

Exercise Sheet 2

Submit until Tuesday, **November 8** at **2:00pm**

Exercise 1 (5 points)

Copy your code from *sheet-01* to a new subfolder *sheet-02*. Extend your code to incorporate BM25 scores, as explained in the lecture. This entails the following:

1. Add BM25 scores to the inverted lists. Pay attention to the implementation advice given in the lecture, and avoid unnecessary complexity.
2. Change your routine for intersecting two lists in a routine for merging two lists. Note that this can be done with a relatively minor change.
3. Sort the results by the aggregated BM25 scores and output the top-3 results. You don't have to implement the sorting algorithm yourself; you can use one of the built-in sorting functions.

There is a TIP file on the Wiki with a suggestion for the structure of your code. The TIP file also contains test cases for some of the functions.

You must implement these test cases, also for the following exercises and whenever they are specified for future exercise sheets. Otherwise your submission will not be graded.

It's for your own good and for the good of the tutors, as explained in the lecture. You are free to extend these test cases or add more tests, but you may not restrict or otherwise weaken them. After all, their purpose is a basic correctness check of your code.

Try to keep your code as simple and as short as possible. Otherwise, adhere to the standards described on the back of Exercise Sheet 1. All this goes without saying from now on.

Exercise 2 (5 points)

Find a good setting for the BM25 parameters by inspecting the results for a variety of queries of your choice. Optionally (= you don't have to do this to get full points), feel free to improve your ranking in any way you see fit; we discussed various possibilities in the lecture (slide 18).

This development phase should be completed **before** you proceed with Exercise 3. Briefly(!) describe your insight from this phase in your *experiences.txt*.

[please turn over without checkstyle errors]

Exercise 3 (10 points)

Evaluate your system on the benchmark provided on the Wiki. The benchmark provides 10 queries, and for each query the ids (line numbers in *movies.txt*) of the relevant documents. Evaluate the measures P@3, P@R, and MAP. Report your results in the table on the Wiki, following the examples already given there. In the last column, provide your BM25 parameter settings + a very brief (not complete) description of any additional feature you might have added.

Once you start testing your system on this benchmark you should not go back anymore to Exercise 2 and change your ranking method or tune your parameters. That would be overfitting. In a real competition, your system would be evaluated by the organizers after the submission deadline, on an undisclosed benchmark.

Commit your *sheet-02* to our SVN. As usual, in your *experiences.txt*, provide a brief account of your experience with this sheet and the corresponding lecture. As a minimum, say how much time you invested and if you had major problems, and if yes, where.