

Continuous characters analyzed as such

Pablo A. Goloboff^{1,*}, Camilo I. Mattoni² and Andrés Sebastián Quinteros³

¹Consejo Nacional de Investigaciones Científicas y Técnicas, Instituto Miguel Lillo, Miguel Lillo 205, 4000 S. M. de Tucumán, Argentina; ²Diversidad Animal I, F.C.E.F.y N., Universidad Nacional de Córdoba, Avenida Velez Sarsfield 299, 5000 Córdoba, Argentina; ³Cátedra de Anatomía Comparada, Facultad de Ciencias Naturales, Universidad Nacional de Salta, Avenida, Bolivia 5150, 4400 Salta, Argentina

Accepted 2 May 2005

Abstract

Quantitative and continuous characters have rarely been included in cladistic analyses of morphological data; when included, they have always been discretized, using a variety of ad hoc methods. As continuous characters are typically additive, they can be optimized with well known algorithms, so that with a proper implementation they could be easily analyzed without discretization. The program TNT has recently incorporated algorithms for analysis of continuous characters. One of the problems that has been pointed out with existing methods for discretization is that they can attribute different states to terminals that do not differ significantly—or vice versa. With the implementation in TNT, this problem is diminished (or avoided entirely) by simply assigning to each terminal a range that goes from the mean minus one (or two) SE to the mean plus one (or two) SE; given normal distributions, terminals that do not overlap thus differ significantly (more significantly if using more than 1 SE). Three real data sets (for scorpions, spiders and lizards) comprising both discrete and quantitative characters are analyzed to study the performance of continuous characters. One of the matrices has a reduced number of continuous characters, and thus continuous characters analyzed by themselves produce only poorly resolved trees; the support for many of the groups supported by the discrete characters alone, however, is increased when the continuous characters are added to the analysis. The other two matrices have larger numbers of continuous characters, so that the results of separate analyses for the discrete and the continuous characters can be more meaningfully compared. In both cases, the continuous characters (analyzed alone) result in trees that are relatively similar to the trees produced by the discrete characters alone. These results suggest that continuous characters carry indeed phylogenetic information, and that (if they have been observed) there is no real reason to exclude them from the analysis.

© The Willi Hennig Society 2006.

Most computer programs for parsimony analysis only accept discrete characters as input. The treatment of discrete characters in computer programs is simpler, because the state sets can be easily stored as sets of bits. This technique was used in the first efficient program for parsimony analysis on personal computers, Hennig86 (Farris, 1988). Other subsequently released programs (PAUP, Swofford, 1993; Nona, Goloboff, 1993c; PAUP*, Swofford, 1998) also used this technique. In the absence of alternative implementations in

tree-searching programs, users interested in the analysis of quantitative characters have proposed different methods for converting the continuous variation inherent to those characters into a series of discrete states that could be accepted as input by standard parsimony programs. Some programs (such as MacClade, Maddison and Maddison, 1992, and Mesquite, Maddison and Maddison, 2005) are capable of optimizing characters on to given trees as continuous, but these programs take single values (instead of ranges) and cannot find optimal trees.

Probably the most widely used method for discretization is Thiele's (1993) method, which divides the observed range R (the difference between maximum and minimum observed values) into as many states, N , as

*Corresponding author:

E-mail address: pablogolo@csnat.unt.edu.ar

allowed by the program to be used. Thus, a taxon with an observed value V will be coded as having discrete state S (with $0 \leq S \leq N$) if:

$$V \geq SR/N \text{ and } V < (S+1)R/N$$

The discretization can be done either by hand, using software developed by Schols et al. (2004), or with simple scripts for TNT (see Appendix 1). Other methods for discretization have been proposed by Mickevich and Johnson (1976), Archie (1985), and Goldman (1988). Farris (1990) offered criticisms of discretization methods that apply equally to all those methods (including Thiele's method, published after Farris' paper): taxa with significantly different values may be assigned to the same state, and/or taxa with non-significant differences may be assigned to different states. Farris (1990) proposed to test for significant differences between the taxa, but then it may be hard to consider the results of such tests in a discretization scheme. As exemplified by Rae (1998), "taxa A and B could be indistinguishable from one another in central tendency, but only taxon A is significantly different from a third taxon C", which "can lead to situations in which two taxa that are not significantly different from one another are assigned different codes.

The problems just discussed result only from discretization—i.e., dividing a continuous series in an arbitrary set of states. Some methods to estimate ancestral states, like squared-changes reconstructions (Rogers, 1984; Maddison, 1991), deal with continuous characters as such, but they spread out changes evenly among branches (Maddison, 1991; see Hormiga et al., 2000; Fig. 3, for an example). However, continuous or meristic characters are best treated as additive characters, and thus nothing precludes application of Farris's (1970) algorithms for optimization, other than the absence of implementation. Farris (1970) defined his algorithm in terms of intervals (that is, differences in the numerical values of the variables being optimized) and therefore no modification whatsoever is needed for application of those algorithms. This method is implemented in MacClade and Mesquite, but these programs only allow for single values in continuous characters, and using suitably chosen ranges for the terminals is necessary to avoid the problem of the significance in differences. When continuous characters are seen as simply additive characters, ranges can be naturally expressed as polymorphisms, analyzed in the usual way. Wiens (2001) proposed to use step-matrices to express the differences between the (mean) states in the terminals, which would produce results identical to Farris optimization using unique states, except that his method (as recognized by Wiens himself) becomes impractical when there is more than a handful of taxa, and is impractical for the treatment of taxa with ranges instead of single values (which, in the case of continuous characters, is all or

most taxa)¹. Wiens's (2001) method can be used with existing software for analysis of step-matrix characters, but a proper implementation of additive optimization is simpler, faster, and more general. These algorithms therefore have been implemented in the program TNT (Goloboff et al., 2003b), allowing for values between 0 and 65 (inclusive), using up to three decimal values. In bifurcating trees, TNT uses Farris's (1970) algorithms for the down-pass optimization, and Goloboff's (1993b) algorithms for the up-pass. Although not directly needed for tree-searches, most parsimonious optimization of polytomous trees presents special difficulties for continuously valued variables; TNT treats those with new algorithms, described below.

When continuous characters are treated as such, rather than distorted through discretization, they seem to carry indeed useful phylogenetic information. This is exemplified with analysis of three separate matrices, containing both continuous and discrete characters, in which there is significant congruence in the results for both types of characters analyzed separately.

Optimization of polytomies

Multifurcations pose a special problem for most parsimonious optimization of continuous characters. In the discussion that follows, an interval is defined as the range of values between a maximum and a minimum value; as defined by Farris (1970), the distance between two intervals is taken to be zero if the intervals overlap, or as the difference between the maximum value of the interval with the lowest maximum and the minimum value of the interval with the highest maximum. In the case of discrete characters, optimization of a multifurcation proceeds by trying each possible assignment to the node being (down or up) optimized; this requires trying a limited number of possible states (32 at the most, for 32-bit implementations). When the character is continuous, there is an infinite number of possible state assignments to each node, and enumerating them becomes impossible (even when TNT uses only 65 001 possible values, enumerating them becomes time consuming). Because of this difficulty, programs that can perform parsimonious optimization of continuous characters on given trees (such as MacClade and Mesquite) handle only strictly bifurcating trees in the case of continuous characters. The optimal state (range) assignments can be found easily, however, taking into account

¹Actually, and despite the fact that polymorphic step-matrix characters could in principle be used to accommodate ranges, Wiens (2001) explicitly recommended assigning to every terminal a single value, corresponding to its mean value, as is done in Thiele's method. Such a recoding is open to Farris's (1990) criticism regarding significant differences—a point that Wiens (2001) never discussed.

