Utilizing Knowledge Graphs for Text-Centric Information Retrieval

Laura Dietz
University of New Hampshire
Durham, NH, USA
dietz@cs.unh.edu

Alexander Kotov
Wayne State University
Detroit, MI, USA
kotov@wayne.edu

Edgar Meij
Bloomberg
London, UK
edgar.meij@acm.org

ABSTRACT

The past decade has witnessed the emergence of several publicly available and proprietary knowledge graphs (KGs). The depth and breadth of content in these KGs made them not only rich sources of structured knowledge by themselves, but also valuable resources for search systems. A surge of recent developments in entity linking and entity retrieval methods gave rise to a new line of research that aims at utilizing KGs for text-centric retrieval applications. This tutorial is the first to summarize and disseminate the progress in this emerging area to industry practitioners and researchers.

CCS CONCEPTS

• **Information systems** → *Presentation of retrieval results*;

KEYWORDS

Knowledge graphs, Entity Linking, Entity Retrieval, Information Retrieval

ACM Reference Format:
Laura Dietz, Alexander Kotov, and Edgar Meij. 2018. Utilizing Knowledge Graphs for Text-Centric Information Retrieval. In *SIGIR '18: 41st International ACM SIGIR Conference on Research and Development in Information Retrieval, July 8-12, 2018, Ann Arbor, MI, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3209978.3210187

1 INTRODUCTION

Large knowledge graphs (KGs) and scalable entity linking technology are powerful tools for deeper understanding of the semantics of text. While these tools can be used for a wide range of tasks, in this tutorial, we focus on how knowledge graphs and entity links are most effectively utilized for text-centric information retrieval (IR).

We use the term *entity* to denote any entry in a KG, while distinguishing it from a *mention* of an entity in text (which was previously referred to as entity in the literature about named entity recognition). As such, we leverage a generic, extended definition of entities to encompass any entry in a KG, which may include, for example, people and places, but also chemical compounds, diseases, as well as intangible concepts, such as “information retrieval”. KGs also provide information on *relations* between entities, which can be typed according to a schema or declared as links without further semantic specification. Such relations are represented as edges (or hyperedges) in the KG. They may be entered manually into the graph, sourced from various data bases, or extracted from text for automatic knowledge graph population. *Entity retrieval*, then, refers to the task of retrieving relevant KG entries in response to a user query. *Entity linking* refers to the annotation of text such that all entity mentions are annotated with identifiers to KG entries.

Starting with the INEX [13], TREC, and TAC KBP initiatives [4, 29], the tasks of entity linking and retrieval have gained momentum. Although it has been previously demonstrated that KGs can be utilzed as a source of expansion terms and smoothing [1, 2, 6, 33, 57], improvements in entity linking and retrieval methods have led to a series of successes in the utilization of entity relations, descriptions and types in ad hoc text-centric retrieval scenarios [12, 38, 49, 55]. Accurate entity linking methods play a critical role in this scenario, as they provide a bridge between unstructred text and structured information about entities in KGs.

These successes led to the emergence of a new line of research on how to effectively utilize entity-centric knowledge repositories to understand textual data and estimate entity-based relevance to a given information need. A large number of recent advances in this field makes this an ideal time to summarize and report the state-of-the-art approaches to the community. Methods and approaches outlined in this tutorial provide a foundation for future advances in several text-centric retrieval tasks, ranging from discovering emerging entities [28], resolving query aspects [50], organizing content into topics [3, 16] as well as entity-aware ad hoc document retrieval [12, 38, 49, 55].

We also touch on the issue of semantic search by providing an overview of novel and recent advances in entity retrieval that are not covered in previous tutorials on this topic. This tutorial focuses on the use of KGs for text-centric information retrieval and, more specifically, on how to leverage different types of data provided by KGs for ad hoc document retrieval and other search systems. We refer to the KG4IR Workshop for ongoing work in the area [17].[1] The tutorial is divided into four parts: a) entity linking, b) entity retrieval, c) utilizing entities in text-centric information retrieval, and d) open research areas, which are discussed in the following sections.

2 ENTITY LINKING

There exists a wide variety of general-purpose encyclopedic knowledge graphs, such as Freebase, DBpedia, WikiData, Yago, Microsoft's

ACM ISBN 978-1-4503-5657-2/18/07.
https://doi.org/10.1145/3209978.3210187

[1] See https://kg4ir.github.io/.

