

Dynamic Programming is a powerful technique that allows one to solve many different types of problems in time $O(n^2)$ or $O(n^3)$ for which a naive approach would take exponential time. In this lecture, we discuss this technique, and present a few key examples. Topics in this lecture include:

- The basic idea of Dynamic Programming.
- Example: Longest Common Subsequence.
- Example: Knapsack.
- Example: Matrix-chain multiplication.
- Example: Optimal Binary Search Trees.
- Example: Breaking Paragraphs into Lines.

1 Introduction

Dynamic Programming is a powerful technique that can be used to solve many problems in time $O(n^2)$ or $O(n^3)$ for which a naive approach would take exponential time. (Usually to get running time below that—if it is possible—one would need to add other ideas as well.) Dynamic Programming is a general approach to solving problems, much like “divide-and-conquer” is a general method, except that unlike divide-and-conquer, the subproblems will typically overlap. This lecture we will present two ways of thinking about Dynamic Programming as well as a few examples.

There are several ways of thinking about the basic idea.

Basic Idea (version 1): What we want to do is take our problem and somehow break it down into a reasonable number of subproblems (where “reasonable” might be something like n^2) in such a way that we can use optimal solutions to the smaller subproblems to give us optimal solutions to the larger ones. Unlike divide-and-conquer (as in mergesort or quicksort) it is OK if our subproblems overlap, so long as there are not too many of them.

2 Example 1: Longest Common Subsequence

Definition 1 *The Longest Common Subsequence (LCS) problem is as follows. We are given two strings: string S of length n , and string T of length m . Our goal is to produce their longest common subsequence: the longest sequence of characters that appear left-to-right (but not necessarily in a contiguous block) in both strings.*

For example, consider:

$S = \text{ABAZDC}$

$T = \text{BACBAD}$

In this case, the LCS has length 4 and is the string **ABAD**. Another way to look at it is we are finding a 1-1 matching between some of the letters in S and some of the letters in T such that none of the edges in the matching cross each other.

For instance, this type of problem comes up all the time in genomics: given two DNA fragments, the LCS gives information about what they have in common and the best way to line them up.

Let's now solve the LCS problem using Dynamic Programming. As subproblems we will look at the LCS of a prefix of S and a prefix of T , running over all pairs of prefixes. For simplicity, let's worry first about finding the *length* of the LCS and then we can modify the algorithm to produce the actual sequence itself.

So, here is the question: say $LCS[i, j]$ is the length of the LCS of $S[1..i]$ with $T[1..j]$. How can we solve for $LCS[i, j]$ in terms of the LCS's of the smaller problems?

Case 1: what if $S[i] \neq T[j]$? Then, the desired subsequence has to ignore one of $S[i]$ or $T[j]$ so we have:

$$LCS[i, j] = \max(LCS[i - 1, j], LCS[i, j - 1]).$$

Case 2: what if $S[i] = T[j]$? Then the LCS of $S[1..i]$ and $T[1..j]$ might as well match them up. For instance, if I gave you a common subsequence that matched $S[i]$ to an earlier location in T , for instance, you could always match it to $T[j]$ instead. So, in this case we have:

$$LCS[i, j] = 1 + LCS[i - 1, j - 1].$$

So, we can just do two loops (over values of i and j), filling in the LCS using these rules. Here's what it looks like pictorially for the example above, with S along the leftmost column and T along the top row.

	B	A	C	B	A	D
A	0	1	1	1	1	1
B	1	1	1	2	2	2
A	1	2	2	2	3	3
Z	1	2	2	2	3	3
D	1	2	2	2	3	4
C	1	2	3	3	3	4

We just fill out this matrix row by row, doing constant amount of work per entry, so this takes $O(mn)$ time overall. The final answer (the length of the LCS of S and T) is in the lower-right corner.

