

Knowledge Graphs: An Information Retrieval Perspective

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

Knowledge Graphs: An Information Retrieval Perspective

Ridho Reinanda

Bloomberg L.P.

UK

rreinanda@bloomberg.net

Edgar Meij

Bloomberg L.P.

UK

emeij@bloomberg.net

Maarten de Rijke

University of Amsterdam & Ahold Delhaize

The Netherlands

m.derijke@uva.nl

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

R. Reinanda, E. Meij and M. de Rijke. *Knowledge Graphs: An Information Retrieval Perspective*. Foundations and Trends[®] in Information Retrieval, vol. 14, no. 4, pp. 289–444, 2020.

ISBN: 978-1-68083-729-2

© 2020 R. Reinanda, E. Meij and M. de Rijke

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 4, 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	2
1.1	Motivation	2
1.2	Aims	3
1.3	Methodology	3
1.4	Scope	4
1.5	Structure	4
2	Preliminaries	6
2.1	Key Concepts	6
2.2	Evaluation	9
3	Background on Entity Linking and Recognition	11
3.1	Entity Linking	11
3.2	Entity Recognition and Classification	21
4	Knowledge Graphs for Information Retrieval	31
4.1	Document Retrieval	32
4.2	Entity Retrieval	43
4.3	Entity Recommendation	51
4.4	Entity Relationship Explanation	59
4.5	Conclusion	62

5	Information Retrieval for Knowledge Graphs	64
5.1	Entity Discovery	67
5.2	Entity Typing	73
5.3	Entity-Centric Document Filtering	78
5.4	Relation Extraction and Link Prediction	83
5.5	KG Quality Estimation	94
5.6	Conclusion	101
6	Applications	102
6.1	Web Search	102
6.2	Knowledge Graph Construction	109
7	Conclusion and Discussion	115
7.1	Conclusion	115
7.2	Discussion	116
	Acknowledgements	121
	Appendices	122
A	Acronyms Used	123
B	Resources	126
B.1	Corpora	126
B.2	Knowledge Graphs	126
B.3	Datasets	126
B.4	Code	130
B.5	Libraries	130
B.6	Tutorials	130
	References	133

Knowledge Graphs: An Information Retrieval Perspective

Ridho Reinanda¹, Edgar Meij² and Maarten de Rijke³

¹*Bloomberg L.P., UK; rreinanda@bloomberg.net*

²*Bloomberg L.P., UK; emej@bloomberg.net*

³*University of Amsterdam & Ahold Delhaize, The Netherlands; m.derijke@uva.nl*

ABSTRACT

In this survey, we provide an overview of the literature on knowledge graphs (KGs) in the context of information retrieval (IR). Modern IR systems can benefit from information available in KGs in multiple ways, independent of whether the KGs are publicly available or proprietary ones. We provide an overview of the components required when building IR systems that leverage KGs and use a task-oriented organization of the material that we discuss. As an understanding of the intersection of IR and KGs is beneficial to many researchers and practitioners, we consider prior work from two complementary angles: leveraging KGs for information retrieval and enriching KGs using IR techniques. We start by discussing how KGs can be employed to support IR tasks, including document and entity retrieval. We then proceed by describing how IR—and language technology in general—can be utilized for the construction and completion of KGs. This includes tasks such as entity recognition, typing, and relation extraction. We discuss common issues that appear across the tasks that we consider and identify future directions for addressing them. We also provide pointers to datasets and other resources that should be useful for both newcomers and experienced researchers in the area.

1

Introduction

1.1 Motivation

A *knowledge graph* (KG) is a repository of entities as well as their relationships and attributes that is represented as a graph. In modern approaches to information access, KGs are ubiquitous (Dalton and Dietz, 2013). Specifically, in information retrieval (IR) KGs are instrumental in enabling semantic search.

There are two hallmarks of semantic search in an IR context: (1) going beyond “ten blue links” in order to return relevant results of any kind (such as direct answers, actionable entities, or relationships) and (2) understanding queries and documents, and improving the matching between them with relevant relationships. Ideally, a search engine is able to directly answer a user’s information need—or at least generate possible interpretations of the information need that is expressed through the query. To achieve this goal, various entity-oriented components that solve specific problems at different stages in the information retrieval pipeline are required, including identifying entities in the query, identifying entities in documents, and methods that leverage entity and relationship information to help identify relevant items to retrieve.

Despite the fact that IR and KGs are increasingly intertwined in the context of modern web and domain-specific search engines, there is no broad treatment in the literature of KGs from an IR perspective, and vice versa. We aim to fill this gap through this survey by providing a task-oriented overview of research in this area.

1.2 Aims

The aim of this survey is to bridge two important components of modern information access: IR and KGs. We summarize research work, group related approaches, and discuss challenges shared across tasks at the interface of IR and KGs. Our contributions in this survey can be summarized as follows: (1) we present an extensive overview of tasks related to KGs from an IR perspective; (2) we provide a thorough review for each task; and (3) we present discussions on common issues that are shared among the tasks.

1.3 Methodology

To meet the aim articulated above and to be able to present the methods described in this survey in a systematic manner, we first identify different sets of tasks related to IR and KGs and group individual tasks that are closely related. The main organizational principle that we use in the survey is to group tasks in two directions: *knowledge graphs for information retrieval* and *information retrieval for knowledge graphs*.

For each task, we trace back its origin, the original motivation, setup, and define the task in a formal fashion. We then identify seminal work or influential approaches as they have been introduced over time. We group approaches based on characteristics that are natural for each task. We also identify related work based on these groupings. We put more emphasis on recent developments concerning the task, how the methods differ from early approaches, and the interesting additional problems that arise over time in the context of the task.

Having examined each task one-by-one, we then proceed to identify the key challenges that we encounter frequently across tasks. We focus on challenging aspects that will be beneficial for future research.

1.4 Scope

We consider over 300 publications published prior to 2020 and spanning the fields of information retrieval, knowledge representation, machine learning, and natural language processing. Due to the broad nature of the survey, we put more emphasis on recent advances involving new tasks and approaches. Thus, some tasks and approaches will be covered in greater detail than others.

In the survey, our view of **IR** is an inclusive one and that incorporates natural language processing and language technology techniques. We also consider tasks that have an origin in those fields, such as entity recognition, relation extraction, and knowledge base (**KB**) completion.

Some of the tasks that we consider cover a broad area. For broad tasks—that may well deserve a survey of their own—we only cover key publications and present the task in a high-level fashion. This includes tasks such as entity recognition, entity linking, and relation extraction. We present an overview of tasks, but refer the reader to existing surveys (if they exist) for details. For emerging, specific tasks we provide more details in addition to a literature review. We present their setup and contrast different approaches with more depth and detail.

Recent interest in the area of semantic search has not only given rise to hundreds of publications but also to attempts to synthesize the material. By now there is a growing number of tutorials in the area, which we enumerate in Appendix **B.6**. While we believe that ours is the first survey to focus on the interaction between **IR** and **KGs**, there is a recent survey on semantic search by Bast *et al.* (2016) that partially overlaps with ours. The most significant differences are that we discuss recent developments on how **KGs** are being leveraged for **IR**, we provide a broader coverage of knowledge graph construction and completion, and finally, we present several applications that involve a combination of the individual tasks and components in our survey.

1.5 Structure

The rest of this survey is organized as follows. In Section **2** we describe the background: definitions of fundamental concepts and notation that

we use throughout the survey. In Section 3 we introduce core entity-related tasks on which we build in the remainder of the survey: entity linking and named entity recognition and classification. The heart of the survey consists of Sections 4 and 5. Section 4 describes how KGs are being leveraged to improve IR tasks. In Section 5 we turn the table and detail how IR and, more generally, language technology is being used for the construction and completion of KGs. Section 6 is meant to provide detailed motivation for the survey by offering treatments of end-to-end tasks at the interface of IR and KGs. We conclude the survey in Section 7 with a look back, with a discussion of the key issues that we identified in the course of the survey, and with potential research directions in at the interface of IR and KGs.

Acronyms used and useful resources used in this survey are listed in appendices to this survey, Appendix A and B, respectively.

