

## Chapter 1

# Sequence Alignment

The goal of pairwise sequence alignment is to establish a correspondence between the elements in a pair of sequences that share a common property, such as common ancestry or a common structural or functional role. In computational biology, the sequences under consideration are typically nucleic acid or amino acid polymers. We will consider three variants of the pairwise sequence alignment problem: global alignment, semi-global alignment, and local alignment.

Global alignment is used in cases where we have reason to believe that the sequences are related along their entire length. If, for example, sequences  $s_1$  and  $s_2$  are two independent sequencing runs of the same PCR product, then they should differ only at those positions where there are sequencing errors. In order to find those sequencing errors, we align all of sequence  $s_1$  with all of sequence  $s_2$ . Other applications of global alignment include finding mutations in closely related gene or protein sequences and identification of single nucleotide polymorphisms (SNPs).

Semi-global alignment is a variant of global alignment that allows for gaps at the beginning and/or the end of one of the sequences. Semi-global alignment is used in situations where we believe that  $s_1$  and  $s_2$  are related along the entire length of the region where they overlap. For example, if  $s_1$  is the open reading frame of a eukaryotic gene with a single exon and  $s_2$  is the mRNA transcript produced when  $s_1$  is expressed, every base in  $s_1$  should correspond to some base in  $s_2$ . Semi-global alignment “jumps” over the 5' untranslated region in  $s_2$  without exacting a penalty, but forces an alignment along the entire length of  $s_1$ .

In contrast, local alignment addresses cases where we only expect to find isolated regions of similarity. One example is alignment of genomic DNA upstream from two co-expressed genes to find conserved regions that may correspond to transcription factor binding sites. Another application is identification of conserved domains<sup>1</sup> in two amino acid sequences

---

<sup>1</sup>A domain is a peptide sequence that encodes a protein module that will fold into its characteristic

that encode proteins that share one or more domains, but are otherwise unrelated.

Prior to introducing algorithms for these pairwise alignment problems, we introduce some notation in Box 1.

## 1.1 Global pairwise alignment

Now that we have some notation to work with, we will introduce a formal definition of a global alignment. Next we will consider how to assign a numerical score to an alignment. The score represents an assessment of the quality of the alignment. Finally, we will introduce an efficient algorithm to find the alignment that is optimal with respect to a scoring function.

Let  $\Sigma' = \Sigma \cup \{-\}$  be the alphabet expanded to include a character to represent gaps. Given sequence  $s_1 \in \Sigma^*$  of length  $n_1$  and sequence  $s_2 \in \Sigma^*$  of length  $n_2$ ,  $\alpha^\kappa(s_1, s_2) = \{s_1^\kappa, s_2^\kappa\}$  is a global alignment of  $s_1$  and  $s_2$  if and only if

- $s_1^\kappa, s_2^\kappa \in (\Sigma')^*$ ,
- $|s_1^\kappa| = |s_2^\kappa| = l^\kappa$ , where  $\max(n_1, n_2) \leq l^\kappa \leq n_1 + n_2$ ,
- $s_1$  is the subsequence obtained by removing '-' from  $s_1^\kappa$  and  $s_2$  is the subsequence obtained by removing '-' from  $s_2^\kappa$ ,
- there is no value of  $i$  for which  $s_1^\kappa[i] = s_2^\kappa[i] = '-'$ .

There are many alignments of  $s_1$  and  $s_2$ . The superscript  $\kappa$  is an index to designate a specific alignment. In situations where only one alignment is under consideration or there is no ambiguity, we use the simpler notation  $\alpha(s_1, s_2) = \{s_1', s_2'\}$ , where the length of  $\alpha(s_1, s_2)$  is simply denoted  $l$ .

Our goal is to find the global alignment that best captures the relationship between  $s_1$  and  $s_2$ . Which alignment best reflects the relationship between  $s_1$  and  $s_2$  is fundamentally a biological question. From a practical perspective, we use a mathematical approach: We introduce an objective criterion that provides a measure of the quality of an alignment and then seek the alignment,  $\alpha^*(s_1, s_2)$ , that optimizes that criterion. There may be more than one.

