

RUNNING HEAD: POINTS OF VIEW

MEASURING SUPPORT AND FINDING UNSUPPORTED RELATIONSHIPS IN
SUPERTRTEES

Mark Wilkinson¹, Davide Pisani^{1,2}, James A. Cotton¹ and Ian Corfe^{1,3}

¹*Department of Zoology, The Natural History Museum, London SW7 5BD, UK*

²*Department of Biology, National University of Ireland, Maynooth, Co. Kildare,*

Ireland

³*Department of Earth Sciences, University of Bristol, Bristol BS8 1RJ, UK*

Address correspondence to:

Dr Mark Wilkinson, Department of Zoology, The Natural History Museum, London,
SW7 5BD, UK.

Tel: +44 020 7942 5164, Fax: +44 020 7942 5054; e-mail: marw@nhm.ac.uk

Supertree methods can combine information in phylogenetic trees to yield novel relationships, but matrix representation with parsimony (MRP) supertree methods (Baum, 1992; Regan, 1992) sometimes return supertrees that include relationships which appear to have no support among the input trees, individually or jointly (Bininda-Emonds and Bryant, 1998; Pisani and Wilkinson, 2002; Wilkinson et al., 2004). Assessing the extent to which this might occur in practice requires a clear conception of how a set of input trees may provide support for relationships in supertrees. Bininda-Emonds (2003) broke new ground in presenting the first explicit conceptual analysis and categorisation of the kinds of correspondence that can occur between relationships in input trees and supertrees, and he investigated the frequency of unsupported relationships in some real supertrees and with simulations. He reported that unsupported clades were completely absent from the real supertrees and very rare in simulations, suggesting that unsupported groups are unlikely to be a problem for MRP in practice.

Here we present an alternative view of the correspondences between relationships in supertrees and input trees, and define associated measures that quantify these correspondences. We review previous work, contrast it with our own, and consider the implications. We draw heavily upon the treatment of analogous problems in the correspondence between characters and phylogenetic trees (Wilkinson, 1998). Following Bininda-Emonds (2003), we focus almost exclusively upon support for supertree clades (components, rooted full splits), as opposed to support for other relationships (e.g. resolved triplets, partial splits, nestings, subtrees) or nestings, but our approach readily generalises to unrooted trees. We thus aim to clarify how a rooted input tree can support or conflict with a supertree clade. All reference to Bininda-Emonds is to his 2003 article, unless otherwise indicated.

SUPPORT, CONFLICT, PERMISSION AND IRRELEVANCE

Support is an important concept in phylogenetic inference. We often speak of particular data supporting a phylogenetic hypothesis, and a number of indices are widely used to quantify support (see e.g. Wilkinson et al., 2003 for a recent discussion). Individual characters can support or conflict with particular relationships in phylogenetic trees, and characters can be treated as corresponding to the trees that they directly support (e.g. Wilkinson, 1998). For example, a parsimony-informative binary character corresponds to, and directly supports, a tree with one internal edge, and a multistate character corresponds to one (ordered) or more (unordered) trees with

more than one internal edge (assuming all states are informative). This correspondence underpins the various pseudocharacter matrix representations of trees (Wilkinson et al., 2004). Supertrees are phylogenetic inferences based on the evidence (the support) provided by a set of input trees. Thus, it should be possible to address the question of how individual input trees support or conflict with clades in the same way that the question of how individual characters support or conflict with clades has been addressed. We would expect treatments of these issues to have a consistent foundation.

Support and conflict may be taken as all or nothing, so that a character or a tree either supports a clade or does not, and a character either conflicts with a clade or does not. Alternatively, characters may be interpreted as having different fits to different trees and comparative or relative support for one tree over another is evaluated in terms of the differences in fit. The difference is analogous to the treatment of character data in clique and parsimony analyses: in the former each character provides a two-rank classification of possible trees (as conflicting or not) whereas the latter enables further ranking of conflicting sub-optimal trees. For a single character the methods agree upon the top rank, but parsimony may, in addition, allow us to assert that one sub-optimal tree is better supported by the character than another (what Wilkinson and Nussbaum (1996) referred to as qualified support). Here we also follow Bininda-Emonds in focussing only upon the simpler, all-or-nothing, interpretation of support and conflict.

In the special case where an input tree has the same leaf set as the supertree and both are fully resolved (i.e. comparing two binary trees of the same size) it is trivial to determine if a clade in the supertree is supported by the input tree. A supertree clade is supported by the input tree if it is present in the input tree. If the clade is not present, then the input tree must conflict with (contradicts, is incompatible with, is incongruent with, disagrees with) the clade. Both polytomies in, and leaves missing from, input trees complicate the situation. Where polytomies are interpreted as hard (Maddison, 1989) the above dichotomy of ‘supports or conflicts’ still holds. If, as here, polytomies are interpreted as soft, a third possibility arises – that of neither directly supporting nor conflicting with the clade.

In the more general case where an input tree has fewer leaves than the supertree then it cannot include any supertree clade. Thus, it cannot by itself support a supertree clade in the strict (Nixon and Carpenter, 1996) sense of including all the

relationships asserted by that clade. Taken in isolation, a single input tree can only strictly support the relationships that it includes, and yet input trees must support supertree clades in some less than strict sense, because supertree clades are inferred from the input trees. Note that input trees may jointly entail, and thus strictly support, novel relationships that are not strictly supported by any single input tree.

The support provided by a single input tree for a larger supertree clade is analogous to the support for a clade provided by a character that has some leaves scored as missing. Strictly speaking such a character does not support any specific clade because it only conveys information on a subset of leaves. However, it supports a subset of trees, those in which it can be mapped with no homoplasy. In this case it is natural to think of the character as potentially supporting a number of clades, each corresponding to the possible replacements of missing entries with character states (e.g. Wilkinson, 1998) and each requiring no homoplasy in the character. The potential clades are those that entail the relationships strictly supported by the character, but each potential clade includes additional information that is not directly supported by the character.