How can we now find the sequence? To find the sequence, we just walk backwards through matrix starting the lower-right corner. If either the cell directly above or directly to the right contains a value equal to the value in the current cell, then move to that cell (if both to, then chose either one). If both such cells have values strictly less than the value in the current cell, then move diagonally up-left (this corresponds to applying Case 2), and output the associated character. This will output the characters in the LCS in reverse order. For instance, running on the matrix above, this outputs DABA.

3 More on the basic idea, and Example 1 revisited

We have been looking at what is called "bottom-up Dynamic Programming". Here is another way of thinking about Dynamic Programming, that also leads to basically the same algorithm, but viewed from the other direction. Sometimes this is called "top-down Dynamic Programming".

Basic Idea (version 2): Suppose you have a recursive algorithm for some problem that gives you a really bad recurrence like $T(n) = 2T(n - 1) + n$. However, suppose that many of the subproblems

you reach as you go down the recursion tree are the *same*. Then you can hope to get a big savings if you store your computations so that you only compute each *different* subproblem once. You can store these solutions in an array or hash table. This view of Dynamic Programming is often called *memoizing*.

For example, for the LCS problem, using our analysis we had at the beginning we might have produced the following exponential-time recursive program (arrays start at 1):

```
LCS(S,n,T,m)
{
  if (n==0 || m==0) return 0;
  if (S[n] == T[m]) result = 1 + LCS(S,n-1,T,m-1); // no harm in matching up
  else result = max( LCS(S,n-1,T,m), LCS(S,n,T,m-1) );
  return result;
}
```

This algorithm runs in exponential time. In fact, if S and T use completely disjoint sets of characters (so that we never have $S[n]=T[m]$) then the number of times that $LCS(S,1,T,1)$ is recursively called equals $\binom{n+m-2}{m-1}$.¹ In the memoized version, we store results in a matrix so that any given set of arguments to LCS only produces new work (new recursive calls) once. The memoized version begins by initializing $arr[i][j]$ to *unknown* for all i, j , and then proceeds as follows:

```
LCS(S,n,T,m)
{
  if (n==0 || m==0) return 0;
  if (arr[n][m] != unknown) return arr[n][m]; // <- added this line (*)
  if (S[n] == T[m]) result = 1 + LCS(S,n-1,T,m-1);
  else result = max( LCS(S,n-1,T,m), LCS(S,n,T,m-1) );
  arr[n][m] = result; // <- and this line (**)
  return result;
}
```

All we have done is saved our work in line (**) and made sure that we only embark on new recursive calls if we haven't already computed the answer in line (*).

In this memoized version, our running time is now just $O(mn)$. One easy way to see this is as follows. First, notice that we reach line (**) at most mn times (at most once for any given value of the parameters). This means we make at most $2mn$ recursive calls total (at most two calls for each time we reach that line). Any given call of LCS involves only $O(1)$ work (performing some equality checks and taking a max or adding 1), so overall the total running time is $O(mn)$.

Comparing bottom-up and top-down dynamic programming, both do almost the same work. The top-down (memoized) version pays a penalty in recursion overhead, but can potentially be faster than the bottom-up version in situations where some of the subproblems never get examined at all. These differences, however, are minor: you should use whichever version is easiest and most intuitive for you for the given problem at hand.

¹This is the number of different “monotone walks” between the upper-left and lower-right corners of an n by m grid. For this course you don't need to know this bound or how to prove it—it just suggests that brute-force is not the way to go. But if you are interested, look up Catalan numbers.

More about LCS: Discussion and Extensions. An equivalent problem to LCS is the “minimum edit distance” problem, where the legal operations are insert and delete. (E.g., the unix “diff” command, where S and T are files, and the elements of S and T are lines of text). The minimum edit distance to transform S into T is achieved by doing $|S| - \text{LCS}(S, T)$ deletes and $|T| - \text{LCS}(S, T)$ inserts.

