

Dynamic Programming: Edit Distance¹

In this lecture, we will look at another class of poster-child problems for dynamic programming: “string” problems. The input to these problems will be strings of the form $s[1 : m]$ where each $s[i]$ will be from some alphabet Σ . This alphabet could be $\{0, 1\}$, the alphabet of ASCII symbols, or $\{A, C, G, T\}$, depending on the applications.

How far is the string algorithm from the string logarithm? How far is apple from banana? How far is the cat-genome from the mouse-genome? Does it make sense to ask these questions? Well, the third question may point one to the importance of such a question. But how do we make the question well-defined so that we can talk about it? The *edit distance* or the *Levenshtein distance* is one way to capture this.

Given two strings $s[1 : m]$ and $t[1 : n]$, the edit distance captures the notion of distance between s and t defined using 3 operations. The first is the *insert* operation, $\text{ins}(s, i, c)$, which inserts character c after $s[i]$, thus making s longer; $\text{del}(s, j)$ deletes $s[j]$ from s making it shorter; and $\text{sub}(s, i, c)$ replaces $s[i]$ with the character c keeping the length the same. Each operation is assumed to cost 1 unit. The edit distance between $s[1 : m]$ and $t[1 : n]$ is the minimum cost sequence of operations of the kind ins , sub , del which can be required to convert s into t . This cost is called the *edit distance* between s and t , and is denoted as $\text{ED}(s, t)$.

For example, if s is `apple` and t is `banana`, then $\text{ED}(s, t) \leq 5$ since one can go from `apple` \rightarrow `bapple` \rightarrow `banple` \rightarrow `banale` \rightarrow `banane` \rightarrow `banana`. The operations are $\text{ins}(s, 0, b)$, $\text{sub}(s, 3, n)$, $\text{sub}(s, 4, a)$, $\text{sub}(s, 5, n)$, and $\text{sub}(s, 6, a)$. Turns out, this is the best one can do for these two strings. Finding the edit distance is the problem we will look at today.

EDIT DISTANCE

Input: Two strings $s[1 : m]$ and $t[1 : n]$.

Output: Return $\text{ED}(s, t)$.

Size: m, n .

At first glance, the problem seems to be rather complicated. For instance, suppose we had to find the edit distance between `apple` and `rallies`. One sees that the sequence of letters (a, l, e) appears in both strings in that order, although not contiguously. It seems like a good idea (not saying the best idea) to not touch them, and then move using insertions and substitutions as `apple` \rightarrow `rapple` \rightarrow `raplle` \rightarrow `raplles` \rightarrow `ralles` \rightarrow `rallies`.

However, the number of such subsequences can be pretty large. Indeed, it corresponds to a *subset* of locations on the strings, and the number of subsets (as in the subset-sum and knapsack) is *exponentially large* (it's 2^n). Going over all of them is a bad idea. Again, dynamic programming will come to our rescue.

As in SUBSET SUM and KNAPSACK, we imagine the “best” solution which takes us from s to t . Note that a solution for an EDIT DISTANCE instance, is a sequence π of operations where each entry is an ins , del , or sub . The cost of the solution, $\text{ED}(s, t)$ is simply the length $|\pi|$ of this sequence. Let π^* be the best sequence (which we don't know), but we want to argue about its structure. In particular, we want to argue this solution contains solutionettes.

To do so, let us focus on the *last* entries of both strings s and t . For example, if $s = \text{apple}$ and $t = \text{rallies}$. The last entries are $s[m] = \text{e}$ and $t[n] = \text{s}$. After we perform the operations in π^* , the character $s[m]$ at the end must become the character $t[n]$ at the end. There are *four* ways this can occur.

¹Lecture notes by Deeparnab Chakrabarty. Last modified : 19th Mar, 2022
These have not gone through scrutiny and may contain errors. If you find any, or have any other comments, please email me at deeparnab@dartmouth.edu. Highly appreciated!