Our concept of support is analogous to Wilkinson's (1998) treatment of the all-or-none support provided for relationships by incomplete characters. It is founded upon the intuition that an input tree *supports* a supertree clade when all the relationships the supertree clade entails of just those leaves present in the input tree are displayed by that tree. Relationships asserted by the supertree clade that could not be present in the input tree (because relevant leaves are not present) are considered irrelevant to this assessment of support. In simple examples we can readily identify supportive correspondences between supertree clades and the typically less inclusive relationships in an input tree that conform to our intuition (e.g. Figs. 1 and 4).

Different supertree clades (that differ only in leaves that are not present in the input tree) can entail the same input tree relationships, and consequently a single input tree clade may simultaneously support more than one supertree clade (see Fig. 1). While the latter support is not strict (because not all leaves are present), we recognise it as somewhat stronger (or more clear cut) than when a supertree clade is but one of many supported by a single input tree clade. This would seem to be the strongest support a single input tree can provide in the typical case where it has fewer leaves than the supertree. That an input tree supports a particular clade in a given supertree does not mean that the input tree does not also support some other clade in some other

supertree. Input trees support a subset of the possible supertrees, those that entail them. This is analogous to ambiguous branch lengths and arbitrary resolutions that can result from incomplete characters that potentially support multiple clades in parsimony analysis (Wilkinson, 1995). Where a single input tree relationship supports multiple supertree clades, weighting so as to distribute the support among the supertree clades might be considered in the quantification of support (see below).

In contrast to support, conflict is more straightforward. Given two trees we can always say whether they conflict or not. They conflict if they assert logically contradictory relationships so that the supertree clade cannot be present in any tree that includes the relationships in the input tree. Conflict is well understood and efficient algorithms to determine the compatibility of sets of input trees have been known for some time (Aho et al., 1981). Although a supertree clade may assert some relationships that cannot be directly contradicted by an input tree (because they pertain to leaves that are absent from the input tree), an input tree *conflicts* with a supertree clade if it contradicts any relationship entailed by the supertree clade.

It is also possible for input trees to neither support nor conflict with a particular supertree clade. We say that an input tree *permits* a supertree clade when it could have supported it or conflicted with it but did not (because it was incompletely resolved). We say an input tree is *irrelevant* to a supertree clade when it could not have supported or conflicted with it (because it does not contain the relevant leaves). Note that irrelevance is always with respect to a particular supertree clade and should not be misinterpreted to imply an input tree is totally irrelevant. Support, conflict, permission and irrelevance are four exhaustive and mutually exclusive categories that describe the relation between any rooted input tree and a supertree clade. Simple exemplars of these are given in Figure 2 and more formal definitions are provided in the next section.

Alternative interpretations are possible. For example, consider instead that an input tree supports a supertree clade if *any* of the resolved triplets entailed by the former are also entailed by the latter, rather than *all* of them. The input tree we see as permitting the supertree clade (Fig. 2) would be interpreted as supporting the clade because both share a resolved triplet (AB)D. However, a consequence of this view is that a single input tree may both support and conflict with the same supertree clade. For example, if B is added next to A in the conflicting input tree in Figure 2 - to give (((AB)D)C), then (AB)D would support and (AD)C and (BD)C would conflict with

(ABC)DEF. This in no way invalidates triplet-based assessment of support and conflict, but suggests an incompatibility with our desire here for categories of support and conflict that are mutually exclusive attributes of whole input trees.

FINDING SUPPORT AND CONFLICT

Let $L(S)$ be the leaf set of the supertree S , and $L(I)$ the leaf set of an input tree I such that $L(I) \subseteq L(S)$ where I and S are both rooted X -trees in the sense of Semple and Steel (2003: 16-17). A supertree clade (or non-trivial split) σ partitions the supertree leaf set $L(S)$ into two non-empty sets. In the case of rooted trees, we can distinguish these two sets as $L(S)_{\sigma M}$ and $L(S)_{\sigma N}$, the members and non-members of the clade, respectively (so that the subtree induced by $L(S)_{\sigma N}$ includes the root). Similarly, σ partitions $L(I)$ into two sets (one of which may be empty) $L(I)_{\sigma N} = L(I) \cap L(S)_{\sigma N}$ and $L(I)_{\sigma M} = L(I) \cap L(S)_{\sigma M}$ that may define a split (clade) on the input tree leaf set.

If $|L(I)_{\sigma N}| > 0$ and $|L(I)_{\sigma M}| > 1$ (which can be determined without considering relationships in the input tree), then the supertree clade induces a parsimony informative split on the leaf set of the input tree, otherwise I is *irrelevant* to σ . If the induced split is present in I then I *supports* σ . This is demonstrated by comparison with the splits in I . I *conflicts* with σ when the induced split contradicts relationships in the input tree. This is demonstrated by pairwise incompatibility (e.g. Semple and Steel, 2003) of the induced split with any of the splits in I or with using the algorithm of Aho et al. (1981). I *permits* σ when the parsimony informative split induced by σ on I is a resolution of a polytomy in I : permission is what remains when other categories are ruled out.

Whether an input tree supports, conflicts with, or permits a supertree clade can also be diagnosed by measuring the parsimony fit of the of the character encoding of the induced split to the input tree. The categories correspond to the three mutually exclusive and exhaustive possible combinations of perfect and imperfect (i.e. extra steps > 0) fits, under soft and hard interpretations of input tree polytomies. A perfect fit under the hard interpretation (entails the same for the soft) diagnoses support, and an imperfect fit under the soft interpretation (entails the same for the hard) diagnoses conflict. A combination of perfect and imperfect fits with polytomies interpreted as soft and hard respectively diagnoses permission. The fourth combination (perfect hard fit and imperfect soft fit) is impossible.

