

1 Computing alignments in only linear space

One of the defects of dynamic programming for all the problems we have discussed is that the dynamic programming tables use $\Theta(nm)$ space when the input strings have length n and m . (When we talk about the space used by a method, we refer to the maximum space ever in use *simultaneously*. Reused space does not add to the count of space use.) It is quite common that the limiting resource in string alignment problems is not time but space. That limit makes it difficult to handle large strings, no matter how long we may be willing to wait for the computation to finish. Therefore, it is very valuable to have methods that reduce the use of space without dramatically increasing the time requirements.

Hirschberg [2] developed an elegant and practical space reduction method that works for many dynamic programming problems. For several string alignment problems, this method reduces the required space from $\Theta(nm)$ to $O(n)$ (for $n < m$) while only *doubling* the worst case time bound. Miller and Myers expanded on the idea and brought it to the attention of the computational biology community [3]. The method has since been extended and applied to many more problems [1]. We illustrate the method using the dynamic programming solution to the problem of computing the optimal weighted global alignment of two strings.

1.1 Space reduction for computing similarity

Recall that the *similarity* of two strings is a *number*, and that under the similarity objective function there is an optimal alignment whose value equals that number. Now *if* we only require the *similarity* $V(n, m)$, and not an actual alignment with that value, then the maximum space needed (in addition to the space for the strings) can be reduced to $2m$. The idea is that when computing V values for row i , the only values needed from previous rows are from row $i - 1$; any rows before $i - 1$ can be discarded. This observation is clear from the recurrences for similarity. Thus, we can implement the dynamic programming solution using only two rows, one called row C for *current*, and one called row P for *previous*. In each iteration, row C is computed using row P , the recurrences and the two strings. When that row C is completely filled in, the values in row P are no longer needed and C gets copied to P to prepare for the next iteration. After n iterations, row C holds the values for row n of the full table and hence $V(n, m)$ is located in the last cell of that row. In this way, $V(n, m)$ can be computed in $O(m)$ space and $O(nm)$ time. In fact, *any* single row of the full table can be found and stored in those same time and space bounds. This ability will be critical in the method to come.

As a further refinement of this idea, the space needed can be reduced to *one* row plus one additional cell (in addition to the space for the strings). So $m + 1$ space is all that is needed. And, if $n < m$ then space use can be further reduced to $n + 1$. We leave the details as an exercise.

1.2 How to find the optimal alignment in linear space

The above idea is fine *if* we only want the similarity $V(n, m)$ or just want to store one preselected row of the dynamic programming table. But what can we do if we actually want an *alignment* that achieves value $V(n, m)$? In most cases it is such an alignment that is sought, not just its value. In the basic algorithm, the alignment would be found by traversing the pointers set while computing the full dynamic programming table for similarity. But the above linear space method does not store the whole table and linear space is not sufficient to store the pointers.

The high-level idea for finding the optimal alignment in only linear space is to do *several* smaller alignment computations, each using only linear space and each determining a bit more about an actual optimal alignment. The net result of these computations is a full description of an optimal alignment. We first describe how the initial piece of the full alignment is found using only linear space.

Definition For any string α , let α^r denote the reverse of string α .

Definition Given strings S_1 and S_2 , define $V^r(i, j)$ as the similarity of the string consisting of the first i characters of S_1^r , and the string consisting of the first j characters of S_2^r . Equivalently, $V^r(i, j)$ is the similarity of the last i characters of S_1 and the last j characters of S_2 (see Figure 1).




Figure 1: The similarity of the first i characters of S_1^r and the first j characters of S_2^r equals the similarity of the last i characters of S_1 and the last j characters of S_2 . (The dotted lines denote the substrings being aligned.)

Clearly, the table of $V^r(i, j)$ values can be computed in $O(nm)$ time, and any single preselected row of that table can be computed and stored in $O(nm)$ time using only $O(m)$ space.

The initial piece of the full alignment is computed in linear space by computing $V(n, m)$ in two parts. The first part uses the original strings, and the second part uses the reverse strings. The details of this two part computation are suggested in the following lemma.

Lemma 1.1 $V(n, m) = \max_{0 \leq k \leq m} [V(n/2, k) + V^r(n/2, m - k)]$.

Proof This is almost obvious, and yet it requires a proof. Recall that $S_1[1..i]$ is the prefix of string S_1 consisting of the first i characters, and $S_1^r[1..i]$ is the reverse of the suffix of S_1 consisting of the last i characters of S_1 . Similar definitions hold for S_2 and S_2^r .

For any fixed position k' in S_2 , there is an alignment of S_1 and S_2 consisting of an alignment of $S_1[1..n/2]$ and $S_2[1..k']$ followed by a disjoint alignment of $S_1[n/2 + 1..n]$ and $S_2[k' + 1..m]$. By definition of V and V^r , the best alignment of the first type has value $V(n/2, k')$ and the best alignment of the second type has value $V^r(n/2, m - k')$, so the combined alignment has value $V(n/2, k') + V^r(n/2, m - k') \leq \max_k [V(n/2, k) + V^r(n/2, m - k)] \leq V(n, m)$.