- One, $t[n]$ was introduced using a *ins* operation somewhere in π^* . This could occur, for instance, if $s = \text{pan}$ and $t = \text{any}$ where π^* is $\text{pan} \rightarrow \text{pany} \rightarrow \text{any}$. The character $t[n] = y$ is inserted in the first step of π^* .
- Two, maybe the character $t[n]$ was already present in the string s . In that case, we get there by using a *del* on $s[m]$ somewhere in π^* . This could occur, for instance, if $s = \text{ate}$ and $t = \text{cat}$ where π^* is $\text{ate} \rightarrow \text{cate} \rightarrow \text{cat}$. The character $t[n] = t$ was already present, and the character $s[m] = e$ was deleted in the second step of π^* .
- Three, maybe the $t[n]$ *substituted* the last entry $s[m]$. This could occur, for instance, if $s = \text{pan}$ and $t = \text{pit}$ where π^* is $\text{pan} \rightarrow \text{pat} \rightarrow \text{pit}$. The character $t[n] = t$ substitutes $s[m] = n$ in the first step of π^* .
- Four, maybe $s[m]$ and $t[n]$ are the *same* characters, in which case we can leave them alone. This could occur, for instance, if $s = \text{bag}$ and $t = \text{cog}$ where π^* is $\text{bag} \rightarrow \text{bog} \rightarrow \text{cog}$.

To make dynamic programming work, we need to see if the *remaining* steps of π^* give an *optimum solution* for a smaller sub-instance of EDIT DISTANCE. If so, then we will have discovered the recursive substructure. Indeed, in all the four cases described above, we do have this.

- In this case, the remaining operations of π^* takes $s[1 : n]$ to $t[1 : m - 1]$. In the example above, it would be $\text{pan} \rightarrow \text{an}$.
- In this case, the remaining operations of π^* takes $s[1 : n - 1]$ to $t[1 : m]$. In the example above, it would be $\text{at} \rightarrow \text{cat}$.
- In this case, the remaining operations of π^* takes $s[1 : n - 1]$ to $t[1 : m - 1]$. In the example above, it would be $\text{pa} \rightarrow \text{pi}$.
- In this case, not the remaining but the *whole* of π^* takes $s[1 : n - 1]$ to $t[1 : m - 1]$. In the example above, it would be $\text{ba} \rightarrow \text{co}$.

Therefore, from the instance $I = (s[1 : m], t[1 : n])$, we get three smaller subinstances $I_1 = (s[1 : m - 1], t[1 : n - 1])$, $I_2 = (s[1 : m], t[1 : n - 1])$, and $I_3 = (s[1 : m - 1], t[1 : n])$, and it seems that (a) a solution of I leads to a solution of I_1 or I_2 or I_3 , and (b) solutions of I_1, I_2, I_3 lead to solutions of I (once again, formal proof comes later – first the idea). If one were to draw the recursion tree more, one would observe that a “typical subinstance” would look like $I' = (s[1 : i], t[1 : j])$. Thus, these are parameterized by $0 \leq i \leq m$ and $0 \leq j \leq n$, and therefore can be arranged in an $(m + 1) \times (n + 1)$ grid. The base case: when $i = 0$ or $j = 0$, that is when one of the strings is empty, in which case the edit distance is the length of the other string. We have all the ingredients for the dynamic programming solution which we now rigorously provide below in our usual six-step procedure.

- a. **Definition:** For any $0 \leq i \leq m$ and $0 \leq j \leq n$, let us use $\text{ED}(i, j)$ to be the edit distance between the strings $s[1 : i]$ and $t[1 : j]$. We are interested in $\text{ED}(m, n)$.

What should $\text{Cand}(i, j)$ be? Since the edit distance is the smallest number of “string operations” (ins/del/sub), let’s define $\text{Cand}(i, j)$ as the all possible sequences π of string operations which take the string $s[1 : i]$ to the string $t[1 : j]$. Armed with this notation, we get

$$\text{ED}(i, j) = \min_{\pi \in \text{Cand}(i, j)} |\pi|$$

- b. **Base Cases:**

$\text{ED}(0, j) = j$ for all $0 \leq j \leq n$ and $\text{ED}(i, 0) = i$ for all $0 \leq i \leq m$. There is only one way to go from an empty string to a string j – keep inserting. There is only one way to from a string of length i to an empty string – keep deleting.

