

Phylogenetics

Improving the efficiency of SPR moves in phylogenetic tree search methods based on maximum likelihood

Wim Hordijk and Olivier Gascuel*

Projet Méthodes et Algorithmes pour la Bioinformatique, LIRMM, UMR CNRS 5506, Université Montpellier 2, 161 rue Ada, 34392 Montpellier, France

Received on July 25, 2005; revised on October 7, 2005; accepted on October 9, 2005

Advance Access publication October 18, 2005

ABSTRACT

Motivation: Maximum likelihood (ML) methods have become very popular for constructing phylogenetic trees from sequence data. However, despite noticeable recent progress, with large and difficult datasets (e.g. multiple genes with conflicting signals) current ML programs still require huge computing time and can become trapped in bad local optima of the likelihood function. When this occurs, the resulting trees may still show some of the defects (e.g. long branch attraction) of starting trees obtained using fast distance or parsimony programs.

Methods: Subtree pruning and regrafting (SPR) topological rearrangements are usually sufficient to intensively search the tree space. Here, we propose two new methods to make SPR moves more efficient. The first method uses a fast distance-based approach to detect the least promising candidate SPR moves, which are then simply discarded. The second method locally estimates the change in likelihood for any remaining potential SPRs, as opposed to globally evaluating the entire tree for each possible move. These two methods are implemented in a new algorithm with a sophisticated filtering strategy, which efficiently selects potential SPRs and concentrates most of the likelihood computation on the promising moves.

Results: Experiments with real datasets comprising 35–250 taxa show that, while indeed greatly reducing the amount of computation, our approach provides likelihood values at least as good as those of the best-known ML methods so far and is very robust to poor starting trees. Furthermore, combining our new SPR algorithm with local moves such as PHYML's nearest neighbor interchanges, the time needed to find good solutions can sometimes be reduced even more.

Availability: Executables of our SPR program and the used datasets are available for download at <http://atgc.lirmm.fr/spr>

Contact: gascuel@lirmm.fr; wim@santafe.edu

1 INTRODUCTION

Maximum likelihood (ML) methods have become very popular for constructing phylogenies from sequence data. Felsenstein brought this framework to nucleotide-based phylogenetic inference (Felsenstein, 1981), and it was later also applied to amino acid

sequences (Kishino *et al.*, 1990). Several variants were proposed, including the widely used Bayesian methods (Rannala and Yang, 1996; Simon and Larget, 2000; Huelsenbeck and Ronquist, 2001). A number of computer studies (Kuhner and Felsenstein, 1994; Huelsenbeck, 1995; Rosenberg and Kumar, 2001; Ranwez and Gascuel, 2002; Guindon and Gascuel, 2003) have shown that ML programs can recover the correct tree from simulated datasets more frequently than other methods can, which supported numerous observations from real data and explains their popularity.

However, the disadvantage of ML methods is that they require much computational effort. Tree inference in the ML setting is computationally hard (Chor and Tuller, 2005), and all practical approaches rely on heuristics. The main idea behind those heuristics is that the space of possible tree topologies is searched for an optimal topology, optimizing edge lengths along the way. But even computing the optimal values of edge lengths on a single tree is not an easy task. This requires heavy numerical optimization techniques (reviewed in Bryant *et al.*, 2005), simply because of the number of parameters ($2n-3$ edges, where n is the number of taxa), and local optima are still possible (Chor *et al.*, 2000).

Despite these computational difficulties, ML methods have become faster and faster. We distinguish three main components in recent approaches:

(1) *Simultaneous optimization of tree topology and edge lengths.* The first hill climbing algorithms that were introduced iterate the following steps: (1) choose a neighboring topology of the current one; (2) optimize the edge lengths of this new topology to obtain its likelihood and (3) if this new topology is better than the current one, then it becomes the new current topology, else the current topology is unchanged. Owing to the size of the topology space and the computational intensity of edge length optimization, this approach becomes impractical for even moderate numbers of taxa and limited topological rearrangements. Stochastic approaches based on MCMC (Simon and Larget, 2000; Huelsenbeck and Ronquist, 2001), simulated annealing (Salter and Pearl, 2001; Stamatakis, 2005) or genetic optimization (Lewis, 1998; Lemmon and Milinkovitch, 2002) simultaneously change edge lengths and tree topology, while recent deterministic methods (Guindon and Gascuel, 2003; Vinh and von Haeseler, 2004; Stamatakis *et al.*,

*To whom correspondence should be addressed.

2005) estimate the likelihood of the new topology using approximate edge lengths, which provides lower bounds of its (fully optimized) likelihood. In all cases, likelihood approximations become more and more precise as the algorithm proceeds and as edge lengths become more accurate. This way expensive computations are avoided and the algorithm converges toward a local (or global in favorable cases) optimum of the likelihood function.

(2) *Local moves.* Nearest neighbor interchange (NNI) is the simplest topological move: it exchanges two subtrees that are connected by a single edge. Using NNIs, as implemented in PHYML (Guindon and Gascuel, 2003), is fast as the change in likelihood for each move can be calculated locally. Moreover, it is shown by Guindon and Gascuel (2003) that the likelihood improving NNI moves, although evaluated independently, can be performed all at once in most cases. Even though NNI moves do not allow the tree space to be searched intensively, this approach is usually sufficient to greatly improve starting trees, such as those obtained by fast distance-based or parsimony programs, and to get high topological accuracy and satisfactory likelihood values (Guindon and Gascuel, 2003).

(3) *Global moves.* Not surprisingly, NNIs can be trapped in (bad) local optima, and this seems to happen frequently with difficult datasets, e.g. those obtained from multiple genes with conflicting signals. In those cases, it occurs that some of the defects of the starting tree (e.g. long branch attraction) are still present in the resulting tree, even when this does not correspond to the likelihood optimum (H. Philippe and J. Pons, personal communication). In a subtree pruning and regrafting (SPR) move, a subtree is pruned and then regrafted at a different position in the remaining tree. With such moves, getting trapped in bad local optima can often be avoided, resulting in a better and more exhaustive search (Felsenstein, 1989; Swofford, 1996). However, SPR moves are much more expensive in terms of computation, as the likelihood has to be re-evaluated over the entire tree for each potential move. Recently, Stamatakis (2005) showed that the amount of computing time for performing an SPR-based search can be reduced by a combination of an efficient implementation of the actual likelihood computations, lazy subtree rearrangements (i.e. pre-evaluating potential moves by some fast likelihood estimate), and by carefully ordering potential SPRs, starting from the regraft positions that are close to the original pruned edge, before evaluating unlikely distant positions.