PRUNING AND GRAFTING

The support provided by one tree for clades in another is most clear-cut when the trees have identical leaf sets. When trees do not we might facilitate comparison by converting them into trees with the same leaf set. As Bininda-Emonds noted, one means of such conversion is to prune those leaves that are not present in the input tree from the supertree (e.g. Creevey et al, 2004). Our use of the splits induced by supertree clades upon input tree leaf sets for defining and finding support and conflict is equivalent to this pruning operation.

An alternative is to graft the missing leaves onto the input tree to produce a supertree-sized extended input tree. For example, there are 13 positions to which D can be grafted onto the input tree in Figure 1, giving a corresponding set of 13 extended input trees that each display (include, contain or entail) the original input tree, nine of which are fully resolved, and four of which include D in a polytomy. We refer to the set of all such extended input trees as the span $\langle I \rangle$ of the input tree I (e.g. Bryant and Steel, 1995).

A supertree clade can be in all of the trees in $\langle I \rangle$ only if $L(I) = L(S)$ and the clade is strictly supported. If the input tree conflicts with the supertree clade then it will not be in any of the trees in $\langle I \rangle$. If an input tree has fewer leaves than the supertree, it is always possible to graft missing leaves onto the input tree (and resolve polytomies) so that some trees in $\langle I \rangle$ conflict with and some strictly support the supertree clade, including when the input tree intuitively supports the supertree clade (as in Fig. 1). The membership of $\langle I \rangle$ can reveal conflict is of no help in recognising support, permission and irrelevance.

MATCHES AND MISMATCHES

In his pioneering treatment, Bininda-Emonds considered support and conflict between supertrees and input trees in terms of their constituent clades, with the overall relation of input tree and supertree clade considered a function of the relations of the individual input tree clades. He attempted to distinguish and define five distinct types of correspondence between a supertree clade and a single input tree clade. He also illustrated his five categories with the examples (Table 1) that we also present graphically in Figure 3.

A *hard match* is synonymous with strict support and requires supertree sized input trees which, as Bininda-Emonds notes, may be expected to be relatively rare and unimportant in supertrees. A *hard mismatch* occurs (p. 840) “when the source tree clade contradicts directly the relationships presented in the supertree clade”. A further

restriction, that a hard mismatch requires all taxa in the supertree clade must be present in the input tree was mistaken (Bininda-Emonds, pers. comm.). If there is a hard mismatch between any input tree clade and the supertree clade then there is a hard mismatch between the input tree and the supertree clade. The concept of hard mismatch is thus synonymous with the well-understood concept of conflict, but is the only one of the five categories that corresponds to one of ours.

All other cases are varieties of *equivocal matches* that (p. 840) "usually result from the presence of missing taxa in the source tree", and which are further divided as follows: "In a *soft match*, addition of the missing taxa may support the supertree clade but never contradict it. Conversely, in a *soft mismatch*, the missing taxa may contradict the supertree clade but never support it. *True equivocal matches* result when the supertree clade contains the source tree clade or vice-versa or when the missing taxa can both support and contradict the supertree clade" (our italics). Summing across the individual input tree clades provides an assessment of the overall relationship between an input tree and a supertree clade. Thus (p. 841) "For a soft match the missing taxa will never contradict the clade (i.e. the number of individual soft matches $>$ number of individual soft mismatches = 0), whereas for a soft mismatch they will never support it (i.e. number of individual soft mismatches $>$ number of individual soft matches = 0). True equivocalness represents all remaining options."

Reference to the addition of missing taxa suggests the grafting operation defining $\langle I \rangle$. However, it is always possible to add a missing leaf to an input tree so as to conflict with an otherwise uncontradicted supertree clade. Soft matches and soft mismatches cannot exist if the addition of missing taxa is understood as the operation defining $\langle I \rangle$. They are possible only if we do not consider all of the possible relationships in $\langle I \rangle$. This occurs because the input tree is first broken down into its constituent clades (using a matrix representation) before the 'addition' of missing leaves which then considers only whether the missing leaves are included in a given clade or not (i.e. are scored as one or zero in the corresponding matrix representation of the input tree clade). In the example we considered earlier (Fig. 1), $\langle I \rangle$ includes 13 trees but Bininda-Emonds' method considers only the four trees in which the missing leaf is attached to form a (hard?) polytomy.

That soft matches and mismatches require decomposition of input trees into clades and consequent consideration of only a subset of the possible relationships of

missing leaves may not be readily apparent from the original exposition because worked examples are only of input trees with single clades (Fig. 3). Even with these examples only a limited number of the possible relationships are considered. Consider the two examples of soft matches (Figure 3, examples 6 and 7). Attention is restricted only to the implications of the missing leaf either being a member of the single input tree clade or not. The possibility that, for example, D is the sister of C, in which cases the input trees would conflict with the supertree clade, is not considered. In our view both of these input trees support the supertree clade by virtue of the fact that the supertree clade entails the relationships in the input trees. Potential strict support (hard matches) or conflict resulting from the addition of missing leaves is considered unimportant, and selective consideration of the addition of missing leaves is considered misleading.

