

Dynamic Programming: Improving Space and/or Time

Lecture 6

September 9, 2016

What is Dynamic Programming?

Every recursion can be memoized. Automatic memoization does not help us understand whether the resulting algorithm is efficient or not.

Dynamic Programming:

A recursion that when memoized leads to an *efficient* algorithm.

Key Questions:

- Given a recursive algorithm, how do we analyze the complexity when it is memoized?
- How do we recognize whether a problem admits a dynamic programming based efficient algorithm?
- How do we further optimize time and space of a dynamic programming based algorithm?

Part I

Edit Distance

Edit Distance

Definition

Edit distance between two words X and Y is the number of letter insertions, letter deletions and letter substitutions required to obtain Y from X .

Example

The edit distance between FOOD and MONEY is at most **4**:

FOOD \rightarrow MOOD \rightarrow MONOD \rightarrow MONED \rightarrow MONEEY

Edit Distance: Alternate View

Alignment

Place words one on top of the other, with gaps in the first word indicating insertions, and gaps in the second word indicating deletions.

F	O	O		D
M	O	N	E	Y

Formally, an **alignment** is a set M of pairs (i, j) such that each index appears at most once, and there is no “crossing”: $i < i'$ and i is matched to j implies i' is matched to $j' > j$. In the above example, this is $M = \{(1, 1), (2, 2), (3, 3), (4, 5)\}$. Cost of an alignment is the number of mismatched columns plus number of unmatched indices in both strings.

Edit Distance Problem

Problem

Given two words, find the edit distance between them, i.e., an alignment of smallest cost.

Edit Distance

Basic observation

Let $X = \alpha x$ and $Y = \beta y$

α, β : strings. x and y single characters.

Possible alignments between X and Y

α	x
β	y

or

α	x
βy	

or

αx	
β	y

Observation

Prefixes must have optimal alignment!

$$EDIST(X, Y) = \min \begin{cases} EDIST(\alpha, \beta) + [x \neq y] \\ 1 + EDIST(\alpha, Y) \\ 1 + EDIST(X, \beta) \end{cases}$$

Subproblems and Recurrence

Each subproblem corresponds to a prefix of X and a prefix of Y

Optimal Costs

Let $\text{Opt}(i, j)$ be optimal cost of aligning $x_1 \cdots x_i$ and $y_1 \cdots y_j$.
Then

$$\text{Opt}(i, j) = \min \begin{cases} [x_i \neq y_j] + \text{Opt}(i - 1, j - 1), \\ 1 + \text{Opt}(i - 1, j), \\ 1 + \text{Opt}(i, j - 1) \end{cases}$$

Base Cases: $\text{Opt}(i, 0) = i$ and $\text{Opt}(0, j) = j$

$X = x_1 x_2 \dots x_m$ and $Y = y_1 y_2 \dots y_n$, we wish to compute $\text{Opt}(m, n)$.

Matrix and DAG of Computation




Figure: Iterative algorithm in previous slide computes values in row order.

Computing in column order to save space




Figure: $M(i, j)$ only depends on previous column values. Keep only two columns and compute in column order.

Optimizing Space

① Recall

$$M(i, j) = \min \begin{cases} [x_i \neq y_j] + M(i-1, j-1), \\ 1 + M(i-1, j), \\ 1 + M(i, j-1) \end{cases}$$

- ② Entries in j th column only depend on $(j-1)$ st column and earlier entries in j th column
- ③ Only store the current column and the previous column reusing space; $N(i, 0)$ stores $M(i, j-1)$ and $N(i, 1)$ stores $M(i, j)$

Space Efficient Algorithm

```
for all  $i$  do  $N[i, 0] = i$ 
for  $j = 1$  to  $n$  do
   $N[0, 1] = j$  (* corresponds to  $M(0, j)$  *)
  for  $i = 1$  to  $m$  do
    
$$N[i, 1] = \min \begin{cases} [x_i \neq y_1] + N[i - 1, 0] \\ 1 + N[i - 1, 1] \\ 1 + N[i, 0] \end{cases}$$

  for  $i = 1$  to  $m$  do
    Copy  $N[i, 0] = N[i, 1]$ 
```

Analysis

Running time is $O(mn)$ and space used is $O(2m) = O(m)$

Finding an Optimum Solution

The DP algorithm finds the minimum edit distance in $O(nm)$ space and time.