a property of additive optimization: when deciding range assignments in the down (or up) pass for a given node *N*, based on the preliminary ranges of the descendants (and the final states of the ancestor) of the node, the limits for the preliminary (or final) interval of *N* will always be taken from a set of possible values, *P*, corresponding to one of the limits present in at least one of the descendants (or the ancestor) of *N*. Then, the number of possible state assignments that have to be tried for each node is greatly reduced; at the most, when finding final assignments, if all the *d* descendants and the ancestor of the node have intervals with different limits, the cardinality of set *P*—that is, the number of values that have to be tried—is only $2(d + 1)$. As in any type of optimization (see Goloboff, 1998, for the case of Sankoff optimization) the up-pass to a node *N* must consider each of the possible state assignments for the ancestor; thus, when ancestor is assigned an optimal value *A*, the cost of assigning state *S* to *N* equals $|A - S|$, plus the sum of the distances from *S* to the range in each of the descendants. All the states in *N* that produce the minimum cost, given *A*, must be included in the final interval for *N*. Considering the distance between the interval in the ancestor, to each possible state *S*, may lead to missing some possible assignments for *N*. In this context, the possible values that have to be tried for an ancestor having minimum and maximum possible values v_{\min} and v_{\max} is not all the possible values between v_{\min} and v_{\max} , but instead all of the values in the set *P* that are between those two values (Fig. 1 is an example).

When optimization of polytomies is done in this way, the time for finding down-pass assignments for a node leading to *d* descendants increases at the most with *d*, and the cost of finding up-pass assignments increases at the most with d^2 . That time dependency is less than the one seen in step-matrix characters (Sankoff and Rousseau, 1975; where both the down-pass and the up-pass for *s* states depends on s^2), particularly as multifurcating nodes also imply that fewer nodes have to be optimized.

Significance

Farris (1990) pointed out that a problem with previous methods for recoding is that it can assign different states to terminals that do not differ significantly, or the same state to terminals that are significantly different. Methods that assign a unique value to terminals (like Thiele's or Wiens' methods), are especially prone to this problem. If continuous characters are treated just like regular additive characters, the problem can be diminished (or entirely avoided) by assigning to the terminals suitably chosen ranges.

The choice of ranges is up to the user. In most cases, a good approximation will be obtained by assigning each

terminal a range going from its mean $- 1$ standard error of the mean (SE) to the mean $+ 1$ SE (the 95% confidence interval of the mean could be used instead of SEM; see Zar, 1974). This will produce ranges that are smaller than the observed ones, but such that the ranges of two terminal taxa will overlap (thus producing a step count of 0) whenever their means are not statistically different.

Scaling

One of the most pervasive problems in the analysis of continuous characters (discretized or not) is scaling. It seems clear that, within a given character, a change from a condition observed in one species to a (significantly different) condition observed in another species, should be proportional to the magnitude of the differences, but the problem arises for the costs of transformations in different characters. Needless to say, different scalings often yield different trees. Although this problem also affects discrete characters, its relevance is even more obvious in the case of continuous characters.

There seems to be no obvious solution for this problem. In a sense, the question itself is problematic; it assumes that the exact cost ratio of state transformations in two different characters can somehow be exactly determined. It is not simply that there is a value of cost ratio that is the true one but hard to know—the very idea that such a value exists is hard to justify. That question is meaningful only within the context of methods justified probabilistically—methods in which the relative costs of transformations are a function of their hypothesized probabilities. Parsimony is not intended to be such a type of method, and therefore the very question of whether a wing loss is exactly as reliable as a change in leg shape is, in a sense, irrelevant. Therefore, a final and general answer to the problem of scaling is unlikely to ever be found.

A way to decrease the problems of scaling is provided by using implied weights (Goloboff, 1993a). If the fit for character *i* is measured as $k/h_i + k$ (where h_i = homoplasy for character *i*, and k = concavity constant²), then h_i will tend to be greater when the character is measured on a larger scale. Therefore, characters with smaller ranges will tend to have higher implied weights, and vice versa. Consider as example two contradictory characters, *A* and *B*, measured on a different scale; comparable transformations in characters *A* and *B* cost 10 and 1 steps, respectively. Adding the first homoplasious step to character *A* (with a cost of 10) decreases tree fit much more than adding the first homoplasious step to

² h_i is calculated, as usual, as $h_i = s_i - m_i$, where m_i = minimum possible number of steps for character *i*, and s_i = number of steps for character *i*.

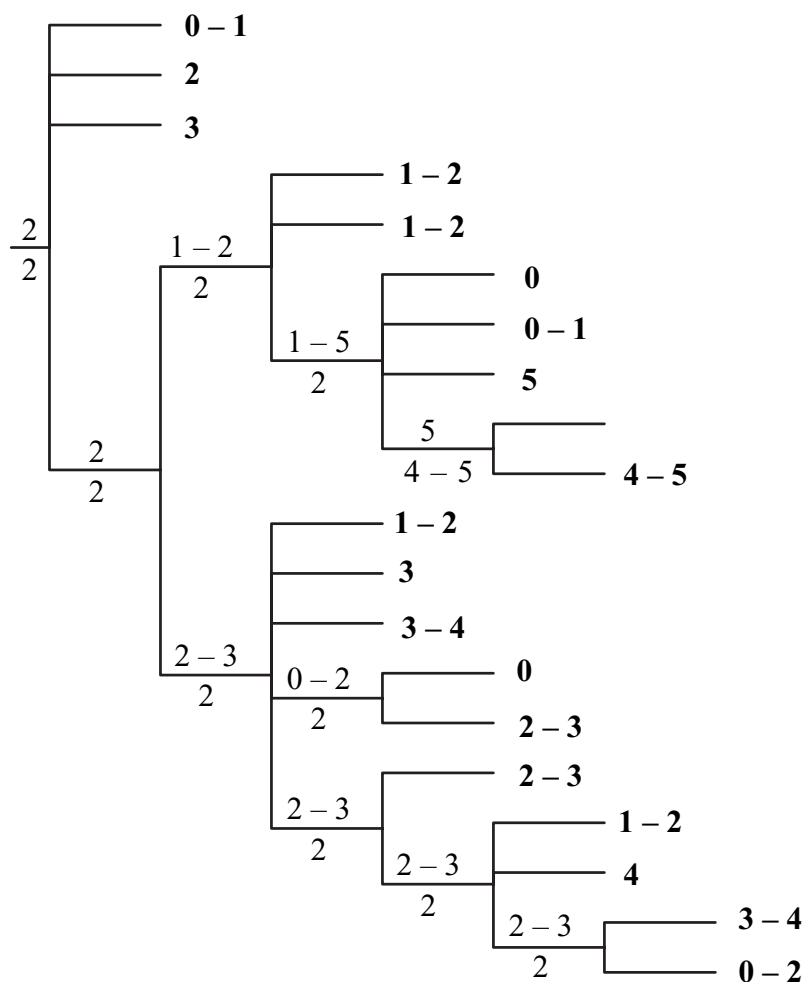


Fig. 1. Example showing how to simplify calculation of most parsimonious state assignments for continuous variables. The ranges indicated above branches correspond to the first pass of the optimization (the ranges resulting from the up-pass are indicated below branches). The limits for the preliminary (down-pass) assignment for polytomous node C are taken from the set of candidates { 0, 1, 5 }. The down-pass cost of 0 is 10, of 1 is 9, and of 5 is 9; all the values intermediate between the values that produce a minimum local cost (1 and 5) also produce a minimum local cost. There is no need to calculate explicitly the down-pass cost of assigning states 2, 3 or 4 to node C, as they do not occur in the limits of any of the descendants. For the other polytomous nodes, the down-pass calculations can also be simplified by selecting candidate values from reduced subsets (for E from { 1, 2, 3, 4 }, for D from { 0, 1, 2, 3, 4 }, for B from { 1, 2, 5 }, and for A from { 0, 1, 2, 3 }). The limits for the final (up-pass) assignments are calculated similarly, but adding to the set of candidate values the limits of the final states for the ancestor of each node, and trying as alternative states for the ancestor only those states included in the final range for the ancestor, which occur in the set of possible candidates (in the example, the range in the ancestor for each of the polytomous nodes is included within the preliminary range of the node itself, so the equivalent of rule 1 of Fitch, 1971 applies).