### 1.1.1 Scoring an alignment

Given sequences  $s_1$  and  $s_2$  and an alignment  $\alpha^\kappa(s_1, s_2) = \{s_1^\kappa, s_2^\kappa\}$ , it is convenient to assign a score to alignment  $\alpha^\kappa$  that quantifies how well  $\alpha^\kappa$  captures the relationship between  $s_1$  and  $s_2$ . This score may be a minimization or a maximization criterion.

---

shape independent of the surrounding amino acid context and that is found in many different proteins.

**Box 1: Notation for pairwise alignment****Alphabet:**

An *alphabet*, denoted by  $\Sigma$ , is a finite, unordered set of symbols; e.g.,

DNA:  $\Sigma_D = \{A, C, G, T\}$

RNA:  $\Sigma_R = \{A, C, G, U\}$

Amino acids:  $\Sigma_{AA} = \{A, C, D, E, F, G, H, I, K, L, M, N, P, Q, R, S, T, V, W, Y\}$

**Sequences or Strings:**

A *sequence* or *string*,  $s$ , is a finite succession of the symbols in  $\Sigma$ .

$\Sigma^*$  denotes the set of all sequences over alphabet  $\Sigma$ , including the empty sequence,  $\emptyset$ . For example,  $\Sigma_R^* = \{\emptyset, A, C, G, U, AA, AC, AG, AU, CA, CC, CG, CU, \dots\}$ .

Given a sequence  $s$  of length  $n$ , we use  $s[1]s[2] \dots s[n]$  to denote the symbols in  $s$ .

**Subsequences:**

A *subsequence* of  $s$  is any sequence obtained by removing zero or more symbols from  $s$ . The sequences CATA and CTG are subsequences of CATTAG. AATTCG is not.

A *proper subsequence* is a subsequence obtained by removing one or more symbols from  $s$ .

**Substrings:**

A *substring* of  $s$  is a subsequence of  $s$  consisting of consecutive symbols in  $s$ . Given a sequence,  $s$ , of length  $n$ , the substring that begins with  $s[i]$  and ends with  $s[j]$  is denoted  $s[i \dots j]$ ,  $1 \leq i \leq j \leq n$ . The sequence CAT is a substring of CATTAG. CATA is not.

A *prefix* of  $s$  is denoted  $s[1 \dots j]$ ,  $j \leq n$ .

A *suffix* of  $s$  is denoted  $s[i \dots n]$ ,  $1 \leq i$ .

Distance scoring: An alignment can be scored using a distance-based metric. This is a minimization criterion: a lower distance indicates a better alignment. We define the distance score of an alignment  $\alpha^\kappa(s_1, s_2) = \{s_1^\kappa, s_2^\kappa\}$  to be

$$\begin{aligned} D(\alpha^\kappa(s_1, s_2)) &= D(s_1^\kappa, s_2^\kappa) \\ &= \sum_{i=1}^{l^\kappa} d(s_1^\kappa[i], s_2^\kappa[i]), \end{aligned} \tag{1.1}$$

where  $d(x, y)$  is the distance between a pair of symbols  $x$  and  $y$  in  $\Sigma'$  and  $l^\kappa$  is the length of the alignment. The optimal alignment, denoted  $\alpha^*$ , is the alignment that minimizes the distance between  $s$  and  $t$ :

$$\alpha^*(s_1, s_2) = \underset{\kappa}{\operatorname{argmin}} D(\alpha^\kappa(s_1, s_2)).$$

The function specifying the distance between pairs of symbols must satisfy the following properties, for all  $x, y$ , and  $z$  in  $\Sigma'$ :

1.  $d(x, x) = 0$
2.  $d(x, y) > 0$
3.  $d(x, y) = d(y, x)$
4.  $d(x, y) \leq d(x, z) + d(z, y)$  (triangle inequality)

Several properties of the distance scoring function are worth noting. First,  $D(s_1^\kappa, s_2^\kappa)$  is a metric; that is, it satisfies the triangle inequality. This means that the penalty for replacing  $x$  with  $y$  is never improved by first replacing  $x$  with  $z$  and then replacing  $z$  with  $y$ . When  $z = \text{'.'}$ , this says that deleting  $x$  and then inserting  $y$  is never an improvement on a direct substitution of  $x$  with  $y$ . One consequence of the triangle inequality is that the cost of a substitution can never be greater than twice the cost of an indel. Intuitively, this makes sense: Alignments obtained by minimizing a function where one substitution costs more than two indels would contain no substitutions. With such a function, an alignment in which  $x$  is aligned with  $y$  could always be replaced by a lower cost alignment in which  $x$  and  $y$  are both aligned with gaps.

