

Round 1 results — Run irc_pmc submitted from IRC


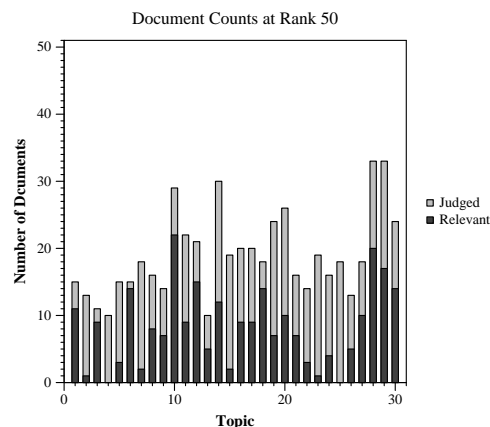

Run Description

As part of TREC-COVID, we submit automatic runs based on (pseudo) relevance feedback in combination with a reranking approach. The reranker is trained on relevance feedback data that is retrieved from PubMed/PubMed Central (PMC). The training data is retrieved with queries using the contents of the ‘<query>’ tags only. For each topic a new reranker is trained. We consider those documents retrieved by the specific topic query as relevant training data, and the documents of the other 29 topics as non-relevant training data. Given a baseline run, the trained system reranks documents. The baseline run is retrieved with the default ranker of Elasticsearch/Lucene (BM25) and queries using the contents of the ‘<query>’ tags only. For our reranker we use GloVe embeddings in combination with the Deep Relevance Matching Model (DRMM). Our three run submissions differ by the way training data is retrieved from PubMed/PMC. ‘irc_pmc’: This run is trained on full text documents retrieved from PMC.


Summary Statistics	
Run ID	irc_pmc
Topic type	automatic
Contributed to judgment sets?	no

Overall measures	
Number of topics	30
Total number retrieved	30000
Total relevant	2352
Total relevant retrieved	1162
MAP	0.0965
Mean Bpref	0.3023
Mean NDCG@10	0.2233

Document Level Averages	
	Precision
At 5 docs	0.3400
At 10 docs	0.2833
At 15 docs	0.2711
At 20 docs	0.2417
At 30 docs	0.2122
R-Precision	
Exact	0.1513




Per-topic difference from median bpref for all Round 1 runs



Per-topic difference from median NDCG@10 for all Round 1 runs

Round 1 results — Run irc_pmc submitted from IRC



Per-topic difference from median P@5 for all Round 1 runs