character B (with a cost of 1). Under $k = 3$, the first homoplasious step to character A decreases fit in $3/3-3/13$ (going from $h = 0$ to $h = 10$), or 0.769. The first homoplasious step to character B decreases fit in $3/3-3/4$, or 0.250. The implied weight ratio (about 3 : 1) is then much less than the ratio implied by the difference in scale (10 : 1). Further, when adding the second step of homoplasy this situation is inverted, because then the decrease in fit for character A (0.100) is less than the decrease in fit for character B (0.15). This effect may possibly balance the overall influence of the different characters, making their contributions more even.

Examples

Continuous characters as such have never been used as evidence in parsimony analysis. Many authors (e.g., Pimentel and Riggins, 1987; Cranston and Humphries, 1988) have even questioned whether the use of such characters is desirable at all. The problem of whether continuous characters provide useful phylogenetic information is, in part, methodological: as they are the result of heritable variation, the principle of considering all the relevant evidence simply forces us to consider the degree to which a phylogeny provides a good explanation of

continuous characters. But the problem is also in part empirical; the phenotypic variation in continuous characters might be so large as to confound the heritable variation. Whether or not this is the case, can be analyzed by considering whether phylogenies produced by continuous characters are congruent with those produced by discrete characters. This was examined in three morphological data sets, comprising both discrete and continuous data, corresponding to the genus *Bothriurus* (Scorpiones, Bothriuridae), the genus *Phymaturus* (Squamata, Liolaemidae), and the subfamily Ischnothelinae (Araneae, Dipluridae). All the cases were analyzed under implied weighting (with $k = 15$).

Scorpions

The scorpion data set comprises 50 taxa (43 species of *Bothriurus* or synonyms, and seven outgroup taxa), 103 discrete characters, and seven continuous characters. With only seven continuous characters, it is unlikely that the continuous data by themselves will allow a reliable estimation of phylogeny. Unsurprisingly, the strict consensus of the discrete characters alone, and the strict consensus of the continuous characters alone, have no groups in common. The difference is due to more than just the placement of a few taxa, because no possible pruning of a few taxa from the consensus improves the resolution of the grand consensus. A method often used to evaluate concordance between trees (e.g., Eulenstein et al., 2004) is the number of taxa in the agreement subtree; the agreement subtree of the two consensus (Fig. 2) has only 12 taxa. Identical relationships for 12 of 50 taxa are obtained under chance alone with relative ease, but (just like an unresolved strict consensus) agreement subtrees of few taxa may easily occur even if the input trees are very similar. Figure 3 shows an example, where the agreement subtree has only 17 of 65

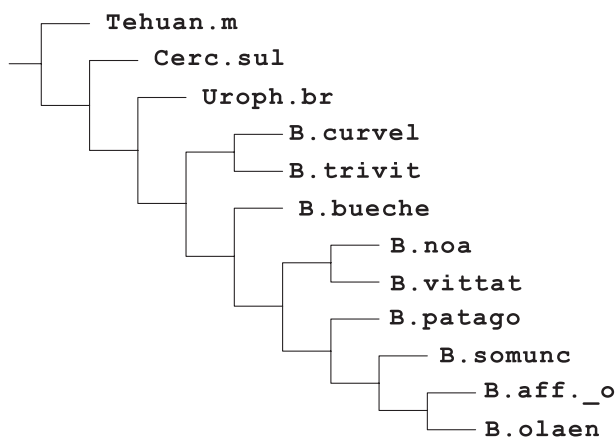


Fig. 2. Agreement subtree for the trees resulting from analysis of discrete characters alone, and continuous characters alone, for the *Bothriurus* data set.

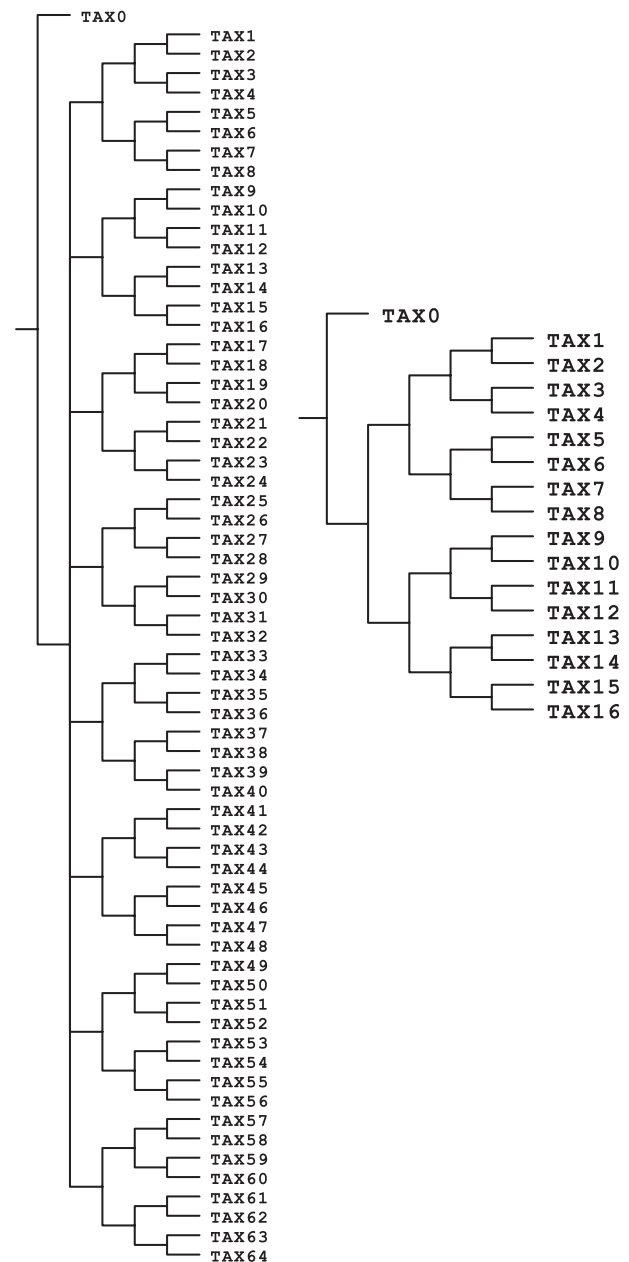


Fig. 3. Example showing that the agreement subtree may have few taxa even when the input trees are relatively similar. The input trees are binary, and have different resolutions for the relationships of the 8-taxon clades. Most of the nodes in the input trees (56 of 63, 88%) are shared, but the agreement subtree displays only 16 taxa (25%).

taxa, yet the input trees are rather similar (as evidenced by the strict consensus with nothing less than 56 nodes). In other words, agreement subtrees with many taxa indicate similarity, but agreement subtrees with few taxa need not indicate difference.

