



Deep Learning Track results — (CIP) Chinese Information Processing Lab


Run Description	
Run ID:	cip_fl_r
Task:	Passage Ranking
Subtask:	Full ranking from the collection
Topic type:	Automatic
Single-stage retrieval?	No
Dense retrieval?	Yes
Baseline run?	No
Used deep nn model?	Pre-trained model
Type of training:	This year's MS MARCO training data; Previous year's MS MARCO training data
Training cost:	Cost about 393 minutes for training on MS MARCO v1 on a single GeForce RTX (24G) 3090 GPU, 364 minutes for training on MS MARCO v2 on a single NVIDIA TITAN RTX (24G) GPU and 1005 minutes for re-ranking model training on six NVIDIA TITAN RTX (24G) GPU. The total GPU minutes is 6,787minutes.
Pre-processing/indexing cost:	Cost 507 minutes for inference on six NVIDIA TITAN RTX (24G) GPUs, 66 minutes for dense index, 57 minutes for docT5query index. The total GPU minutes is 3042 minutes.
Query processing cost:	Cost 84 minutes for dense retrieval, 189 minutes for docT5query ensemble and 1 minutes and 49 seconds for re-ranking on six NVIDIA TITAN RTX (24G) GPUs. The total GPU minutes is 11 minutes.

Overall measures		Document Level Averages	
Number of topics	76		Precision
Total number retrieved	6792	At 5 docs	0.4500
Total relevant	4613	At 10 docs	0.4461
Total relevant retrieved	1167	At 15 docs	0.4026
MAP	0.1599	At 20 docs	0.3757
Mean NDCG@10	0.5007	At 30 docs	0.3237
Mean Reciprocal Rank	0.6428	R-Precision	
		Exact	0.2425







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