequence of some third, unmeasured variable (Shipley 2000b). By using pre-existing knowledge of the system gained through observation and/or experimentation, however, one can make an informed hypothesis about the causal structure of A, B and other variables that are thought to mediate their relationship. SEM allows for the direct test of this supposed causal structure. In this way, SEM is a departure from tradi-

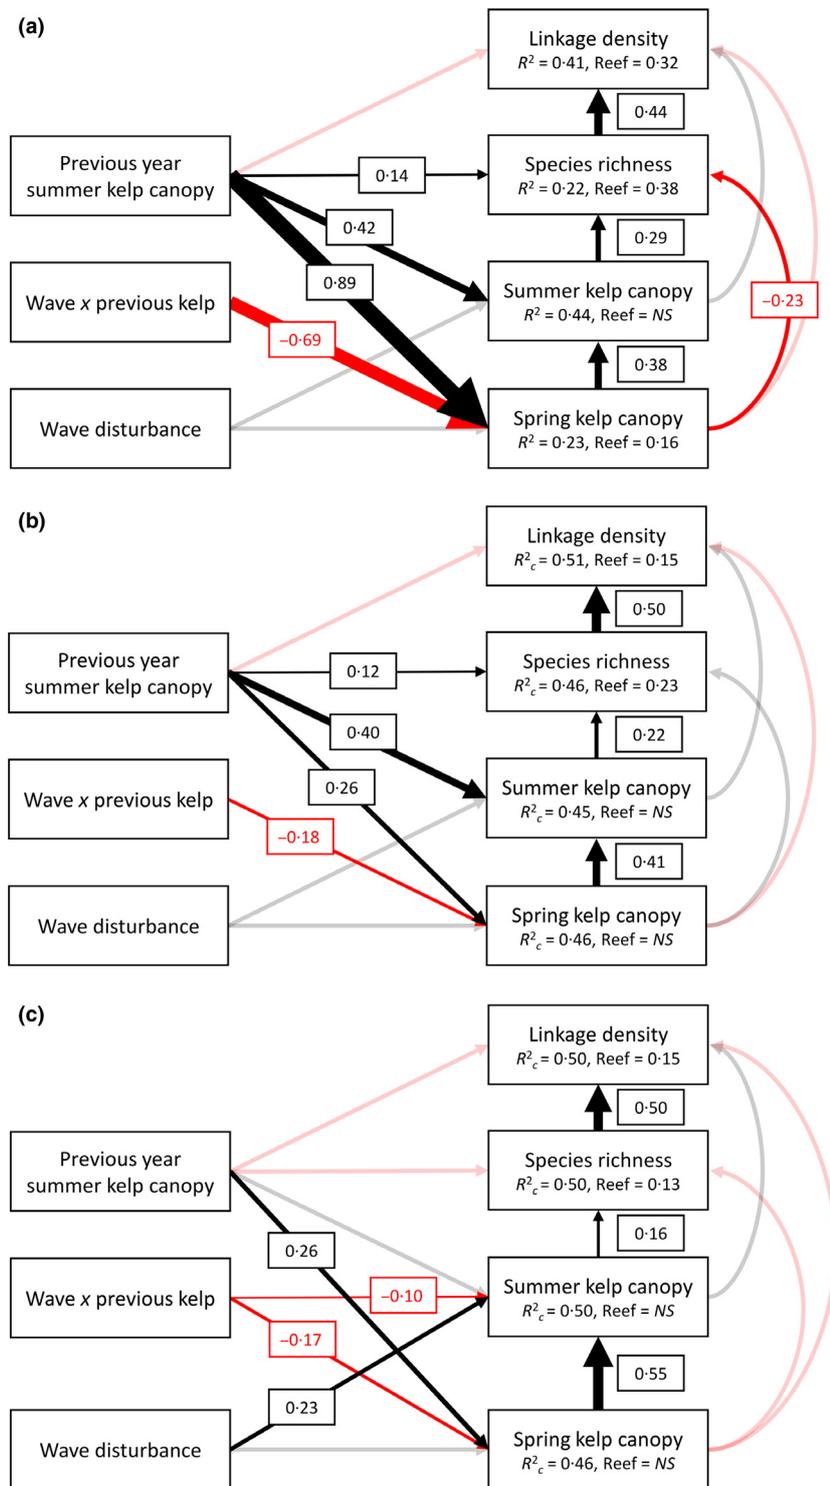*Correspondence author. E-mail: jslefche@vims.edu