In computational biology applications, often one has a more general notion of sequence alignment. Many of these different problems all allow for basically the same kind of Dynamic Programming solution.

4 Example #2: The Knapsack Problem

Imagine you have a homework assignment with different parts labeled A through G. Each part has a “value” (in points) and a “size” (time in hours to complete). For example, say the values and times for our assignment are:

	A	B	C	D	E	F	G
value	7	9	5	12	14	6	12
time	3	4	2	6	7	3	5

Say you have a total of 15 hours: which parts should you do? If there was partial credit that was proportional to the amount of work done (e.g., one hour spent on problem C earns you 2.5 points) then the best approach is to work on problems in order of points-per-hour (a greedy strategy).

But, what if there is no partial credit? In that case, which parts should you do, and what is the best total value possible?²

Exercise: Give an example where using the greedy strategy will get you less than 1% of the optimal value (in the case there is no partial credit).

The above is an instance of the *knapsack problem*, formally defined as follows:

Definition 2 *In the knapsack problem we are given a set of n items, where each item i is specified by a size s_i and a value v_i . We are also given a size bound S (the size of our knapsack). The goal is to find the subset of items of maximum total value such that sum of their sizes is at most S (they all fit into the knapsack).*

We can solve the knapsack problem in exponential time by trying all possible subsets. With Dynamic Programming, we can reduce this to time $O(nS)$.

Let’s do this top down by starting with a simple recursive solution and then trying to memoize it. Let’s start by just computing the best possible *total value*, and we afterwards can see how to actually extract the items needed.

```
// Recursive algorithm: either we use the last element or we don't.
Value(n,S)    // S = space left, n = # items still to choose from
{
    if (n == 0) return 0;
    if (s_n > S) result = Value(n-1,S); // can't use nth item
    else result = max{v_n + Value(n-1, S-s_n), Value(n-1, S)};
}
```

²Answer: In this case, the optimal strategy is to do parts A, B, F, and G for a total of 34 points. Notice that this doesn’t include doing part C which has the most points/hour!

```

    return result;
}

```

Right now, this takes exponential time. But, notice that there are only $O(nS)$ *different* pairs of values the arguments can possibly take on, so this is perfect for memoizing. As with the LCS problem, let us initialize a 2-d array `arr[i][j]` to “unknown” for all i, j .

```

Value(n,S)
{
    if (n == 0) return 0;
    if (arr[n][S] != unknown) return arr[n][S]; // <- added this
    if (s_n > S) result = Value(n-1,S);
    else result = max{v_n + Value(n-1, S-s_n), Value(n-1, S)};
    arr[n][S] = result; // <- and this
    return result;
}

```

Since any given pair of arguments to `Value` can pass through the array check only once, and in doing so produces at most two recursive calls, we have at most $2n(S+1)$ recursive calls total, and the total time is $O(nS)$.

So far we have only discussed computing the *value* of the optimal solution. How can we get the items? As usual for Dynamic Programming, we can do this by just working backwards: if `arr[n][S] = arr[n-1][S]` then we *didn't* use the n th item so we just recursively work backwards from `arr[n-1][S]`. Otherwise, we *did* use that item, so we just output the n th item and recursively work backwards from `arr[n-1][S-s_n]`. One can also do bottom-up Dynamic Programming.

Exercise: The *fractional* knapsack problem is the one where you can add $\delta_i \in [0, 1]$ fraction of task i to the knapsack, using $\delta_i s_i$ space and getting $\delta_i v_i$ value. The greedy algorithm adds items in decreasing order of v_i/s_i . Prove that this greedy algorithm produces the optimal solution for the *fractional* problem.

5 Example #3: Matrix product parenthesization

Our final example for Dynamic Programming is the *matrix product parenthesization* problem.