As to possible reading orders of the material in the survey, we recommend the following. Readers who are relatively new to the area should simply read all sections in their natural order: 1, 2, . . . , 7. Readers who are already familiar with the area can move ahead to the core of the survey in Section 4. Alternatively, they can freshen up on notation and terminology in Section 2, catch up on the background material on entity linking and entity recognition and classification in Section 3, sample from Section 6, and then continue with the remaining material. See Figure 1.1.




Figure 1.1: Possible reading orders.

References

- Abhishek, A., A. Anand, and A. Awekar (2017). “Fine-grained entity type classification by jointly learning representations and label embeddings”. In: *EACL '17*. ACL.
- Agichtein, E. and L. Gravano (2000). “Snowball: Extracting relations from large plain-text collections”. In: *DL '00*. ACM.
- Alfonseca, E., K. Filippova, J.-Y. Delort, and G. Garrido (2012). “Pattern learning for relation extraction with a hierarchical topic model”. In: *ACL '12*. ACL.
- Alhelbawy, A. and R. Gaizauskas (2014). “Graph ranking for collective named entity disambiguation”. In: *ACL '14*. ACL.
- Allan, J. (2002). “Introduction to topic detection and tracking”. In: *Topic Detection and Tracking*. Kluwer Academic Publishers. 1–16.
- Asahara, M. and Y. Matsumoto (2003). “Japanese named entity extraction with redundant morphological analysis”. In: *HLT-NAACL '03*. ACL.
- Bach, N. and S. Badaskar (2007). *A Survey on Relation Extraction*. Language Technologies Institute, Carnegie Mellon University.
- Balog, K. (2018). *Entity-Oriented Search*. Springer.
- Balog, K. and R. Neumayer (2013). “A test collection for entity search in DBpedia”. In: *SIGIR '13*. ACM.
- Balog, K. and H. Ramampiaro (2013). “Cumulative citation recommendation: Citation vs. ranking”. In: *SIGIR '13*. ACM.

- Balog, K., L. Azzopardi, and M. de Rijke (2006). “Formal models for expert finding in enterprise corpora”. In: *SIGIR '06*. ACM.
- Balog, K., L. Azzopardi, and M. de Rijke (2009a). “A language modeling framework for expert finding”. *Information Processing & Management*. 45(1): 1–19.
- Balog, K., P. Serdyukov, A. P. de Vries, P. Thomas, and T. Westerveld (2009b). “Overview of the TREC 2009 entity track”. In: *TREC*.
- Balog, K., M. Bron, and M. de Rijke (2011). “Query modeling for entity search based on terms, categories, and examples”. *ACM Transactions on Information Systems*. 29(4): 22:1–22:31.
- Balog, K., Y. Fang, M. de Rijke, P. Serdyukov, and L. Si (2012). “Expertise retrieval”. *Foundations and Trends in Information Retrieval*. 6(2–3): 127–256.
- Balog, K., H. Ramampiaro, N. Takhirov, and K. Nørvgåg (2013). “Multi-step classification approaches to cumulative citation recommendation”. In: *OAIR '13*. Le Centre De Hautes Etudes Internationales D’Informatique Documentaire. 121–128.
- Banko, M., M. J. Cafarella, S. Soderland, M. Broadhead, and O. Etzioni (2007). “Open information extraction from the web”. In: *IJCAI '07*. ACM.
- Bao, J., N. Duan, M. Zhou, and T. Zhao (2014). “Knowledge-based question answering as machine translation”. In: *ACL '14*. ACL.
- Bast, H., B. Buchhold, and E. Haussmann (2015). “Relevance scores for triples from type-like relations”. In: *SIGIR '15*. ACM.
- Bast, H., B. Buchhold, and E. Haussmann (2016). “Semantic search on text and knowledge bases”. *Foundations and Trends in Information Retrieval*. 10(2–3): 119–271.
- Berant, J., A. Chou, R. Frostig, and P. Liang (2013). “Semantic parsing on freebase from question-answer pairs”. In: *EMNLP '13*.
- Bi, B., H. Ma, B.-J. Hsu, W. Chu, K. Wang, and J. Cho (2015). “Learning to recommend related entities to search users”. In: *WSDM '15*. ACM.
- Bikel, D. M., R. Schwartz, and R. M. Weischedel (1999). “An algorithm that learns what’s in a name”. *Machine Learning*. 34(1–3): 211–231.

- Bing, L., W. Lam, and T.-L. Wong (2013). “Wikipedia entity expansion and attribute extraction from the web using semi-supervised learning”. In: *WSDM '13*. ACM.
- Blanco, R., B. B. Cambazoglu, P. Mika, and N. Torzec (2013). “Entity recommendations in web search”. In: *ISWC '13*. Springer-Verlag.
- Blanco, R., G. Ottaviano, and E. Meij (2015). “Fast and space-efficient entity linking for queries”. In: *WSDM '15*. ACM.
- Bonnefoy, L., V. Bouvier, and P. Bellot (2013). “A weakly-supervised detection of entity central documents in a stream”. In: *SIGIR '13*. ACM.
- Bordes, A., J. Weston, R. Collobert, and Y. Bengio (2011). “Learning structured embeddings of knowledge bases”. In: *AAAI '11*. AAAI Press.
- Bordes, A., N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko (2013). “Translating embeddings for modeling multi-relational data”. In: *NIPS '13*. JLMR.
- Bordes, A., X. Glorot, J. Weston, and Y. Bengio (2014). “A semantic matching energy function for learning with multi-relational data”. *Machine Learning*. 94(2): 233–259.
- Bordino, I., G. De Francisci Morales, I. Weber, and F. Bonchi (2013a). “From Machu_Picchu to ‘rafting the Urubamba river’: Anticipating information needs via the entity-query graph”. In: *WSDM '13*. ACM.
- Bordino, I., Y. Mejova, and M. Lalmas (2013b). “Penguins in sweaters, or serendipitous entity search on user-generated content”. In: *CIKM '13*. ACM.
- Borisov, A., I. Markov, M. de Rijke, and P. Serdyukov (2016). “A neural click model for web search”. In: *WWW '16*. ACM.
- Borthwick, A. E. (1999). “A maximum entropy approach to named entity recognition”. *PhD thesis*. New York, NY, USA: New York University.
- Bota, H., K. Zhou, and J. M. Jose (2016). “Playing your cards right: The effect of entity cards on search behaviour and workload”. In: *CHIIR '16*. ACM.
- Brambilla, M., S. Ceri, E. Della Valle, R. Volonterio, and F. X. Acero Salazar (2017). “Extracting emerging knowledge from social media”. In: *WWW '17*. ACM.

- Brambilla, M., S. Ceri, F. Daniel, M. Di Giovanni, A. Mauri, and G. Ramponi (2018). “Iterative knowledge extraction from social networks”. In: *WWW '18 Companion*. ACM.
- Brin, S. (1998). “Extracting patterns and relations from the world wide web”. In: *WebDB '98*. Springer-Verlag.
- Bron, M., K. Balog, and M. de Rijke (2010). “Ranking related entities: Components and analyses”. In: *CIKM '10*. ACM.
- Bronnenberg, W., H. Bunt, J. Landsbergen, R. Scha, W. Schoenmakers, and E. van Utteren (1980). “The question answering system Phliqa1”. In: *Natural Language Question Answering Systems*. Ed. by L. Bolc. MacMillan. 217–305.
- Bunescu, R. C. and R. J. Mooney (2005). “A shortest path dependency kernel for relation extraction”. In: *HTL-EMNLP '05*. ACL.
- Bunescu, R. and M. Pasca (2006). “Using encyclopedic knowledge for named entity disambiguation”. In: *EACL*. ACL.
- Cai, R., H. Wang, and J. Zhang (2015). “Learning entity representation for named entity disambiguation”. In: *ACL'13*. ACL.
- Cano, I., S. Singh, and C. Guestrin (2014). “Distributed non-parametric representations for vital filtering: UW at TREC KBA 2014”. In: *TREC 2014*. NIST.
- Carmel, D., M.-W. Chang, E. Gabrilovich, B.-J. Hsu, and K. Wang (2014). “ERD'14: Entity recognition and disambiguation challenge”. In: *SIGIR '14*. ACM.
- Ceccarelli, D., C. Lucchese, S. Orlando, R. Perego, and S. Trani (2013). “Learning relatedness measures for entity linking”. In: *CIKM '13*. ACM.
- Charton, E., M.-J. Meurs, L. Jean-Louis, and M. Gagnon (2014). “Mutual disambiguation for entity linking”. In: *ACL '14*. ACL.
- Chen, M., Y. Tian, M. Yang, and C. Zaniolo (2016). “Multi-lingual knowledge graph embeddings for cross-lingual knowledge alignment”. *CoRR*. abs/1611.03954.
- Chuklin, A., I. Markov, and M. de Rijke (2015). *Click Models for Web Search*. Morgan & Claypool Publishers.
- Collins, M. and Y. Singer (1999). “Unsupervised models for named entity classification”. In: *EMNLP '99*. ACL.

