

Lab 8 – Wikipedia Search

CC5212-1

April 27, 2016

Today we're going to build a search engine for Wikipedia.

First we're going to index the keywords in the title and abstract of all the articles in Spanish Wikipedia using Apache Lucene (an open source package for doing inverted indexing of text in documents). Then we'll implement the search function: you type in a keyword, you get back the articles that match.

We'll be working on the local computer with Eclipse.

- Download the code project from <http://aidanhogan.com/teaching/cc5212-1-2016/lab/08/mdp-lab08.zip> and open it in Eclipse.
- Start downloading the data from <http://aidanhogan.com/teaching/data/wiki/es/es-wiki-articles.tsv.gz> and <http://aidanhogan.com/teaching/data/wiki/es/es-wiki-articles-1k.tsv>. The first file contains 964,843 lines, with an article encoded in each line (*there is no need to unzip this file; we can read a GZIP input stream from the Java program later*). The second file contains the first 1,000 lines, so you can open it to see the format. Each line has the article URL, the title, and the abstract (delimited by tab).
- While that's downloading, open up the code project again. The first Java class we'll work with will be `BuildWikiIndex`, which will index the large data file into the Lucene inverted index. The second class we'll work with will be `SearchWikiIndex`, which will take your keyword searches and use the index to find relevant Wikipedia articles. To help you in this, there were some Lucene code examples in the slides from last Monday's lecture. The examples are a bit different so you cannot copy them symbol for symbol, but the main steps are exemplified there.
- So first open up `BuildWikiIndex`. The main method is already written for you. From the command line, it takes an input file (which will be `es-wiki-abstracts.tsv.gz`) and an output directory (an empty existing directory) to write the inverted index to. Jump to the `indexTitleAndAbstract(...)` method. You'll see some code structure given for you.
 - Open the directory at `indexDir`, set the analyser for `Version.LUCENE_48`, and open a new `IndexWriter` configured for writing. Instead of a `StandardAnalyzer` which is generally for English, let's use a `SpanishAnalyzer` (it takes the same arguments).
 - For each line in the input, if the line is not empty (or just whitespace), split the line by tab. The first split contains the article URL, the second the title, the third the abstract. If the array does not contain at least a URL and a title, skip it.
 - * Create a new Lucene document for every valid line.
 - * For the following field names, use the `FieldNames` enum, for example, `FieldNames.URL.name()`.
 - * Index the URL as a string-field with store set to yes.
 - * Index the title as a text-field with store set to yes.
 - * If an abstract is present and non-empty (some articles don't have one), index the abstract as a text-field with store set to yes.
 - * Index the time the document was indexed (current time) as a long field with store set to no.
 - * Add the document to the writer.

- Don't forget to close the writer at the end.
 - Once finished, run the class with `-i [INPUT-FILE] -igz -o [INDEX-DIR]`, where `[INDEX-DIR]` is a fresh output directory. You can run it over the big file (`-igz` indicates that the input is GZipped).
 - Assuming all runs well, congrats, you've made a search engine index! But now to see that it works.
- Open up the `SearchWikiIndex` class.
 - Again the `main` method is already done. Jump to `startSearchApp(.)`. This will take the directory of the index you just wrote and accepts keyword searches from the command line over that index. Unlike the example in the lecture, we want to read the searches from the input prompt. Again there's some structure and hints given.
 - * First open an index reader for the directory name passed as input.
 - * Open a searcher over the index. Make sure to use the same (type of) analyser as before!
 - * We want to search over both title AND abstract so we want to create a multi-field query parser. We also want title matches to score more highly: you can pass the `BOOSTS` object at the top of the class for this. Again, make sure to use the same analyser as before!
 - * Next we need to read lines from the input. There's some code left there for you. Inside the `try` you need to ...
 - parse the query (also print the raw query and the parsed query to the console),
 - get the hits from the searcher (get `DOCS_PER_PAGE` number of hits),
 - for each result, print the details of the document (the url, title, etc.; you can print anything that had `store` set to `true` during indexing).
 - * Once you're finished, time to test it! Run the class with the argument `-i [INDEX-DIR]`. Try a search like "obama". Do the results make sense? Play around with some other searches.
 - * If you are having trouble with accents, make sure to set the encoding in Eclipse to `utf-8`. Go to `Run Configurations > Common` and select `Other [UTF-8]` under encoding.
 - * OPTIONAL: Try vary the boost factors between title and abstract and see how it affects the results of a few queries. How do the results change?
 - Submit your two classes (`BuildWikiIndex` and `SearchWikiIndex`) to u-cursos by next Monday.