# An Efficient Algorithm for Statistical Multiple Alignment on Arbitrary Phylogenetic Trees 

G.A. LUNTER, I. MIKLÓS, Y.S. SONG, and J. HEIN


#### Abstract

We present an efficient algorithm for statistical multiple alignment based on the TKF91 model of Thorne, Kishino, and Felsenstein (1991) on an arbitrary $k$-leaved phylogenetic tree. The existing algorithms use a hidden Markov model approach, which requires at least $O\left(\sqrt{5}^{k}\right)$ states and leads to a time complexity of $O\left(5^{k} L^{k}\right)$, where $L$ is the geometric mean sequence length. Using a combinatorial technique reminiscent of inclusion/exclusion, we are able to sum away the states, thus improving the time complexity to $O\left(2^{k} L^{k}\right)$ and considerably reducing memory requirements. This makes statistical multiple alignment under the TKF91 model a definite practical possibility in the case of a phylogenetic tree with a modest number of leaves.


Key words: statistical alignment, multiple alignment, phylogeny, maximum likelihood.

## 1. INTRODUCTION

Amajor way by which biological sequences evolve is mutation. Three main types of mutation events are substitutions, insertions, and deletions of amino acids or nucleotides. The latter two events, insertions and deletions, introduce the problem of aligning sequences, so that homologous positions appear in the same column of an alignment. When the aim is to find sequences from a database which are homologous to a query sequence, current alignment techniques perform quite well. The alignment problem in phylogenetics, however, is still a serious challenge (Lee, 2001; Goldman, 1998). For example, the phylogenies inferred from different, but equally good, alignments can be quite different (Goldman, 1998). Some researchers have suggested that the regions in the sequences which are responsible for this variance in the inferred phylogeny (usually referred to as unalignable regions [Lee, 2001]) should be omitted. It is unclear, however, how much information is lost by ignoring such regions. Moreover, several studies have suggested that accurate evolutionary parameters cannot be obtained using only a single "best" alignment (Thorne et al., 1991; Hein et al., 2000), even if this alignment is seemingly reliable (Hein et al., 2000). A more robust approach is to take into account all possible alignments, or a large subset of those, in a statistical framework.

Since insertions and deletions, like substitutions, are rare random events, it seems natural to model them by a continuous time stochastic process, with rates for the three types of mutations. Stochastic models

Department of Statistics, University of Oxford, Oxford, OX1 3TG, UK.
for the substitution process are well known and widely used to obtain maximum-likelihood evolutionary parameters (Felsenstein, 1981; Whelan et al., 2001). The model of Thorne, Kishino, and Felsenstein (the TKF91 model, hereafter) of sequence evolution incorporates such a mutation model, and moreover allows a nucleotide to spawn new nucleotides adjacent to itself, and to die. To model evolution on a phylogenetic tree, the TKF91 model is applied to each branch of the tree, and likelihood calculation of multiple alignments then becomes possible.

Since the TKF91 model is reversible (and the sequence distribution at the root is assumed to be at equilibrium), the root placement is immaterial. In the case of two sequences, one of the sequences may therefore be considered the ancestral sequence and the other its descendant. In that framework, the number of possible alignments is finite. When the TKF91 model is applied to a tree, however, an infinite number of sequences may appear at internal nodes, for any given set of sequences which appear at the leaf nodes. Thus, such an extension is not straightforward, and several methods have been proposed. For star-shaped trees, Steel and Hein (2001) have constructed an algorithm with $O\left(L^{2 k}\right)$ running time for $k$ sequences with geometric mean length $L$, and it has subsequently been extended, with similar time complexity, to binary trees (Hein, 2001). Thenceforth, the time complexity has been reduced to $O\left(4^{k} L^{k}\right)$ for star-shaped trees (Miklós, 2002) and to $O\left(5^{k} L^{k}\right)$ for arbitrary binary trees (Hein et al., 2002) (see Appendix B). In the latter paper, the TKF91 model has been reformulated as a multiple hidden Markov model, in which likelihood calculations can be performed using forward and backward algorithms known in the HMM literature (Durbin et al., 1998). Similar ideas have appeared for pairwise and triplewise statistical alignment as well, e.g., Holmes and Bruno (2001) give HMM formulations for alignment on two- and three-leaved trees and use it for sampling larger trees. In the Markov models, there are $O\left(\sqrt{5}^{k}\right)$ states, and therefore the running time of these algorithms is indeed $O\left(5^{k} L^{k}\right)$, in accordance with the general theory of HMMs (Durbin et al., 1998). The memory usage is of the order $O\left(\sqrt{5}^{k} L^{k}\right)$ if the entire dynamic programming table is retained and $O\left(\sqrt{5}^{k} L^{k-1}\right)$ if not (Hirshberg, 1977).

Because of the large number of states, the aforementioned algorithms are quite slow, even for a moderate number of sequences. Hein et al. (2000), however, have developed an algorithm for statistical alignment of two sequences which needs neither different states of a HMM nor partition of probabilities into different types of alignments (Thorne et al., 1991). This is in contrast to the original formulation of the TKF91 model, which uses three states. The algorithm of Hein et al. (2000) is, both in time and space complexity, as simple as the traditional distance- or similarity-based dynamic programming algorithms (Needleman and Wunsch, 1970; Sankoff and Kruskall, 1983). Henceforward, we refer to it as the one-state recursion.

In this paper, we present a one-state recursion for statistical alignment on arbitrary phylogenetic trees. Essentially, it combines the idea of Hein et al. (2000) generalized to trees and Felsenstein's reverse traversal algorithm (1981). The first idea allows us to reduce the number of states to one, and the second to compute in linear time the exponential number of terms that occur in the transition factors. The resulting algorithm has time complexity $O\left(2^{k} L^{k}\right)$ and space complexity $O\left(L^{k}\right)$ if the entire dynamic programming table is stored in memory, $O\left(L^{k-1}\right)$ if not. This represents a great saving in both space and time compared to the hidden Markov recursion and makes statistical alignment on a phylogenetic tree of modest size a definite practical possibility. Indeed, we have implemented both methods for three and four sequences and achieved a considerable acceleration. One of us (IM) was able to perform a likelihood ratio test (Felsenstein, 1981; Whelan et al., 2001) for more than 250 triplets of yeast protein sequences in a day, using triplewise statistical alignments. Moreover, this one-state recursion can be coupled with corner-cutting techniques (Hein et al., 2000), which provides further reduction both in time and space complexity.

The organization of this paper is as follows. We briefly describe the TKF model in Section 2 and discuss in Section 3 the one-state recursion in the case of two sequences. The two-sequence example is simple, but well illustrates our general approach. We discuss in Section 4 our algorithm for the one-state recursion and consider a specific application in Section 5. The main ideas underlying our proof are sketched in Section 6. The general one-state recursion is described in Section 7 and the computation of the transition factors which occur in the recursion is discussed in Section 8. In Appendix C we draw a connection with a hidden Markov model and describe how optimal alignments can be obtained. We conclude with some remarks and discussion in Section 9. In the appendix, we show that the number of Markov states in the HJP recursion is $O\left(\sqrt{5}^{k}\right)$ and give a proof of Lemma 1 from Section 7.

## 2. THE TKF MODEL

The TKF91 model is a continuous time reversible Markov model for the evolution of nucleotide (or amino acid) sequences. It models three of the main processes in sequence evolution, namely substitutions, insertions, and deletions of nucleotides, approximating these as single-nucleotide processes. A nucleotide sequence is represented by an alternating string of nucleotides and links connecting the nucleotides, and this string both starts and terminates with a link. We adopt the view that insertions originate from links and add a nucleotide-link pair to the right of the original link; deletions originate from nucleotides and have the effect of removing the nucleotide and its right link. (This view is slightly different but equivalent to the original description; see Thorne et al. [1991].) In this way, nucleotide subsequences evolve independently of each other, and the evolution of a nucleotide sequence is the sum of the evolutions of individual nucleotide-link pairs. The leftmost link of the sequence has no corresponding nucleotide to its left; hence, it is never deleted, and for this reason it is called the immortal link.

Since subsequences evolve independently, it suffices to describe the evolution of a single nucleotide-link pair. In a given time span $\tau$, this evolves into a sequence of nucleotides of finite length. The first nucleotide of this sequence may be homologous to the original one, while subsequent ones will be nonhomologous. Table 1 summarizes the corresponding probabilities. The parameters $\lambda$ and $\mu$ are the birth rate per link and the death rate per nucleotide, respectively, and in order to have a finite equilibrium sequence length, we require $\lambda<\mu$. For brevity, we write

$$
\beta(\tau)=\frac{1-e^{(\lambda-\mu) \tau}}{\mu-\lambda e^{(\lambda-\mu) \tau}}
$$

On the right-hand side of the arrow in the column labeled "Fate," $C$ denotes a nucleotide homologous to the original nucleotide, whereas \#'s denote nonhomologous nucleotides. The immortal link is denoted by $\star$, and other links are suppressed. All final arrangements can be thought of as being built from five basic "processes" which we call Birth, Extinction, Homologous, New (or Nonhomologous) and Initial (or Immortal). These processes are labeled by their initials, and each corresponds to a specific probability factor as follows:

$$
\begin{gather*}
B_{\tau}=\lambda \beta(\tau) \quad E_{\tau}=\mu \beta(\tau) \quad I_{\tau}=1-\lambda \beta(\tau)  \tag{1}\\
H_{\tau}=e^{-\mu \tau}(1-\lambda \beta(\tau)) \quad N_{\tau}=\left(1-e^{-\mu \tau}-\mu \beta(\tau)\right)(1-\lambda \beta(\tau))
\end{gather*}
$$

In a tree, time flows forward from the root to the leaves, and to each node of the tree we associate a time parameter $\tau$ which is set equal to the length of the incoming branch. For the root, $\tau=\infty$ by assumption of stationarity at the root, and the resulting equilibrium length distribution of the immortal link sequence is geometric with parameter $B_{\infty}=\lambda / \mu$ (where length 0 is possible); other links will have left no descendants since $H_{\infty}=N_{\infty}=0$.

In the original TKF91 model, a simple substitution process known as the "Felsenstein81 model" (Felsenstein, 1981) has been used. It is straightforward to replace this by more general models for substitutions of nucleotides or amino acids (Hein et al., 2000). In the present paper, when a new nonhomologous character appears at a node (as the result of a $B$ or $N$ process), it is always drawn from the equilibrium distribution; if a character at a node is homologous to the character at its immediate ancestral node, then the probability of this event is given by the chosen substitution model.

