

Deep Learning for Matching in Search and Recommendation

Other titles in Foundations and Trends® in Information Retrieval

Information Retrieval: The Early Years

Donna Harman

ISBN: 978-1-68083-584-7

Bandit Algorithms in Information Retrieval

Dorota Glowacka

ISBN: 978-1-68083-574-8

Neural Approaches to Conversational AI

Jianfeng Gao, Michel Galley and Lihong Li

ISBN: 978-1-68083-552-6

An Introduction to Neural Information Retrieval

Bhaskar Mitra and Nick Craswell

ISBN: 978-1-68083-532-8

Deep Learning for Matching in Search and Recommendation

Jun Xu

Gaoling School of Artificial Intelligence
Renmin University of China
China
junxu@ruc.edu.cn

Xiangnan He

School of Information Science and Technology
University of Science and Technology of China
China
hexn@ustc.edu.cn

Hang Li

Bytedance AI Lab
China
lihang.lh@bytedance.com

now

the essence of knowledge

Boston — Delft

Foundations and Trends[®] in Information Retrieval

Published, sold and distributed by:

now Publishers Inc.
PO Box 1024
Hanover, MA 02339
United States
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:

now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is

J. Xu, X. He and H. Li. *Deep Learning for Matching in Search and Recommendation*. Foundations and Trends[®] in Information Retrieval, vol. 14, no. 2–3, pp. 102–288, 2020.

ISBN: 978-1-68083-707-0

© 2020 J. Xu, X. He and H. Li

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: www.copyright.com

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; www.nowpublishers.com; sales@nowpublishers.com

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, www.nowpublishers.com; e-mail: sales@nowpublishers.com

Foundations and Trends[®] in Information Retrieval

Volume 14, Issue 2–3, 2020

Editorial Board

Editors-in-Chief

Maarten de Rijke

University of Amsterdam
and Ahold Delhaize
The Netherlands

Yiqun Liu

Tsinghua University
China

Diane Kelly

University of Tennessee
USA

Editors

Barbara Poblete

University of Chile

Claudia Hauff

Delft University of Technology

Doug Oard

University of Maryland

Ellen M. Voorhees

*National Institute of Standards and
Technology*

Fabrizio Sebastiani

Consiglio Nazionale delle Ricerche, Italy

Hang Li

Bytedance Technology

Helen Huang

University of Queensland

Isabelle Moulinier

Capital One

Jaap Kamps

University of Amsterdam

Jimmy Lin

University of Waterloo

Leif Azzopardi

University of Glasgow

Lynda Tamine

University of Toulouse

Mark Sanderson

RMIT University

Rodrygo Luis Teodoro Santos

Universidade Federal de Minas Gerais

Ruihua Song

Microsoft Xiaoice

Ryen White

Microsoft Research

Shane Culpepper

RMIT

Soumen Chakrabarti

Indian Institute of Technology

Xuanjing Huang

Fudan University

Zi Helen Huang

University of Queensland

Editorial Scope

Topics

Foundations and Trends® in Information Retrieval publishes survey and tutorial articles in the following topics:

- Applications of IR
- Architectures for IR
- Collaborative filtering and recommender systems
- Cross-lingual and multilingual IR
- Distributed IR and federated search
- Evaluation issues and test collections for IR
- Formal models and language models for IR
- IR on mobile platforms
- Indexing and retrieval of structured documents
- Information categorization and clustering
- Information extraction
- Information filtering and routing
- Metasearch, rank aggregation and data fusion
- Natural language processing for IR
- Performance issues for IR systems, including algorithms, data structures, optimization techniques, and scalability
- Question answering
- Summarization of single documents, multiple documents, and corpora
- Text mining
- Topic detection and tracking
- Usability, interactivity, and visualization issues in IR
- User modelling and user studies for IR
- Web search

Information for Librarians

Foundations and Trends® in Information Retrieval, 2020, Volume 14, 5 issues. ISSN paper version 1554-0669. ISSN online version 1554-0677. Also available as a combined paper and online subscription.

Contents

1	Introduction	6
1.1	Search and Recommendation	6
1.2	Unifying Search and Recommendation from Matching Viewpoint	8
1.3	Mismatching Challenge in Search	10
1.4	Mismatching Challenge in Recommendation	11
1.5	Recent Advances	12
1.6	About This Survey	13
2	Traditional Matching Models	16
2.1	Learning to Match	16
2.2	Matching Models in Search and Recommendation	22
2.3	Latent Space Models in Search	25
2.4	Latent Space Models in Recommendation	28
2.5	Further Reading	31
3	Deep Learning for Matching	33
3.1	Overview of Deep Learning	33
3.2	Overview of Deep Learning for Matching	49
4	Deep Matching Models in Search	55
4.1	Matching Based on Representation Learning	58

4.2	Matching Based on Matching Function Learning	75
4.3	Discussions and Further Reading	93
5	Deep Matching Models in Recommendation	102
5.1	Matching Based on Representation Learning	102
5.2	Matching Based on Matching Function Learning	134
5.3	Further Reading	145
6	Conclusion and Future Directions	150
6.1	Summary of the Survey	150
6.2	Matching in Other Tasks	151
6.3	Open Questions and Future Directions	152
	Acknowledgements	155
	References	156

Deep Learning for Matching in Search and Recommendation

Jun Xu¹, Xiangnan He² and Hang Li³

¹*Gaoling School of Artificial Intelligence, Renmin University of China, China; junxu@ruc.edu.cn*

²*School of Information Science and Technology, University of Science and Technology of China; hexn@ustc.edu.cn*

³*Bytedance AI Lab, China; lihang.lh@bytedance.com*

ABSTRACT

Matching is a key problem in both search and recommendation, which is to measure the relevance of a document to a query or the interest of a user to an item. Machine learning has been exploited to address the problem, which learns a matching function based on input representations and from labeled data, also referred to as “learning to match”. In recent years, efforts have been made to develop deep learning techniques for matching tasks in search and recommendation. With the availability of a large amount of data, powerful computational resources, and advanced deep learning techniques, deep learning for matching now becomes the state-of-the-art technology for search and recommendation. The key to the success of the deep learning approach is its strong ability in learning of representations and generalization of matching patterns from data (e.g., queries, documents, users, items, and contexts, particularly in their raw forms).

This survey gives a systematic and comprehensive introduction to the deep matching models for search and recommendation developed recently. It first gives a unified view of matching in search and recommendation. In this way, the solutions from the two fields can be compared under one framework. Then, the survey categorizes the current deep learning solutions into two types: methods of representation learning and methods of matching function learning. The fundamental problems, as well as the state-of-the-art solutions of query-document matching in search and user-item matching in recommendation, are described. The survey aims to help researchers from both search and recommendation communities to get in-depth understanding and insight into the spaces, stimulate more ideas and discussions, and promote developments of new technologies.

Matching is not limited to search and recommendation. Similar problems can be found in paraphrasing, question answering, image annotation, and many other applications. In general, the technologies introduced in the survey can be generalized into a more general task of matching between objects from two spaces.

List of Acronyms

PLS	Partial Least Square
RMLS	Regularized Matching in Latent Space
SSI	Supervised Semantic Indexing
BMF	Biased Matrix Factorization
FISM	Factored Item Similarity Model
FM	Factorization Machine
FFN	Feedforward Neural Network
MLP	Multilayer Perceptron
CNN	Convolutional Neural Networks
RNN	Recurrent Neural Networks
GAN	Generative Adversarial Network
AE	Autoencoders
DAE	Denoising Autoencoder
CBOW	Continuous Bag of Words
SG	Skip Gram
BERT	Bidirectional Encoder Representations from Transformers
DSSM	Deep Structured Semantic Models
CLSM	Convolutional Latent Semantic Model
CNTN	Convolutional Neural Tensor Network
LSTM-RNN	Recurrent Neural Networks with Long Short-Term Memory cells
NVSM	Neural Vector Space Model

SNRM Standalone Neural Ranking Model

ACMR Adversarial Cross Modal Retrieval

ARC-II Convolutional Matching Model II

DRMM Deep Relevance Matching Model

K-NRM Kernel Based Neural Ranking Model

DeepMF Deep Matrix Factorization

CDAE Collaborative Denoising Auto-Encoder

NAIS Neural Attentive Item Similarity

NARM Neural Attentive Recommendation Machine

DeepCoNN Deep Cooperative Neural Networks

NARRE Neural Attention Regression with Review-level Explanation

VBPR Visual Bayesian Personalized Ranking

CDL Comparative Deep Learning

ACF Attentive Collaborative Filtering

NGCF Neural Graph Collaborative Filtering

KGAT Knowledge Graph Attention Network

KPRN Knowledge Path Recurrent Network

NCF Neural Collaborative Filtering

ConvNCF Convolutional Neural Collaborative Filtering

GMF Generalized Matrix Factorization

NeuMF Neural Matrix Factorization

CML Collaborative Metric Learning

TransRec Translation-based Recommendation

LRML Latent Relational Metric Learning

NFM Neural Factorization Machine

AFM Attentional Factorization Machine

1

Introduction

1.1 Search and Recommendation

With the rapid growth of the internet, one of the fundamental problems in information science becomes even more critical today, that is, how to identify the information satisfying a user's need from a usually huge pool of information. The goal is to present the user only the information that is of interest and relevance, at the right time, place, and context. Nowadays, two types of information accessing paradigms, search and recommendation, are widely used in a great variety of scenarios.

In search, documents (e.g., web documents, Twitter posts, or E-commerce products) are first pre-processed and indexed in the search engine. After that, the search engine takes a query (a number of keywords) from the user. The query describes the user's information need. Relevant documents are retrieved from the index, matched with the query, and ranked according to their relevance to the query. For example, if a user is interested in news about quantum computing, the query "quantum computing" may be submitted to a search engine and get news articles about the topic will be returned.

Different from search, a recommendation system typically does not take a query. Instead, it analyzes the user’s profile (e.g., demographics and contexts) and historical interactions on items, and then makes recommendation on items to the user. The user features and item features are indexed and stored in the system in advance. The items are ranked according to the likelihood that the user is interested in them. For example, on a news website, when a user browses and clicks a new article, several news articles with similar topics or news articles that other users have clicked together with the current one may be shown.

