5

Profile HMMs for sequence families

So far we have concentrated on the intrinsic properties of single sequences, such as CpG islands in DNA, or on pairwise alignment of sequences. However, functional biological sequences typically come in families, and many of the most powerful sequence analysis methods are based on identifying the relationship of an individual sequence to a sequence family. Sequences in a family will have diverged from each other in their primary sequence during evolution, having separated either by a duplication in the genome, or by speciation giving rise to corresponding sequences in related organisms. In either case they normally maintain the same or a related function. Therefore, identifying that a sequence belongs to a family, and aligning it to the other members, often allows inferences about its function.

If you already have a set of sequences belonging to a family, you can perform a database search for more members using pairwise alignment with one of the known family members as the query sequence. To be more thorough, you could even search with all the known members one by one. However, pairwise searching with any one of the members may not find sequences distantly related to the ones you have already. An alternative approach is to use statistical features of the whole set of sequences in the search. Similarly, even when family membership is clear, accurate alignment can be often be improved significantly by concentrating on features that are conserved in the whole family.

How, in brief, do we identify such features? Just as a pairwise alignment captures much of the relationship between two sequences, a multiple alignment can show how the sequences in a family relate to each other. Figure 5.1 shows a multiple alignment of seven sequences from the large globin family (hundreds of globin sequences are available in the protein sequence databases). The three dimensional structure has been obtained for each protein in the alignment shown, and the sequences have been aligned on the basis of aligning the eight alpha helices of the conserved globin fold, and also on the basis of aligning certain key residues in the sequences, such as two conserved histidines (H) which are the residues which interact with an oxygen-binding heme prosthetic group in the globin active site.

It is clear that some positions in the globin alignment are more conserved than others. In general the helices are more conserved than the loop regions between

```
Helix                     AAAAAAAAAAAAAAAA  BBBBBBBBBBBBBBBBBCCCCCCCCCCCC
HBA_HUMAN    ---------VLSPADKTNVKAAWGKVGA--HAGEYGAEALERMFLSFPTTKTYFPHF
HBB_HUMAN    --------VHLTPEEKSAVTALWGKV----NVDEVGGEALGRLLVVYPWTQRFFESF
MYG_PHYCA    ---------VLSEGEWQLVLHVWAKVEA--DVAGHGQDILIRLFKSHPETLEKFDRF
GLB3_CHITP   ----------LSADQISTVQASFDKVKG-----DPVGILYAVFKADPSIMAKFTQF
GLB5_PETMA   PIVDTGSVAPLSAAEKTKIRSAWAPVYS--TYETSGVDILVKFFTSTPAAQEFFPKF
LGB2_LUPLU   --------GALTESQAALVKSSWEEFNA--NIPKHTHRFFILVLEIAPAAKDLFS-F
GLB1_GLYDI   ---------GLSAAQRQVIAATWKDIAGADNGAGVGKDCLIKFLSAHPQMAAVFG-F
Consensus             Ls....  v a W kv .  .    g . L.. f . P .   F  F


Helix        DDDDDDDEEEEEEEEEEEEEEEEEEEEEE              FFFFFFFFFFFF
HBA_HUMAN    -DLS-----HGSAQVKGHGKKVADALTNAVAHV---D--DMPNALSALSDLHAHKL-
HBB_HUMAN    GDLSTPDAVMGNPKVKAHGKKVLGAFSDGLAHL---D--NLKGTFATLSELHCDKL-
MYG_PHYCA    KHLKTEAEMKASEDLKKHGVTVLTALGAILKK----K-GHHEAELKPLAQSHATKH-
GLB3_CHITP   AG-KDLESIKGTAPFETHANRIVGFFSKIIGEL--P---NIEADVNTFVASHKPRG-
GLB5_PETMA   KGLTTADQLKKSADVRWHAERIINAVNDAVASM--DDTEKMSMKLRDLSGKHAKSF-
LGB2_LUPLU   LK-GTSEVPQNNPELQAHAGKVFKLVYEAAIQLQVTGVVVTDATLKNLGSVHVSKG-
GLB1_GLYDI   SG----AS---DPGVAALGAKVLAQIGVAVSHL--GDEGKMVAQMKAVGVRHKGYGN
Consensus    . t    .. . v..Hg kv. a   a...l   d   .a l. l   H  .


Helix        FFGGGGGGGGGGGGGGGGGGGG    HHHHHHHHHHHHHHHHHHHHHHHHHHHH
HBA_HUMAN    -RVDPVNFKLLSHCLLVTLAAHLPAEFTPAVHASLDKFLASVSTVLTSKYR------
HBB_HUMAN    -HVDPENFRLLGNVLVCVLAHHFGKEFTPPVQAAYQKVVAGVANALAHKYH------
MYG_PHYCA    -KIPIKYLEFISEAIIHVLHSRHPGDFGADAQGAMNKALELFRKDIAAKYKELGYQG
GLB3_CHITP   --VTHDQLNNFRAGFVSYMKAHT--DFA-GAEAAWGATLDTFFGMIFSKM-------
GLB5_PETMA   -QVDPQYFKVLAAVIADTVAAG--------DAGFEKLMSMICILLRSAY-------
LGB2_LUPLU   --VADAHFPVVKEAILKTIKEVVGAKWSEELNSAWTIAYDELAIVIKKEMNDAA---
GLB1_GLYDI   KHIKAQYFEPLGASLLSAMEHRIGGKMNAAAKDAWAAAYADISGALISGLQS-----
Consensus    v.  f  l . ..  ....    f   . aa. k. .    l sky
```

Figure 5.1 *An alignment of seven globins from Bashford, Chothia & Lesk [1987].* To the left is the protein identifier in the SWISS-PROT database [Bairoch & Apweiler 1997]. The eight alpha helices are shown as A–H above the alignment. A consensus line below the alignment indicates residues that are identical among at least six of the seven sequences in upper case, ones identical in four or five sequences in lower case, and positions where there is a residue identical in three sequences with a dot.

them, and certain residues are particularly strongly conserved. When identifying a new sequence as a globin, it would be desirable to concentrate on checking that these more conserved features are present. How to obtain and use such information will be the subject of this chapter.

As might be expected, our approach to consensus modelling will be to make a probabilistic model. In particular, we will develop a particular type of hidden Markov model well suited to modelling multiple alignments. We call these *profile HMMs* after standard *profiles*, which are closely related non-probabilistic structures introduced previously for the same purpose by Gribskov, McLachlan & Eisenberg [1987]. Profile HMMs are probably the most popular application of hidden Markov models in molecular biology at the moment [Eddy 1996].