Table 1. Probabilities of the TKF91 Model (See Text)

| Fate | Probability | Label |
| :--- | :--- | :--- |
| $C \rightarrow C \#^{n-1}$ | $e^{-\mu \tau}(1-\lambda \beta(\tau))(\lambda \beta(\tau))^{n-1}$ | $H_{\tau} B_{\tau}^{n-1}$ |
| $C \rightarrow \#^{n}$ | $\left(1-e^{-\mu \tau}-\mu \beta(\tau)\right)(1-\lambda \beta(\tau))(\lambda \beta(\tau))^{n-1}$ | $N_{\tau} B_{\tau}^{n-1}$ |
| $C \rightarrow-$ | $\mu \beta(\tau)$ | $E_{\tau}$ |
| $\star \rightarrow \star^{n}$ | $(1-\lambda \beta(\tau))(\lambda \beta(\tau))^{n}$ | $I_{\tau} B_{\tau}^{n}$ |

Note that we do not sum over possible alignments, but over what we call evolutionary histories (see section 7.1), which is a discrete summary of the actual evolutionary events in the model specifying the fate of all ancestral nucleotides. As was pointed out by Thorne et al. (1991), for two sequences, there is almost a correspondence between alignment and evolutionary history (except that gaps in two sequences that immediately follow each other may be interchanged without changing the interpretation as alignments whereas the evolutionary interpretation in the TKF model is different.) This almost-correspondence breaks down for trees with larger number of leaves, since information on what occurs at internal nodes is lost in an alignment, which only records the homology structure of observed nucleotides.

## 3. ONE-STATE RECURSION: A GRAPHICAL PROOF FOR TWO SEQUENCES

Hein et al. (2000) give a proof for the one-state recursion in the case of two-sequence alignment. Here we provide a different proof for this particular case. Our proof for the general case follows a similar line of ideas.

Suppose $A^{1}$ is the ancestral sequence observed at the root and $A^{2}$ is the descendant sequence which results from the TKF91 process after time $\tau$. Let $A_{i}^{k}$ (respectively, $a_{i}^{k}$ ) denote the $i$-long prefix (respectively, the $i$ th character) of sequence $A^{k}$. Let $P(i, j)$ be the joint probability of observing $A_{i}^{1}$ and $A_{j}^{2}$. Given that $P(i-1, j), P(i, j-1)$, and $P(i-1, j-1)$ are known, we want to compute $P(i, j)$.

The possible evolutionary histories (or alignments) are customarily (Thorne et al., 1991) classified into three groups, $S^{0}, S^{1}$, and $S^{2}$, according to whether the rightmost root link has 0,1 , or at least 2 descendants at the leaf. Another way to define these sets is by the last columns in the alignment; symbolically, $S^{0}=\left\{\begin{array}{c}\# \\ -\end{array}\right\}$, $S^{1}=\left\{\begin{array}{c}\#, \\ \#, \\ - \\ \#\end{array}\right\}$, and $S^{2}=\left\{\begin{array}{c}\#-\underset{\#}{\#}, \underset{\#}{-} \\ \#\end{array}\right\}$, where the upper symbols represent nucleotides at the root. These sets are associated to the evolutionary processes as follows. After the $B_{\infty} E_{\tau}$ process, the rightmost root link has 0 descendants, so the history is in $\left\{\begin{array}{c}\# \\ -\end{array}\right\}$. After $B_{\infty} H_{\tau}$ or $B_{\infty} N_{\tau}$, it is in $\left\{\begin{array}{c}{\left[\begin{array}{l}\# \\ \#\end{array},-\begin{array}{l}\#\end{array}\right\} \text {, while after a }}\end{array}\right.$ birth at the leaf node $\left(B_{\tau}\right)$, the alignment ends in either of $\left\{\begin{array}{l}\#- \\ \#, ~-~ \# ~ \# ~\end{array}\right\}$. These processes may happen in any order, except that a transition from $\left\{\begin{array}{c}\# \\ -\end{array}\right\}$ to $\left\{\begin{array}{c}\#-- \\ \#,-- \\ \#\end{array}\right\}$ is not allowed: new births at the leaf $\left(B_{\tau}\right.$ processes) are allowed except after an extinction $\left(E_{\tau}\right)$ process; see Table 1. A graphical depiction of the three-state recursion that implements this restriction is provided in Fig. 1. The labeled circle segment represents probabilities of observing the associated sequence prefixes, provided the alignment ends according to the labeling.

If we combine the three states at each position, by summing their probabilities, we can still compute the contribution of the vertical and diagonal transitions to the probabilities at $(i, j)$, since the transition probabilities ( $p_{0}$ and $p_{1}$ ) do not depend on the state. But, for the horizontal transition we would include a contribution from the illegal transition $\left\{\begin{array}{c}\# \\ -\end{array}\right\} \rightarrow\left\{\begin{array}{c}\#-\underset{\# \#}{\#}, \underset{\#}{-}, ~\end{array}\right\}$ (denoted by a dashed arrow), leading to an overestimation of the probability at $(i, j)$. This overestimation can be exactly corrected for, however, since the segment labeled ${ }_{-}^{\#}$ at $(i, j-1)$ is the result of a $B_{\infty} E_{\tau}$ process from position $(i-1, j-1)$ (the dot-dashed arrows in Fig. 1). In summary, if we let $P(i, j)$ be the sum of the three conditional probabilities at $(i, j)$, we obtain the following one-state recursion:

$$
\begin{aligned}
P(i, j)= & P(i-1, j) \times B_{\infty} \pi\left(a_{i}^{1}\right) E_{\tau} \\
& +P(i-1, j-1) \times B_{\infty} \pi\left(a_{i}^{1}\right)\left[N_{\tau} \pi\left(a_{j}^{2}\right)+H_{\tau} p\left(a_{i}^{1} \rightarrow a_{j}^{2}\right)\right] \\
& +\left[P(i, j-1)-P(i-1, j-1) B_{\infty} \pi\left(a_{i}^{1}\right) E_{\tau}\right] \times B_{\tau} \pi\left(a_{j}^{2}\right)
\end{aligned}
$$

Here $\pi(\cdot)$ denotes the equilibrium nucleotide distribution, and $p(\alpha \rightarrow \gamma)$ is the probability that nucleotide $\alpha$ evolves into $\gamma$ in time $\tau$. After rearranging terms, we arrive at the recursion depicted in Fig. 2.

Although it is possible, in principle, to manually derive analogous recursions for an arbitrary number of sequences following the same line of reasoning, such an approach rapidly gets very tedious, as the number of terms which one needs to consider grows exponentially with the number of sequences. The goal of the present paper is to provide a more constructive method that allows us to prove that a one-state recursion indeed exists for arbitrary trees. Of more practical importance, the proof immediately suggests an efficient method for calculating the transition factors, through a post-order tree traversal algorithm, thus allowing computational derivation of the desired one-state recursion for arbitrary trees.


FIG. 1. A graphical representation of the three-state recursion for two-sequence alignment. The $p_{k}$ are transition probabilities associated to $S^{k}$ (see text): $p_{0}=B_{\infty} E_{\tau}, p_{1}=B_{\infty} H_{\tau}+B_{\infty} N_{\tau}$, and $p_{2}=B_{\tau}$. Circles represent positions in the dynamic programming table; circle segments correspond to the three states. The \# and - symbols represent residues and gaps in the last (two) column(s) of the partial alignment.


FIG. 2. The one-state recursion for two sequences. The factors next to the arrows are "transition factors," not probabilities. Due to combining states, negative terms appear to compensate for illegal transitions implicit in other transition factors.

## 4. ALGORITHM

Using the TKF model, it is easy to compute the likelihood of a specific evolutionary history, given as a sequence of events. In effect, this integrates out all possible evolutionary histories on the edges of the phylogenetic tree, while fixing the nucleotide sequences and their mutual alignment at internal nodes. The likelihood of observing the given nucleotide sequences at the leaves is obtained by summing this likelihood over all possible sequences of events. Since we usually do not have information about nucleotide sequences at the internal nodes, such a computation should sum out the internal nucleotide sequences as well. This is what our algorithm computes.

In order to write down the main theorem, we need some notation. Set $T$ is the set of nodes of the phylogenetic tree relating the sequences, and $r$ is its root. If $n \in T$ is a node, then $n_{l}$ and $n_{r}$ denote the left and right immediate descendants of $n$. For a node $n, \tau(n)$ denotes the length of the incoming branch,
and we abbreviate $B_{n}:=B_{\tau(n)}$ etc. For the root, $\tau(r)=\infty$ by assumption. Finally, $p_{n}(\alpha \rightarrow \gamma)$ is the probability that nucleotide $\alpha$ evolves into $\gamma$ in time $\tau(n)$.

Theorem 1. Suppose $A^{1}, \ldots, A^{k}$ are sequences related by a phylogenetic tree $T$ with $k$ leaves and root $r$. Let $\mathbf{K}=\left(\mathrm{K}_{1}, \ldots, \mathrm{~K}_{k}\right)$ and let $P(\mathbf{K})$ be the probability of producing, under the TKF91 model, the prefixes of $A^{1}, \ldots, A^{k}$ up to position $\mathrm{K}_{1}, \ldots, \mathrm{~K}_{k}$, respectively, where the numbering starts at 0 with the immortal link. The following algorithm computes $P(\mathbf{K})$ :

$$
\begin{aligned}
P(\mathbf{0}) & =\frac{\prod_{n \in T}\left(1-B_{n}\right)}{G_{\mathbf{0}}^{\mathbf{0}}(r,-)}, \\
P(\mathbf{K}) & =\frac{1}{G_{\mathbf{K}}^{\mathbf{0}}(r,-)} \sum_{\mathbf{v} \in\{0,1\}^{k} \backslash \mathbf{0}}\left(-G_{\mathbf{K}}^{\mathbf{v}}(r,-)\right) P(\mathbf{K}-\mathbf{v}) .
\end{aligned}
$$