Table 1.1 summarizes the differences between search and recommendation. The fundamental mechanism of search is “pull”, because users first make specific requests (i.e., submit queries) and then receive information. The fundamental mechanisms of recommendation is “push”, because users are provided information which they do not specifically request (e.g., submit queries). Here “beneficiary” means the people whose interests are to be met in the task. In a search engine, the results are typically created solely based on the user’s needs, and thus the beneficiary is the users. In a recommendation engine, the results usually need to satisfy both the users and providers, and thus the beneficiary is all of them. However, the distinction is becoming blurred recently. For example, some search engines mix search results with paid advertisements, which benefits both the users and the providers. As for “serendipity”, it means that conventional search focuses more on information that is clearly relevant. Conventional recommendation, on the other hand, is allowed to offer unexpected but useful information.

Table 1.1: Information-providing mechanisms of search and recommendation

	Search	Recommendation
Query available	Yes	No
Delivery model	Pull	Push
Beneficiary	User	User and provider
Serendipity	No	Yes




Figure 1.1: Unified view of matching in search and recommendation.

1.2 Unifying Search and Recommendation from Matching Viewpoint

Garcia-Molina *et al.* (2011) pointed out that the fundamental problems in search and recommendation are to identify information objects satisfying users' information needs. It is also indicated that search (information retrieval) and recommendation (information filtering) are the two sides of the same coin, having strong connections and similarities (Belkin and Croft, 1992). Figure 1.1 illustrates the unified matching view of search and recommendation. The goal in common is to present to the users the information they need.

Search is a retrieval task, which aims to retrieve the documents that are relevant to the query. In contrast, recommendation is a filtering task, which aims to filter out the items that are of interest to the user (Adomavicius and Tuzhilin, 2005). As such, search can be considered as conducting matching between queries and documents, and recommendation can be considered as conducting matching between users and items. More formally, both the matching in search and recommendation can be considered as constructing a matching model $f: \mathcal{X} \times \mathcal{Y} \mapsto \mathcal{R}$ which calculates the matching degree between two input objects x and y , where \mathcal{X} and \mathcal{Y} denote two object spaces. \mathcal{X} and \mathcal{Y} are the spaces of queries and documents in search, or the spaces of users and items in recommendation.

Under the unified matching view in Figure 1.1, we use the term information objects to denote the documents/items to retrieve/recommend, and use information needs to denote the queries/users in the respective task. By unifying the two tasks under the same view of matching and

comparably reviewing existing techniques, we can provide deeper insights and more powerful solutions to the problems. Moreover, unifying the two tasks also has practical and theoretical implications.

Search and recommendation have already been combined in some practical applications. For example, at some E-commerce sites, when the user submits a query, a ranking list of products are presented based on not only relevance (query-product matching) but also user interest (user-product matching). In some lifestyle apps, when the user searches for restaurants, the results are returned based on both relevance (query-restaurant matching) and user interest (user-restaurant matching). There is a clear trend that search and recommendation will be integrated into a single system at certain scenarios to meet users' needs better, where matching plays an essential role.

Search and recommendation already have many shared technologies because of their similarities in matching. Some search problems can be solved by using recommendation techniques (Zamani *et al.*, 2016), and vice versa (Costa and Roda, 2011), on the basis of matching. With the use of deep learning technologies, the matching models for search and recommendation bear even more resemblance in architecture and methodology, as reflected in the techniques: embedding the inputs (queries, users, documents, and items) as distributed representations, combining neural network components to represent the matching function, and training the model parameters in an end-to-end manner. Moreover, search and recommendation can be jointly modeled and optimized if they share the same set of information objects (as in the above examples of E-commerce sites and lifestyle apps) (Schedl *et al.*, 2018; Zamani and Croft, 2018a, 2020). Therefore, in order to develop more advanced ones, it is necessary and advantageous to take a unified matching view to analyze and compare existing search and recommendation technologies.

The matching tasks in search and in recommendation face different challenges in practice. The underlying problem is essentially the same, however, that is, the mismatch challenge. Next, we introduce the key challenges of the two tasks, respectively.

1.3 Mismatching Challenge in Search

In search, queries and documents (usually their titles) are taken as texts. The relevance of a document to a query is mainly represented by the matching degree between the two. The document is considered relevant to the query if the matching degree is high. Natural language understanding by computer is still challenging, and thus the calculation of matching degree is still limited to the text level but not at the semantic level. A high match degree at the text level does not necessarily mean high relevance at the semantic level, and vice versa. Moreover, queries are issued by users, while documents are compiled by editors. Due to the ambiguity of natural language, users and editors are likely to use different language styles and expressions for presenting the same concepts or topics. As a result, the search system may suffer from the so-called query-document mismatch problem. Specifically, when the users of a search engine and the editors of the documents use different texts to describe the same concept (e.g., “ny times” vs. “new york times”), query-document mismatch may occur. This is still one of the main challenges for search. Moving to the cross-modal IR (e.g., using text queries to retrieve image documents), the query-document mismatch problem becomes even more severe, because different modalities have different types of representations. In cross-modal retrieval, one major challenge is how to construct a matching function that can bridge the “heterogeneity gap” amongst the modalities.

To address the query-document mismatch challenge, methods have been proposed to perform matching at the semantic level, referred to as semantic matching. The key idea in the solutions is either to perform more query and document understanding to better represent the meanings of the query and document, or to construct more powerful matching functions that can bridge the semantic gap between the query and document. Both traditional machine learning approaches (Li and Xu, 2014) and deep learning approaches (Guo *et al.*, 2019b; Mitra and Craswell, 2018; Onal *et al.*, 2018) have been developed for semantic matching.

1.4 Mismatching Challenge in Recommendation

The mismatching problem is even more severe in recommendation. In search, queries and documents consist of terms in the same language,¹ making it at least meaningful to conduct direct matching on their terms. In recommendation, however, users and items are usually represented by different types of features, for example, the features of users can be the user ID, age, income level, and recent behaviors, while the features for items can be the item ID, category, price, and brand name. Since the features of users and items are from the spaces of different semantics, the naive approaches based on the matching of superficial features do not work for recommendation. More challengingly, the items can be described by multi-modal features, e.g., images of clothing products and cover images of movies, which could play a pivotal role in affecting the decision-making of users. In such visually-aware scenarios, we need to consider the cross-modal matching between users and multi-modal content.

To address the mismatching challenge in recommendation, the collaborative filtering principle has been proposed (Shi *et al.*, 2014). Collaborative Filtering (CF), which works as the fundamental basis of almost all personalized recommender systems, assumes that a user may like (consume) the items that are liked (consumed) by the similar users, for which the similarity is judged from the historical interactions (Sarwar *et al.*, 2001). However, directly evaluating the similarity between users (items) suffers from the sparsity issue, since a user only consumed a few items in the whole item space. A typical assumption to address the sparsity issue is that the user-item interaction matrix is low-rank, which thus can be estimated from low-dimensional user (and item) latent feature matrix. Then the user (item) similarity can be more reliably reflected in the latent feature matrix. This leads to the effectiveness of matrix factorization for collaborative filtering (Koren *et al.*, 2009; Rendle *et al.*, 2009), which becomes a strong CF method and an essential design for many recommender models. Besides matrix factorization, many other types of CF methods have been developed

¹Here we do not consider cross-language information retrieval.

like neural network-based methods (He *et al.*, 2017c; Liang *et al.*, 2018) and graph-based methods (Wang *et al.*, 2019b; Ying *et al.*, 2018).

To leverage the various side information beyond the interaction matrix, such as user profiles, item attributes, and the current contexts, many generic recommender models that follow the standard supervised learning paradigm have been proposed. These models can be used in the (re-)ranking stage of a recommendation engine, e.g., by predicting the click-through rate (CTR) of an item. A representative model is factorization machine (FM) (Rendle, 2010), which extends the low-rank assumption of matrix factorization to model feature interactions. Since the expressiveness of FM is limited by its linearity and second-order interaction modeling, many later efforts complement it with neural networks for nonlinear and higher-order interaction modeling (He and Chua, 2017; Lian *et al.*, 2018; Zhou *et al.*, 2018). These neural network models have now been intensively used in industrial applications. Batmaz *et al.* (2019) and Zhang *et al.* (2019) reviewed deep learning methods for recommendation systems.

Please note that though query-document matching and user-item matching are critical for search engines and recommendation systems, these systems also include other important components. Besides matching, web search engines also include crawling, indexing, document understanding, query understanding, and ranking, etc. Recommendation systems also include components such as user modeling (profiling), indexing, caching, diversity controlling, and online exploration, etc.

1.5 Recent Advances

Though traditional machine learning was successful for matching in search and recommendation, recent advances in deep learning have brought even more significant progress to the area with a large number of deep matching models proposed. The power of deep learning models lies in the ability to learn distributed representations from the raw data (e.g., text) for the matching problem, to avoid many limitations of hand-crafted features, and to learn the representations and matching networks in an end-to-end fashion. Moreover, deep neural networks have sufficient capacity to model complicated matching tasks. They have

the flexibility of extending to cross-modal matching naturally, where the common semantic space is learned to represent data of different modalities universally. All these characteristics are helpful in handling the complexity of search and recommendation.

In search, the mismatch between query and document is more effectively addressed by deep neural networks, including the feed-forward neural networks (FFNs), convolutional neural networks (CNNs), and Recurrent neural networks (RNNs), because they have stronger capabilities in representation learning and matching function learning. Most notably, Bidirectional Encoder Representations from Transformers (BERT) has significantly enhanced the accuracy of matching in search and stands out as the state-of-the-art technique now.

In recommendation, recent focus has shifted from behavior-centric collaborative filtering to information-rich user-item matching as in sequential, context-aware, and knowledge graph enhanced recommendations, which are all practical scenario-driven. In terms of techniques, graph neural networks (GNNs) become an emerging tool for representation learning (Wang *et al.*, 2019a,b), because recommendation data can be naturally organized in a heterogeneous graph and GNNs have the capability to exploit such data. To handle user behavior sequence data, self-attention and BERT are also adopted, which demonstrates promising results in sequential recommendation (Sun *et al.*, 2019; Yuan *et al.*, 2020).

1.6 About This Survey

This survey focuses on the fundamental problems of matching in search and recommendation. State-of-the-art matching solutions using deep learning are described. A unified view of search and recommendation from matching is provided. The ideas and solutions explained may motivate industrial practitioners to turn the research results into products. The methods and the discussions may help academic researchers to develop new approaches. The unified view may bring researchers in the search and the recommendation communities together and inspire them to explore new directions.