Another measure of tree similarity can be obtained by calculating the (minimum) number of SPR moves necessary to convert a tree onto the other; this often produces

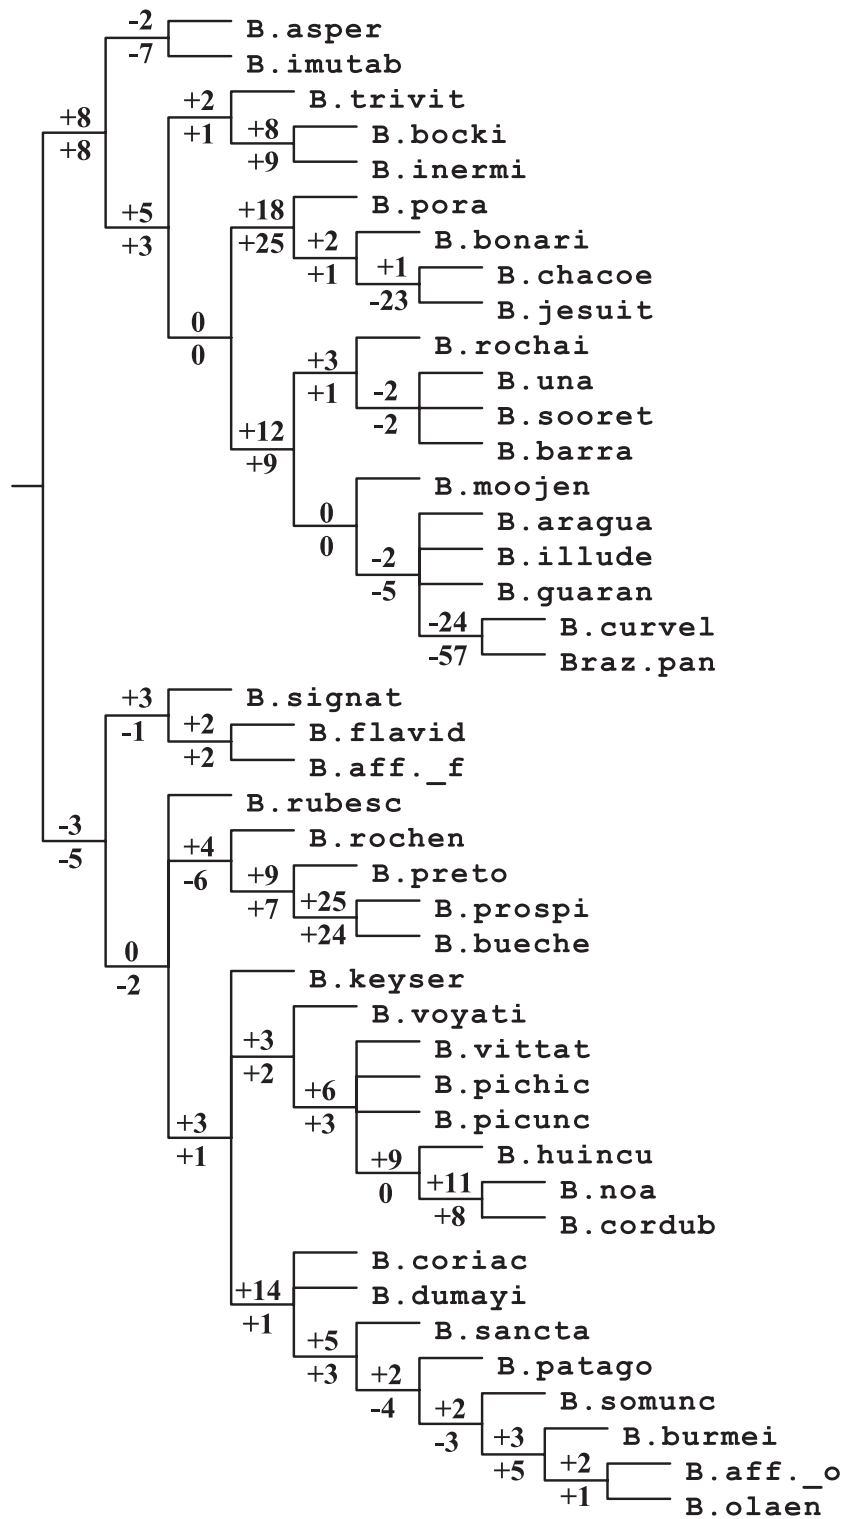


Fig. 4. Differences in group support, in the *Bothriurus* part of the tree, for the groups supported by the discrete-only data set, when the continuous characters are added. The differences correspond to supports calculated under implied weighting, with 500 replications of symmetric resampling ($P = 0.33$). The labels show the difference in support measured with the raw frequency of the group (above branches), and the difference in support measured using frequency differences (below branches; GC values; see Goloboff et al., 2003a). A plus sign indicates that the support is increased when adding the continuous characters, a minus sign indicates it is decreased.

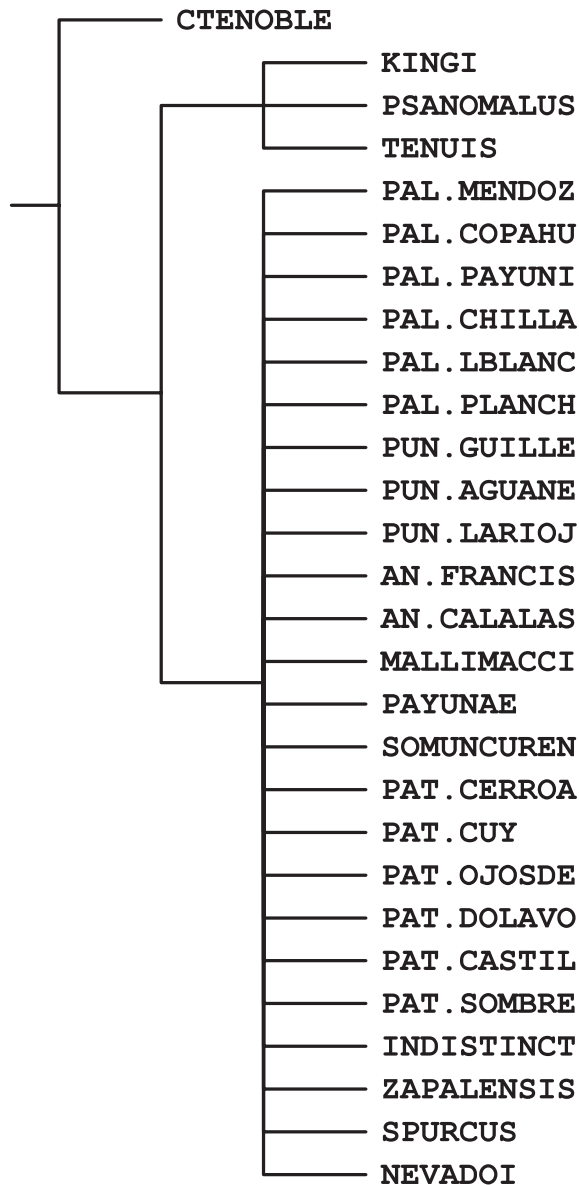


Fig. 5. Strict consensus of the groups supported by discrete characters alone, and by continuous characters alone, for the *Phymaturus* data set.

better assessments of tree similarity than agreement subtrees³. Transforming the consensus for the discrete data set into the consensus for the continuous data set requires a total of 31 moves. This indicates a large number of differences between the two trees, although much less

³Calculating SPR-distances between two trees is an NP-hard problem (Bordewich and Semple, 2005). The method used here (implemented in TNT) is a heuristic, which will be described in detail elsewhere (Goloboff, unpub. data). In the case of unresolved trees, the implementation in TNT considers the distance between the closest resolutions of the two trees.

than would be expected for random trees (out of 1000 pairs of random trees for 50 taxa, 0.04% differed in 44 moves, and 99.96% differed in 45 or more moves).