We will assume for the purposes of this chapter that we are given a correct multiple alignment, from which we will build a model that can be used to find and score potential matches to new sequences. The multiple alignment could

be built from structural information, like the globin alignment shown here, or it could come from a sequence-based alignment procedure, such as those discussed in Chapter 6.

Much of this chapter makes use of the theory presented in Chapter 3 for general HMMs. The most important algorithms will be presented again in the specific form relevant to profile HMMs. There is also an extensive discussion of how to estimate optimal probability parameters from multiple sequence alignments.

5.1   Ungapped score matrices

One general feature of protein family multiple alignments, which can be seen in Figure 5.1, is that gaps tend to line up with each other, leaving solid blocks where there are no insertions or deletions in any of the sequences. We will start by considering models for these ungapped regions.

As an example, consider the E helix of Figure 5.1. A natural probabilistic model for such a region would be to specify independent probabilities $e_i(a)$ of observing amino acid a in position i (we use letter e because these will turn out to be the *emission probabilities* of the hidden Markov model when we introduce gaps). The probability of a new sequence x according to this model is then

$$P(x|M) = \prod_{i=1}^{L} e_i(x_i),$$

where L is the length of the block, 21 in this case. As usual, we are in fact more interested in the ratio of this probability to the probability of x under a random model, and so to test for membership in the family we evaluate the log-odds ratio

$$S = \sum_{i=1}^{L} \log \frac{e_i(x_i)}{q_{x_i}}.$$

The values $\log \frac{e_i(a)}{q_a}$ behave like elements in a score matrix $s(a,b)$, where the second index is position i, rather than amino acid b. For this reason, such an approach is known as a *position specific score matrix* (PSSM). A PSSM can be used to search for a match in a longer sequence x of length N by evaluating the score S_j for each starting point j in x from 1 to $N - L + 1$, where L is the length of the PSSM.

5.2   Adding insert and delete states to obtain profile HMMs

Although a PSSM captures some conservation information, it is clearly an inadequate representation of all the information in a multiple alignment of a protein

family. We have to find some way to take account of gaps. It is possible to combine the scores of multiple ungapped block models, and this is the approach taken by Henikoff & Henikoff [1991] in the BLOCKS database. However, we will pursue here the aim of developing a single probabilistic model for the whole extent of the alignment.

One approach is to allow gaps at each position in the alignment, using the same gap score $\gamma(g)$ at each position, as in pairwise alignment. However, this is also ignoring information, because the alignment gives us explicit indications of where gaps are more and less likely. We want to capture this information to give us position sensitive gap scores, just as the emission probabilities gave us position sensitive substitution scores.

The approach we take is to build a hidden Markov model (HMM), with a repetitive structure of states, but different probabilities in each position. This will provide a full probabilistic model for sequences in the sequence family. We start off by observing that the PSSM can be viewed as a trivial HMM with a series of identical states that we will call *match* states, separated by transitions of probability 1.

Alignment is trivial because there is no choice of transitions. We rename the emission probabilities for the match states to $e_{M_i}(a)$.

The next step is to deal with gaps. We must treat insertions and deletions separately. To handle insertions, i.e. portions of x that do not match anything in the model, we introduce a set of new states I_i, where I_i will be used to match insertions after the residue matching the ith column of the multiple alignment. The I_i have emission distribution $e_{I_i}(a)$, but these are normally set to the background distribution q_a, just as for seeing an unaligned inserted residue in a pairwise alignment. We need transitions from M_i to I_i, a loop transition from I_i to itself, to accommodate multi-residue insertions, and a transition back from I_i to M_{i+1}. Here is a single insert state of this kind:


We denote insert states in our diagrams by diamonds. The log-odds cost of an insert is the sum of the costs of the relevant transitions and emissions. Assuming that $e_{I_i}(a) = q_a$ as described above, there is no log-odds contribution from the emission, and the score of a gap of length k is

$$\log a_{M_j I_j} + \log a_{I_j M_{j+1}} + (k-1)\log a_{I_j I_j}.$$

From this you can see that the type of insert state shown corresponds to an affine gap scoring model.

Deletions, i.e. segments of the multiple alignment that are not matched by any residue in x, could be handled by forward 'jump' transitions between non-neighbouring match states:

However, to allow arbitrarily long gaps in a long model this way would require a lot of transitions. Instead we introduce silent states D_j as described in Section 3.4:

Because the silent states do not emit any residues, it is possible to use a sequence of them to get from any match state to any later one, between two residues in the sequence. The cost of a deletion will then be the sum of the costs of an $M \rightarrow D$ transition followed by a number of $D \rightarrow D$ transitions, then a $D \rightarrow M$ transition. This is at first sight exactly analogous to the cost of an insert, although the path through the model looks different. In detail, it is possible that the $D \rightarrow D$ transitions will have different probabilities, and hence contribute differently to the score, whereas all the $I \rightarrow I$ transitions for one insert involve the same state, and so are guaranteed to have the same cost.

The full resulting HMM has the structure shown in Figure 5.2. This form of model, which we call a profile HMM, was first introduced in Haussler *et al.* [1993] and Krogh *et al.* [1994]. We have added transitions between insert and delete states, as they did, although these are usually very improbable. Leaving them out has negligible effect on scoring a match, but can create problems when building the model.

Figure 5.2 *The transition structure of a profile HMM. We use diamonds to indicate the insert states and and circles for the delete states.*

We I
used
it is
whic

Lı
then
com
sequ
to in
valu
prol
prol

Ir
tainı
one
find
sam
scril
stan
belc

Altl
proș
intr
stru
to c
the
sus

Τ
mu
illu
me·

A r
& I

Eis

Profile HMMs generalise pairwise alignment

We have seen how the costs of using gap states in a profile HMM mirror those used in pairwise alignment with affine gaps. To help make clear the relationship, it is useful to consider the degenerate case where the multiple alignment from which we build the HMM contains just one sequence.

Let us compare Figure 5.2 with Figure 4.2. If we call the example sequence y, then Figure 5.2 is an unrolled version of Figure 4.2, with the y_j emissions each coming from a separate copy of the pair HMM. The states M_j correspond to a sequence of match states M, the I_j to corresponding incarnations of X, and the D_j to incarnations of Y. To achieve as close a correspondence as possible, the natural values for the match emission probabilities $e_{M_i}(a)$ are $p_{y_i a}/q_{y_i}$, the conditional probabilities of seeing a given y_i in a pairwise alignment, and for the transition probabilities $a_{M_i I_i} = a_{M_i D_{i+1}} = \delta$ and $a_{I_i I_i} = a_{D_i D_{i+1}} = \varepsilon$ for all i.

