

Run Description	
Run ID:	ielab-TILDEv2
Task:	Passage Ranking
Subtask:	Full ranking from the collection
Topic type:	Automatic
Single-stage retrieval?	No
Dense retrieval?	No
Used deep nn model?	Pre-trained model
Type of training:	Previous year's MS MARCO training data
Pre-processing/indexing cost:	requires expanding the original passage collection (append additional tokens to each passage); this process took about 1 hour 40 mins per 2 million passages on a single Tesla v100 16G Gpu, we use 6 gpus for this part, thus it took around 20 hours in total. After expanding the whole collection, TILDEv2 required encoding the expanded collection which costs around 40 mins per 2 million passages, we again use 6 gpus to do this. Finally, we used Anserini to indexing the encoded collection which took us around 3.5 hours on a 2018 Mac Mini machine with 3.2 GHz 6-Core Intel Core i7 CPU.
Query processing cost:	it uses a BERT tokenizer to process each query which takes only 0.1 ms per test query on cpu. This is followed by searching in our custom index (implemented using Hashtable) to compute the final scores and rerank the 1,000 passages retrieved by BM25, which takes around 11 ms per query on CPU.

Overall measures		Document Level Averages	
Number of topics	53		Precision
Total number retrieved	5300	At 5 docs	0.5623
Total relevant	3427	At 10 docs	0.4906
Total relevant retrieved	1123	At 15 docs	0.4503
MAP	0.2112	At 20 docs	0.4170
Mean NDCG@10	0.5825	At 30 docs	0.3667
Mean Reciprocal Rank	0.6926	R-Precision	
		Exact	0.2875



