

MPII at the NTCIR-14 WWW-2 Task

Andrew Yates

Max Planck Institute for Informatics

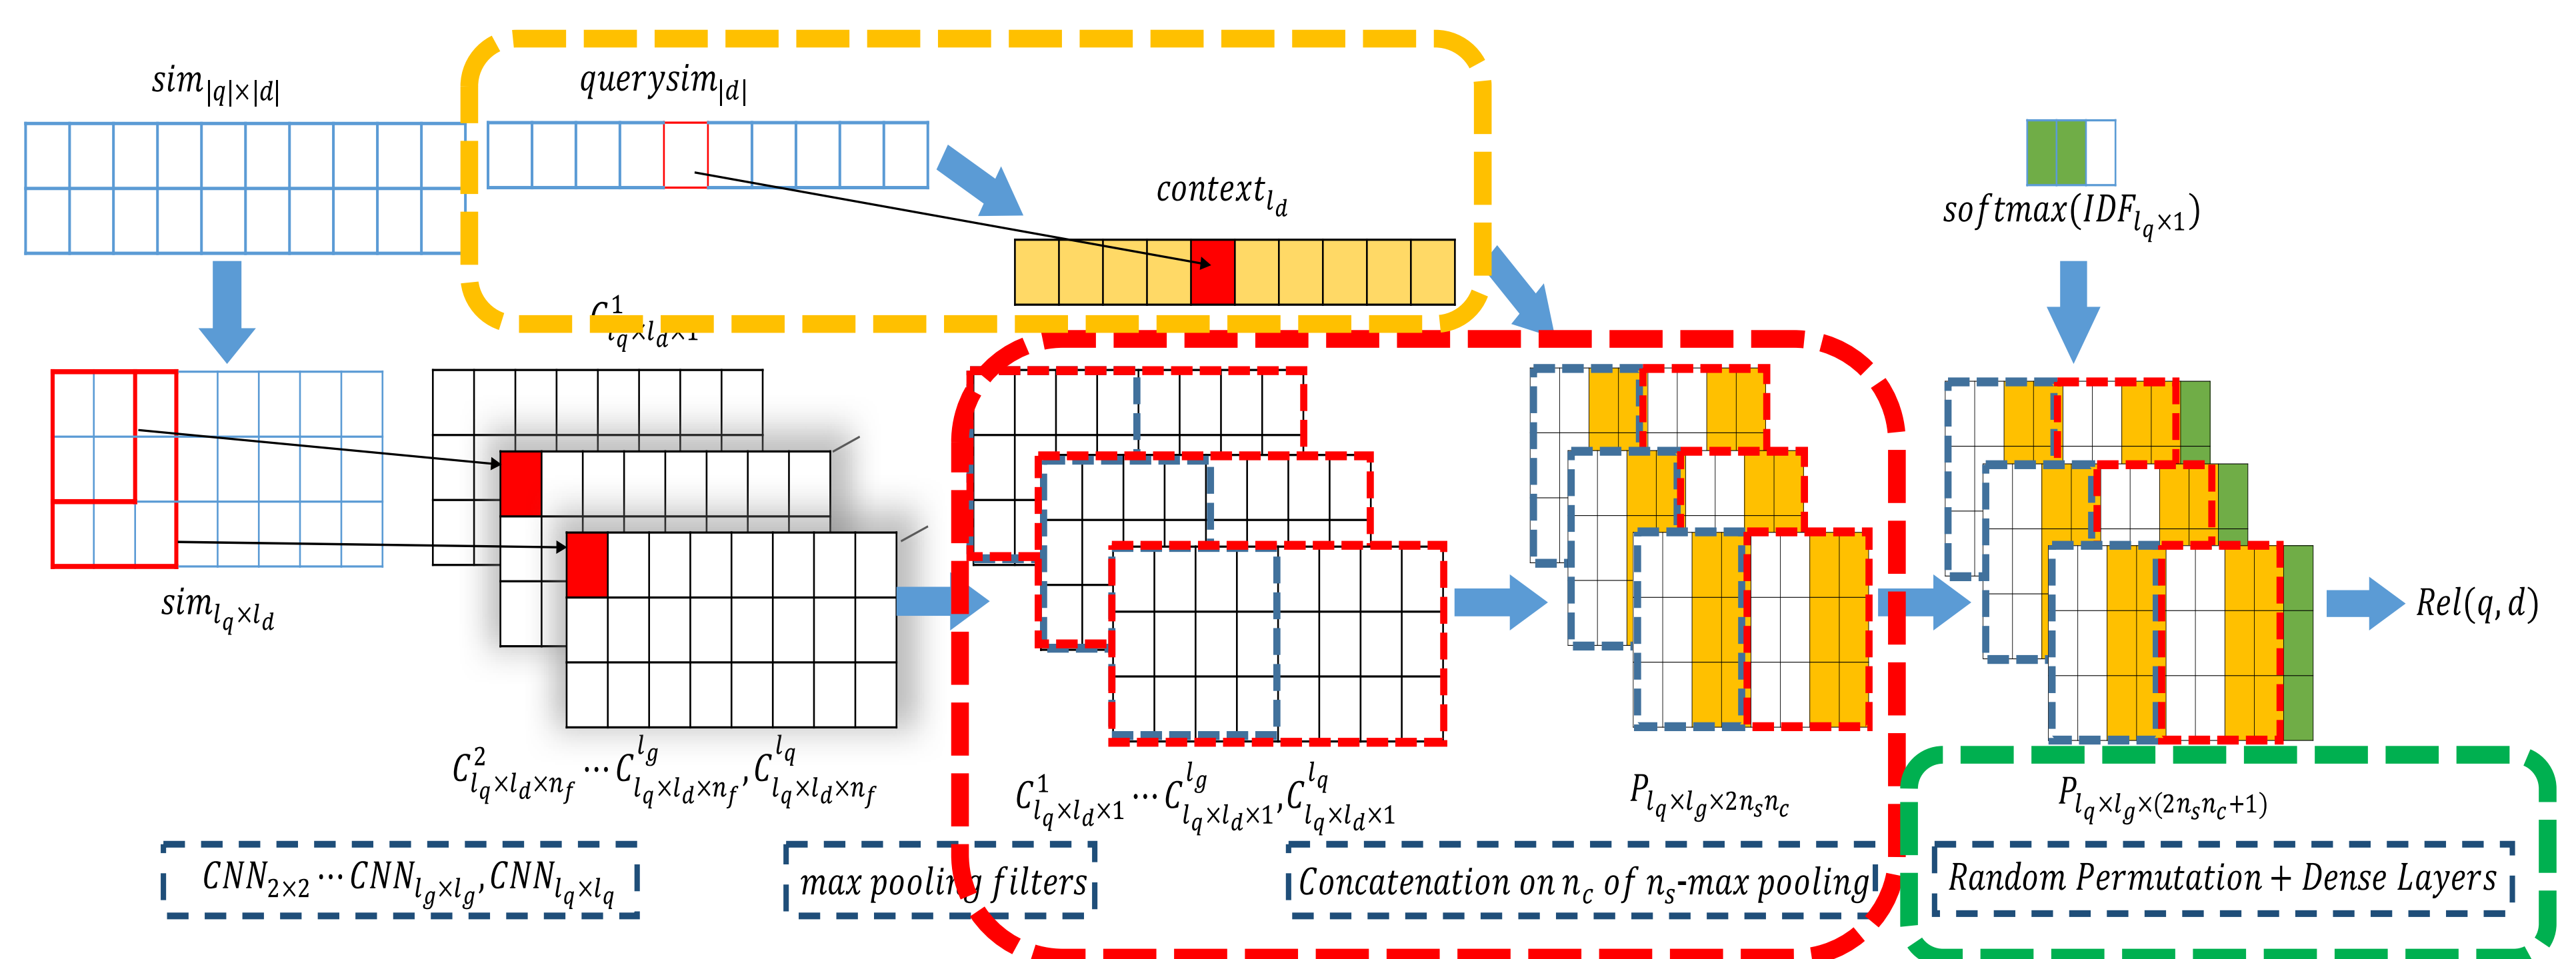
ayates@mpi-inf.mpg.de @andrewyates

Overview


MPII participated in the English subtask of WWW-2 [4] in order to evaluate several variants of the PACRR [2] and Co-PACRR [3] neural re-ranking models, including a modified query term combination strategy. No significant differences were found between any pair of runs, however.

Methodology

We evaluated variants of the PACRR and Co-PACRR [2,3] models, including cascade pooling. [1] All variants summed per-query-term scores rather than using a LSTM [2] or query-wide fully connected layer. [3]



Cascade relevance model



first half of the doc complete document

Run	Cascade pooling	k -max pooling	Hidden layer size
MPII-E-CO-NU-Base-1	25%, 50%, 75%, 100%	5	1
MPII-E-CO-NU-Base-2	25%, 50%, 75%, 100%	15	8
MPII-E-CO-NU-Base-3	25%, 50%, 75%, 100%	5	8
MPII-E-CO-NU-Base-4	100%	5	1
MPII-E-CO-NU-Base-5	100%	15	8

Table 1. Model hyperparameters that varied across runs. *Cascade pooling* indicates the document positions considered, *k-max pooling* indicates the number of term matches considered for each query term, and *hidden layer size* indicates the size of the hidden layer used in the combination component.

Results

