


Tutorial Exercise 4: Dynamic Programming

1. **Bellman-Ford Examples.** This exercise illustrates the execution of two different versions of the Bellman-Ford algorithm, in cases with and without negative t-connected cycles.


1a) First, consider the algorithm described on slide 6 of the Dynamic Programming in Graphs lecture notes. Trace the execution of this algorithm on the directed graph $G = (V, E, w)$ shown below. Suppose the iterations through the vertices V occur in “alphabetic” order (i.e., in the sequence (a, b, \dots, s, t)). Show the value $M(2, v)$ after the **first two** iterations through all the vertices.



You can mark the value of $M(i, v)$, after the i^{th} iteration, beside each vertex (crossing out any previous values when the value changed). Recall that at the beginning of the i^{th} iteration, $M(i-1, v) = OPT(i-1, v)$, which is the minimum cost of a $v-t$ path that uses at most $i-1$ edges. And, at the end of the i^{th} iteration, $M(i, v) = OPT(i, v)$. If you like, you could also keep track of the successor to each vertex v , which is defined to be the next vertex on a $v-t$ path that achieves this minimum distance $M(i, v)$.

1b) In the lectures we noted that this algorithm could terminate when it first detects that either: a) $M(i, v) = M(i-1, v)$ for all vertices v , or b) it completes the iteration for $i = |V|$ (i.e., one more iteration than on slide 6, since this allows us to detect cycles). After what iteration number, i , does this algorithm terminate when given the graph in part (1a) as input?

1c) Next suppose we change the weight of the edge (a, b) from 5 to 1 (which is shown below). Show $M(i, v)$ on the figure below after the fifth and sixth iteration. Note that there is at least one vertex v at which $M(i, v) < M(i-1, v)$ for all $i \leq 6$. Draw all the successor edges at the end of the sixth iteration. Is this “successor” graph acyclic?




1d) Consider the push-based implementation that was briefly described in the lecture notes (slide 9 of Dynamic Programming in Graphs). A detailed function definition is given below. We assume that the loops over the vertices are executed in the same order as before, namely in the order (a, b, c, d, s, t) . (This order now matters, as we shall see.)

```

Push-Based-Shortest-Path( $G, t$ )
  foreach node  $v \in V$  {
     $M[v] \leftarrow \infty$ ;  $C\_prev[v] \leftarrow false$ 
     $successor[v] \leftarrow \phi$ 
  }
   $M[t] = 0$ ;  $C\_prev[t] = true$ 
  for  $i = 1$  to  $|V|$  {
     $C\_cur(v) = false$  for all  $v \in V$ 
    foreach node  $w \in V$  with  $(C\_prev[w] \vee C\_cur[w])$  true {
      foreach node  $v$  such that  $(v, w) \in E$  {
        if  $(M[v] > M[w] + c_{vw})$  {
           $M[v] \leftarrow M[w] + c_{vw}$ 
           $C\_cur(v) \leftarrow true$ 
           $successor[v] \leftarrow w$ 
        }
      }
    }
    If  $C\_cur[w]$  is false for all  $w \in V$ , break.
     $C\_prev[w] \leftarrow C\_cur[w]$  for all  $w \in V$ 
  }
  return  $M, successor, any(C\_cur)$ 

```

After just two iterations through the loop over i , show the $M(v)$ computed by this push-based implementation applied to the graph in part (1c). (For your convenience we have redrawn this graph below.) Compare this with the result $M(2, v)$ of two iterations of the original implementation (now applied to this new graph). You can indicate $M(2, v)$ on the second copy of the graph below. After each iteration, is it true that $M(v) = OPT(i, v)$? How about $M(v) \leq OPT(i, v)$?



2. **Longest Increasing Subsequence.** Given an array of n integers $[x(1), \dots, x(n)]$ define an **increasing subsequence** of x , to be a subsequence $x(j(1)) < x(j(2)) < \dots < x(j(p))$ where the indicies $j(i)$ are also strictly increasing with i , that is, $1 \leq j(1) < j(2) < \dots < j(p) \leq n$. We are interested in a **longest increasing subsequence** of x , i.e., one of the ones with the maximum length p . Define $LLIS(x)$ to be the **length** of a longest increasing subsequence of x (there may be more than one increasing subsequence with this length).

For example, given the array $x = [22, 5, 8, -3, 10, 1]$ the longest increasing subsequence is $[5, 8, 10]$ which has length is 3. Therefore, for this example, $LLIS(x) = 3$.

In addition, define the array $[q(1), \dots, q(n)]$ to have elements $q(k)$ which equal the length of the longest increasing subsequence of x **ending with the value** $x(k)$. In particular, such a subsequence only uses values from the first k elements of x , namely $[x(1), \dots, x(k)]$ and must end with $x(k)$ itself.

For the example array x above, $[q(1), \dots, q(6)] = [1, 1, 2, 1, 3, 2]$.

2a) Give an equation (a recurrence relation) which expresses $q(k)$ for $1 < k \leq n$ in terms of the values of $q(j)$ for various $j < k$ and possibly other simple expressions involving elements of the array x . Explain.

2b) Express $LLIS(x)$ in terms of the array q . Explain.