A better way to compare the extent of the conflict or concordance between the two subsets of data, which goes beyond simple topological comparisons, is to examine whether the support for the groups supported by the (more numerous) discrete characters alone is increased or decreased when the continuous characters are added to the data set (Fig. 4; note that some of the groups shown in that figure might actually be contradicted by the combined data). The measurements taken in this study were intended to help resolve relationships within the genus *Bothriurus*, as the outgroup comprises other genera. Within *Bothriurus*, the support (measured as the frequency under resampling) is decreased in five groups (with a sum of support differences of 33), increased in 25 (with a sum of support differences of 162), and unchanged in three. The general trend in group support within *Bothriurus*, when adding the continuous characters, is clearly an increase. Therefore, the continuous characters seem to provide additional support for many of the groups supported by the discrete characters. This suggests that, although continuous and discrete characters support different trees, the conflict between the two types of data is not very strong.

Phymaturus

The second example examined corresponds to the lizard genus *Phymaturus*. The matrix comprises 28 taxa, 39 continuous characters, and 90 discrete ones. The continuous and discrete characters analyzed by themselves produce different strict consensus; Fig. 5 shows

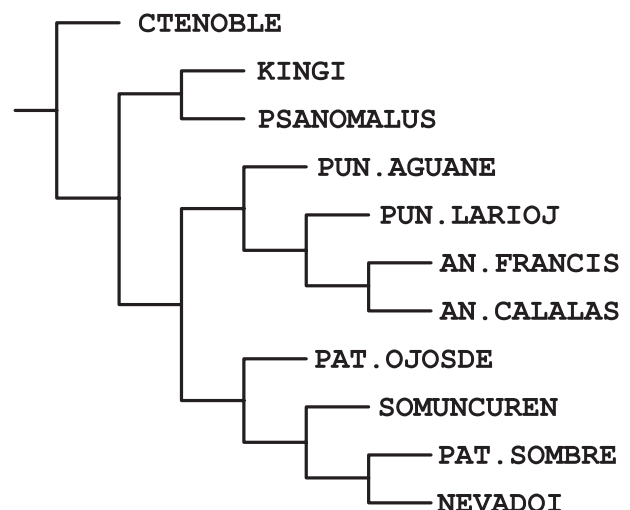


Fig. 6. Agreement subtree for the trees resulting from analysis of discrete characters alone, and continuous characters alone, for the *Phymaturus* data set.

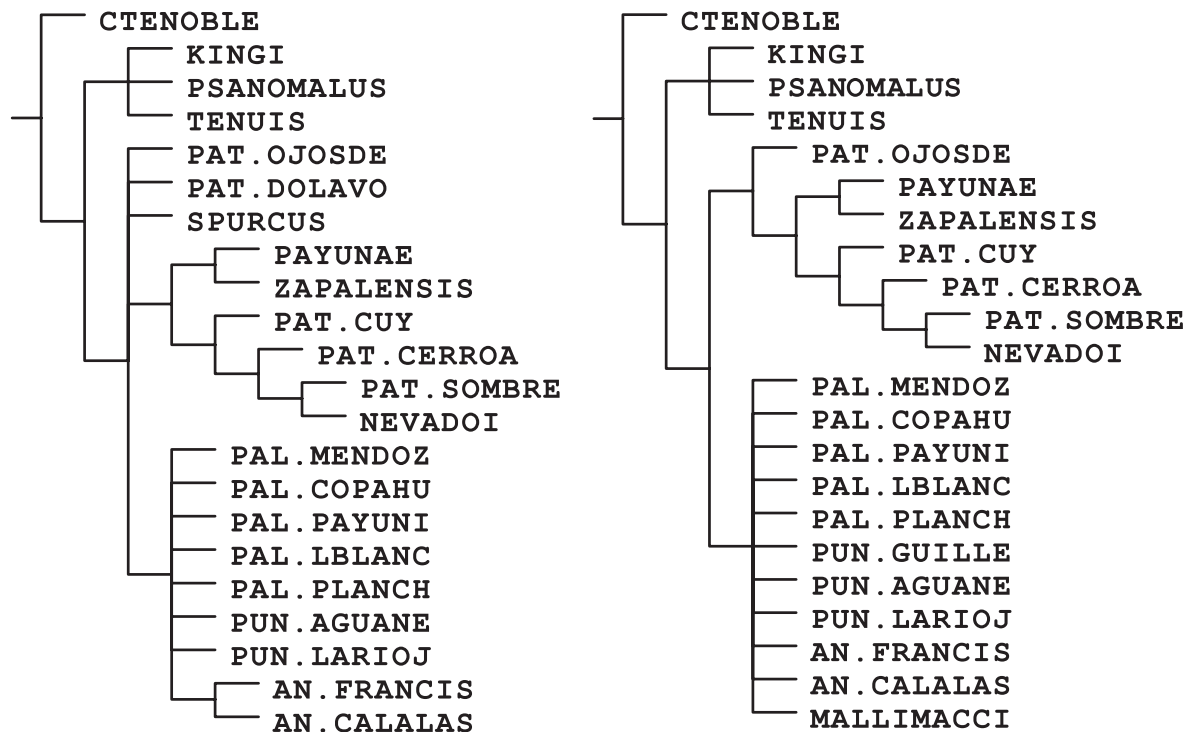


Fig. 7. Combinable components consensus of the strict consensuses resulting from analysis of discrete characters alone, and continuous characters alone. Two different sets of the six most conflictive species are excluded from the consensuses. The tree in the left excludes *pal.chilla*, *pun.guille*, *mallimacci*, *somunucuren*, *pat.castil*, and *indistin*. The tree in the right excludes *pal.chilla*, *somuncuren*, *pat.dolavo*, *pat.castil*, *indistinct* and *spurcus*.

the two consensuses and their grand consensus. The agreement subtree (Fig. 6), with 11 of 28 taxa, has a greater proportion of taxa than in the *Bothriurus* data set. When generating 1000 random pairs of trees for 28 taxa, 98.1% of the cases had eight or fewer taxa in the agreement subtree; not a single case had as many as 11 taxa.

The combinable component consensus of the two consensuses (Fig. 7) has 10 nodes when the six most conflictive taxa are excluded⁴. The results for the discrete and continuous data sets show much more topological concordance than in the case of *Bothriurus*. The number of moves necessary to convert the tree for one subset of data on to the tree for the other subset is 13. Generating 1000 pairs of (binary) random trees for 28 taxa, produced 99.8% of cases with 16 moves or more, and 99.6% of cases with 17 moves or more. The bulk of the cases (78.9%) had 20–23 moves. The degree of tree similarity for the *Phymaturus* data

subsets is, clearly, more than could be obtained by chance alone.

The group supports (for the groups supported by the discrete-only data subset, measured with frequency under resampling, with the positions of the six most conflictive taxa disregarded) when adding the continuous characters is increased or unchanged for 12 groups, and decreased for only four (Fig. 8).