- Cornolti, M., P. Ferragina, and M. Ciaramita (2013). “A framework for benchmarking entity-annotation systems”. In: *WWW '13*. ACM.
- Cornolti, M., P. Ferragina, M. Ciaramita, S. Rüd, and H. Schütze (2016). “A Piggyback system for joint entity mention detection and linking in web queries”. In: *WWW '16*. ACM.
- Corro, L. D., A. Abujabal, R. Gemulla, and G. Weikum (2015). “FINET: Context-aware fine-grained named entity typing”. In: *EMNLP '15*. ACL.
- Craswell, N., A. P. de Vries, and I. Soboroff (2005). “Overview of the TREC-2005 enterprise track”. In: *TREC 2005 Conference Notebook*. NIST.
- Croft, B., D. Metzler, and T. Strohman (2009). *Search Engines: Information Retrieval in Practice*. Addison-Wesley Publishing Company.
- Cucchiarelli, A. and P. Velardi (2001). “Unsupervised named entity recognition using syntactic and semantic contextual evidence”. *Computational Linguistics*. 27(1): 123–131.
- Cucerzan, S. (2007). “Large-scale named entity disambiguation based on Wikipedia data”. In: *EMNLP '07*. ACL.
- Culotta, A. and J. Sorensen (2004). “Dependency tree kernels for relation extraction”. In: *ACL '04*. ACL.
- Daiber, J., M. Jakob, C. Hokamp, and P. N. Mendes (2013). “Improving efficiency and accuracy in multilingual entity extraction”. In: *Semantic Systems '13*. ACM.
- Dalton, J. and L. Dietz (2013). “Constructing query-specific knowledge bases”. In: *AKBC '13*. ACM.
- Dalton, J., J. Allan, and D. A. Smith (2011). “Passage retrieval for incorporating global evidence in sequence labeling”. In: *CIKM '11*. ACM.
- Dalton, J., L. Dietz, and J. Allan (2014). “Entity query feature expansion using knowledge base links”. In: *SIGIR '14*. ACM.
- Dalvi, B., A. Mishra, and W. W. Cohen (2016). “Hierarchical semi-supervised classification with incomplete class hierarchies”. In: *WSDM '16*. ACM.
- Dasgupta, S. S., S. N. Ray, and P. P. Talukdar (2018). “HyTE: Hyperplane-based temporally aware knowledge graph embedding”. In: *EMNLP '18*. ACL.

- de Vries, A. P., A.-M. Vercoustre, J. A. Thom, N. Craswell, and M. Lalmas (2008). “Overview of the INEX 2007 entity ranking track”. In: *Focused Access to XML Documents*. Ed. by N. Fuhr, J. Kamps, M. Lalmas, and A. Trotman. Springer-Verlag.
- Devlin, J., M. Chang, K. Lee, and K. Toutanova (2018). “BERT: Pre-training of deep bidirectional transformers for language understanding”. *NAACL '18*. ACL.
- Dietz, L. and J. Dalton (2013). “UMass at TREC 2013 knowledge base acceleration track”. In: *TREC 2013*. NIST.
- Dietz, L., A. Kotov, and E. Meij (2017). “Utilizing knowledge graphs in text-centric information retrieval”. In: *WSDM '17*. ACM.
- Dietz, L., A. Kotov, and E. Meij (2018). “Utilizing knowledge graphs for text-centric information retrieval”. In: *SIGIR '18*. ACM.
- Doddington, G., A. Mitchell, M. Przybocki, L. Ramshaw, S. Strassel, and R. Weischedel (2004). “The automatic content extraction (ACE) program – tasks, data, and evaluation”. In: *LREC '04*. ACL.
- Dong, X. L., E. Gabrilovich, G. Heitz, W. Horn, K. Murphy, S. Sun, and W. Zhang (2014a). “From data fusion to knowledge fusion”. In: *VLDB '14*. VLDB Endowment.
- Dong, X., E. Gabrilovich, G. Heitz, W. Horn, N. Lao, K. Murphy, T. Strohmman, S. Sun, and W. Zhang (2014b). “Knowledge vault: A web-scale approach to probabilistic knowledge fusion”. In: *KDD '14*. ACM.
- Dong, L., F. Wei, H. Sun, M. Zhou, and K. Xu (2015a). “A hybrid neural model for type classification of entity mentions”. In: *IJCAI '15*. AAAI Press.
- Dong, X. L., E. Gabrilovich, K. Murphy, V. Dang, W. Horn, C. Lugaresi, S. Sun, and W. Zhang (2015b). “Knowledge-based trust: Estimating the trustworthiness of web sources”. In: *VLDB '15*. VLDB Endowment.
- Dutta, S. and G. Weikum (2015). “C3EL: A joint model for cross-document co-reference resolution and entity linking”. In: *EMNLP '15*. ACL.
- Efron, M., C. Willis, and G. Sherman (2014). “Learning sufficient queries for entity filtering”. In: *SIGIR '14*. ACM.

- Ellis, J., J. Getman, and S. Strassel (2014). “Overview of linguistic resources for the TAC KBP 2014 evaluations: Planning, execution, and results”. In: *TAC '14*. LDC.
- Ensan, F. and E. Bagheri (2017). “Document retrieval model through semantic linking”. In: *WSDM '17*.
- Etzioni, O., M. Cafarella, D. Downey, S. Kok, A.-M. Popescu, T. Shaked, S. Soderland, D. S. Weld, and A. Yates (2004). “Web-scale information extraction in Knowitall: (Preliminary results)”. In: *WWW '04*. ACM.
- Etzioni, O., A. Fader, J. Christensen, S. Soderland, and Mausam (2011). “Open information extraction: The second generation”. In: *IJCAI '11*. AAAI Press.
- Fader, A., S. Soderland, and O. Etzioni (2011). “Identifying relations for open information extraction”. In: *EMNLP '11*. ACL.
- Fang, L., A. D. Sarma, C. Yu, and P. Bohannon (2011). “REX: Explaining relationships between entity pairs”. In: *VLDB '11*. VLDB Endowment.
- Fathalla, S., S. Vahdati, S. Auer, and C. Lange (2017). “Towards a knowledge graph representing research findings by semantifying survey articles”. In: *TPDL '17*. Springer-Verlag.
- Ferragina, P. and U. Scaiella (2010). “TAGME: On-the-fly annotation of short text fragments (by Wikipedia entities)”. In: *CIKM '10*. ACM.
- Fetahu, B., K. Markert, and A. Anand (2015). “Automated news suggestions for populating Wikipedia entity pages”. In: *CIKM '15*. ACM.
- Finkel, J. R., T. Grenager, and C. Manning (2005). “Incorporating non-local information into information extraction systems by Gibbs sampling”. In: *ACL '05*. ACL.
- Fissaha Adafre, S. and M. de Rijke (2007). “Ask the crowd to find out what’s important”. In: *ICDM'07 Workshop on Data Mining in Web 2.0 Environments*.
- Flekova, L., O. Ferschke, and I. Gurevych (2014). “What makes a good biography?” In: *WWW '14*. ACM.
- Foley, J., B. O’Connor, and J. Allan (2016). “Improving entity ranking for keyword queries”. In: *CIKM '16*. ACM.