Similarly, in each of the two examples of soft mismatches (Figure 3, examples 3 and 8) whether it is impossible to graft the leaf in such a way as to support the supertree clade (as is required of a soft mismatch) or not depends upon the interpretation of polytomies in the input tree – it is only with the hard interpretation that these input trees can never support the supertree clade. In both the examples, the input tree clade does not directly conflict with the supertree clade. In our view the input trees permit the supertree clade, with potential conflict (and strict support) considered unimportant and selective consideration of this unhelpful. Of the three examples of true equivocal matches (Fig 3, examples, 1, 5 and 9), we interpret the first two of these as examples of input trees that permit the supertree clade and the third as a case of support rather than of equivocation.

In summary, there are substantial differences between the alternative treatments of the kinds of relations that can pertain between an input tree and a supertree clade. Only one category, conflict, is the same in both. Whereas Bininda-Emonds does not distinguish between support, permission and irrelevance in the senses we have described, we see no need for additional categories, or reason to consider only the subset of possible relationships of missing leaves upon which they depend. The differences impinge on the behaviour of measures based upon these alternative foundations.

MEASURING SUPPORT AND CONFLICT

Building on his qualitative categories, Bininda-Emonds devised a quantitative measure for supertree clades. Input trees are scored +1 for hard matches, + 0.5 for soft

matches, 0 for equivocal matches, -0.5 for soft mismatches and -1 for hard mismatches. These scores are averaged across all input trees to give the qualitative support (QS) for a particular clade that ranges from +1 to -1 (with $QS = -1$ distinguished as *hard conflict*). He suggested that average QS across all supertree clades provides a measure of overall support for the supertree.

Our alternative formulation also lends itself to simple quantitative measures. Let t be the number of input trees, s the number of input trees supporting a supertree clade, r the number of input trees that are irrelevant to the supertree clade, q the number of input trees that conflict with the supertree clade, and p the number of input trees that permit the supertree clade, so that $t = p + q + r + s$. Where several supertree clades are supported by the same input tree relationships, we think it useful to spread the support provided by the input tree across the supertree clades by assigning a weight of $1/b$ to the support that an input tree provides to a supertree clade, where b is the number of supertree clades that entail the same parsimony informative split on the input tree leaf set. The sum across all input trees is ws , the weighted support for the supertree clade. We define ss , the strongest support for a supertree clade, as the number of input trees that support the supertree clade with $b = 1$. The number of input trees supporting a supertree clade, s is analogous to a measure of support provided for splits by character data (Wilkinson, 1998).

As Bininda-Emonds recognised in developing QS , it may also be useful to have a measure of the overall quality of a supertree clade that tells us something of the extent to which the input trees support or conflict with it. We call V the *value* of a supertree clade, where $V = (s - q)/(s + q)$, and zero divided by anything is taken to be zero. Both permitting and irrelevant input trees are treated as unimportant to a supertree clade's value. If $V = 1$ this tells us that all important input trees support the supertree clade (whereas $QS = 1$ tells us that all input trees strictly support the supertree clade). If $V = -1$ we know all important input trees conflict with the supertree clade (whereas $QS = -1$, tells us all input trees conflict with the supertree clade). If $V = 0$ then there are equal number of input trees supporting and conflicting with the supertree clade (whereas if $QS = 0$ we do not know the relative numbers of input trees supporting and conflicting with the supertree clade). Two simple variants, V^+ and V^- , reflect alternative interpretations of p , the number of input trees that permit the supertree clade, should they be considered relevant in interpreting the value of a supertree clade. $V^+ = (s - q + p)/(s + q + p)$, so that the failure of an input tree to

contradict the supertree clade when it could have done so is a vote in favour of the supertree clade, and $V = (s - q - p)/(s + q + p)$, so that a failure of an input tree to support a supertree clade is taken as a vote against the supertree clade. These measures are analogous to QS in that they have the same range as QS and similar intent, and the average across all supertree clades might be used as an overall measure of supertree ‘quality’ but they are otherwise quite different. Average V does not provide any basis for choosing among supertrees not least because it does not penalise polytomies.

A SIMPLE COMPARISON

We use a simple example from Gordon’s (1986) seminal work to highlight some differences between our and Bininda-Emonds’ measures. The two input trees do not conflict with each other or with their strict component consensus supertree. In fact all the novel relationships in the supertree (those not present in any input tree) are jointly entailed by the input trees, and we would prefer measures of support that recognise them as supported. Table 2 summarises the dramatically different assessments of the supertree clades. Using Bininda-Emonds’ scheme, only three of the nine supertree clades receive any sort of support (in the form of soft matches) from the first input tree, with the six others registering soft mismatches. In contrast, our assessment is that the first input tree is totally irrelevant to two of the supertree clades (H and I) and supports all the others. In the case of the second input tree, Bininda-Emonds’ approach finds six supertree clades have soft matches, one is truly equivocal and two have soft mismatches. In contrast, we see the input tree as irrelevant to one supertree clade (A) and as supporting the others. Overall, QS for the supertree clades ranges from -0.5 to +0.5. In contrast the value of every supertree clade is +1, telling us that each is supported by all the relevant input trees. Note that the input trees jointly entail all of the clades in the supertree, and thus we would hope any measure would indicate that all the relationships to be supported.

This example demonstrates that QS can give results that are quite contrary to our intuitions. Certainly, Gordon’s example has never been considered to involve conflict (mismatches) of any sort. Mean QS (approximately 0.028) is far lower than one might expect for a case where there is no conflict at all and all supertree clades are strictly supported through joint entailment. We believe that the failure to match our intuitions is because QS and the categories it is based on are not well-founded.

