

# Assignment 1 – Getting Started

CS499 - Intro to NLP  
Antonios Anastasopoulos

January 2021

Due Date: 2/5/2021 (eod)  
Instructions: see at the end of the assignment.

## 1 Preliminaries (1 credit total)

[1 credit] Answer all four questions as they appear below.

Read in the Hamlet corpus. You can read it easily with

```
1 with open("hamlet.txt") as inp:  
2     lines = inp.readlines()
```

Use the NLTK package<sup>1</sup> (which you can install with `pip install nltk`) to split the dataset into sentences. Use the `sent_tokenize()` function. (a) How many sentences are there? (ignore empty lines and empty sentences)

Now, we want to split into words. First, split the sentences on whitespace (e.g. using the python `.split(' ')` function over each sentence. (b) How many words do you find?

However, this tokenization approach is not ideal. We are still treating words that are attached to punctuation as different tokens. It's better to do a more informed word tokenization. Use NLTK's `word_tokenize()` function, to split into words. Now the first sentence should have 13 tokens. (c) How many tokens are there now?

Last, lowercase everything, to avoid skewing your counts due to the (arbitrary) rule of uppercasing the first letter of the first word of each sentence.

Compute the number of types and tokens in the corpus. (d) How many types and how many tokens appear in Hamlet?

## 2 Herdan's Law(1 credit total)

[1 credit] According to Herdan's Law (or Heaps' Law), the relationship between the number of types  $|V|$  and number of tokens in  $N$  in a corpus should be

---

<sup>1</sup><https://www.nltk.org/>

described by the following equation:


$$|V| = kN^\beta.$$

For  $\beta = 0.7$  (a standard value typically found for English), what is the value for  $k$ , for your lowercased, tokenized Hamlet corpus?

### 3 Zipf's Law (2 credits total)

Zipf's law<sup>2</sup> states that given a large sample of text, the frequency of any word is inversely proportional to its rank in the frequency table. Zipf's law is most easily observed by plotting the data on a log-log graph, with the axes being the log rank order of a type and log frequency of a type. For example, “the” is the most common type in English, so it should appear at  $x = \log(1)$ .<sup>3</sup> If it appears 70000 times, then its y-axis value should be  $y = \log(70000)$ . The data conform to Zipf's law to the extent that the plot is linear.

For example, look at this chart (from Wikipedia: [https://en.wikipedia.org/wiki/Zipf%27s\\_law](https://en.wikipedia.org/wiki/Zipf%27s_law)) that plots the rank versus the frequency for the first 10 million words over 30 wikipedias in a log-log scale:



We can describe the law more generally, using a general formula that has parameters that will vary slightly across languages. Formally, let:

- $|V|$  be the number of types (hyperparameter);
- $s$  be the value of the exponent characterizing the distribution (hyperparameter);
- $k$  be the rank of a type.

<sup>2</sup>Named after the linguist George Kingsley Zipf, who first sought to explain it.

<sup>3</sup>Note that “the” is actually the third most common type in the lowercased, tokenized Hamlet corpus, because two other tokens have higher frequencies.

Zipf's law then predicts that out of a population of  $N$  elements, the normalized frequency of the element of rank  $k$ ,  $f(k; s, |V|)$ , is:

$$f(k; s, |V|) = \frac{1/k^s}{\sum_{n=1}^{|V|} (1/n^s)}.$$

[1 credit] Fill the following table based on the (lowercased, tokenized) Hamlet corpus. Use  $s = 1$ . Report (relative) frequencies up to 5 decimal digits. The observed relative frequency should be the type's frequency as observed in your corpus. The predicted one should be the prediction of the formula above.

Type	Rank	Observed Relative Frequency	Predicted Relative Frequency
the			
hamlet			
rose			
royally			
honourable			

[1 credit] Create a plot that will illustrate Zipf's Law over your dataset. We want a plot similar to the example above. You can use packages like matplotlib, seaborn, ggplot, etc. Remember to label your axes!