To improve current likelihood-based search methods even more, and most notably to make them less prone to starting tree defects, it is desirable to combine the advantages of both types of moves: the efficiency of local moves and the more elaborate search of global moves. We propose two methods to make SPR moves more efficient. The first method uses a fast distance approach based on the minimum evolution principle (Desper and Gascuel, 2002). Candidate SPR moves that are unpromising in terms of the minimum evolution criterion are simply discarded, which acts as a first filtering stage. The second method involves updating a limited number of partial likelihoods along the path between the prune and regraft positions [similar to what is done in RAxML (A. Stamatakis, personal communication)] and estimating edge lengths by distance-based analytical formulae; this enables a local approximate (but accurate) calculation of the change in likelihood for candidate SPR moves as opposed to a global (and expensive) computation. This approximate likelihood value is often sufficient to detect truly improving SPRs and, when none are found, is used to filter again

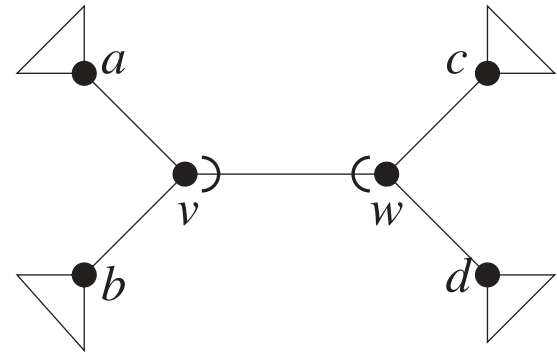


Fig. 1. An example of a phylogenetic tree. Triangles represent (possibly empty) subtrees, and the half-circles on the edge (*v*, *w*) represent the two partial likelihoods associated with this edge.

potential SPR moves before full likelihood calculations are performed. Our experiments with real datasets indicate that, while indeed greatly reducing the amount of computation, the results of this approach are at least as good as those of the best-known ML methods so far and show excellent robustness to poor starting trees.

The remainder of this paper is organized as follows. The next section will review the basic mathematical notation and equations involved in likelihood computations and the minimum evolution principle on phylogenetic trees, in particular in the context of SPR moves. Section 3 shows how the minimum evolution criterion can be used efficiently to filter out unpromising SPR moves. Section 4 then describes how updating only relevant partial likelihoods on the path between the prune and regraft positions and estimating relevant edge lengths suffice for enabling a fast but accurate estimation of the change in likelihood for an SPR move. Our complete SPR algorithm is presented in Section 5, and Section 6 shows the results of this algorithm, including the achieved reduction in computation time. Finally, Section 7 summarizes the main results and conclusions.

2 NOTATION AND MATHEMATICAL CONCEPTS

Consider a phylogenetic tree as shown in Figure 1. The triangles attached to nodes *a*, *b*, *c* and *d* represent (possibly empty) subtrees. The half-circles on the edge (*v*, *w*) connecting nodes *v* and *w* represent partial likelihoods. For example, the half-circle next to node *v* represents the partial likelihood Lk_p of the subtree rooted at node *v* with child nodes *a* and *b*. The likelihood $Lk(T, i)$ of the complete tree *T* at site *i* can be calculated locally on the edge (*v*, *w*) given these partial likelihoods and the edge length:

$$Lk(T, i) = \sum_{x, y \in \sigma} \pi_x Lk_p(i(v) = x) Lk_p(i(w) = y) P_{xy}(l), \quad (1)$$

where $\sigma = \{A, C, G, T\}$, π_x is the a priori probability of observing nucleotide *x*, $i(v)$ is the state (nucleotide) at site *i* at node (taxon) *v* and $P_{xy}(l)$ is the probability of a substitution from nucleotide *x* to *y* given a time period *l*, which is the length $d(v, w)$ of edge (*v*, *w*). The tree likelihood is then equal to the product of the likelihoods of the sites, as induced by the site independence assumption (Felsenstein, 1981):

$$Lk(T) = \prod_i Lk(T, i). \quad (2)$$

Note that Equation (1) is given for nucleotide sequences, but of course holds for protein sequences as well by replacing the

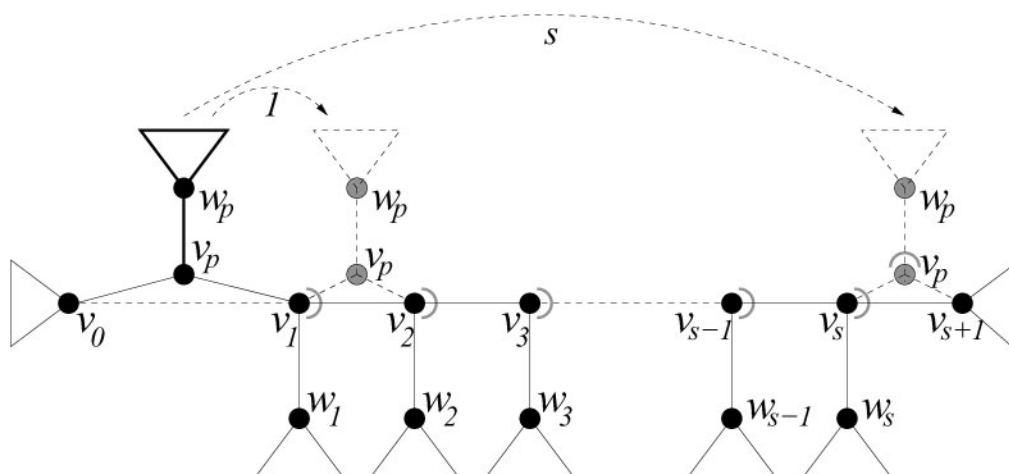


Fig. 2. An example of an SPR move of distance 1 (equivalent to an NNI move) and of distance s in general. The subtree indicated in bold is pruned (at node v_p), and regrafted between nodes v_s and v_{s+1} (dashed subtrees). The nodes v_0 and v_1 (at the prune position) will now be connected, indicated with a dashed edge.

alphabet σ over which the sum is taken. This equation is also easily adapted to incorporate site-to-site variation using a discrete rate (e.g. gamma) distribution. The full likelihood of a given site is then obtained by summing, over the rate categories, the likelihoods of the site according to each rate, weighted by the probability of each rate category (Yang, 1994). Finally, every edge defines two partial likelihoods (one in each direction), and all partial likelihoods within the tree can be calculated and updated efficiently (in linear time per site) by performing a pre-order [analogous to Felsenstein's pruning algorithm (Felsenstein, 1981)] and then a post-order depth-first traversal (Berry and Gascuel, 2000) of the tree, starting from the leaf partial likelihoods that are equal to 0 or 1 depending on the site value for the given taxon [see Bryant *et al.* (2005) for more on likelihood calculations in phylogenetics]. Note that these partial likelihood calculations (based on double tree traversals), along with the data structure of Figure 1 (which involves two partial likelihoods per edge), are important implementation features of MOLPHY (Adachi and Hasegawa, 1996) and PHYML (and possibly of other programs, but implementation descriptions are usually sparse). The most standard approach [e.g. FastDNAML (Olsen *et al.*, 1994) and RAXML] involves rooting the tree and having partial likelihoods only for the subtrees not containing the root, i.e. one per edge. The standard approach is then more economical in terms of memory space; but the likelihood of the full tree can be computed with the tree root only, using classical Felsenstein's pruning (Felsenstein, 1981). The 'double' approach requires twice as much memory, and partial likelihood calculations are twice as long compared with running pre-order search only, but after these calculations the tree likelihood can be computed locally and quickly on any branch of the tree using Equations (1) and (2). This makes branch length optimization and NNI local moves much faster, and we shall see that this also accelerates SPR computations.