Using simulations, Bininda-Emonds showed that QS for clades is positively correlated with their MRP bootstrap support (i.e. bootstrapping the matrix elements rather than the input trees) but the present example furnishes a case where there is no correlation. Consider what would be expected of bootstrap support for the supertree clades in the contrasting cases of (1) having just the two input trees, and (2) having each input tree repeated an arbitrarily large number of times. In the former case, MRP bootstrap proportions are expected to be less than maximal (they range from 60 to 76) and s will be small, because there are so few input trees. In the latter, bootstrap proportions are expected to be maximal and s large. Comparing the two cases, bootstrap proportions and s behave as we would expect: they increase as the strength of support (the number of input trees) increases. In contrast, QS and V are unchanged in both cases. Clearly neither measure captures all aspects of support. V measures the extent to which the available evidence supports a supertree clade irrespective of whether that evidence is sufficient to support high bootstrap proportions. Thus in both cases it correctly tells us that all the supertree clades are supported and there is no conflict among the input trees. In contrast, QS gives us a picture of support for supertree clades and of overall support that we find confusing and potentially misleading in both cases.

UNSUPPORTED GROUPS

Bininda-Emonds used QS and his conceptualisation to assess how frequent unsupported groups are in MRP supertrees using both simulations and published supertrees. But what exactly is an unsupported group? For us a supertree clade is *unsupported* precisely when it is not supported by any input tree, i.e. when $s = 0$, and such clades may be considered problematic precisely because they lack any support. Supertree clades would be more objectionable if, in addition to lacking support, they conflicted with any input tree ($q > 0$). They are more objectionable still to the extent that they conflict with more of the input trees (up to $q = t$) or with more of the input trees that they could conflict with (up to $q = t - r$). The latter has been briefly discussed as a weakened co-Pareto supertree axiom that MRP does not obey (Wilkinson et al., 2004).

Bininda-Emonds' assessment of the frequency of unsupported groups in real supertrees and simulations, counted supertree clades as unsupported only when all input trees conflict with the supertree clade, something he termed *hard conflict*. This is rather high in the hierarchy of objectionable relationships that might occur in

supertrees, and we see it as an unnecessarily severe restriction on what is construed as unsupported. In particular, if any input tree is irrelevant to a supertree clade then even universal conflict among the more relevant trees will not be counted as an instance of an unsupported supertree clade. Consequently, for many supertree analyses, such “unsupported groups” are impossible. We do not know the number of groups with no support in previously reported simulations and empirical examples, rather we know only the subset of such groups that were in conflict with every input tree, which is likely to be an underestimate of the frequency of unsupported groups.

We evaluated nine real supertrees, including six reported on by Bininda-Emonds (Table 3). We find approximately 5% of the supertree clades are unsupported in the sense we have defined. Over a third of these unsupported clades also conflict with at least one input tree, one conflicts with all relevant input trees, but none conflict with all the input trees and two-thirds are from a single study. In simulations, Bininda-Emonds found unsupported groups to be commonest when there were few input trees, and here also it is the two studies with the fewest input trees that have the most unsupported groups presumably because there is less support to be had when there are fewer input trees. Our results show that unsupported groups are far from ubiquitous (four supertrees are entirely free of them) and suggest that good sampling may help minimise unsupported groups in practice.

DISCUSSION

As supertree construction has become more commonplace, the need for measures of the support for relationships in supertrees has increased. Although a laudable first attempt, we are not convinced that *QS* fulfils this need. We have highlighted a number of problems with the definitions and exclusivity of the categories of support and conflict upon which *QS* is based. Comparison of our and Bininda-Emonds’ alternatives using simple hypothetical examples and data from real supertrees demonstrates that the latter gives confusing and counterintuitive results and that *QS* underestimates the frequency of groups lacking support. We believe these problems undermine previous conclusions as to the rarity and unimportance of unsupported groups in MRP supertrees and call into question the utility of *QS* and its foundations, which should not be used uncritically, if they are used at all.

Although the frequency of unsupported groups might seem important in determining to what extent they are a problem for MRP in practice, we think it plausible that they are only the most obvious manifestation of undesirable aspects of

how MRP resolves conflict and that they may not be unconnected to reported biases of MRP with respect to input tree size (Purvis, 1995; Bininda-Emonds and Bryant, 1998) and shape (Wilkinson et al., 2001; 2005). In our view, these problems should not be dismissed as unimportant even if they prove to be uncommon. Rather, we see the further development of supertree methods that are designed not to have these undesirable properties as the surest means of avoiding them.

We hope that our analysis and measures will prove helpful in distinguishing seemingly well and poorly supported relationships but we recognise that altogether better measures based on fundamentally different approaches may be developed. We also stress that they are intended to complement, and not to substitute for, method-based measures of support such as bootstrap proportions (e.g., Creevey et al., 2004) and measures based on differences in fit.

Acknowledgments

We thank Olaf Bininda-Emonds for providing treefiles, for detailed comments on earlier drafts of this paper and for providing the impetus to our study. We also thank Andy Purvis, Vincent Savolainen, Rod Page and an anonymous reviewer for their helpful comments. This work was funded by BBSRC grant 40/G18385 and an NHM MRF award. Software to compute the measures we have introduced is available from <http://taxonomy.zoology.gla.ac.uk/~jcotton/>.

REFERENCES

- Aho, A. V., Y. Sagiv, T. G. Szymanski, and J. D. Ullman. 1981. Inferring a tree from lowest common ancestors with an application to the optimization of relational expressions. *SIAM J. Computing* 10: 405-421.
- Baum, B. R. 1992. Combining trees as a way of combining data sets for phylogenetic inference, and the desirability of combining gene trees. *Taxon* 41:3-10.
- Bininda-Emonds, O.R.P. 2003. Novel versus unsupported clades: assessing the qualitative support for clades in MRP supertrees. *Syst. Biol.* 52:839-848.
- Bininda-Emonds, O. R. P., and H. N. Bryant. 1998. Properties of matrix representation with parsimony analyses. *Syst. Biol.* 47: 497-508.
- Bininda-Emonds, O.R.P., J. L. Gittleman, and A. Purvis. 1999. Building large trees by combining phylogenetic information: A complete phylogeny of the extant Carnivora (Mammalia). *Biol. Rev.* 74: 143-175.