## 4 Edit Distance over Gazetteers (6 credits total)

"Gazetteer" basically means "index". In NLP, we typically refer to them to denote mappings of named entities<sup>4</sup> across languages. The assignment package provides a sample gazetteer with mappings of entities between English and Kinyarwanda. Kinyarwanda is one of the official languages of Rwanda, spoken by at least 12 million people in Rwanda, Congo, and Uganda.

Here's an example of the gazetteer, where each line follows the format [english] ||| [kinyarwanda]:

```
1 Iran ||| Irani
2 Eggplant ||| Urutoryi
3 Bible ||| Bibiliya
4 Flag ||| Ibendera
5 Rose ||| Iroza
```

Programmatically, you can obtain the contents of the file line by line as follows:

```
1 with open("gazetteer.eng-kin.txt") as inp:
2     lines = inp.readlines()
3
4 for l in lines:
```

<sup>4</sup>A named entity is a real-world object, such as persons, locations, organizations, products, etc., that can be denoted with a proper name.

```

5 l = l.strip().split(' ||| ')
6 eng = l[0].strip()
7 kin = l[1].strip()
8 ...

```

The goal of this assignment will be to study how similar the English and Kinyarwanda named entities are. To do this, we will rely on the Edit Distance algorithm, as described in Section 3.11 of the SLP book.

[1 **credit**] First, implement a function that computes the minimum edit distance of two strings. Pseudo-code can be found in Figure 3.25 of the book. Print the computation matrix for the pair in the first line of the dataset.

*Note: there exist Python packages that implement this function. I highly advise you to implement it from scratch, because you will need to modify it for the following questions. If you are unsure about whether your implementation is correct, you can certainly use such a package to confirm that your outputs are correct.*

[1 **credit**] What is the average edit distance over the whole dataset? What is the minimum edit distance? What is the maximum edit distance? How many pairs have the minimum, and how many have the maximum edit distance?

[1 **credit**] What is the edit distance for the Kinyarwanda phrases corresponding to the English “Ankara” (the capital of Turkey) and “Ankara Province”? How do you explain the second value?

[1 **credit**] Next, implement backtracing (also described in section 3.11 of the book) on top of your function. To do this, follow the instructions of the book:

In the first step, we augment the minimum edit distance algorithm to store backpointers in each cell. The backpointer from a cell points to the previous cell (or cells) that were extended from in entering the current cell. [...] In the second step, we perform a backtrace. In a backtrace, we start from the last cell (at the final row and column), and follow the pointers back through the dynamic programming matrix. Each complete path between the final cell and the initial cell is a minimum distance alignment.

For the first pair that has a minimum edit distance of 5 (which country is it?), print the edit distance matrix **and** the backtrace (the sequence of cells in the matrix – from top right to bottom left (as shown in Figure 3.27 in the book) – that yields this minimum cost.

[1 **credit**] Modify your code to also keep track of exactly which characters get *inserted* when translating from English to Kinyarwanda. What are the two most commonly inserted characters and how many times did this insertion occur in the dataset?

*Note: this question is specifically about insertions, not substitutions. Although you could consider a substitution as a sequence of two operations (one deletion, one insertion), here we do not take this approach.*

[1 credit] These two most commonly inserted characters often have very “short” pronunciation in Kinyarwanda. As a result, one could argue that they don’t really impact the understanding of these terms, so inserting them shouldn’t have such a high cost when computing the edit distance.

Modify your code so that when you compute the edit distance, the cost of inserting either of these two characters is 0.5 instead of 1. What is the average edit distance over the dataset with this modified scoring scheme?