Say we want to multiply three matrices X , Y , and Z . We could do it like $(XY)Z$ or like $X(YZ)$. Which way we do the multiplication doesn't affect the final outcome but it *can* affect the running time to compute it. For example, say X is 100×20 , Y is 20×100 , and Z is 100×20 . So, the end result will be a 100×20 matrix. If we multiply using the usual algorithm, then to multiply an $\ell \times m$ matrix by an $m \times n$ matrix takes time $O(\ell mn)$. So in this case, which is better, doing $(XY)Z$ or $X(YZ)$?

Answer: $X(YZ)$ is better because computing YZ takes $20 \times 100 \times 20$ steps, producing a 20×20 matrix, and then multiplying this by X takes another $20 \times 100 \times 20$ steps, for a total of $2 \times 20 \times 100 \times 20$. But, doing it the other way takes $100 \times 20 \times 100$ steps to compute XY , and then multiplying this with Z takes another $100 \times 20 \times 100$ steps, so overall this way takes 5 times longer. More generally, what if we want to multiply a series of n matrices?

Definition 3 *The Matrix Product Parenthesization problem is as follows. Suppose we need to multiply a series of matrices: $A_1 \times A_2 \times A_3 \times \dots \times A_n$. Given the dimensions of these matrices, what is the best way to parenthesize them, assuming for simplicity that standard matrix multiplication is to be used (e.g., not Strassen)?*

There are an exponential number of different possible parenthesizations, in fact $\binom{2(n-1)}{n-1}/n$, so we don't want to search through all of them.³ Dynamic Programming gives us a better way.

As before, let's first think: how might we do this recursively? One way is that for each possible split for the final multiplication, recursively solve for the optimal parenthesization of the left and right sides, and calculate the total cost (the sum of the costs returned by the two recursive calls plus the lmn cost of the final multiplication, where "m" depends on the location of that split). Then take the overall best top-level split.

For Dynamic Programming, the key question is now: in the above procedure, as you go through the recursion, what do the subproblems look like and how many are there? Answer: each subproblem looks like "what is the best way to multiply some sub-interval of the matrices $A_i \times \dots \times A_j$?" So, there are only $O(n^2)$ *different* subproblems.

The second question is now: how long does it take to solve a given subproblem assuming you've already solved all the smaller subproblems (i.e., how much time is spent inside any *given* recursive call)? Answer: to figure out how to best multiply $A_i \times \dots \times A_j$, we just consider all possible middle points k and select the one that minimizes:

$$\begin{array}{ll}
 \text{optimal cost to multiply } A_i \dots A_k & \leftarrow \text{already computed} \\
 + \text{ optimal cost to multiply } A_{k+1} \dots A_j & \leftarrow \text{already computed} \\
 + \text{ cost to multiply the results.} & \leftarrow \text{get this from the dimensions}
 \end{array}$$

This just takes $O(1)$ work for any given k , and there are at most n different values k to consider, so overall we just spend $O(n)$ time per subproblem. So, if we use Dynamic Programming to save our results in a lookup table, then since there are only $O(n^2)$ subproblems we will spend only $O(n^3)$ time overall.

If you want to do this using bottom-up Dynamic Programming, you would first solve for all subproblems with $j - i = 1$, then solve for all with $j - i = 2$, and so on, storing your results in an n by n matrix. The main difference between this problem and the two previous ones we have seen is that any *given* subproblem takes time $O(n)$ to solve rather than $O(1)$, which is why we get $O(n^3)$ total running time. It turns out that by being very clever you can actually reduce this to $O(1)$ amortized time per subproblem, producing an $O(n^2)$ -time algorithm, but we won't get into that here.⁴


³For the purposes of this course, you need not know how to derive this bound — it's only to suggest the futility of brute-force. But if you are interested, refer to 15-251 notes on counting, or to Catalan numbers.

⁴For details, see Knuth (insert ref).

6 Optimal Binary Search Trees

The keys $\{1, 2, \dots, n\}$ are stored in a binary search tree (BST). You're given a list of frequencies $[f_1, f_2, \dots, f_n]$ corresponding to the keys. The problem we address here is to compute the optimal binary search tree for these keys and corresponding frequencies.