Second, the symmetric property,  $d(x, y) = d(y, x)$ , implies that there is no directionality in the scoring system. This is because, given a column with symbols  $x$  and  $y$  in a pairwise alignment, we have no way of knowing whether the ancestral symbol was an  $x$  that was later replaced by a  $y$ , or vice versa. It is also possible that the ancestor was neither  $x$  nor  $y$  and some combination of substitutions gave rise to  $x$  in one sequence and to  $y$  in the other. Similarly, when a symbol,  $x$ , in sequence  $s_1$  is aligned with a gap in sequence  $s_2$  (or vice versa), there is no way to know whether  $x$  was inserted in  $s_1$  or deleted from  $s_2$ . For this reason, gaps are also called “indels.”

If  $d(x, y) = 1$  and  $d(x, -) = 1, \forall x, y$ , then  $D(\alpha^*(s_1, s_2))$  corresponds to the minimum number of operations required to transform  $s_1$  into  $s_2$ , where the operations are substitution, insertion, and deletion. This is called the *edit distance*. If  $d(x, y) > 1$  or  $d(x, -) > 1$  or both, then  $D(\alpha^*(s_1, s_2))$  is called the *weighted edit distance*.

Similarity scoring: Alignments can also be scored with similarity measures. These are maximization criteria: a higher score indicates a better alignment. The similarity score of  $\alpha^\kappa(s_1, s_2) = \{s_1^\kappa, s_2^\kappa\}$  is

$$S(\alpha^\kappa(s_1, s_2)) = \sum_{i=1}^{l^\kappa} p(s_1^\kappa[i], s_2^\kappa[i]), \quad (1.2)$$

where  $p(x, y)$  is a score that reflects the similarity of  $x$  and  $y$  and  $p(x, -)$  is the gap score. The optimal alignment is the alignment that maximizes the similarity between  $s_1$  and  $s_2$ :

$$\alpha^*(s_1, s_2) = \underset{\kappa}{\operatorname{argmax}} S(\alpha^\kappa(s_1, s_2)).$$

In general, amino acid alignments are scored with substitution matrices that assign a different similarity score to each pair of amino acid residues. Typically, pairs of amino acids with similar properties have higher scores than pairs with divergent properties. Examples of substitution matrices used to score alignments include the PAM and BLOSUM matrices. We will discuss how such substitution matrices are derived later in the semester. For now, we consider a simple similarity scoring function that treats all symbols in  $\Sigma$  equally. This simple scoring function has just three values; a score for matching symbols ( $M$ ), a score for a mismatch ( $m$ ), and a gap score ( $g$ ):

$$\begin{aligned} p(x, x) &= M, \\ p(x, y) &= m, \\ p(x, -) &= g. \end{aligned} \quad (1.3)$$

In order to obtain alignments that make sense with similarity scoring, several constraints are imposed on the values of  $M$ ,  $m$ , and  $g$ . First, we require that  $M > m$ , because matches are preferred over mismatches. We further require that a substitution be preferred over two gaps (i.e.,  $m > 2g$ ). A scoring function with  $m < 2g$  would exclude the possibility of an alignment with substitutions, because the score of any substitution could be improved by replacing it with two gaps.

Note that for both distance and similarity scoring, the score of an alignment is defined to be the sum of the scores for the individual positions in the alignment (Equations 1.1 and 1.2), which implies that each position in the alignment is independent of neighboring positions. This assumption is unrealistic: In real biomolecular sequences, there can be interactions between neighboring, or even distant, residues in the sequence. However, scoring functions that assume positional independence are widely used because they greatly simplify the calculation of alignment scores and other mathematical analyses.

