Maarten de Rijke

derijke@uva.nl University of Amsterdam

Amsterdam, The Netherlands



SIGIR 2017 Workshop on Neural Information Retrieval (Neu-IR'17)

Nick Craswell nickcr@microsoft.com Microsoft Bellevue, US W. Bruce Croft croft@cs.umass.edu University of Massachusetts Amherst, US

Jiafeng Guo guojiafeng@ict.ac.cn Chinese Academy of Sciences Beijing, China Bhaskar Mitra bmitra@microsoft.com Microsoft Cambridge, UK

to identify shared needs for test collections, evaluation tasks, and software infrastructure—and that will be the focus of Neu-IR:'17.

2 SCOPE

Neu-IR:'17 will be a forum for setting agenda and resources development relating to deep learning and other neural network based approaches to IR. The purpose is to provide an opportunity for academic and industrial researchers—working at the intersection of IR and neural networks—to come together and define shared research agenda, which may include:

- What are the core resources (evaluation, training material, shared implementations) that are most urgently needed? Specific issues that emerge here include:
 - Identifying appropriate large scale benchmark collections appropriate for training and evaluating DNNs with millions (or billions) of parameters.
 - Standardizing metrics and tasks appropriate for evaluating DNNs.
 - Building a central shared model repository without enforcing the use of any specific NN toolkit.
- Now that some working models have been proposed in the area, what are the core research themes related to neural methods for IR that we should focus on?
 - Specific issues that emerge here include:
 - Dense representations for long documents. Rare words. Compositionality of vector representations. Which representations for which IR tasks?
 - Which architectures fit which retrieval tasks? What are the most effective architectures for ranking tasks? How to model learning to rank as a DNN objective?

These are just a few of the many deep learning needs that are highly relevant to the IR community. SIGIR is uniquely positioned to host a workshop that would (1) help shape shared research directions and (2) help identify and strategize on shared evaluation and infrastructure needs regarding deep learning and IR.

3 INTERACTION FORMAT

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

Neu-IR will feature presentations of the following types:

Website: http://neu-ir.weebly.com

1 MOTIVATION

In recent years, deep neural networks (DNNs) have yielded significant performance improvements in application areas such as speech recognition, computer vision, and machine translation. We have now seen the first successful applications of DNNs in short-text matching [6–8, 13, 15, 16], in ad-hoc retrieval [5, 10], in a range of recommendation tasks [17], and in the analysis of users' information interaction behavior [1]. Interestingly, successful architectures used in computer vision may have 100+ layers, while in language technology more shallow networks, with at most a dozen layers, but usually fewer, have proven to be more effective. No clear picture has emerged yet on the type of architecture that is most appropriate for ad-hoc retrieval. In fact, due to the lack of availability of large scale training data much of the explorations in this area is constrained to the application of word embeddings within existing information retrieval (IR) models [3, 4, 11, 14, 18–20].

What has become clear, though, is that identifying information that satisfies a user's information need is fundamentally different from finding regularities in texts or multimedia documents. New architectures are needed, new loss functions, and new optimization techniques, catering specifically to IR problems. But there is more to the relation between DNNs and IR than algorithm development. Since DNNs operate differently from traditional IR methods, they can lead to new insights into core IR problems. This has already happened as documented in the work listed in [9, 12]. Finally, new evaluation resources, especially large scale benchmark collections, are needed for sufficient learning and evaluation of DNNs for IR.

Neu-IR:'17 (pronounced "*new IR*") will be the second workshop on neural IR, after the highly successful 2016 edition [2], which brought together more than 120 researchers from academia and industry and which gave rise to a special issue of the *Information Retrieval Journal*. Deep learning will have a major impact on IR as a field. The time is right to discuss a shared research agenda,

SIGIR '17, August 7-11, 2017, Shinjuku, Tokyo, Japan

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).

^{© 2017} Copyright held by the owner/author(s). 978-1-4503-5022-8/17/08. DOI: 10.1145/3077136.3084373

- Invited keynotes (industrial and academic)
- Contributed proposals on generating large scale benchmark collections, building a shared model repository, and standardizing frameworks appropriate for evaluating DNNs
- Contributed poster presentations and demos.
- Multiple breakout sessions on research agenda, evaluation, data, infrastructure related to neural methods for IR.

4 LIST OF ORGANIZERS

All the organizers from Neu-IR 2016 will be involved in the organization of the second edition of the workshop.

Nick Craswell is a Principal Science Manager 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 past 25 years.

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 IR. Some of their recent work uses DNNs for text similarity, entity search, product search, session analysis, and click modeling.

Jiafeng Guo is a 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 modeling. His current research is focused on representation learning and deep models for IR and information filtering.

Bhaskar Mitra is a Senior Applied Scientist in Microsoft Bing, where he has worked since 2007 (then called Live Search). His research interests include representation learning and neural networks, and their applications to IR. He co-organized multiple tutorials (WSDM, SIGIR, and RUSSIR) in 2017 on neural IR, and served as guest editor for the IRJ special issue on the same topic.

