

Round 2 results — Run irc_logreg_tfidf submitted from IRC




Run Description

This run submission combines a BM25 baseline with a logistic regression based reranker trained on tfidf-features in combination with relevance judgments of the first round. The baseline run is retrieved with the default ranker of Elasticsearch/Lucene (BM25) and queries using the contents of the <query>, <question>, and <narrative> tags. Documents are reranked for those topics where relevance judgments are available (1-30), otherwise the baseline ranking remains unaltered (31-35).


Summary Statistics	
Run ID	irc_logreg_tfidf
Topic type	feedback
Contributed to judgment sets?	no

Overall measures	
Number of topics	35
Total number retrieved	35000
Total relevant	3002
Total relevant retrieved	1604
MAP	0.1666
Mean Bpref	0.4025
Mean NDCG@10	0.4170
Mean RBP(p=0.5)	0.4841 +0.2341

Document Level Averages	
	Precision
At 5 docs	0.5371
At 10 docs	0.4543
At 15 docs	0.4133
At 20 docs	0.3829
At 30 docs	0.3381
R-Precision	
Exact	0.2382




Per-topic difference from median bpref for all Round 2 runs




Per-topic difference from median NDCG@10 for all Round 2 runs

Round 2 results — Run irc_logreg_tfidf submitted from IRC



Per-topic difference from median P@5 for all Round 2 runs



Per-topic difference from median RBP(p=0.5) for all Round 2 runs