### 1.1.2 A dynamic programming algorithm to align a pair of sequences

We now have a formal definition of an alignment and a way of assigning a numerical score to any given alignment. How do we find the alignment with the optimal score? We could generate all possible alignments, score each one, and choose the alignment with the best score. However, the computational cost would be prohibitive, since the size of the space of all possible alignments of  $s_1$  and  $s_2$  is  $O(2^{n_1+n_2})$ . (Convince yourself this is the case.)

Dynamic programming can be used to find the optimal alignment efficiently. This strategy takes advantage of the fact that every prefix of an optimal pairwise alignment is the optimal alignment of a prefix of  $s_1$  and a prefix of  $s_2$ . This means that the optimal alignment of pairs of progressively longer prefixes of  $s_1$  and  $s_2$  can be obtained by extending the optimal alignment of shorter prefixes of  $s_1$  and  $s_2$ . It is not necessary to examine a suboptimal alignment of prefixes in order to find the optimal alignment of the full length strings.

The dynamic programs for all three sequence alignment problems compute a matrix  $\mathcal{A}$ , where  $\mathcal{A}[i, j]$  is the score of the optimal alignment of the prefixes  $s_1[1..i]$  and  $s_2[1..j]$ , that is, the prefixes of  $s_1$  and  $s_2$  that end at positions  $i$  and  $j$ , respectively.

Dynamic programming algorithms for sequence alignment have four components:

- Initialization of the first row and column of  $\mathcal{A}$ .
- A recurrence relation that specifies how to calculate the value of  $\mathcal{A}[i, j]$ ,  $i > 0, j > 0$ , from the values of neighboring cells.
- Determination of the score of the optimal alignment from the entries in matrix  $\mathcal{A}$ .
- A procedure to trace back through the matrix to obtain the optimal alignment. This is assisted by a traceback matrix,  $\mathcal{T}$ , which stores information on how to construct the optimal alignment from matrix  $\mathcal{A}$ .

The details of each of these steps are what differentiate global, semi-global, and local alignment. In all cases, the dynamic program proceeds from the upper left to the lower right corner of  $\mathcal{A}$ , calculating the cost of progressively longer optimal alignments of prefixes. The details of dynamic programming for global alignment are given below for both distance and similarity scoring functions.

#### Global alignment with distance scoring:

*Input:*

Sequences  $s_1$  and  $s_2$  of lengths  $n_1$  and  $n_2$ , respectively.

*Initialization:*

$$\begin{aligned}\mathcal{A}[0, j] &= \mathcal{A}[0, j - 1] + d(-, s_2[j]) && \text{(top row)} \\ \mathcal{A}[i, 0] &= \mathcal{A}[i - 1, 0] + d(s_1[i], -) && \text{(left column)}\end{aligned}$$

*Recurrence relation:*

$$\mathcal{A}[i, j] = \min \begin{cases} \mathcal{A}[i - 1, j] + d(s_1[i], -) \\ \mathcal{A}[i - 1, j - 1] + d(s_1[i], s_2[j]) \\ \mathcal{A}[i, j - 1] + d(-, s_2[j]) \end{cases} \quad (1.4)$$

Store the indices of the entry (or entries) in  $\mathcal{A}$  that minimize the right hand side of Equation 1.4 in the  $n_1 \times n_2$  matrix,  $\mathcal{T}$ , which we call the traceback matrix.

$$\mathcal{T}[i, j] = \operatorname{argmin} \begin{cases} \mathcal{A}[i - 1, j] + d(s_1[i], -) \\ \mathcal{A}[i - 1, j - 1] + d(s_1[i], s_2[j]) \\ \mathcal{A}[i, j - 1] + d(-, s_2[j]) \end{cases} \quad (1.5)$$

*Optimal alignment score:*

The score of the optimal alignment is  $\mathcal{A}[n_1, n_2]$ .

*Trace back:*

Follow the pointers from  $\mathcal{T}[n_1, n_2]$  to  $\mathcal{T}[0, 0]$  to obtain the optimal alignment.

If the entry at  $\mathcal{T}[i, j]$  is

$$\begin{array}{ll} [i - 1, j] & \text{then align } s_1[i] \text{ with a gap '-' } \\ [i - 1, j - 1] & \text{then align } s_1[i] \text{ with } s_2[j] \text{ '-' } \\ [i, j - 1] & \text{then align } s_2[j] \text{ with a gap '-' } \end{array}$$

*Output:*

The optimal alignment score,  $\mathcal{A}[n_1, n_2]$ .