In formal terms our profile HMM is effectively the hidden Markov model obtained by conditioning the pair HMM of Figure 4.2 on emitting sequence y as one of the sequences in its alignment. Because of this, the Viterbi equations for finding the most probable alignment of x to our profile HMM are essentially the same as those for the most probable alignment of x and y to the pair HMM described in Chapter 4. If we convert them into log-odds ratio form we recover our standard affine gap cost pairwise alignment equations of (2.16), as we will see below. Any differences are due to slightly different Begin and End arrangements.

5.3 Deriving profile HMMs from multiple alignments

Although it is nice to see that the profile HMM is doing the same sort of dynamic programming as we have used before for pairwise alignment, this is not why we introduced them. The key idea behind profile HMMs is that we can use the same structure as shown in Figure 5.2, but set the transition and emission probabilities to capture specific information about each position in the multiple alignment of the whole family. Essentially, we want to build a model representing the consensus sequence for the family, not the sequence of any particular member.

There are a number of different ways to derive the parameter values from a multiple alignment of the sequences in the family. To provide an example for illustrating these methods, Figure 5.3 shows a short section of the globin alignment shown in Figure 5.1.

Non-probabilistic profiles

A model similar to the profile HMM was first introduced by Gribskov, McLachlan & Eisenberg [1987] who coined the name 'profile' (see also Gribskov, Lüthy & Eisenberg [1990]). However, they did not have an underlying probabilistic model,

```
HBA_HUMAN     ...VGA--HAGEY...
HBB_HUMAN     ...V----NVDEV...
MYG_PHYCA     ...VEA--DVAGH...
GLB3_CHITP    ...VKG------D...
GLB5_PETMA    ...VYS--TYETS...
LGB2_LUPLU    ...FNA--NIPKH...
GLB1_GLYDI    ...IAGADNGAGV...
              ***  *****
```

Figure 5.3 *Ten columns from the multiple alignment of seven globin protein sequences shown in Figure 5.1. The starred columns are ones that will be treated as 'matches' in the profile HMM.*

but rather directly assigned position specific scores for each match state and gap penalty, for use in standard 'best match' dynamic programming. They set the scores for each consensus position to the averages of the standard substitution scores from all the residues seen in the corresponding multiple alignment column. For example, they would set the score for residue a in column 1 of our example to be

$$\tfrac{5}{7}s(V,a) + \tfrac{1}{7}s(F,a) + \tfrac{1}{7}s(I,a)$$

where $s(a,b)$ is the standard substitution matrix. They also set gap penalties for each column using a heuristic equation that decreased the cost of a gap (either insertion or deletion) according to the length of the longest gap observed in the multiple alignment spanning the column.

Although this seems an intuitively obvious way to combine information, and it has been used effectively by many people for finding new members of families, it does produce anomalies. For example, column 1 is much more strongly conserved than column 2 in the example shown in Figure 5.3, but the information in column 1 will be smeared out just as much by the substitution matrix as that in column 2. If we had an alignment with 100 sequences, all with a cysteine (C) at some position, then the implicit probability distribution for that column for an 'average' profile would be exactly the same as would be derived from a single sequence. This does not correspond to our expectation that the likelihood of a cysteine should go up as we see more confirming examples.

In addition to these observations about substitution scores, the scores for gaps do not behave as expected. For example, from the alignment in Figure 5.3 the score for a deletion would be set to be the same in column 2, where there is a deletion in one sequence, HBB_HUMAN, as in column 4, where there is a deletion opening in five of the seven sequences. It would be more reasonable to set the probability of a new gap opening to be higher in column 4.

Changes have been made to non-probabilistic profiles to address these and

other problems [Thompson, Higgins & Gibson 1994b; Gribskov & Veretnik 1996], and we shall return to some of these later.

Basic profile HMM parameterisation

Let us turn back to hidden Markov model profiles. Like all HMMs, these have emission and transition probabilities. Assuming that these probabilities are non-zero, a profile HMM can model any possible sequence of residues from the given alphabet. It therefore defines a probability distribution over the whole space of sequences. The aim of the parameterisation process it to make this distribution peak around members of the family.

The parameters we have available to control the shape of the distribution are the values of the probabilities, and also the length of the model. There is a lot to say about setting these optimally. We give here the basic methods from Krogh *et al.* [1994]. After sections on database searching and variants for local alignment, we will return to an extended discussion of alternative parameter estimation techniques.

The choice of length of the model corresponds more precisely to a decision on which multiple alignment columns to assign to match states, and which to assign to insert states. The profile HMM we derived above from the single sequence y had a match state for each residue y_i. However, looking at Figure 5.3 it seems clear that the consensus sequence for this region should only have eight residues, and that the two non-starred residues in GLB1_GLYDI should be treated as an insertion with respect to the consensus. For the time being we will use a heuristic rule to decide which columns should correspond to match states, and which to inserts. A simple rule that works well is that columns that are more than half gap characters should be modelled by inserts.

The second problem is how to assign the probability parameters. We regard the alignment as providing a set of independent samples of alignments of sequences x to our HMM. Since the alignments are given, we can estimate the parameters directly using equations (3.18) from Section 3.3. We just count up the number of times each transition or emission is used, and assign probabilities according to

$$a_{kl} = \frac{A_{kl}}{\sum_{l'} A_{kl'}} \quad \text{and} \quad e_k(a) = \frac{E_k(a)}{\sum_{a'} E_k(a')}$$

where k and l are indices over states, and a_{kl} and e_k are the transition and emission probabilities, and A_{kl} and E_k are the corresponding frequencies.

In the limit of having a very large number of sequences in our training alignment, this will give an accurate and consistent estimate of the probabilities. However, it has problems when there are only a few sequences. A major difficulty is that some transitions or emissions may not be seen in the training alignment, and so would acquire zero probability, which would mean they would never be

Figure 5.4 *A hidden Markov model derived from the small alignment shown in Figure 5.3 using Laplace's rule. Emission probabilities are shown as bars opposite the different amino acids for each match state, and transition probabilities are indicated by the thickness of the lines. The* I → I *transition probabilities times 100 are shown in the insert states. (Figure generated automatically using the* SAM *package.)*

allowed in the future. More broadly, we are not using any previous knowledge about protein alignments, as the earlier non-probabilistic methods did implicitly, by using an independently derived substitution matrix. As a minimal approach to avoid zero probabilities, we can add pseudocounts to the observed frequencies (as in Chapters 1 and 3). The simplest pseudocount method is Laplace's rule: to add one to each frequency. We discuss better ways to choose the pseudocount values, and other approaches to estimating the parameters, at greater length below in Section 5.6.