Can find minimum edit distance in $O(m + n)$ space and $O(mn)$ time.

Previous Exercise: Find an optimum alignment in $O(mn)$ space and time.

Finding an Optimum Solution

The DP algorithm finds the minimum edit distance in $O(nm)$ space and time.

Can find minimum edit distance in $O(m + n)$ space and $O(mn)$ time.

Previous Exercise: Find an optimum alignment in $O(mn)$ space and time.

Today: Finding an optimum alignment and cost in $O((m + n) \log(m + n))$ space and $O(mn)$ time.

Divide and Conquer Approach

Fix an optimum alignment between $X[1..m]$ and $Y[1..n]$

In this optimum alignment $X[1..\frac{m}{2}]$ is aligned with $Y[1..h]$ for some h where $1 \leq h \leq n$. (Need not be unique but we can choose smallest such h). Call this $\text{Half}(X, Y)$

Suppose we can find $h = \text{Half}(X, Y)$ in time $O(mn)$ time and $O(m + n)$ space, that is, in the same time as finding $\text{Opt}(m, n)$ the optimum value of the alignment between X and Y .

Divide and Conquer Algorithm

Linear-Space-Alignment($X[1..m], Y[1..n]$)

If $m = 1$ use basic algorithm in $O(n)$ time and $O(n)$ space

If $n = 1$ use basic algorithm in $O(m)$ time and $O(n)$ space

Compute $h = \text{Half}(X, Y)$ in $O(mn)$ time and $O(m + n)$ space

Linear-Space-Alignment($X[1..m/2], Y[1..h]$)

Linear-Space-Alignment($X[m/2 + 1..m], Y[h + 1..n]$)

Output combination of the two alignments

Divide and Conquer Algorithm

Linear-Space-Alignment($X[1..m]$, $Y[1..n]$)

If $m = 1$ use basic algorithm in $O(n)$ time and $O(n)$ space

If $n = 1$ use basic algorithm in $O(m)$ time and $O(n)$ space

Compute $h = \text{Half}(X, Y)$ in $O(mn)$ time and $O(m + n)$ space

Linear-Space-Alignment($X[1..m/2]$, $Y[1..h]$)

Linear-Space-Alignment($X[m/2 + 1..m]$, $Y[h + 1..n]$)

Output combination of the two alignments

Correctness: Clear based on definition of $\text{Half}(X, Y)$.

$T(m, n)$: time bound for above algorithm

$S(m, n)$: space bound

Claim: $T(m, n) = O(mn)$ and $S(m, n) = O(m + n)$.

Time bound

$$T(m, n) \leq \begin{cases} cm & \text{if } n \leq 1 \\ cn & \text{if } m \leq 1 \\ T(m/2, h) + T(m/2, n - h) + cmn & \text{otherwise} \end{cases}$$

Time bound

$$T(m, n) \leq \begin{cases} cm & \text{if } n \leq 1 \\ cn & \text{if } m \leq 1 \\ T(m/2, h) + T(m/2, n - h) + cmn & \text{otherwise} \end{cases}$$

Claim: $T(m, n) \leq 2cmn$ by induction on $m + n$.

Time bound

$$T(m, n) \leq \begin{cases} cm & \text{if } n \leq 1 \\ cn & \text{if } m \leq 1 \\ T(m/2, h) + T(m/2, n - h) + cmn & \text{otherwise} \end{cases}$$

Claim: $T(m, n) \leq 2cmn$ by induction on $m + n$.

Inductive step:

$$\begin{aligned} T(m, n) &\leq 2chm/2 + 2c(n - h)m/2 + cmn \\ &\leq 2cnm \end{aligned}$$

Space bound

$$S(m, n) \leq \begin{cases} cm & \text{if } n \leq 1 \\ cn & \text{if } m \leq 1 \\ \max\{S(m/2, h), S(m/2, n - h), c(m + n)\} + O(1) \end{cases}$$

We can reuse space for computing **EDIST**(X, Y) and **Half**(X, Y) and storing the alignment can be accounted separately as $O(m + n)$.