5 NAMES OF POTENTIAL PROGRAM COMMITTEE MEMBERS

Potential PC members for reviewing paper submissions:

- Alessandro Moschitti, University of Trento
- Alexey Borisov, University of Amsterdam
- Aliaksei Severyn, Google Research
- Carsten Eickhoff, ETH Zurich
- Christina Lioma, University of Copenhagen
- Christophe Van Gysel, University of Amsterdam
- Fabrizio Silvestri, Facebook
- Hamed Zamani, University of Massachusetts Amherst

- Hang Li, Huawei
- Mostafa Dehghani, University of Amsterdam
- Pavel Serdyukov, Yandex
- Peng Zhang, Tianjin University
- Piotr Mirowski, Google DeepMind
- Qi Zhang, Fudan University
- Rishabh Mehrotra, University College London
- Yiqun Liu, Tsinghua University
- Zhicheng Dou, Renmin University of China

REFERENCES

- Alexey Borisov, Ilya Markov, Maarten de Rijke, and Pavel Serdyukov. 2016. A neural click model for web search. In *Proceedings of the 25th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 531–541.
- [2] Nick Craswell, W Bruce Croft, Jiafeng Guo, Bhaskar Mitra, and Maarten de Rijke. 2016. Report on the SIGIR 2016 Workshop on Neural Information Retrieval (Neu-IR). 50, 2 (2016), 96–103.
- [3] Fernando Diaz, Bhaskar Mitra, and Nick Craswell. 2016. Query Expansion with Locally-Trained Word Embeddings. In *Proc. ACL*.
- [4] Debasis Ganguly, Dwaipayan Roy, Mandar Mitra, and Gareth JF Jones. 2015. Word Embedding based Generalized Language Model for Information Retrieval. In Proc. SIGIR. ACM, 795–798.
- [5] Jiafeng Guo, Yixing Fan, Qingyao Ai, and W Bruce Croft. 2016. A Deep Relevance Matching Model for Ad-hoc Retrieval. In Proc. CIKM. ACM, 55–64.
- [6] Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen. 2014. Convolutional neural network architectures for matching natural language sentences. In Proc. NIPS. 2042–2050.
- [7] Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning deep structured semantic models for web search using clickthrough data. In *Proc. CIKM*. ACM, 2333–2338.
- [8] Zhengdong Lu and Hang Li. 2013. A deep architecture for matching short texts. In Advances in Neural Information Processing Systems. 1367–1375.
- [9] Bhaskar Mitra and Nick Craswell. 2017. Neural Models for Information Retrieval. arXiv preprint arXiv:1705.01509 (2017).
- [10] Bhaskar Mitra, Fernando Diaz, and Nick Craswell. 2017. Learning to Match Using Local and Distributed Representations of Text for Web Search. In Proc. WWW. 1291-1299.
- [11] Bhaskar Mitra, Eric Nalisnick, Nick Craswell, and Rich Caruana. 2016. A Dual Embedding Space Model for Document Ranking. arXiv preprint arXiv:1602.01137 (2016).
- [12] Kezban Dilek Onal, Ye Zhang, and others. 2017. Neural Information Retrieval: At the End of the Early Years. *Information Retrieval Journal* (2017). Submitted.
- [13] Liang Pang, Yanyan Lan, Jiafeng Guo, Jun Xu, Shengxian Wan, and Xueqi Cheng. 2016. Text Matching as Image Recognition. In Proc. AAAI.
- [14] Dwaipayan Roy, Debjyoti Paul, Mandar Mitra, and Utpal Garain. 2016. Using word embeddings for automatic query expansion. In *Neu-IR: The SIGIR 2016* Workshop on Neural Information Retrieval.
- [15] Aliaksei Severyn and Alessandro Moschitti. 2015. Learning to rank short text pairs with convolutional deep neural networks. In *Proc. SIGIR*. ACM, 373–382.
 [16] Yelong Shen, Xiaodong He, Jianfeng Gao, Li Deng, and Gregoire Mesnil. 2014.
- [16] Yelong Shen, Xiaodong He, Jianfeng Gao, Li Deng, and Gregoire Mesnil. 2014. A latent semantic model with convolutional-pooling structure for information retrieval. In *Proc. CIKM*. ACM, 101–110.
- [17] Chao-Yuan Wu, Amr Ahmed, Alex Beutel, Alexander J. Smola, and How Jing. 2017. Recurrent Recommender Networks. In Proceedings of the Tenth ACM International Conference on Web Search and Data Mining (WSDM '17). ACM, New York, NY, USA, 495–503. DOI:http://dx.doi.org/10.1145/3018661.3018689
- [18] Hamed Zamani and W Bruce Croft. 2016. Embedding-based query language models. In Proceedings of the 2016 ACM on International Conference on the Theory of Information Retrieval. ACM, 147–156.
- [19] Hamed Zamani and W Bruce Croft. 2016. Estimating embedding vectors for queries. In Proceedings of the 2016 ACM on International Conference on the Theory of Information Retrieval. ACM, 123–132.
- [20] Guido Zuccon, Bevan Koopman, Peter Bruza, and Leif Azzopardi. 2015. Integrating and evaluating neural word embeddings in information retrieval. In Proceedings of the 20th Australasian Document Computing Symposium. ACM, 12.