Example: Parameters for an HMM based on Figure 5.3

Let us assume that we use Laplace's rule to obtain parameters for an HMM corresponding to the alignment in Figure 5.3. Then $e_{M_1}(V) = 6/27, e_{M_1}(I) = e_{M_1}(F) = 2/27$, and $e_{M_1}(a) = 1/27$ for all residue types a other than V, I, F. Similarly, $a_{M_1 M_2} = 7/10, a_{M_1 D_2} = 2/10$ and $a_{M_1 I_1} = 1/10$ (following column 1 there are six transitions from match to match, one transition to a delete state, in HBB_HUMAN, and no insertions). Figure 5.4 shows the complete set of parameters for the HMM in diagrammatic form.   □

5.4   Searching with profile HMMs

One of the main purposes of developing profile HMMs is to use them to detect potential membership in a family by obtaining significant matches of a sequence to the profile HMM. We will assume for now that we are looking for global matches.

In practice, as for pairwise alignment, one of the local alignment methods may be more sensitive for finding distant matches. We discuss these in the next section.

We have a choice of ways to score a match to a hidden Markov model. We can either use the Viterbi equations to give the most probable alignment π^* of a sequence x together with its probability $P(x, \pi^* | M)$, or the forward equations to calculate the full probability of x summed over all possible paths $P(x|M)$.

In either case, for practical purposes the result we want to consider when evaluating potential matches is the log-odds ratio of the resulting probability to the probability of x given our standard random model

$$P(x|R) = \prod_i q_{x_i}.$$

We therefore show here versions of the Viterbi and forward algorithms that are designed specifically for profile HMMs, and which result directly in the desired log-odds values. Note that changing to log-odds does not change the result; we could have subtracted the random model log score afterwards. However, it is cleaner and more efficient. Another practical reason for working in log-odds units is to avoid problems of underflow when working with raw probabilities, as we discussed in Section 3.6.

Viterbi equations

Let $V_j^M(i)$ be the log-odds score of the best path matching subsequence $x_{1...i}$ to the submodel up to state j, ending with x_i being emitted by state M_j. Similarly $V_j^I(i)$ is the score of the best path ending in x_i being emitted by I_j, and $V_j^D(i)$ for the best path ending in state D_j. Then we can write

$$
\begin{aligned}
V_j^M(i) &= \log \frac{e_{M_j}(x_i)}{q_{x_i}} + \max \begin{cases} V_{j-1}^M(i-1) + \log a_{M_{j-1}M_j}, \\ V_{j-1}^I(i-1) + \log a_{I_{j-1}M_j}, \\ V_{j-1}^D(i-1) + \log a_{D_{j-1}M_j}; \end{cases} \\
V_j^I(i) &= \log \frac{e_{I_j}(x_i)}{q_{x_i}} + \max \begin{cases} V_j^M(i-1) + \log a_{M_j I_j}, \\ V_j^I(i-1) + \log a_{I_j I_j}, \\ V_j^D(i-1) + \log a_{D_j I_j}; \end{cases} \quad (5.1) \\
V_j^D(i) &= \max \begin{cases} V_{j-1}^M(i) + \log a_{M_{j-1}D_j}, \\ V_{j-1}^I(i) + \log a_{I_{j-1}D_j}, \\ V_{j-1}^D(i) + \log a_{D_{j-1}D_j}. \end{cases}
\end{aligned}
$$

These are the general equations. In a typical case, there is no emission score $e_{I_j}(x_i)$ in the equation for $V_j^I(i)$ because we assume that the emission distribution from the insert states I_j is the same as the background distribution, so the probabilities cancel in the log-odds form. Also, the $D \rightarrow I$ and $I \rightarrow D$ transition terms may not be present, as discussed above.

We need to take a little care over initialisation and termination of the dynamic programming. We want to allow the alignment to start and end in a delete or insert state, in case the beginning or end of the sequence does not match the first or the last match state of the model. The simplest way to ensure this mechanistically is to rename the Begin state as M_0 and set $V_0^M(0) = 0$ (as we did in Chapter 3). We then allow transitions to I_0 and D_1. Similarly, at the end we can collect together possible paths ending in insert and delete states by renaming the End state to M_{L+1} and using the top relation without the emission term to calculate $V_{L+1}^M(n)$ as the final score.

If these recurrence equations are compared with those for standard gapped dynamic programming in (2.16), it can be seen that apart from renaming of variables this is the same algorithm, but with the substitution, gap-open and gap-extend scores all depending on position in the model, j.

Forward algorithm

The recurrence equations for the forward algorithm are similar to the Viterbi equations, but with the max() operation replaced by addition. We define variables $F_j^M(i)$, $F_j^I(i)$ and $F_j^D(i)$ for the partial full log-odds ratios, corresponding to $V_j^M(i)$, $V_j^I(i)$ and $V_j^D(i)$. The recurrence equations are then:

$$F_j^M(i) = \log \frac{e_{M_j}(x_i)}{q_{x_i}} + \log \left[a_{M_{j-1}M_j} \exp \left(F_{j-1}^M(i-1) \right) \right.$$
$$\left. + a_{I_{j-1}M_j} \exp \left(F_{j-1}^I(i-1) \right) + a_{D_{j-1}M_j} \exp \left(F_{j-1}^D(i-1) \right) \right];$$

$$F_j^I(i) = \log \frac{e_{I_j}(x_i)}{q_{x_i}} + \log \left[a_{M_j I_j} \exp \left(F_j^M(i-1) \right) \right.$$
$$\left. + a_{I_j I_j} \exp \left(F_j^I(i-1) \right) + a_{D_j I_j} \exp \left(F_j^D(i-1) \right) \right];$$

$$F_j^D(i) = \log \left[a_{M_{j-1}D_j} \exp \left(F_{j-1}^M(i) \right) + a_{I_{j-1}D_j} \exp \left(F_{j-1}^I(i) \right) \right.$$
$$\left. + a_{D_{j-1}D_j} \exp \left(F_{j-1}^D(i) \right) \right].$$

Initialisation and termination conditions are handled as for the Viterbi case, with $F_0^M(0)$ being initialised to 0.

Although these appear a little complicated, in a practical implementation the operation $\log(e^x + e^y)$ can be performed efficiently to adequate accuracy by function lookup and interpolation; see Section 3.6.