Space bound

$$S(m, n) \leq \begin{cases} cm & \text{if } n \leq 1 \\ cn & \text{if } m \leq 1 \\ \max\{S(m/2, h), S(m/2, n - h), c(m + n)\} + O(1) \end{cases}$$

We can reuse space for computing **EDIST**(X, Y) and **Half**(X, Y) and storing the alignment can be accounted separately as $O(m + n)$.

Claim: $S(m, n) \leq c(m + n) + O(\log m)$.

Computing **Half**(X, Y)

Want to find h such that

$$\begin{aligned} \text{EDIST}(X, Y) = & \text{EDIST}(X[1..m/2], Y[1..h]) \\ & + \text{EDIST}(X[(m/2 + 1)..m], Y[(h + 1)..n]) \end{aligned}$$

Computing Half(X, Y)

Want to find h such that

$$\begin{aligned}\text{EDIST}(X, Y) = & \text{EDIST}(X[1..m/2], Y[1..h]) \\ & + \text{EDIST}(X[(m/2 + 1)..m], Y[(h + 1)..n])\end{aligned}$$

Instead compute for all k where $1 \leq k \leq n$,

$\text{EDIST}(X[1..m/2], Y[1..k])$ &
 $\text{EDIST}(X[(m/2 + 1)..m], Y[(k + 1)..n])$

and compute h as

$$\min_k (\text{EDIST}(X[1..m/2], Y[1..k]) + \text{EDIST}(X[(m/2 + 1)..m], Y[(k + 1)..n]))$$

Computing Half(X, Y)

Compute for all k where $1 \leq k \leq n$, $\text{EDIST}(X[1..m/2], Y[1..k])$

Computing Half(X, Y)

Compute for all k where $1 \leq k \leq n$, $\text{EDIST}(X[1..m/2], Y[1..k])$

Claim: All values available if we compute $\text{EDIST}(X[1..m/2], Y[1..n])$ which we can do in $O(mn)$ time.

Can we do it in $O(m + n)$ space?

Yes! Use the space saving trick in computing edit distance and store the last row!

Computing $\text{Half}(X, Y)$

Compute for all k where $1 \leq k \leq n$,
 $\text{EDIST}(X[(m/2 + 1)..m], Y[(k + 1)..n])$

If we compute $\text{EDIST}(X[(m/2 + 1)..m], Y[1..n])$ we get the values $\text{EDIST}(X[(m/2 + 1)..m], Y[1..k])$ for $1 \leq k \leq n$ which is not what we quite want.

Observation: $\text{EDIST}(X, Y) = \text{EDIST}(X^{\text{rev}}, Y^{\text{rev}})$.

Computing Half(X, Y)

Compute for all k where $1 \leq k \leq n$,
 $\text{EDIST}(X[(m/2 + 1)..m], Y[(k + 1)..n])$

If we compute $\text{EDIST}(X[(m/2 + 1)..m], Y[1..n])$ we get the values $\text{EDIST}(X[(m/2 + 1)..m], Y[1..k])$ for $1 \leq k \leq n$ which is not what we quite want.

Observation: $\text{EDIST}(X, Y) = \text{EDIST}(X^{\text{rev}}, Y^{\text{rev}})$.

Hence compute $\text{EDIST}(A, B)$ where A is reverse of $X[(m/2 + 1)..m]$ and B is reverse of $Y[1..n]$ and this will give all the desired values.

Part II

Longest Increasing Subsequence

Sequences

Definition

Sequence: an ordered list a_1, a_2, \dots, a_n . **Length** of a sequence is number of elements in the list.

Definition

a_{i_1}, \dots, a_{i_k} is a **subsequence** of a_1, \dots, a_n if
 $1 \leq i_1 < i_2 < \dots < i_k \leq n$.