The survey is organized as follows: Section 2 describes the traditional machine learning approaches to matching for search and recommendation; Section 3 gives a general formulation of deep matching methods; Section 4 and Section 5 describe the details of the deep learning approaches to search and recommendation respectively. Each section includes the representation learning-based approaches and matching function learning-based approaches; Section 6 summarizes the survey and discusses open problems. Sections 2, 3, 4, and 5 are self-contained, and the readers can choose to read on the basis of their interest and need.

Note that deep learning for search and recommendation is a very hot topic of research. As such, this survey does not try to cover all related works in the fields of information retrieval and recommender systems. Instead, we discuss the most representative approaches of the two fields from the viewpoint of matching, aiming to summarize their key ideas which are general and essential. In particular, this survey covers the representative work before 2019.

Several previous FnTIR issues have given detailed introductions to related topics. One issue (Li and Xu, 2014) introduces the traditional machine learning approaches to the semantic matching problem, particularly in web search. Our survey in this issue is very different from it in the sense that (1) it focuses on the newly developed deep learning methods, and (2) it considers both search and recommendation. Mitra and Craswell (2018) conducted a comprehensive survey on deep neural networks for information retrieval, referred to as Neural IR. Bast *et al.* (2016) carries out a survey on the techniques and systems of semantic search, which means search with keyword queries, structured queries, and natural language queries, to documents, knowledge bases, and their combinations.

Several surveys and tutorials have been made on deep learning for information retrieval and recommendation. For example, Onal *et al.* (2018) have explained neural models for ad-hoc retrieval, query understanding, question answering, sponsored search, and similar item retrieval. Zhang *et al.* (2019) reviews deep learning-based recommendation methods according to the taxonomy of deep learning techniques, e.g., MLP, CNN, RNN, autoencoder-based, and so on. Other related

surveys and tutorials include Kenter *et al.* (2017), Li and Lu (2016), Guo *et al.* (2019b), Batmaz *et al.* (2019), and Zhang *et al.* (2019). They all quite differ from this survey, which summarizes existing work from the perspective of matching (e.g., input representations and the way for matching).

This survey focuses on state-of-the-art matching techniques using deep learning. We expect that the readers have a certain knowledge of search and recommendation. Those who are not familiar with the areas may consult existing materials (e.g., Adomavicius and Tuzhilin, 2005; Croft *et al.*, 2009; Li and Xu, 2014; Liu, 2009; Ricci *et al.*, 2015). We also assume that the readers have sufficient knowledge of machine learning, particularly deep learning.

References

- Adomavicius, G. and A. Tuzhilin (2005). “Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions”. *IEEE Transactions on Knowledge and Data Engineering*. 17(6): 734–749.
- Ai, Q., K. Bi, J. Guo, and W. B. Croft (2018). “Learning a deep listwise context model for ranking refinement”. In: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval. SIGIR '18*. Ann Arbor, MI, USA: Association for Computing Machinery. 135–144.
- Andrew, G., R. Arora, J. Bilmes, and K. Livescu (2013). “Deep canonical correlation analysis”. In: *Proceedings of the 30th International Conference on International Conference on Machine Learning – Volume 28. ICML'13*. Atlanta, GA, USA: JMLR.org. III-1247–III-1255. URL: <http://dl.acm.org/citation.cfm?id=3042817.3043076>.
- Ba, J. L., J. R. Kiros, and G. E. Hinton (2016). “Layer normalization”. *CoRR*. abs/1607.06450. arXiv: [1607.06450](https://arxiv.org/abs/1607.06450).
- Bahdanau, D., K. Cho, and Y. Bengio (2015). “Neural machine translation by jointly learning to align and translate”. In: *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA*. URL: <http://arxiv.org/abs/1409.0473>.

- Bai, B., J. Weston, D. Grangier, R. Collobert, K. Sadamasu, Y. Qi, O. Chapelle, and K. Weinberger (2009). “Supervised semantic indexing”. In: *Proceedings of the 18th ACM Conference on Information and Knowledge Management. CIKM '09*. Hong Kong, China: ACM. 187–196.
- Bai, B., J. Weston, D. Grangier, R. Collobert, K. Sadamasu, Y. Qi, O. Chapelle, and K. Weinberger (2010). “Learning to rank with (a lot of) word features”. *Information Retrieval*. 13(3): 291–314.
- Bai, T., J.-R. Wen, J. Zhang, and W. X. Zhao (2017). “A neural collaborative filtering model with interaction-based neighborhood”. In: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management. CIKM '17*. Singapore, Singapore: ACM. 1979–1982.
- Balaneshin-Kordan, S. and A. Kotov (2018). “Deep neural architecture for multi-modal retrieval based on joint embedding space for text and images”. In: *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining. WSDM '18*. Marina Del Rey, CA, USA: ACM. 28–36.
- Baltrunas, L., K. Church, A. Karatzoglou, and N. Oliver (2015). “Frappe: Understanding the usage and perception of mobile app recommendations in-the-wild”. *CoRR*. abs/1505.03014. arXiv: [1505.03014](https://arxiv.org/abs/1505.03014).
- Bast, H., B. Björn, and E. Haussmann (2016). “Semantic search on text and knowledge bases”. *Foundations and Trends in Information Retrieval*. 10(2–3): 119–271.
- Batmaz, Z., A. Yurekli, A. Bilge, and C. Kaleli (2019). “A review on deep learning for recommender systems: Challenges and remedies”. *Artificial Intelligence Review*. 52(1): 1–37.
- Belkin, N. J. and W. B. Croft (1992). “Information filtering and information retrieval: Two sides of the same coin?” *Communications of the ACM*. 35(12): 29–38.
- Bello, I., S. Kulkarni, S. Jain, C. Boutilier, E. H. Chi, E. Eban, X. Luo, A. Mackey, and O. Meshi (2018). “Seq2Slate: Re-ranking and slate optimization with RNNs”. In: *Proceedings of the Workshop on Negative Dependence in Machine Learning at the 36th International Conference on Machine Learning*. Long Beach, CA. PMLR 97, 2019.

- Bendersky, M., W. B. Croft, and D. A. Smith (2011). “Joint annotation of search queries”. In: *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies – Volume 1. HLT ’11*. Portland, OR, USA: Association for Computational Linguistics. 102–111. URL: <http://dl.acm.org/citation.cfm?id=2002472.2002486>.
- Berg, R. van den, T. N. Kipf, and M. Welling (2017). “Graph convolutional matrix completion”. *CoRR*. abs/1706.02263. arXiv: [1706.02263](https://arxiv.org/abs/1706.02263).
- Berger, A. and J. Lafferty (1999). “Information retrieval as statistical translation”. In: *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR ’99*. Berkeley, CA, USA: ACM. 222–229.
- Bergsma, S. and Q. I. Wang (2007). “Learning noun phrase query segmentation”. In: *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*. Prague, Czech Republic: Association for Computational Linguistics. 819–826. URL: <https://www.aclweb.org/anthology/D07-1086>.
- Beutel, A., P. Covington, S. Jain, C. Xu, J. Li, V. Gatto, and E. H. Chi (2018). “Latent cross: Making use of context in recurrent recommender systems”. In: *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining. WSDM ’18*. Marina Del Rey, CA, USA: ACM. 46–54.
- Bowman, S. R., G. Angeli, C. Potts, and C. D. Manning (2015). “A large annotated corpus for learning natural language inference”. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal: Association for Computational Linguistics. 632–642. URL: <https://www.aclweb.org/anthology/D15-1075>.
- Brill, E. and R. C. Moore (2000). “An improved error model for noisy channel spelling correction”. In: *Proceedings of the 38th Annual Meeting on Association for Computational Linguistics. ACL ’00*. Hong Kong: Association for Computational Linguistics. 286–293.

- Burges, C. J. (2010). “From RankNet to LambdaRank to LambdaMART: An overview”. *Technical report*. MSR-TR-2010-82. <https://www.microsoft.com/en-us/research/publication/from-ranknet-to-lambda-rank-to-lambdamart-an-overview/>.
- Cao, Y., J. Xu, T.-Y. Liu, H. Li, Y. Huang, and H.-W. Hon (2006). “Adapting ranking SVM to document retrieval”. In: *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '06*. Seattle, Washington, DC, USA: ACM. 186–193.
- Chen, J., H. Zhang, X. He, L. Nie, W. Liu, and T.-S. Chua (2017a). “Attentive collaborative filtering: Multimedia recommendation with item- and component-level attention”. In: *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '17*. Shinjuku, Tokyo, Japan: ACM. 335–344.
- Chen, Q., X. Zhu, Z.-H. Ling, S. Wei, H. Jiang, and D. Inkpen (2017b). “Enhanced LSTM for natural language inference”. In: *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vancouver, Canada: Association for Computational Linguistics. 1657–1668. URL: <https://www.aclweb.org/anthology/P17-1152>.
- Chen, C., M. Zhang, Y. Liu, and S. Ma (2018a). “Neural attentional rating regression with review-level explanations”. In: *Proceedings of the 2018 World Wide Web Conference. WWW '18*. Lyon, France. 1583–1592.
- Chen, H., F. X. Han, D. Niu, D. Liu, K. Lai, C. Wu, and Y. Xu (2018b). “MIX: Multi-channel information crossing for text matching”. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. KDD '18*. London, UK: ACM. 110–119.
- Cheng, H.-T., L. Koc, J. Harmsen, T. Shaked, T. Chandra, H. Aradhye, G. Anderson, G. Corrado, W. Chai, M. Ispir, R. Anil, Z. Haque, L. Hong, V. Jain, X. Liu, and H. Shah (2016). “Wide & deep learning for recommender systems”. In: *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems. DLRS 2016*. Boston, MA, USA: ACM. 7–10.