Ischnothelinae

The last example examined here is the spider subfamily Ischnothelinae. This matrix has 25 taxa, 44 discrete characters and 46 continuous characters (measurements or meristic). Coyle (1995) reported 19 additional ratios, which were not considered here, as most of these are ratios between the other measurements (thus showing logical dependence). Nine of the 44 discrete characters had been discretized by Coyle (1995); three of these nine are measurements, and six of them are ratios (only the measurements are shared between the two data sets). Whether or not the nine characters discretized by Coyle are included in the discrete data set, there is a very high concordance between the strict consensus for the discrete and the continuous data sets (Fig. 9). The concordance becomes even more evident when the

⁴Note that here, and in all other cases of taxon exclusion, the analyses themselves included all taxa; the taxa were subsequently excluded only from the trees. In this way, the characters scored for the excluded taxa could potentially influence the relationships of the other taxa.

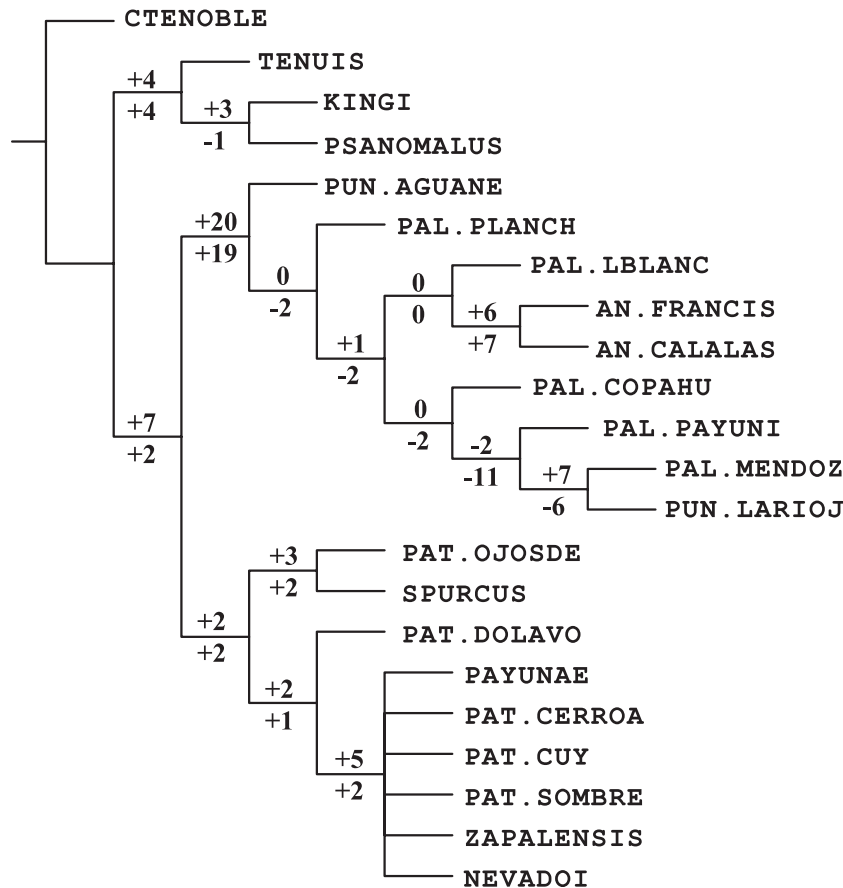


Fig. 8. Differences in group support, calculated as in Fig. 4. The tree excludes the same six species as the left tree in Fig. 7 (*pal.chilla*, *pun.guille*, *mallimacci*, *somunucuren*, *pat.castil* and *indistin*).

positions of the three most conflictive species (two of which are known only from females) are excluded from the consensus (Figs 10 and 11).

The number of SPR moves necessary to convert the tree for Coyle's original data on to the tree for the continuous characters is only 6. Of 1000 random pairs of trees for 24 taxa, 99.5% of cases had 13 moves or more, and 97.5% had 14 moves or more (the bulk of cases, totaling 69.0%, had 16–18 moves). Given that the results for the discrete and continuous characters are so similar, it is not surprising that the supports for the groups supported by the discrete characters alone are increased for the majority of the groups (and decreased only for three) when the continuous characters are added (Fig. 12).

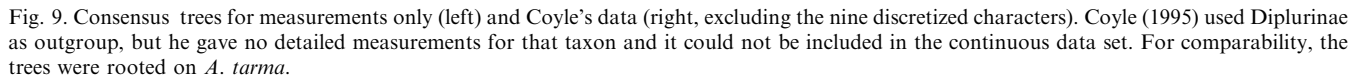
Discussion

The examples show that continuous characters may well carry useful phylogenetic information. Dismissing characters just because of their continuous nature is therefore obviously unwarranted. This should not be

taken as a claim that all attempts to consider quantitative characters are justified, or (even less) that just taking numerous randomly defined sets of measures would be useful to establish phylogeny. The infamous *Pemphigus* case attests against that idea⁵. The measurements used in the data sets analyzed here were not randomly chosen; most of them were selected for detailed measurements precisely because they seemed to correlate well with taxonomic groupings. Obviously, caution is also necessary to avoid using correlated measures (e.g., absolute dimensions of different parts, in animals of different size and isometric growth).

Among the potential problems of continuous characters is that the heritability of the traits measured is never perfect, but only the heritable aspect is of interest to the

⁵In the 70s, Robert Sokal and colleagues studied morphometrics of *Pemphigus* aphids (plant pests). They predefined a large number of measurements to take from all their specimens. Subsequently, an aphid taxonomist found out that what had been considered as a single species before was in fact more than one. The pheneticists had preserved only the measurements (not the specimens), but the measurements themselves were unable to even discriminate the species.



Pimentel and Riggins (1987) and Cranston and Humphries (1988) have questioned the use of continuous characters, on the grounds that means have no cladistic meaning. That is true (*contra* Rae, 1998), if the means correspond to individuals with genetic differences. For example, in a matrix that includes mammals and other vertebrates, what is the meaning of the mean weight of a mammal? At the most, rather than mixing whales and elephants with mice and shrews in the same count, one should try to determine the best weight estimate for the ancestor of mammals, and represent "mammals" by that estimate in the matrix. Assuming

This radical idea, however, creates more problems than it solves, because if all the morphological characters are considered as continuous then the problem of scaling becomes an issue for all the characters in the matrix. Additionally, the idea of having to take precise measurements (and calculate means and SD) in order for a researcher to determine whether a bird has wings, or whether a spider has spinnerets, or whether an arthropod has an exoskeleton, is simply not very

