International Conference on Information and Knowledge Management, Proceedings, 2018, pages 1173-1182

Attentive neural architecture for ad-hoc structured document retrieval

Balaneshinkordan S., Kotov A., Nikolaev F.
Kazan Federal University, 420008, Kremlevskaya 18, Kazan, Russia

Abstract

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM. The problem of ad-hoc structured document retrieval arises in many information access scenarios, from Web to product search. Yet neither deep neural networks, which have been successfully applied to ad-hoc information retrieval and Web search, nor the attention mechanism, which has been shown to significantly improve the performance of deep neural networks on natural language processing tasks, have been explored in the context of this problem. In this paper, we propose a deep neural architecture for ad-hoc structured document retrieval, which utilizes attention mechanism to determine important phrases in keyword queries as well as the relative importance of matching those phrases in different fields of structured documents. Experimental evaluation on publicly available collections for Web document, product and entity retrieval from knowledge graphs indicates superior retrieval accuracy of the proposed neural architecture relative to both state-of-the-art neural architectures for ad-hoc document retrieval and probabilistic models for ad-hoc structured document retrieval.

http://dx.doi.org/10.1145/3269206.3271801

Keywords

Attention, Deep Neural Networks, Structured Document Retrieval

References

[1] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In Proceedings of ICLR.

[2] Saeid Balaneshin-kordan and Alexander Kotov. 2016. A study of document expansion using translation models and dimensionality reduction methods. In Proceedings of ACM ICTIR. 233-236.

[3] Saeid Balaneshin-kordan and Alexander Kotov. 2017. Embedding-based Query Expansion for Weighted Sequential Dependence Retrieval Model. In Proceedings of ACM SIGIR. 1213-1216.

[4] Saeid Balaneshin-kordan and Alexander Kotov. 2018. Deep Neural Architecture for Multi-Modal Retrieval based on Joint Embedding Space for Text and Images. In Proceedings of ACM WSDM. 28-36.

[5] Marco Baroni, Georgiana Dinu, and Germ n Kruszewski. 2014. Don't count, predict! A systematic comparison of context-counting vs. context-predicting semantic vectors.. In Proceedings of ACL. 238-247.

[6] Alexey Borisov, Ilya Markov, Maarten de Rijke, and Pavel Serdyukov. 2016. A neural click model for web search. In Proceedings of WWW. 531-541.

[7] Zhuyun Dai, Chenyan Xiong, Jamie Callan, and Zhiyuan Liu. 2018. Convolutional Neural Networks for Soft-Matching N-Grams in Ad-hoc Search. In Proceedings of ACM WSDM. 126-134.

[8] Mostafa Dehghani, Hamed Zamani, Aliaksei Severyn, Jaap Kamps, and W Bruce Croft. 2017. Neural Ranking Models with Weak Supervision. arXiv preprint arXiv:1704.08803 (2017).

[9] Jiafeng Guo, Yixing Fan, Qingyao Ai, and W Bruce Croft. 2016. A Deep Relevance Matching Model for Ad-hoc Retrieval. In Proceedings of ACM CIKM. 55-64.

[10] Faegheh Hasibi, Fedor Nikolaev, Chenyan Xiong, Krisztian Balog, Svein Erik Bratsberg, Alexander Kotov, and Jamie Callan. 2017. DBpedia-Entity v2: A Test Collection for Entity Search. In Proceedings of ACM SIGIR. 1265-1268.

[11] Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen. 2014. Convolutional neural network architectures for matching natural language sentences. In Proceedings of NIPS. 2042-2050.

[12] Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning deep structured semantic models for web search using click-through data. In Proceedings of ACM CIKM. 2333-2338.

[13] Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In Proceedings of ACL. 655-665.

[14] Jinyoung Kim, Xiaobing Xue, and W Bruce Croft. 2009. A probabilistic retrieval model for semistructured data. In Proceedings of ECIR. 228-239.

[15] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014).

[16] Alexander Kotov, Vineeth Rakesh, Eugene Agichtein, and Chandan K Reddy. 2015. Geographical latent variable models for microblog retrieval. In Proceedings of ECIR. 635-647.

[17] Zhengdong Lu and Hang Li. 2013. A deep architecture for matching short texts. In Proceedings of NIPS. 1367-1375.

[18] Irina Matveeva, Chris Burges, Timo Burkard, Andy Laucius, and Leon Wong. 2006. High accuracy retrieval with multiple nested ranker. In Proceedings of ACM SIGIR. 437-444.

[19] Donald Metzler and W Bruce Croft. 2005. A Markov random field model for term dependencies. In Proceedings of ACM SIGIR. 472-479.

[20] Donald Metzler and W Bruce Croft. 2007. Linear feature-based models for information retrieval. Information Retrieval (2007), 257-274.