(a)

(b)

(c)

Fig. 1. Structural equation models (SEM) from Byrnes *et al.* (2011) exploring the effects of storm frequency (wave disturbance) on kelp forest community structure and food web complexity (linkage density). Boxes represent measured variables. Arrows represent unidirectional relationships among variables. Black arrows denote positive relationships, and red arrows negatives ones. Arrows for non-significant paths ($P \geq 0.05$) are semi-transparent. The thickness of the significant paths has been scaled based on the magnitude of the standardized regression coefficient, given in the associated box. R^2s for component models are given in the boxes of response variables (for panels b and c, this is reported as the conditional R^2_c based on the variance of both the fixed and random effects). The variable 'Reef habitat' has been omitted for clarity and the path coefficient is instead reported in the corresponding box of the response, as in Byrnes *et al.* (2011) (NS = not significant). (a) Original analysis using variance–covariance SEM. (b) The same model in a fitted using piecewise SEM and incorporating a random effect of *Site*. (c) The piecewise model from panel b, with an additional autocorrelation term for *Year*.

tional linear modelling by explicitly testing the hypothesis that A *causes* B. While an in-depth treatment of this idea is beyond the scope of this paper, further discussion of causality and how it relates to SEM can be found in: Pearl (2012) and Bollen & Pearl (2013).

2. *Variables can appear as both predictors and responses.* By allowing one variable to serve as a response in one path and as a predictor in another, SEM is useful for testing and quantifying indirect or cascading effects that

would otherwise go unrecognized by any single model (e.g. Grace *et al.* 2007).

Structural equation models demands a shift in how ecological and evolutionary questions are structured and tested, with an emphasis on the simultaneous evaluation of multiple causal hypotheses within a single network.

Historically, SEMs have been estimated using a maximum-likelihood approach to select parameter values that best reproduce the entirety of the observed variance–covariance
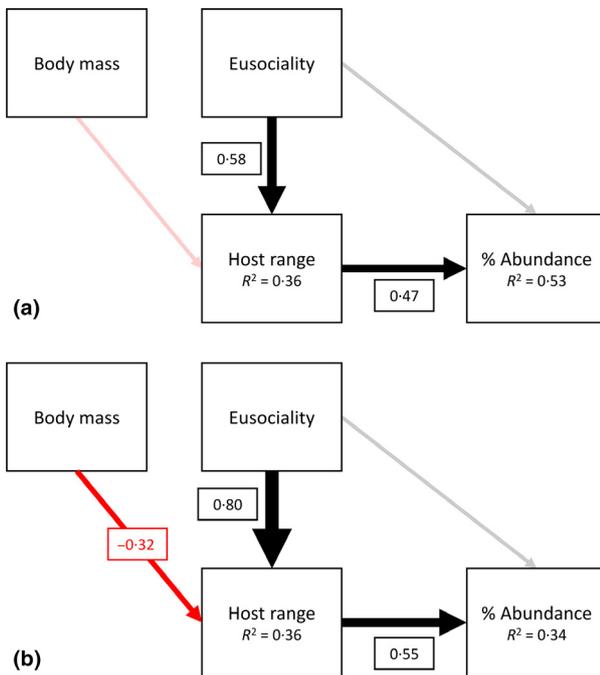
Fig. 2. Structural equation models (SEM) derived from hypotheses in Duffy & Macdonald (2010) exploring the relationships among eusociality, body size, host range size, and proportional regional abundance for eusocial *Synalpheus* shrimps. Arrows represent unidirectional relationships among variables. Black arrows denote positive relationships, and red arrows negatives ones. Arrows for non-significant paths ($P \geq 0.05$) are semi-transparent. The thickness of the significant paths has been scaled based on the magnitude of the standardized regression coefficient, given in the associated box. R^2s for component models are given in the boxes of response variables. (a) Analysis using variance–covariance SEM. (b) The same model as in a fit using piecewise SEM, and additionally incorporating a fixed correlation structure based on phylogenetic distances obtained from a molecular phylogeny.

matrix. The goodness-of-fit of the SEM can then be evaluated using a chi-square test comparing the estimated to the observed covariance matrix (Grace 2006). This approach, however, assumes that all observations are independent, and all variables follow a (multivariate) normal distribution (Grace 2006). It also restricts the minimum number of observations necessary to fit the SEM, since there needs to be sufficient degrees of freedom to estimate the whole variance–covariance matrix (the '*t* rule,' Grace 2006).