- Cheng, Z., Y. Ding, X. He, L. Zhu, X. Song, and M. S. Kankanhalli (2018). “A3NCF: An adaptive aspect attention model for rating prediction”. In: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*. 3748–3754.
- Cohen, D., L. Yang, and W. B. Croft (2018). “WikiPassageQA: A benchmark collection for research on non-factoid answer passage retrieval”. In: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval. SIGIR '18*. Ann Arbor, MI, USA: ACM. 1165–1168.
- Costa, A. and F. Roda (2011). “Recommender systems by means of information retrieval”. In: *Proceedings of the International Conference on Web Intelligence, Mining and Semantics. WIMS '11*. Sogndal, Norway: ACM. 57:1–57:5.
- Covington, P., J. Adams, and E. Sargin (2016). “Deep neural networks for youtube recommendations”. In: *Proceedings of the 10th ACM Conference on Recommender Systems*. 191–198.
- Croft, W. B., D. Metzler, and T. Strohman (2009). *Search Engines: Information Retrieval in Practice*. 1st Edn. USA: Addison-Wesley Publishing Company. I–XXV, 1–524.
- Dai, Z., C. Xiong, J. Callan, and Z. Liu (2018). “Convolutional neural networks for soft-matching N-grams in ad-hoc search”. In: *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining. WSDM '18*. Marina Del Rey, CA, USA: ACM. 126–134.
- Dehghani, M., H. Zamani, A. Severyn, J. Kamps, and W. B. Croft (2017). “Neural ranking models with weak supervision”. In: *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '17*. Shinjuku, Tokyo, Japan: Association for Computing Machinery. 65–74.
- Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova (2019). “BERT: Pre-training of deep bidirectional transformers for language understanding”. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Minneapolis, Minnesota: Association for Computational Linguistics. 4171–4186. URL: <https://www.aclweb.org/anthology/N19-1423>.

- Dolan, B. and C. Brockett (2005). “Automatically constructing a corpus of sentential paraphrases”. In: *Third International Workshop on Paraphrasing (IWP2005)*. Asia Federation of Natural Language Processing. URL: <https://www.microsoft.com/en-us/research/publication/automatically-constructing-a-corpus-of-sentential-paraphrases/>.
- Eisenschat, A. and L. Wolf (2017). “Linking image and text with 2-way nets”. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1855–1865.
- Eksombatchai, C., P. Jindal, J. Z. Liu, Y. Liu, R. Sharma, C. Sugnet, M. Ulrich, and J. Leskovec (2018). “Pixie: A system for recommending 3+ Billion items to 200+ Million users in real-time”. In: *Proceedings of the 2018 World Wide Web Conference on World Wide Web, WWW 2018, Lyon, France*. 1775–1784.
- Elkahky, A. M., Y. Song, and X. He (2015). “A multi-view deep learning approach for cross domain user modeling in recommendation systems”. In: *Proceedings of the 24th International Conference on World Wide Web*. Republic and Canton of Geneva, CHE. 278–288.
- Fan, Y., J. Guo, Y. Lan, J. Xu, C. Zhai, and X. Cheng (2018). “Modeling diverse relevance patterns in ad-hoc retrieval”. In: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval. SIGIR '18*. Ann Arbor, MI, USA: ACM. 375–384.
- Fan, W., Y. Ma, Q. Li, Y. He, E. Zhao, J. Tang, and D. Yin (2019). “Graph neural networks for social recommendation”. In: *The World Wide Web Conference. WWW '19*. San Francisco, CA, USA: Association for Computing Machinery. 417–426.
- Gao, J., J.-Y. Nie, G. Wu, and G. Cao (2004). “Dependence language model for information retrieval”. In: *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '04*. Sheffield, UK: ACM. 170–177.
- Gao, L., H. Yang, J. Wu, C. Zhou, W. Lu, and Y. Hu (2018). “Recommendation with multi-source heterogeneous information”. In: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization. 3378–3384.

- Gao, C., X. Chen, F. Feng, K. Zhao, X. He, Y. Li, and D. Jin (2019a). “Cross-domain recommendation without sharing user-relevant data”. In: *The World Wide Web Conference. WWW’19*. New York, NY, USA: Association for Computing Machinery. 491–502.
- Gao, C., X. He, D. Gan, X. Chen, F. Feng, Y. Li, and T.-S. Chua (2019b). “Neural multi-task recommendation from multi-behavior data”. In: *Proceedings of IEEE 35th International Conference on Data Engineering (ICDE)*, Macao, China. 1554–1557.
- Garcia-Molina, H., G. Koutrika, and A. Parameswaran (2011). “Information seeking: Convergence of search, recommendations, and advertising”. *Communications of the ACM*. 54(11): 121–130.
- Gong, Y., H. Luo, and J. Zhang (2018). “Natural language inference over interaction space”. In: *6th International Conference on Learning Representations, ICLR 2018*.
- Goodfellow, I., Y. Bengio, and A. Courville (2016). *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- Graves, A., S. Fernández, and J. Schmidhuber (2007). “Multi-dimensional recurrent neural networks”. In: *Artificial Neural Networks – ICANN 2007*. Berlin, Heidelberg: Springer Berlin Heidelberg. 549–558.
- Guo, J., G. Xu, H. Li, and X. Cheng (2008). “A unified and discriminative model for query refinement”. In: *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR ’08*. Singapore, Singapore: ACM. 379–386.
- Guo, J., Y. Fan, Q. Ai, and W. B. Croft (2016). “A deep relevance matching model for ad-hoc retrieval”. In: *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management. CIKM ’16*. Indianapolis, IN, USA: ACM. 55–64.
- Guo, H., R. Tang, Y. Ye, Z. Li, and X. He (2017). “DeepFM: A factorization-machine based neural network for CTR prediction”. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence. IJCAI’17*. Melbourne, Australia: AAAI Press. 1725–1731.
- Guo, Y., Z. Cheng, L. Nie, X. Xu, and M. S. Kankanhalli (2018). “Multi-modal preference modeling for product search”. In: *Proceedings of the 26th ACM International Conference on Multimedia*. 1865–1873.

- Guo, J., Y. Fan, X. Ji, and X. Cheng (2019a). “MatchZoo: A learning, practicing, and developing system for neural text matching”. In: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR’19*. Paris, France: ACM. 1297–1300.
- Guo, J., Y. Fan, L. Pang, L. Yang, Q. Ai, H. Zamani, C. Wu, W. B. Croft, and X. Cheng (2019b). “A deep look into neural ranking models for information retrieval”. *Information Processing and Management: Information Retrieval*. URL: <https://doi.org/10.1016/j.ipm.2019.102067>.
- Gysel, C. V., M. de Rijke, and E. Kanoulas (2018). “Neural vector spaces for unsupervised information retrieval”. *ACM Transactions on Information Systems*. 36(4): 38:1–38:25.
- Haddad, D. and J. Ghosh (2019). “Learning more from less: Towards strengthening weak supervision for ad-hoc retrieval”. In: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR’19*. Paris, France: Association for Computing Machinery. 857–860.
- Hardoon, D. R. and J. Shawe-Taylor (2003). “KCCA for different level precision in content-based image retrieval”. Event Dates: 22–24 September 2004. URL: <https://eprints.soton.ac.uk/259596/>.
- Hardoon, D. R., S. R. Szedmak, and J. R. Shawe-Taylor (2004). “Canonical correlation analysis: An overview with application to learning methods”. *Neural Computation*. 16(12): 2639–2664.
- He, R. and J. McAuley (2016a). “VBPR: Visual Bayesian personalized ranking from implicit feedback”. In: *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence. AAAI’16*. Phoenix, AZ: AAAI Press. 144–150. URL: <http://dl.acm.org/citation.cfm?id=3015812.3015834>.
- He, R. and J. J. McAuley (2016b). “Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering”. In: *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada*. 507–517.
- He, X. and T.-S. Chua (2017). “Neural factorization machines for sparse predictive analytics”. In: *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR ’17*. Shinjuku, Tokyo, Japan: ACM. 355–364.

- He, X., M.-Y. Kan, P. Xie, and X. Chen (2014). “Comment-based multi-view clustering of web 2.0 items”. In: *Proceedings of the 23rd International Conference on World Wide Web. WWW '14*. Seoul, Korea: ACM. 771–782.
- He, K., X. Zhang, S. Ren, and J. Sun (2016a). “Deep residual learning for image recognition”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 770–778.
- He, X., H. Zhang, M.-Y. Kan, and T.-S. Chua (2016b). “Fast matrix factorization for online recommendation with implicit feedback”. In: *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '16*. Pisa, Italy: ACM. 549–558.
- He, R., W.-C. Kang, and J. McAuley (2017a). “Translation-based recommendation”. In: *Proceedings of the Eleventh ACM Conference on Recommender Systems. RecSys '17*. Como, Italy: ACM. 161–169.
- He, X., M. Gao, M.-Y. Kan, and D. Wang (2017b). “BiRank: Towards ranking on bipartite graphs”. *IEEE Transactions on Knowledge and Data Engineering*. 29(1): 57–71.
- He, X., Z. He, J. Song, Z. Liu, Y. Jiang, and T. Chua (2018a). “NAIS: Neural attentive item similarity model for recommendation”. *IEEE Transactions on Knowledge and Data Engineering*. 30(12): 2354–2366.
- He, X., X. Du, X. Wang, F. Tian, J. Tang, and T.-S. Chua (2018b). “Outer product-based neural collaborative filtering”. In: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization. 2227–2233.
- He, X., L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua (2017c). “Neural collaborative filtering”. In: *Proceedings of the 26th International Conference on World Wide Web. WWW '17*. Perth, Australia. 173–182.
- He, X., K. Deng, X. Wang, Y. Li, Y. Zhang, and M. Wang (2020). “LightGCN: Simplifying and powering graph convolution network for recommendation”. In: *The 43rd International ACM SIGIR Conference on Research & Development in Information Retrieval. SIGIR '20*. New York, NY, USA.

- Hidasi, B., A. Karatzoglou, L. Baltrunas, and D. Tikk (2016). “Session-based recommendations with recurrent neural networks”. In: *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico*. URL: <http://arxiv.org/abs/1511.06939>.
- Hinton, G. E. and R. R. Salakhutdinov (2006). “Reducing the dimensionality of data with neural networks”. *Science*. 313(5786): 504–507. URL: <https://science.sciencemag.org/content/313/5786/504>.
- Hofmann, T. (1999). “Probabilistic latent semantic indexing”. In: *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '99*. Berkeley, CA, USA: ACM. 50–57.
- Hornik, K. (1991). “Approximation capabilities of multilayer feedforward networks”. *Neural Networks*. 4(2): 251–257. URL: <http://www.sciencedirect.com/science/article/pii/089360809190009T>.
- Hsieh, C.-K., L. Yang, Y. Cui, T.-Y. Lin, S. Belongie, and D. Estrin (2017). “Collaborative metric learning”. In: *Proceedings of the 26th International Conference on World Wide Web. WWW '17*. Perth, Australia. 193–201.
- Hu, B., Z. Lu, H. Li, and Q. Chen (2014). “Convolutional neural network architectures for matching natural language sentences”. In: *Advances in Neural Information Processing Systems 27*. Curran Associates, Inc. 2042–2050. URL: <http://papers.nips.cc/paper/5550-convolutional-neural-network-architectures-for-matching-natural-language-sentences.pdf>.
- Huang, P.-S., X. He, J. Gao, L. Deng, A. Acero, and L. Heck (2013). “Learning deep structured semantic models for web search using clickthrough data”. In: *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management. CIKM '13*. San Francisco, CA, USA: ACM. 2333–2338.
- Huang, J., S. Yao, C. Lyu, and D. Ji (2017). “Multi-granularity neural sentence model for measuring short text similarity”. In: *Database Systems for Advanced Applications*. Cham: Springer International Publishing. 439–455.
- Huang, Y., Q. Wu, W. Wang, and L. Wang (2018). “Image and sentence matching via semantic concepts and order learning”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*: 1–1.

