

Deep Natural Language Processing for Search and Recommender Systems

Weiwei Guo, Huiji Gao, Jun Shi, Bo Long, Liang Zhang, Bee-Chung Chen, Deepak Agarwal

LinkedIn, Mountain View, California

{wguo,hgao,jshi,blong,lizhang,bchen,dagarwal}@linkedin.com

ABSTRACT

Search and recommender systems share many fundamental components including language understanding, retrieval and ranking, and language generation. Building powerful search and recommender systems requires processing natural language effectively and efficiently. Recent rapid growth of deep learning technologies has presented both opportunities and challenges in this area. This tutorial offers an overview of deep learning based natural language processing (NLP) for search and recommender systems from an industry perspective. It first introduces deep learning based NLP technologies, including language understanding and language generation. Then it details how those technologies can be applied to common tasks in search and recommender systems, including query and document understanding, retrieval and ranking, and language generation. Applications in LinkedIn production systems are presented. The tutorial concludes with discussion of future trend.

KEYWORDS

Deep Learning, Natural Language Understanding/Generation, Search Engine, Recommender System

ACM Reference Format:

Weiwei Guo, Huiji Gao, Jun Shi, Bo Long, Liang Zhang, Bee-Chung Chen, Deepak Agarwal. 2019. Deep Natural Language Processing for Search and Recommender Systems. In *The 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '19)*, August 4–8, 2019, Anchorage, AK, USA. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3292500.3332290>

1 INTRODUCTION

Search and recommender systems process rich natural language text data such as user queries and documents. Achieving high-quality search and recommendation results requires processing and understanding such information effectively and efficiently, where natural language processing (NLP) technologies are widely deployed. In recent years, the rapid development of deep learning models has been proven successful for improving various NLP tasks, indicating their great potential of promoting search and recommender systems.

In this tutorial, we summarize the current effort of deep learning for NLP in search/recommender systems. We first give an overview of search/recommender systems with NLP, then introduce basic concept of deep learning for NLP [4, 9, 17, 19, 26, 27], covering

state-of-the-art technologies in both language understanding and language generation. After that, we share our hands-on experience with LinkedIn applications. In the end, we highlight several important future trends.

We focus on three fundamental components of search and recommender systems: 1). **query and document understanding** extracts and infers relevant information, such as intent prediction [16], entity tagging and disambiguation [13, 24], topic understanding [1, 18], and opinion mining [11, 14]; 2). **retrieval and ranking** methodologies designed with strong latency restrictions [6, 8, 15] and various matching strategies [3, 10, 12]; and 3). **language generation** techniques designed to proactively guide/interact with users to further resolve ambiguity, including query reformulation, (i.e., query suggestion [7, 23], auto-completion [20], spell correction [5]) and conversational recommendation [2, 25]. At last, we illustrate end-to-end real-world examples in LinkedIn, sharing our experiences in algorithm development and infrastructure building, especially a conversation AI [21, 22] usage for help center and natural language search at LinkedIn.

Targeted Audience The tutorial is suitable for academic and industrial researchers, graduate students, and practitioners. After the tutorial, we expect the audience have learnt concepts and principles of applying state-of-the-art deep learning for NLP in search/recommender systems, and gained real-world experiences through illustrative examples of end-to-end systems.

2 TUTORIAL OUTLINE

1. Introduction (15 mins)

- (1) Overview of Search and Recommender Systems
- (2) NLP in Search and Recommender Systems

2. Deep Learning for NLP (45 mins)

- (1) Preliminaries
- (2) Language Understanding
- (3) Language Generation

3. Deep NLP in Search and Recommender Systems (90 mins)

- (1) Query and Document Understanding
 - Intent Prediction
 - Entity Tagging and Disambiguation
 - Topic Understanding
 - Opinion Mining
- (2) Retrieval and Ranking
 - Efficient Candidate Retrieval
 - Deep Ranking Models
- (3) Language Generation for Interaction with Users
 - Query Reformulation
 - Conversational Recommendation

4. Real-world Examples (50 mins)

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).

KDD '19, August 4–8, 2019, Anchorage, AK, USA

© 2019 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-6201-6/19/08.

<https://doi.org/10.1145/3292500.3332290>

- (1) An End-to-end Example of Deep NLP in LinkedIn Job Search/Recommender Systems
- (2) Conversational AI
 - Chatbot for LinkedIn Help Center
 - Natural Language Search at LinkedIn

5. Future Trends and Conclusions (10 mins)

3 PRESENTERS' BIOGRAPHY

Dr. Weiwei Guo is a senior software engineer at LinkedIn where he leads several efforts to apply the deep learning models into search productions. He obtained his Ph.D. with focus on NLP from Columbia University. Weiwei has published over 20 peer-reviewed papers in top conferences including ACL, EMNLP, NAACL, SIGIR with 1000+ citations.

Dr. Huiji Gao leads the AI Algorithms Foundation team at LinkedIn. He has broad interests in machine learning/AI and its applications, including search/recommender systems, computational advertising, and NLP. He received Ph.D. in Computer Science from Arizona State University. He has filed 10+ U.S. patents and published 40+ publications in top journals and conferences including KDD, AAAI, WWW, ICDM, DMKD with thousands of citations.

Dr. Jun Shi is a staff software engineer at LinkedIn, leading various efforts on promoting deep-learning based NLP technologies. He is a co-recipient of 2009 IEEE Communications Society & Information Theory Society Joint Paper Award.