2c) Given the arrays x and q defined above, provide pseudo-code for an algorithm which extracts a longest increasing subsequence. (In this case precisely worded English sentences describing simple steps are allowed in your pseudo-code.) In the previous notation, the algorithm should return $[x(j(1)), x(j(2)), \dots, x(j(p))]$ for the maximum p . The $j(k)$'s themselves do not need to be returned, just the corresponding x values (in order).

2d) Briefly explain why your algorithm is correct.

3. **Pseudo-Polynomial Time.** In class we briefly discussed the fact that the dynamic programming solution to the Knapsack problem is not a polynomial time solution. We revisit that here.

The issue is that a polynomial time algorithm must formally have a runtime that is bounded by a polynomial in the size of the input. What's a simple measure for the size of the input for an instance of the Knapsack problem? The input, say I , includes a list of n pairs of values and weights, say $\{(v_i, w_i)\}_{i=1}^n$, along with another integer, W , which is the capacity of the backpack. If all these integers were specified in B bits, then the input size for n , W and all pairs, would be $(2n+2)B$ bits. (This will suffice for defining the input size. Fortunately, due to approximation properties of big-Oh, we will not need a more accurate estimate of the number of bits needed to specify the input. But, for the insatiably curious, see Information Theory.)

We say that an algorithm runs in polynomial time iff there is some constant $q > 0$ such that, for all input I , the runtime $T(I)$ satisfies

$$T(I) \in O(|I|^q). \quad (1)$$

That is, as the size of the input $|I|$ goes to infinity, the runtime must remain bounded by a constant times this monomial $|I|^q$.

In class we showed the dynamic programming solution to Knapsack runs in time $\Theta(nW)$. Show that there is no q for which (1) is true.

Hint 1: You can treat n as fixed and consider only the effect of increasing B .

Hint 2: Look up l'Hospital's rule.

4. **Checkerboard Pebbling.** Suppose you have something similar to a checkerboard, with every square on the board coloured black or red, and with these two colours alternating along every row and column. (As a result, the colours are the same on diagonals.) Moreover, this board has an integer valued "score" listed on each square, say $S(i, j)$, which can be both positive, negative or zero. For this problem, we assume the board has only four rows, and $n > 0$ columns, and you have $2n$ pebbles.

You can select a square, and gain the score for that square, by placing a pebble on it. The constraints on placing your pebbles is that you cannot place two pebbles on two squares in the same row and neighbouring columns (say, row i and both columns j and $j+1$), nor on two squares in the same column and neighbouring rows (say, column j and both rows i and $i+1$). (Note that placing pebbles along any diagonal in the board is allowed, as is placing two pebbles on just the black squares in a column, as are many other configurations.) Pebble placements that satisfy these constraints are said to be **feasible**.

The problem is to place some or all of your pebbles on the board in a feasible way such that you maximize the sum of the scores $S(i, j)$ for the squares which have a pebble on them. Moreover, we wish to use a dynamic programming approach to find a solution of this problem.

4a) Consider the following choice of subproblems. For each $1 \leq i \leq 4$, $1 \leq j \leq n$ and $u = 0, 1$, define

$$\begin{aligned} OPT(i, j, u) \equiv & \text{Maximum score for feasibly placing pebbles in the first } j-1 \text{ columns; and} \\ & \text{in the first } i-1 \text{ rows (if any) of column } j; \text{ and, if } u = 1, \text{ then placing a pebble on} \\ & \text{the square } (i, j), \text{ otherwise leaving that square empty.} \end{aligned} \quad (2)$$

Assume that $OPT(i, 0, 0) = 0$ and $OPT(i, 0, 1) = -\infty$ for all i . You can assume the score is $-\infty$ for any illegal placement.

Either: A) derive a recurrence in terms of just these sub-problems and the scores $S(i, j)$; or, B) if you find that additional subproblems are needed then clearly explain why and explain what they might be. In case (B) you do not need to derive the recurrence relation for these new subproblems.

4b) An alternative approach to this problem is to first define all feasible configurations for one column, and then do dynamic programming over these column configurations.

Specifically, suppose we denote the location of pebbles on one column using a four bit vector, $\vec{b} = (b_1, b_2, b_3, b_4)$, where $b_i = 1$ iff there is a pebble on row i in this column. Then there are eight feasible binary patterns (i.e., which don't have two 1's as neighbours). We can denote these bit patterns as $\vec{b}(k)$ for $k = 1, 2, \dots, 8$, and we will use the notation $b_i(k)$ to denote the i^{th} bit of $\vec{b}(k)$. We assume that the first such pattern is $\vec{b}(1) = (0, 0, 0, 0)$.

What is the recurrence relation for the following choice of subproblems? Define

$$\begin{aligned} OPT(k, j) \equiv & \text{Maximum score for placing pebbles in the first } j - 1 \text{ columns, and} \\ & \text{placing pebbles in column } j \text{ according to the pattern } \vec{b}(k). \end{aligned} \quad (3)$$

Here you can assume $OPT(1, 0) = 0$, $OPT(k, 0) = -\infty$ for $1 < k \leq 8$, and the score for any infeasible arrangement is $-\infty$.

If it is not possible to define such a recurrence relation, explain why. Otherwise, explain why your recurrence relation is correct and state the runtime order for solving this problem.