If $n$ is an internal node, then

$$
\begin{aligned}
G_{\mathbf{K}}^{\mathbf{v}}(n, \alpha)= & {\left[E_{n_{l}} G_{\mathbf{K}}^{\mathbf{v}}\left(n_{l},-\right)+\sum_{\gamma}\left(H_{n_{l}} p_{n_{l}}(\alpha \rightarrow \gamma)+N_{n_{l}} \pi(\gamma)\right) G_{\mathbf{K}}^{\mathbf{v}}\left(n_{l}, \gamma\right)\right] } \\
& \times\left[E_{n_{r}} G_{\mathbf{K}}^{\mathbf{v}}\left(n_{r},-\right)+\sum_{\gamma}\left(H_{n_{r}} p_{n_{r}}(\alpha \rightarrow \gamma)+N_{n_{r}} \pi(\gamma)\right) G_{\mathbf{K}}^{\mathbf{v}}\left(n_{r}, \gamma\right)\right], \\
G_{\mathbf{K}}^{\mathbf{v}}(n,-)= & G_{\mathbf{K}}^{\mathbf{v}}\left(n_{l},-\right) G_{\mathbf{K}}^{\mathbf{v}}\left(n_{r},-\right)-B_{n} \sum_{\gamma} \pi(\gamma) G_{\mathbf{K}}^{\mathbf{v}}(n, \gamma) .
\end{aligned}
$$

If $n$ is a leaf node, then

$$
\begin{aligned}
G_{\mathbf{K}}^{\mathbf{v}}(n, \alpha) & = \begin{cases}1, & \text { if } \mathrm{v}_{n}=1 \text { and } a_{\mathrm{K}_{n}-1}^{n}=\alpha, \\
0, & \text { otherwise. }\end{cases} \\
G_{\mathbf{K}}^{\mathbf{v}}(n,-) & = \begin{cases}1, & \text { if } \mathrm{v}_{n}=0, \\
-B_{n} \pi\left(a_{\mathrm{K}_{n}-1}^{n}\right), & \text { if } \mathrm{v}_{n}=1 .\end{cases}
\end{aligned}
$$

By $\mathrm{K}_{n}$, we mean the component of the vector $\mathbf{K}$ corresponding to leaf $n$, and similarly for $\mathrm{v}_{n}$. By $a_{\mathrm{K}_{n}-1}^{n}$ we denote the character at position $\mathrm{K}_{n}-1$ in the sequence at leaf $n$.

The numbers $G_{\mathbf{K}}^{\mathbf{v}}(r,-)$ are the transition factors and play the same role as the transition factors in Fig. 2. They are sums and differences of probabilities, designed to exactly cancel all contributions of illegal transitions due to the merging of states into a single probability, analogous to the two-sequence case described in Section 3. Since we now have $k$ leaves, the dynamic programming recursion for $P$ refers back to $2^{k}-1$ entries, instead of just three.

Note that $G_{\mathbf{K}}^{\mathbf{0}}$ does not depend on $\mathbf{K}$, so a small acceleration can be achieved by computing the prefactor $1 / G_{\mathbf{K}}^{\mathbf{0}}$ only once. A more substantial acceleration is achieved by noting that $G_{\mathbf{K}}^{\mathbf{v}}(n, \alpha)$ and $G_{\mathbf{K}}^{\mathbf{v}}(n,-)$ depend only on the components of $\mathbf{v}$ which correspond to the leaves contained in the subtree of $n$. Therefore, if the tree with root $n$ has $s$ leaves, it is enough to compute $G_{\mathbf{K}}^{\mathbf{v}}(n, \alpha)$ for $2^{s}$ values of $\mathbf{v}$. Once these are computed, any entry for the ancestral node can be computed in constant time. Therefore, the overall running time of the algorithm is $2^{k} L^{k}$, where $k$ is the number of sequences and $L$ is the geometric mean of the sequence lengths.

## 5. AN EXAMPLE

As a concrete example, we have used our algorithm to find the most likely evolutionary tree which relates the following four globins: human $\alpha 1$ and $\beta$ hemoglobins, human myoglobin, and leghemoglobin from the bean Canavalia lineata. There are three topologically different unrooted trees with four leaves, and we have estimated maximum likelihood parameters for each of them. The most likely tree is, as expected, the one that groups human $\alpha 1$ and $\beta$ hemoglobin together; the likelihoods of the other two trees are about a factor 400 smaller (see Fig. 3). After obtaining the maximum likelihood tree, the maximum likelihood alignment can be obtained using the Viterbi algorithm. This alignment is shown in Table 2. For a more detailed discussion on the Viterbi algorithm, see Appendix C.

We have optimized the likelihood as a function of the five branch length parameters and $\lambda$; the deletion rate parameter $\mu$ has been adjusted to make the expected sequence length $\frac{\lambda}{\mu-\lambda}$ equal to the average sequence length. By using corner cutting methods, we could confine the calculation to the parts of the


FIG. 3. Maximum likelihood trees relating human $\alpha 1$ and $\beta$ hemoglobins, myoglobin, and bean leghemoglobin, for all three topologically distinct trees; total likelihood values and insertion rate; and estimated standard deviations. The numbers refer to branch lengths in units of expected number of substitutions per site. As substitution rate matrix we used Dayhoff's PAM matrix.

Table 2. The Maximum Likelihood Alignment for the First Pedigree in Fig. $3^{\mathrm{a}}$
Hba1: MV--LSPADKTNVKAAWGKVGAHAGEYGAEALERMFLSFPTTKTYFPHF--DLS-H-----GSAQVKGHGKKVAD-AL-TNAHbb: MV-HLTPEEKSAVTALWGK--NVDEVGGEALGRLLVVYPWTQRFFESF-GDLSTPDAVM-GNPKVKAHGKKVLG-AF-SDGMyo: MG--LSDGEWQLVLNVWGKVEADIPGHGQEVLIRLFKGHPETLEKFDKFK-HLKSEDE-MKASEDLKKHGATVLT-AL-GGI-
Legh: MGA-FSEKQESLVKSSWEAFKQNVPHHSAVFYTLIIEKAPAAQNMFS-F---LSNGVD-P-NNPKLKAHAEKVFKMTVDSAVQ


#### Abstract

VAHVDDMPNALSALSDLHAHKLRVDPVNFK-LLSHCLLVTLAAHLPAEFTPAVHASLDKFLASVSTVL-TS-K---YR-LAHLDNLKGTFATLSELHCDKLHVDPENFR-LLGNVLVCVLAHHFGKEFTPPVQAAYQKVVAGVANAL-AH-K---YH-LKKKGHHEAEIKPLAQSHATKHKI-PVKYLEFISECIIQVLQSKHPGDFGADAQGAMNKALELFRKDMASNYKELGFQG LRAKGEVVLADPTLGSVHVQKGVLDP-HFL-VVKEALLKTFKEAVGDKWNDELGNAWEVAYDELAAAI-KK-A-MGSA-


${ }^{\mathrm{a}}$ The log-likelihood of this alignment is -1593.223 .
dynamic programming table which contribute noninfinitesimally to the final probability, resulting in a speedup of about a factor of 500 . At the extremal point we estimated the second derivative to get a covariance matrix; see Fig. 3 for the standard deviations in the directions of the eigenvectors of this matrix.

## 6. IDEA OF PROOF

Our aim is to give a "one-state" dynamic programming algorithm for calculating the joint likelihood of observing a set of sequences at the leaves of an evolutionary tree. This likelihood is the sum of the probabilities of all evolutionary histories which produce the observed sequences at the leaves.

We introduce the concept of an event which corresponds to a nucleotide birth at a node of the phylogenetic tree and its subsequent fate down its descendant subtree. More precisely, in the process of going from one node to a descendant node, one of the following three things may happen: The nucleotide may survive as a homologous nucleotide, it may die leaving at least one surviving new nucleotide, or it may go extinct altogether. These three possibilities are labeled H, N, and E, respectively, as discussed in Section 2. In the first two cases, other nonhomologous descendants may have been born, but these possibilities are dealt with in subsequent events. Hence, in this model, evolutionary histories of nucleotide sequences correspond to sequences of events. However, this correspondence is not one to one for two reasons. One reason is that we need to honor the TKF model, which does not allow extinct nucleotides to spawn new ones. (In Table 1, this is reflected by the third entry, $C \rightarrow-$, which is the only one that may not be followed by a string of \#'s.) The other reason is that different sequences of events may represent the same evolutionary history, since the relative ordering of events on disjoint subtrees has no meaning. In a Markov chain approach, these two requirements are met by defining states and by ruling out certain events (i.e., transitions, in the Mealy machine view [Durbin et al., 1998]) in certain states. Essentially that is how we also approach the problem. We define a "state" and use it to rule out events which violate the TKF model. Our choice of state (see Fig. 4) still allows overcounting of histories. To overcome this, we require the events to be ordered in a specific way depending on the state.

The ordering we use has the property that in properly ordered sequences of events, an event which violates the TKF model can be recognized by comparing it to its immediate predecessor. This means that we do not need the state to decide which events are allowed, but the state is still used to define the proper ordering. We can now use a recursion which, as a first approximation, allows all events. This includes all legal sequences of events, as well as some sequences of events which end in an illegal pair, either because they are not properly ordered, or because they violate the TKF model. As a second approximation, we subtract all sequences of events which end in particular illegal pairs. In the same way as before, this includes some illegal triplets which should not have been subtracted, and these are added in again, etc., in a procedure reminiscent of inclusion-exclusion. We have dubbed the pairs, triplets, etc. involved in this procedure the disallowed sequences.

It turns out that there exist only finitely many disallowed sequences, so that the inclusion-exclusion procedure stops after finitely many steps (bounded by the number of nodes in the tree). Moreover, disallowed sequences emit at most one nucleotide per leaf, and that leads to a recursion which refers back only to the $2^{k}-1$ nearest predecessors in the $k$-dimensional dynamic programming table, where $k$ is the number of leaves.


FIG. 4. (a) A 6-leaved phylogenetic binary tree. (b) A possible state. Filled circles • (resp., open circles o) represent the nodes which are contained (resp., not contained) in the state. (c) A possible event. Any node without a label is not contained in the subtree corresponding to the event.

The ordering, and hence the disallowed sequences we have to sum over, still depends on the state. However, the state turns out to influence only the ordering of the elements of these disallowed sequences, not the set of events themselves. More precisely, for each state, there is a one-to-one correspondence between such sets (that we call illegal sets) and disallowed sequences. The terms involved in the summation do not depend on the ordering, and this enables us to sum away the dependence on the state. A small technical problem is that some events do not emit any character (known in the HMM literature as "silent states"). These events lead to self-references in the recursion, turning it into a (trivial, one-dimensional) linear equation. We can sum over infinitely many possible silent events using the standard "wing folding" technique (Eddy, 2001), which amounts to solving the linear equation.