- c. **Recursive Formulation:** Since we need to know whether the last characters of the strings are equal or not, let us introduce a piece of notation which will help us. Let $\mathbf{1}_{i,j}$ be the indicator variable of whether $s[i] = t[j]$ formally defined as

$$\mathbf{1}_{i,j} = \begin{cases} 1 & \text{if } s[i] = t[j] \\ 0 & \text{otherwise} \end{cases}$$

For all $i > 0, j > 0$:

$$\text{ED}(i, j) = \min(1 + \text{ED}(i - 1, j), 1 + \text{ED}(i, j - 1), (1 - \mathbf{1}_{i,j}) + \text{ED}(i - 1, j - 1))$$

- d. **Formal Proof:**

(\geq): Let π^* be the sequence of operations which took $s[1 : i]$ to $t[1 : j]$ and $|\pi^*| = \text{ED}(i, j)$. Note that, in π^* , either $s[i]$ is deleted from the end, or $t[j]$ is inserted at the end, and if neither of these two occur, we must either substitute $s[i]$ and $t[j]$, or these characters are the same. In the first case, consider the sequence of operations π which is π^* without the deletion. Observe, that π acting on $s[1 : i - 1]$ would take us to $t[1 : j]$. Thus, $|\pi^*| = 1 + |\pi| \geq 1 + \text{ED}(i - 1, j)$. Similarly, in the second case, consider the sequence π which is π^* without the insertion. π takes us from $s[1 : i]$ to $t[1 : j - 1]$, and thus, in this case, $|\pi^*| \geq 1 + \text{ED}(i, j - 1)$. Finally, if neither of the above two occur, then either $s[i] = t[j]$ in which case π^* actually takes $s[1 : i - 1]$ to $t[1 : j - 1]$. That is, $\pi^* \geq \text{ED}(i - 1, j - 1) = (1 - \mathbf{1}_{i,j}) + \text{ED}(i - 1, j - 1)$ since $s[i] = t[j]$. Or, $s[i] \neq t[j]$, and there is a substitution. And in this case, π defined as π^* minus that substitution takes $s[1 : i - 1]$ to $t[1 : j - 1]$. Again giving, $\pi^* \geq 1 + \text{ED}(i - 1, j - 1) = (1 - \mathbf{1}_{i,j}) + \text{ED}(i - 1, j - 1)$ in this case. In sum, in all of the possible cases, $\text{ED}(i, j) = |\pi^*|$ is larger than one of the things in the RHS parenthesis.

(\leq): Let π be the sequence of operations in $\text{Cand}(i - 1, j)$ of length $\text{ED}(i - 1, j)$. Consider the sequence of operations $\pi' = \text{del}(s, i) \circ \pi$, which first *deletes* the last entry of $s[1 : i]$ to get $s[1 : i - 1]$, and then follows the sequence of operations in π to get to $s[1 : j]$. Then, π' takes $s[1 : i]$ to $t[1 : j]$ and thus, $\pi' \subseteq \text{Cand}(i, j)$. Therefore, $|\pi'| \geq \text{ED}(i, j)$. Since $|\pi'| = 1 + |\pi| = 1 + \text{ED}(i - 1, j)$, we get that $\text{ED}(i, j) \leq 1 + \text{ED}(i - 1, j)$. Similarly, one can show

$ED(i, j) \leq 1 + ED(i, j - 1)$; the only difference is that we would $\text{ins}(t[1 : j - 1], t[j], j)$ at the end of doing π .