- Hui, K., A. Yates, K. Berberich, and G. de Melo (2017). “PACRR: A position-aware neural IR model for relevance matching”. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Copenhagen, Denmark: Association for Computational Linguistics. 1049–1058. URL: <http://aclweb.org/anthology/D17-1110>.
- Hui, K., A. Yates, K. Berberich, and G. de Melo (2018). “Co-PACRR: A context-aware neural IR model for ad-hoc retrieval”. In: *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining. WSDM '18*. Marina Del Rey, CA, USA: ACM. 279–287.
- Jiang, J.-Y., M. Zhang, C. Li, M. Bendersky, N. Golbandi, and M. Najork (2019a). “Semantic text matching for long-form documents”. In: *The World Wide Web Conference. WWW '19*. San Francisco, CA, USA: Association for Computing Machinery. 795–806.
- Jiang, R., S. Gowal, Y. Qian, T. A. Mann, and D. J. Rezende (2019b). “Beyond greedy ranking: Slate optimization via list-CVAE”. In: *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA*. OpenReview.net. URL: <https://openreview.net/forum?id=r1xX42R5Fm>.
- Joachims, T. (2002). “Optimizing search engines using clickthrough data”. In: *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '02*. Edmonton, Alberta, Canada: ACM. 133–142.
- Joachims, T., A. Swaminathan, and T. Schnabel (2017). “Unbiased learning-to-rank with biased feedback”. In: *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining. WSDM '17*. Cambridge, UK. 781–789.
- Kabbur, S., X. Ning, and G. Karypis (2013). “FISM: Factored item similarity models for top-N recommender systems”. In: *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '13*. Chicago, IL, USA: ACM. 659–667.
- Kang, W. and J. J. McAuley (2018). “Self-attentive sequential recommendation”. In: *IEEE International Conference on Data Mining (ICDM), Singapore*. 197–206.

- Karatzoglou, A., X. Amatriain, L. Baltrunas, and N. Oliver (2010). “Multiverse recommendation: N-dimensional tensor factorization for context-aware collaborative filtering”. In: *Proceedings of the Fourth ACM Conference on Recommender Systems. RecSys '10*. Barcelona, Spain: ACM. 79–86.
- Karpathy, A. and F. Li (2015). “Deep visual-semantic alignments for generating image descriptions”. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA*. IEEE Computer Society. 3128–3137.
- Karpathy, A., A. Joulin, and L. Fei-Fei (2014). “Deep fragment embeddings for bidirectional image sentence mapping”. In: *Proceedings of the 27th International Conference on Neural Information Processing Systems – Volume 2. NIPS'14*. Montreal, Canada: MIT Press. 1889–1897. URL: <http://dl.acm.org/citation.cfm?id=2969033.2969038>.
- Kenter, T., A. Borisov, C. Van Gysel, M. Dehghani, M. de Rijke, and B. Mitra (2017). “Neural networks for information retrieval”. In: *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '17*. Shinjuku, Tokyo, Japan: ACM. 1403–1406.
- Kingma, D. P. and M. Welling (2014). “Auto-encoding variational Bayes”. In: *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings*. URL: <http://arxiv.org/abs/1312.6114>.
- Koren, Y. (2008). “Factorization meets the neighborhood: A multifaceted collaborative filtering model”. In: *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '08*. Las Vegas, NV, USA: ACM. 426–434.
- Koren, Y., R. Bell, and C. Volinsky (2009). “Matrix factorization techniques for recommender systems”. *Computer*. 42(8): 30–37.
- Kulis, B. (2013). “Metric learning: A survey”. *Foundations and Trends in Machine Learning*. 5(4): 287–364.
- Le, Q. and T. Mikolov (2014). “Distributed representations of sentences and documents”. In: *Proceedings of the 31st International Conference on International Conference on Machine Learning – Volume 32. ICML'14*. Beijing, China: JMLR.org. II–1188–II–1196.

- Lei, C., D. Liu, W. Li, Z. Zha, and H. Li (2016). “Comparative deep learning of hybrid representations for image recommendations”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA*. 2545–2553.
- Lei, W., X. He, Y. Miao, Q. Wu, R. Hong, M.-Y. Kan, and T.-S. Chua (2020). “Estimation-action-reflection: Towards deep interaction between conversational and recommender systems”. In: *Proceedings of the 13th ACM International Conference on Web Search and Data Mining. WSDM '20*. New York, NY, USA: ACM.
- Li, H. (2011). “Learning to rank for information retrieval and natural language processing”. *Synthesis Lectures on Human Language Technologies*. 4(1): 1–113.
- Li, H. and J. Xu (2014). “Semantic matching in search”. *Foundations and Trends in Information Retrieval*. 7(5): 343–469.
- Li, H. and Z. Lu (2016). “Deep learning for information retrieval”. In: *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '16*. Pisa, Italy: ACM. 1203–1206.
- Li, L., W. Chu, J. Langford, and R. E. Schapire (2010). “A contextual-bandit approach to personalized news article recommendation”. In: *Proceedings of the 19th International Conference on World Wide Web. WWW '10*. Raleigh, NC, USA: Association for Computing Machinery. 661–670.
- Li, S., J. Kawale, and Y. Fu (2015). “Deep collaborative filtering via marginalized denoising auto-encoder”. In: *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management. CIKM '15*. Melbourne, Australia: ACM. 811–820.
- Li, J., P. Ren, Z. Chen, Z. Ren, T. Lian, and J. Ma (2017). “Neural attentive session-based recommendation”. In: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management. CIKM '17*. Singapore, Singapore: ACM. 1419–1428.
- Li, C., C. Quan, L. Peng, Y. Qi, Y. Deng, and L. Wu (2019). “A capsule network for recommendation and explaining what you like and dislike”. In: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21–25, 2019*. 275–284.

- Lian, J., X. Zhou, F. Zhang, Z. Chen, X. Xie, and G. Sun (2018). “xDeepFM: Combining explicit and implicit feature interactions for recommender systems”. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. KDD '18*. London, UK: ACM. 1754–1763.
- Liang, D., L. Charlin, J. McInerney, and D. M. Blei (2016). “Modeling user exposure in recommendation”. In: *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11–15, 2016*. 951–961.
- Liang, D., R. G. Krishnan, M. D. Hoffman, and T. Jebara (2018). “Variational autoencoders for collaborative filtering”. In: *Proceedings of the 2018 World Wide Web Conference. WWW '18*. Lyon, France: International World Wide Web Conferences Steering Committee. 689–698.
- Liu, B., D. Niu, H. Wei, J. Lin, Y. He, K. Lai, and Y. Xu (2019a). “Matching article pairs with graphical decomposition and convolutions”. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics. 6284–6294. URL: <https://www.aclweb.org/anthology/P19-1632>.
- Liu, F., Z. Cheng, C. Sun, Y. Wang, L. Nie, and M. S. Kankanhalli (2019b). “User diverse preference modeling by multimodal attentive metric learning”. In: *Proceedings of the 27th ACM International Conference on Multimedia*. 1526–1534.
- Liu, T.-Y. (2009). “Learning to rank for information retrieval”. *Foundations and Trends in Information Retrieval*. 3(3): 225–331.
- Liu, Y., Y. Guo, E. M. Bakker, and M. S. Lew (2017). “Learning a recurrent residual fusion network for multimodal matching”. In: *2017 IEEE International Conference on Computer Vision (ICCV)*. 4127–4136.
- Ma, L., Z. Lu, L. Shang, and H. Li (2015). “Multimodal convolutional neural networks for matching image and sentence”. In: *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV). ICCV '15*. Washington, DC, USA: IEEE Computer Society. 2623–2631.

- Masci, J., U. Meier, D. Cireşan, and J. Schmidhuber (2011). “Stacked convolutional auto-encoders for hierarchical feature extraction”. In: *Proceedings of the 21th International Conference on Artificial Neural Networks – Volume Part I. ICANN’11*. Espoo, Finland: Springer-Verlag. 52–59. URL: <http://dl.acm.org/citation.cfm?id=2029556.2029563>.
- Mikolov, T., I. Sutskever, K. Chen, G. Corrado, and J. Dean (2013). “Distributed representations of words and phrases and their compositionality”. In: *Proceedings of the 26th International Conference on Neural Information Processing Systems – Volume 2. NIPS’13*. Lake Tahoe, Nevada: Curran Associates Inc. 3111–3119. URL: <http://dl.acm.org/citation.cfm?id=2999792.2999959>.
- Mitra, B. and N. Craswell (2018). “An introduction to neural information retrieval”. *Foundations and Trends in Information Retrieval*. 13(1): 1–126.
- Mitra, B. and N. Craswell (2019). “Duet at Trec 2019 deep learning track”. In: *Proceedings of the Twenty-Eighth Text REtrieval Conference, TREC 2019, Gaithersburg, MD, USA*. URL: <https://trec.nist.gov/pubs/trec28/papers/Microsoft.DL.pdf>.
- Mitra, B., F. Diaz, and N. Craswell (2017). “Learning to match using local and distributed representations of text for web search”. In: *Proceedings of the 26th International Conference on World Wide Web. WWW ’17*. Perth, Australia. 1291–1299.
- Nallapati, R. (2004). “Discriminative models for information retrieval”. In: *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR ’04*. Sheffield, UK: ACM. 64–71.
- Naumov, M., D. Mudigere, H. M. Shi, J. Huang, N. Sundaraman, J. Park, X. Wang, U. Gupta, C. Wu, A. G. Azzolini, D. Dzhulgakov, A. Malleovich, I. Cherniavskii, Y. Lu, R. Krishnamoorthi, A. Yu, V. Kondratenko, S. Pereira, X. Chen, W. Chen, V. Rao, B. Jia, L. Xiong, and M. Smelyanskiy (2019). “Deep learning recommendation model for personalization and recommendation systems”. *CoRR*. abs/1906.00091. arXiv: [1906.00091](https://arxiv.org/abs/1906.00091).

