

Neu-IR: The SIGIR 2016 Workshop on Neural Information Retrieval

Nick Craswell Microsoft Bellevue, US nickcr@microsoft.com W. Bruce Croft University of Massachusetts Amherst, US croft@cs.umass.edu Jiafeng Guo Chinese Academy of Sciences Beijing, China guojiafeng@ict.ac.cn

Bhaskar Mitra Microsoft Cambridge, UK bmitra@microsoft.com Maarten de Rijke University of Amsterdam Amsterdam, The Netherlands derijke@uva.nl

Keywords

Neural networks; deep learning; information retrieval

Website: http://research.microsoft.com/neuir2016

1. MOTIVATION

In recent years, deep neural networks have yielded significant performance improvements on speech recognition and computer vision tasks [4, 7], as well as led to exciting breakthroughs in novel application areas such as automatic voice translation [8], image captioning [2, 12], and conversational agents [11]. Despite demonstrating good performance on natural language processing (NLP) tasks (e.g., language modelling [6] and machine translation [1]), the performance of deep neural networks on information retrieval (IR) tasks has had relatively less scrutiny. Recent work in this area has mainly focused on word embeddings [3, 9, 14] and neural models for short text similarity [5, 10].

The lack of many positive results in this area of information retrieval is partially due to the fact that IR tasks such as ranking are fundamentally different from NLP tasks, but also because the IR and neural network communities are only beginning to focus on the application of these techniques to core information retrieval problems. Given that deep learning has made such a big impact, first on speech processing and computer vision and now, increasingly, also on computational linguistics, it seems clear that deep learning will have a major impact on information retrieval and that this is an ideal time for a workshop in this area.

Neu-IR (pronounced "*new IR*") will be a forum for new research relating to deep learning and other neural network based approaches to IR. The purpose is to provide an opportunity for people to present new work and early results, compare notes on neural network toolkits, share best practices, and discuss the main challenges facing this line of research.

SIGIR '16 July 17-21, 2016, Pisa, Italy

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4069-4/16/07.

DOI: http://dx.doi.org/10.1145/2911451.2917762

2. SCOPE

The Neu-IR workshop will be a gathering of academic and industrial researchers working at the intersection of IR and neural networks. We solicit submission of papers of two to six pages, representing reports of original research, preliminary research results, proposals for new work, descriptions of neural network based toolkits tailored for IR, and position papers. Papers presented at the workshop will be required to be uploaded to arXiv.org but will be considered nonarchival, and may be submitted elsewhere (modified or not), although the workshop site will maintain a link to the arXiv versions. This makes the workshop a forum for the presentation and discussion of current work, without preventing the work from being published elsewhere. We are interested in submissions relevant to the following main themes:

- The application of neural network models in IR tasks, including but not limited to:
 - Full text document retrieval, passage retrieval, question answering
 - Web search, searching social media, distributed information retrieval, entity ranking
 - Learning to rank combined with neural network based representation learning
 - User / task modelling, personalization, diversity
 - Query formulation assistance, query recommendation, conversational search
 - Multimedia retrieval
- Fundamental modelling challenges faced in such applications, including but not limited to:
 - Learning dense representations for long documents
 - Dealing with rare queries and rare words
 - Modelling text at different granularities (character, word, passage, document)
 - Compositionality of vector representations
 - Jointly modelling queries, documents, entities and other structured/knowledge data
- Best practices for research and development in the area, dealing with concerns such as:
 - Finding sufficient publicly-available training data
 - Baselines, test data, avoiding overfitting
 - Neural network toolkits
 - Real-world use cases, deployment at scale

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

3. INTERACTION FORMAT

Neu-IR will be a highly interactive full day workshop, featuring a mix of presentation and interaction formats, appropriate for the main themes identified in §2. The workshop will feature presentations of the following types:

- Invited keynotes
- Invited toolkit walkthrough on CNTK [13]
- Contributed paper presentations
- An interactive "Lessons from the Trenches" session to share best practices and lessons learned
- Breakout sessions and panel discussions on issues that emerge from the contributed papers

4. ORGANIZERS

Nick Craswell is a Principal Lead Researcher in Microsoft Bing, working on core relevance. And leads a small team of Applied Researchers who are embedded in Microsoft Research Cambridge in the UK. He has published a number of papers relating to clicks and their use, and he ran the first successful WSDM workshop on Web Search Click data in 2009, and co-organized WSCD workshops in 2012 and 2013.

