

I&C SCI 46 Fall 2022 Project 3: Lewis Carroll Distance

Due Wednesday October 26 7:30 AM. This project is eligible for late submissions.

Introduction

In 1879, Lewis Carroll proposed the following puzzle to the readers of *Vanity Fair*: transform one English word into another by going through a series of intermediate English words, where each word in the sequence differs from the next by only one substitution. To transform *head* to *tail*, one can use four intermediates, assuming we are using a standard English dictionary to determine words: *head* → *heal* → *teal* → *tell* → *tall* → *tail*. We refer to the smallest number of substitutions necessary to transform one word to another as the *Lewis Carroll* distance between the two words.

[Optional] Choosing a project partner

You *have the option* to work with a second person for this assignment. If you do so, I expect you to work via pair programming. That is, you **may not** split the assignment, such as by having one person implement the Wordset while the other person implements the function that does the substitutions, and the two are stitched together later. I reserve the right to ask one or both project partners about the implementation and adjust the score accordingly. Similarly, any academic dishonesty arising from a group will be treated as an offense by both partners.

To declare a partnership, **both partners** must fill out this form by **Saturday, October 22, 11:59 PM** Irvine time. Failure to fill out the form by one or both partners, or an incorrect filling out of the form, may be grounds for the partnership to be voided or for one or both partners to have a reduced or zeroed grade for the project. The form will ask for your UCINetID. For most of you, that is the prefix of your UCI email address. ***It is not your ID number. It is not the email address that includes @uci.edu. If you do not know what your UCI Net ID is, please find out before filling out the form.***

https://docs.google.com/forms/d/e/1FAIpQLScGcVtKAG_sy5q8pVLZ0RgRZsj9mHUcrvaggMO9ZMfpj7MXxA/viewform

Please refer to the “Partnerships” section of the class lab manual (https://www.ics.uci.edu/~mikes/ics46/Projects/ICS46_Lab_Manual_F22.pdf) for information on adding a partner to your Gitlab project. The project specified in the partnership form is the one that will be graded.

Requirements

- Fill in WordSet.cpp
 - You will need to fill in the function `polynomialHashFunction`, which takes three parameters.
 - This function will interpret the string parameter, which will contain only lower-case letters, as coefficients for a polynomial of degree equal to the length of the string. You will evaluate it at the point “base,” which can be thought of (if you prefer) as interpreting the string as a base-that integer. Treat ‘a’ as 1, ‘b’ as 2, and so on. It returns the smallest positive integer that is a representation of that string, in that base, mod the given modulus.
 - For example, `polynomialHashFunction("abz", 5, 10)` should return the result of $(1 * 5^2 + 2 * 5 + 26) \% 10$.
 - Be careful about when you take the modulus within the hash function.
 - Be careful about when you take the modulus outside the hash function.
 - Do not assume any parameters match the named constants in the file -- we will use those later.
 - The provided code DOES NOT accurately compute this function, but it does show you how to easily access numeric values for the letters.
 - Your implementation **must be** done via a Cuckoo hash table, as described in lecture. Note that we *are not* creating a Cuckoo filter.
 - The first table’s hash function is to be done with `base=BASE_H1`, with a similar rule for the second hash function for `BASE_H2`.
 - You must use a dynamically-allocated C-style array, not a `std::vector` or similar container, as the basis in your `WordSet`.
 - You should start your array at size *initialCapacity*.
 - Resize only when the element about to be inserted cannot be added to the existing structure. An element cannot be inserted when *evictionThreshold* evictions have been made during its attempted insertion. That value is an optional second parameter to the constructor.
 - When you resize and rehash, the next table size should be the *smallest* prime number that is no smaller than twice the current table size. For example, if your current table size is 11, your next one is 23. If your current table size were somehow 13, your next one would be 29.
 - Your implementation must fit the interface given.
 - Your implementation **does not need to be templated** -- nor should it be, for the purposes of this assignment.
 - In fact, doing so will cause an issue for some of our provided tests.
 - You **do** need to implement the destructor. Memory leaks will cause a grade penalty.
 - Do not hard code your table for the uses we’ll have in the next section.

- Write function `std::vector<std::string> convert(const std::string & s1, const std::string & s2, const WordSet & words)` in `convert.cpp`
 - This function will return the conversion between `s1` and `s2`, according to the lowest *Lewis Carroll* distance. The first element of the vector should be `s1`, the last `s2`, and each consecutive should be one letter apart. Each element should be a valid word. If there are two or more equally least Lewis Carroll distance ways to convert between the two words, you may return any of them.
 - If there is no path between `s1` and `s2`, return an empty vector.
 - It is recommended that you compute the distance via a breadth-first search. To visualize this, imagine a graph where the words are vertices and two vertices share an (undirected) edge if they are one letter apart.
 - If you do not know what this means, please ask -- Shindler would be happy to explain.
 - You *may* use `std::queue` -- you do not need to write your own
 - A good thing to do the first time you see a word in the previous part is to place it into a `std::unordered_map<string, string>`, where the key is the word you just saw and the value is the word *that you saw immediately before it*. This will allow you to later produce the path: you know the last word, and you know the prior word for each word in the path except the first. Furthermore, if the key isn't in that map, this tells you that you haven't seen it before.
 - Your implementation does not have to be the most efficient thing ever, but it cannot be “too slow.” In general, any test case that takes over a minute on the grader’s computer may be deemed a wrong answer, even if it will later return a correct one.

You **may not** use any classes that are part of the C++ standard template library in implementing your Wordset. You *may* use simple functions, such as `std::swap` or computing a logarithm if you so choose.

You may use `std::queue`, `std::map`, `std::stack`, and/or `std::set` in `convert.cpp`, but you *may not* use them in place of your WordSet: you must use that in the appropriate places. *A solution for convert.cpp that does not use WordSet to determine if something is a word will receive no credit.*