Satori, and Google's Knowledge graph, and domain-specific knowledge bases, such as the Unified Medical Language System. Linguistic and other knowledge can also be encoded in controlled vocabularies and semantic networks, such as MeSH, WordNet, Babelnet and ConceptNet. Most knowledge graphs have an underlying ontology that specifies types, relations, and meta data, while others are simply structured as networks of concepts or collections of entities.

While each KG has unique characteristics, KG entities are typically associated with different names and, possibly, types from a taxonomy or category system, as well as have relations with other entities. Some KGs also incorporate explicit textual descriptions for entities and/or links to textual documents that are "about" each entity. Throughout the tutorial, we discuss how each of these different types of information can be used to: a) retrieve a set of entities for an information need formulated as a keyword query or a question, or more broadly: how to assess the relevance of KB elements to a given topic, b) how to recognize mentions of entities from a KG in a textual fragment and c) how to utilize these mentions to assess the relevance of a textual fragment.

Entity linking [48] is the task of identifying entity mentions in text and aligning them with their corresponding entities in the knowledge graph. Entity linking systems are typically structured as a pipeline. The first step is to identify linkable phrases, i.e., text segments that could mention an entity. In the second step, a candidate set for each such phrase is retrieved, of course the possibility that the mentioned entity is not contained in the knowledge graph (so-called NIL entities) must be considered. The final step is to disambiguate which of the candidate entities are actually referred in the mention based on the context of the mention. We discuss a variety of best practices and methods, such as topic models and word embeddings.

A range of entity linking toolkits are available for documents [20], queries [24, 39], and microblog posts [8]. In addition, large collections of entity link annotation for ClueWeb [21] led to reproducible research on retrieval models that utilize knowledge graphs.

3 ENTITY RETRIEVAL

A large fraction of queries posed to Web search system aim at finding an entity or a set of entities, which can be directly retrieved from a KG [47]. Such queries may refer to the target entities by their names, attributes or related entities and be expressed in the form of keywords or a question [32]. The resulting ranked entities can be either presented to the user directly or utilized as a source of query expansion terms for text-centric retrieval.

Since knowledge graph entities are not the same as text documents or Web pages, new retrieval models are required, which are often referred to as object retrieval models. A canonical approach is to combine heterogeneous and semi-structured information about an entity (e.g. its name aliases, attributes, categories, outgoing/incoming links, and content) into a static [42, 45, 58] or dynamic [23] multi-field entity representation. Such entity representations can be retrieved using specialized structured document retrieval models, such as the Fielded Sequential Dependence Model [58] and its feature based variant [45]. Entity retrieval models can also effectively utilize entity links in queries [24] or type hints [22, 25]. Furthermore, retrieved entities can be diversified by taking the distance in the knowledge graph into consideration [34].

Entities can also be retrieved using a corpus-based pseudo-relevance feedback approach, in which feedback documents are analyzed for entity links [51]. Using entities to retrieve relevant text can be viewed as an inverse problem to retrieving entities through relevant text [11].

4 UTILIZING ENTITIES IN TEXT RETRIEVAL

In this tutorial we focus on three core angles of text-centric IR systems: a) keyword matching and smoothing models, b) query expansion models using pseudo-relevance feedback and query logs, and c) components for diversification and redundancy removal. Most work on these fronts operate at the level of terms and phrases. However, recent developments in entity linking algorithms and object retrieval make it feasible to efficiently tap into the rich information provided by KGs.

Research on vertical, composite, and aggregate search provides an alternative perspective on the problem, where the main task is to combine information from various resources. One central idea is the formation of information bundles [9] by using of entities as pivots of information and for diversification. In some cases, a knowledge base is interpreted as a further vertical for retrieval [44].

Previously proposed systems successfully leverage knowledge bases to improve ad-hoc document retrieval. These systems combine the notion of entity retrieval and semantic search on one hand, with text retrieval models and entity linking on the other. Sometimes users may find it helpful to explicitly include KG entities into their free text queries [5], or track an entity over time [15].

KG-aware document retrieval models incorporate matches of entity names, contextual terms, and entity links. Together with approaches for finding relevant entities these give rise to an effective generalizable retrieval approach.