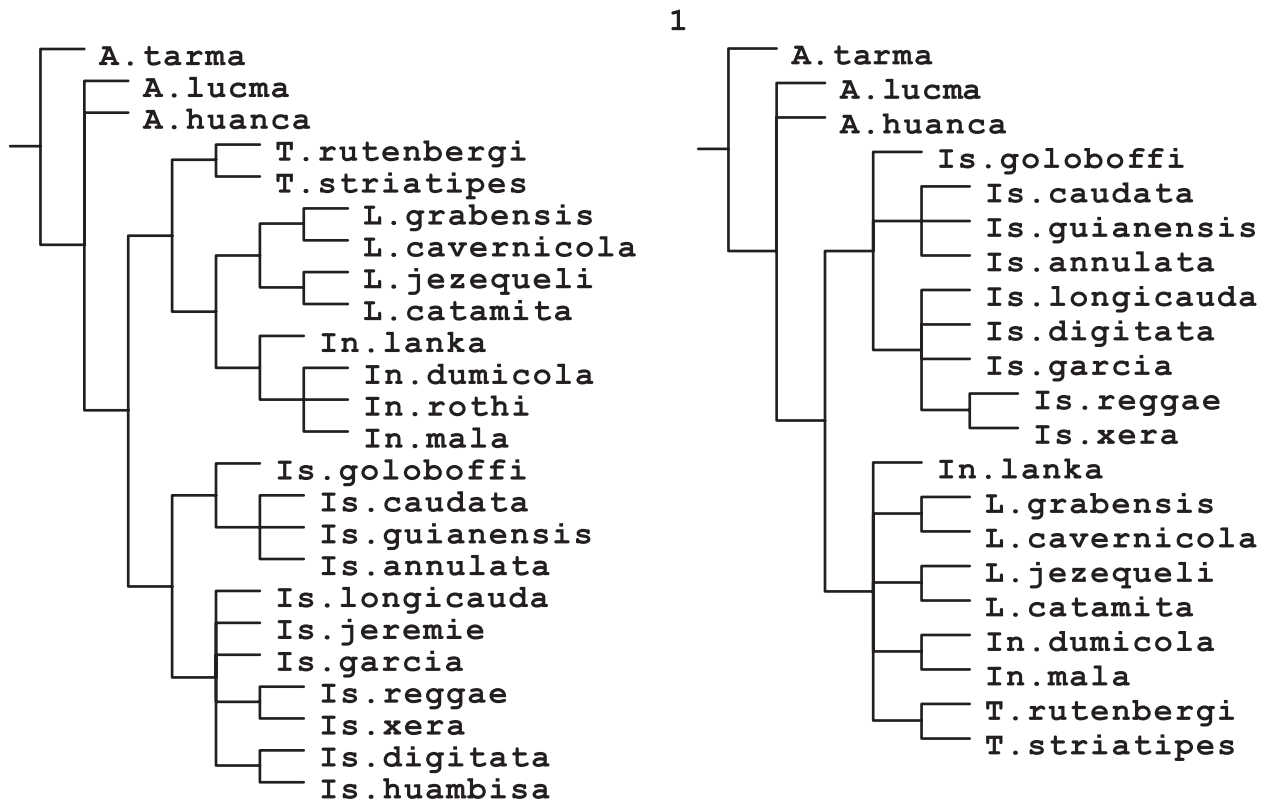


Fig. 10. Strict consensus for Coyle's original matrix (left, which includes nine discretized characters, out of 44), and the strict consensus of that tree with the measurements-only data set (right), when *In. rothi*, *Is. huambisa* and *Is. jeremie*, are excluded (the former two species are known only from females).

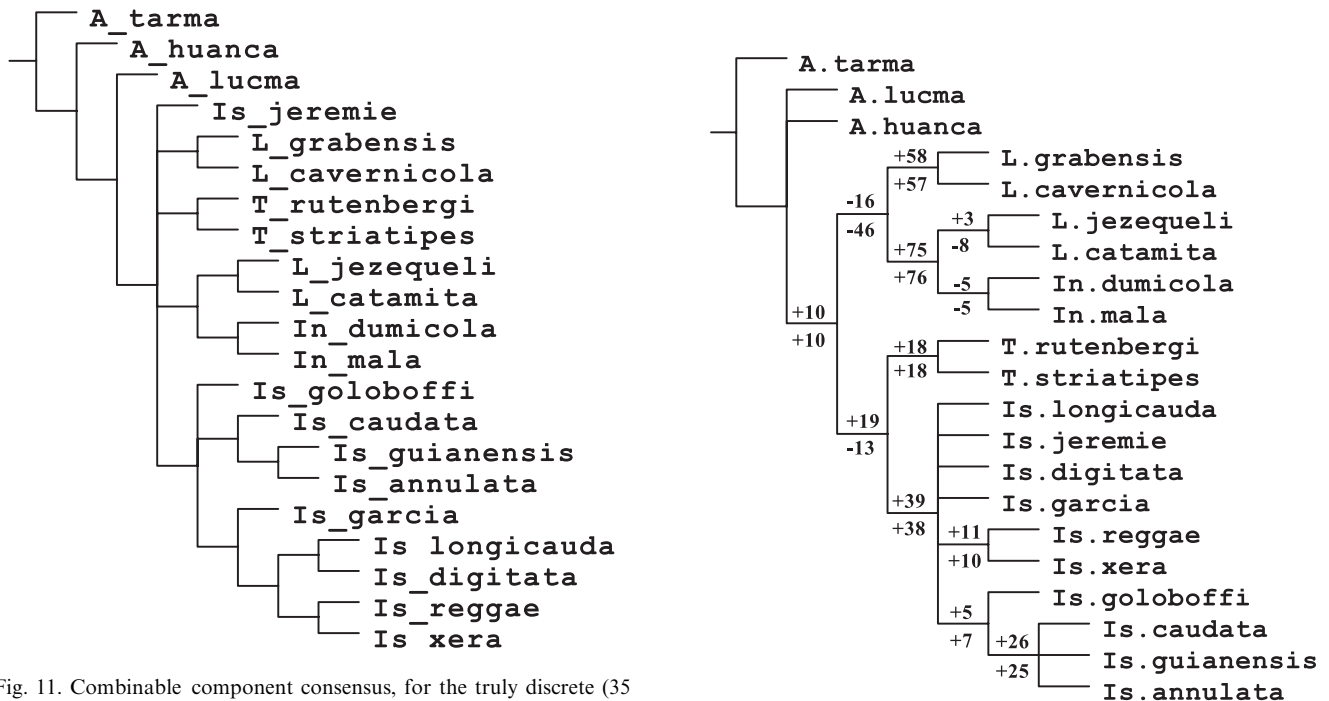


Fig. 11. Combinable component consensus, for the truly discrete (35 characters) and measurements only data sets, excluding three species (*Is. huambisa*, *In. rothi* and *In. lanka*; the former two are known only from females, so their taxonomic placement is expected to be ambiguous).

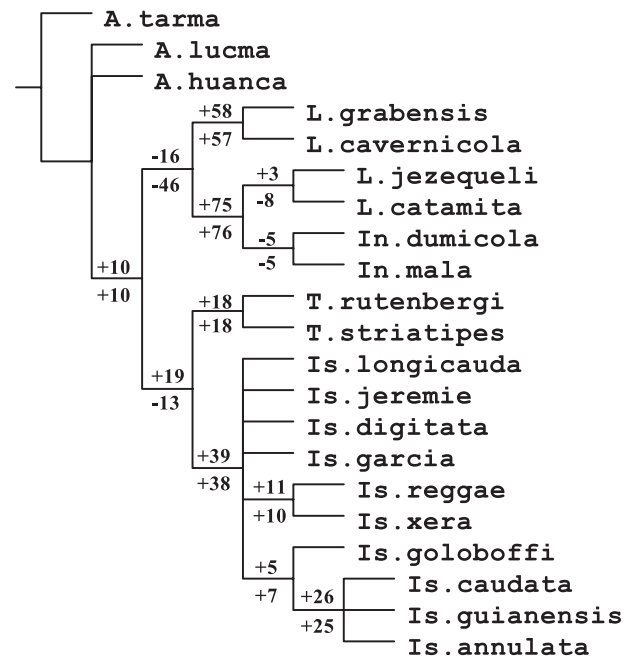


Fig. 12. Differences in group support, when adding continuous characters, to Coyle's original data, calculated as in Fig. 4. The same three species as in Fig. 11 were excluded from the consensus.

appealing—precisely because the outcome of such an exercise is so easily predictable. When characters show well-defined, discrete variation (i.e., well delimited states, gaps, little variation in terminal taxa) they should be coded as such—which is what phylogeneticists have been doing all along. By the same token, when characters that seem to be of taxonomic interest appear as continuous, they should be coded as such, instead of simply ignored or discretized.