There are various ways of estimating or optimizing the edge lengths of a given tree. The distance between nodes v and w can be optimized accurately using the partial likelihoods $Lk_p(i(v) = x)$ and $Lk_p(i(w) = y)$ and Equations (1) and (2) above so that the overall likelihood $Lk(T)$ is maximal [an iterative optimization method such as, e.g. Newton-Raphson or Brent (1973) can be used for this]. Another, much faster but less optimal, way (providing a lower

bound of tree likelihood) of estimating $d(v, w)$ is based on the average subtree distance [reviewed in Desper and Gascuel (2005)]. The average subtree distance $\Delta_{v|w}$ between two (non-overlapping) subtrees rooted at nodes v and w , respectively, is the average distance between all taxa in the first subtree and all taxa in the second subtree, and is recursively defined as (referring to Fig. 1):

$$\Delta_{v|w} = \frac{1}{2}(\Delta_{v|c} + \Delta_{v|d}) = \frac{1}{2}(\Delta_{a|w} + \Delta_{b|w}), \quad (3)$$

where $\Delta_{v|w}$ is the actual distance if both v and w are taxa. Note that this definition is the balanced version of average subtree distance, where each subtree has equal weight regardless of the number of taxa it contains. Given the matrix of pairwise evolutionary distances (as obtained by any model, e.g. K2P or HKY), the average subtree distance between all pairs of non-overlapping subtrees can be calculated in $O(n^2)$ time, where n is the total number of taxa (Desper and Gascuel, 2002). The length of edge (v, w) can now be estimated as

$$d(v, w) = \Delta_{v|w} - \frac{1}{2}\Delta_{a|b} - \frac{1}{2}\Delta_{c|d}. \quad (4)$$

An example of an SPR move of (topological) distance s is shown in Figure 2. The subtree indicated in bold is pruned at node v_p , and nodes v_0 and v_1 will now be connected by a new edge (indicated by the dashed line between them), while the edges originally connecting node v_p with nodes v_0 and v_1 will be removed. Next, the pruned subtree is regrafted at a distance s from its original position, between nodes v_s and v_{s+1} . This new situation is indicated with the dashed subtrees (one for $s = 1$ and one for general values of s). The original edge between nodes v_s and v_{s+1} is removed, and two new edges connecting node v_p with nodes v_s and v_{s+1} are inserted. An SPR move of distance $s = 1$ is, in fact, equivalent to an NNI move.

Since an SPR move changes the topology, and thus the relative subtree positions, many average subtree distances as calculated before the SPR move are not correct anymore. However, in most cases they can be updated in constant time to reflect the new topology, without having to recalculate them all (which would take quadratic time). In particular, two average subtree distances at the

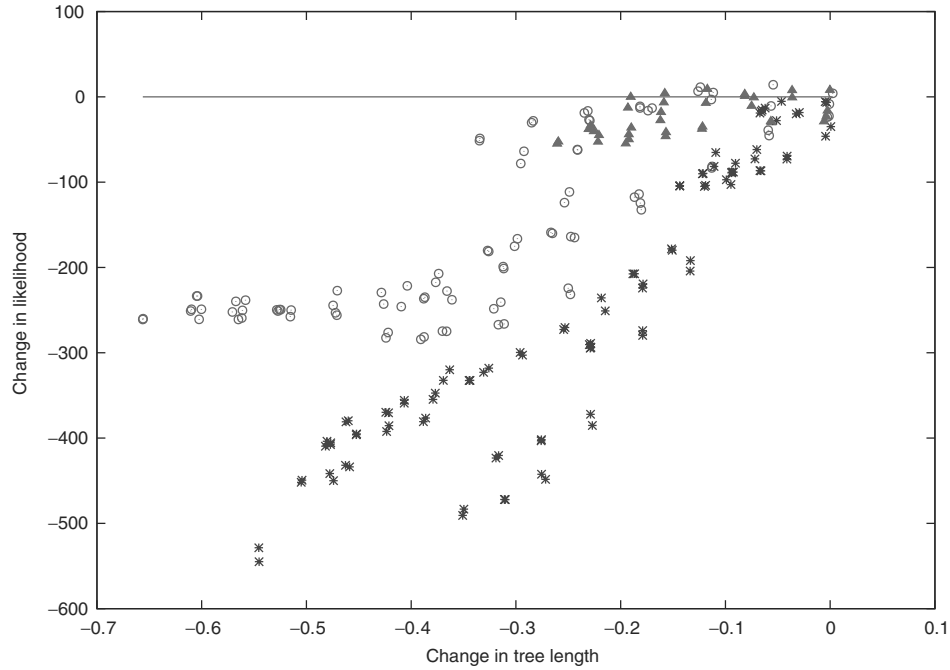


Fig. 3. Correlation between change in tree length and change in log likelihood. Each symbol represents data for one particular pruned subtree and all its candidate regraft positions.

regraft position that will be useful later on, $\Delta_{v_s|w_p}$ and $\Delta_{v_{s-1}|w_s}$, can be updated recursively as follows:

$$\Delta_{v_s|w_p}^* = \frac{1}{2}(\Delta_{v_{s-1}|w_p}^* + \Delta_{w_s|w_p}),$$

with

$$\Delta_{v_0|w_p}^* = \Delta_{v_0|w_p}, \quad (5)$$

and

$$\Delta_{v_{s-1}|w_s}^* = \Delta_{v_{s-1}|w_s} - \left(\frac{1}{2}\right)^s \Delta_{w_p|w_s} + \left(\frac{1}{2}\right)^s \Delta_{v_0|w_s}.$$

Since $\Delta_{v_s|w_p}^*$ is defined recursively in s , in our SPR algorithm (presented in full in Section 5) we consider possible regraft positions, given a pruned subtree, recursively in (increasing) distance s .

3 SPR MOVES AND TREE LENGTH

A criterion in phylogenetic inference that is conceptually related to parsimony, and based on average subtree distances, is the minimum evolution principle. Several variants of this principle have been proposed, but here we follow the definition and notation of Desper and Gascuel (2002, 2005). In particular, we use the concept of tree length $L(T)$, which is defined as the sum of the edge lengths of a tree T , and consider the balanced version, first introduced by Pauplin (2000). The minimum evolution tree is then that tree T which minimizes $L(T)$.

Consider again an SPR move of distance s (as in Fig. 2). The change in tree length dL_1 resulting from an SPR move of distance $s = 1$ (i.e. an NNI move), is equal to [Desper and Gascuel (2002), Equation 12]:

$$dL_1 = \frac{1}{4}[(\Delta_{v_0|w_p} + \Delta_{v_2|w_1}) - (\Delta_{v_0|w_1} + \Delta_{v_2|w_p})].$$

From this, a recursive formula dL_s for the change in tree length for an SPR move of distance s can be derived:

$$dL_s = dL_{s-1} + \frac{1}{4}[(\Delta_{v_{s-1}|w_p}^* + \Delta_{v_{s+1}|w_s}) - (\Delta_{v_{s-1}|w_s}^* + \Delta_{v_{s+1}|w_p})],$$

where the Δ^* s are as defined in equations (5).

