



Deep Learning Track results — (TU_Vienna) TU Wien


Run Description	
Run ID:	TUW_TAS-B_ANN
Task:	Passage Ranking
Subtask:	Full ranking from the collection
Topic type:	Automatic
Single-stage retrieval?	Yes
Dense retrieval?	Yes
Used deep nn model?	Pre-trained model
Type of training:	Previous year's MS MARCO training data
Pre-processing/indexing cost:	Encoding: 400 GPU-minutes on a single TITAN RTX GPU (Using a single newer generation A40 GPU we observe 320 GPU-Minutes); Indexing for an HNSW index with 96 neighbors per vector: 500*16 cores = 8,000-CPU Minutes; The HNSW index storage requirement is 162GB (needed in CPU-memory)
Query processing cost:	On average: 1.7 GPU-ms to encode a single query + 216 CPU-ms for HNSW lookup

Overall measures		Document Level Averages	
Number of topics	53		Precision
Total number retrieved	5300	At 5 docs	0.4906
Total relevant	3427	At 10 docs	0.4566
Total relevant retrieved	1158	At 15 docs	0.4327
MAP	0.1932	At 20 docs	0.4009
Mean NDCG@10	0.5426	At 30 docs	0.3585
Mean Reciprocal Rank	0.7015	R-Precision	
		Exact	0.2625







Per-topic difference from median for NDCG@10 for Passage Ranking runs



Per-topic difference from median for Prec@10 for Passage Ranking runs



Per-topic difference from median for Average Precision for Passage Ranking runs



Per-topic difference from median for Reciprocal Rank for Passage Ranking runs