Definition

A sequence is **increasing** if $a_1 < a_2 < \dots < a_n$. It is **non-decreasing** if $a_1 \leq a_2 \leq \dots \leq a_n$. Similarly **decreasing** and **non-increasing**.

Sequences

Example...

Example

- ① Sequence: **6, 3, 5, 2, 7, 8, 1, 9**
- ② Subsequence of above sequence: **5, 2, 1**
- ③ Increasing sequence: **3, 5, 9, 17, 54**
- ④ Decreasing sequence: **34, 21, 7, 5, 1**
- ⑤ Increasing subsequence of the first sequence: **2, 7, 9.**

Longest Increasing Subsequence Problem

Input A sequence of numbers a_1, a_2, \dots, a_n

Goal Find an **increasing subsequence** $a_{i_1}, a_{i_2}, \dots, a_{i_k}$ of maximum length

Longest Increasing Subsequence Problem

Input A sequence of numbers a_1, a_2, \dots, a_n

Goal Find an **increasing subsequence** $a_{i_1}, a_{i_2}, \dots, a_{i_k}$ of maximum length

Example

- 1 Sequence: 6, 3, 5, 2, 7, 8, 1
- 2 Increasing subsequences: 6, 7, 8 and 3, 5, 7, 8 and 2, 7 etc
- 3 Longest increasing subsequence: 3, 5, 7, 8

Recursive Algorithm

Definition

LISEnding($A[1..n]$): length of longest increasing sub-sequence that *ends* in $A[n]$.

Question: can we obtain a recursive expression?

Recursive Algorithm

Definition

LISEnding($A[1..n]$): length of longest increasing sub-sequence that *ends* in $A[n]$.

Question: can we obtain a recursive expression?

$$\text{LISEnding}(A[1..n]) = \max_{i:A[i]<A[n]} \left(1 + \text{LISEnding}(A[1..i]) \right)$$

Example

Sequence: $A[1..8] = 6, 3, 5, 2, 7, 8, 1, 9$

Recursive Algorithm

```
LIS_ending_alg( $A[1..n]$ ):  
  if ( $n = 0$ ) return 0  
   $m = 1$   
  for  $i = 1$  to  $n - 1$  do  
    if ( $A[i] < A[n]$ ) then  
       $m = \max(m, 1 + \text{LIS\_ending\_alg}(A[1..i]))$   
  return  $m$ 
```

```
LIS( $A[1..n]$ ):  
  return  $\max_{i=1}^n \text{LIS\_ending\_alg}(A[1 \dots i])$ 
```

Recursive Algorithm

```
LIS_ending_alg( $A[1..n]$ ):  
  if ( $n = 0$ ) return 0  
   $m = 1$   
  for  $i = 1$  to  $n - 1$  do  
    if ( $A[i] < A[n]$ ) then  
       $m = \max(m, 1 + \text{LIS\_ending\_alg}(A[1..i]))$   
  return  $m$ 
```

```
LIS( $A[1..n]$ ):  
  return  $\max_{i=1}^n \text{LIS\_ending\_alg}(A[1 \dots i])$ 
```

- How many distinct sub-problems will **LIS_ending_alg**($A[1..n]$) generate?

Recursive Algorithm

```
LIS_ending_alg( $A[1..n]$ ):  
  if ( $n = 0$ ) return 0  
   $m = 1$   
  for  $i = 1$  to  $n - 1$  do  
    if ( $A[i] < A[n]$ ) then  
       $m = \max(m, 1 + \text{LIS\_ending\_alg}(A[1..i]))$   
  return  $m$ 
```

```
LIS( $A[1..n]$ ):  
  return  $\max_{i=1}^n \text{LIS\_ending\_alg}(A[1 \dots i])$ 
```

- How many distinct sub-problems will **LIS_ending_alg**($A[1..n]$) generate? $O(n)$

Recursive Algorithm

```
LIS_ending_alg( $A[1..n]$ ):  
  if ( $n = 0$ ) return 0  
   $m = 1$   
  for  $i = 1$  to  $n - 1$  do  
    if ( $A[i] < A[n]$ ) then  
       $m = \max(m, 1 + \text{LIS\_ending\_alg}(A[1..i]))$   
  return  $m$ 
```

```
LIS( $A[1..n]$ ):  
  return  $\max_{i=1}^n \text{LIS\_ending\_alg}(A[1 \dots i])$ 
```

- How many distinct sub-problems will **LIS_ending_alg**($A[1..n]$) generate? $O(n)$
- What is the running time if we memoize recursion?