Dr. Bo Long leads LinkedIn's AI Foundations team. He has 15 years of experience in data mining and machine learning with applications to web search, recommendation, and social network analysis. He holds dozens of innovations and has published peer-reviewed papers in top conferences and journals including ICML, KDD, ICDM, AAAI, SDM, CIKM, and KAIS. He has served as reviewers, workshops co-organizers, conference organizer committee members, and area chairs for multiple conferences, including KDD, NIPS, SIGIR, ICML, SDM, CIKM, JSM etc.

Dr. Liang Zhang is a Director of Machine Learning Engineering at LinkedIn, leading Search and Sales Solution AI efforts. During his time at LinkedIn, Liang has led a lot of critical AI projects in many different products (Ads, Feed, Email and notification, Jobs, and Search) to success and brought great improvements of experiences to the 500M+ professional users of LinkedIn through the cutting-edge AI technology. Liang has published extensively in top-tier conferences and journals, and coauthored 20+ AI-related patents. Liang also served as the Program Committee members for various data mining and machine learning venues.

Dr. Bee-Chung Chen is a Principal Staff Engineer at LinkedIn with extensive industrial and research experience in recommender systems, some of which is summarized in the book titled Statistical Methods for Recommender Systems. He currently leads the development of LinkedIn's machine learning technology. He was a key designer of the recommendation algorithms that power LinkedIn news feed, Yahoo! homepage, Yahoo! News and other sites. His research interests include recommender systems, machine learning and big data processing.

Dr. Deepak Agarwal is the VP of Artificial Intelligence at LinkedIn. He is an expert in Artificial Intelligence technologies and

engineering leadership with more than twenty years of experience developing and deploying state-of-the-art machine learning and statistical methods for improving the relevance of web applications. He is a Fellow of the American Statistical Association, Member Board of Directors for SIGKDD, program chair of KDD in the past and associate editor of two top-tier journals in Statistics. He regularly serves on senior program committees of top-tier conferences like KDD, NIPS, CIKM, ICDM, SIGIR, WSDM.

REFERENCES

- [1] Trapit Bansal, David Belanger, and Andrew McCallum. 2016. Ask the gru: Multi-task learning for deep text recommendations. In *RecSys*.
- [2] Konstantina Christakopoulou, Alex Beutel, Rui Li, Sagar Jain, and Ed H. Chi. 2018. Q&R: A two-stage approach toward interactive recommendation. In *KDD*.
- [3] W. Bruce Croft, Donald Metzler, and Trevor Strohman. 2010. *Search engines: Information retrieval in practice*. Reading: Addison-Wesley.
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*.
- [5] Shaona Ghosh and Per Ola Kristensson. 2017. Neural Networks for Text Correction and Completion in Keyboard Decoding. In *arXiv preprint arXiv:1709.06429*.
- [6] Song Han, Huizi Mao, and William J Dally. 2016. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *ICLR*.
- [7] Yunlong He, Jiliang Tang, Hua Ouyang, Changsung Kang, Dawei Yin, , and Yi Chang. 2016. Personalized Language Model for Query Auto-Completion. In *CIKM*.
- [8] Matthew Henderson, Rami Al-Rfou, Brian Strope, Yun hsuan Sung, Laszlo Lukacs, Ruiqi Guo, Sanjiv Kumar, Balint Miklos, and Ray Kurzweil. 2017. Efficient natural language response suggestion for smart reply. In *arXiv preprint arXiv:1705.00652*.
- [9] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. In *Neural computation*.
- [10] Baotian Hu, Zhengdong Lu, Hang Li, , and Qingcai Chen. 2014. Convolutional Neural Network Architectures for Matching Natural Language Sentences. In *NIPS*.
- [11] Mingqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. In *KDD*.
- [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 clickthrough data. In *CIKM*.
- [13] Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. In *arXiv preprint arXiv:1508.01991*.
- [14] Ozan Irsoy and Claire Cardie. 2014. Opinion mining with deep recurrent neural networks. In *EMNLP*.
- [15] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2017. Billion-scale similarity search with GPUs. *arXiv preprint arXiv:1702.08734* (2017).
- [16] Yoon Kim. 2014. Convolutional Neural Networks for Sentence Classification. In *EMNLP*.
- [17] Yann LeCun and Yoshua Bengio. 1995. Convolutional networks for images, speech, and time series. In *The handbook of brain theory and neural networks*.
- [18] Yishu Miao, Lei Yu, and Phil Blunsom. 2016. Neural variational inference for text processing. In *ICML*.
- [19] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*.
- [20] Dae Hoon Park and Rikio Chiba. 2017. A neural language model for query auto-completion. In *SIGIR*.
- [21] Chen Qu, Liu Yang, W. Bruce Croft, Johanne R. Trippas, Yongfeng Zhang, and Minghui Qiu. 2018. Analyzing and Characterizing User Intent in Information-seeking Conversations. In *SIGIR*.
- [22] Filip Radlinski and Nick Craswell. 2017. A theoretical framework for conversational search. In *CHIIR*.
- [23] Gary Ren, Xiaochuan Ni, Manish Malik, and Qifa Ke. 2018. Conversational Query Understanding Using Sequence to Sequence Modeling. In *WWW*.
- [24] Wei Shen, Jianyong Wang, and Jiawei Han. 2014. Entity linking with a knowledge base: Issues, techniques, and solutions. *IEEE Transactions on Knowledge and Data Engineering* (2014).
- [25] Yueming Sun and Yi Zhang. 2018. Conversational recommender system. In *SIGIR*.
- [26] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. In *NIPS*.
- [27] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*.