As it turns out, there is a strong correlation between the change in tree length and the change in likelihood for a given SPR move. Figure 3 shows three typical cases. Each collection of dots of one particular symbol represents data for one particular pruned subtree and all its possible regraft positions (for some given tree). The data shown here (using a real dataset of 55 taxa) are representative for this type of data in general, i.e. when considering other trees or datasets as well. Along the horizontal axis the change in tree length (dL_s) is shown, while the vertical axis shows the change in log likelihood. Clearly, the regraft positions with the highest change in likelihood for a given pruned subtree also have the highest change in tree length values and vice versa. Of course, in most cases an arbitrary SPR move will not improve the likelihood of a tree, as for example the collection of stars shows that all possible regraft positions result in a reduction in likelihood (negative change). However, if given a pruned subtree there are regraft positions that result in an improvement in likelihood (as with the circles and triangles), they always seem to be among the best ones in terms of change in tree length.

Based on this observation, computing changes in likelihood can be avoided altogether for the majority of potential regraft positions. For a given pruned subtree, first the change in tree length dL_s can be calculated quickly for all potential regraft positions (recursively in distance s), and the actual change in likelihood computations are

then done only for the most promising ones. This constitutes our first efficiency improving method.

4 SPR MOVES AND LIKELIHOOD ESTIMATION

One of the main advantages of local moves like NNI is that the tree topology is changed only locally, and the likelihood of the new tree can be calculated efficiently based on Equations (1) and (2) (see Guindon and Gascuel, 2003, for details). With SPR moves, on the contrary, the tree topology is often changed significantly, and generally the whole tree has to be re-evaluated (including most or all partial likelihoods). However, to estimate the change in likelihood for a (potential) SPR move to see if it would result in an improved tree, it actually suffices to only update a limited number of relevant partial likelihoods.

To calculate the likelihood of the new tree locally at edge (v_p, w_p) [using Equations (1) and (2)] after an SPR move, the partial likelihood $Lk_p(i(v_p) = x)$ [represented by the gray half-circle on the edge (v_p, w_p) in Fig. 2] needs to be updated to reflect the new tree topology [note that $Lk_p(i(w_p) = x)$ is not changed by the SPR move]. To update $Lk_p(i(v_p) = x)$, however, $Lk_p(i(v_s) = x)$ [represented by the gray half-circle on the edge (v_s, v_{s+1})] needs to be updated, and so on, all the way back to node v_1 :

$$Lk_p(i(v_k) = x) = \left(\sum_{y \in \sigma} Lk_p(i(v_{k-1}) = y) P_{xy}(d(v_{k-1}, v_k)) \right) \times \left(\sum_{z \in \sigma} Lk_p(i(w_k) = z) P_{xz}(d(v_k, w_k)) \right). \quad (6)$$

Moreover, w_k is replaced by v_{s+1} and v_k becomes v_s when $k = p$. So, by updating the partial likelihoods on the path between the prune and regraft positions (as indicated by the gray half-circles in Fig. 2), the change in likelihood induced by the SPR move can now be calculated locally. Consequently, for an SPR move of distance s , only $s + 1$ partial likelihoods need to be updated, instead of having to re-evaluate the entire tree. RAXML actually uses a similar approach (Stamatakis, 2004, 2005) but in a rooted version, which implies that more partial likelihoods might have to be computed depending on the root position. As discussed earlier, RAXML stores fewer partial likelihoods (one per edge, instead of two), which is more economical in terms of memory requirements but likely to be slower in terms of CPU time. Moreover, this method of updating partial likelihoods along the path between the prune and regraft positions has (to our knowledge) not yet been described fully and explicitly anywhere else.

Since updating partial likelihoods and estimating changes in tree likelihood depend on edge lengths, accurate estimates for the lengths of the relevant edges at the prune and regraft positions are also needed (while all other edge lengths are kept constant, at least during the first selection stages, see Section 5). An iterative optimization method (as mentioned in Section 2) can be used for this, but this is computationally expensive, especially if it has to be performed for each candidate regraft position. Therefore, we use estimates based on the average subtree distances as explained in Section 2. Although slightly less accurate than the iterative optimization estimates, they can be calculated much faster.

First consider the new edge connecting nodes v_0 and v_1 at the prune position. A naive and straightforward estimation for the length $d(v_0, v_1)$ of this edge is to simply take the sum of the lengths

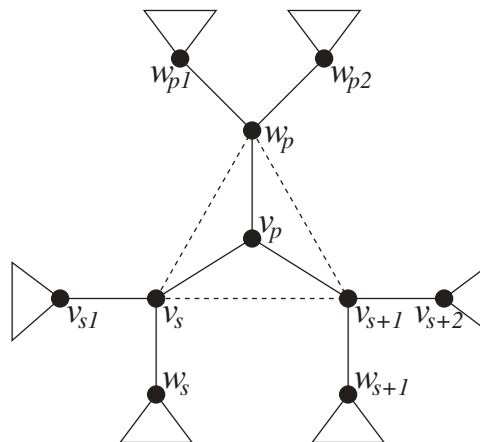


Fig. 4. Local topology at the regraft position of an SPR move of distance s .

of the original edges (v_p, v_0) and (v_p, v_1) before pruning, i.e.

$$d(v_0, v_1) = d(v_p, v_0) + d(v_p, v_1).$$

An estimate based on the average subtree distances is obtained as follows (neglecting taxa within the pruned subtree):

$$d(v_0, v_1) = \Delta_{v_0|v_1} - \frac{1}{2} \Delta_{v_{0a}|v_{0b}} - \frac{1}{2} \Delta_{w_1|v_2},$$

where v_{0a} and v_{0b} are the subtrees rooted at the children of v_0 , unless v_0 is a taxon (leaf node), in which case $\Delta_{v_{0a}|v_{0b}} = 0$. Compared with the (more accurate) iterative optimization method, one of these values usually gives an overestimate, while the other usually gives an underestimate. So, as a ‘compromise’, we actually take the average of the two values to get a fairly accurate yet easy to calculate estimate for the length of edge (v_0, v_1) .

Next consider the regraft position. The relevant nodes and edges are shown in Figure 4. The edges for which the lengths need to be estimated are (v_p, v_s) , (v_p, v_{s+1}) and (v_p, w_p) . For $d(v_p, v_s)$ and $d(v_p, v_{s+1})$ we can compute a simple and naive estimate (similar to the one used in RAXML) by taking half of the length of the original edge (v_s, v_{s+1}) (indicated by the dashed line between v_s and v_{s+1}). The length of edge (v_p, w_p) can then be estimated as

$$d(v_p, w_p) = d(v_{s+1}, w_p) - d(v_p, v_{s+1}).$$

The calculation of $d(v_{s+1}, w_p)$ (indicated by the dashed line between v_{s+1} and w_p) is given below.