These restrictions led to the parallel development of directed acyclic, or piecewise, SEMs based on applications from graph theory. In piecewise SEM, the path diagram is translated to a set of linear (structured) equations, which are then evaluated individually. The switch from *global estimation*, where equations are solved simultaneously, to *local estim*ation, where each equation is solved separately, allows for the fitting of a wide range of distributions and sampling designs (Shipley 2000a, 2009). It also, in theory, permits the fitting of smaller data sets, since there only need be enough degrees of freedom to fit any given component model (Shipley 2000a) (but see Discussion: Limitations). Finally, it can incorporate distances obtained from taxonomy or phylogeny to address the potentially

confounding effects of shared evolutionary history (Von Hardenberg & Gonzalez-Voyer 2013). Because piecewise SEM does not yet incorporate latent or composite variables, it is often and more correctly referred to as confirmatory path analysis. I will, however, continue to refer to it broadly as SEM *sensu* Grace *et al.* (2012), who include local estimation under their definition of 'third-generation SEM.'

Since piecewise SEM produces no valid global covariance matrix, alternative goodness-of-fit tests are necessary. The typical approach uses Shipley's test of *directed separation*. This procedure tests the assumption that all variables are conditionally independent. In simplest terms, conditional independence implies that there are no missing relationships among unconnected variables (Shipley 2000a). The first step in the test of direct separation is to derive the minimum set of conditional independence claims associated with the hypothesized path diagram, known as the *basis set*. The basis set can be translated into a set of linear equations, each of which can be solved like any other linear model. The significance of any given independence claim, that is its *P*-value, can be estimated and extracted. The test of directed separation is conducted by combining all *P*-values across the basis set in a test statistic, Fisher's *C*, using the following equation:

$$C = -2 \sum_{i=1}^{k} \ln(p_i) \qquad \text{eqn 1}$$

where P_i is the *i*th independence claim in a basis set consisting of *k* claims. *C* can then be compared to a chi-square distribution with $2k$ degrees of freedom. The hypothesized relationships are considered to be consistent with the data when there is weak support for the sum of the conditional independence claims, that is where the collection of such relationships represented by *C* could have easily occurred by chance, in which case *P* for the chi-square test is greater than the chosen significance threshold (typically $\alpha = 0.05$). Several approachable examples of the derivation of basis sets can be found in Shipley (2000a, 2009).

Shipley (2013) showed that the Fisher's *C* statistic can be used to obtain a value of Akaike's information criterion (AIC) using the following equation:

$$\text{AIC} = C + 2K \qquad \text{eqn 2}$$

where *C* is from eqn (1), and *K* is the likelihood degrees of freedom (not to be confused with *k*, the number of independence claims in the basis set). Because this estimator is not derived from maximum likelihood, it is sometimes referred to as the *C* statistic information criterion (CIC, *sensu* Cardon *et al.* 2011). Equation (2) can also be extended to small sample sizes (AIC$_c$), typically when the number of parameters exceeds the total sample size $n/40$, with an additional correction: AIC$_c = C + 2K(n/n - K - 1)$.

The implementation of piecewise SEM is limited by the correct specification and evaluation of the basis set, which can be prohibitive to obtain by hand, especially for very complex models. To that end, I provide a fully documented and open-source package PIECEWISESEM (https://github.com/jslefche/piecewiseSEM) for the R statistical

language to aid in the calculation of piecewise SEM by constructing the basis set, conducting goodness-of-fit tests for both the full and component models, calculating AIC scores, returning (scaled) parameter estimates, plotting partial correlations, and generating predictions. SEMs are built using a list of structured equations, which can be specified using most common linear modelling functions in R, and thus can accommodate non-normal distributions, hierarchical structures and different estimation procedures. In this paper, I present two worked examples: the first incorporating mixed effects models and temporally correlated observations, and the second involving non-independence via phylogenetically independent contrasts. The data and R code to reproduce all analyses are given in the Supporting information.