- Franco-Salvador, M., P. Rosso, and R. Navigli (2014). “A knowledge-based representation for cross-language document retrieval and categorization”. In: *EACL '14*. ACL.
- Frank, J. R., M. Kleiman-Weiner, D. A. Roberts, F. Niu, Z. Ce, R. Christopher, and I. Soboroff (2012). “Building an entity-centric stream filtering test collection for TREC 2012”. In: *TREC 2012*. NIST.
- Frank, J. R., M. Kleiman-Weiner, D. A. Roberts, E. Voorhees, and I. Soboroff (2014). “TREC KBA overview”. In: *TREC 2014*. NIST.
- Fuxman, A. (2015). “In situ insights”. In: *SIGIR '15*. 655–664.
- Gabrilovich, E. and S. Markovitch (2007). “Computing semantic relatedness using Wikipedia-based explicit semantic analysis”. In: *IJCAI '07*. AAAI Press.
- Gabrilovich, E. and N. Usunier (2016). “Constructing and mining web-scale knowledge graphs”. In: *SIGIR '16*. ACM.
- Gabrilovich, E., M. Ringgaard, and A. Subramanya (2013). “FACC1: Freebase annotation of ClueWeb corpora, Version 1”. *Tech. rep.* Google Research.
- Galárraga, L. A., N. Preda, and F. M. Suchanek (2013). “Mining rules to align knowledge bases”. In: *AKBC '13*. ACM.
- Ganea, O.-E., M. Horlescu, A. Lucchi, C. Eickhoff, and T. Hofmann (2015). “Probabilistic bag-of-hyperlinks model for entity linking”. In: *WWW '16*. ACM.
- Ganjisaffar, Y., R. Caruana, and C. Lopes (2011). “Bagging gradient-boosted trees for high precision, low variance ranking models”. In: *SIGIR '11*. ACM.
- Gao, J., P. Pantel, M. Gamon, X. He, and L. Deng (2014). “Modeling interestingness with deep neural networks”. In: *EMNLP '14*. ACL.
- Gao, J., X. Li, Y. E. Xu, B. Sisman, X. L. Dong, and J. Yang (2019). “Efficient knowledge graph accuracy evaluation”. In: *VLDB '19*. VLDB Endowment.
- Garcia-Duran, A., S. Dumancic, and M. Niepert (2018). “Learning sequence encoders for temporal knowledge graph completion”. In: *EMNLP '18*.

- Gardner, M. and T. Mitchell (2015). “Efficient and expressive knowledge base completion using subgraph feature extraction”. In: *EMNLP '15*. ACL. 1488–1498.
- Gardner, M., P. P. Talukdar, B. Kisiel, and T. Mitchell (2013). “Improving learning and inference in a large knowledge-base using latent syntactic cues”. In: *EMNLP '13*. ACL.
- Gerritse, E. J., F. Hasibi, and A. P. de Vries (2020). “Graph-embedding empowered entity retrieval”. In: *Advances in Information Retrieval*. Ed. by J. M. Jose, E. Yilmaz, J. Magalhães, P. Castells, N. Ferro, M. J. Silva, and F. Martins. Cham: Springer International Publishing. 97–110.
- Gillick, D., N. Lazic, K. Ganchev, J. Kirchner, and D. Huynh (2014). “Context-dependent fine-grained entity type tagging”. *CoRR* abs/1412.1820.
- Girju, R., P. Nakov, V. Nastase, S. Szpakowicz, P. Turney, and D. Yuret (2007). “SemEval-2007 task 04: Classification of semantic relations between nominals”. In: *SemEval '07*. ACL.
- Globerson, A., N. Lazic, S. Chakrabarti, A. Subramanya, M. Ringgaard, and F. Pereira (2016). “Collective entity resolution with multi-focal attention”. In: *ACL '16*. ACM.
- Graus, D., M. Tsagkias, L. Buitinck, and M. de Rijke (2014). “Generating pseudo-ground truth for predicting new concepts in social streams”. In: *ECIR '14*. Springer-Verlag.
- Graus, D., M. Tsagkias, W. Weerkamp, E. Meij, and M. de Rijke (2016). “Dynamic collective entity representations for entity ranking”. In: *WSDM '16*. ACM.
- Graus, D., D. Odijk, and M. de Rijke (2018). “The birth of collective memories: Analyzing emerging entities in text streams”. *Journal of the Association for Information Science and Technology*. 69(6): 773–786.
- Green, B. F., A. K. Wolf, C. Chomsky, and K. Laughery (1963). “Baseball: An automatic question answerer”. In: *Computers and Thought*. Ed. by E. Figenbaum and J. Fledman. McGraw-Hill. 207–216.
- Grishman, R. and B. Sundheim (1996). “Message understanding conference-6: A brief history”. In: *COLING '96*. ACL.

- Guo, Y., B. Qin, T. Liu, and S. Li (2013). “Microblog entity linking by leveraging extra posts”. In: *EMNLP '13*. ACL.
- GuoDong, Z., S. Jian, Z. Jie, and Z. Min (2005). “Exploring various knowledge in relation extraction”. In: *ACL '05*. ACL.
- Gupta, N., S. Singh, and D. Roth (2017). “Entity linking via joint encoding of types, descriptions, and context”. In: *EMNLP '17*. ACL.
- Hachey, B., J. Nothman, and W. Radford (2014). “Cheap and easy entity evaluation”. In: *ACL '14*. ACL.
- Hakkani-Tür, D., A. Celikyilmaz, L. Heck, G. Tur, and G. Zweig (2014). “Probabilistic enrichment of knowledge graph entities for relation detection in conversational understanding”. In: *Interspeech 2014*. ACM.
- Hasibi, F., K. Balog, and S. E. Bratsberg (2015). “Entity linking in queries: Tasks and evaluation”. In: *ICTIR '15*. ACM.
- Hasibi, F., K. Balog, and S. E. Bratsberg (2016). “Exploiting entity linking in queries for entity retrieval”. In: *ICTIR '16*. ACM.
- Hegde, M. and P. P. Talukdar (2015). “An entity-centric approach for overcoming knowledge graph sparsity”. In: *EMNLP '15*. ACL.
- Heindorf, S., M. Potthast, B. Stein, and G. Engels (2015). “Towards vandalism detection in knowledge bases: Corpus construction and analysis”. In: *SIGIR '15*. ACM.
- Heindorf, S., M. Potthast, B. Stein, and G. Engels (2016). “Vandalism detection in Wikidata”. In: *CIKM '16*. ACM.
- Hoffart, J., M. A. Yosef, I. Bordino, H. Fürstenau, M. Pinkal, M. Spaniol, B. Taneva, S. Thater, and G. Weikum (2011). “Robust disambiguation of named entities in text”. In: *EMNLP '11*. ACL.
- Hoffart, J., Y. Altun, and G. Weikum (2014). “Discovering emerging entities with ambiguous names”. In: *WWW '14*. ACM.
- Hoffart, J., D. Milchevski, G. Weikum, A. Anand, and J. Singh (2016). “The knowledge awakens: Keeping knowledge bases fresh with emerging entities”. In: *WWW '16 Companion*. ACM.
- Hoffmann, R., C. Zhang, X. Ling, L. Zettlemoyer, and D. S. Weld (2011). “Knowledge-based weak supervision for information extraction of overlapping relations”. In: *ACL-HLT '11*. ACL.