Alternatives to log-odds scoring

In some of the earlier papers on HMMs, rather than calculating the log-odds score relative to a random model, the logarithm of the probability of the sequence given the model was used directly. This was called the LL score for 'log likelihood': $LL(x) = \log P(x|M)$. The LL score is strongly length dependent, so for searching

Figure 5.5 *To the left the length-normalized LL score is shown as a function of sequence length. The right plot shows the same for the log-odds score.*

it is not good enough to use a simple threshold. It is better to use LL divided by the sequence length, but even that is not always perfect, because the dependence between LL and sequence length is not linear (see example below).

A way to get around this is to estimate an average score and a standard deviation as a function of length and then use the number of standard deviations each sequence is away from the average. This is called the Z-score, and is also illustrated in the example below.

Example: Modelling and searching for globins

From 300 randomly picked globin sequences a profile HMM was estimated from scratch, i.e. starting from unaligned sequences using procedures we will explain in Chapter 6. A simple pseudocount regulariser was used. The estimation was done several times and the model with the highest overall LL score was picked. (We used the default settings of the SAM package, version 1.2; Hughey & Krogh [1996]).

With this model a database of about 60 000 proteins (SWISS-PROT release 34; Bairoch & Apweiler [1997]) was searched using the forward algorithm. The LL and log-odds scores were found for each sequence. For the null model we used the amino acid frequencies of the 300 sequences in the training set. In Figure 5.5 the length-normalised scores are shown for all the globins in the training set, all the other globins in the database and all the rest of the proteins with lengths up to 300 amino acids.[1] The globin sequences are clearly separated from the non-globins apart from a few in the 'twilight zone.'

The main difference between the two is in the variance of the score for non-globins, which is lower for the log-odds score, and therefore the separation is clearer. However, just choosing a cut-off of zero for the log-odds would miss a

---

[1] A few dubious globins and other strange sequences were removed from these data.
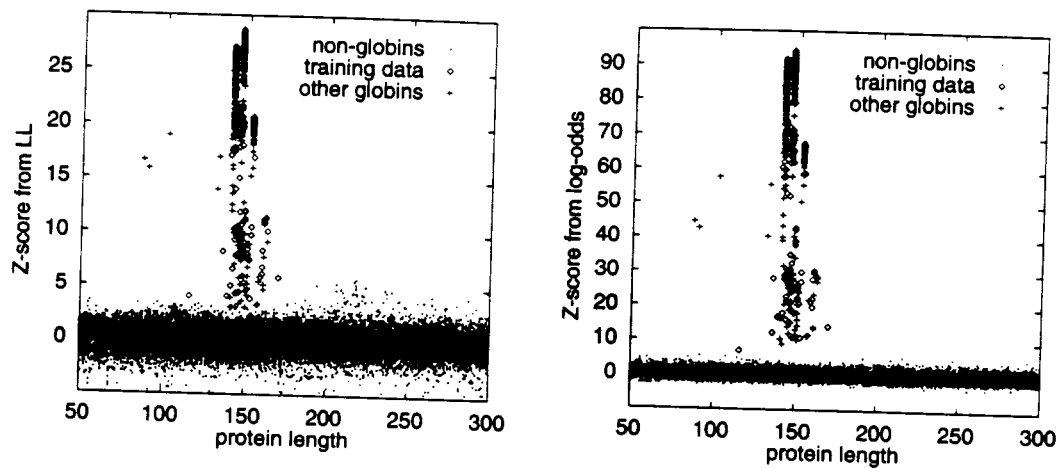
Figure 5.6 *The Z-score calculated from the LL scores (left) and the log-odds (right).*

lot of real globins in the search. This is because the profile HMM is not broad enough: it is too concentrated on a subset of the globins. Although there are ways to address this problem directly that we will return to later in the chapter, it is also possible to take a pragmatic approach to the separation of signal from noise given the results of the search, and calculate Z-scores for each hit.

To calculate Z-scores, a smooth curve is fitted to the LL or log-odds score of the non-globin sequences (a method is outlined in Krogh *et al.* [1994]). A standard deviation is then estimated for each length (or rather a little interval around it), and for each score the distance from the smooth curve is calculated in units of the standard deviation. This is the Z-score. The result (still as a function of sequence length) is shown in Figure 5.6.[2]

It is evident that it is now possible to find a threshold which will separate most globins from all other sequences. It is also clear that the score based on log-odds is much better for discrimination, with approximately three times the signal to noise ratio of the LL score. The reason for this is that dividing by the probability of the random model adjusts for the residue composition of the sequence. Without doing that, sequences with similar residue compositions as globins will tend to score more highly than sequences containing different residues, increasing the variance of the noise.

□

Alignment

Aside from finding matches, the other principal use of profile HMMs is to give an alignment of a sequence to the family, or more precisely to add it into the multiple alignment of the family. This is primarily the subject of the next chapter,

---

[2] There is no analytical result about the shape of these score distributions. The global alignment distribution is probably not exactly a Gaussian [Waterman 1995], but it appears to be a good approximation. For local alignments the extreme value distribution may be more reasonable, as discussed in Chapter 2.

on multiple alignment methods, which covers alignment with profile HMMs at length. For now, we will just point out that the natural solution is to take the highest scoring, or Viterbi, alignment. This is obtained by tracing back on the Viterbi variables $V_j^*(i)$, exactly as with pairwise alignment. Beyond this, all the methods of Chapter 4 can be applied, to explore variants, and to assess the reliability of the alignment.

5.5 Profile HMM variants for non-global alignments

We have seen that there is a very close relationship between the Viterbi alignment of a sequence to a profile HMM and the global dynamic programming comparison between two sequences using affine gap penalties, which we described in Chapter 2. It is therefore possible to generalise all the variations of dynamic programming, such as those that find local, repeat and overlap matches, to use profile HMMs.

However, we have developed probabilistic models much more fully since Chapter 2, and this time we want to take more care to ensure that the result of converting to a local algorithm remains a proper probabilistic model, i.e. that we assign each sequence a true probability so that the sum over all sequences $\sum_x P(x|M) = 1$. Our approach to doing this is to specify a new model for the complete sequence x, which incorporates the original profile HMM together with one or more copies of a simple self-looping model that is used to account for the regions of unaligned sequence. These behave very like the insert states that we added to the profile itself. We call them *flanking model* states, because they are used to model the flanking sequences to the actual profile match itself.

The model for local (Smith–Waterman style) alignment is shown here:

The flanking model states are shown as shaded diamonds. Notice that as well as specifying the emission probabilities of the new states, which will normally of course be q_a, we must specify a number of new transition probabilities. The looping probability on the flanking states should be close to 1, since they must

account for long stretches of sequence. Let us set these to $(1 - \eta)$. Note also that we have made use of silent states, shown as shaded circles, as 'switching points' to reduce the total number of transitions.

