Greedy Algorithms
Module 4: Techniques

Overview. Sometimes we can solve optimization problems with a technique called *greedy*. A greedy algorithm picks the option that looks best according to a *quick* evaluation criterion, without evaluating each choice recursively. Greedy algorithms are intuitive and may seem much easier to understand than dynamic programming. We'll see that it's quite easy to come up with greedy strategies. However, many times (even: most of the time) greedy strategies do not lead to optimal solutions. Before we look at a specific example we underline the similarities and differences between greedy and dynamic programming. Overall, the greedy technique is an important tool to know and gives simple and elegant solutions to important problems.

Dynamic programming:

- The problem must have the *optimal substructure property*: the optimal solution to the problem contains within it optimal solutions to subproblems. This allows for a recursive solution.

- Dynamic programming: Evaluate *all* options *recursively*, pick the best and recurse. An option is evaluated recursively, meaning all its options are evaluated and all their options are evaluated, and so on; basically to evaluate an option, you have to go through the whole recursion tree for that option.

- Dynamic programming is essentially smart recursion (recursion without repetition). The subproblems are stored in a table to ensure that they are computed at most once. Typically the number of different subproblems is small, but the recursive algorithm implementing the recursive formulation of the problem is exponential. This is because of overlapping calls to same subproblem. So if you don't use the table to cache partial solutions you incur a significant penalty (Often the difference is polynomial vs. exponential).

- Sometimes we might want to take dynamic programming a step further, and eliminate the recursion —- this is purely for eliminating the overhead of recursion, and does not change the $\Theta()$ of the running time. Iterative dynamic programming is filling the table bottom-up, without recursion.

The greedy technique:

- As with dynamic programming, in order to be solved with the greedy technique, the problem must have *optimal substructure*.

- The problems that can be solved with the greedy method are a (very small!) subset of those that can be solved with dynamic programming.

- The idea of the greedy technique is the following: Instead of evaluating *all* options *recursively* and picking the best one, we use a quick way to pick what looks locally like the best choice, add it to the solution, and repeat. So basically a greedy algorithm picks the *locally* optimal choice at each step, in the hopes to get to a solution that is *globally* optimal.

- Coming up with greedy heuristics is easy, but proving that a heuristic gives the optimal solution can be tricky.

Let's look at an example.

Example: Interval scheduling (Activity selection)

Problem: Given a set $A = \{A_1, A_2, \cdots, A_n\}$ of n activities with start and finish times (s_i, f_i), $1 \le i \le n$, find a maximal set S of non-overlapping activities.

- There exists a more general problem, the *weighted-interval scheduling* problem, where all intervals have different weights. Activity selection is a special case where all intervals have the same weight.

- A possible practical interpretation: corresponds to scheduling the maximal number of classes (given their start and finish times) in one classroom. Or more exciting: get your money's worth at Disney Land! you are in the park, you have lots of rides to chose from, you know start and end time for each ride, you want to ride as many rides as possible. Take Algorithms!

- Let's try dynamic programming: Look at the first activity, try a solution with it and one without it, and see which one is better. We can make it run in $O(n \lg n)$.

- Note that with dynamic programming we evaluate recursively both options (current activity is in or out), and pick the best. With a greedy solution, we would find a quick way to pick one or the other option.

- Here are some possible greedy algorithms for activity selection:

  Greedy 1: Pick the shortest activity, eliminate all activities that overlap with it, and recurse.

  Clearly all we need to do is sort the activities, so this would run in $O(n \lg n)$ time. Does this work? Yeah it works in the sense that it comes up with a set of activities that are non-overlapping, but is this set of activities optimal (i.e. largest set possible)? Either we need to argue that this algorithm always gives an optimal solution; or, we need to give a counter-example (an instance where this strategy does NOT give the optimal solution).

  In this case it is fairly simple to find a counter-example:

- Greedy 2: Pick the activity that starts first (eliminate all activities that overlap with it, and recurse).

  Counter-example:

- Greedy 3: Pick the activity that ends first (eliminate all activities that overlap with it, and recurse).

  Counter-example:

A greedy solution for Activity Selection

Turns out that picking activities in order of their finish time gives the correct optimal solution. We'll argue below why. The intuition is that by picking what ends first, we maximize the remaining time. First let's spell out the idea a bit more:

1. Sort activity by finish time and renumber the activities so that $A_1, A_2, \cdots, A_n$ denotes sorted sequence.

2. Pick first activity A_1 (and insert it in the output set of non-overlapping activities)

3. Compute B = set of activities in A that do not overlap with A_1.

4. Recursively solve problem on B.

Correctness

- Before we think about running time, we need to argue that this algorithm is correct. We know that Greedy outputs a set of activities. 1. Are they non-overlapping? 2. Is this set the largest possible? It's easy to see that 1 is true. For 2, the crux of the proof is the following:

- **Claim:** Let A_1 be the activity with earliest finish time. Then there exists an optimal solution that includes A_1.