It remains to compute the transition factors, which involve sums over illegal sets of events that lead to specific emissions. The number of illegal sets is still exponential in the number of nodes in the phylogenetic tree. However, using a correspondence between illegal sets and labelings of the tree, it is possible to compute this sum in linear time using a reverse-traversal algorithm. This recursion also deals with the nucleotide assignments in the internal nodes of the tree, similar to Felsenstein's postorder tree traversal algorithm (Felsenstein, 1981).

## 7. THE GENERAL ONE-STATE RECURSION

In this section, we introduce the general one-state recursion. The main result is presented in Section 7.2; to maintain the flow of discussion, we defer the proof until Appendix A. In Section 8, to arrive at the algorithm discussed in Section 4, we combine the result of the present section with a fast algorithm for computing the transition factors.

### 7.1. Definitions

Let $T$ denote the set of nodes of a rooted phylogenetic binary tree. Note that we suppose the tree edges to be given and fixed and that we only include the nodes in our representation $T$. Although the tree is rooted, the root position is immaterial because of reversibility of the model (Felsenstein's pulley principle, [1981]). The root may coincide with an internal node, giving a zero-length edge and somewhat simplifying the recursion, but for clarity we do not consider that case here.

A subtree $t$ of $T$ is a set of nodes with the properties that (i) for any node $n \in t$, its descendants are in $t$, and (ii) there is a unique node, called the root of $t$ and denoted by $r(t)$, without a parent. A labeling of a tree $t$ is a function $l: t \rightarrow\left\{B^{\alpha},-, H^{\alpha}, N^{\alpha} \mid \alpha \in \mathcal{A}\right\}$, where $\mathcal{A}$ is the alphabet, which is $\{A, C, G, T\}$ in the case of DNA sequences. Symbols $B, H$, and $N$ stand for birth, homologous, and new, respectively, as discussed in Section 2, and $\alpha$ is the character associated to the process. The symbol "-" stands for "no character" and means that an $E_{\tau}$ (extinction) process has occurred somewhere higher up in the tree.

An event $e$ is a pair ( $e^{t}, e^{l}$ ), where $e^{t}$ is a subtree of $T$ and $e^{l}$ a labeling of $e^{t}$, with the additional properties that (i) $e^{l}(n)=B^{\alpha}$ (any $\alpha$ ) if and only if $n=r\left(e^{t}\right)$, and (ii) if $e^{l}(n)=-$ for any $n \in e^{t}$, then $e^{l}\left(n^{\prime}\right)=-$ for all descendants $n^{\prime}$ of $n$. In this way, an event represents the birth of a nucleotide at $r\left(e^{t}\right)$


Event $e$


Event $e^{t}$


Event $e^{\prime \prime}$

FIG. 5. Three possible events.
and its subsequent fate down the phylogenetic tree. An event carries a definite probability $p(e)$ which can be calculated from the TKF91 model, and we will concern ourselves with that in Section 8. Note that an event gives rise to at most one new nucleotide at each leaf sequence, and that such nucleotides may or may not be homologous to each other.

We consider sequences of events, denoted by $E=\left(e_{1}, e_{2}, \ldots, e_{N}\right)$. These give rise to sequences of states $\left(S_{1}, S_{2}, \ldots, S_{N}\right)$, which are defined recursively as follows: $S_{1}:=T$ and $S_{i+1}:=\left(S_{i} \backslash e_{i}^{t}\right) \cup\{n \mid n \in$ $e_{i}^{t}$ and $\left.e_{i}^{l}(n) \neq-\right\}$. The state $S_{i}$ determines at which nodes new nucleotides may be born at event $i$. Figures 4 and 5 illustrate these definitions.

We introduce some notation for convenience: $t_{i}:=e_{i}^{t}, r_{i}:=r\left(e_{i}^{t}\right), l_{i}:=e_{i}^{l}$. These stand for the tree, the root, and the labeling of event $e_{i}$. We define the context $c_{i}$ by $c_{i}:=F$ if $r_{i} \in S_{i}$, and $c_{i}:=I$ if $r_{i} \notin S_{i}$, meaning fertile if the birth of event $e_{i}$ occurred at a live node, infertile if not. An event $e_{i}$ is called legal if $c_{i}=F$. A sequence of events is called legal if all events are legal. We define those (and only those) states that may result from a sequence of legal events to be legal states. Equivalently, we may define legal states as those which are the result of a single event at the root, or yet differently as $T$ with a number of proper subtrees removed. Examples of legal and illegal states are provided in Fig. 6.

In a sequence $E$ of events, the ordering of the events is important insofar as it determines the ordering of events pertaining to each individual leaf sequence. Permutations do not change the evolutionary history represented by a sequence as long as, for all leaves $n$, the ordering in the subsequence of those events which


Legal State
FIG. 6. Sequences of states corresponding to the sequences $\left(e, e^{\prime}\right)$ and $\left(e, e^{\prime \prime}\right)$ of events, where $e, e^{\prime}, e^{\prime \prime}$ are as in Fig. 5. The sequence $\left(e, e^{\prime}\right)$ is a legal sequence resulting is a legal state, whereas $\left(e, e^{\prime \prime}\right)$ is an illegal sequence giving rise to an illegal state. Note that $\left(e, e^{\prime}, e^{\prime \prime}\right)$ is a legal sequence.
intersect the path from root to $n$ is undisturbed. More precisely, two sequences $E=\left(e_{1}, \ldots, e_{N}\right), E^{\prime}=$ $\left(e_{1}^{\prime}, \ldots, e_{N}^{\prime}\right)$ represent the same history iff there exists a permutation $\sigma$ of $\{1, \ldots, N\}$ such that

$$
e_{\sigma(i)}=e_{i}^{\prime} \quad \text { and } \quad i<j, \quad t_{i} \cap t_{j} \neq \varnothing \Rightarrow \sigma(i)<\sigma(j) .
$$

We call these $\sigma$ allowed permutations (for $E$ ). This defines an equivalence relation on the set of sequences of events. We call the equivalence class of any sequence of events the (evolutionary) history defined by that sequence.

Although in general the state depends on the ordering of events, the context $c_{j}$ depends only on the history, since it depends on whether $r_{j} \in S_{j}$, which in turn depends on $l_{i}\left(r_{j}\right)$, for the unique $i<j$ with $r_{j} \in t_{i}$ and $r_{j} \notin t_{k}$ for all $i<k<j$. The pair ( $j, i$ ) so defined is equivariant under allowed permutations; i.e., the indices of the pair transform according to the permutation. This allows us to define a legal history to be a history represented by a legal sequence of events.

### 7.2. Summing away the states

We denote by $E(\boldsymbol{K})$ a set of legal sequences of events which emit the prefixes of the given sequences $A^{1}, \ldots, A^{k}$ up to position $\boldsymbol{K}=\left(K_{1}, \ldots, K_{k}\right)$ and which contains precisely one representative of each history compatible with the emissions. The quantity of interest is the likelihood

$$
P(\boldsymbol{K}):=\prod_{n \in T}\left(1-B_{n}\right) \sum_{\left(e_{1}, \ldots, e_{N}\right) \in E(\boldsymbol{K})} p\left(e_{1}\right) \cdots p\left(e_{N}\right) .
$$

To state the main result of this section, we need to introduce one more definition, which admittedly comes out of the blue at this point. The connection is made at the end of the proof in Appendix A.

Definition 1 (set of nested events). A set of events $\left\{e_{1}, \ldots, e_{N}\right\}$ is called $a$ set of nested events if $t_{i} \supseteq t_{j} \Rightarrow l_{i}\left(r_{j}\right)=-$.

Lemma 1. The likelihood $P(\boldsymbol{K})$ satisfies the following equation:

$$
\begin{aligned}
P(\boldsymbol{K})= & \sum_{e} P\left(\boldsymbol{K}-\mathbf{v}_{e}\right) p(e) \\
& -\sum_{\left\{e_{1}, e_{2}\right\} \text { nested events }} P\left(\boldsymbol{K}-\mathbf{v}_{e_{1}}-\mathbf{v}_{e_{2}}\right) p\left(e_{1}\right) p\left(e_{2}\right) \\
& +\sum_{\left\{e_{1}, e_{2}, e_{3}\right\} \text { nested events }} P\left(\boldsymbol{K}-\mathbf{v}_{e_{1}}-\mathbf{v}_{e_{2}}-\mathbf{v}_{e_{3}}\right) p\left(e_{1}\right) p\left(e_{2}\right) p\left(e_{3}\right)
\end{aligned}
$$

$$
\begin{equation*}
-\cdots \tag{2}
\end{equation*}
$$

Here $\mathbf{v}_{e}$ is the emission vector corresponding to the event $e$, that is, a $k$-dimensional vector whose component corresponding to each leaf is 0 if the leaf is labeled "-" (or if it is outside $t(e)$ and hence unlabeled), and 1 otherwise.

Proof. See Appendix A.

### 7.3. Eliminating the silent events

Note that, for a nontrivial tree $T$, the set $E(\boldsymbol{K})$ of legal sequences of events emitting a certain prefix is infinite, even for $\boldsymbol{K}=\mathbf{0}$. This is due to the so-called silent events, i.e., events $e$ without emissions at the leaves ( $\mathbf{v}_{e}=\mathbf{0}$ ). This results in the appearance of the term $P(\boldsymbol{K})$ on the right-hand side of (2). However, the contribution of these silent events can easily be calculated using a trick analogous to the "wing folding" technique in the HMM literature (Eddy, 2001). A similar trick can be found in Steel and Hein (2001).

Note that by definition, the set $\{e\}$ is a set of nested events, and therefore (2) can be rewritten as follows:

$$
\begin{equation*}
P(\boldsymbol{K})=\sum_{M \neq \varnothing \text { nested events, }} P\left(\boldsymbol{K}-\sum_{e_{i} \in M} \mathbf{v}_{e_{i}}\right)(-1)^{|M|+1} \prod_{e_{i} \in M} p\left(e_{i}\right) . \tag{3}
\end{equation*}
$$

We can decompose the sum as

$$
\begin{aligned}
P(\boldsymbol{K})= & \sum_{\substack{M \neq \varnothing \text { nested events, } \\
\text { not all } \mathbf{v}_{e_{i}}=\mathbf{0}}} P\left(\boldsymbol{K}-\sum_{e_{i} \in M} \mathbf{v}_{e_{i}}\right)(-1)^{|M|+1} \prod_{e_{i} \in M} p\left(e_{i}\right) \\
& +P(\boldsymbol{K}) \sum_{\substack{M \neq \varnothing \text { nested events, } \\
\text { all } \mathbf{v}_{e_{i}}=\mathbf{0}}}(-1)^{|M|+1} \prod_{e_{i} \in M} p\left(e_{i}\right) .
\end{aligned}
$$

Solving this equation for $P(\boldsymbol{K})$ gives the correct answer and includes histories with $0,1,2, \ldots$ silent events. A heuristic way to see this is to observe that the equation $X=a+s X$, where $s$ and $a$ are the silent and nonsilent contributions, respectively, is solved by $X=a /(1-s)=a\left(1+s+s^{2}+\cdots\right)$. This expansion clearly shows the separate contributions of chains of silent events of length $0,1,2, \ldots$. The reasoning in the proof of Lemma 1 makes sure that only sequences of legal events are included, and that applies to the silent states as well.