- Hovy, D. (2014). “How well can we learn interpretable entity types from text?” In: *ACL '14*. ACL.
- Jatowt, A. and S. Yamamoto (2017). “Overview of NTCIR-13 actionable knowledge graph (AKG) task”. In: *NCTIR*. NTCIR.
- Jenatton, R., N. L. Roux, A. Bordes, and G. Obozinski (2012). “A latent factor model for highly multi-relational data”. In: *NIPS '12*. JLMR.
- Ji, H., R. Grishman, and H. T. Dang (2011). “Overview of the TAC 2011 knowledge base population task”. In: *Proceedings of the TAC-KBP 2011 Workshop*. NIST.
- Ji, G., S. He, L. Xu, K. Liu, and J. Zhao (2015). “Knowledge graph embedding via dynamic mapping matrix”. In: *ACL '15*. ACL.
- Jiang, J. and C.-Y. Lin (2014). “MSR KMG at TREC 2014 KBA track vital filtering task”. In: *TREC 2014*. NIST.
- Jin, Y., E. Kiciman, K. Wang, and R. Loynd (2014). “Entity linking at the tail: Sparse signals, unknown entities, and phrase models”. In: *WSDM '14*. ACM.
- Joshi, M., U. Sawant, and S. Chakrabarti (2014). “Knowledge graph and corpus driven segmentation and answer inference for telegraphic entity-seeking queries”. In: *EMNLP*. ACL. 1104–1114.
- Kambhatla, N. (2004). “Combining lexical, syntactic, and semantic features with maximum entropy models for extracting relations”. In: *ACL '04*. ACL.
- Kang, C., S. Vadrevu, R. Zhang, R. van Zwol, L. G. Pueyo, N. Torzecz, J. He, and Y. Chang (2011). “Ranking related entities for web search queries”. In: *WWW '11*. ACM.
- Kenter, T., M. Wevers, P. Huijnen, and M. de Rijke (2015). “Ad hoc monitoring of vocabulary shifts over time”. In: *CIKM '15*. ACM.
- Kim, J., X. Xue, and W. B. Croft (2009). “A probabilistic retrieval model for semistructured data”. In: *ECIR '09*. Springer-Verlag.
- Kipf, T. N. and M. Welling (2016). “Semi-supervised classification with graph convolutional networks”. arXiv preprint arXiv:1609.02907.
- Kolitsas, N., O.-E. Ganea, and T. Hofmann (2018). “End-to-end neural entity linking”. In: *CoNLL '18*.
- Kong, F., R. Zhang, H. Guo, S. Mensah, Z. Hu, and Y. Mao (2019). “A neural bag-of-words modelling framework for link prediction in knowledge bases with sparse connectivity”. In: *WWW '19*. ACM.

- Kotnis, B., P. Bansal, and P. Talukdar (2015). “Knowledge base inference using bridging entities”. In: *EMNLP '15*. ACL.
- Kripke, S. (1980). *Naming and Necessity*. Wiley-Blackwell.
- Krishna, R., Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, M. Bernstein, and L. Fei-Fei (2016). “Visual genome: Connecting language and vision using crowdsourced dense image annotations”. arXiv preprint arxiv:1602.07332. Feb.
- Kulkarni, S., A. Singh, G. Ramakrishnan, and S. Chakrabarti (2009). “Collective annotation of Wikipedia entities in web text”. In: *KDD '09*. ACM.
- Kuzey, E., J. Vreeken, and G. Weikum (2014). “A fresh look on knowledge bases: Distilling named events from news”. In: *CIKM '14*. ACM.
- Lagun, D., C.-H. Hsieh, D. Webster, and V. Navalpakkam (2014). “Towards better measurement of attention and satisfaction in mobile search”. In: *SIGIR '14*. ACM.
- Lample, G., M. Ballesteros, S. Subramanian, K. Kawakami, and C. Dyer (2016). “Neural architectures for named entity recognition”. In: *HLT-NAACL '16*. ACL. 260–270.
- Lao, N. and W. W. Cohen (2010). “Relational retrieval using a combination of path-constrained random walks”. *Machine Learning*. 81(1): 53–67.
- Lazic, N., A. Subramanya, M. Ringgaard, and F. Pereira (2015). “Plato: A selective context model for entity resolution”. *Transactions of the Association for Computational Linguistics*. 3: 503–515.
- Leblay, J. and M. W. Chekol (2018). “Deriving validity time in knowledge graph”. In: *WWW '18*. ACM.
- Lee, J., A. Fuxman, B. Zhao, and Y. Lv (2015). “Leveraging knowledge bases for contextual entity exploration”. In: *KDD '15*. ACM.
- Lehmann, J., R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P. N. Mendes, S. Hellmann, M. Morsey, P. van Kleef, S. Auer, and C. Bizer (2015). “DBpedia – A large-scale, multilingual knowledge base extracted from Wikipedia”. *Semantic Web Journal*. 6(2): 167–195.

- Li, X., G. Tur, D. Hakkani-Tür, and Q. Li (2014). “Personal knowledge graph population from user utterances in conversational understanding”. In: *SLT '14*. IEEE.
- Li, X., J. Tang, T. Wang, Z. Luo, and M. de Rijke (2015). “Automatically assessing Wikipedia article quality by exploiting article-editor networks”. In: *ECIR '15*. Springer-Verlag.
- Li, F., X. Dong, A. Langen, and Y. D. Li (2017). “Knowledge verification for LongTail verticals”.
- Liao, Z., X. Song, Y. Shen, S. Lee, J. Gao, and C. Liao (2017). “Deep context modeling for web query entity disambiguation”. In: *CIKM '17*. ACM.
- Lin, T., Mausam, and O. Etzioni (2012). “No noun phrase left behind: Detecting and typing unlinkable entities”. In: *EMNLP-CoNLL '12*. ACL.
- Lin, Y., Z. Liu, M. Sun, Y. Liu, and X. Zhu (2015). “Learning entity and relation embeddings for knowledge graph completion”. In: *AAAI '15*. AAAI Press.
- Ling, X. and D. S. Weld (2012). “Fine-grained entity recognition”. In: *AAAI '12*. AAAI Press.
- Liu, X. and H. Fang (2015). “Latent entity space: A novel retrieval approach for entity-bearing queries”. *Information Retrieval Journal*. 18(6): 473–503.
- Liu, X., J. Darko, and H. Fang (2013). “A related entity based approach for knowledge base acceleration”. In: *TREC 2013*. NIST.
- Liu, Q., L. Jiang, M. Han, Y. Liu, and Z. Qin (2016). “Hierarchical random walk inference in knowledge graphs”. In: *SIGIR '16*. ACM.
- Liu, Z., C. Xiong, M. Sun, and Z. Liu (2018). “Entity-duet neural ranking: Understanding the role of knowledge graph semantics in neural information retrieval”. In: *ACL '18*. ACL.
- Luo, G., X. Huang, C.-Y. Lin, and Z. Nie (2015a). “Joint entity recognition and disambiguation”. In: *EMNLP '15*. ACL.
- Luo, Y., Q. Wang, B. Wang, and L. Guo (2015b). “Context-dependent knowledge graph embedding”. In: *EMNLP '15*. ACL.
- Ma, H. and Y. Ke (2015). “An introduction to entity recommendation and understanding”. In: *WWW' 15 Companion*. ACM.

- Ma, W., M. Zhang, Y. Cao, W. Jin, C. Wang, Y. Liu, S. Ma, and X. Ren (2019). “Jointly learning explainable rules for recommendation with knowledge graph”. In: *WWW '19*. ACM.
- Manning, C. D., P. Raghavan, and H. Schütze (2008). *Introduction to Information Retrieval*. Cambridge University Press.
- McCallum, A. and W. Li (2003). “Early results for named entity recognition with conditional random fields, feature induction and web-enhanced lexicons”. In: *CONLL '03*. ACL.
- Medelyan, O., I. H. Witten, and D. Mile (2008). “Topic indexing with Wikipedia”. In: *WIKAI '08*.
- Meij, E., M. Bron, B. Huurnink, L. Hollink, and M. de Rijke (2009). “Learning semantic query suggestions”. In: *ISWC '09*. Springer-Verlag.
- Meij, E., M. Bron, L. Hollink, B. Huurnink, and M. de Rijke (2011). “Mapping queries to the linking open data cloud: A case study using DBpedia”. *Web Semantics: Science, Services and Agents on the World Wide Web*. 9(4): 418–433.
- Meij, E., W. Weerkamp, and M. de Rijke (2012). “Adding semantics to microblog posts”. In: *WSDM '12*. ACM.
- Meij, E., K. Balog, and D. Odijk (2013). “Entity linking and retrieval”. In: *SIGIR '13*. ACM.
- Meij, E., K. Balog, and D. Odijk (2014). “Entity linking and retrieval for semantic search”. In: *WSDM '14*. ACM.
- Melo, G. de and N. Tandon (2016). “Seeing is believing: The quest for multimodal knowledge”. *SIGWEB Newsletter*. (Spring): 4:1–4:9.
- Metzler, D. and W. B. Croft (2005). “A Markov random field model for term dependencies”. In: *SIGIR '05*. ACM.
- Metzler, D. and W. B. Croft (2007). “Latent concept expansion using Markov random fields”. In: *SIGIR '07*. ACM.
- Mihalcea, R. and A. Csomai (2007). “Wikify!: Linking documents to encyclopedic knowledge”. In: *CIKM '07*. ACM.
- Mikolov, T., K. Chen, G. Corrado, and J. Dean (2013). “Distributed representations of words and phrases and their compositionality”. In: *EMNLP '04*. ACL.