Estimates based on the average subtree distances are slightly more complicated than at the prune position, but can be calculated from estimates for the lengths of the dashed ‘edges’ in Figure 4, i.e. $d(v_s, w_p)$, $d(v_{s+1}, w_p)$ and $d(v_s, v_{s+1})$. For the third one, $d(v_s, v_{s+1})$, we can take the length of the original edge (v_s, v_{s+1}) . The distance between nodes v_{s+1} and w_p can be directly estimated as

$$d(v_{s+1}, w_p) = \Delta_{v_{s+1}|w_p} - \frac{1}{2} \Delta_{v_{s+2}|w_{s+1}} - \frac{1}{2} \Delta_{w_{p1}|w_{p2}}.$$

Finally, the distance between nodes v_s and w_p can be calculated similarly as

$$d(v_s, w_p) = \Delta_{v_s|w_p}^* - \frac{1}{2} \Delta_{v_{s-1}|w_s}^* - \frac{1}{2} \Delta_{w_{p1}|w_{p2}},$$

where the Δ^* s are again as defined in Section 2, and $\Delta_{w_{p1}|w_{p2}} = 0$ if w_p is a leaf.

Once these three distances are calculated, the required edge length estimates are simply

$$\begin{aligned}d(v_p, w_p) &= \frac{1}{2} [+d(v_s, w_p) + d(v_{s+1}, w_p) - d(v_s, v_{s+1})], \\d(v_p, v_s) &= \frac{1}{2} [+d(v_s, w_p) - d(v_{s+1}, w_p) + d(v_s, v_{s+1})], \\d(v_p, v_{s+1}) &= \frac{1}{2} [-d(v_s, w_p) + d(v_{s+1}, w_p) + d(v_s, v_{s+1})].\end{aligned}$$

At the regraft position, as at the prune position, the naive estimates and the average subtree distance-based estimates also tend to give over- and underestimates compared with the more accurate iterative edge length optimization. So here too, we take the average of both estimates for each of the three relevant edges.

With these edge length estimates, and the updated partial likelihoods along the path between the prune and regraft position, an estimate for the change in likelihood of a given SPR move can be obtained very quickly. This results in an underestimate, but in practice it turns out that it is usually not too far off from the actual change in likelihood (when edge lengths are fully optimized). Moreover, this estimation method avoids a significant amount of computation and constitutes our second efficiency improving method.

5 THE SPR ALGORITHM

The main idea of our SPR algorithm is to consider candidate SPR moves, accepting those that result in an improvement in likelihood, but using the methods introduced in the previous sections to discard unpromising moves and to reduce the amount of computation necessary for those moves that are not already filtered out at the first stage. Briefly, each subtree t in the current tree T is considered for pruning in turn. Given a pruned subtree t , all potential regraft positions for t are considered recursively in increasing distance s . The change in tree length dL_s is calculated for each candidate position (as explained in Section 3), and a list `rgrft_cand` of the `N_RGRFT` best ones is maintained and updated along the way.

Next, for each candidate in the `rgrft_cand` list (starting with the best one in terms of the dL_s value), the relevant edge lengths are estimated (as explained in the previous section) and the partial likelihoods along the path from the prune to the regraft positions are updated. The change in likelihood is then estimated [locally using Equations (1) and (2)], and as soon as an improvement is found, this candidate regraft position is accepted and the remaining ones are discarded. This procedure is then repeated for the next subtree t (but now on a possibly improved and updated tree T). So, for each next pruned subtree t a new list `rgrft_cand` of regraft position candidates is constructed.

During the estimation of the change in likelihood for the various candidate SPR moves, a second list `optim_cand` is maintained and updated with the `N_OPTIM` best candidate moves in terms of these changes in likelihood. Thus, this second list is constructed over all possible pruned subtrees t and their respective potential regraft positions. If improvements were already found during the previous stage (i.e. calculating dL_s values and estimating changes in likelihood), then the algorithm will exit and return `true`, indicating that the tree was improved, and this second list is ignored. However, if no improvement was found at all during the previous stage, the

`optim_cand` list (which now obviously contains only candidates with negative changes in likelihood) is considered in the following way. For each candidate in the list (starting with the best, although negative, one), perform the SPR move as indicated by this candidate, use an iterative method (Brent optimization in our implementation) to optimize the lengths of the three relevant edges at the regraft position and calculate the new likelihood of the tree. If this results in an improvement, accept the move, discard the remaining candidates, exit and return `true`. Otherwise, try the next candidate in the list. The idea behind this procedure is that an iterative optimization method, although computationally expensive, usually gives slightly better estimates for the relevant edge lengths than the fast average subtree distance-based estimates. It can, and does, happen that with these more accurate edge length estimates the change in likelihood of one of these `N_OPTIM` candidates does indeed become positive, and thus give rise to an improved tree after all.

Finally, if this ‘local optimization’ procedure does not result in any improvements either, the best `N_GLOBL` candidates from the `optim_cand` list are considered for ‘global optimization’. For each candidate in the list (again starting with the best one based on the local optimizations from the previous stage), perform the SPR move as indicated by this candidate, use an iterative method to optimize the lengths of all the edges in the tree and calculate the new likelihood of the tree. If this results in an improvement, accept the move, discard the remaining candidates, exit and return `true`. Otherwise, try the next candidate in the list. Here, again, the idea is that an improvement might be found by doing an even more elaborate and complete edge length optimization. However, since this global edge length optimization is computationally very expensive, it is only done for the `N_GLOBL` best candidates, where usually `N_GLOBL` \ll `N_OPTIM`. If this last procedure does not result in an improvement, then the algorithm will exit and return `false`, indicating no improvements could be found. So, at each next stage of the algorithm, fewer candidates are considered, while the corresponding change in likelihood estimations become more accurate but also require more intensive computation.

To perform an actual tree search, the algorithm can be performed repeatedly until it returns `false`. Or, alternatively, it can be combined with a local search method such as NNI, where rounds of NNI moves and SPR moves are alternated, for example. The complete SPR algorithm as described above is presented more formally next.

SPR ()

- (1) Calculate the average subtree distances for all pairs of (non-overlapping) subtrees.
- (2) For each next subtree t do:
 - (a) Clear the `rgrft_cand` list.
 - (b) Prune t and estimate $d(v_0, v_1)$.
 - (c) Recursively consider potential regraft positions at increasing distance s . At each recursion step do:
 - (i) Calculate the relevant Δ^* values and compute dL_s .
 - (ii) If dL_s is within the `N_RGRFT` best ones so far, store the current regraft position together with its dL_s value, the current path from the prune position, and the Δ^* values in the `rgrft_cand` list, sorted in decreasing dL_s values.