The next issue is how to set all the transition probabilities from the left flanking state to different start points in the model. One option is to give them equal probabilities, η/L. Another is to assign more probability to starting at the beginning of the model. The default option in the HMMER package for profile HMMs [Eddy 1996] assigns probability $\eta/2$ to the start of the profile, and $\eta/(2(L-1))$ to the other positions, favouring matches that start at the beginning of the model.

If all the probability is assigned to the first model state, then it forces this model to match only complete copies of the profile in the searched sequence, ensuring a type of 'overlap' match constraint. This can be appropriate when, for example, the HMM represents a protein domain that you expect to find either present as a whole or absent. However, to allow for rare cases where the first residue might be missing, it may be wise in such cases to allow a direct transition from the flanking state into a delete state, as shown here:

It is clear that by tinkering with the transition connections and probabilities a wide variety of different models can be produced, each potentially useful in different circumstances. A final example similar to the first model for local matches is

which allows repeat matches to subsections of the profile model, like the repeat algorithm variant in Chapter 2.

Note that all these variants of transition connectivity and probability assignment affect not only the types of match that are allowed, but also the score. More

restrictive transition distributions will give higher match scores if a good match is found, so are preferable if they can be designed to represent the types of correct matches that are expected.

Exercises

5.1 Show that if the random model is the same as that described in Chapter 4 (a succession of two states looping on themselves with probability $(1 - \eta)$), with η the same as in the flanking models, the local alignment model gives update equations like those of equation (2.9).

5.2 Explain the reasons for any differences.

5.6 More on estimation of probabilities

As promised above, we now return to the subject of parameter estimation at greater length. Although our discussion for most of this section will be focused on the emission probabilities, analogous methods can be used for the transition probabilities. The aim here is to introduce methods that can be used. A more detailed mathematical discussion about the estimation of probabilites from sample counts is given in Chapter 11 (p. 311).

The most straightforward approach to parameter estimation would be to give the maximum likelihood estimates for the parameters. We will change notation slightly from that used before. Given observed frequencies c_{ja} of residue a in position j of the alignment, maximum likelihood estimates for $e_{M_j}(a)$, the corresponding model parameters, are

$$e_{M_j}(a) = \frac{c_{ja}}{\sum_{a'} c_{ja'}}. \tag{5.2}$$

As we described above, a clear problem with this is that if there are no observed examples of a particular outcome then its probability is estimated as zero. This will frequently occur. For example, in the alignment of Figure 5.3 only V, I and F are present in the first column. However, it is quite likely that other amino acids will occur in that position amongst all the other globin sequences in biology. The easiest way to deal with this problem is to add pseudocounts to the observed counts c_{ja}. Below, we first discuss the pseudocount approach at greater length, then give some more complex alternatives.

Simple pseudocounts

A very simple and much-used pseudocount method is to add a constant to all the counts, which prevents the problem with zero probabilities. When the constant is one, as we used above in our example, this is called 'Laplace's rule'. A slightly

more sophisticated method is to add a quantity proportional to the background distribution, giving

$$e_{M_j}(a) = \frac{c_{ja} + Aq_a}{\sum_{a'} c_{ja'} + A},$$ (5.3)

where c_{ja} are the real counts, and A is the weight put on the pseudocounts as compared to the real counts. Values of A of around twenty seem to work well for protein alignments.

This form of regularisation has the appealing feature that $e_{M_j}(a)$ is approximately equal to q_a if very little data is available, i.e. all the real counts are very small compared to A. At the other extreme, where a large amount of data is available, the effect of the regulariser becomes insignificant and $e_{M_j}(a)$ is essentially equal to the maximum likelihood solution. So, at this intuitive level, pseudocounts make a lot of sense.

Adding pseudocounts amounts to adding some fake imagined data into the alignment, based on our general knowledge of proteins, to represent all the other things that might happen. They thus correspond to prior information about protein families, before having seen the specific data for the family in the form of the alignment. This statement can be formalised in a Bayesian framework. Bayes' equation tells us how to combine data, D, with a prior probability distribution over the parameters $P(\theta)$ to give a posterior distribution over θ, from which we can take either the maximum or the mean as our best estimate,

$$P(\theta|D) = \frac{P(D|\theta)P(\theta)}{P(D)}.$$

In our case the parameters θ are our model probabilities. The pseudocount method given above corresponds in this Bayesian framework to assuming a Dirichlet prior distribution with parameters $\alpha_a = Aq_a$ over the probabilities; see Chapter 11 for mathematical details.

Dirichlet mixtures

The problem with the simple pseudocounts, as compared to the substitution matrix based methods, is that only the most rudimentary prior knowledge can be contained in a single pseudocount vector. For this reason we need a lot of example data in the alignment to get good estimates of the parameters. Experience suggests that to achieve good discrimination typically fifty or more examples are desirable when modelling proteins.

In order to include better prior information, it was therefore suggested by Brown *et al.* [1993] that one should use a *mixture* of Dirichlet distributions as the prior. The idea is that there might be several different sets of pseudocount priors $\alpha_{\bullet}^1, \ldots, \alpha_{\bullet}^K$ corresponding to different types of alignment environments, where

α_a^k corresponds to Aq_a in the example above. One set might be relevant for exposed loop environments, one for buried small residue environments, etc. Given our counts c_{ja} we first estimate how likely each prior distribution k is (based on how well it fits the observed data), then combine their effects according to these posterior probabilities:

$$e_{M_j}(a) = \sum_k P(k|c_j) \frac{c_{ja} + \alpha_a^k}{\sum_{a'}(c_{ja'} + \alpha_{a'}^k)},$$

where the $P(k|c_i)$ are the *posterior mixture coefficients*. We calculate these by Bayes' rule,

$$P(k|c_j) = \frac{p_k P(c_j|k)}{\sum_{k'} p_{k'}' P(c_j|k')}$$

where the p_k are the prior probabilities of each mixture component, and $P(c_j|k)$ is the probability of the data according to Dirichlet mixture k. The equation for $P(c_j|k)$ has a frightening looking form, which is in fact fairly simple to calculate:

$$P(c_j|k) = \frac{(\sum_a c_{ja})!}{\prod_a c_{ja}!} \frac{\Gamma\left(\sum_a c_{ja} + \alpha_a^k\right)}{\prod_a \Gamma(c_{ja} + \alpha_a^k)} \frac{\Gamma\left(\sum_a \alpha_a^k\right)}{\prod_a \Gamma(\alpha_a^k)},$$

where $\Gamma(x)$ is the gamma function, a standard function over the reals related to the factorial function on the integers. For further details and an explanation of this equation, see Chapter 11, where we also describe how the mixture component distributions $\alpha_\bullet^k$ are obtained.