Finally, suppose π was a sequence of operations that took $s[1 : i - 1]$ to $t[1 : j - 1]$ and whose length was $ED(i - 1, j - 1)$. If $s[i] = t[j]$, and thus $\mathbf{1}_{i,j} = 1$, then π also takes $s[1 : i]$ to $t[1 : j]$. And so, $|\pi| \geq ED(i, j)$ implying $ED(i, j) \leq (1 - \mathbf{1}_{i,j}) + ED(i - 1, j - 1)$. If $s[i] \neq t[j]$, and thus $\mathbf{1}_{i,j} = 0$, then consider the sequence $\pi' = \text{sub}(s, t[j], i) \circ \pi$. π' takes $s[1 : i]$ to $t[1 : j]$ and thus $|\pi'| \geq ED(i, j)$. Since $|\pi'| = 1 + |\pi| = (1 - \mathbf{1}_{i,j}) + ED(i - 1, j - 1)$, we get $ED(i, j) \leq (1 - \mathbf{1}_{i,j}) + ED(i - 1, j - 1)$.

e. **Pseudocode for computing** $ED(m, n)$.

```

1: procedure ED( $s[1 : m], t[1 : n]$ ):
2:   ▷ Returns the edit distance between  $s$  and  $t$ .
3:   Allocate space  $E[0 : m, 0 : n]$  ▷  $E[i, j]$  will contain the edit distance between  $s[1 : i]$  and
    $t[1 : j]$ .
4:    $E[0, j] \leftarrow j$  for all  $0 \leq j \leq n$  and  $E[i, 0] \leftarrow i$  for all  $0 \leq i \leq m$ . ▷ Base Cases.
5:   for  $1 \leq i \leq m$  do:
6:     for  $1 \leq j \leq n$  do:
7:        $E[i, j] \leftarrow \min( E[i - 1, j] + 1, E[i, j - 1] + 1, E[i - 1, j - 1] + (1 - \mathbf{1}_{i,j}) )$ 
8:     return  $E[m, n]$ .
9:   ▷  $E[m, n]$  now contains the value of the edit distance
10:  ▷ Below we show the “recovery” pseudocode which shows one way to get from  $s$  to  $t$ 
   in  $E[m, n]$  moves.

11:   $i \leftarrow m; j \leftarrow n; \pi = []$ 
12:  ▷ Invariant:  $|\pi| + E[i, j] = E[m, n]$ 
13:  while  $i > 0$  and  $j > 0$  do:
14:    if  $E[i, j] = E[i - 1, j - 1] + (1 - \mathbf{1}_{i,j})$  then:
15:      if  $\mathbf{1}_{i,j} \neq 1$  that if  $s[i] \neq t[j]$  then: ▷ Substitution
16:        Append  $\text{sub}(s, i, t[j])$  to  $\pi$ 
17:         $i \leftarrow i - 1; j \leftarrow j - 1$ 
18:      else if  $E[i, j] = 1 + E[i - 1, j]$  then:
19:        ▷ We must have deleted  $s[i]$ 
20:        Append  $\text{del}(s, i, s[i])$  to  $\pi$ 
21:         $i \leftarrow i - 1$ 
22:      else: ▷ We must have that  $E[i, j] = 1 + E[i, j - 1]$  and we must have inserted  $t[j]$  at the
   end
23:        Append  $\text{ins}(s, i, t[j])$  to  $\pi$ 
24:         $j \leftarrow j - 1$ 
25:    ▷ At this point either  $i = 0$  or  $j = 0$ . Depending on which, we need to do a few more
   operations.
26:    while  $i > 0$  do:
27:      Append  $\text{del}(s, i, s[i])$  to  $\pi$ 
28:       $i \leftarrow i - 1$ 
29:    while  $j > 0$  do:
30:      Append  $\text{ins}(s, i, t[j])$  to  $\pi$ 
31:       $j \leftarrow j - 1$ 
32:  return Reverse of  $\pi, E[m, n]$ 

```

f. **Running time and space** The above pseudocode take $O(mn)$ time and space.

Theorem 1. The EDIT DISTANCE between two strings can be found in $O(nm)$ time and space.