- Bryant, D., and M. A. Steel. 1995. Extension operation on sets of leaf-labelled trees. *Adv. Appl. Math.* 16: 425-453.
- Creevey, C. J., D. A. Fitzpatrick, G. K. Philip, R. J. Kinsella, M. J. O'Connell, M. M. Pentony, S. A. Travers, M. Wilkinson, and J. O. McInerney. 2004. Does a tree-like phylogeny only exist at the tips in the prokaryotes? *Proc. R. Soc. B.* 271: 2551-2558.
- Gordon, A. D. 1986. Consensus supertrees: the synthesis of rooted trees containing overlapping sets of labeled leaves. *J. Classif.* 3: 31-39.
- Kennedy, M., and R. D. M. Page. 2002. Seabird Supertrees: Combining partial estimates of Procellariiform phylogeny. *Auk* 199:88-108.
- Maddison, W. P. 1989. Reconstructing character evolution on polytomous cladograms. *Cladistics* 5: 365-377.
- Nixon, K. C. and J. M. Carpenter. 1996. On consensus, collapsibility and clade concordance. *Cladistics* 12: 305-321.
- Pisani, D., and M. Wilkinson. 2002. MRP, taxonomic congruence and total evidence. *Syst. Biol.* 51:151-155.
- Pisani D., A. M. Yates, M. C. Langer, and M. J. Benton. 2002. A genus-level supertree of the Dinosauria. *Proc. R. Soc. B* 269: 915-921.
- Purvis, A. 1995. A modification to Baum and Ragan's method for combining phylogenetic trees. *Syst. Biol.* 44: 251-255.
- Ragan, M. A. 1992. Phylogenetic inference based on matrix representation of trees. *Mol. Phy. Evol.* 1:53-58.
- Semple, C. and M. Steel. 2003. *Phylogenetics*. Oxford University Press. Oxford.
- Stoner, C.J., O. R. P. Bininda-Emonds, and T. M. Caro. 2003. The adaptive significance of coloration in lagomorphs. *Biol. J. Linn. Soc.* 79: 309-328.
- Wilkinson, M. 1995. Arbitrary resolutions, missing entries and the problem of zero-length branches in parsimony analysis. *Syst. Biol.* 44: 108-111.
- Wilkinson, M. 1998. Split support and split conflict randomization tests in phylogenetic inference. *Syst. Biol.* 47: 673-695.
- Wilkinson, M. and R. A. Nussbaum. 1996. On the phylogenetic position of the Uraeotyphlidae (Amphibia: Gymnophiona). *Copeia* 1996: 550-562.
- Wilkinson, M., J. L. Thorley, D. T. J. Littlewood, and R. A. Bray. 2001. Towards a Phylogenetic Supertree for the Platyhelminthes? Pages 292-301 *in*

Interrelationships of the Platyhelminthes (D. T. J. Littlewood and R. A. Bray, eds.) Chapman-Hall, London.

Wilkinson, M., F. J. Lapointe, and D. J. Gower. 2003. Branch lengths and support. *Syst. Biol.* 52:127-130.

Wilkinson, M., J. L. Thorley, D. Pisani, F.-J. Lapointe, and J. O. McInerney. 2004. Some desiderata for liberal supertrees. Pages 227-246 *in* *Phylogenetic Supertrees: Combining Information to Reveal the Tree of Life*. (O. R. P. Bininda-Emonds, ed.) Kluwer Academic, Dordrecht, The Netherlands.

Wilkinson, M., J. A. Cotton, C. Creevey, O. Eulenstein, S. R. Harris, F.-J. Lapointe, C. Levasseur, J. O. Mcinerney, D. Pisani, and J. L. Thorley. 2005. The shape of supertrees to come: tree shape related properties of fourteen supertree methods. *Syst. Biol.* in press.

FIGURE LEGENDS

Figure 1. Correspondences between a supertree and an input tree. Solid and open circles indicate possible locations in the input tree where the missing leaf (D) can be grafted to produce binary or polytomous trees respectively. Dashed lines indicate corresponding nodes and arrows indicate logical entailment of input tree relationships by supertree clades. Two of the supertree clades (2 and 3) are simultaneously supported by the same input tree clade (II) while the other supertree clades (1 and 4) are each supported by a different input tree clade (I and III).

Figure 2. Summary diagram of the four mutually exclusive relationships between input tree as a whole and supertree clades. The supportive input tree is the only one with a branch that corresponds to (is entailed by) the supertree clade.

Figure 3. Examples of correspondences between input trees and a supertree clade according to Bininda-Emonds (2003: Table 1). Zeros and ones indicate the character encodings of the trees. ($s > i$) indicates that the supertree clade includes the input tree clade, ($i > s$) indicates the opposite.

Figure 4. Two compatible input trees and their strict component consensus supertree (after Gordon, 1986). Input tree clades supporting particular supertree clades are given in parentheses. Note that all the supertree clades are entailed by the input trees together in that any supertree that displays both input trees must include these clades. Supertree clades can in this way be strictly supported by sets of trees that cannot strictly support the clade individually.

Table 1. Classification of correspondences between relationships in the input tree clades represented by binary character encodings and those in the supertree clade (111000) in Figure 3. After Bininda-Emonds (2003: Table 1) with addition of suggested corresponding categories from compatibility analysis in parentheses. Compatible and incompatible are usually taken to be mutually exclusive and exhaustive categories (Semple and Steel, 2003) undermining the claimed correspondences.