The keys must be stored in left to right (symmetric) order in the tree. The following picture shows two possible trees that store the keys $\{1, 2, 3, 4\}$. In the situation where $f_4 = 1000$ and all the other frequencies were 1, the first tree would be better than the second.



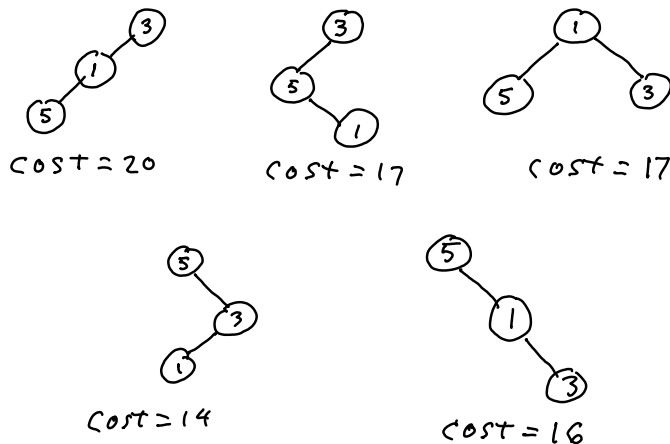
Ground Rules:

- The quality of a tree is measured by the number of comparisons done.
- Comparisons are “3-way” in the sense that a single comparison between a desired key v (being searched for) and a key k stored in a node in the tree tells us if the $v < k$, or $v = k$, or $v > k$.
- We assume that all searches are successful. That is, the only keys we are searching for are ones that are in the tree.

These assumptions are not crucial, but they serve to simplify the exposition of the algorithm.

The “depth” of a node x refers to the number of comparisons required to get from the root of the tree to x (including x itself).

To make sure these definitions are clear, consider three frequencies $[f_1, f_2, f_3] = [5, 1, 3]$. The following diagram shows the five possible trees, along with their costs incurred by doing a sequence of searches with the given frequencies.



To solve the problem, we first define $f_{i,j} = \sum_{k=1}^j f_k$. (We'll talk about computing this later.)


We also define $C_{i,j}$ to be the total cost (number of comparisons) of the optimal BST for frequencies $[f_i, f_{i+1}, \dots, f_j]$. The following recurrence defines $C_{i,j}$:

$$C_{i,j} = \begin{cases} 0 & \text{if } i > j \\ f_i & \text{if } i = j \\ \min_{i \leq k \leq j} f_{i,j} + C_{i,k-1} + C_{k+1,j} & \text{if } i < j \end{cases}$$

Let's see where this recurrence comes from. First note that

$$f_{i,j} + C_{i,k-1} + C_{k+1,j} = (f_{i,k-1} + C_{i,k-1}) + f_k + (f_{k+1,j} + C_{k+1,j})$$

This is the cost of the optimal tree (with root k) among the keys $\{i, i+1, \dots, j\}$. The three terms correspond to (1) the cost of the searches that end up in the left subtree, (2) the cost of searches for k , and (3) the cost of searches that end up in the right subtree. In case (1) there are $f_{i,k-1}$ comparisons to the node containing k , then there are $C_{i,k-1}$ additional comparisons that take place within the optimal tree formed from $\{i, \dots, k-1\}$. Case (3) is analogous. In the second case each of the f_k searches terminates here after just one comparison.



Computing each value of $C_{i,j}$ takes $O(n)$ time and there are $O(n^2)$ such values. So the running time of this algorithm is $O(n^3)$.