Different machine learning approaches aid in solving this task. Concept Feedback [33] uses a feature-based system with graph walks. Latent Entity Space [38] uses generative language models. EsdRank [55] and Entity Query Feature Expansion [12] integrate entity retrieval, text retrieval, and different indicators from KG-based query expansion with a supervised learning-to-rank approach. Inference on semantic networks, latently relevant entities, entity types can be integrated into such systems [31, 57]. Language models are built over uncertain entity links [49]. Furthermore, statistical term association graphs with knowledge bases provide additional sources for query expansion [2]. Utilizing relation extraction into text retrieval bears potential but remains a challenge [30, 52].

Several approaches to neural networks have for information retrieval been introduced lately, including extentions for incorporating entity annotations. An example is the combination of entity and document predictions in a duet model [56].

5 OPEN RESEARCH AREAS

5.1 Graph structure and Graph Walks

Many knowledge bases contain both untyped hyperlinks as well as typed relational facts between entities, the former appearing in abundance and the latter being often sparse and biased to entities of particular types. The graph structure can help understand the context, as long as concept drift can be avoided [33, 35, 43]. However, modern knowledge graphs cover a wide range of relations—some noteworthy, some ideosyncratic. As a result, many spurious edges in the knowledge graph lead to severe concept drift when graph walk algorithms are employed. While machine learning-based methods for focused exploration of term graphs and semantic networks have been previously proposed [1], filtering out non-relevant edges when traversing knowledge graphs for query expansion remains an open problem.

5.2 Relations and Relation Extractions

Most progress towards utilizing relations, i.e., the edges in the knowledge graph, have been made in question answering [7]. Although even question answering benefits from more accurate prediction of relevant aspects and types of relations [36].

Relation extraction systems [40] can provide us with additional relations as extracted from text as well as textual evidences we might want to retrieve. Additionally, with the advent of schemaless, so called "open information extraction" methods [19], more links with term-associations are becoming available. Schuhmacher et al. [52] found that schema-based relation extraction can be used to find relevant relations for a query, but is applicable only to 60% of web queries. In contrast, Voskarides et al [54] focuses on the inverse problem of retrieving support passages for given relations. Kadry and Dietz [30] demonstrate that, for retrieval of support passages for entity relevance, open relation extraction improves precision. However, a range of limitations of relation extraction technology affects performance of retrieval systems.

5.3 Entity Aspects

Many entities have different aspects [37, 41, 50] of which only one needs to be relevant in order to render the entity relevant for the query. For example, the United Kingdom might be known for being a European country with a constitutional monarchy in some contexts, or as a financial metropolis, or even as a country that appreciates punting as a spare time activity. Even with perfect disambiguation choices of the entity linking algorithms, it remains crucial to understand which aspect is relevant and how they are expressed in text to assess relevance for the information need.

Liu and Fang [38] explore a range of contextual language models to model query-relevant aspects. Duan and Zhai [18] estimate coordinated intents associated with entities for a given information need. In contrast, Nanni et al [41] harvest headings from Wikipedia articles as explicit aspects.

5.4 Query Subtopics

Research on diversification relies on the identification of different subtopics within query-relevant material. Entity-centric approaches can be applied to topic detection [46]. Especially for complex information needs, it becomes more important to organize topics for coherent presentation [3, 16]. It seems sensible that knowledge graphs can help here, but more work on utilizing KGs in the identification of query sub-topics is needed.

5.5 Conversational Search

Conversational search and dialog systems for information seeking would likely benefit from utilizing knowledge graphs. The KB-InfoBot is a dialog system that helps users find entities of interest [14]. Knowledge plays an important role in telling a story in exploratory search systems [53]. Identifying relevant entities for the user query helps to find relevant information to produce an utterance through natural language generation [10]. Many open questions center around how to use past interactions to estimate a user-specific knowledge graph [26], and inquire information about yet unknown entities and relations [27].

6 SUPPORTING MATERIALS

The supporting materials include: a) a collections of tools, b) collections of data and annotation sets, 3) lecture notes, and 4) an annotated bibliography.[2] We are also moderating a "kg4ir" google group mailinglist for follow-up questions and discussions.

7 CONCLUSION

The recent progress in entity linking and retrieval ensured robust access of IR systems to vast amounts of information stored in KGs. Since utilization of this information has been recently shown to yield improvements in many IR tasks, the main goal of this tutorial is to educate the community about these important results.