## 8. PRUNING THE TREE

In this section we shall present Theorem 1. The recursion (2) enables us to compute $P(\boldsymbol{K})$, but it requires summing over a number of terms exponential in the number of sequences. A similar problem occurs when calculating the likelihood of nucleotide emissions on a tree under a simple substitution model, in which case one needs to sum over an exponential number of nucleotide assignments to internal nodes. This can be computed in linear time by Felsenstein's linear time post-order tree traversal algorithm, which is again a dynamic programming algorithm, now on a tree instead of the more familiar square lattice. We proceed in a similar way, by finding a correspondence between certain labelings of $T$ and sets of nested events. The summation over all such labelings, including the $(-1)^{k+1}$ sign arising from the inclusion-exclusion argument, can be performed in linear time by a dynamic programming algorithm similar to Felsenstein's.

The probability factor associated to an event is the product of conditional factors at the nodes of the phylogenetic tree, and the conditional factor at a node depends only on the labeling of the node and its ancestral node. The label determines both the process and the associated nucleotide, except for "-" which implies an Extinction event only if its ancestral node is not labeled "-". Symbolically,

$$
\begin{array}{cl} 
& p\left(B^{\alpha} \mid-\right)=B_{n} \pi(\alpha), \\
p\left(N^{\alpha} \mid X^{\gamma}\right)=N_{n} \pi(\alpha), & p\left(N^{\alpha} \mid-\right)=0, \\
p\left(H^{\alpha} \mid X^{\gamma}\right)=H_{n} p_{n}(\gamma \rightarrow \alpha), & p\left(H^{\alpha} \mid-\right)=0, \\
p\left(-\mid X^{\gamma}\right)=E_{n}, & p(-\mid-)=1,
\end{array}
$$

where $\alpha, \gamma \in \mathcal{A}$ (c.f. Section 7.1). The symbol $X$, denoting the process in the ancestral node, may be anything except " - ". The probability factors $B_{n}, N_{n}, H_{n}$, and $E_{n}$ are subscripted with a node and are defined (c.f. (1)) in terms of the length of the incoming branch ( $\infty$ if $n=r(T)$ ). Note that $p\left(B^{\alpha} \mid X^{\gamma}\right)$ will never occur.

This allows us to calculate probability factors of events, but we want to sum over nested events directly. It follows from Definition 1 that for each node $n \in T$, at most one event in the nested set has $n$ labeled with a symbol other than "-". Furthermore, since the roots of events are uniquely identified as the only nodes labeled $B^{\alpha}$, there is a one-to-one correspondence between sets of nested events and labelings of $T$


FIG. 7. The tree labeling corresponding to the set of nested events $\left\{e, e^{\prime \prime}\right\}$, where the events $e$ and $e^{\prime \prime}$ are as in Fig. 5.
with $B^{\alpha}, H^{\alpha}, N^{\alpha}$, "-" which obey the following rule: the root and the immediate descendants of " - " are labeled with $B^{\alpha}$ or "-". See Fig. 7 for an example. Indeed, such labelings can be decomposed into sets of nested events by recursively removing subtrees containing exactly one $B^{\alpha}$ label at their roots and resetting the labelings to "-". Note that the all-"-" tree has a proper labeling according to these rules and corresponds to the empty set of nested events.

The factor associated to a nested set $S$ is the product of the separate probability factors of the event in $S$, with a $\operatorname{sign}(-1)^{|S|+1}$. The exponent $|S|+1$ is just one more than the number of $B^{\alpha}$,s occurring in the associated tree labeling. Hence, analogous to computing the probability factor of an event, the factor associated to a tree labeling may be calculated (except for a single overall minus sign) as the product of several "conditional factors" as follows:

$$
\begin{aligned}
f\left(B^{\alpha} \mid X^{\gamma}\right) & =-E_{n} B_{n} \pi(\alpha), & & f\left(B^{\alpha} \mid-\right)=-B_{n} \pi(\alpha), \\
f\left(N^{\alpha} \mid X^{\gamma}\right) & =N_{n} \pi(\alpha), & & f\left(N^{\alpha} \mid-\right)=0, \\
f\left(H^{\alpha} \mid X^{\gamma}\right) & =H_{n} p_{n}(\gamma \rightarrow \alpha), & & f\left(H^{\alpha} \mid-\right)=0, \\
f\left(-\mid X^{\gamma}\right) & =E_{n}, & & f(-\mid-)=1 .
\end{aligned}
$$

If we group terms in (3) according to their total emission vector $\mathbf{v}=\mathbf{v}_{e_{1}}+\cdots+\mathbf{v}_{e_{m}}$, the sum of contributing factors may be obtained by a Felsenstein-like recursion:

$$
\begin{aligned}
F(n, \alpha)= & {\left[E_{n_{l}} F_{\mathbf{K}}^{\mathbf{v}}\left(n_{l},-\right)+\sum_{\gamma} F\left(n_{l}, \gamma\right)\left(H_{n_{l}} p_{n_{l}}(\alpha \rightarrow \gamma)+N_{n_{l}} \pi(\gamma)-E_{n_{l}} B_{n_{l}} \pi(\gamma)\right)\right] } \\
& \times\left[E_{n_{r}} F\left(n_{r},-\right)+\sum_{\gamma} F\left(n_{r}, \gamma\right)\left(H_{n_{r}} p_{n_{r}}(\alpha \rightarrow \gamma)+N_{n_{r}} \pi(\gamma)-E_{n_{r}} B_{n_{r}} \pi(\gamma)\right)\right], \\
F(n,-)= & {\left[F^{\mathbf{v}}\left(n_{l},-\right)-\sum_{\gamma} B_{n_{l}} F^{\mathbf{v}}\left(n_{l}, \gamma\right) \pi(\gamma)\right] \times\left[F^{\mathbf{v}}\left(n_{r},-\right)-\sum_{\gamma} B_{n_{l}} F^{\mathbf{v}}\left(n_{r}, \gamma\right) \pi(\gamma)\right], }
\end{aligned}
$$

where $n_{l}$ and $n_{r}$ denotes the left and right child node of $n$. Here, $F(n, \alpha)$ (respectively, $F(n,-)$ ) is the sum of all products of conditional factors on the subtree with root $n$, given that $n$ is labeled $X^{\alpha}$ (respectively, "-").

Note that we suppressed the dependence on $\mathbf{v}$. This dependence surfaces only when $n$ is a leaf, in which case the recursion terminates with $F(n, \alpha)=1$ if $\mathrm{v}_{n}=1$ and $\alpha$ is the character in sequence $A^{n}$ at the current position, and $F(n,-)=1$ if $\mathrm{v}_{n}=0 ; F=0$ in all other cases. The transition factor associated to $\mathbf{v}$ is finally

$$
\sum_{\substack{\left\{e_{1}, \ldots, e_{k}\right\} \text { nested set } \\ \mathbf{v}_{e_{1}}+\cdots+\mathbf{v}_{e_{k}}=\mathbf{v}}}(-1)^{k+1} \prod_{i} p\left(e_{i}\right)=-\left(F(r,-)-\sum_{\alpha} F(r, \alpha) B_{r} \pi(\alpha)\right)
$$

where $r$ is the root of $T$. A final useful simplification occurs if we set $G(n, \alpha)=F(n, \alpha)$ and $G(n,-)=$ $F(n,-)-\sum_{\gamma} B_{n} \pi(\gamma) F(n, \gamma)$, in which case the recursion for $G$ takes the form given in Section 4 and the transition factor is simply $-G(r,-)$.

To compute $P(\mathbf{K})$, we have to eliminate the silent events by solving a linear equation, as described in Section 7.3. This amounts to dividing each transition factor by

$$
\begin{equation*}
1-\sum_{\substack{\varnothing \neq\left\{e_{1}, \ldots, e_{k}\right\} \text { nested set } \\ \mathbf{v}_{e_{1}}+\cdots+\mathbf{v}_{e_{k}}=\mathbf{0}}}(-1)^{k+1} \prod_{i} p\left(e_{i}\right) \tag{4}
\end{equation*}
$$

one minus the transition factor associated to the null emission. Since the recursion for $F$ also includes the "-"-labeled tree which gets assigned a term +1 (corresponding to also including the empty set in (4)), in fact $G(r,-)$ for $\mathbf{v}=\mathbf{0}$ is precisely equal to (4).

It remains to compute the initial value $P(\mathbf{0})$, the probability of emitting no nucleotides to any sequence and being left with only the immortal link. According to the TKF91 model, this is the product of the immortal link prefactors $I_{n}=1-B_{n}$, where the product extends over all nodes in $T$, and silent events are included by dividing this by $G(r,-)$. This completes the proof of Theorem 1.

## 9. DISCUSSION

The importance of incorporating statistical analysis into biological studies has become abundantly clear over the years. In particular, many interesting problems which arise in the field of bioinformatics have been successfully addressed using statistical models. For instance, the evolution of biological sequences, which is the focus of the present paper, has been formulated in a statistical framework in which mutation events are seen as stochastic processes. A closer examination of the past progress reveals, however, that although substitution processes have been modeled as continuous-time evolutionary processes for more than three decades (Jukes and Cantor, 1969) and have been widely used in phylogenetics and genealogy (Felsenstein, 2001), modeling insertion and deletion processes based on evolution has only recently been generalized to an arbitrary number of sequences (Steel and Hein, 2001; Hein, 2001; Miklós, 2002; Hein et al., 2002). Furthermore, implementations of such generalizations have hitherto required quite a large running time. Alternative probabilistic approaches to sequence alignment exist (for example, see Mitchison [1999]), but insertion and deletion processes in such models are not explicitly based on evolution.