Recursive Algorithm

```
LIS_ending_alg( $A[1..n]$ ):  
  if ( $n = 0$ ) return 0  
   $m = 1$   
  for  $i = 1$  to  $n - 1$  do  
    if ( $A[i] < A[n]$ ) then  
       $m = \max(m, 1 + \text{LIS\_ending\_alg}(A[1..i]))$   
  return  $m$ 
```

```
LIS( $A[1..n]$ ):  
  return  $\max_{i=1}^n \text{LIS\_ending\_alg}(A[1 \dots i])$ 
```

- How many distinct sub-problems will **LIS_ending_alg**($A[1..n]$) generate? $O(n)$
- What is the running time if we memoize recursion? $O(n^2)$ since each call takes $O(n)$ time

Recursive Algorithm

```
LIS_ending_alg( $A[1..n]$ ):  
  if ( $n = 0$ ) return 0  
   $m = 1$   
  for  $i = 1$  to  $n - 1$  do  
    if ( $A[i] < A[n]$ ) then  
       $m = \max(m, 1 + \text{LIS\_ending\_alg}(A[1..i]))$   
  return  $m$ 
```

```
LIS( $A[1..n]$ ):  
  return  $\max_{i=1}^n \text{LIS\_ending\_alg}(A[1 \dots i])$ 
```

- How many distinct sub-problems will **LIS_ending_alg**($A[1..n]$) generate? $O(n)$
- What is the running time if we memoize recursion? $O(n^2)$ since each call takes $O(n)$ time
- How much space for memoization?

Recursive Algorithm

```
LIS_ending_alg( $A[1..n]$ ):  
  if ( $n = 0$ ) return 0  
   $m = 1$   
  for  $i = 1$  to  $n - 1$  do  
    if ( $A[i] < A[n]$ ) then  
       $m = \max(m, 1 + \text{LIS\_ending\_alg}(A[1..i]))$   
  return  $m$ 
```

```
LIS( $A[1..n]$ ):  
  return  $\max_{i=1}^n \text{LIS\_ending\_alg}(A[1 \dots i])$ 
```

- How many distinct sub-problems will **LIS_ending_alg**($A[1..n]$) generate? $O(n)$
- What is the running time if we memoize recursion? $O(n^2)$ since each call takes $O(n)$ time
- How much space for memoization? $O(n)$

Removing recursion to obtain iterative algorithm

Typically, after finding a dynamic programming recursion, we often convert the recursive algorithm into an *iterative* algorithm via *explicit memoization* and *bottom up* computation.

Why?

Removing recursion to obtain iterative algorithm

Typically, after finding a dynamic programming recursion, we often convert the recursive algorithm into an *iterative* algorithm via *explicit memoization* and *bottom up* computation.

Why? Mainly for further optimization of running time and space.

Removing recursion to obtain iterative algorithm

Typically, after finding a dynamic programming recursion, we often convert the recursive algorithm into an *iterative* algorithm via *explicit memoization* and *bottom up* computation.

Why? Mainly for further optimization of running time and space.

How?

- First, allocate a data structure (usually an array or a multi-dimensional array that can hold values for each of the subproblems)
- Figure out a way to order the computation of the sub-problems starting from the base case.