Using this type of approach, it seems that good profile HMMs can be fit to alignments with as few as ten or twenty examples [Sjölander *et al.* 1996].

Substitution matrix mixtures

An alternative approach to using a mixture of Dirichlets is to adjust the pseudocounts in a single Dirichlet formulation, using information from the observed counts and a substitution matrix. This is not a theoretically well-founded approach, but it makes intuitive sense as a heuristic, combining features of the non-probabilistic profile methods and the Dirichlet pseudocount methods.

The first step is to convert the matrix entries $s(a,b)$ into conditional probabilities $P(b|a)$. If we assume that the substitution matrix entries are derived as log-odds ratios, as in Chapter 2, then $s(a,b) = \log(P(a,b)/q_a q_b)$, which is the same as $\log(P(b|a)/P(b))$, so $P(b|a) = q_b e^{s(a,b)}$. We can in fact derive $P(b|a)$ values from an arbitrary score matrix $s(a,b)$ given background probabilities q_a; see below.

Given conditional probabilities $P(b|a)$ we can generate pseudocounts as follows. Let f_{ja} be the maximum likelihood probabilities derived from the counts,

so $f_{ja} = c_{ja}/\sum_{a'} c_{ja'}$. Using these we set pseudocount values with

$$\alpha_{ja} = A \sum_b f_{jb} P(a|b),$$

where A is a positive constant comparable to the one we used with simple pseudo-counts [Tatusov, Altschul & Koonin 1994; Claverie 1994; Henikoff & Henikoff 1996]. We then use essentially the same equation as (5.3) to obtain the model parameters:

$$e_{M_j}(a) = \frac{c_{ja} + \alpha_{ja}}{\sum_{a'} c_{ja'} + \alpha_{ja'}}.$$

There is no obvious statistical interpretation for this type of pseudocount, but the idea is quite natural: amino acid i contributes to pseudocount j in proportion to its abundance in the column and the probability of its changing to amino acid j. The formula interpolates between the treatment of pairwise alignments and the maximum likelihood solution. The substitution matrix term dominates if there are small numbers of sequences (especially if $A \gg 1$), and values close to the maximum likelihood estimate are obtained when the number of counts is large (more precisely when the total number of counts $C_j \gg A$).

There are various choices for the scaling constant A of the pseudocounts. For instance $A = 1$ was used in Lawrence et al. [1993], but this appears to be too weak in practice. Claverie [1994] suggests $A = \min(20, N)$, and Henikoff & Henikoff [1996] suggest $A = 5R$, where R is the number of different residue types observed in the column (i.e. the number of a for which $c_{ja} > 0$).

Deriving $P(b|a)$ from an arbitrary matrix

Even if a score matrix $s(a, b)$ was not derived as a log-odds matrix, as long as certain conditions are fulfilled it is possible to find a scale factor λ such that $\lambda s(a, b)$ will behave correctly when interpreted as a log-odds matrix [Altschul 1991]. The conditions are that the matrix is negatively biased, i.e. $\sum_{ab} q_a q_b s(a, b) < 0$, and that it contains at least one positive entry.

What we want is a set of values r_{ij} for which

$$s(a, b) = \frac{1}{\lambda} \log \frac{r_{ab}}{q_a q_b},$$

where r_{ab} can be interpreted as the probability for the pair a, b. This equation is easily inverted, so we get the pair probabilities expressed in terms of the substitution matrix $r_{ab} = q_a q_b \exp(\lambda s(a, b))$. To be legitimate probabilities the r_{ab} have to sum to one. We therefore need to find a λ such that

$$f(\lambda) = \sum_{a,b} q_a q_b e^{\lambda s(a,b)} - 1 = 0. \tag{5.4}$$

One such value is $\lambda = 0$, but clearly this is not what we want. The two conditions

we gave above turn out to be sufficient to ensure there is another, positive solution to this equation; see the exercises below.

The resulting value of λ is called the natural scaling factor of the substitution matrix. This probabilistic interpretation of the substitution matrix leads to an entropy measure for the matrix of $\sum_{ab} r_{ab} \log(r_{ab}/q_a q_b)$, which is a useful quantity for characterising and comparing substitution matrices [Altschul 1991].

Exercises

5.3   Use the negative bias condition to show that $f(\lambda)$ is negative for small enough λ. Hint: calculate $f'(0)$, the derivative of $f(\lambda)$ at $\lambda = 0$.

5.4   Use the second condition, that there is at least one positive $s(a, b)$, to show that $f(\lambda)$ becomes positive for large enough λ.

5.5   Finally, show that the second derivative of $f(\lambda)$ is positive, and from this and the results of the previous two exercises that there is one and only one positive value of λ satisfying (5.4).

Estimation based on an ancestor

There is a more principled and direct way to use the information in substitution matrices for estimating the HMM probabilities than that described above. This approach does not use pseudocounts. Instead, it assumes that all the observed sequences have been derived independently from a common ancestor, and generates an estimate of the residue present in a given position in that common ancestor (or rather a posterior probability distribution for what that residue was). From this we can estimate the probability of seeing each residue in a new descendant of the ancestor, different from those in the sample.

Assume we have example sequences x^k with residues x_j^k in column j of the alignment (we have adjusted our notation slightly; this x_j^k is not the jth residue in sequence x^k if there are gaps, but it is a convenient notation for what we need here). Once again, we need the conditional probabilities $P(b|a)$ derived from the substitution matrix. Let the residue in the common ancestor be y_j. Then we can use Bayes rule to calculate the posterior probability that $y_j = a$

$$P(y_j = a | \text{alignment}) = \frac{q_a \prod_k P(x_j^k|a)}{\sum_{a'} q_{a'} \prod_k P(x_j^k|a')}. \tag{5.5}$$

Note that we needed a prior distribution for residues at the common ancestor, which we set to q_a because that is our background probability for amino acids in the absence of further information.