8 PRESENTERS

Prof. Dr. Laura Dietz is an Assistant Professor at University of New Hampshire, where she teaches Information Retrieval and Data Science. Before that she was working in the Data and Web Science group at Mannheim University, with Prof. Bruce Croft and Prof. Andrew McCallum at University of Massachusetts, and obtained her Ph.D. from the Max Planck Institute for Informatics. Her research focuses on text processing and information retrieval with KGs. Her scientific contributions span from entity linking to the prediction of influences in citation graphs. In this tutorial, she will cover her seminal publication on entity query feature expansion and her work on finding relevant relations.

Prof. Dr. Alexander Kotov is an Assistant Professor in the Department of Computer Science at Wayne State University. His general research interests lie at the intersection of information retrieval, textual data mining and health informatics. Before joining Wayne State, he was a post-doctoral fellow at Emory University working with Prof. Eugene Agichtein. Dr. Kotov obtained his PhD from the University of Illinois at Urbana-Champaign, under the supervision of Professor ChengXiang Zhai. At Wayne State, he has been teaching graduate courses on Information Retrieval and NoSQL databases as well as undergraduate courses. In this tutorial, he will cover his recent work on entity retrieval from knowledge graphs along with the methods for entity representation and ranking.

Dr. Edgar Meij is a senior scientist at Bloomberg where he also leads a team focusing on graph analytics and semantic technologies. Before this, he was a research scientist at Yahoo Labs and a postdoc at the University of Amsterdam, where he also obtained

[2] See https://kg4ir.github.io/tutorial.

his Ph.D. He regularly teaches at the (post-)graduate level, including university courses and conference tutorials, e.g., at EACL, ICTIR, SIGIR, WWW, and WSDM. His research focuses on all applications and aspects of knowledge graphs, entity linking, and semantic search. This tutorial will cover his contributions on entity linking, entity aspect mining, and finding supporting passages for entity relations.

Acknowledgements

L. Dietz is supported by a scholarship of the Eliteprogramm for Postdocs of the Baden-Württemberg Stiftung.

References

[1] Saeid Balaneshin-kordan and Alexander Kotov. 2016. Sequential query expansion using concept graph. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management.* 155–164.

[2] Saeid Balaneshinkordan and Alexander Kotov. 2016. An Empirical Comparison of Term Association and Knowledge Graphs for Query Expansion. In *ECIR.*

[3] Niranjan Balasubramanian and Silviu Cucerzan. 2010. Beyond Ranked Lists in Web Search: Aggregating Web Content into Topic Pages. *International Journal of Semantic Computing* 4, 4 (2010), 509–534.

[4] Krisztian Balog, Pavel Serdyukov, and Arjen P de Vries. 2010. Overview of the TREC 2010 entity track. In *TREC.*

[5] Hannah Bast, Florian Bäurle, Björn Buchhold, and Elmar Haußmann. 2014. Semantic full-text search with broccoli. In *SIGIR.* 1265–1266.

[6] Michael Bendersky, Donald Metzler, and W Bruce Croft. 2012. Effective query formulation with multiple information sources. In *Proc. of WSDM-12.* 443–452.

[7] Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. Semantic parsing on freebase from question-answer pairs. In *EMNLP.* 1533–1544.

[8] Roi Blanco, Giuseppe Ottaviano, and Edgar Meij. 2015. Fast and Space-Efficient Entity Linking for Queries. In *WSDM.*

[9] Horatiu Bota, Ke Zhou, Joemon M Jose, and Mounia Lalmas. 2014. Composite retrieval of heterogeneous web search. In *Proceedings of the 23rd international conference on World wide web.* ACM, 119–130.

[10] Kevin K Bowden, Shereen Oraby, Jiaqi Wu, Amita Misra, and Marilyn Walker. 2017. Combining Search with Structured Data to Create a More Engaging User Experience in Open Domain Dialogue. In *SCAI Workshop at ICTIR.*

[11] Jeffrey Dalton and Laura Dietz. 2012. *Bi-directional linkability from Wikipedia to documents and back again: UMass at TREC 2012 knowledge base acceleration track.* Technical Report. DTIC Document.