- Miliaraki, I. and R. Blanco (2015). “From Selena Gomez to Marlon Brando: Understanding explorative entity search”. In: *WWW '15*. ACM.
- Milne, D. and I. H. Witten (2008). “Learning to link with Wikipedia”. In: *CIKM '08*. ACM.
- Minard, A.-L., M. Speranza, R. Urizar, B. Altuna, M. van Erp, A. Schoen, and C. van Son (2016). “MEANTIME, the NewsReader multilingual event and time corpus”. In: *LREC '16*. ELRA.
- Mintz, M., S. Bills, R. Snow, and D. Jurafsky (2009). “Distant supervision for relation extraction without labeled data”. In: *ACL '09*. ACL.
- Mitra, B. and N. Craswell (2018). “An introduction to neural information retrieval”. *Foundations and Trends in Information Retrieval*. 13(1): 1–126.
- Mohapatra, H., S. Jain, and S. Chakrabarti (2013). “Joint bootstrapping of corpus annotations and entity types”. In: *EMNLP '13*. ACL.
- Nakashole, N., T. Tylenda, and G. Weikum (2013). “Fine-grained semantic typing of emerging entities”. In: *ACL '13*. ACL.
- Navalpakkam, V., L. Jentzsch, R. Sayres, S. Ravi, A. Ahmed, and A. Smola (2013). “Measurement and modeling of eye-mouse behavior in the presence of nonlinear page layouts”. In: *WWW '13*. ACM.
- Nickel, M., V. Tresp, and H.-P. Kriegel (2011). “A three-way model for collective learning on multi-relational data”. In: *ICML '11*. Omnipress.
- Nickel, M., V. Tresp, and H.-P. Kriegel (2012). “Factorizing YAGO: Scalable machine learning for linked data”. In: *WWW '12*. ACM.
- Nickel, M., K. Murphy, V. Tresp, and E. Gabrilovich (2016). “A review of relational machine learning for knowledge graphs”. *Proceedings of the IEEE*. 104(1): 11–33.
- Nikolaev, F., A. Kotov, and N. Zhiltsov (2016). “Parameterized fielded term dependence models for ad-hoc entity retrieval from knowledge graph”. In: *SIGIR '16*. ACM.
- Odijk, D., E. Meij, I. Sijaranamual, and M. de Rijke (2015). “Dynamic query modeling for related content finding”. In: *SIGIR '15*. ACM.

- Ojha, P. and P. P. Talukdar (2017). “KGEval – Estimating accuracy of automatically constructed knowledge graphs”. In: *EMNLP '17. ACL*.
- 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. de Rijke, and M. Lease (2018). “Neural information retrieval: At the end of the early years”. *Information Retrieval Journal*. 21(2–3): 111–182.
- Oosterhuis, H. and M. de Rijke (2018). “Ranking for relevance and display preferences in complex presentation layouts”. In: *SIGIR '18. ACM*.
- Pantel, P. and A. Fuxman (2011). “Jigs and Lures: Associating web queries with structured entities”. In: *ACL-HLT '11. ACL*.
- Pasca, M., D. Lin, J. Bigham, A. Lifchits, and A. Jain (2006). “Organizing and searching the world wide web of facts – step one: The one-million fact extraction challenge”. In: *AAAI'06*. AAAI Press.
- Passos, A., V. Kumar, and A. McCallum (2014). “Lexicon infused phrase embeddings for named entity resolution”. In: *CoNLL '14. ACL*.
- Petkova, D. and W. B. Croft (2008). “Hierarchical language models for expert finding in enterprise corpora”. *International Journal on Artificial Intelligence Tools*. 17: 5–18.
- Potthast, M., B. Stein, and R. Gerling (2008). “Automatic vandalism detection in Wikipedia”. In: *ECIR '08*. Springer-Verlag.
- Pound, J., P. Mika, and H. Zaragoza (2010). “Ad-hoc object retrieval in the web of data”. In: *WWW '10. ACM*.
- Pound, J., A. K. Hudek, I. F. Ilyas, and G. Weddell (2012). “Interpreting keyword queries over web knowledge bases”. In: *CIKM '12. ACM*.
- Prokofyev, R., G. Demartini, and P. Cudré-Mauroux (2014). “Effective named entity recognition for idiosyncratic web collections”. In: *WWW '14. ACM*.
- Pujara, J., E. Augustine, and L. Getoor (2017). “Sparsity and noise: Where knowledge graph embeddings fall short”. In: *EMNLP '15. ACL*.
- Ran, C. and J. Wang (2018). “An attention factor graph model for tweet entity linking”. In: *WWW '18. ACM*.

- Ratinov, L., D. Roth, D. Downey, and M. Anderson (2011). “Local and global algorithms for disambiguation to Wikipedia”. In: *ACL-HLT '11*. ACL.
- Raviv, H., O. Kurland, and D. Carmel (2016). “Document retrieval using entity-based language models”. In: *SIGIR '16*. ACM.
- Reinanda, R., E. Meij, and M. de Rijke (2015). “Mining, ranking and recommending entity aspects”. In: *SIGIR '15*. ACM.
- Reinanda, R., E. Meij, and M. de Rijke (2016). “Document filtering for long-tail entities”. In: *CIKM '16*. ACM.
- Ren, X., A. El-Kishky, C. Wang, F. Tao, C. R. Voss, and J. Han (2015). “ClusType: Effective entity recognition and typing by relation phrase-based clustering”. In: *KDD '15*. ACM.
- Ren, X., W. He, M. Qu, L. Huang, H. Ji, and J. Han (2016a). “AFET: Automatic fine-grained entity typing by hierarchical partial-label embedding”. In: *EMNLP '16*. ACL.
- Ren, X., W. He, M. Qu, C. R. Voss, H. Ji, and J. Han (2016b). “Label noise reduction in entity typing by heterogeneous partial-label embedding”. In: *KDD '16*. ACM.
- Riedel, S., L. Yao, and A. McCallum (2010). “Modeling relations and their mentions without labeled text”. In: *ECML PKDD '10*. Springer-Verlag.
- Riedel, S., L. Yao, B. M. Marlin, and A. McCallum (2013). “Relation extraction with matrix factorization and universal schemas”. In: *HLT-NAACL '13*. ACL.
- Ribeiro, M. T., S. Singh, and C. Guestrin (2016). “‘Why should I trust you?’: Explaining the predictions of any classifier”. In: *KDD '16*. ACM.
- Riloff, E. and R. Jones (1999). “Learning dictionaries for information extraction by multi-level bootstrapping”. In: *AAAI '99*. AAAI Press.
- Sanderson, M. (2015). “Test collection based evaluation of information retrieval systems”. *Foundations and Trends in Information Retrieval*. 4(4): 247–375.
- Sarawagi, S. and W. W. Cohen (2004). “Semi-Markov conditional random fields for information extraction”. In: *NIPS '04*. JLMR.
- Sarmiento, L., V. Jijkoun, M. de Rijke, and E. Oliviera (2007). “‘More like these’: Growing entity classes from seeds”. In: *CIKM '07*. ACM.