Conversely, consider an optimal alignment of S_1 and S_2 . Let k' be the rightmost position in S_2 that is aligned with a character at or before position $n/2$ in S_1 . Then the optimal alignment of S_1 and S_2 consists of an alignment of $S_1[1..n/2]$ and $S_2[1..k']$ followed by an alignment of $S_1[n/2 + 1..n]$ and $S_2[k' + 1..m]$. Let the value of the first alignment be denoted p and the value of the second alignment be denoted q . Then p must be equal to $V(n/2, k')$, for if $p < V(n/2, k')$ we could replace the alignment of $S_1[1..n/2]$ and $S_2[1..k']$ with the alignment of $S_1[1..n/2]$ and $S_2[1..k']$ that has value $V(n/2, k')$. That would create an alignment of S_1 and S_2 whose value is larger than the claimed optimal. Hence $p = V(n/2, k')$. By similar reasoning, $q = V^r(n/2, m - k')$. So $V(n, m) = V(n/2, k') + V^r(n/2, m - k') \leq \max_k [V(n/2, k) + V^r(n/2, m - k)]$.

Having shown both sides of the inequality, we conclude that $V(n, m) = \max_k [V(n/2, k) + V^r(n/2, m - k)]$. \square

Definition: Let k^* be a position k that maximizes $[V(n/2, k) + V^r(n/2, m - k)]$.

By Lemma 1.1, there is an optimal alignment whose traceback path in the full dynamic programming table (if one had filled in the full n by m table) goes through cell $(n/2, k^*)$. Another way to say this is that there is an optimal (longest) path L from node $(0, 0)$ to node (n, m) in the alignment graph which goes through node $(n/2, k^*)$. That is the key feature of k^* .

Definition Let $L_{n/2}$ be the subpath of L that starts with the last node of L in row $n/2 - 1$, and ends with the first node of L in row $n/2 + 1$.

Lemma 1.2 *A position k^* in row $n/2$ can be found in $O(nm)$ time and $O(m)$ space. Moreover, a subpath $L_{n/2}$ can be found and stored in those time and space bounds.*

Proof First, execute dynamic programming to compute the optimal alignment of S_1 and S_2 , but stop after iteration $n/2$, i.e., after the values in row $n/2$ have been computed. Moreover, when filling in row $n/2$, establish and save the normal traceback pointers for the cells in that row. At this point, $V(n/2, k)$ is known for every $0 \leq k \leq m$. Following the earlier discussion, only $O(m)$ space is needed to obtain the values and pointers in rows $n/2$. Second, begin computing the optimal alignment of S_1^r and S_2^r but stop after iteration $n/2$. Save both the values for cells in row $n/2$ along with the traceback pointers for those cells. Again, $O(m)$ space suffices

and value $V^r(n/2, m - k)$ is known for every k . Now, for each k , add $V(n/2, k)$ to $V^r(n/2, m - k)$, and let k^* be an index k that gives the largest sum. These additions and comparisons take $O(m)$ time.

Using the first set of saved pointers, follow any traceback path from cell $(n/2, k^*)$ to a cell k_1 in row $n/2 - 1$. This identifies a subpath that is on an optimal path from cell $(0, 0)$ to cell $(n/2, k^*)$. Similarly, using the second set of traceback pointers, follow any traceback path from cell $(n/2, k^*)$ to a cell k_2 in row $n/2 + 1$. That path identifies a subpath of an optimal path from $(n/2, k^*)$ to (n, m) . These two subpaths taken together form the subpath $L_{n/2}$ that is part of an optimal path L from $(0, 0)$ to (n, m) . Moreover, that optimal path goes through cell $(n/2, k^*)$. Overall, $O(nm)$ time and $O(m)$ space is used to find k^*, k_1, k_2 and $L_{n/2}$. \square

In order to analyze the full method to come, we will express the time needed to fill in the dynamic programming table of size p by q as cpq , for some unspecified constant c , rather than as $O(pq)$. In that view, the $n/2$ row of the first dynamic program computation is found in $cnm/2$ time, as is the $n/2$ row of the second computation. Thus, a total of cnm time is needed to obtain and store both rows.

The key point to note is that with a cnm -time and $O(m)$ -space computation, the algorithm learns k^*, k_1, k_2 and $L_{n/2}$. This specifies part of an optimal *alignment* of S_1 and S_2 , and not just the value $V(n, m)$. By Lemma 1.1 it learns that there is an optimal alignment of S_1 and S_2 consisting of an optimal alignment of the first $n/2$ characters of S_1 with the first k^* characters of S_2 , followed by an optimal alignment of the last $n/2$ characters of S_1 with the last $m - k^*$ characters of S_2 . In fact, since the algorithm has also learned the subpath (subalignment) $L_{n/2}$, the problem of aligning S_1 and S_2 reduces to two smaller alignment problems, one for the strings $S_1[1..n/2 - 1]$ and $S_2[1..k_1]$, and one for the strings $S_1[n/2 + 1..n]$ and $S_2[k_2..m]$. We call the first of the two problems the *top* problem and the second the *bottom* problem. Note that the top problem is an alignment problem on strings of lengths at most $n/2$ and k^* , while the bottom problem is on strings of lengths at most $n/2$ and $m - k^*$.