We can now calculate the HMM emission probabilities as the predicted probabilities for a new sequence

$$e_{M_j}(a) = \sum_{a'} P(a|a')P(y_j = a'|\text{alignment}). \tag{5.6}$$

One problem with this approach is that, as we noticed above, different columns vary widely in their degree of conservation. Indeed, that is one of the properties that we wanted to exploit when using alignments to estimate profile HMMs. However, using a single substitution matrix implies assuming a fixed degree of conservation. As we discussed in Chapter 2, matrices typically come in families varying in their level of implied conservation. Examples are the PAM [Dayhoff, Schwartz & Orcutt 1978] and the BLOSUM [Henikoff & Henikoff 1992] series of matrices. We can therefore significantly improve the approach in (5.5) and (5.6) if we optimise over choice of matrix from a family. This way, a very conserved column might use a conservative matrix, such as PAM30, and a very varied column would use a divergent matrix, such as PAM500.

How do we choose the optimal matrix? A natural approach is to maximise the likelihood of the observed data

$$P(x_j^1,\ldots,x_j^N|t) = \sum_a q_a \prod_k P(x_j^k|a,t) \qquad (5.7)$$

where t is the matrix family parameter (t for evolutionary *time*). It would also be possible to use a Bayesian approach here, proposing a prior distribution over t, then combining this with (5.7) in Bayes' rule to obtain a posterior distribution for t, and summing over this in (5.6). However, that would require signficantly more computation.

The maximum likelihood time-dependent approach is closely related to the 'evolutionary weights' method in the PROFILE package [Gribskov & Veretnik 1996]. However, that method estimates different evolutionary times t for each possible ancestral amino acid, and also adjusts the resulting weights with respect to a set of baseline probabilities; for details see Gribskov & Veretnik [1996]. There are also strong connections between the methods of this subsection and those discussed later in Chapter 8 when building phylogenetic trees using maximum likelihood methods.

Testing the pseudocount methods

All the methods mentioned above have been tested in various ways. Direct tests, in which profiles were constructed and used for searching, were carried out extensively by Henikoff & Henikoff [1996]. The best method turned out to be the substitution matrix based method (5.6), with $A = 5R$ as described above; the Dirichlet mixture regulariser came a reasonably close second. Other tests gave different results [Tatusov, Altschul & Koonin 1994; Karplus 1995], so it is not clear which method is best, and it is likely that this will depend on the application and the details of the mixture components or substitution matrix used.

An interesting method was for testing various regularisers was given by Karplus [1995]. Instead of performing a huge number of database searches, he

asked the following question:[3] How well can an amino acid distribution be approximated from a small sample? Columns were extracted from a large set of deep alignments (the BLOCKS database; Henikoff & Henikoff [1991]). Imagine we take a small sample of size n with counts s_a from a column with complete counts C_a. From the sample counts s_a we can estimate the probabilities $e_s(a)$ of other symbols that might occur in the same column, using one of the methods described above (we use a subscript s to remind ourselves that this estimation is dependent on the sample counts). We can also estimate the probabilities of other symbols directly from the frequencies with which they occur in all columns of the database together with the probability $P(s|C)$ of drawing s from a column C (given by the multinomial distribution). This estimate is given by:

$$P(a|s) = \frac{\sum_{\text{columns } C} P(s|C)C_a}{\sum_{\text{columns } C} P(s|C)|C|},$$

where $|C|$ denotes the number of symbols in the column C. $P(a|s)$ can only be calculated up to a sample size of $n = 5$, but this is also the most interesting regime, because it is for small sample sizes that regularisation is most crucial. We can now use the relative entropy $-\sum_a P(a|s)\log e_s(a)$ to compare the 'ideal' probability $P(a|s)$ with that given by the regulariser. Summing over all samples s of size n gives a measure

$$E_n = \sum_{s,|s|=n} P(s)\left(-\sum_a P(a|s)\log e_s(a)\right), \tag{5.8}$$

where $P(s)$ is the probability of drawing the sample s averaged over all columns in the database. This can be calculated using $P(s) = \sum_C P(s|C)|C|/\sum_C |C|$.

Karplus proposed that a good regulariser should minimise E_n. He showed that several of the more complex regularisers described above resulted in estimators that were very close to optimal, in the sense that E_n was very small up to $n = 5$. Of course, we are ultimately interested in database searches, and it is not evident that the regulariser obtaining the lowest value of E_n will actually be best for searching. It is likely that the typical similarities in the source alignment database are not the same as the ones that we will be searching for with our HMM.

As well as evaluating methods, Karplus' approach can also be used to set the free parameters in the various methods described above, for example the total number of pseudocounts A to use in (5.3). For any value of A we can calculate E_n from our database of columns, either directly or by some sort of random sampling, and in fact we can also calculate the gradient of the relative entropy with respect to A. We can therefore find the value of A that minimises this average relative entropy, using gradient descent methods [Press *et al.* 1992], or by

---

[3] This page has been rewritten for the second printing.

other optimisation methods. In principle this can be done for any sample size n, yielding parameters dependent on n.

5.7   Optimal model construction

When we first discussed the parameterisation of profile HMMs, we pointed out that as well as estimating the probability parameters, it is necessary to decide which columns of the alignment should be assigned to insert states, and which to match states. We call this process model construction. At the time we proposed a simple heuristic, but we can do better than that. There is an efficient dynamic programming algorithm which can find the column assignments that maximise the posterior probability of the mode, at the same time as fitting optimal probability parameters.

In the profile HMM formalism, it is assumed that an aligned column of symbols corresponds either to emissions from the same match state or to emissions from the same insert state. It therefore suffices to mark which columns come from match states to specify a profile HMM architecture and the state paths for all the sequences in the alignment, as shown in Figure 5.7. In a marked column, symbols are assigned to match states and gaps are assigned to delete states. In an unmarked column, symbols are assigned to insert states and gaps are ignored. State transition and symbol emission counts are obtained from the state paths, and these counts can be used to estimate probability parameters by one of the methods in the previous section. In passing, we note that this model estimation procedure implicitly assumes that the multiple alignment is correct, i.e. that the implied state paths have probability one and all other state paths have probability zero, which is akin to a Viterbi assumption. The next chapter addresses issues of simultaneous alignment and model estimation.

There are 2^L combinations of markings for an alignment of L columns, and hence 2^L different profile HMMs to choose from. There are at least three ways to determine the marking. In *manual* construction, the user marks alignment columns by hand. This is perhaps the simplest way to allow users to manually specify the model architecture to use for a given alignment. In *heuristic* construction, a rule is used to decide whether a column should be marked. For instance, a column might be marked when the proportion of gap symbols in it is below a certain threshold. In *MAP* construction, a maximum *a posteriori* choice is determined by dynamic programming. A description of this algorithm follows.

MAP match–insert assignment

The MAP construction algorithm recursively calculates a number S_j, which is the log probability of the optimal model for the alignment up to and including column