- (d) For each candidate in the `rgrft_cand` list (starting at the first, or best, one) do:
 - (i) Update the partial likelihoods Lk_p along the path between the prune and candidate regraft positions.
 - (ii) Estimate edge lengths at the candidate regraft position using the average subtree distances and stored Δ^* values.
 - (iii) Regraft t at the candidate regraft position, set the relevant edge lengths, update the $Lk_p(i(v_p) = x)$ s using Equation (6), and calculate the likelihood $Lk(T)$ using Equations (1) and (2).
 - (iv) If the new $Lk(T)$ value is an improvement, accept the current move, update the complete tree (partial likelihoods, edge lengths, average subtree distances, etc.), discard the remaining candidate regraft positions and go back to Step 2 to try the next subtree.
 - (v) If no improvement, undo the regraft operation and reset the edge lengths and partial likelihoods. If the change in likelihood (although negative) is within the `N_OPTIM` best ones so far, store the current prune and regraft positions together with other relevant data (path, Δ^* s, etc.) in the `optim_cand` list, sorted in decreasing order.
 - (vi) Go to Step 2(d) and try the next candidate.
- (e) Go back to Step 2 to try the next subtree.
- (3) If an improvement was found in Step 2, exit and return `true`.
- (4) Otherwise, for each candidate in the `optim_cand` list (starting at the first, or best, one) do:
 - (a) Perform the SPR move as indicated by the current candidate and update the partial likelihoods along the path between the prune and regraft positions.
 - (b) Optimize the relevant edge lengths at the regraft position using an iterative optimization method.
 - (c) Calculate the new likelihood and if it results in an improvement, accept the current move, update the complete tree, exit and return `true`.
 - (d) If no improvement, update the likelihood of the current move in the `optim_cand` list, undo the move and go back to Step 4 to try the next candidate.
- (5) For the best `N_GLOBL` candidates in the `optim_cand` list do:
 - (a) Perform the SPR move as indicated by the current candidate.
 - (b) Update all partial likelihoods and perform global edge length optimization using an iterative optimization method.
 - (c) If an improvement in likelihood results, accept the current move, exit and return `true`.
 - (d) If no improvement, undo the move and go back to Step 5 to try the next candidate.
- (6) Exit and return `false`.

Note that there are various parameters in the SPR algorithm, in particular `N_RGRFT`, `N_OPTIM` and `N_GLOBL`. Of course the

performance of the algorithm depends largely on the chosen values for these parameters. Currently it is not obvious what the best values are, as this in turn depends on the datasets considered, and they will have to be found by trial and error. The next section will show the best results we have obtained so far, but more investigation needs to be done in terms of finding good parameter values, or better yet, a systematic way to set them depending on the given data.

6 RESULTS

In this section, we show results of applying our SPR algorithm to some real datasets from the benchmark set used by Stamatakis *et al.* (2005). In particular, datasets with a number of taxa of 101 (`101_SC`), 150 (`150_ARB`) and 250 (`250_ARB`) were considered. Additionally, we considered the four datasets that were provided with publications in a recent issue of *Systematic Biology*. These datasets contain 132 (Vogler *et al.*, 2005), 42 (Yuan *et al.*, 2005), 39 (McCracken and Sorenson, 2005), and 35 (Winkworeth *et al.*, 2005) taxa, respectively.

Our SPR algorithm has been implemented in the PHYML program (Guindon and Gascuel, 2003), in such a way that it can be used either by itself or in combination with the NNI algorithm provided by PHYML. In the latter case, the SPR algorithm is called once each time the NNI moves are stuck on a local optimum. If the SPR algorithm is able to improve the tree and get out of the local optimum, the NNI algorithm is applied again until it is trapped in another local optimum, etc. If the SPR algorithm cannot find an improvement anymore either, the program terminates.

We applied both versions (SPR and PHYML+SPR) on the mentioned datasets and used both random trees and maximum parsimony trees as starting points. Random trees are useful to check that the algorithm is not affected by potentially poor starting trees, while starting with parsimony trees corresponds to standard use, notably regarding run time. The maximum parsimony trees were created using RAxML (Stamatakis *et al.*, 2005), and the results of our algorithms are compared with this program. The HKY model of nucleotide substitution was used in all cases, and the transition/transversion parameter is estimated by all programs during the search.

Table 1 shows the likelihood values found by the different programs. The results for random trees and parsimony trees are averages over 10 trees each. The rows labeled ‘best’ shows the best likelihood score found by each program over all 20 starting trees (10 random + 10 parsimony). Since there is a difference in the way likelihood values are calculated between PHYML and RAxML, all final trees found by RAxML were re-evaluated using PHYML (without doing any further optimization, but just evaluating the likelihood of the given tree) to enable a direct comparison.

As the table shows, there is not one program that is always the best, but they are all able to find the same best solution (although not necessarily on the same starting tree). The most striking difference is that SPR and PHYML+SPR are more robust with poor (random) starting trees than RAxML and should thus be more resistant to artifacts (e.g. long branch attraction) that sometimes affect the fast (distance-based or parsimony) methods being used to build initial trees. For example, with 250 taxa SPR and PHYML+SPR are on average about 400 log-likelihood points better than RAxML, which is considerable. With less taxa, the difference is lower but still quite

Table 1. Search results of the different programs on the given datasets

#Taxa	Start	SPR	PHYML+SPR	RAxML
35	rnd/prs	-1084/-1084	-1084/-1084	-1087/-1084
	best	-1084	-1084	-1084
39	rnd/prs	-2947/-2946	-2946/-2946	-2953/-2947
	best	-2945	-2945	-2945
42	rnd/prs	-5732/-5734	-5734/-5734	-5743/-5734
	best	-5731	-5731	-5731
101	rnd/prs	-73 869/-73 874	-73 878/-73 869	-73 926/-73 870
	best	-73 862	-73 862	-73 862
132	rnd/prs	-50 496/-50 539	-50 502/-50 539	-50 529/-50 530
	best	-50 477	-50 481	-50 477
150	rnd/prs	-76 856/-76 856	-76 855/-76 858	-77 161/-76 853
	best	-76 849	-76 849	-76 849
250	rnd/prs	-130 920/-130 940	-130 902/-130 920	-131 325/-130 897
	best	-130 894	-130 891	-130 889

Results for random (rnd) and parsimony (prs) trees are averages over 10 trees each. The 'best' score is the best likelihood value found over all (20) starting trees.

Table 2. Average running times of the three programs on the random (rnd) and parsimony (prs) starting trees

#Taxa	Start	SPR	PHYML+SPR	RAxML
35	rnd/prs	0:02/0:02	0:02/0:01	0:04/0:04
39	rnd/prs	0:06/0:05	1:03/0:04	0:12/0:10
42	rnd/prs	0:13/0:10	0:16/0:06	0:15/0:11
101	rnd/prs	7:33/6:45	9:50/6:00	23:52/15:30
132	rnd/prs	14:39/16:00	10:19/5:13	20:56/9:31
150	rnd/prs	15:29/12:50	21:00/13:00	36:00/14:30
250	rnd/prs	2:10:00/2:00:00	1:25:00/1:21:00	2:16:00/0:51:44

significant. Surprisingly, SPR and PHYML+SPR even seem to perform somewhat better (on average) on random trees than on parsimony trees, but more results are needed to confirm this tendency.

Table 2 shows the running times of the different programs on the random and parsimony trees (again averaged over the 10 trees). The runs were performed on an unloaded linux PC with an Intel Pentium 4 (3.0 GHz) processor and 512 MB of RAM. As the table shows, the running times are of the same order of magnitude for the different programs. However, a direct comparison of running times between these two programs (although useful for practical purposes) does not give much insight into the improvements in efficiency that is achieved by our proposed two methods. For example, the SPR algorithm in RAxML has a very efficient but also very specific (nucleotide data only) implementation for storing trees and calculating likelihoods (Stamatakis *et al.*, 2002 and A. Stamatakis, personal communication), whereas PHYML (and thus our SPR implementation) is more generic (both nucleotide and protein data) and consequently less optimized for speed. So, the efficiency of RAxML comes to a large extent from implementation optimizations, whereas here we focus more on algorithmic aspects. In the end, of course, it would pay off to combine both approaches.