All runs were trained on TREC 2009-2013 Web Track data with WT14 and NTCIR WWW-1 reserved for validation. Models re-ranked the BM25 baseline provided by the WWW-2 organizers. All variants used a maximum n-gram size of 3 and 16 kernels with each CNN.

Run	nDCG@10	Q@10	nERR@10
MPII-E-CO-NU-Base-1	0.3204	0.3009	0.4541
MPII-E-CO-NU-Base-2	0.3394	0.3255	0.4590
MPII-E-CO-NU-Base-3	0.3413	0.3183	0.4658
MPII-E-CO-NU-Base-4	0.3336	0.3265	0.4723
MPII-E-CO-NU-Base-5	0.3293	0.3110	0.4584

Results from WWW-2. There are no significant differences between pairs. [4]

[1] Craswell, N., Zoeter, O., Taylor, M., Ramsey, B.: An experimental comparison of click position-bias models. In: Proceedings of the 2008 International Conference on Web Search and Data Mining (2008).
 [2] Hui, K., Yates, A., Berberich, K., de Melo, G.: PACRR: A position-aware neural IR model for relevance matching. In: Proceedings of EMNLP (2017).
 [3] Hui, K., Yates, A., Berberich, K., de Melo, G.: Co-PACRR: A context-aware neural IR model for ad-hoc retrieval. In: Proceedings of the 2018 International Conference on Web Search and Data Mining (2018).
 [4] Mao, J., Sakai, T., Luo, C., Xiao, P., Liu, Y., Dou, Z.: Overview of the NTCIR-14 We Want Web task. In: Proceedings of the 14th NTCIR Conference on Evaluation of Information Access Technologies (2019).