Example 1: Storm frequency and kelp forest food webs

In this first example, I use data from Byrnes *et al.* (2011), who examined the role of storm events on the diversity and food web structure of kelp forests in California, USA. They combined biological surveys of kelp forests over 35 different sites and 8 years, potential food web linkages derived from the literature, data on wave height and period from physical monitoring stations, and kelp canopy cover from satellite imagery. They summarized these variables in a single causal network derived using *a priori* knowledge of the system and results from experimental manipulation (Fig. 1). They then evaluated this model using traditional variance–covariance SEM.

Byrnes *et al.* hypothesized that the wave disturbance generated by winter storms would be contingent on the amount of existing kelp, which interactively affect the spring canopy cover. Spring canopy cover would in turn inform summer canopy cover, which is also subject to physical forcing. The amount of canopy cover, spring or summer, would provide structural habitat for various species, such as algae, sessile invertebrates and their consumers. Total species richness would finally determine the number of potential trophic links in the observed food web (linkage density, or the mean number of feeding links per observed species).

The results of their original analysis are reproduced in Fig. 1a using the LAVAAN package (Rosseel 2012). The model was an adequate fit to the data based on output from a chi-square goodness-of-fit test ($\chi_5^2 = 8.784$, $P = 0.118$). Byrnes *et al.* (2011) saw that spring canopy cover was strongly influenced by the interaction between wave disturbance and previous kelp cover: as the previous year's cover increased, the effect of wave disturbance on the current spring's canopy cover became more negative. Spring canopy cover had both a direct negative effect on species richness, and an indirect positive effect mediated through summer kelp cover. Species richness in turn enhanced food web complexity. However, they noted that the direct negative effect of spring canopy cover on species richness had a larger magnitude (standardized β = −0.23) than the indirect effect, which is obtained by multiplying the path coefficients ($0.38 \times 0.29 = 0.11$). Thus, they concluded that the

removal of spring canopy by winter storms actually increased species richness (by reducing the stronger direct negative effect), ultimately increasing food web complexity in the short term. However, given the effect of losing kelp, total species richness should decline if reefs experienced multiple years of wave disturbance in a row.

Their analysis, however, treated each observation as independent. In reality, sites that are proximate are likely to share similar characteristics, and within a site, observations closer in time are likely to be more similar than those that are farther apart. To address both of these concerns, I re-fit their original model using piecewise SEM. In the first re-analysis, I addressed the non-independence of sampling sites by fitting each response to a general linear mixed effects models using the function lme from the NLME package (Pinheiro *et al.* 2013). I chose to log-transform the variables as in Byrnes *et al.* (2011) instead of fitting integer responses to a Poisson distribution in order to facilitate direct comparisons to the original analysis, although this is possible using PIECEWISESEM. For each component model, I fit a random effect of *Site* and allowed only its intercept to vary. I then added the component models to a list and passed the list to the function sem.fit, which returns the tests of directed separation, Fisher's *C* statistic and AIC values for the SEM. I then recovered the standardized regression coefficients (scaled by mean and variance, as in Byrnes *et al.*) using the sem.coefs function.

The piecewise SEM based on mixed models reproduced the data equally well as the output from LAVAAN, based on comparison of the Fisher's *C* statistic to a chi-square distribution ($C_{10} = 15.64$, $P = 0.11$). The results from this re-analysis are given in Fig. 1b. In general, the models explained a larger proportion of variance on average than the traditional SEM, based on R^2 values derived from the variance of both fixed and random effects (Nakagawa & Schielzeth 2012), obtained using the function sem.model.fits.