In the present paper, we have constructed several algorithmic improvements which make multiple statistical alignment computationally tractable. More precisely, we have proved that the one-state recursion exists for an arbitrary number of sequences, thus reducing the space complexity by a factor of $O\left(\sqrt{5}^{k}\right)$, where $k$ is the number of sequences. Furthermore, we have developed a reverse traversal algorithm which calculates, with time complexity linear in the number of sequences, each transition factor appearing in the one-state dynamic programming algorithm. This reduces the time complexity of the entire algorithm to $O\left(k 2^{k} L^{k}\right)$, where $L$ is the geometric mean sequence length. In fact, most of the transition factors appearing in the recursion share a common part, which hence needs to be calculated only once. Therefore, with a more careful implementation, the running time of the final algorithm can be reduced to $O\left(2^{k} L^{k}\right)$.

The achieved improvements are not only theoretical, but make multiple statistical alignment on trees possible in practice, as we have shown in the globin example. Compared to the ordinary hidden Markov model implementation on a four-sequence tree, our method gives rise to a speedup of about a factor 50, and as well as reduces the number of states from 45 to 1 , with an associated reduction in memory usage.

Our algorithm is useful mainly for trees with small number of leaves and may be used for hypothesis testing or as a basis for the quartet method. Even for small trees, however, corner cutting techniques are necessary in practical applications. The example of Section 5 requires computing a table of size $4.8 \times 10^{8}$ in a straightforward implementation, while corner cutting can reduce this by a factor of about 500 . Because the probability mass is highly concentrated around the alignment region, we lose almost no contribution to the total probability, even though only $0.2 \%$ of the entire table is actually visited. For more than, say, six sequences, our method will cease to be practical, and MCMC methods will be more useful. Some promising attempts have already been made in this direction (Holmes and Bruno, 2001; Jensen and Hein,
2002). Our algorithm can be used to compare exact likelihood calculations and approximations, and hence it might still be useful for developing new sampling techniques.

The underlying sequence evolution model has the drawback of treating only single nucleotide insertions and deletions, as was already mentioned in the original paper (Thorne et al., 1991). A model that can handle insertion and deletion events of whole subsequences is expected to perform better (Thorne et al., 1992), just as affine gaps improve alignments in score-based methods. Similarly, models that incorporate some structure information will probably result in better multiple alignments as well. We hope to incorporate these ideas in the future.

## APPENDIX A. PROOFS

The main goal of this section is to prove Lemma 1 from Section 7.2. In the first two subsections, we first establish some results which are used in the proof of Lemma 1.

## A.1. Unique legal sequences

To define a unique ordering of events for legal histories, we introduce a condition on sequences of events. We suppose the nodes of $T$ to be ordered. For purposes later on, we choose the specific ordering which satisfies, for every subtree $t$ of $T$,

$$
\{\text { left subtree of } t\}<\{\text { right subtree of } t\}<\operatorname{root} \text { of } t
$$

where left and right are arbitrary but fixed. An example of our choice of ordering is shown in Fig. 8. We also order (context, node) pairs as follows:

$$
(c, n)<\left(c^{\prime}, n^{\prime}\right) \text { if } c=I \text { and } c^{\prime}=F, \text { or } c=c^{\prime} \text { and } n<n^{\prime}
$$

To specify a unique ordering of events, we introduce condition $\Delta$ for a sequence of events $E=\left(e_{1}, \ldots, e_{N}\right)$ as follows:

$$
\begin{gathered}
\Delta(i): t_{i} \cap t_{i+1} \neq \varnothing \text { or }\left(c_{i}, r_{i}\right)<\left(c_{i+1}, r_{i+1}\right), \\
\Delta: \Delta(i) \text { holds, } \forall 1 \leq i<N .
\end{gathered}
$$

Lemma 2. Condition $\Delta$ holds for precisely one representative of each legal history.

Proof. First we prove existence. Let $k$ be the first index for which $\Delta(k)$ is false. Interchanging $e_{k}$ and $e_{k+1}$ is an allowed permutation and makes $\Delta(k)$ true, but might render $\Delta(k-1)$ false. In this case, interchanging $e_{k-1}$ and $e_{k}$ is allowed and renders $\Delta(k-1)$ and $\Delta(k)$ true, the latter because $\Delta(k-1)$ was true originally. Continuing backwards, eventually all $\Delta(i)$ for $i \leq k$ are true. By induction on $k$, we find an allowed permutation $\pi$ such that $\pi(E)$ satisfies $\Delta$.

We now prove uniqueness. Assume there are two representatives for which condition $\Delta$ holds, $E$ and $E^{\prime}$. Let $\pi$ be the allowed permutation such that $\pi(E)=E^{\prime}$. We can find an $i<j$ such that $\pi(j)+1=\pi(i)$


FIG. 8. A 6-leaved tree with its nodes labeled according to our convention.
(Christie, 1996). Since $\pi$ is an allowed permutation, $t_{i} \cap t_{j}=\varnothing$. $E^{\prime}$ satisfies $\Delta$, therefore $t_{j}$ and $t_{i}$ are part of the left and right subtree of a node $C_{0}$, respectively. Moreover, $i+1 \neq j$, because this would contradict $\Delta(i)$ for $E$.

We prove by induction that for each $k \geq 0$, a node $C_{k}$ exists such that $t_{j-k}$ is in the left subtree of $C_{k}, t_{i}$ is in the right subtree of $C_{k}$, and $t_{j}$ is contained in either. This means that there is no $k$ such that $i+1=j-k$; therefore, there is only one representative that satisfies the $\Delta$ condition.

We already proved that the induction hypothesis holds for $k=0$. Assume it holds for certain $k$. Four possible cases are to be considered.

1. $t_{j-k-1} \cap t_{i} \neq \varnothing$ and $t_{j-k-1} \cap t_{j-k} \neq \varnothing$. This implies that $C_{k} \in t_{j-k-1}$ and hence $t_{j-k-1} \cap t_{j} \neq \varnothing$. Since $\pi$ is an allowed permutation, $\pi(i)<\pi(j-k-1)$ and $\pi(j-k-1)<\pi(j)$, but this is impossible.
2. $t_{j-k-1} \cap t_{i} \neq \varnothing$ and $t_{j-k-1} \cap t_{j-k}=\varnothing$. This means that $t_{j-k-1}$ is on the right subtree of $C_{k}$, and therefore $r_{j-k-1}>r j-k$, but this contradicts the $\Delta$ condition, since all $c_{j}=F$ because the history is legal.
3. $t_{j-k-1} \cap t_{i}=\varnothing$ and $t_{j-k-1} \cap t_{j-k} \neq \varnothing$. This means that $t_{j-k-1}$ is on the left part of $C_{k}$. For $C_{k+1}=C_{k}$, the condition of the induction holds.
4. $t_{j-k-1} \cap t_{i}=\varnothing$ and $t_{j-k-1} \cap t_{j-k}=\varnothing$. In this case, $r_{j-k-1}<r_{j-k}<r_{i}$. The first inequality comes from the $\Delta$ condition, while the second one is from the induction hypothesis. However, these inequalities guarantee the existence of a proper $C_{k+1}$.

In properly ordered sequences, legal sequences may be recognized by looking at pairs of events.
Definition 2. A pair of events $\left(e_{1}, e_{2}\right)$ is called pairwise legal if $t_{1} \supseteq t_{2} \Rightarrow l_{1}\left(r_{2}\right) \neq-$.
Lemma 3. The set

$$
\left\{\left(e_{1}, \ldots, e_{N}\right) \mid \Delta \text { holds, }\left(e_{i}, e_{i+1}\right) \text { pairwise legal for } i=1, \ldots, N-1\right\}
$$

consists of precisely one representative of each legal history.
Proof. $\subseteq$ : Since $\Delta$ holds, at most one of each legal history is included. We prove that illegal histories are not included. Let $e_{1}, \ldots, e_{N}$ be an illegal sequence in the set. Let $e_{i}$ be the first event with $c_{i}=I$, that is, $r_{i} \notin S_{i}$. Since $c_{j}=F$ for $j<i$, in particular $S_{i-1}$ is legal. Suppose $t_{i-1} \cap t_{i} \neq \varnothing$. If $t_{i-1} \subset t_{i}$, then $r_{i} \notin S_{i}$ implies $r_{i} \notin S_{i-1}$, but $r_{i-1} \in S_{i-1}$ and $r_{i-1}$ is a descendant of $r_{i}$, in contradiction with $S_{i-1}$ legal. So $t_{i-1} \supseteq t_{i}$, but then $r_{i} \notin S_{i}$ implies ( $e_{i-1}, e_{i}$ ) is not pairwise legal. So $t_{i-1} \cap t_{i}=\varnothing$. Note that $c_{i-1}=F$ and $c_{i}=I$, so $\Delta$ is false, which is the required contradiction.
$\supseteq$ : For all legal histories there exists a sequence obeying $\Delta$ by Lemma 2, and all neighboring pairs in legal sequences are pairwise legal.

This result enables us to determine the sequences of events we want to sum over, by looking only at neighboring pairs. To use the inclusion/exclusion trick, we next have to characterize sequences of event pairs that do not obey our condition.

## A.2. Disallowed sequences and illegal sets

Definition 3. Suppose $S_{i}$ is a legal state, and suppose that for all $j=i, \ldots, k-1$ we have $\left(e_{j}, e_{j+1}\right)$ not pairwise legal or $\Delta(j)$ false. We call such (sub)sequences disallowed sequences.

Lemma 4. Let $\left(e_{i}, e_{i+1}, \ldots, e_{k}\right)$ be a disallowed sequence starting from state $S_{i}$. Then:
(a) The sequence $\left(c_{j}, r_{j}\right), i \leq j \leq k$, is strictly decreasing.
(b) All events $e_{j}$ with $c_{j}=F$ are disjoint.
(c) For each node $n$, there is at most one event $e_{j}$ with $n \in t_{j}$ and $l_{j}(n) \neq-$.

Note that statement (c) implies that there is at most one emission per leaf.

Proof. If $\left(e_{j}, e_{j+1}\right)$ is not pairwise legal, then $r_{j+1}$ is a descendant of $r_{j}$ and $l_{j}\left(r_{j+1}\right)=-$. This means that $r_{j}>r_{j+1}$ and $c_{j+1}=I$, so $\left(c_{j}, r_{j}\right)>\left(c_{j+1}, r_{j+1}\right)$. If $\Delta(j)$ is false, then it directly follows from the definition of condition $\Delta$ that $\left(c_{j}, r_{j}\right)>\left(c_{j+1}, r_{j+1}\right)$.