[21] Bhaskar Mitra and Nick Craswell. 2017. Neural Models for Information Retrieval. arXiv preprint arXiv:1705.01509 (2017).

[22] Eric Nalisnick, Bhaskar Mitra, Nick Craswell, and Rich Caruana. 2016. Improving Document Ranking with Dual Word Embeddings. In Proceedings of WWW. 83-84.

[23] Fedor Nikolaev, Alexander Kotov, and Nikita Zhiltsov. 2016. Parameterized fielded term dependence models for ad-hoc entity retrieval from knowledge graph. In Proceedings of ACM SIGIR. 435-444.

[24] Paul Ogilvie and Jamie Callan. 2003. Combining document representations for known-item search. In Proceedings of ACM SIGIR. 143-150.

[25] Liang Pang, Yanyan Lan, Jiafeng Guo, Jun Xu, Shengxian Wan, and Xueqi Cheng. 2016. Text matching as image recognition. In Proceedings of AAAI. 12-17.

[26] Benjamin Piwowarski and Patrick Gallinari. 2003. A machine learning model for information retrieval with structured documents. Machine Learning and Data Mining in Pattern Recognition (2003), 425-438.

[27] Navid Rekabsaz, Mihai Lupu, Allan Hanbury, and Hamed Zamani. 2017. Word Embedding Causes Topic Shifting; Exploit Global Context!. In Proceedings of ACM SIGIR. 1105-1108.

[28] Stephen Robertson, Hugo Zaragoza, and Michael Taylor. 2004. Simple BM25 extension to multiple weighted fields. In Proceedings of ACM CIKM. 42-49.

[29] Alexander M Rush, Sumit Chopra, and Jason Weston. 2015. A neural attention model for abstractive sentence summarization. In Proceedings of EMNLP. 379-389.

[30] Aliaksei Severyn and Alessandro Moschitti. 2015. Learning to rank short text pairs with convolutional deep neural networks. In Proceedings of ACM SIGIR. 373-382.

[31] Daniel Sheldon, Milad Shokouhi, Martin Szummer, and Nick Craswell. 2011. LambdaMerge: merging the results of query reformulations. In Proceedings of ACM WSDM. 795-804.

[32] Yelong Shen, Xiaodong He, Jianfeng Gao, Li Deng, and Gr goire Mesnil. 2014. A latent semantic model with convolutional-pooling structure for information retrieval. In Proceedings of ACM CIKM. 101-110.

[33] Rupesh K Srivastava, Klaus Greff, and J rgen Schmidhuber. 2015. Training very deep networks. In Proceedings of NIPS. 2377-2385.

[34] Krysta M Svore and Christopher JC Burges. 2009. A machine learning approach for improved BM25 retrieval. In Proceedings of ACM CIKM. 1811-1814.

[35] Kai Sheng Tai, Richard Socher, and Christopher D Manning. 2015. Improved semantic representations from tree-structured long short-term memory networks. In Proceedings of ACL. 1556-1566.

[36] Andrew Trotman. 2004. Optimal Structure Weighted Retrieval. In Proceedings of ADCS.

[37] Andrew Trotman. 2004. Searching structured documents. Information Processing & Management 40, 4 (2004), 619-632.

[38] Chenyan Xiong, Zhuyun Dai, Jamie Callan, Zhiyuan Liu, and Russell Power. 2017. End-to-end neural ad-hoc ranking with kernel pooling. In Proceedings of ACM SIGIR. 55-64.

[39] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In Proceedings of ICML. 2048-2057.

[40] Liu Yang, Qingyao Ai, Jiafeng Guo, and W Bruce Croft. 2016. aNMM: Ranking short answer texts with attention-based neural matching model. In Proceedings of ACM CIKM. 287-296.

[41] Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alexander J Smola, and Ed-uard H Hovy. 2016. Hierarchical Attention Networks for Document Classification.. In HLT-NAACL. 1480-1489.

[42] Wenpeng Yin and Hinrich Sch tze. 2015. MultiGranCNN: An Architecture for General Matching of Text Chunks on Multiple Levels of Granularity.. In Proceedings of ACL. 63-73.

[43] Haonan Yu, Jiang Wang, Zhiheng Huang, Yi Yang, and Wei Xu. 2016. Video paragraph captioning using hierarchical recurrent neural networks. In Proceedings of IEEE CVPR. 4584-4593.

[44] Hamed Zamani, Bhaskar Mitra, Xia Song, Nick Craswell, and Saurabh Tiwary. 2018. Neural Ranking Models with Multiple Document Fields. In Proceedings of ACM WSDM. 700-708.

[45] Nikita Zhiltsov, Alexander Kotov, and Fedor Nikolaev. 2015. Fielded sequential dependence model for Ad-Hoc entity retrieval in the web of data. In Proceedings of ACM SIGIR. 253-262.