Considering the results as presented in Tables 1 and 2, it can be concluded that the search performance of our algorithm is at least comparable with that of RAxML [which produced the best-known trees to date on a benchmark set of real alignments (Stamatakis *et al.*, 2005)] with parsimony starting trees, and better with (poor) random starting trees. However, the results do depend on the parameter settings that were used. The results shown here were the best ones obtained after trying a (limited) number of different parameter settings. In all cases we used $N_{OPTIM} = 100$, and for N_{RGRFT} and N_{GLOBL} it seems that values of about 15–20% and 10% (respectively) of the number of edges in the tree give the best results. In practice, we also used a cutoff MAX_DIST for the potential regraft positions considered: regraft positions more than a distance MAX_DIST from the prune position were not evaluated at all. This parameter (also used by RAxML, although slightly differently) provides another way of reducing the total number of likelihood computations performed by the algorithm. In our experiments this parameter was set to 10% of the number of edges, which is a relatively high value regarding the diameter of phylogenies (Gascuel, 2000) and resulted in a low influence on the output trees. However, as already mentioned, more tests need to be done to find optimal parameter settings, or some systematic way of setting these values.

Having verified that our SPR algorithm can live up to the best-known performance so far, of course the main question is how much the efficiency of the SPR search has been improved by our two introduced methods. Table 3 shows to what extent likelihood computations have been avoided and reduced in our experiments. The second column in this table shows the (average) number of dL_s (change in tree length) computations performed during a search of our SPR algorithm, averaged over the 10 parsimony starting trees (results on random trees are similar). The third column shows the number of actual change in likelihood estimations that were computed by only considering the N_{RGRFT} best candidate regraft positions for each pruned subtree. The fourth column shows the number of times local edge length optimization (at the regraft position) has been performed if no improvements could be found with the analytical estimates. Finally, the fifth column shows the

Table 3. The reduction in likelihood computations achieved by our algorithm

# Taxa	Tree length ^a	Likelihood ^b	Local ^c	Global ^d
35	22 776	9878	291	6
39	28 382	12 619	291	10
42	28 185	8932	237	10
101	402 676	92 989	256	22
132	1 100 011	242 133	792	45
150	844 105	142 422	279	31
250	4 519 332	979 891	662	63

The columns show the number of candidate SPR moves.

^aThe change in tree length is calculated.

^bThe change in likelihood is estimated.

^cLocal edge length optimization (at the regraft position) is performed to estimate the change in likelihood.

^dGlobal edge length optimization is performed.

number of times global edge length optimization was performed if no improvements in likelihood could be found with local optimization. As the table clearly shows, our progressive filtering method has an enormous impact.

To measure this impact on the actual running time of the SPR algorithm, we compared the following variants of the algorithm on two datasets of 55 and 101 taxa, respectively:

- (1) The ‘regular’ algorithm as presented in Section 5, with $N_{\text{RGRFT}}=15$ (~15% of the number of edges) for the 55-taxon dataset, and $N_{\text{RGRFT}}=40$ (20% of the number of edges) for the 101-taxon dataset.
- (2) As variant 1, but with $N_{\text{RGRFT}}=100\%$ of the number of edges for both datasets. In other words, this variant estimates the change in likelihood for all candidate SPR moves considered.
- (3) As variant 1, but performing global likelihood computations (i.e. updating and re-evaluating the entire tree) instead of updating only relevant partial likelihoods and performing a local likelihood computation.
- (4) A combination of variants 2 and 3, i.e. setting $N_{\text{RGRFT}}=100\%$ and performing global likelihood computations.

As starting trees we used the BIONJ tree as computed by PHYML for the 55-taxon dataset and one of the 10 parsimony trees used above for the 101-taxon dataset. All four variants of the algorithm found the best-known likelihood values for these datasets (−24583 and −73862, respectively), and the times necessary to find these solutions are presented in Table 4.

As the table clearly shows, the reduction in runtime is very significant for both methods, while the algorithm is still able to find the best-known solutions. Avoiding likelihood computations by only estimating likelihood changes for the most promising candidate moves based on the change in tree length (comparing variant 3 with variant 4) results in a decrease in runtime by a factor of about 3. Reducing the amount of computation by estimating the change in likelihood locally after updating only relevant partial likelihoods and quickly estimating edge lengths (comparing variant 2 with variant 4) improves the runtime by a factor of about 10–20.

Table 4. The running times for the four variants of the SPR algorithm on two datasets of 55 and 101 taxa, respectively

# Taxa	55		101	
Variant	Running time	Factor	Running time	Factor
1	0:00:26	36.0	0:07:42	60.3
2	0:01:13	12.8	0:30:13	21.6
3	0:05:06	3.1	2:46:15	2.8
4	0:15:36	1.0	7:44:21	1.0

Combining the two methods together (comparing variant 1 with variant 4) gives a factor of 36 and 60, respectively, in runtime improvement. Finally, combining the SPR algorithm with local (NNI) moves can sometimes provide even more of a speedup. For example, on the 101-taxon dataset and the same parsimony starting tree, the combination of PHYML+SPR found the best solution in just under 5 min.

7 CONCLUSIONS AND FURTHER WORK

We have introduced two methods that can be used to drastically reduce or even avoid expensive likelihood computations in SPR-based search algorithms on phylogenetic trees using ML. The first method avoids likelihood computations altogether by only calculating changes in likelihood for the most promising SPR moves in terms of the change in tree length, a statistic based on the minimum evolution principle that can be calculated efficiently. The second method reduces the amount of (remaining) likelihood computation by updating only relevant partial likelihoods, estimating relevant edge lengths using analytical formulae and then calculating the change in likelihood for candidate SPR moves locally, instead of having to re-evaluate the entire tree. Our results show that this indeed gives rise to a significant reduction in both the number of likelihood computations and the running time of the algorithm, while maintaining a search performance comparable with that of the best-known methods to date.

It is expected that even better results can be achieved by investigating more thoroughly the various parameters that are involved in the algorithm. The speed of the algorithm can possibly still be improved by finding optimal parameter values, or better yet, a systematic way of setting these values based on the input data. Moreover, combining our (algorithmic) methods with a highly optimized implementation for performing likelihood computations (such as RAxML) will most likely result in an extremely efficient program for performing SPR searches on phylogenetic trees using ML.

ACKNOWLEDGEMENTS

We thank Stéphane Guindon for comments and help with PHYML, Alexandros Stamatakis for comments and details on the RAxML implementation and Rick Desper for discussions on distance-based SPR moves. We also thank the ACI IMPBIO (French Ministry of Research) for financial support for this project.

Conflict of Interest: none declared.