We should also mention how to compute the $f_{i,j}$ values. One way to do this is build a table of size n^2 and fill in the values using the following recurrence in $O(n^2)$ time.

$$f_{i,j} = \begin{cases} f_i & \text{if } i = j \\ f_{i,j-1} + f_j & \text{if } i < j \end{cases}$$

Another more efficient way is to compute $[s_0, s_1, \dots, s_n]$ where $s_i = \sum_{k \leq i} f_k$. This can be computed in linear time using $s_i = s_{i-1} + f_i$. Furthermore $f_{i,j} = s_j - s_{i-1}$.

7 Breaking Paragraphs into Lines

You're given a sequence of words $[w_1, w_2, \dots, w_n]$ which form a paragraph. To display or print this out we need to decide where to put the linebreaks within this sequence of words. (We ignore the issue of hyphenation.) This is the problem of "breaking a paragraph into lines".

For clarity we encapsulate all of the aesthetic criteria into a single function $B(i, j)$. This is a real number, and it's a measure of the "badness" of having a line comprised of the words w_i, w_{i+1}, \dots, w_j . (For example, it would take into account if the words are too long for a line, or too short, as well as the special case of this being the last line of the paragraph.) Our goal is to find the way of breaking the sequence of words up into lines that minimizes the total sum of the badnesses of all of the lines.

The way to set this up as a recurrence is to observe that any solution must have a last line. So we simply try all possible beginnings of the last line. Given that we have chosen the last line, we can remove those words and solve an analogous smaller problem of breaking the first part into lines.

This gives the following recurrence for C_j , the minimum cost way of breaking words w_1, w_2, \dots, w_j into lines.

$$C_i = \begin{cases} 0 & \text{if } i = 0 \\ \min_{0 \leq j \leq i-1} C_j + B(j+1, i) & \text{otherwise} \end{cases}$$

The running time of the algorithm resulting from this recurrence is $O(n^2)$.

8 High-level discussion of Dynamic Programming

What kinds of problems can be solved using Dynamic Programming? One property these problems have is that if the optimal solution involves solving a subproblem, then it uses the *optimal* solution to that subproblem. For instance, say we want to find the shortest path from A to B in a graph, and say this shortest path goes through C . Then it must be using the shortest path from C to B . Or, in the knapsack example, if the optimal solution does not use item n , then it is the optimal solution for the problem in which item n does not exist. The other key property is that there should be only a polynomial number of different subproblems. These two properties together allow us to build the optimal solution to the final problem from optimal solutions to subproblems.

In the top-down view of dynamic programming, the first property above corresponds to being able to write down a recursive procedure for the problem we want to solve. The second property corresponds to making sure that this recursive procedure makes only a polynomial number of *different* recursive calls. In particular, one can often notice this second property by examining the arguments to the recursive procedure: e.g., if there are only two integer arguments that range between 1 and n , then there can be at most n^2 different recursive calls.

Sometimes you need to do a little work on the problem to get the optimal-subproblem-solution property. For instance, suppose we are trying to find paths between locations in a city, and some intersections have no-left-turn rules (this is particularly bad in San Francisco). Then, just because the fastest way from A to B goes through intersection C , it doesn't necessarily use the fastest way to C because you might need to be coming into C in the correct direction. In fact, the right way to model that problem as a graph is not to have one node per intersection, but rather to have one node per $\langle \textit{intersection}, \textit{direction} \rangle$ pair. That way you recover the property you need.

8.1 Why “Dynamic Programming”?

The term dynamic programming comes from Richard Bellman, who also gave us the Bellman-Ford algorithm from the next lecture. In his autobiography he says

“Let’s take a word that has an absolutely precise meaning, namely dynamic, in the classical physical sense. It also has a very interesting property as an adjective, and that is it’s impossible to use the word, dynamic, in a pejorative sense. Try thinking of some combination that will possibly give it a pejorative meaning. [...] Thus, I thought dynamic programming was a good name. It was something not even a Congressman could object to. So I used it as an umbrella for my activities.”

Bellman clearly had a way with words, coining another super-well-known phrase, “*the curse of dimensionality*”.