There are several major differences between the models in Fig. 1a,b. First, the magnitudes of the main effect of the previous year's kelp canopy cover and the interaction between this variable and wave disturbance were both reduced by about two-thirds, although they retained the same signs. Most consequential for the original interpretation is that the negative relationship between spring canopy cover and species richness was non-significant. By nesting observations based on their hierarchical structure, variation that was formerly assumed to be generated by canopy cover was reallocated to random (spatial) variation. Thus, based on the output from the piecewise SEM, wave disturbance both directly and indirectly reduces spring canopy cover, which indirectly *reduces* food web complexity as a consequence of cascading positive relationships between spring and summer canopy cover, summer canopy cover and species richness, and finally species richness and linkage density.

In the second re-analysis, I addressed both the non-independence of sites and any potential temporal autocorrelation by retaining the same random structure as above, and additionally modelling the correlation among sampling years using a

continuous autoregressive 1 autocorrelation structure from the CAR1 function from the NLME package (Pinheiro *et al.* 2013). This analysis also reproduced the data well based on comparison of the Fisher's *C* statistic to a chi-square distribution ($C_8 = 7.84$, $P = 0.45$). The results from this re-analysis are given in Fig. 1c.

There is slightly greater amount of variance explained for each component model vs. the piecewise SEM without the autocorrelation structure. There are, however, fewer notable differences between the two piecewise models. The path between spring canopy cover and species richness is still non-significant. There is now a significant positive path between summer canopy cover and wave disturbance, and the formerly significant path between the previous year's canopy cover and species richness is now non-significant. However, comparison of the two piecewise SEMs using AIC reveals that the model additionally incorporating the CAR1 autocorrelation structure is considerably less likely model than the one with only the hierarchical random structure ($AIC_c = 97.69$ *cf*. 81.44).

In sum, this re-analysis has revealed that modelling the hierarchical structure of the data leads to a different interpretation of the original data: wave disturbance *decreases* food web complexity, principally by removing habitat. This interpretation, however, supports the overall conclusions of Byrnes *et al.* (2011) that (repeated) storm events (i.e. wave disturbance) should decrease food web complexity, although I show this effect is mediated through the removal of habitat upon the first occurrence of disturbance and not necessarily a decrease in species richness after repeated disturbance events, as suggested by Byrnes *et al.* (2011). Additionally, AIC model comparisons revealed that modelling potential temporal autocorrelation does not add to our ability to understand this system of interactions.

Further exploration of the models from Byrnes *et al.* (2011) decomposing total species richness into trophic components revealed that canopy cover significantly reduced algal but not sessile invertebrate or mobile consumer species richness, as in their original analysis (see supplementary code). Modelling the random effect of *Site* likely absorbed some of the variation in algae-rich vs. algae-poor sites, making it more difficult to see the algae richness contribution to total species richness in the simpler piecewise model (Fig. 1b,c). This additional analysis confirms that the deeper exploration by Byrnes *et al.* (2011) was warranted to reconcile the statistical output with the biology of the system.

Example 2: Eusociality and ecological success in sponge-dwelling shrimp

In this second example, I use population and ecological data from a genus of sponge-dwelling shrimps (*Synalpheus*) to explore the drivers of ecological success. Species in this genus exhibit a range of social structures, from pair-forming to truly eusocial, with a single reproducing female per colony. It has been hypothesized that complex social structures like those exhibited by certain *Synalpheus* species are ecologically advan-

tageous in fostering greater competitive ability and/or resource acquisition. To answer this question, Duffy & Macdonald (2010) collated data on female body mass, number of host species used (host range), and proportional regional abundance for 20 species of *Synalpheus* in Belize. They additionally calculated an index of eusociality for each species. They hypothesized that more eusocial species (i.e. larger colonies with a single breeding female) would occupy a wider range of hosts, which would lead to greater success in defending those hosts (i.e. achieve higher relative abundance in the study area). They additionally hypothesized that the effect of host range might be confounded by body size, since most eusocial species are small-bodied.

As a first pass, I fit a traditional SEM using the sem function from the LAVAAN package (Rosseel 2012), assuming independence among all 20 data points (species). The model reproduced the data well ($\chi_1^2 = 0.653$, $P = 0.419$), and the results are given in Fig. 2a. There are two significant paths of interest: a strong positive effect of eusociality on host range accounting for body mass (standardized $\beta = 0.58$), and a positive effect of host range and relative abundance (0.47). There was not, however, a significant direct relationship between eusociality and abundance. Thus, it appears that the success of eusocial species is largely a consequence of their ability to occupy a wide range of hosts. Because of this generalist habitat use, they then also make up a larger percentage of total regional abundance, but the model does not support the hypothesis that eusociality confers a direct advantage in defending and holding onto a particular habitat resource.