[12] Jeffrey Dalton, Laura Dietz, and James Allan. 2014. Entity Query Feature Expansion Using Knowledge Base Links. In *SIGIR.*

[13] Gianluca Demartini, Tereza Iofciu, and Arjen P. de Vries. 2009. Overview of the INEX 2009 Entity Ranking Track. In *INEX.*

[14] Bhuwan Dhingra, Lihong Li, Xiujun Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmed, and Li Deng. 2017. Towards End-to-End Reinforcement Learning of Dialogue Agents for Information Access. In *ACL.* 484–495.

[15] Laura Dietz and Jeffrey Dalton. 2013. UMass at TREC 2013 Knowledge Base Acceleration Track: Bi-directional Entity Linking and Time-aware Evaluation.. In *TREC.*

[16] Laura Dietz, Manisha Verma, Filip Radlinski, and Nick Craswell. 2017. TREC Complex Answer Retrieval Overview. In *TREC.*

[17] Laura Dietz, Chenyan Xiong, and Edgar Meij (Eds.). 2017. *Proceedings of the First Workshop on Knowledge Graphs and Semantics for Text Retrieval and Analysis (KG4IR 2017) co-located with the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2017), Shinjuku, Tokyo, Japan, August 11, 2017.* CEUR Workshop Proceedings, Vol. 1883. CEUR-WS.org. http://ceur-ws.org/Vol-1883

[18] Huizhong Duan and ChengXiang Zhai. 2015. Mining Coordinated Intent Representation for Entity Search and Recommendation. In *CIKM.* ACM, 333–342.

[19] Anthony Fader, Stephen Soderland, and Oren Etzioni. 2011. Identifying relations for open information extraction. In *EMNLP.*

[20] Paolo Ferragina and Ugo Scaiella. 2010. Tagme: on-the-fly annotation of short text fragments (by wikipedia entities). In *CIKM.* 1625–1628.

[21] Evgeniy Gabrilovich, Michael Ringgaard, and Amarnag Subramanya. 2013. FACC1: Freebase annotation of ClueWeb corpora, Version 1.

[22] Darío Garigliotti and Krisztian Balog. 2017. On Type-Aware Entity Retrieval. In *SIGIR.* 27–34.

[23] David Graus, Manos Tsagkias, Wouter Weerkamp, Edgar Meij, and Maarten de Rijke. 2016. Dynamic collective entity representations for entity ranking. In *WSDM.*

[24] Faegheh Hasibi, Krisztian Balog, and Svein Erik Bratsberg. [n. d.]. Exploiting entity linking in queries for entity retrieval.

[25] Faegheh Hasibi, Fedor Nikolaev, Chenyan Xiong, Krisztian Balog, Svein Erik Bratsberg, Alexander Kotov, and Jamie Callan. 2017. DBpedia-Entity v2: A Test Collection for Entity Search. In *SIGIR.* 1265–1268.

[26] Jiyin He and Marc Bron. 2017. Measuring Demonstrated Potential Domain Knowledge with Knowledge Graphs. In *KG4IR Workshop at SIGIR.*

[27] Ben Hixon, Peter Clark, and Hannaneh Hajishirzi. 2015. Learning knowledge graphs for question answering through conversational dialog. In *NAACL-HLT.* 851–861.

[28] Johannes Hoffart, Dragan Milchevski, and Gerhard Weikum. 2014. STICS: Searching with Strings, Things, and Cats. In *SIGIR.*

[29] Heng Ji, Joel Nothman, Ben Hachey, et al. 2014. Overview of tac-kbp2014 entity discovery and linking tasks. In *TAC.* 1333–1339.

[30] Amina Kadry and Laura Dietz. 2017. Open Relation Extraction for Support Passage Retrieval: Merit and Open Issues. In *SIGIR.* 1149–1152.

[31] Rianne Kaptein, Pavel Serdyukov, Arjen P. de Vries, and Jaap Kamps. 2010. Entity ranking using Wikipedia as a pivot. In *CIKM.* 69–78.

[32] Alexander Kotov. 2016. Knowledge Graph Entity Representation and Retrieval. *tutorial chapter, RuSSIR* (2016).

[33] Alexander Kotov and ChengXiang Zhai. 2012. Tapping into knowledge base for concept feedback: leveraging conceptnet to improve search results for difficult queries. In *WSDM.*

[34] Adit Krishnan, P Deepak, Sayan Ranu, and Sameep Mehta. 2017. Leveraging semantic resources in diversified query expansion. *World Wide Web* (2017), 1–27.