Now, that the sequence $\left(c_{j}, r_{j}\right)$ is strictly decreasing implies that, for some $m$, we have $c_{i}=c_{i+1}=$ $\cdots=c_{m}=F$ and $c_{m+1}=c_{m+2}=\cdots=c_{k}=I$. For $i \leq j<m$, the pair $\left(e_{j}, e_{j+1}\right)$ is pairwise legal since $c_{j+1}=F$; therefore, we must have $\Delta(j)$ false, which implies that $t_{j} \cap t_{j+1}=\varnothing$. Furthermore, $t_{j} \cap t_{j+1}=\varnothing$ and $r_{j}>r_{j+1}$ together imply that there exists a subtree $t$ of $T$ such that $t_{j}$ is in a right subtree of $t$, whereas $t_{j+1}$ is in a left subtree of $t$. Applying this reasoning sequentially, we conclude that for all $i \leq a<b \leq m$ there exists a subtree $t$ such that $t_{a}\left(t_{b}\right)$ is in a right (left) subtree of $t$. Hence, $t_{a} \cap t_{b}=\varnothing$ and it follows that all of $t_{i}, t_{i+1}, \ldots, t_{m}$ are pairwise disjoint.

Note that $S_{m+1}$ is a legal state since $S_{i}$ is a legal state and all events $e_{i}, e_{i+1}, \ldots, e_{m}$ are legal events. For every $m<j \leq k$, that the event $e_{j}$ is illegal implies that $r_{j} \notin S_{j}$. Now, suppose $S_{j} \cap t_{j} \neq \varnothing$. Then, $S_{j}$ must contain at least one descendant of $r_{j}$, which is possible only if, since $S_{m+1}$ is a legal state and $r_{j} \notin S_{j}$, at least one of $e_{m+1}, e_{m+2}, \ldots, e_{j-1}$ has a birth at a descendant node of $r_{j}$. But this contradicts the fact that $r_{m+1}>r_{m+2}>\cdots>r_{k}$, and therefore we must have $S_{j} \cap t_{j}=\varnothing$ for every $m<j \leq k$.

Finally, the results from the previous two paragraphs together imply (c).
Lemma 5. For a given legal initial state $S$, there is a one-to-one correspondence between illegal sets and disallowed sequences starting from $S$.

Proof. We prove the correspondence by establishing two injections.
From disallowed sequence to illegal set: The events with $c_{j}=F$ are pairwise disjoint, and their trees not the subtree of any other because the initial state is legal. The trees with $c_{j}=I$ may be subtrees of another $t_{k}$, in which case $l_{k}\left(r_{j}\right)=-$, so the events form an illegal set. There is only one disallowed sequence giving rise to this set, since the sequence $\left(c_{j}, r_{j}\right)$ corresponding to the disallowed sequence is ordered.

From illegal set to disallowed sequence: Take the events whose trees are not a subset of another and whose root is in $S$, sorted decreasing by root. Follow them by the other events, sorted decreasing by root. A disallowed sequence results.

## A.3. Proof of Lemma 1

For ease of reference, we here restate the lemma. We denote by $E(\mathbf{K})$ a set of legal sequences of events which emit the prefixes of the given sequences $A^{1}, \ldots, A^{k}$ up to position $\mathbf{K}=\left(\mathrm{K}_{1}, \ldots, \mathrm{~K}_{k}\right)$ and which contains precisely one representative of each history. The quantity of interest is the likelihood

$$
P(\mathbf{K}):=\sum_{\left(e_{1}, \ldots, e_{N}\right) \in E(\mathbf{K})} p\left(e_{1}\right) \cdots p\left(e_{N}\right) .
$$

Lemma 1. The likelihood $P(\mathbf{K})$ satisfies the following equation:

$$
\begin{aligned}
P(\mathbf{K})= & \sum_{e} P\left(\mathbf{K}-\mathbf{v}_{e}\right) p(e) \\
& -\sum_{\left\{e_{1}, e_{2}\right\} \text { illegal }} P\left(\mathbf{K}-\mathbf{v}_{e_{1}}-\mathbf{v}_{e_{2}}\right) p\left(e_{1}\right) p\left(e_{2}\right) \\
& +\sum_{\left\{e_{1}, e_{2}, e_{3}\right\} \text { illegal }} P\left(\mathbf{K}-\mathbf{v}_{e_{1}}-\mathbf{v}_{e_{2}}-\mathbf{v}_{e_{3}}\right) p\left(e_{1}\right) p\left(e_{2}\right) p\left(e_{3}\right)
\end{aligned}
$$

Here, $\mathbf{v}_{e}$ is the emission vector corresponding to the event $e$, that is, a $k$-dimensional vector whose component corresponding to each leaf is 0 if the leaf is labeled - , or 1 otherwise.

Proof. Let $S(e)$ denote the state after event $e$ in state $S$. Define $P_{S}(\mathbf{K})$ to be the likelihood of emitting legal sequences up to position $\mathbf{K}$ and ending up in state $S$. We allow illegal states $S$, in which case $P_{S}(K)=0$ always. For brevity, we denote by $C\left(S_{0}, e_{0}, S_{1}, e_{1}, \ldots, e_{n-1}, S_{n}\right)$ the condition $S_{0}\left(e_{0}\right)=$ $S_{1} \wedge \ldots \wedge S_{n-1}\left(e_{n-1}\right)=S_{n}$, and by $\operatorname{Dis}\left(S, e_{0}, \ldots, e_{n}\right)$ the condition that $\left(e_{0}, \ldots, e_{n}\right)$ is a disallowed sequence when starting from state $S$. Then

$$
\begin{aligned}
& P_{S}(\mathbf{K})=\sum_{\left(S^{\prime}, e\right): C\left(S^{\prime}, e, S\right)} P_{S^{\prime}}\left(\mathbf{K}-\mathbf{v}_{e}\right) p(e) \\
& -\sum_{\substack{\left(S^{\prime}, e, S^{\prime \prime}, e^{\prime}\right): \\
C\left(S^{\prime \prime}, e^{\prime}, S^{\prime}, e, S, S\right) \wedge \operatorname{Dis}\left(S^{\prime \prime}, e^{\prime}, e\right)}} P_{S^{\prime \prime}}\left(\mathbf{K}-\mathbf{v}_{e}-\mathbf{v}_{e^{\prime}}\right) p\left(e^{\prime}\right) p(e) \\
& +\sum_{\substack{\left(S^{\prime}, e, S^{\prime \prime}, e^{\prime}, S^{\prime \prime \prime}, e^{\prime \prime}\right): \\
C\left(S^{\prime \prime \prime}, e^{\prime \prime}, S^{\prime \prime}, e^{\prime},, e^{\prime}, e, S\right) \wedge \\
\wedge \operatorname{Dis}\left(S^{\prime \prime}, e^{\prime \prime}, e^{\prime}, e\right)}} P_{S^{\prime \prime}}\left(\mathbf{K}-\mathbf{v}_{e}-\mathbf{v}_{e^{\prime}}-\mathbf{v}_{e^{\prime \prime}}\right) p\left(e^{\prime \prime}\right) p\left(e^{\prime}\right) p(e) \\
& \text { - ... }
\end{aligned}
$$

This recursion sums over all sequences of events which do not have disallowed subsequences, since all terms in the first line that include a disallowed subsequence (of length 2 ) are subtracted in the second line; this also subtracts terms involving disallowed subsequences of length 3 that were not included in the first place, which are added in again in the third line, and so on. By Lemma 3, this implies that we indeed sum over all sequences of legal events (which end up in state $S$ ). Note that for illegal states $S, P_{S}(\mathbf{K})=0$ always, so we may define $\operatorname{Dis}\left(S, e_{0}, \ldots\right)$ arbitrarily for illegal states $S$.

Now we sum the recursion equation over all states $S$. The left-hand side becomes $P(\mathbf{K})$. The first summation on the right-hand side turns into

$$
\sum_{S} \sum_{\left(S^{\prime}, e:: S^{\prime}(e)=S\right.} P_{S^{\prime}}\left(\mathbf{K}-\mathbf{v}_{e}\right) p(e)=\sum_{\left(S^{\prime}, e\right)} P_{S^{\prime}}\left(\mathbf{K}-\mathbf{v}_{e}\right) p(e)=\sum_{e} P\left(\mathbf{K}-\mathbf{v}_{e}\right) p(e),
$$

which is the desired term. Similarly, the second summation becomes

$$
\sum_{\left(S^{\prime \prime}, e^{\prime}, e\right): \operatorname{Dis}\left(S^{\prime \prime}, e^{\prime}, e\right)} P_{S^{\prime \prime}}\left(\mathbf{K}-\mathbf{v}_{e}-\mathbf{v}_{e^{\prime}}\right) p\left(e^{\prime}\right) p(e)
$$

Observe that the summand is independent of the ordering of events $e^{\prime}, e$. Hence, we can also sum over illegal sets of events and over $S^{\prime \prime}$, using Lemma 5. This turns the summation into

$$
\sum_{S^{\prime \prime}} \sum_{\left\{e^{\prime}, e\right\} \text { illegal }} P_{S^{\prime \prime}}\left(\mathbf{K}-\mathbf{v}_{e}-\mathbf{v}_{e^{\prime}}\right) p\left(e^{\prime}\right) p(e)=\sum_{\left\{e^{\prime}, e\right\} \text { illegal }} P\left(\mathbf{K}-\mathbf{v}_{e}-\mathbf{v}_{e^{\prime}}\right) p\left(e^{\prime}\right) p(e),
$$

which is the desired term. All other summations are dealt with analogously. Note that only finitely many summands contribute to the recursion, since illegal sets have a bounded number of elements.

## APPENDIX B. STATES FOR THE HJP RECURSION

For a fixed number $k$ of sequences, the total number of Markov states ${ }^{1}$ in the HJP recursion depends on the tree topology. Let $T$ be a $k$-leaved rooted binary tree. Let $w(n)$ denote the number of leaves joined to the internal node $n \in T$, and let $q(T)$ be the number of internal nodes $n \in T$ with $w(n)=2$.

[^0]If $q(T)=2$ and the $\operatorname{root}^{2} r$ of the tree has $w(r)=2$, then the number of Markov states in the HJP recursion for such a tree is equal to

$$
\left(2^{k}+\sum_{j=0}^{k-1} 2^{j}\right)+2\left[\sum_{n=0}^{k-4}\left(2^{k-2-n}+\sum_{j=0}^{k-2-n-1} 2^{j}\right)\right]+\sum_{n=1}^{k-4} \sum_{m=0}^{n-1}\left(2^{n-m}+\sum_{j=0}^{n-m-1} 2^{j}\right) .
$$

As the tree topology we just considered has the least number of states, the above formula is a lower bound on the number of Markov states for any $k$-leaved rooted binary tree of arbitrary topology.