Of course, Duffy & Macdonald (2010) correctly point out that the data points are not independent because some species are more related than others. To address this issue, I re-fit the SEM in Fig. 1a but additionally fixed the model correlation matrix based on genetic distances derived from a phylogeny of *Synalpheus* in the region (Hultgren & Duffy 2012). I obtained the model correlations from the phylogenetic tree using the function `corBrownian` from the APE package (Paradis, Claude & Strimmer 2004), and fit the component models using the function `gls` from the NLME package (Pinheiro *et al.* 2013). I stored the component models in a list and then evaluated the SEM using sem.fit. As before, the model reproduced the data well ($C_8 = 0.57$, $P = 0.751$), and the results are given in Fig. 2b.

The striking difference between the two SEMs in Fig. 2 is that the phylogenetic SEM recovers a significant negative effect of body mass on host range (-0.32), supporting the expectation that body size has a confounding influence. Even in the presence of a body size effect, there is a significant positive effect of eusociality on host range (indeed it is substantially stronger: 0.80). As with the previous SEM (Fig. 1a), there was no direct effect of eusociality on proportional regional abundance. Again, this relationship was mediated through an increase in host range. Repeating this analysis using the function `pgls` from the CAPER package (Orme *et al.* 2013), which estimates an additional scaling parameter λ, yielded nearly identical results (see supplementary code).

In their original paper, Duffy & Macdonald (2010) used multiple linear regression to explore relationships among these four variables. In their analysis, they showed that eusociality had a strong positive relationship with both relative abundance and host range size, after accounting for differences in body size and shared evolutionary history. Here, in a re-analysis of their data using SEM, I show the relationship between eusociality and relative abundance is not direct, but rather an indirect consequence of occupying a wider number of hosts, an insight that was simply not possible to infer from the individual multiple regressions. The extension of phylogenetic methods to SEM facilitates the testing of more complex, multivariate hypotheses in evolutionary ecology and, as shown here, can yield substantial additional insight.

Discussion

In this paper, I briefly introduce the concepts behind piecewise SEM, and apply piecewise SEM to two existing analyses. In both cases, acknowledging the non-independence of data points by incorporating random variation or phylogenetic distances yielded substantially different inferences than multiple regression or even traditional variance–covariance SEM. I also demonstrate how a new R package, PIECEWISESEM, can be used to quickly and easily implement complex local estimation. Indeed, this package has already been used to explore the planetary drivers of ecosystem functioning in eelgrass beds (Duffy *et al.* 2015), disentangle the influence of functional diversity across trophic levels in experimental estuarine mesocosms (Lefcheck & Duffy 2015) and quantify the biotic and abiotic drivers of grassland multifunctionality (Jing *et al.* 2015).

BROADER APPLICATIONS

The piecewise SEM package contains a number of additional functions that may be of general interest to users. `sem.model.fits,` for example, generates R^2, pseudo-R^2 and AIC values for component models based on methods in Nakagawa & Schielzeth (2012) and Johnson (2014). `sem.predict` is a wrapper for the generic `predict` function, and additionally implements standard errors on predictions from models constructed using `lme(r)` based on the variance of the fixed effects only http://glmm.wikidot.com/faq), although these approaches are controversial because they do not take into account the uncertainty of the random effects into account, and thus, estimates of error on predictions from mixed models should be interpreted with caution. `partial.resid` returns the partial correlation plot between two variables in a single model having accounted for the effects of covariates, and is an intuitive way to visualize the partial effects returned from `sem.coefs` or, more generally, `summary`. Exploration of partial correlations also allows for the identification of previously unrecognized nonlinear relationships, which can then be incorporated into the model structure.