The optimal global alignment of  $s_1$  and  $s_2$  with respect to distance function,  $D$ .

The dynamic programming algorithm for global alignment with similarity scoring has the same general structure as the global alignment algorithm for distance scoring. However, the details of the initialization and recurrence differ.

Global alignment with similarity scoring:

*Input:*

Sequences  $s_1$  and  $s_2$  of lengths  $n_1$  and  $n_2$ , respectively.

*Initialization:*

$$\begin{aligned} \mathcal{A}[0, j] &= \mathcal{A}[0, j - 1] + g && \text{(top row)} \\ \mathcal{A}[i, 0] &= \mathcal{A}[i - 1, 0] + g && \text{(left column)} \end{aligned}$$

*Recurrence relation:*

$$\mathcal{A}[i, j] = \max \begin{cases} \mathcal{A}[i - 1, j] + g \\ \mathcal{A}[i - 1, j - 1] + p(s_1[i], s_2[j]) \\ \mathcal{A}[i, j - 1] + g \end{cases} \quad (1.6)$$

Store the indices of the entry (or entries) in  $\mathcal{A}$  that maximize the right hand side of Equation 1.6 in the traceback matrix,  $\mathcal{T}$ .

$$\mathcal{T}[i, j] = \operatorname{argmax} \begin{cases} \mathcal{A}[i - 1, j] + g \\ \mathcal{A}[i - 1, j - 1] + p(s_1[i], s_2[j]) \\ \mathcal{A}[i, j - 1] + g \end{cases} \quad (1.7)$$

*Optimal alignment score:*

The score of the optimal alignment is  $\mathcal{A}[n_1, n_2]$ .

*Trace back:*

Follow the pointers from  $\mathcal{T}[n_1, n_2]$  to  $\mathcal{T}[0, 0]$  to obtain the optimal alignment.

*Output:*

The optimal alignment score,  $\mathcal{A}[n_1, n_2]$ .

The optimal global alignment of  $s_1$  and  $s_2$  with respect to similarity function,  $S$  (Equation 1.2).

At each step, the algorithm computes the value of  $\mathcal{A}[i, j]$  from the values of  $\mathcal{A}[i - 1, j]$ ,  $\mathcal{A}[i, j - 1]$ , and  $\mathcal{A}[i - 1, j - 1]$ . With distance scoring, all entries in  $\mathcal{A}$  are non-negative, since  $d(x, y) \geq 0, \forall x, y$ . With a similarity scoring function, the entries in  $\mathcal{A}$  may be positive or negative. The indices of the entry in  $\mathcal{A}$  that optimize the right hand side of the recurrence (Equations 1.4 and 1.6) are stored in a matrix,  $\mathcal{T}$ . These pointers are used to reconstruct the alignment that gave the optimal score. The algorithm continues until all entries in

the matrix  $\mathcal{A}$  have been assigned values. This algorithm computes the scores of all pairs of prefixes in  $O(n_1 \cdot n_2)$  time. The trace back through the alignment matrix to obtain the optimal alignment requires  $O(n_1 + n_2)$  time. Note that there may be more than one optimal alignment.

## 1.2 Semi-global alignment

Global alignment seeks the best, full length alignment of a pair of sequences; that is, the best way to match up two sequences along their entire length. For some applications, it is desirable to relax this requirement and not penalize gaps at the beginning and/or end of an alignment. For example, for sequence assembly, we seek sequence fragments that overlap; that is, we expect to be able to align the end of one fragment with the beginning of another. Very occasionally, we may find sequence fragments that start and end at the same position, but, in general, we expect some gaps at the beginning and at the end of the alignment. Another example is aligning cDNA's with genomic DNA to identify gene structure. Because the cDNA corresponds to a small region in the genome, the cDNA fragment will be flanked by gaps at both ends when aligned with the genomic DNA.