Iterative Algorithm via Memoization

Compute the values **LIS_ending_alg**($A[1..i]$) iteratively in a bottom up fashion.

```
LIS_ending_alg( $A[1..n]$ ):  
  Array  $L[1..n]$  (*  $L[i]$  = value of LIS_ending_alg( $A[1..i]$ ) *)  
  for  $i = 1$  to  $n$  do  
     $L[i] = 1$   
    for  $j = 1$  to  $i - 1$  do  
      if ( $A[j] < A[i]$ ) do  
         $L[i] = \max(L[i], 1 + L[j])$   
  return  $L$ 
```

```
LIS( $A[1..n]$ ):  
   $L = \text{LIS\_ending\_alg}(A[1..n])$   
  return the maximum value in  $L$ 
```

Iterative Algorithm via Memoization

Simplifying:

```
LIS(A[1..n]):  
  Array L[1..n]  (* L[i] stores the value LISEnding(A[1..i]) *)  
  m = 0  
  for i = 1 to n do  
    L[i] = 1  
    for j = 1 to i - 1 do  
      if (A[j] < A[i]) do  
        L[i] = max(L[i], 1 + L[j])  
    m = max(m, L[i])  
  return m
```

Correctness: Via induction following the recursion

Running time: $O(n^2)$

Space: $\Theta(n)$

Improving run time

Want to improve run time to $O(n \log n)$ from $O(n^2)$. How?

Idea: Use data structures to improve run-time of computing

$$\text{LISEnding}(i) = \max_{j < i: A[j] < A[i]} 1 + \text{LISEnding}(j)$$

Improving run time

Want to improve run time to $O(n \log n)$ from $O(n^2)$. How?

Idea: Use data structures to improve run-time of computing

$$\text{LISEnding}(i) = \max_{j < i: A[j] < A[i]} 1 + \text{LISEnding}(j)$$

- When computing $\text{LISEnding}(i)$ we want to focus only on indices j such that $A[j] < A[i]$
- We need to store $\text{LISEnding}(j)$ with each value $A[j]$ stored in the data structure

Augmented Balanced Binary Search Tree

Assume for simplicity that a_1, a_2, \dots, a_n are distinct numbers.

- Store a_1, a_2, \dots, a_k in a dynamic balanced binary search tree T and when iteration $k + 1$ is processed we add a_{k+1} to the tree. This ensures that when considering **LISEnding**(i) the tree T only has a_1, \dots, a_{i-1} .

Augmented Balanced Binary Search Tree

Assume for simplicity that a_1, a_2, \dots, a_n are distinct numbers.

- Store a_1, a_2, \dots, a_k in a dynamic balanced binary search tree T and when iteration $k + 1$ is processed we add a_{k+1} to the tree. This ensures that when considering **LISEnding(i)** the tree T only has a_1, \dots, a_{i-1} .
- We can search for a_i in T to obtain in $O(\log n)$ time a set of subtrees such that each subtree has only numbers smaller than a_i . Precisely what we want!

Augmented Balanced Binary Search Tree

Assume for simplicity that a_1, a_2, \dots, a_n are distinct numbers.

- Store a_1, a_2, \dots, a_k in a dynamic balanced binary search tree T and when iteration $k + 1$ is processed we add a_{k+1} to the tree. This ensures that when considering **LISEnding**(i) the tree T only has a_1, \dots, a_{i-1} .
- We can search for a_i in T to obtain in $O(\log n)$ time a set of subtrees such that each subtree has only numbers smaller than a_i . Precisely what we want!
- We store with the root of each subtree of T the max **LISEnding** value for all indices represented in that subtree.

Augmented Balanced Binary Search Tree

Assume for simplicity that a_1, a_2, \dots, a_n are distinct numbers.

- Store a_1, a_2, \dots, a_k in a dynamic balanced binary search tree T and when iteration $k + 1$ is processed we add a_{k+1} to the tree. This ensures that when considering **LISEnding**(i) the tree T only has a_1, \dots, a_{i-1} .
- We can search for a_i in T to obtain in $O(\log n)$ time a set of subtrees such that each subtree has only numbers smaller than a_i . Precisely what we want!
- We store with the root of each subtree of T the max **LISEnding** value for all indices represented in that subtree.
- Updating tree after computing **LISEnding**(i) requires inserting a_i into the tree T and also updating the **LISEnding** values. Can be done in $O(\log n)$ time. Thus, overall $O(n \log n)$ time.

Example

A better algorithm

Using only two arrays. Elegant, fast. See Wikipedia article https://en.wikipedia.org/wiki/Longest_increasing_subsequence

Not a first-cut solution.