[35] Ni Lao and William W Cohen. 2010. Relational retrieval using a combination of path-constrained random walks. *Machine Learning* 81, 1 (2010), 53–67.

[36] Hongyu Li, Chenyan Xiong, and Jamie Callan. 2017. Natural Language Supported Relation Matching for Question Answering with Knowledge Graphs. In *KG4IR Workshop at SIGIR.*

[37] Peng Li, Jing Jiang, and Yinglin Wang. 2010. Generating templates of entity summaries with an entity-aspect model and pattern mining. In *ACL.* 640–649.

[38] Xitong Liu and Hui Fang. 2015. Latent entity space: a novel retrieval approach for entity-bearing queries. *Information Retrieval Journal* 18, 6 (2015), 473–503.

[39] Edgar Meij, Wouter Weerkamp, and Maarten de Rijke. 2012. Adding Semantics to Microblog Posts. In *WSDM.*

[40] Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *ACL.*

[41] Federico Nanni, Simone Paolo Ponzetto, and Laura Dietz. 2018. Entity-aspect linking: providing fine-grained semantics of entities in context. In *JCDL.*

[42] Robert Neumayer, Krisztian Balog, and Kjetil Nørvåg. 2012. On the modeling of entities for ad-hoc entity search in the web of data. In *Proc. ECIR.* 133–145.

[43] Lan Nie, Brian D. Davison, and Xiaoguang Qi. 2006. Topical Link Analysis for Web Search. In *SIGIR.* New York, NY, USA, 91–98.

[44] Zaiqing Nie, Ji-Rong Wen, and Wei-Ying Ma. 2007. Object-level Vertical Search.. In *CIDR.* 235–246.

[45] Fedor Nikolaev, Alexander Kotov, and Nikita Zhiltsov. 2016. Parameterized Fielded Term Dependence Models for Ad-hoc Entity Retrieval from Knowledge Graph. In *SIGIR.*

[46] Sandeep Panem, Romil Bansal, Manish Gupta, and Vasudeva Varma. 2014. Entity Tracking in Real-Time using Sub-Topic Detection on Twitter. In *Advances in Information Retrieval.* Springer, 528–533.

[47] Jeffrey Pound, Peter Mika, and Hugo Zaragoza. 2010. Ad-hoc object retrieval in the web of data. In *WWW.*

[48] Delip Rao, Paul McNamee, and Mark Dredze. 2013. Entity linking: Finding extracted entities in a knowledge base. In *Multi-source, Multilingual Information Extraction and Summarization.*

[49] Hadas Raviv, Oren Kurland, and David Carmel. 2016. Document Retrieval Using Entity-Based Language Models. In *SIGIR.*

[50] Ridho Reinanda, Edgar Meij, and Maarten de Rijke. 2015. Mining, ranking and recommending entity aspects. In *SIGIR.*

[51] Michael Schuhmacher, Laura Dietz, and Simone Paolo Ponzetto. 2015. Ranking Entities for Web Queries through Text and Knowledge. In *CIKM.*

[52] Michael Schuhmacher, Benjamin Roth, Simone Paolo Ponzetto, and Laura Dietz. 2016. Finding Relevant Relations in Relevant Documents. In *ECIR.*

[53] Svitlana Vakulenko, Ilya Markov, and Maarten de Rijke. 2017. Conversational Exploratory Search via Interactive Storytelling. In *SCAI.*

[54] Nikos Voskarides, Edgar Meij, Manos Tsagkias, Maarten de Rijke, and Wouter Weerkamp. 2015. Learning to explain entity relationships in knowledge graphs. In *ACL-IJCNLP.*

[55] Chenyan Xiong and Jamie Callan. 2015. Esdrank: Connecting query and documents through external semi-structured data. In *CIKM.*

[56] Chenyan Xiong, Jamie Callan, and Tie-Yan Liu. 2017. Word-Entity Duet Representations for Document Ranking. In *SIGIR.* 763–772.

[57] Chenyan Xiong and James P. Callan. 2015. Query Expansion with Freebase. In *ICTIR.*

[58] Nikita Zhiltsov, Alexander Kotov, and Fedor Nikolaev. 2015. Fielded Sequential Dependence Model for Ad-Hoc Entity Retrieval in the Web of Data. In *SIGIR.*