For $q(T) \gg 2$, the total number of states chiefly depends on whether or not the nodes $n$ with $w(n)=2$ are assigned "\#" (which stands for having a nucleotide). The number of states associated to such choices is

$$
\begin{align*}
\sum_{m=0}^{q}\left[\binom{q}{m}\left(2^{k-2 m}+\sum_{j=0}^{k-2 m-1} 2^{j}\right)\right] & =2^{k} \sum_{m=0}^{q}\binom{q}{m}\left(2^{-2 m}+\frac{2^{k-2 m}-2}{2^{k}}\right) \\
& \approx 2 \cdot 2^{k}\left(\frac{5}{4}\right)^{q} . \tag{5}
\end{align*}
$$

For $q(T)=\lfloor k / 2\rfloor$, i.e., in the case of a "balanced" tree, the expression in (5) is of order $O\left(\sqrt{5}^{k}\right)$.

## APPENDIX C. A HIDDEN MARKOV MODEL

In many cases, it is useful to obtain the maximum likelihood alignment using the Viterbi algorithm. This cannot be done with the one-state recursion, and we need to use the full Markov chain corresponding to the TKF91 model on a tree. We give two versions here. The first is not properly a Markov chain, in that the sum of exit probabilities from states do not sum to 1 , but it does assign the right probability to full paths through the chain. This is enough for Viterbi to work. Secondly, we show how to turn this simpler chain into a proper Markov chain, which can be used for sampling alignments.

We have used the word state in a different way than is usual for Markov chains. In this section, we write State with a capital letter if we use it in the Markov chain sense.

We view our chain as a Mealy machine (Durbin et al., 1998), i.e., events and their emissions correspond to transitions, and paths through the chain correspond to histories. Lemma 3 gives a one-to-one correspondence between histories and paths through our chain. Given a sequence of events $e_{0}, \ldots, e_{1}$, a new event $e_{i+1}$ is allowed if $\left(e_{i}, e_{i+1}\right)$ is pairwise legal, and if condition $\Delta(i)$ holds. Both can be determined if we know the state $S_{i+1}$ and the root of $e_{i}^{t}$, so as a first approximation to our Markov State we can use the pair $\left(S_{i+1}, e_{i}^{t}\right)$. We say that the nodes $r(t)$ and $r\left(t^{\prime}\right)$ are directly related if $t \cap t^{\prime} \neq \varnothing$, i.e., one is the ancestor of the other, or they coincide. Then, the transition rule is that from State ( $S, v$ ), event $e$ is allowed if (1) $r\left(e^{t}\right) \in S$ and (2) $r\left(e^{t}\right)$ is directly related to $v$ or $r\left(e^{t}\right) \geq v$. Now, if we set

$$
S_{i+1}^{\prime}:=S_{i+1} \backslash\left\{n \mid n \text { indirectly related to } r\left(e_{i}^{t}\right) \text { and } n<r\left(e_{i}^{t}\right)\right\},
$$

the transition rule simplifies to: $e_{i+1}$ is allowed if $r\left(e_{i+1}^{t}\right) \in S_{i+1}^{\prime}$, and it can be shown that the update rule

$$
\begin{aligned}
S_{i+1}^{\prime}:= & \left(S_{i}^{\prime} \backslash\left(e_{i}^{t} \cup\left\{n \in S_{i}^{\prime} \mid n \text { indirectly related to } r\left(e_{i}^{t}\right) \text { and } n<r\left(e_{i}^{t}\right)\right\}\right)\right) \\
& \cup\left\{n \in e_{i}^{t} \mid e_{i}^{l}(n) \neq-\right\}
\end{aligned}
$$

correctly updates $S_{i}^{\prime}$ to $S_{i+1}^{\prime}$.
Let $p(e)$ be the probability factor associated to event $e$. The States $S^{\prime}$ and transition and update rules give rise to a graph with states $S^{\prime}$ as vertices, and we associate the probability factors $p(e)$ to the edges.

[^1]To complete the graph, we introduce the initial and end States $A$ and $\Omega$. Furthermore, we introduce a single edge (with $p=\prod_{n \in T} I_{n}$ ) between $A$ and the State $T$ and join each State except $A$ to $\Omega$ by a single edge (with $p=1$ ). In this way, the product of probability factors along any path from $A$ to $\Omega$ is just the probability of the associated sequence of events, or history, under the TKF91 model. Since by our choice of States, each permissible history is counted exactly once, the probabilities sum up to 1 .

For sampling, it is necessary to have a proper Markov chain, i.e., have the probabilities of outgoing edges sum to 1 , and the graph just defined can be turned into one by a scalar base change. Let $P_{S}$ be the probability of being in State $S$, i.e., the sum of probabilities of all paths ending in $S$. Then, the state equation can be written as $P_{S}=\sum_{S^{\prime}} P_{S^{\prime}} t_{S^{\prime} S}$, where $t_{S^{\prime} S}$ is the transition matrix. If we introduce new variables $\hat{P}_{S}=f_{S} P_{S}$, then we can write $\hat{P}_{S}=\sum_{S^{\prime}} \hat{P}_{S^{\prime}} \hat{t}_{S^{\prime}} S$, where $\hat{t}_{S^{\prime} S}=t_{S^{\prime} S} f_{S} / f_{S^{\prime}}$. The graph becomes a Markov chain if $\sum_{S} \hat{t}_{S^{\prime} S}=1$. A tedious but straightforward calculation shows that $f_{S}=\left(\prod_{n \in S} I_{n}\right)^{-1}$ solves this equation, and the transition probabilities become

$$
S_{i} \xrightarrow{e} S_{i+1}: \frac{f_{S_{i+1}}}{f_{S_{i}}} p(e),
$$

where $f_{A}=f_{\Omega}=1$.

## ACKNOWLEDGMENTS

The authors wish to thank Rune Lyngsø and Alexei Drummond for helpful discussions and the referees for their thorough reading and detailed comments. Just before publication, Jens Ledet Jensen sent us an alternative proof of the results presented here. ${ }^{3}$ This research is supported by EPSRC under grant HAMJW and by MRC under grant HAMKA. Y.S. Song is supported in part by a grant from the Danish Natural Science Foundation (SNF-5503-13370).

## REFERENCES

Christie, D.A. 1996. Sorting permutations by block-interchanges. Inf. Proc. Letters 60, 165-169.
Durbin, R., Eddy, S., Krogh, A., and Mitchison, G. 1998. Biological Sequence Analysis, Cambridge University Press. Eddy, S. 2001. HMMER: Profile hidden markov models for biological sequence analysis (hmmer.wustl.edu/).
Felsenstein, J. 1981. Evolutionary trees from DNA sequences: A maximum likelihood approach. J. Mol. Evol. 17, 368-376.
Felsenstein, J. 2001. Phylip: Phylogeny Inference Package, University of Washington.
Goldman, N. 1998. Effects of sequence alignment procedures on estimates of phylogeny. BioEssays 20, 287-290.
Hein, J. 2001. An algorithm for statistical alignment of sequences related by a binary tree. Pac. Symp. Biocomp., 179-190.
Hein, J., Jensen, J.L., and Pedersen, C.N.S. 2002. Recursions for statistical multiple alignment. Technical Report 425, Dept. of Theor. Stat., Univ. of Aarhus.
Hein, J., Wiuf, C., Knudsen, B., Møller, M.B., and Wibling, G. 2000. Statistical alignment: Computational properties, homology testing and goodness-of-fit. J. Mol. Biol. 302, 265-279.
Hirshberg, D.S. 1977. Algorithms for the longest common subsequence problem. J. ACM 24, 664-675.
Holmes, I., and Bruno, W.J. 2001. Evolutionary HMMs: A Bayesian approach to multiple alignment. Bioinformatics 17(9), 803-820.
Jensen, J.L., and Hein, J. 2002. Gibbs sampler for statistical multiple alignment. Technical Report 429, Dept. of Theor. Stat., Univ. of Aarhus.
Jukes, T.H., and Cantor, C.R. 1969. Mammalian Protein Metabolism, chapter Evolution of protein molecules, 21-132, Academic Press.
Lee, M.S.Y. 2001. Unalignable sequences and molecular evolution. Trends in Ecology and Evolution 16, 681-685.

[^2]Miklós, I. 2002. An improved algorithm for statistical alignment of sequences related by a star tree. Bull. Math. Biol. 64, 771-779.
Mitchison, G. 1999. A probabilistic treatment of phylogeny and sequence alignment. J. Mol. Evol. 49, 11-22.
Needleman, S.B., and Wunsch, C.D. 1970. A general method applicable to the search for similarities in the amino acid sequences in two proteins. J. Mol. Biol. 48, 443-453.
Sankoff, D., and Kruskall, J.B. 1983. Time Warps, String Edits and Macromolecules: The Theory and Practice of Sequence Comparison, Addison-Wesley, Reading, PA.
Steel, M., and Hein, J. 2001. Applying the Thorne-Kishino-Felsenstein model to sequence evolution on a star-shaped tree. Appl. Math. Let. 14, 679-684.
Thorne, J.L., Kishino, H., and Felsenstein, J. 1991. An evolutionary model for maximum likelihood alignment of DNA sequences. J. Mol. Evol. 33, 114-124.
Thorne, J.L., Kishino, H., and Felsenstein, J. 1992. Inching toward reality: An improved likelihood model of sequence evolution. J. Mol Evol. 34, 3-16.
Whelan, S., Lió, P., and Goldman, N. 2001. Molecular phylogenetics: State-of-the-art methods for looking into the past. Trends Genet. 17, 262-272.

Address correspondence to:
G.A. Lunter

Department of Statistics
University of Oxford
Oxford, OX1 3TG, UK
E-mail: lunter@stats.ox.ac.uk


[^0]:    ${ }^{1}$ We refer the reader to Hein et al. (2002) for their definition of "state."

[^1]:    ${ }^{2}$ In Hein et al. (2002), the root of a tree $T$ is one of the internal nodes $n \in T$ of degree 3 .

[^2]:    ${ }^{3}$ This proof has appeared in: Lunter, G.A., Miklos, I., Drummond, A.J., Jensen, J.L., Hein, J. 2003. Bayesian Phylogenetic Inference under a Statistical Insertion-Deletion Model. In: Benson, G., and Page, R. (eds.), WABI 2003, LNBI 2812, 228-244. (Note added in proof.)