- Nguyen, T., M. Rosenberg, X. Song, J. Gao, S. Tiwary, R. Majumder, and L. Deng (2016). “MS MARCO: A human generated MACHine Reading COmprehension dataset”. In: *Proceedings of the Workshop on Cognitive Computation: Integrating Neural and Symbolic Approaches 2016 Co-Located with the 30th Annual Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain*.
- Nie, Y., A. Sordoni, and J.-Y. Nie (2018). “Multi-level abstraction convolutional model with weak supervision for information retrieval”. In: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval. SIGIR '18*. Ann Arbor, MI, USA: ACM. 985–988.
- Nogueira, R. and K. Cho (2019). “Passage re-ranking with BERT”. *CoRR*. abs/1901.04085. arXiv: [1901.04085](https://arxiv.org/abs/1901.04085).
- Nogueira, R., W. Yang, K. Cho, and J. Lin (2019). “Multi-stage document ranking with BERT”. *CoRR*. abs/1910.14424. arXiv: [1910.14424](https://arxiv.org/abs/1910.14424).
- Onal, K. D., Y. Zhang, I. S. Altingovde, M. M. Rahman, P. Karagoz, A. Braylan, B. Dang, H.-L. Chang, H. Kim, Q. Mcnamara, A. Angert, E. Banner, V. Khetan, T. McDonnell, A. T. Nguyen, D. Xu, B. C. Wallace, M. Rijke, and M. Lease (2018). “Neural information retrieval: At the end of the early years”. *Information Retrieval Journal*. 21(2–3): 111–182. ISSN: 1573-7659.
- Palangi, H., L. Deng, Y. Shen, J. Gao, X. He, J. Chen, X. Song, and R. Ward (2016). “Deep sentence embedding using long short-term memory networks: Analysis and application to information retrieval”. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*. 24(4): 694–707.
- Pan, F., S. Li, X. Ao, P. Tang, and Q. He (2019). “Warm up cold-start advertisements: Improving CTR predictions via learning to learn ID embeddings”. In: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France*. 695–704.
- Pang, L., Y. Lan, J. Guo, J. Xu, and X. Cheng (2016a). “A study of MatchPyramid models on ad-hoc retrieval”. *CoRR*. abs/1606.04648. arXiv: [1606.04648](https://arxiv.org/abs/1606.04648).

- Pang, L., Y. Lan, J. Guo, J. Xu, S. Wan, and X. Cheng (2016b). “Text matching as image recognition”. In: *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence. AAAI’16*. Phoenix, AZ: AAAI Press. 2793–2799. URL: <http://dl.acm.org/citation.cfm?id=3016100.3016292>.
- Pang, L., Y. Lan, J. Guo, J. Xu, J. Xu, and X. Cheng (2017a). “Deep-Rank: A new deep architecture for relevance ranking in information retrieval”. In: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management. CIKM ’17*. Singapore, Singapore: ACM. 257–266.
- Pang, L., Y. Lan, J. Xu, J. Guo, S.-X. Wan, and X. Cheng (2017b). “A survey on deep text matching”. *Chinese Journal of Computers*. 40(4): 985–1003. URL: <http://cjc.ict.ac.cn/online/onlinepaper/pl-201745181647.pdf>.
- Pang, L., J. Xu, Q. Ai, Y. Lan, X. Cheng, and J.-R. Wen (2020). “SetRank: Learning a permutation-invariant ranking model for information retrieval”. In: *The 43rd International ACM SIGIR Conference on Research & Development in Information Retrieval. SIGIR ’20*. Association for Computing Machinery.
- Parikh, A., O. Täckström, D. Das, and J. Uszkoreit (2016). “A decomposable attention model for natural language inference”. In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, TX: Association for Computational Linguistics. 2249–2255. URL: <http://www.aclweb.org/anthology/D16-1244>.
- Pasricha, R. and J. McAuley (2018). “Translation-based factorization machines for sequential recommendation”. In: *Proceedings of the 12th ACM Conference on Recommender Systems. RecSys ’18*. Vancouver, British Columbia, Canada: Association for Computing Machinery. 63–71.
- Pasumarthi, R. K., S. Bruch, X. Wang, C. Li, M. Bendersky, M. Najork, J. Pfeifer, N. Golbandi, R. Anil, and S. Wolf (2019). “TF-ranking: Scalable TensorFlow library for learning-to-rank”. In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. KDD ’19*. Anchorage, AK, USA: ACM. 2970–2978.

- Pearl, J. (2019). “The seven tools of causal inference, with reflections on machine learning”. *Communications of the ACM*. 62(3): 54–60.
- Pei, C., Y. Zhang, Y. Zhang, F. Sun, X. Lin, H. Sun, J. Wu, P. Jiang, J. Ge, W. Ou, and D. Pei (2019). “Personalized re-ranking for recommendation”. In: *Proceedings of the 13th ACM Conference on Recommender Systems. RecSys '19*. Copenhagen, Denmark: Association for Computing Machinery. 3–11.
- Pennington, J., R. Socher, and C. Manning (2014). “Glove: Global vectors for word representation”. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha, Qatar: Association for Computational Linguistics. 1532–1543. URL: <https://www.aclweb.org/anthology/D14-1162>.
- Peters, M., M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer (2018). “Deep contextualized word representations”. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. New Orleans, LA: Association for Computational Linguistics. 2227–2237. URL: <https://www.aclweb.org/anthology/N18-1202>.
- Qiao, Y., C. Xiong, Z. Liu, and Z. Liu (2019). “Understanding the behaviors of BERT in ranking”. *CoRR*. abs/1904.07531. arXiv: [1904.07531](https://arxiv.org/abs/1904.07531).
- Qiu, R., J. Li, Z. Huang, and H. Yin (2019). “Rethinking the item order in session-based recommendation with graph neural networks”. In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3–7, 2019*. 579–588.
- Qiu, X. and X. Huang (2015). “Convolutional neural tensor network architecture for community-based question answering”. In: *Proceedings of the 24th International Conference on Artificial Intelligence. IJCAI'15*. Buenos Aires, Argentina: AAAI Press. 1305–1311. URL: <http://dl.acm.org/citation.cfm?id=2832415.2832431>.
- Radford, A., K. Narasimhan, T. Salimans, and I. Sutskever (2018). “Improving language understanding by generative pre-training”. *Technical report*, OpenAI. URL: https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf.

- Radford, A., J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever (2019). “Language models are unsupervised multitask learners”. *Technical report*, OpenAI. URL: https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf.
- Ranzato, M. A., Y.-L. Boureau, and Y. LeCun (2007). “Sparse feature learning for deep belief networks”. In: *Proceedings of the 20th International Conference on Neural Information Processing Systems. NIPS’07*. Vancouver, British Columbia, Canada: Curran Associates Inc. 1185–1192. URL: <http://dl.acm.org/citation.cfm?id=2981562.2981711>.
- Rasiwasia, N., J. Costa Pereira, E. Coviello, G. Doyle, G. R. Lanckriet, R. Levy, and N. Vasconcelos (2010). “A new approach to cross-modal multimedia retrieval”. In: *Proceedings of the 18th ACM International Conference on Multimedia. MM ’10*. Firenze, Italy: ACM. 251–260.
- Reimers, N. and I. Gurevych (2019). “Sentence-BERT: Sentence embeddings using siamese BERT-networks”. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, Hong Kong, China*. Association for Computational Linguistics. 3982–3992. URL: <http://arxiv.org/abs/1908.10084>.
- Rendle, S. (2010). “Factorization machines”. In: *Proceedings of the 2010 IEEE International Conference on Data Mining. ICDM ’10*. Washington, DC, USA: IEEE Computer Society. 995–1000.
- Rendle, S., C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme (2009). “BPR: Bayesian personalized ranking from implicit feedback”. In: *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence. UAI ’09*. Montreal, Quebec, Canada: AUAI Press. 452–461. URL: <http://dl.acm.org/citation.cfm?id=1795114.1795167>.
- Rendle, S., C. Freudenthaler, and L. Schmidt-Thieme (2010). “Factorizing personalized Markov chains for next-basket recommendation”. In: *Proceedings of the 19th International Conference on World Wide Web. WWW ’10*. Raleigh, NC, USA: ACM. 811–820.
- Ricci, F., L. Rokach, and B. Shapira (2015). *Recommender Systems Handbook*. 2nd Edn. Springer Publishing Company, Incorporated.

- Richardson, M., R. Agrawal, and P. M. Domingos (2003). “Trust management for the semantic web”. In: *The Semantic Web - ISWC 2003, Second International Semantic Web Conference, Sanibel Island, FL, USA*. 351–368.
- Robertson, S., H. Zaragoza, and M. Taylor (2004). “Simple BM25 extension to multiple weighted fields”. In: *Proceedings of the Thirteenth ACM International Conference on Information and Knowledge Management. CIKM '04*. Washington, D.C., USA: ACM. 42–49.
- Rosipal, R. and N. Krämer (2006). “Overview and recent advances in partial least squares”. In: *Proceedings of the 2005 International Conference on Subspace, Latent Structure and Feature Selection. SLSFS'05*. Bohinj, Slovenia: Springer-Verlag. 34–51.
- Salakhutdinov, R. and A. Mnih (2007). “Probabilistic matrix factorization”. In: *Proceedings of the 20th International Conference on Neural Information Processing Systems. NIPS'07*. Vancouver, British Columbia, Canada: Curran Associates Inc. 1257–1264. URL: <http://dl.acm.org/citation.cfm?id=2981562.2981720>.
- Sarwar, B., G. Karypis, J. Konstan, and J. Riedl (2001). “Item-based collaborative filtering recommendation algorithms”. In: *Proceedings of the 10th International Conference on World Wide Web. WWW '01*. Hong Kong, Hong Kong: ACM. 285–295.
- Schedl, M., H. Zamani, C.-W. Chen, Y. Deldjoo, and M. Elahi (2018). “Current challenges and visions in music recommender systems research”. *International Journal of Multimedia Information Retrieval*. 7(2): 95–116.
- Sedhain, S., A. K. Menon, S. Sanner, and L. Xie (2015). “AutoRec: Autoencoders meet collaborative filtering”. In: *Proceedings of the 24th International Conference on World Wide Web. WWW '15 Companion*. Florence, Italy: ACM. 111–112.
- Shan, Y., T. R. Hoens, J. Jiao, H. Wang, D. Yu, and J. Mao (2016). “Deep crossing: Web-scale modeling without manually crafted combinatorial features”. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '16*. San Francisco, CA, USA: ACM. 255–262.