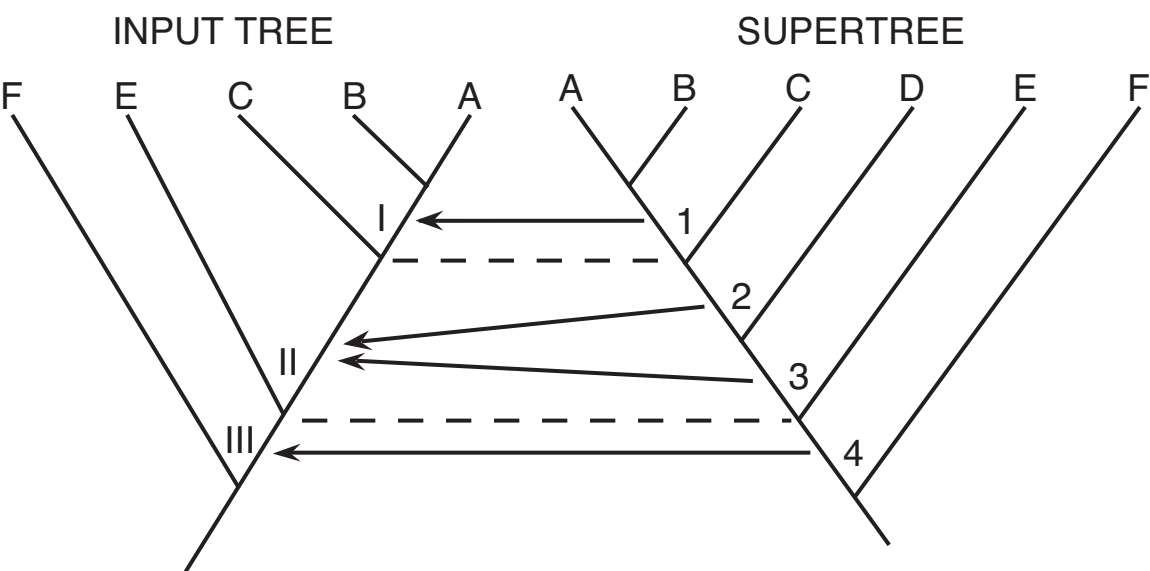
N	Input tree clade	Supports (compatible)	Does not support (not compatible)	Contradicts (incompatible)	Does not contradict (not incompatible)	Interpretation
1	110000		X		X	Equivocal
2	110100		X	X		Hard mismatch
3	110?00		X	X	X	Soft mismatch
4	111000*	X			X	Hard match
5	111100		X		X	Equivocal
6	111?00	X	X		X	Soft match
7	11?000	X	X		X	Soft match
8	11?100		X	X	X	Soft mismatch
9	11??00	X	X	X	X	Equivocal

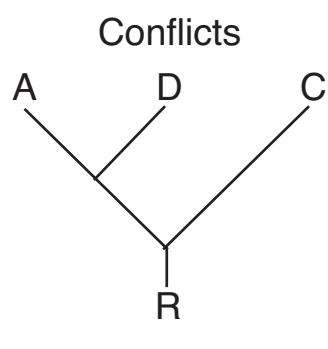
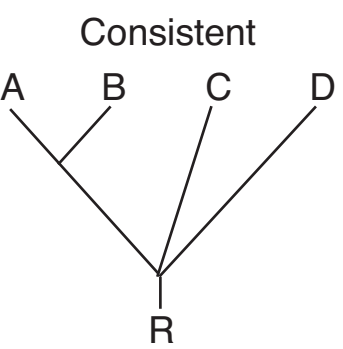

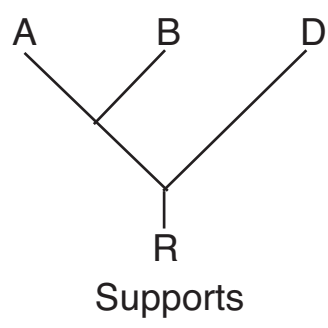
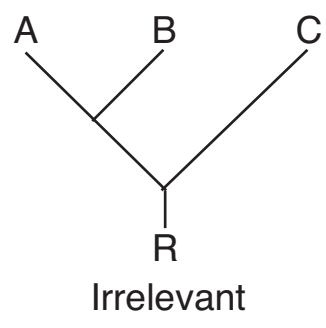
Table 2. Comparative assessment of the support and conflict provided by the input tree and its clades for the supertree clades in the example in Figure 4, using the method of Bininda-Emonds and our alternative approach. Cat is our categorisation of the relation between input tree and supertree clade, with s indicating support and r indicating irrelevance. Soft matches, soft mismatches and equivocal matches are indicated by +, - and = respectively. QS is qualitative support, V is the value of the supertree clade, and ws is the weighted support. QS is given separately for each input tree and for the input trees combined. Overall QS is 0.028.