Semi-global alignment is a modification of global alignment that allows the user to specify that gaps will be penalty-free at the beginning of one of the sequences and/or at the end of one of the sequences. Given sequences  $s_1$  and  $s_2$ , there are eight possible cases to consider:

1. Gaps are penalty-free at the beginning of  $s_1$ ; *e.g.*,

```
s_1:  _ _ D O
s_2:  R E D O
```

2. Gaps are penalty-free at the beginning of  $s_2$ ; *e.g.*,

```
s_1:  R E D O
s_2:  _ _ D O
```

3. Gaps are penalty-free at the end of  $s_1$ ; *e.g.*,

```
s_1:  D O _ _
s_2:  D O N E
```

4. Gaps are penalty-free at the end of  $s_2$ ; *e.g.*,

```
s_1:  D O N E
s_2:  D O _ _
```

5. Gaps are penalty-free at the beginning and end of  $s_1$ ; *e.g.*,

```
s_1:  _ _ D O _ _
s_2:  R E D O N E
```

6. Gaps are penalty-free at the beginning and end of  $s_2$ ; *e.g.*,

```
s_1:  R E D O N E
s_2:  _ _ D O _ _
```

7. Gaps are penalty-free at the beginning of  $s_1$  and at the end of  $s_2$ ; *e.g.*,

```
s_1:  _ _ D O N E
s_2:  R E D O _ _
```

8. Gaps are penalty-free at the beginning of  $s_2$  and at the end of  $s_1$ ; *e.g.*,

```
s_1:  R E D O _ _
s_2:  _ _ D O N E
```

In semi-global alignment, we do not allow penalty-free gaps at the beginning of  $s_1$  and the beginning of  $s_2$  in the same alignment. Nor do we allow penalty-free gaps at the end of  $s_1$  and the end of  $s_2$ . Why not?

Like global alignment, the optimal semi-global alignment can be found using dynamic programming with either distance or similarity scoring. Below, we describe the modifications that are required to adapt the dynamic program for global pairwise alignment to the semi-global alignment problem. These modifications are described in terms of alignment with similarity scoring. Similar modifications can be made to obtain a semi-global alignment algorithm that uses distances.

Semi-global alignment with similarity scoring:

*Initialization:*

To allow penalty-free gaps at the beginning of  $s_1$  (Case 1), set the first row to zero and initialize the first column as in global alignment.

$$\begin{aligned} \mathcal{A}[0, j] &= 0, \forall_j && \text{(top row)} \\ \mathcal{A}[i, 0] &= \mathcal{A}[i - 1, 0] + g && \text{(left column)} \end{aligned}$$

To allow penalty-free gaps at the beginning of  $s_2$  (Case 2), set the first column to zero and initialize the first row as in global alignment.

$$\begin{aligned} \mathcal{A}[0, j] &= \mathcal{A}[0, j - 1] + g && \text{(top row)} \\ \mathcal{A}[i, 0] &= 0, \forall_i && \text{(left column)} \end{aligned}$$

*Recurrence relation:*

Same as global.

*Optimal alignment score and trace back:*

To avoid trailing gap penalties at the end of  $s_1$  (Case 3), we define the optimal score to be the optimal score in the bottom row

$$\max_j \mathcal{A}[n_1, j].$$

Trace back from  $\mathcal{T}[n_1, j^*]$ , where  $j^* = \operatorname{argmax}_j \mathcal{A}[n_1, j]$ . In other words, trace back from the cell(s) in the last row with optimal score.

To avoid trailing gap penalties at the end of  $s_2$  (Case 4), we define the optimal score to be the optimal score in the last column

$$\max_i \mathcal{A}[i, n_2].$$

Trace back from  $\mathcal{T}[i^*, n_2]$ , where  $i^* = \operatorname{argmax}_i \mathcal{A}[i, n_2]$ . In other words, trace back from the cell(s) in the last column with optimal score.

Note that when the first row (or column) of the matrix is initialized to zero, the traceback will end in the first row (or column), but not necessarily in the cell  $\mathcal{A}[0, 0]$ .

Like global alignment, either distance or similarity scoring can be used for semiglobal alignment. There may be more than one optimal semiglobal alignment.

### 1.3 Local pairwise alignment

Global and semiglobal alignment are used in cases where we expect that  $s_1$  and  $s_2$  are related from end to end. Semi-global allows for some gaps at the beginning and/or end of one sequence, but the underlying assumption is the same:  $s_1$  and  $s_2$  share a relationship within the entire aligned region. In contrast, local alignment is used in cases where  $s_1$  and  $s_2$  share one or more local regions that are related, but are not related from end to end.