In terms of the dynamic programming table, the top problem is computed in section A of the original n by m table shown in Figure 2, and the bottom problem is computed in section B of the table. The rest of the table can be ignored. Again, we can determine the values in the middle row of A (or B) in time proportional to the total size of A (or B). Hence the middle row of the top problem can be determined at most $ck^*n/2$ time, and the middle row in the bottom problem can be determined in at most $c(m - k^*)n/2$ time. These two times add to $cnm/2$. This leads to the full idea for computing the optimal alignment of S_1 and S_2 .

1.3 The full idea: Use recursion

Having reduced the original n by m alignment problem (for S_1 and S_2) to two smaller alignment problems (the top and bottom problems) using $O(nm)$ time and $O(m)$ space, we now solve the top and bottom problems by a *recursive* application of this




Figure 2: After finding k^* , the alignment problem reduces to finding an optimal alignment in section A of the table, and another optimal alignment in section B of the table. The total area of subtables A and B is at most $cnm/2$. The subpath $L_{n/2}$ through cell $(n/2, k^*)$ is represented by a dashed path.

reduction. (For now, we ignore the space needed to save the subpaths of L). Applying exactly the same idea as was used to find k^* in the n by m problem, the algorithm uses $O(m)$ space to find the best column in row $n/4$ to break up the top $n/2$ by k_1 alignment problem. Then it *reuses* $O(m)$ space to find the best column to break up the bottom $n/2$ by $m - k_2$ alignment problem. Stated another way, we have two alignment problems, one on a table of size at most $n/2$ by k^* and another on a table of size at most $n/2$ by $m - k^*$. We can therefore find the best column in the middle row of each of the two subproblems in at most $cnk^*/2 + cn(m - k^*)/2 = cnm/2$ time, and recurse from there with four subproblems.

Continuing in this recursive way, we can find an optimal alignment of the two original strings with $\log_2 n$ levels of recursion, and at no time do we ever use more than $O(m)$ space. For convenience, assume that n is a power of two so that each successive halving gives a whole number. At each recursive call, we also find and store a subpath of an optimal path L , but these subpaths are edge disjoint, and so their total length is $O(n + m)$. In summary, the recursive algorithm we need is:

Hirschberg's Linear space optimal alignment algorithm

Procedure $OPTA(l, l', r, r')$;

begin

$h := (l' - l)/2$;

In $O(l' - l) = O(m)$ space, find an index k^* between l and l' , inclusively, such that there is an optimal alignment of $S_1[l..l']$ and $S_2[r..r']$ consisting of an optimal alignment of $S_1[l..h]$ and $S_2[r..k^*]$ followed by an optimal alignment of $S_1[h + 1..l']$ and

$S_2[k^* + 1..r']$. Also find and store the subpath L_h that is part of an optimal (longest) path L' from cell (l, r) to cell (l', r') , and that begins with the the last cell k_1 on L' in row $h - 1$ and ends with the first cell k_2 on L' in row $h + 1$. This is done as described earlier.

Call $OPTA(l, h - 1, r, k_1)$; {new top problem}
 Output subpath L_h ;
 Call $OPTA(h + 1, l', k_2, r')$; {new bottom problem }
 end.

The call that begins the computation is to $OPTA(1, n, 1, m)$. Note that the subpath L_h is output between the two $OPTA$ calls, and that the top problem is called before the bottom problem. The effect is that the subpaths are output in order of *increasing* h value, so that their concatenation describes an optimal path L from $(0, 0)$ to (n, m) , and hence an optimal alignment of S_1 and S_2 .

1.4 Time analysis

We have seen that the first level of recursion uses cnm time and the second level uses at most $cnm/2$ time. At the i 'th level of recursion, we have 2^{i-1} subproblems, each of which has $n/2^{i-1}$ rows but a variable number of columns. However, the columns in these subproblems are distinct so the total size of all the problems is at most the total number of columns, m , times $n/2^{i-1}$. Hence the total time used at the i 'th level of recursion is at most $cnm/2^{i-1}$. The final dynamic programming pass to describe the optimal alignment takes cnm time. Therefore,

Theorem 1.1 *Using Hirschberg's procedure $OPTA$, an optimal alignment of two strings of length n and m can be found in $\sum_{i=1}^{\log n} cnm/2^{i-1} \leq 2cnm$ time and $O(m)$ space.*

For comparison, recall that cnm time is used by the original method of filling in the full n by m dynamic programming table. Hirschberg's method reduces the space use from $\Theta(nm)$ to $\Theta(m)$ while only doubling the worst case time needed for the computation.

References

- [1] K. M. Chao, R. Hardison, and W. Miller. Recent developments in linear-space alignment methods: a mini survey. *J. of Comp. Biology*, 1:271–291, 1994.
- [2] Daniel S. Hirschberg. Algorithms for the longest common subsequence problem. *J. ACM*, 24:664–675, 1977.
- [3] Eugene W. Myers and Webb Miller. Optimal alignments in linear space. *Comp. Apps. in the BioSciences*, 4:11–17, 1988.