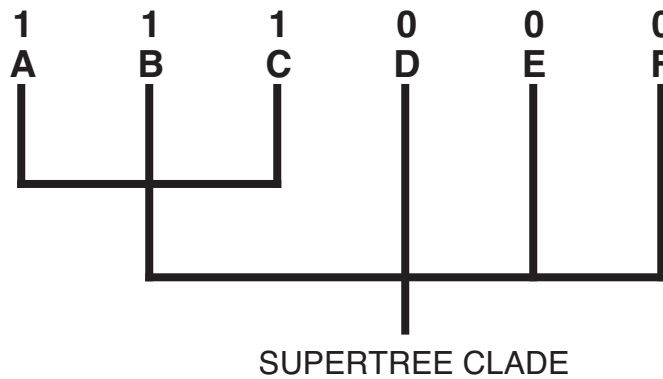




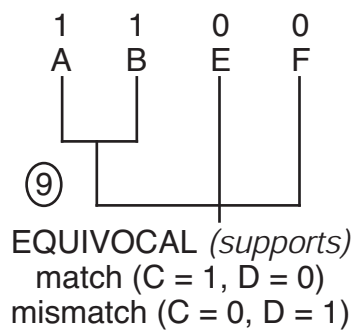



S_C	Input tree 1									Input tree 2									Input trees combined			
	clades								QS	Cat	clades								QS	Cat	QS	V
J	K	L	M	N	O	P	Q	R			S	T	U	V	W	X						
A	+	=	=	=	=	=	=	+0.5	s	-	-	-	-	-	-	-	-	-0.5	r	0	+1	1
B	-	=	-	-	-	-	-	-0.5	s	=	-	-	-	-	-	-	-	-0.5	s	-0.5	+1	2
C	=	=	=	+	=	=	=	+0.5	s	=	+	=	=	=	=	=	=	+0.5	s	+0.5	+1	2
D	-	-	-	-	-	=	-	-0.5	s	=	=	+	-	-	-	-	-	=0	s	-0.25	+1	2
E	-	-	-	-	-	-	=	-0.5	s	=	=	=	+	=	=	=	=	+0.5	s	0	+1	1.33
F	-	-	-	-	-	-	=	-0.5	s	=	=	=	=	+	=	=	=	+0.5	s	0	+1	1.33
G	=	=	=	=	=	=	+	+0.5	s	=	=	=	=	=	+	=	=	+0.5	s	+0.5	+1	1.33
H	-	-	-	-	-	-	-	-0.5	r	=	=	=	=	=	=	+	=	+0.5	s	0	+1	1
I	-	-	-	-	-	-	-	-0.5	r	=	=	=	=	=	=	=	+	+0.5	s	0	+1	1

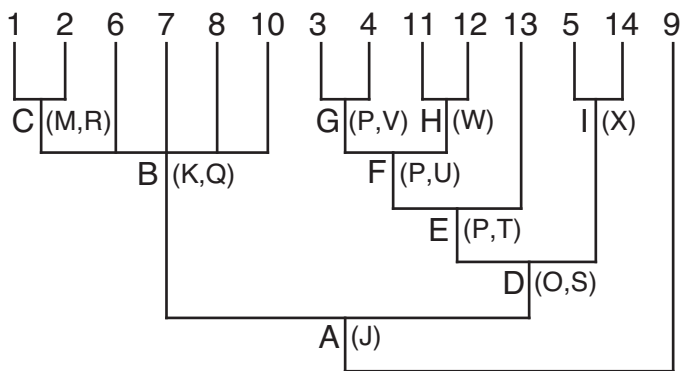
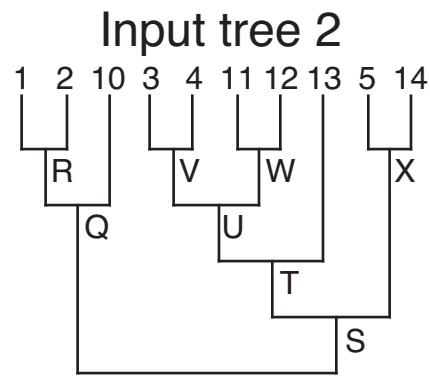
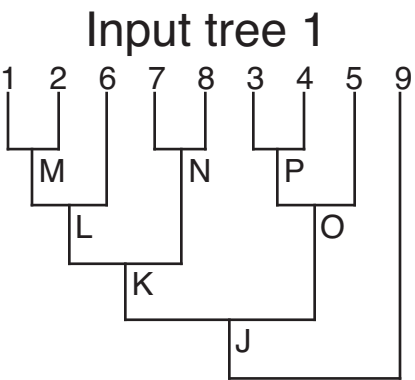
Table 3. Measures of clade support and occurrences of unsupported clades in nine MRP supertrees ranked by decreasing average value (V). The Dinosauria and Seabird studies are those of Pisani et al. (2002) and Kennedy and Page (2002, strict consensus), and the Lagomorpha that of Stoner et al., 2003 (with W indicating an analysis in which input trees were differentially weighted on the basis of an assessment of their robustness). All other are from Bininda-Emonds et al., (1999). I = no. of input trees, L = no. of leaves, C = coverage (average proportion of leaves in the input tree), SC = number of supertree clades. U = no. of unsupported supertree clades, U* = no. of unsupported supertree clades that conflict with at least one input tree, U** no. of unsupported clades conflicting with all relevant input trees. QS = average qualitative support for supertree clades. Figures in parentheses are ranges.

	I	L	C	SC	U	U*	U**	QS	V
Dinosauria	134	276	0.237 (0.014, 0.996)	208	2	2	0	0.009	0.756 (-1, 1)
Seabirds	7	122	0.254 (0.114, 0.738)	76	19	2	1	-0.201	0.571 (-1, 1)
Mustelidae	28	45	0.399 (0.067, 1)	31	0	0	0	-0.143	0.521 (-0.5, 1)
Lagomorpha	147	80	0.223 (0.038, 1)	57	1	1	0	-0.104	0.340 (-1, 1)
Canidae	36	34	0.408 (0.083, 1)	22	0	0	0	-0.146	0.259 (-0.455, 1)
Viverridae	9	34	0.618 (0.118, 1)	31	4	4	0	-0.045	0.253 (-1, 1)
Carnivora	62	12	0.548 (0.25, 1)	10	0	0	0	-0.029	0.199 (-0.556, 0.818)
Lagomorpha (W)	147	80	0.223 (0.038, 1)	76	2	2	0	-0.109	0.168 (-1, 1)
Felidae	40	36	0.494 (0.083, 1)	33	0	0	0	-0.219	0.022 (-0.789, 1)









Supertree