References

- Archie, J., 1985. Methods for coding variable morphological features for numerical taxonomic analysis. *Syst. Zool.* 34, 326–345.
- Bordewich, M., Semple, C., 2005. On the computational complexity of the rooted subtree prune and regraft distance. *Ann. Combinatorics* 8, 409–423.
- Coyle, F., 1995. A revision of the funnelweb Mygalomorph spider subfamily Ischnothelinae (Araneae, Dipluridae). *Bull. Am. Mus. Nat. History* 226, 1–133.
- Cranston, P., Humphries, C., 1988. Cladistics and computers: a chironomid conundrum? *Cladistics* 4, 72–92.
- Eulenstein, O., Chen, D., Burleigh, G., Fernández-Baca, D., Sander-son, M., 2004. Performance of flip supertree construction with a heuristic algorithm. *Syst. Biol.* 53, 299–308.
- Farris, J., 1970. Methods for computing Wagner trees. *Syst. Zool.* 19, 83–92.
- Farris, J.S., 1988. Hennig86. Program and documentation. Port Jefferson Station, distributed by the author.
- Farris, J.S., 1990. Phenetics in camouflage. *Cladistics* 6, 91–100.
- Fitch, W., 1971. Toward defining the course of evolution: minimum change for a specific tree topology. *Syst. Zool.* 20, 406–416.
- Goldman, N., 1988. Methods of discrete coding of morphological characters for numerical analysis. *Cladistics* 4, 59–71.
- Goloboff, P., 1993a. Estimating character weights during tree search. *Cladistics* 9, 83–91.
- Goloboff, P., 1993b. Character optimization and calculation of tree lengths. *Cladistics* 9, 433–436.
- Goloboff, P., 1993c. Nona: A tree searching program. Program and documentation available at <http://www.zmuc.dk/public/phylogeny/Nona-Pee wee>
- Goloboff, P., 1998. Tree searches under Sankoff parsimony. *Cladistics* 14, 229–237.
- Goloboff, P., Farris, J., Källersjö, M., Oxelmann, B., Ramírez, M., Szumik, C., 2003a. Improvements to resampling measures of group support. *Cladistics* 19, 324–332.
- Goloboff, P., Farris, J., Nixon, K., 2003b. T.N.T. Tree Analysis Using New Technology. Program and documentation, available at <http://www.zmuc.dk/public/phylogeny/tnt>
- Hormiga, G., Scharff, N., Coddington, J., 2000. The phylogenetic basis of sexual size dimorphism in orb-weaving spiders (Araneae, Orbiculariae). *Syst. Biol.* 49, 435–462.
- Maddison, W., 1991. Squared-change parsimony reconstructions of ancestral states for continuous-valued characters on a phylogenetic tree. *Syst. Zool.* 40, 304–314.
- Maddison, W., Maddison, D., 1992. Mac-Clade, Version 3.0. Analysis of Phylogeny and Character Evolution. Sinauer, Sunderland, Massachusetts.
- Maddison, W., Maddison, D., 2005. Mesquite, a Modular System for Evolutionary Analysis, Version 1.06 <http://mesquiteproject.org>
- Mickevich, M., Johnson, M., 1976. Congruence between morphological and allozyme data in evolutionary inference and character evolution. *Syst. Zool.* 25, 260–270.
- Pimentel, R., Riggins, R., 1987. The nature of cladistic data. *Cladistics* 3, 201–209.
- Rae, T., 1998. The logical basis for the use of continuous characters in phylogenetic systematics. *Cladistics* 14, 221–228.
- Rogers, J., 1984. Deriving phylogenetic trees from allele frequencies. *Syst. Zool.* 33, 52–63.
- Sankoff, D., Rousseau, P., 1975. Locating the vertices of a Steiner tree in an arbitrary space. *Math. Program.* 9, 240–246.
- Schols, P., Dhondt, C., Geuten, K., Merckx, V., Janssens, S., Smets, E., 2004. MorphoCode: coding quantitative data for phylogenetic analysis. *Phyloinformatics*, 1, 4.
- Stevens, P.F., 1991. Character states, morphological variation, and phylogenetic analysis: A review. *Syst. Bot.* 16, 553–583.
- Swofford, D., 1993. PAUP: Phylogenetic analysis using parsimony. Vers. 3.0. Sinauer, Sunderland, Massachusetts.
- Swofford, D., 1998. PAUP*: Phylogenetic analysis using parsimony (*and other methods), Beta Version, 4.0. Sinauer, Sunderland, Massachusetts.
- Thiele, K., 1993. The Holy Grail of the perfect character: the cladistic treatment of morphometric data. *Cladistics* 9, 275–304.
- Wiens, J., 2001. Character analysis in morphological phylogenetics: problems and solutions. *Syst. Biol.* 50, 689–699.
- Zar, J.H., 1974. Biostatistical Analysis. Prentice-Hall, Englewood Cliffs, New Jersey.

Appendix 1

Script to make TNT discretize continuous characters, using Thiele's method. Thiele's method uses a single value, therefore if a taxon is assigned a range, the script considers the mean value.

```
macro= ;
if ( !largnumber )
errmsg Give number of states in which to divide continuous characters.; end
var: mins maxs range numsta uplim numcont interv thismean lolim
didsome ;
resettime ;
set numsta %1 ;
set numcont 0 ;
loop 0 nchar
if ( !iscont[#1] ) set numcont #1 ; endloop end
stop
if ( !'numcont' )
errmsg Matrix has no continuous characters. Nothing to recode!; end
set numcont -- ;
report- ;
if ( !windows ) sil = con ; end
lquote [ ;
lquote=;
macfloat 3 ;
loop 0 'numcont'
progress #1 'numcont' Working... ;
set mins 10000 ;
set maxs 0 ;
loop 0 ntax
if ( contmaxs[#1 #2] == 65.535 ) continue; end
set thismean ( contmins[#1 #2] + contmaxs[#1 #2] ) / 2 ;
if ( 'thismean' < 'mins' ) set mins 'thismean' ; end
if ( 'thismean' > 'maxs' ) set maxs 'thismean' ; end
stop
set range 'maxs' - 'mins' ;
set interv 'range' / 'numsta' ;
quote &l0&l0CHARACTER #1: interval size is 'interv' ;
loop 0 ( 'numsta' - 1 )
set lolim 'mins' + ( #2 * 'interv' ) ;
set uplim 'mins' + ( #2 + 1 * 'interv' ) ;
quote &l0 - state #2 ('lolim' - 'uplim'); ;
set didsome 0 ;
loop 0 ntax
if ( contmaxs[#1 #3] == 65.535 ) continue ; end
set thismean ( contmins[#1 #3] + contmaxs[#1 #3] ) / 2 ;
if ( ( ( 'thismean' == 'maxs' ) && ( #2 == ( 'numsta'-1 ) ) ) ||
( ( 'thismean' >= 'lolim' ) && ( 'thismean' < 'uplim' ) ) )
quote &32 $taxon #3 ;
set didsome ++ ;
xread = #1 #3 #2 ;
end
end
stop
if ( !'didsome' ) quote &32 (none!) ; end
stop
stop
quote &l0 Done! ;
lquote];
lquote- ;
set 0 time ;
quote TIME USED: '0' secs. ;
if ( !windows ) sil = con ; end
progress/;
report= ;
proc/ ;
```