The alignment of any substring  $s_1[h \dots i]$  of  $s_1$  and any substring  $s_2[j \dots k]$  of  $s_2$  is a local alignment of  $s_1$  and  $s_2$ . The optimal alignment of  $s_1[h \dots i]$  and  $s_2[j \dots k]$  is the highest scoring *global* alignment of those substrings, where  $1 \leq h \leq i \leq n_1$ , and  $1 \leq j \leq k \leq n_2$  and  $S$  is the similarity scoring function defined in Equation 1.2. Note that there may be more than one. The optimal *local* alignment of  $s_1$  and  $s_2$  is the highest scoring optimal alignment of all possible substrings of  $s_1$  and  $s_2$ , that is,

$$\alpha^*(s_1, s_2) = \operatorname{argmax}_{h,i,j,k} S(\alpha^*(s_1[h \dots i], s_2[j \dots k])),$$

where  $1 \leq h \leq i \leq n_1$ , and  $1 \leq j \leq k \leq n_2$  and  $S$  is the similarity scoring function defined in Equation 1.2. Note that there may be more than one optimal local alignment. High scoring sub-optimal alignments may also be of interest.

For local alignment, the pairwise alignment dynamic programming algorithm must be modified to allow the alignment to start and stop anywhere in  $s_1$  and  $s_2$ . Unlike the dynamic programs for global and semiglobal alignment, the local alignment recurrence (Equation 1.8) has a fourth term that sets the score  $\mathcal{A}[i, j]$  to zero whenever adding a substitution or a gap to the alignment results in a negative score. This is what allows the local alignment algorithm to consider all possible starting positions in  $s_1$  and in  $s_2$ .

Local alignment with similarity scoring:

*Initialization:*

Set the first row and column to zero:

$$\begin{aligned}\mathcal{A}[0, j] &= 0, \forall_j \\ \mathcal{A}[i, 0] &= 0, \forall_i\end{aligned}$$

*Recurrence relation:*

$$\mathcal{A}[i, j] = \max \begin{cases} \mathcal{A}[i-1, j] + g \\ \mathcal{A}[i-1, j-1] + p(s_1[i], s_2[j]) \\ \mathcal{A}[i, j-1] + g \\ 0 \end{cases} \quad (1.8)$$

For non-zero entries in  $\mathcal{A}$ , store the indices of the entry (or entries) in  $\mathcal{A}$  that maximize the right hand side of Equation 1.8 in traceback matrix,  $\mathcal{T}$ .

*Optimal alignment score:*

The score of the optimal alignment is

$$\max_{i,j} \mathcal{A}[i, j],$$

where the maximum is taken over all  $i$  and all  $j$ .

*Trace back:*

Trace back starting at  $\mathcal{T}[i^*, j^*]$ , where  $(i^*, j^*) = \operatorname{argmax}_{i,j} \mathcal{A}[i, j]$  is the cell corresponding to the maximum score in the entire matrix  $\mathcal{A}$ . End the trace back at the first cell with value zero encountered.

The point at which  $\mathcal{A}[i, j]$  drops below zero depends on the scoring function and critically determines what the resulting alignment will look like. For this reason, scoring functions for local alignment are subject to more stringent constraints than scoring functions for global and semi-global alignment.

In order to find biologically meaningful conserved regions, a scoring function for local pairwise alignment must satisfy the following requirements:

- $M > m > 2g$ .
- The scoring function must be a similarity function. The local alignment that minimizes the edit distance (weighted or unweighted) is the empty alignment, which tells us nothing.
- There must be at least one pair of residues,  $x$  and  $y$ , for which the similarity score  $p(x, y)$  is positive. Otherwise, the optimal alignment is always the empty alignment.
- The expected alignment score of a pair of randomly generated sequences (i.e., sequences sampled from a background distribution) must be negative.

These rules apply to general scoring functions, where the score for each pair of residues can have a different numerical value. For our simple similarity function, we require a positive score for matches ( $M > 0$ ) and a negative score for mismatches ( $m < 0$ ) and gaps ( $g < 0$ ).