LIMITATIONS

While it has been suggested that piecewise SEM can be used to circumvent restrictions on sample size (Shipley 2000a), it is important to note that small sample sizes may still have severe consequences for the analysis. In particular, tests of directed separation may substantiate a 'good fitting model' only because the tests lacked sufficient power to reject the null (i.e. *P*-values for missing paths are all >0·05). This outcome would be increasingly common as models increase in complexity, but not replication. Ideally, investigators should devise the hypothesized model beforehand and use it to inform data collection, ensuring sufficient replication from the start. As a general rule, Grace, Scheiner & Schoolmaster (2015) propose that the ratio of the total number of samples to the number of variables (*d*) should not fall below *d* = 5. It is also critical to examine the fits of the component models: if the overall SEM has an adequate fit but the component models have low explanatory power, then it is not acceptable (or particularly useful) to draw inferences from the SEM. Finally, users may find themselves with the opposite problem, where large sample size drives statistical but not biological significance, leading to rejection of the basis set on the basis of biologically inconsequential effect sizes. In this case, implementing a more stringent cut-off for statistical significance may alleviate the issue.

It is also worth noting that *P*-values derived from the LMERTEST package (Kuznetsova, Brockhoff & Christensen 2013) are somewhat unstable at the time of writing, and can often lead to errors in the `sem.fit` function. Estimates from NLME appear to be more reliable, and I recommend users should construct their models using NLME when LMERTEST produces an error, assuming the response is normally distributed.

While the piecewise SEM approach represents a considerable leap forward in addressing the assumptions of real-world data, its infancy relative to traditional SEM has led to some limitations. For instance, there is no real implementation of correlated errors: relationships that are bidirectional and assumed to be caused by a shared underlying driver. PIECEWISESEM implements a crude approximation of correlated errors by allowing the user to exclude them from the basis set (since there is no presumed direction of causality), and then running a simple test of significance on the bivariate correlation; however, other methods have been proposed (Shipley 2003), and may be incorporated in future iterations. Piecewise SEM also cannot disentangle cyclic relationships (e.g. A → B → C → A), making it impossible to evaluate feedbacks (Shipley 2009). Similarly, this method cannot evaluate reciprocal relationships in the same model (A → B and B → A, not to be confused with a bidirectional arrow indicating a correlated error). Finally, there is no formal integration of latent variables – those that are not directly measured, but inferred through a combination of observed variables (Grace 2006) – into piecewise SEM as of yet. It would be possible to derive predictions approximating a latent variable using exploratory factor analysis, or through the application of MCMC

estimation. However, there has yet to be a thorough investigation and application of factor analysis to piecewise SEM. With luck, future developments will relax some of these limitations.

Data accessibility

All data are available in the Supporting Information. The latest version of PIECEWISE SEM can be found on CRAN (https://cran.r-project.org/web/packages/piecewiseSEM/) and on GitHub (https://github.com/jslefche/piecewiseSEM.)

References

Bollen, K.A. & Pearl, J. (2013) Eight myths about causality and structural equation models. *Handbook of Causal Analysis for Social Research* (ed. S.L. Morgan), pp. 301–328. Springer, Ithaca, NY.

Byrnes, J.E., Reed, D.C., Cardinale, B.J., Cavanaugh, K.C., Holbrook, S.J. & Schmitt, R.J. (2011) Climate-driven increases in storm frequency simplify kelp forest food webs. *Global Change Biology*, **17**, 2513–2524.

Cardon, M., Loot, G., Grenouillet, G. & Blanchet, S. (2011) Host characteristics and environmental factors differentially drive the burden and pathogenicity of an ectoparasite: a multilevel causal analysis. *Journal of Animal Ecology*, **80**, 657–667.

Duffy, J.E. & Macdonald, K.S. (2010) Kin structure, ecology and the evolution of social organization in shrimp: a comparative analysis. *Proceedings of the Royal Society B*, **277**, 575–584.

Duffy, J.E., Reynolds, P.L., Boström, C., Coyer, J.A., Cusson, M., Donadi, S. *et al.* (2015) Biodiversity mediates top-down control in eelgrass ecosystems: a global comparative-experimental approach. *Ecology Letters*, **18**, 696–705.

Grace, J.B. (2006) *Structural Equation Modeling and Natural Systems.* Cambridge University Press, New York, NY.