- Shen, Y., X. He, J. Gao, L. Deng, and G. Mesnil (2014). “A latent semantic model with convolutional-pooling structure for information retrieval”. In: *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management. CIKM '14*. Shanghai, China: ACM. 101–110.
- Shi, Y., M. Larson, and A. Hanjalic (2014). “Collaborative filtering beyond the user-item matrix: A survey of the state of the art and future challenges”. *ACM Computing Surveys*. 47(1): 3:1–3:45.
- Socher, R., D. Chen, C. D. Manning, and A. Y. Ng (2013). “Reasoning with neural tensor networks for knowledge base completion”. In: *Proceedings of the 26th International Conference on Neural Information Processing Systems – Volume 1. NIPS'13*. Lake Tahoe, Nevada: Curran Associates Inc. 926–934. URL: <http://dl.acm.org/citation.cfm?id=2999611.2999715>.
- Srivastava, N., G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov (2014). “Dropout: A simple way to prevent neural networks from overfitting”. *Journal of Machine Learning Research*. 15(1): 1929–1958. URL: <http://jmlr.org/papers/v15/srivastava14a.html>.
- Sun, F., J. Liu, J. Wu, C. Pei, X. Lin, W. Ou, and P. Jiang (2019). “BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer”. In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management. CIKM '19*. Beijing, China: ACM. 1441–1450.
- Surdeanu, M., M. Ciaramita, and H. Zaragoza (2011). “Learning to rank answers to non-factoid questions from web collections”. *Computational Linguistics*. 37(2): 351–383.
- Tan, C., F. Wei, W. Wang, W. Lv, and M. Zhou (2018). “Multiway attention networks for modeling sentence pairs”. In: *Proceedings of the 27th International Joint Conference on Artificial Intelligence. IJCAI'18*. Stockholm, Sweden: AAAI Press. 4411–4417.
- Tang, J., X. Du, X. He, F. Yuan, Q. Tian, and T. Chua (2020). “Adversarial training towards robust multimedia recommender system”. *IEEE Transactions on Knowledge and Data Engineering*. 32(5): 855–867.

- Tang, J. and K. Wang (2018). “Personalized top-N sequential recommendation via convolutional sequence embedding”. In: *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining. WSDM '18*. Marina Del Rey, CA, USA: Association for Computing Machinery. 565–573.
- Tao, Z., X. Wang, X. He, X. Huang, and T.-S. Chua (2019). “HoAFM: A High-Order Attentive Factorization Machine for ctr Prediction”. *Information Processing & Management*: 102076. URL: <http://www.sciencedirect.com/science/article/pii/S0306457319302389>.
- Tay, Y., L. Anh Tuan, and S. C. Hui (2018a). “Latent relational metric learning via memory-based attention for collaborative ranking”. In: *Proceedings of the 2018 World Wide Web Conference. WWW '18*. Lyon, France. 729–739.
- Tay, Y., A. T. Luu, and S. C. Hui (2018b). “Co-stack residual affinity networks with multi-level attention refinement for matching text sequences”. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels, Belgium: Association for Computational Linguistics. 4492–4502. URL: <https://www.aclweb.org/anthology/D18-1479>.
- Tay, Y., A. T. Luu, and S. C. Hui (2018c). “Hermitian co-attention networks for text matching in asymmetrical domains”. In: *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization. 4425–4431.
- Tay, Y., L. A. Tuan, and S. C. Hui (2018d). “Multi-cast attention networks”. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. KDD '18*. New York, NY, USA: Association for Computing Machinery. 2299–2308.
- Van Gysel, C., M. de Rijke, and E. Kanoulas (2016a). “Learning latent vector spaces for product search”. In: *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management. CIKM '16*. Indianapolis, IN, USA: ACM. 165–174.
- Van Gysel, C., M. de Rijke, and E. Kanoulas (2017). “Structural regularities in text-based entity vector spaces”. In: *Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval. ICTIR '17*. Amsterdam, The Netherlands: ACM. 3–10.

- Van Gysel, C., M. de Rijke, and E. Kanoulas (2018). “Mix ’N match: Integrating text matching and product substitutability within product search”. In: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management. CIKM ’18*. Torino, Italy: ACM. 1373–1382.
- Van Gysel, C., M. de Rijke, and M. Worring (2016b). “Unsupervised, efficient and semantic expertise retrieval”. In: *Proceedings of the 25th International Conference on World Wide Web. WWW ’16*. Montreal, Quebec, Canada. 1069–1079.
- Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin (2017). “Attention is all you need”. In: *Advances in Neural Information Processing Systems 30*. Curran Associates, Inc. 5998–6008. URL: <http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>.
- Vincent, P., H. Larochelle, Y. Bengio, and P.-A. Manzagol (2008). “Extracting and composing robust features with denoising autoencoders”. In: *Proceedings of the 25th International Conference on Machine Learning. ICML ’08*. Helsinki, Finland: ACM. 1096–1103.
- Wan, S., Y. Lan, J. Guo, J. Xu, L. Pang, and X. Cheng (2016a). “A deep architecture for semantic matching with multiple positional sentence representations”. In: *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence. AAAI’16*. Phoenix, AZ: AAAI Press. 2835–2841. URL: <http://dl.acm.org/citation.cfm?id=3016100.3016298>.
- Wan, S., Y. Lan, J. Xu, J. Guo, L. Pang, and X. Cheng (2016b). “Match-SRNN: Modeling the recursive matching structure with spatial RNN”. In: *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence. IJCAI’16*. New York, NY, USA: AAAI Press. 2922–2928. URL: <http://dl.acm.org/citation.cfm?id=3060832.3061030>.
- Wang, B., Y. Yang, X. Xu, A. Hanjalic, and H. T. Shen (2017a). “Adversarial cross-modal retrieval”. In: *Proceedings of the 25th ACM International Conference on Multimedia. MM ’17*. Mountain View, CA, USA: ACM. 154–162.

- Wang, H., F. Zhang, J. Wang, M. Zhao, W. Li, X. Xie, and M. Guo (2018a). “RippleNet: Propagating user preferences on the knowledge graph for recommender systems”. In: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*. New York, NY, USA: Association for Computing Machinery. 417–426.
- Wang, J., A. P. de Vries, and M. J. T. Reinders (2006). “Unifying user-based and item-based collaborative filtering approaches by similarity fusion”. In: *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '06*. Seattle, Washington, DC, USA: ACM. 501–508.
- Wang, L., Y. Li, J. Huang, and S. Lazebnik (2018b). “Learning two-branch neural networks for image-text matching tasks”. *IEEE Transactions on Pattern Analysis and Machine Intelligence*: 1–1.
- Wang, L., Y. Li, and S. Lazebnik (2016). “Learning deep structure-preserving image-text embeddings”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Vol. 00. 5005–5013. URL: doi.ieeecomputersociety.org/10.1109/CVPR.2016.541.
- Wang, X., X. He, Y. Cao, M. Liu, and T. Chua (2019a). “KGAT: Knowledge graph attention network for recommendation”. In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2019, Anchorage, AK, USA, August 4–8, 2019*. 950–958.
- Wang, X., X. He, F. Feng, L. Nie, and T. Chua (2018c). “TEM: Tree-enhanced embedding model for explainable recommendation”. In: *Proceedings of the 2018 World Wide Web Conference on World Wide Web*. 1543–1552.
- Wang, X., X. He, L. Nie, and T.-S. Chua (2017b). “Item silk road: Recommending items from information domains to social users”. In: *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '17*. Shinjuku, Tokyo, Japan: ACM. 185–194.

- Wang, X., X. He, M. Wang, F. Feng, and T.-S. Chua (2019b). “Neural graph collaborative filtering”. In: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR’19*. Paris, France: Association for Computing Machinery. 165–174.
- Wang, X., D. Wang, C. Xu, X. He, Y. Cao, and T. Chua (2019c). “Explainable reasoning over knowledge graphs for recommendation”. In: *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019*. 5329–5336. URL: <https://aaai.org/ojs/index.php/AAAI/article/view/4470>.
- Wang, X., Q. Huang, A. Celikyilmaz, J. Gao, D. Shen, Y.-F. Wang, W. Y. Wang, and L. Zhang (2019d). “Reinforced cross-modal matching and self-supervised imitation learning for vision-language navigation”. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Wang, Z., W. Hamza, and R. Florian (2017c). “Bilateral multi-perspective matching for natural language sentences”. In: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*. 4144–4150.
- Wang, Z., G. Xu, H. Li, and M. Zhang (2011). “A fast and accurate method for approximate string search”. In: *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies – Volume 1. HLT ’11*. Portland, OR, USA: Association for Computational Linguistics. 52–61. URL: <http://dl.acm.org/citation.cfm?id=2002472.2002480>.
- Wei, X. and W. B. Croft (2006). “LDA-based document models for ad-hoc retrieval”. In: *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR ’06*. Seattle, Washington, DC, USA: ACM. 178–185.
- Wei, Y., X. Wang, L. Nie, X. He, R. Hong, and T.-S. Chua (2019). “MMGCN: Multi-modal graph convolution network for personalized recommendation of micro-video”. In: *Proceedings of the 27th ACM International Conference on Multimedia. MM ’19*. Nice, France: ACM. 1437–1445.

- Wu, B., X. He, Z. Sun, L. Chen, and Y. Ye (2019a). “ATM: An attentive translation model for next-item recommendation”. *IEEE Transactions on Industrial Informatics*: 1–1.
- Wu, C.-Y., A. Ahmed, A. Beutel, A. J. Smola, and H. Jing (2017). “Recurrent recommender networks”. In: *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. New York, NY, USA. 495–503.
- Wu, L., P. Sun, Y. Fu, R. Hong, X. Wang, and M. Wang (2019b). “A neural influence diffusion model for social recommendation”. In: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21–25, 2019*. 235–244.
- Wu, Q., H. Wang, Q. Gu, and H. Wang (2016a). “Contextual bandits in a collaborative environment”. In: *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '16*. Pisa, Italy. 529–538.
- Wu, S., Y. Tang, Y. Zhu, L. Wang, X. Xie, and T. Tan (2019c). “Session-based recommendation with graph neural networks”. In: *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, Honolulu, HI, USA, 2019*. 346–353.
- Wu, W., H. Li, and J. Xu (2013a). “Learning query and document similarities from click-through bipartite graph with metadata”. In: *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining. WSDM '13*. Rome, Italy: ACM. 687–696.
- Wu, W., Z. Lu, and H. Li (2013b). “Learning bilinear model for matching queries and documents”. *Journal of Machine Learning Research*. 14(1): 2519–2548. URL: <http://dl.acm.org/citation.cfm?id=2567709.2567742>.
- Wu, Y., C. DuBois, A. X. Zheng, and M. Ester (2016b). “Collaborative denoising auto-encoders for top-N recommender systems”. In: *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining. WSDM '16*. San Francisco, CA, USA: ACM. 153–162.