- Sarrafzadeh, B., O. Vechtomova, and V. Jokic (2014). “Exploring knowledge graphs for exploratory search”. In: *III'X '14*. ACM.
- Schuhmacher, M., L. Dietz, and S. Paolo Ponzetto (2015). “Ranking entities for web queries through text and knowledge”. In: *CIKM '15*. ACM.
- Sekine, S. (1998). “NYU: Description of the Japanese NE system used for MET-2”. In: *MUC-7*. ACL.
- Sekine, S. (2009). *Named Entities: Recognition, Classification and Use*. John Benjamin Publishings.
- Sekine, S. and C. Nobata (2004). “Definition, dictionaries and tagger for extended named entity hierarchy”. In: *LREC*. ACL.
- Seufert, S., K. Berberich, S. J. Bedathur, S. K. Kondreddi, P. Ernst, and G. Weikum (2016). “ESPRESSO: Relationships between entity sets”. In: *CIKM '16*. ACM.
- Sevgili, Ö., A. Panchenko, and C. Biemann (2019). “Improving neural entity disambiguation with graph embeddings”. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*. Florence, Italy: Association for Computational Linguistics. 315–322. URL: <https://www.aclweb.org/anthology/P19-2044>.
- Shen, J., J. Xiao, X. He, J. Shang, S. Sinha, and J. Han (2018). “Entity set search of scientific literature: An unsupervised ranking approach”. In: *SIGIR '18*.
- Shimaoka, S., P. Stenetorp, K. Inui, and S. Riedel (2016). “An attentive neural architecture for fine-grained entity type classification”. In: *AKBC '16*. ACL.
- Shimaoka, S., P. Stenetorp, K. Inui, and S. Riedel (2017). “Neural architectures for fine-grained entity type classification”. In: *EACL '17*. ACL.
- Shirakawa, M., K. Nakayama, T. Hara, and S. Nishio (2013). “Probabilistic semantic similarity measurements for noisy short texts using Wikipedia entities”. In: *CIKM '13*. ACM.
- Shokouhi, M. and Q. Guo (2015). “From queries to cards: Re-ranking proactive card recommendations based on reactive search history”. In: *SIGIR '15*. ACM.

- Sil, A. and A. Yates (2013). “Re-ranking for joint named-entity recognition and linking”. In: *CIKM '13*. ACM.
- Singh, S., A. Subramanya, F. Pereira, and A. McCallum (2012). “Wikilinks: A large-scale cross-document coreference corpus labeled via links to Wikipedia”. *Tech. rep.* No. UM-CS-2012-015. University of Massachusetts, Amherst.
- Singh, J., J. Hoffart, and A. Anand (2016). “Discovering entities with just a little help from you”. In: *CIKM '16*. ACM.
- Socher, R., B. Huval, C. D. Manning, and A. Y. Ng (2012). “Semantic compositionality through recursive matrix-vector spaces”. In: *CoNLL '12*. ACL.
- Socher, R., D. Chen, C. D. Manning, and A. Y. Ng (2013). “Reasoning with neural tensor networks for knowledge base completion”. In: *NIPS '13*. Curran Associates Inc.
- Strubell, E., P. Verga, D. Belanger, and A. McCallum (2017). “Fast and accurate entity recognition with iterated dilated convolutions”. In: *ACL '17*. ACL.
- Suchanek, F. M., G. Kasneci, and G. Weikum (2007). “Yago: A core of semantic knowledge”. In: *WWW '07*. ACM.
- Sun, H., H. Ma, W.-T. Yih, C.-T. Tsai, J. Liu, and M.-W. Chang (2015). “Open domain question answering via semantic enrichment”. In: *WWW '15*. ACM.
- Surdeanu, M., S. Gupta, J. Bauer, D. McClosky, A. X. Chang, V. I. Spitzkovsky, and C. D. Manning (2011). “Stanford’s distantly-supervised slot-filling system”. In: *Proceedings of the TAC-KBP 2011 Workshop*. NIST.
- Surdeanu, M., J. Tibshirani, R. Nallapati, and C. D. Manning (2012). “Multi-instance multi-label learning for relation extraction”. In: *CoNLL '12*. ACL.
- Sutskever, I., J. B. Tenenbaum, and R. R. Salakhutdinov (2009). “Modelling relational data using bayesian clustered tensor factorization”. In: *Advances in Neural Information Processing Systems. NIPS '09*. 1821–1828.
- Tan, C. H., E. Agichtein, P. Ipeirotis, and E. Gabrilovich (2014). “Trust, but verify: Predicting contribution quality for knowledge base construction and curation”. In: *WSDM '14*. ACM.

- Tang, J., Z. Fang, and J. Sun (2015). “Incorporating social context and domain knowledge for entity recognition”. In: *WWW '15*. ACM.
- Tao, F., B. Zhao, A. Fuxman, Y. Li, and J. Han (2015). “Leveraging pattern semantics for extracting entities in enterprises”. In: *WWW '15*. ACM.
- Tjong Kim Sang, E. F. and F. De Meulder (2003). “Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition”. In: *CoNLL '03*. ACL.
- Tong, H. and C. Faloutsos (2006). “Center-piece subgraphs: Problem definition and fast solutions”. In: *KDD '06*. ACM.
- Tonon, A., G. Demartini, and P. Cudré-Mauroux (2012). “Combining inverted indices and structured search for ad-hoc object retrieval”. In: *SIGIR '12*. ACL.
- Toutanova, K. and D. Chen (2015). “Observed versus latent features for knowledge base and text inference”. In: *3rd Workshop on Continuous Vector Space Models and Their Compositionality*. ACL.
- Toutanova, K., D. Chen, P. Pantel, H. Poon, P. Choudhury, and M. Gamon (2015). “Representing text for joint embedding of text and knowledge bases”. In: *EMNLP '15*. ACL.
- Usbeck, R., M. Röder, A.-C. Ngonga Ngomo, C. Baron, A. Both, M. Brümmer, D. Ceccarelli, M. Cornolti, D. Cherix, B. Eickmann, P. Ferragina, C. Lemke, A. Moro, R. Navigli, F. Piccinno, G. Rizzo, H. Sack, R. Speck, R. Troncy, J. Waitelonis, and L. Wesemann (2015). “GERBIL: General entity annotator benchmarking framework”. In: *WWW '15*. ACM.
- Van Gysel, C., M. de Rijke, and E. Kanoulas (2016a). “Learning latent vector spaces for product search”. In: *CIKM '16*. ACM.
- Van Gysel, C., M. de Rijke, and M. Worring (2016b). “Unsupervised, efficient and semantic expertise retrieval”. In: *WWW '16*. ACM.
- Van Gysel, C., M. de Rijke, and E. Kanoulas (2017a). “Semantic entity retrieval toolkit”. In: *SIGIR 2017 Workshop on Neural Information Retrieval*. ACM.
- Van Gysel, C., M. de Rijke, and E. Kanoulas (2017b). “Structural regularities in text-based entity vector spaces”. In: *ICTIR '17*. ACM.