  Proof: Suppose O is an optimal solution (a non-overlapping subset of A of max size).

  - If $A_1 \in O$, we are done
  - If $A_1 \notin O$:

        * Let first activity in O be A_k

        * Make new solution by removing A_k and using A_1 instead: $O - \{A_k\} + \{A_1\}$

        * This is valid solution (because $f_1 < f_k$), and it has maximal size ($|O| - 1 + 1 = |O|$

- So this tells us that the first greedy choice is correct. What about the second one?

- Once the first greedy choice is made, we find an optimal solution for the remaining problem (the activities that don't conflict with A_1); this is the optimal substructure of the problem. The second activity chosen by the greedy solution is the activity that finishes first among all the activities that don't conflict with A_1, so applying the claim again to the remaining problem we know there must exist an optimal solution that includes this second greedy choice.

  And so on, it follows that at every step, the greedy choice stays ahead, and there exists an optimal solution that consists entirely of greedy choices.

Implementing the greedy idea

- The greedy idea above can be implemented in quadratic time: Sorting takes $O(n \lg n)$ time; step 2 takes $O(1)$ time; step 3 can be implemented in $O(n)$ time, and in step 4 we may recurse on $n - 1$ activities, so overall this can be implemented in $O(n^2)$ time.

- However, if we are a little more careful, once the activities are sorted, we can implement the greedy idea in $O(n)$ time. The crucial observation is that the solution consists of greedy choices that are compatible with previous greedy choices; so we pick the first activity A_1, then we traverse $A_2, A_3, ..$ in order, and discard all those that conflict with A_1; we pick the first activity that does not conflict with A_1. And so on. This leads to the following algorithm:

  - Sort A by finish time.
  - Schedule the first activity
  - Then go through remaining activities in order and schedule the first activity that starts after the previously scheduled activity finishes.
  - Repeat.

GreedyActivitySelection (A[1..n])

    Sort A by finish time

    $S = \{A_1\}$

    $j = 1$

    FOR $i = 2$ to n DO

        IF $s_i \geq f_j$ THEN

            $S = S \cup \{A_i\}$

            $j = i$

Analysis: Running time is $O(n \lg n) + O(n) = O(n \lg n)$.

Example: Trace the algorithm on the following 11 activities (already sorted by finish time):

$$(1, 4), (3, 5), (0, 6), (5, 7), (3, 8), (5, 9), (6, 10), (8, 11), (8, 12), (2, 13), (12, 14)$$

Comments, etc

- A complete greedy correctness proof has three parts:

  1. Prove that there exists an optimal solution which contains the first greedy choice.
  2. Prove optimal sub-structure.
  3. Induction on the solution size to prove that there exists an optimal solution that consists entirely of greedy choices.

  Part 2 and 3 are usually omitted (because it's a fairly similar proof for all problems). To prove that a greedy algorithm is correct it suffices to prove part 1, namely that there exists an optimal solution which contains the first greedy choice.

- A greedy algorithm chooses what looks like the best solution at any given moment; its choice is "local" and does not depend on solution to subproblems. (Greediness is shortsightedness: Always go for seemingly next best thing, optimizing the present without regard for the future, and never change past choices).

- In theory, we are only interested in greedy algorithms that are provably optimal. Most optimization problems cannot be solved greedily. Or put differently, the number of problems that have greedy solutions is very small (so chances are that if you came up with a greedy algorithms its probably wrong).

- The greedy technique can be applied to pretty much any optimization problem and is very popular in AI. There one deals with exponential problems or infinite search spaces, and one cannot solve these problems optimally, so the expectations are different. To emphasize that they are not optimal they are usually referred to as greedy *heuristics*.

- It is hard to figure out when being greedy gives the optimal solution! Problems that look very similar may have very different solutions.

  Example:

  - 0 − 1 KNAPSACK PROBLEM: Given n items, with item i being worth \$ v_i and having weight w_i pounds, fill knapsack of capacity w pounds with maximal value. Greedy does not work.

  - FRACTIONAL KNAPSACK PROBLEM it's the same, but we can take fractions of items. Can be solved greedily:
    * Compute value per pound $\frac{v_i}{w_i}$ for each item
    * Sort items by value per pound.
    * Fill knapsack greedily (take objects in order)
    * Runs in $O(n \log n)$ time
    * Can be shown that solution is optimal.