Grace, J.B., Scheiner, S.M. & Schoolmaster, D.R. (2015) Structural equation modeling: building and evaluating causal models. *Ecological Statistics: From Principles to Applications* (eds G.A. Fox, S. Negrete-Yanlelevich & V.J. Sosa), pp. 168–199. Oxford University Press, New York, NY.

Grace, J.B., Michael Anderson, T., Smith, M.D., Seabloom, E., Andelman, S.J., Meche, G. *et al.* (2007) Does species diversity limit productivity in natural grassland communities? *Ecology Letters*, **10**, 680–689.

Grace, J.B., Schoolmaster, D.R. Jr, Guntenspergen, G.R., Little, A.M., Mitchell, B.R., Miller, K.M. & Schweiger, E.W. (2012) Guidelines for a graph-theoretic implementation of structural equation modeling. *Ecosphere*, **3**, 1–44.

Hultgren, K.M. & Duffy, J.E. (2012) Phylogenetic community ecology and the role of social dominance in sponge-dwelling shrimp. *Ecology Letters*, **15**, 704–713.

Jing, X., Sanders, N.J., Shi, Y., Chu, H., Classen, A.T., Zhao, K. *et al.* (2015) The links between ecosystem multifunctionality and above- and belowground biodiversity are mediated by climate. *Nature Communications*, **6**, 8159.

Johnson, P.C.D. (2014) Extension of Nakagawa & Schielzeth's R2 GLMM to random slopes models. *Methods in Ecology and Evolution*, **5**, 944–946.

Kuznetsova, A., Brockhoff, B. & Christensen, H.B. (2013) lmerTest: Tests for random and fixed effects for linear mixed effect models (lmer objects of the lme4 package).

Lefcheck, J.S. & Duffy, J.E. (2015) Multitrophic functional diversity predicts ecosystem functioning in experimental assemblages of estuarine consumers. *Ecology*, **96**, 2973–2983.

Nakagawa, S. & Schielzeth, H. (2012) A general and simple method for obtaining R2 from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, **4**, 133–142.

Orme, D., Freckleton, R., Thomas, G., Petzoldt, T., Fritz, S., Isaac, N. & Pearse, W. (2013) caper: Comparative Analyses of Phylogenetics and Evolution in R.

Paradis, E., Claude, J. & Strimmer, K. (2004) APE: analyses of phylogenetics and evolution in R language. *Bioinformatics*, **20**, 289–290.

Pearl, J. (2012) The Causal Foundations of Structural Equation Modeling. *Handbook of Structural Equation Modeling* (ed R.H. Hoyle), pp. 68–91. Guilford Press, New York, NY.

Pinheiro, J., Bates, D., DebRoy, S., Sarkar, D. & R Core Team. (2013) nlme: Linear and Nonlinear Mixed Effects Models.

Rosseel, Y. (2012) lavaan: an R package for structural equation modeling. *Journal of Statistical Software*, **48**, 1–36.

Shipley, B. (2000a) A new inferential test for path models based on directed acyclic graphs. *Structural Equation Modeling*, **7**, 206–218.

Shipley, B. (2000b) *Cause and Correlation in Biology*. Cambridge University Press, Cambridge, UK.

Shipley, B. (2003) Structural equation modeling: a confirmatory analysis of computer self-efficacy. *Structural Equation Modelling*, **10**, 214–221.

Shipley, B. (2009) Confirmatory path analysis in a generalized multilevel context. *Ecology*, **90**, 363–368.

Shipley, B. (2013) The AIC model selection method applied to path analytic models compared using a d-separation tests. *Ecology*, **94**, 560–564.

Von Hardenberg, A. & Gonzalez-Voyer, A. (2013) Disentangling evolutionary cause–effect relationships with phylogenetic confirmatory path analysis. *Evolution*, **67**, 378–387.

Supporting Information

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

Appendix S1. R code necessary to reproduce all examples in the text (.R).

Appendix S2. Kelp forest food web data from Byrnes et al. (2011) (.CSV).

Appendix S3. Synalpheus phylogeny from Hultgren & Duffy (2012) (.TXT).

Appendix S4. Synalpheus abundance data from Duffy & Macdonald (2010) (.CSV).