- Van Gysel, C., E. Kanoulas, and M. de Rijke (2018). “A neural vector space model”. *ACM Transactions on Information Systems*. 36(4): 38.
- Vannella, D., D. Jurgens, D. Scarfini, D. Toscani, and R. Navigli (2014). “Validating and extending semantic knowledge bases using video games with a purpose”. In: *ACL '14*. ACL.
- Voorhees, E. M. (2005). “The TREC robust retrieval track”. *SIGIR Forum*. 39(1): 11–20.
- Voskarides, N., E. Meij, M. Tsagkias, M. de Rijke, and W. Weerkamp (2015). “Learning to explain entity relationships in knowledge graphs”. In: *ACL '15*. ACL.
- Voskarides, N., E. Meij, and M. de Rijke (2017). “Generating descriptions of entity relationships”. In: *ECIR '17*. Springer-Verlag.
- Wang, J., D. Song, C. Lin, and L. Liao (2013a). “BIT and MSRA at TREC KBA CCR track 2013”. In: *TREC*. NIST.
- Wang, J., D. Song, Q. Wang, Z. Zhang, L. Si, L. Liao, and C.-Y. Lin (2015a). “An entity class-dependent discriminative mixture model for cumulative citation recommendation”. In: *SIGIR '15*. ACM.
- Wang, R. C. and W. W. Cohen (2007). “Language-independent set expansion of named entities using the web”. In: *ICDM '07*. IEEE.
- Wang, T.-X., K.-Y. Tsai, and W.-H. Lu (2014a). “Identifying real-life complex task names with task-intrinsic entities from microblogs”. In: *ACL '14*. ACL.
- Wang, X., D. Wang, C. Xu, X. He, Y. Cao, and T.-S. Chua (2018). “Explainable reasoning over knowledge graphs for recommendation”. In: *AAAI '18*. AAAI Press.
- Wang, X., X. L. Dong, and A. Meliou (2015b). “Data X-ray: A diagnostic tool for data errors”. In: *SIGMOD '15*. ACM.
- Wang, Z., J. Li, Z. Wang, and J. Tang (2012). “Cross-lingual knowledge linking across Wiki knowledge bases”. In: *WWW '12*. ACM.
- Wang, Z., J. Zhang, J. Feng, and Z. Chen (2014b). “Knowledge graph embedding by translating on hyperplanes”. In: *AAAI '14*. AAAI Press.
- Wang, Z., Z. Li, J. Li, J. Tang, and J. Z. Pan (2013b). “Transfer learning based cross-lingual knowledge extraction for Wikipedia”. In: *ACL '13*. ACL.

- West, R., E. Gabrilovich, K. Murphy, S. Sun, R. Gupta, and D. Lin (2014). “Knowledge base completion via search-based question answering”. In: *WWW '14*. ACM.
- Woods, W. A. (1977). “Lunar rocks in natural english: Explorations in natural language question answering”. In: *Linguistic Structures Processing*. Ed. by A. Zampoli. Elsevier North-Holland. 521–569.
- Wu, F. and D. S. Weld (2010). “Open information extraction using Wikipedia”. In: *ACL '10*. ACL.
- Wu, Q., D. Teney, P. Wang, C. Shen, A. R. Dick, and A. van den Hengel (2016a). “Visual question answering: A survey of methods and datasets”. arXiv preprint arXiv:1607.05910.
- Wu, Z., Y. Song, and C. L. Giles (2016b). “Exploring multiple feature spaces for novel entity discovery”. In: *AAAI '16*. AAAI Press.
- Xie, X., J. Mao, M. de Rijke, R. Zhang, M. Zhang, and S. Ma (2018). “Constructing an interaction behavior model for web image search”. In: *SIGIR '18*. ACM.
- Xiong, C. and J. Callan (2015a). “EsdRank: Connecting query and documents through external semi-structured data”. In: *CIKM '15*. ACM.
- Xiong, C. and J. Callan (2015b). “Query expansion with freebase”. In: *ICTIR '15*. ACM.
- Xiong, C., J. Callan, and T.-Y. Liu (2017a). “Word-entity duet representations for document ranking”. In: *SIGIR '17*. ACM.
- Xiong, C., R. Power, and J. Callan (2017b). “Explicit semantic ranking for academic search via knowledge graph embedding”. In: *WWW '17*. ACM.
- Xu, P. and D. Barbosa (2018). “Neural fine-grained entity type classification with hierarchy-aware loss”. In: *EMNLP '18*. ACL.
- Yaghoobzadeh, Y. and H. Schutze (2015). “Corpus-level fine-grained entity typing using contextual information”. In: *EMNLP '15*. ACL.
- Yahya, M., D. Barbosa, K. Berberich, Q. Wang, and G. Weikum (2016). “Relationship queries on extended knowledge graphs”. In: *WSDM '16*. ACM.
- Yang, B., W. Yih, X. He, J. Gao, and L. Deng (2015). “Embedding entities and relations for learning and inference in knowledge bases”. In: *ICLR '15*. ICLR.

- Yao, X. and B. Van Durme (2014). “Information extraction over structured data: Question answering with freebase”. In: *ACL '14*. ACL.
- Yao, L., S. Riedel, and A. McCallum (2010). “Collective cross-document relation extraction without labelled data”. In: *EMNLP '10*. ACL.
- Yih, W.-T., M.-W. Chang, X. He, and J. Gao (2015). “Semantic parsing via staged query graph generation: Question answering with knowledge base”. In: *ACL '15*. ACL.
- 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*. ACM. 974–983.
- Yogatama, D., D. Gillick, and N. Lazic (2015). “Embedding methods for fine grained entity type classification”. In: *ACL '15*. ACL.
- Yosef, M. A., S. Bauer, J. Hoffart, M. Spaniol, and G. Weikum (2012). “HYENA: Hierarchical type classification for entity names”. In: *COLING '12*. ACL.
- Yu, X., H. Ma, B.-J. P. Hsu, and J. Han (2014a). “On building entity recommender systems using user click log and freebase knowledge”. In: *WSDM '14*. ACM.
- Yu, X., X. Ren, Y. Sun, Q. Gu, B. Sturt, U. Khandelwal, B. Norick, and J. Han (2014b). “Personalized entity recommendation: A heterogeneous information network approach”. In: *WSDM '14*. ACM.
- Zelenko, D., C. Aone, and A. Richardella (2002). “Kernel methods for relation extraction”. In: *EMNLP '02*. ACL.
- Zeng, D., K. Liu, Y. Chen, and J. Zhao (2015). “Distant supervision for relation extraction via piecewise convolutional neural networks”. In: *EMNLP '15*. ACL.
- Zhang, C., G. Zhou, Q. Lu, and F. Chang (2017a). “Graph-based knowledge reuse for supporting knowledge-driven decision-making in new product development”. *International Journal of Production Research*. 55(23): 7187–7203.
- Zhang, F., N. J. Yuan, D. Lian, X. Xie, and W. Y. Ma (2016a). “Collaborative knowledge base embedding for recommender systems”. In: *KDD '16*. ACM.

- Zhang, L., A. Rettinger, and J. Zhang (2016c). “A knowledge base approach to cross-lingual keyword query interpretation”. In: *ISWC '16*. Springer-Verlag.
- Zhang, L., M. Färber, and A. Rettinger (2016b). “XKnowSearch!: Exploiting knowledge bases for entity-based cross-lingual information retrieval”. In: *CIKM '16*. ACM.
- Zhang, W., B. Paudel, L. Wang, J. Chen, H. Zhu, W. Zhang, A. Bernstein, and H. Chen (2019). “Iteratively learning embeddings and rules for knowledge graph reasoning”. In: *WWW '19*. ACL.
- Zhang, Y., T. Paradis, L. Hou, J. Li, J. Zhang, and H. Zheng (2017b). “Cross-lingual infobox alignment in Wikipedia using entity-attribute factor graph”. In: *ISWC '17*. Springer-Verlag.
- Zhao, S. and R. Grishman (2005). “Extracting relations with integrated information using kernel methods”. In: *ACL '05*. ACL.
- Zhao, S. and Y. Zhang (2014). “Tailor knowledge graph for query understanding: Linking intent topics by propagation”. In: *EMNLP '14*. ACL.
- Zhiltsov, N. and E. Agichtein (2013). “Improving entity search over linked data by modeling latent semantics”. In: *CIKM '13*. ACM.
- Zhiltsov, N., A. Kotov, and F. Nikolaev (2015). “Fielded sequential dependence model for ad-hoc entity retrieval in the web of data”. In: *SIGIR '15*. ACM.
- Zhong, H., J. Zhang, Z. Wang, H. Wan, and Z. Chen (2015). “Aligning knowledge and text embeddings by entity descriptions”. In: *EMNLP '15*. ACL.
- Zhou, M. and K. C.-C. Chang (2013). “Entity-centric document filtering: Boosting feature mapping through meta-features”. In: *CIKM '13*. ACM.
- Zhu, G. and C. A. Iglesias (2018). “Exploiting semantic similarity for named entity disambiguation in knowledge graphs”. *Expert Systems with Applications*. 101(July): 8–24.
- Zhu, Y., C. Zhang, C. Ré, and F.-F. Li (2015). “Building a large-scale multimodal knowledge base for visual question answering”. arXiv preprint arXiv:1507.05670.