Bruce Croft is a Distinguished Professor in the College of Information and Computer Sciences at the University of Massachusetts Amherst. He leads the Center for Intelligent Information Retrieval, which has produced results in many areas of IR for the past 25 years.

Jiafeng Guo is an Associate Professor at Institute of Computing Technology, Chinese Academy of Sciences. He has worked on a number of topics related to Web search and data mining, including query representation and understanding, learning to rank, and topic modelling. His current research is focused on representation learning and deep models for information retrieval and information filtering.

Bhaskar Mitra is a Senior Applied Scientist at Bing in Microsoft Research Cambridge. He started at Bing in 2007 (then called Live Search) at the Search Technology Center in India. He is currently pursuing a (part-time) doctorate at University College London. His research interests include representation learning and neural networks, and in particular their application to modern day information retrieval.

Maarten de Rijke is a Professor of Computer Science at the Informatics Institute of the University of Amsterdam. Together with a team of PhD students and postdocs he works on problems on semantic search and on- and offline learning to rank for information retrieval. Some of their recent work uses deep neural networks for text similarity, entity search, product search, session analysis, and click modelling.

5. PROGRAM COMMITTEE

- Carsten Eickhoff, ETH Zurich
- Debasis Ganguly, Dublin City University
- Katja Hoffman, Microsoft Research
- Hang Li, Huawei Technologies
- Piotr Mirowski, Google DeepMind
- Alessandro Moschitti, Qatar Computing Research Inst.

- Pavel Serdyukov, Yandex
- Fabrizio Silvestri, Yahoo Labs
- Alessandro Sordoni, University of Montreal

References

- D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473, 2014.
- [2] H. Fang, S. Gupta, F. Iandola, R. Srivastava, L. Deng, P. Dollár, J. Gao, X. He, M. Mitchell, J. Platt, et al. From captions to visual concepts and back. arXiv preprint arXiv:1411.4952, 2014.
- [3] D. Ganguly, D. Roy, M. Mitra, and G. J. Jones. Word embedding based generalized language model for information retrieval. In *Proc. SIGIR*, pages 795–798. ACM, 2015.
- [4] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *Signal Processing Magazine*, *IEEE*, 29(6):82–97, 2012.
- [5] P.-S. Huang, X. He, J. Gao, L. Deng, A. Acero, and L. Heck. Learning deep structured semantic models for web search using clickthrough data. In *Proc. CIKM*, pages 2333–2338. ACM, 2013.
- [6] R. Jozefowicz, O. Vinyals, M. Schuster, N. Shazeer, and Y. Wu. Exploring the limits of language modeling. arXiv preprint arXiv:1602.02410, 2016.
- [7] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.
- [8] W. D. Lewis. Skype translator: Breaking down language and hearing barriers. In *Translating and the Computer (TC37)*, 2015.
- [9] B. Mitra, E. Nalisnick, N. Craswell, and R. Caruana. A dual embedding space model for document ranking. arXiv preprint arXiv:1602.01137, 2016.
- [10] A. Severyn and A. Moschitti. Learning to rank short text pairs with convolutional deep neural networks. In *Proc. SIGIR*, pages 373–382. ACM, 2015.
- [11] O. Vinyals and Q. Le. A neural conversational model. arXiv preprint arXiv:1506.05869, 2015.
- [12] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. arXiv preprint arXiv:1411.4555, 2014.
- [13] D. Yu, A. Eversole, M. Seltzer, K. Yao, et al. An introduction to computational networks and the computational network toolkit. Technical report, Tech. Rep. MSR, Microsoft Research, 2014, https://cntk.ai/, 2014.
- [14] G. Zheng and J. Callan. Learning to reweight terms with distributed representations. In *Proc. SIGIR*, pages 575–584. ACM, 2015.