## [BONUS] Mongo Morphology (3 credits total)

Mongo (also known as Nkundo or Lomongo) is a Bantu language spoken by the Mongo Peoples of the central Democratic Republic of the Congo. Presently, there are around 400,000 native speakers spread out over a large area around the Congo River. Below is a table showing a few verb conjugations in Lomongo.  $\text{d}\text{̥}\text{z}$  is a consonant pronounced like the **j** in the English word **jump**.  $\text{ŋ}$  is a consonant pronounced like the **ng** at the end of the English word **sing**.<sup>5</sup>

Imperative	2nd Singular	3rd Singular	3rd Plural	English
bota	oota	aota	baota	‘beget’
kamba	okamba	akamba	bakamba	‘work’
imed $\text{̥}\text{z}$ a	wimed $\text{̥}\text{z}$ a	imed $\text{̥}\text{z}$ a	bimed $\text{̥}\text{z}$ a	‘consent’
usa	wusa	usa	busa	‘throw’
bata	oata	aata	baata	‘get’
ɛna	wɛna	ɛna	bɛna	‘see’
isa	wisa	isa	bisa	‘hide’
d $\text{̥}\text{z}$ ila	od $\text{̥}\text{z}$ ila	ad $\text{̥}\text{z}$ ila	bad $\text{̥}\text{z}$ ila	‘wait’
ina	wina	ina	bina	‘hate’
bina	oina	aina	baina	‘dance’
asa	wasas	asa	basas	‘search’
sangga	osangga	asangga	basangga	‘say’

[2 credits] Explain how these Mongo verb forms work by filling in the blanks below:

<sup>5</sup>The symbols  $\text{d}\text{̥}\text{z}$  and  $\text{ŋ}$  are part of the International Phonetic Alphabet (IPA; [https://en.wikipedia.org/wiki/International\\_Phonetic\\_Alphabet](https://en.wikipedia.org/wiki/International_Phonetic_Alphabet)). This is the alphabet phoneticians use to describe the *sounds* of a language, and the symbols correspond to phonemes and phones. You can see the whole IPA and hear the corresponding sounds for the symbols in this interactive chart: <https://www.ipachart.com/>.

Each Mongo verb has a root form. The 4 verb forms shown here are formed by adding a prefix before the root form. The prefix for the imperative<sup>6</sup> form is \_\_\_\_\_, the prefix for the 2nd singular form is \_\_\_\_\_, the prefix for the 3rd singular form is \_\_\_\_\_, and the prefix for the 3rd plural form is \_\_\_\_\_.

(Note that some of these prefixes may be empty; to note this, write the symbol  $\emptyset$ —).

However, we are not done yet: to get the final verb form, we must apply some sound change rules. The relevant rules are:

1. If there are two \_\_\_\_\_ in a row, delete \_\_\_\_\_.
2. Delete \_\_\_\_\_ when it appears between two \_\_\_\_\_.
3. Change \_\_\_\_\_ to \_\_\_\_\_ when it appears before a \_\_\_\_\_.

There is one final wrinkle: the order that these rules are applied in matters. The rules must be applied in this order: First apply rule \_\_\_\_\_, then rule \_\_\_\_\_, then rule \_\_\_\_\_.

[1 credit] Fill in the blanks in the table below.

Imperative	2nd Singular	3rd Singular	3rd Plural	English
bakisa	wanda	aponama	basolola	‘add’
solola				‘begin’
ponama	‘chat’			
	oowa	aalusa		‘elect’
				‘cure’
longga				‘turn’
				‘win’

## Submission Instructions

This assignment must be done individually. Discussing the assignment on the class forum is absolutely ok – posting the solutions is not. If you use someone else’s code (e.g. taken from a book or a website), make sure to properly cite it in your report.

Please submit all of the following in a gzipped tar archive (.tgz; not .zip or .rar) via Blackboard. If you’re making a full submission, please name your file **gmuid-hw1.tgz** (for example, the instructor’s submission would be **antonis-hw1.tgz**). If you’re making a partial submission, please name your file **gmuid-hw1-part.tgz** where part is the part (1, 2, 3, or 4) that you’re submitting. Note that submitting two files with the same name will overwrite one of them!

<sup>6</sup>1 The imperative form of a verb is used to give a command, warning, advice, or request.

Your submission should contain:

- A PDF file (not .doc or .docx) with your responses to the instructions/questions above.
- All of the code that you wrote.
- A brief README file with instructions on how to build/run your code.