

FUB at TREC 2009 Relevance Feedback Track: Diversifying Feedback Documents

(Extended Abstract)

Andrea Bernardini, Claudio Carpineto, Edgardo Ambrosi

Fondazione Ugo Bordoni

{a.berna, carpinet, eambrosi}@fub.it

The focus of our participation was optimal selection and use of diverse feedback documents. Assuming that the query has a topical structure and that the user is interested only in some query topics and assuming also that only a small amount of feedback information will be made available, the goal was to select topic representatives to be used as feedback documents and then exploit the feedback information to bias the set of results towards the topics of interest.

Our method consists of the following steps. The selection of topic representatives was performed by running a first-pass retrieval on the original query and extracting the most novel and relevant documents from the obtained results. Such documents were returned for manual evaluation and the provided feedback was given as an input to an improved version of the relevance feedback system that we developed at TREC 2008 Relevance Feedback track. As this system was designed for explicitly dealing with both positive and negative relevance feedback, we reckoned that in principle it would be able to use the diverse topic representatives to filter out the highly-ranked documents belonging to the unwanted topics in an effective manner.

Unfortunately, due to resource and time limitations, we were not able to fully implement and test the method outlined above. We had to make a number of approximations and simplifications that did not allow us to evaluate its true potential for improving relevance feedback. Likewise, we could not test our main hypothesis that this method is suitable for dealing with multi-topic queries.