- Xiao, J., H. Ye, X. He, H. Zhang, F. Wu, and T.-S. Chua (2017). “Attentional factorization machines: Learning the weight of feature interactions via attention networks”. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence. IJCAI’17*. Melbourne, Australia: AAAI Press. 3119–3125. URL: <http://dl.acm.org/citation.cfm?id=3172077.3172324>.
- Xin, X., B. Chen, X. He, D. Wang, Y. Ding, and J. Jose (2019a). “CFM: Convolutional factorization machines for context-aware recommendation”. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence. IJCAI’19. International Joint Conferences on Artificial Intelligence Organization*. 3119–3125. URL: <https://doi.org/10.24963/ijcai.2019/545>.
- Xin, X., X. He, Y. Zhang, Y. Zhang, and J. M. Jose (2019b). “Relational collaborative filtering: Modeling multiple item relations for recommendation”. In: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, Paris, France, July 21–25, 2019*. 125–134.
- Xiong, C., Z. Dai, J. Callan, Z. Liu, and R. Power (2017). “End-to-end neural ad-hoc ranking with kernel pooling”. In: *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR ’17*. Shinjuku, Tokyo, Japan: ACM. 55–64.
- Xue, F., X. He, X. Wang, J. Xu, K. Liu, and R. Hong (2019). “Deep item-based collaborative filtering for top-N recommendation”. *ACM Transactions on Information Systems*. 37(3).
- Xue, H.-J., X. Dai, J. Zhang, S. Huang, and J. Chen (2017). “Deep matrix factorization models for recommender systems”. In: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*. 3203–3209.
- Yan, F. and K. Mikolajczyk (2015). “Deep correlation for matching images and text”. In: *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 3441–3450.

- Yang, L., Q. Ai, J. Guo, and W. B. Croft (2016). “aNMM: Ranking short answer texts with attention-based neural matching model”. In: *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management. CIKM '16*. Indianapolis, IN, USA: ACM. 287–296.
- Yang, P., H. Fang, and J. Lin (2018). “Anserini: Reproducible ranking baselines using lucene”. *J. Data and Information Quality*. 10(4): 16:1–16:20.
- Yang, R., J. Zhang, X. Gao, F. Ji, and H. Chen (2019a). “Simple and effective text matching with richer alignment features”. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics. 4699–4709. URL: <https://www.aclweb.org/anthology/P19-1465>.
- Yang, W., H. Zhang, and J. Lin (2019b). “Simple applications of BERT for ad hoc document retrieval”. *CoRR*. abs/1903.10972. arXiv: [1903.10972](https://arxiv.org/abs/1903.10972).
- Yang, Y., S. W.-T. Yih, and C. Meek (2015). “WikiQA: A challenge dataset for open-domain question answering”. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. ACL – Association for Computational Linguistics. URL: <https://www.microsoft.com/en-us/research/publication/wikiqa-a-challenge-dataset-for-open-domain-question-answering/>.
- Yang, Z., Z. Dai, Y. Yang, J. Carbonell, R. R. Salakhutdinov, and Q. V. Le (2019c). “XLNet: Generalized autoregressive pretraining for language understanding”. In: *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc. 5753–5763. URL: <http://papers.nips.cc/paper/8812-xlnet-generalized-autoregressive-pretraining-for-language-understanding.pdf>.
- Yi, X. and J. Allan (2009). “A comparative study of utilizing topic models for information retrieval”. In: *Proceedings of the 31th European Conference on IR Research on Advances in Information Retrieval. ECIR '09*. Toulouse, France: Springer-Verlag. 29–41.

- Yin, W. and H. Schütze (2015). “MultiGranCNN: An architecture for general matching of text chunks on multiple levels of granularity”. In: *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Beijing, China: Association for Computational Linguistics. 63–73. URL: <https://www.aclweb.org/anthology/P15-1007>.
- Yin, W., H. Schütze, B. Xiang, and B. Zhou (2016). “ABCNN: Attention-based convolutional neural network for modeling sentence pairs”. *Transactions of the Association for Computational Linguistics*. 4: 259–272.
- Ying, R., R. He, K. Chen, P. Eksombatchai, W. L. Hamilton, and J. Leskovec (2018). “Graph convolutional neural networks for web-scale recommender systems”. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. KDD '18*. London, UK: ACM. 974–983.
- Yuan, F., X. He, H. Jiang, G. Guo, J. Xiong, Z. Xu, and Y. Xiong (2020). “Future data helps training: Modeling future contexts for session-based recommendation”. In: *Proceedings of the Web Conference 2020. WWW '20*. Taipei, Taiwan: Association for Computing Machinery. 303–313.
- Yuan, F., A. Karatzoglou, I. Arapakis, J. M. Jose, and X. He (2019). “A simple convolutional generative network for next item recommendation”. In: *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining. WSDM '19*. Melbourne VIC, Australia: Association for Computing Machinery. 582–590.
- Zamani, H. and W. B. Croft (2016). “Estimating embedding vectors for queries”. In: *Proceedings of the 2016 ACM International Conference on the Theory of Information Retrieval. ICTIR '16*. Newark, DE, USA: Association for Computing Machinery. 123–132.
- Zamani, H. and W. B. Croft (2017). “Relevance-based word embedding”. In: *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '17*. Shinjuku, Tokyo, Japan: Association for Computing Machinery. 505–514.

- Zamani, H. and W. B. Croft (2018a). “Joint modeling and optimization of search and recommendation”. In: *Proceedings of the First Biennial Conference on Design of Experimental Search & Information Retrieval Systems. DESIRES '18*. Bertinoro, Italy: CEUR-WS. 36–41. URL: <http://ceur-ws.org/Vol-2167/paper2.pdf>.
- Zamani, H. and W. B. Croft (2018b). “On the theory of weak supervision for information retrieval”. In: *Proceedings of the 2018 ACM SIGIR International Conference on Theory of Information Retrieval. ICTIR '18*. Tianjin, China: Association for Computing Machinery. 147–154.
- Zamani, H. and W. B. Croft (2020). “Learning a joint search and recommendation model from user-item interactions”. In: *Proceedings of the 13th International Conference on Web Search and Data Mining. WSDM '20*. Houston, TX, USA: Association for Computing Machinery. 717–725.
- Zamani, H., W. B. Croft, and J. S. Culpepper (2018a). “Neural query performance prediction using weak supervision from multiple signals”. In: *The 41st International ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '18*. Ann Arbor, MI, USA: Association for Computing Machinery. 105–114.
- Zamani, H., J. Dadashkarimi, A. Shakery, and W. B. Croft (2016). “Pseudo-relevance feedback based on matrix factorization”. In: *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management. CIKM '16*. Indianapolis, IN, USA: ACM. 1483–1492.
- Zamani, H., M. Deghani, W. B. Croft, E. Learned-Miller, and J. Kamps (2018b). “From neural re-ranking to neural ranking: Learning a sparse representation for inverted indexing”. In: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management. CIKM '18*. Torino, Italy: Association for Computing Machinery. 497–506.
- Zamani, H., B. Mitra, X. Song, N. Craswell, and S. Tiwary (2018c). “Neural ranking models with multiple document fields”. In: *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining. WSDM '18*. Marina Del Rey, CA, USA: ACM. 700–708.

- Zhang, S., L. Yao, A. Sun, and Y. Tay (2019). “Deep learning based recommender system: A survey and new perspectives”. *ACM Computing Surveys*. 52(1): Article 5.
- Zhang, X., H. Xie, H. Li, and J. C. S. Lui (2020). “Conversational contextual bandit: Algorithm and application”. In: *Proceedings of the Web Conference 2020. WWW '20*. Taipei, Taiwan: Association for Computing Machinery. 662–672.
- Zhang, Y., Q. Ai, X. Chen, and W. B. Croft (2017). “Joint representation learning for top-N recommendation with heterogeneous information sources”. In: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management. CIKM '17*. Singapore, Singapore: ACM. 1449–1458.
- Zhang, Y., X. Chen, Q. Ai, L. Yang, and W. B. Croft (2018). “Towards conversational search and recommendation: System ask, user respond”. In: *Proceedings of the 27th ACM International Conference on Information and Knowledge Management. CIKM '18*. Torino, Italy: ACM. 177–186.
- Zhao, W. X., G. He, K. Yang, H. Dou, J. Huang, S. Ouyang, and J. Wen (2019). “KB4Rec: A data set for linking knowledge bases with recommender systems”. *Data Intelligence*. 1(2): 121–136.
- Zheng, L., C. Lu, F. Jiang, J. Zhang, and P. S. Yu (2018a). “Spectral collaborative filtering”. In: *Proceedings of the 12th ACM Conference on Recommender Systems, RecSys 2018, Vancouver, BC, Canada, October 2–7, 2018*. 311–319.
- Zheng, L., V. Noroozi, and P. S. Yu (2017). “Joint deep modeling of users and items using reviews for recommendation”. In: *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining. WSDM '17*. Cambridge, UK: ACM. 425–434.
- Zheng, Y., Z. Fan, Y. Liu, C. Luo, M. Zhang, and S. Ma (2018b). “Sogou-QCL: A new dataset with click relevance label”. In: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval. SIGIR '18*. Ann Arbor, MI, USA: ACM. 1117–1120.

- Zhou, G., X. Zhu, C. Song, Y. Fan, H. Zhu, X. Ma, Y. Yan, J. Jin, H. Li, and K. Gai (2018). “Deep interest network for click-through rate prediction”. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. KDD '18*. London, UK: Association for Computing Machinery. 1059–1068.
- Zhu, M., A. Ahuja, W. Wei, and C. K. Reddy (2019). “A hierarchical attention retrieval model for healthcare question answering”. In: *The World Wide Web Conference. WWW '19*. San Francisco, CA, USA: Association for Computing Machinery. 2472–2482.