REFERENCES

- Adachi, J. and Hasegawa, M. (1996) Molphy, version 2.3 programs for molecular phylogenetics based on maximum likelihood. In Ishiguro, M., Kitagawa, G., Ogata, Y., Takagi, H., Tamura, Y. and Tsuchiya, T. (eds), *Computer Science Monographs 28*. Institute of Statistical Mathematics, Tokyo, pp. 1–150.
- Berry, V. and Gascuel, O. (2000) Inferring evolutionary trees with strong combinatorial evidence. *Theor. Comput. Sci.*, **240**, 271–298.
- Brent, R. (1973) *Algorithms for Minimization Without Derivatives*. Prentice-Hall, Englewood Cliffs, NJ.
- Bryant, D., Galtier, N. and Poursat, M.-A. (2005) Likelihood calculations in phylogenetics. In Gascuel, O. (ed.), *Mathematics of Evolution and Phylogeny*. Oxford University Press, Oxford, pp. 33–62.
- Chor, B. et al. (2000) Multiple maxima of likelihood in phylogenetic trees: an analytic approach. *Mol. Biol. Evol.*, **17**, 1529–1541.
- Chor, B. and Tuller, T. (2005) Maximum likelihood of evolutionary trees is hard. In *Proceedings of 9th Annual International Conference on Research in Computational Molecular Biology*, Cambridge, MA, pp. 296–310.
- Desper, R. and Gascuel, O. (2002) Fast and accurate phylogeny reconstruction algorithms based on the minimum-evolution principle. *J. Comput. Biol.*, **9**, 687–705.
- Desper, R. and Gascuel, O. (2005) The minimum evolution distance-based approach to phylogenetic inference. In Gascuel, O. (ed.), *Mathematics of Evolution and Phylogeny*. Oxford University Press, Oxford, pp. 1–32.
- Felsenstein, J. (1981) Evolutionary trees from DNA sequences: a maximum likelihood approach. *J. Mol. Evol.*, **17**, 368–376.
- Felsenstein, J. (1989) PHYLIP—phylogeny inference package (version 3.2). *Cladistics*, **5**, 164–166.
- Gascuel, O. (2000) Evidence for a relationship between algorithmic scheme and shape of inferred trees. In Gaul, W., Opitz, O. and Schader, M. (eds), *Data Analysis, Scientific Modeling and Practical Applications*. Springer, Berlin, pp. 157–168.
- Guindon, S. and Gascuel, O. (2003) A simple, fast, and accurate algorithm to estimate large phylogenies by maximum likelihood. *Syst. Biol.*, **52**, 696–704.
- Huelsenbeck, J.P. (1995) Performance of phylogenetic methods in simulation. *Syst. Biol.*, **44**, 1748.
- Huelsenbeck, J.P. and Ronquist, F. (2001) MrBayes: Bayesian inference of phylogeny. *Bioinformatics*, **17**, 754–755.
- Kishino, H. et al. (1990) Maximum likelihood inference of protein phylogeny and the origin of chloroplasts. *J. Mol. Evol.*, **31**, 151–160.
- Kuhner, M.K. and Felsenstein, J. (1994) A simulation comparison of phylogeny algorithms under equal and unequal evolutionary rates. *Mol. Biol. Evol.*, **11**, 459–468.
- Lemmon, A. and Milinkovitch, M. (2002) The metapopulation genetic algorithm: an efficient solution for the problem of large phylogeny estimation. *Proc. Nat Acad. Sci. USA*, **99**, 10516–10521.
- Lewis, P. (1998) A genetic algorithm for maximum likelihood phylogeny inference using nucleotide sequence data. *Mol. Biol. Evol.*, **15**, 277–283.
- McCracken, K.G. and Sorenson, M.D. (2005) Is homoplasy or lineage sorting the source of incongruent mtDNA and nuclear gene trees in the stiff-tailed ducks (nomonyx-oxymura). *Syst. Biol.*, **54**, 35–55.
- Olsen, G. et al. (1994) fastDNAm1: a tool for construction of phylogenetic trees of DNA sequences using maximum likelihood. *Comput. Appl. Biosci.*, **10**, 41–48.
- Pauplin, Y. (2000) Direct calculation of a tree length using a distance matrix. *J. Mol. Evol.*, **51**, 41–47.
- Rannala, B. and Yang, Z. (1996) Probability distribution of molecular evolutionary trees: A new method of phylogenetic inference. *J. Mol. Evol.*, **43**, 304–311.
- Ranwez, V. and Gascuel, O. (2002) Improvement of distance-based phylogenetic methods by a local maximum likelihood approach using triplets. *Mol. Biol. Evol.*, **19**, 1952–1963.
- Rosenberg, M. and Kumar, S. (2001) Traditional phylogenetic reconstruction methods reconstruct shallow and deep evolutionary relationship equally well. *Mol. Biol. Evol.*, **19**, 1823–1827.
- Salter, L. and Pearl, D. (2001) Stochastic search strategy for estimation of maximum likelihood phylogenetic trees. *Syst. Biol.*, **50**, 717.
- Simon, D. and Larget, B. (2000) Bayesian analysis in molecular biology and evolution (BAMBE), version 2.03beta.
- Stamatakis, A. (2004) Distributed and parallel algorithms and systems for inference of huge phylogenetic trees based on the maximum likelihood method. PhD Thesis. Technische Universität München, Germany.
- Stamatakis, A. (2005) An efficient program for phylogenetic inference using simulated annealing. In *Proceedings of 19th IEEE/ACM International Parallel and Distributed Processing Symposium (IPDPS'05)-Workshop 7*, pp 198b (electronic edition).
- Stamatakis, A. et al. (2005) RAXML-III: a fast program for maximum likelihood-based inference of large phylogenetic trees. *Bioinformatics*, **21**, 456–463.
- Stamatakis, A., Ludwig, T., Meier, H. and Wolf, M.J. (2002) AxML: a fast program for sequential and parallel phylogenetic tree calculations based on the maximum likelihood method. In *Proceedings of 1st IEEE Computer Society Bioinformatics Conference (CSB2002)*, pp. 21–(electronic edition).
- Swofford, D. (1996) *PAUP*—Phylogenetic Analysis Using Parsimony (*and other methods)*. Sinauer, Sunderland.
- Vinh, S. and von Haeseler, A. (2004) IQPNNI: moving fast through tree space and stopping in time. *Mol. Biol. Evol.*, **21**, 1565–1571.
- Vogler, A.P. et al. (2005) Exploring rate variation among and within sites in a densely sampled species tree: species level phylogenetics of north american tiger beetles (genus *Cicindela*). *Syst. Biol.*, **54**, 4–20.
- Winkwoeth, R.C. et al. (2005) Biogeographic interpretation of splits graphs: least squares optimisation of branch lengths. *Syst. Biol.*, **54**, 56–65.
- Yang, Z. (1994) Maximum likelihood phylogenetic estimation from DNA sequences with variable rates over sites: approximate methods. *J. Mol. Evol.*, **39**, 306–314.
- Yuan, Y.-M. et al. (2005) Phylogeny and biogeography of *exacum* (gentianaceae): a disjunctive distribution in the Indian Ocean basin resulting from long distance dispersal and extensive radiation. *Syst. Biol.*, **54**, 21–34.