An Entity-Oriented Approach for Answering Topical Information Needs

Shubham Chatterjee(✉)

University of New Hampshire, Durham, USA
shubham.chatterjee@unh.edu

Abstract. In this dissertation, we adopt an entity-oriented approach to identify relevant materials for answering a topical keyword query such as "Cholera". To this end, we study the interplay between text and entities by addressing three related prediction problems: (1) Identify knowledge base entities that are relevant for the query, (2) Understand an entity's meaning in the context of the query, and (3) Identify text passages that elaborate the connection between the query and an entity. Through this dissertation, we aim to study some overarching questions in entity-oriented research such as the importance of query-specific entity descriptions, and the importance of entity salience and context-dependent entity similarity for modeling the query-specific context of an entity.

1 Introduction

Wikipedia is useful for users seeking information on topics such as "Cholera"; however, it is mostly focused on recent and popular topics. Through this dissertation, we address the first step in answering topical queries: the retrieval of relevant materials (text and entities) that constitutes the answer. We envision a downstream system to utilize these relevant materials to automatically construct a Wikipedia-like article for such topical queries.

Humans usually think about topics in terms of entities, and the background stories and roles of these entities with respect to the topic. Motivated by this, we adopt an entity-oriented approach to identify relevant material for answering such topical queries by addressing three related prediction problems: (1) **Entity Retrieval:** Identify entities that are relevant for a discussion about the query, (2) **Entity Aspect Linking:** Link the mention of an entity in a given context (e.g., sentence) to the aspect (from a catalog) that best captures the meaning of the entity in that context, and (3) **Entity Support Passage Retrieval:** Find passages that explain why the entity is relevant for the query.

Research Questions. An entity such as "Oyster" may be referred to in multiple contexts in text, e.g., "Cultivation", "Ecosystem Services", etc. – each is called an *aspect* of the entity "Oyster". Often, features for entity ranking are derived from entity links found in query-relevant documents [10,16,20]. While entity linking can disambiguate the different mentions of "Oyster" in a text (animal

M. Hagen et al. (Eds.): ECIR 2022, LNCS 13186, pp. 463–472, 2022.
https://doi.org/10.1007/978-3-030-99739-7_57

Query: Cholera
Entity: Oyster

Most cholera cases in developed countries are a result of transmission by food. Food transmission can occur when people harvest seafood such as **oysters** in waters infected with sewage, as Vibrio cholerae accumulates in planktonic crustaceans and the **oysters** eat the zooplankton.

Oyster is the common name for a number of different families of salt-water bivalve molluscs that live in marine or brackish habitats. **Oysters** influence nutrient cycling, water filtration, habitat structure, biodiversity, and food web dynamics.

Cholera is an infection of the small intestine by some strains of the bacterium Vibrio cholerae. The classic symptom is large amounts of watery diarrhea that lasts a few days. Vomiting and muscle cramps may also occur. Cholera can be caused by eating **oysters**.

Fig. 1. Example query and entity with support passages. **Left:** The passage is relevant to the query and entity, and the entity is salient in the passage. The passage clarifies how oysters may cause cholera. Hence, this is a good support passage for the query-entity pair. **Middle:** The passage is relevant to the entity but not to the query. **Right:** The passage is relevant to the query but not to the entity as the entity is not salient in the passage. The passages in the middle and right are not good support passages.

versus place), it cannot identify the different aspects of "Oyster". Entity aspect linking [24] can remedy this by using a unique aspect id to resolve the different meanings (aspects) of an entity in the context of the query. Hence, we study the following research question: **(RQ1)** How can we leverage fine-grained entity aspects for entity retrieval, and to what extent are they useful?

Although we study the utility of entity aspects in the context of Web search, applications such as question-answering and recommender systems that aim to understand the subtleties in the human language would also benefit from research on entity aspect linking. Moreover, since entity aspect linking aims to match an entity's context to a candidate aspect, approaches to entity aspect linking would also be applicable to other text similarity problems. The current entity aspect linking system from Nanni et al. [24] has scope for improvement (see Sect. 3). Hence, we study the following research question: **(RQ2)** How can we improve the current entity aspect linking system?

We use entity aspect linking to obtain a fine-grained understanding of the entity in the context of the query. Alternatively, *entity support passages* may also be used for this purpose. Entity support passages [5,18] are paragraph-size text passages that explain why an entity, e.g., "Oyster", from the entity ranking for a query, e.g., "Cholera", is relevant to the query (Fig. 1). Entity support passages may serve as text to be summarized when generating an answer to a topical query. There are two challenges in finding a good support passage: (1) Support passages must be relevant to both, the query and the entity, and (2) The entity must be *salient*, i.e., *central* to the discussion in the text and not just mentioned as an aside. Hence, we study the following research questions: **(RQ3)** How can we model the joint relevance of a paragraph to an entity and a query? To what extent is entity salience helpful? **(RQ4)** How can we leverage support passages for entity retrieval and to what extent are they useful?

An important consideration in entity-oriented research is regarding the construction of entity descriptions. Often, entity descriptions are constructed without considering the query, by using knowledge bases [8,15,19], the entity's

Wikipedia article [10,21,22], or by collecting documents from a corpus that mention the entity [1,9,28]. As a result, such query-independent descriptions may contain information about the entity that is non-relevant in the context of the query. For example, although the entity "Oyster" is relevant to the query "Cholera", the Wikipedia page of "Oyster" does not even mention Cholera. We study the utility of entity aspects and entity support passages as *query-specific* entity descriptions for learning query-specific entity embeddings. In this regard, we study RQ1 and RQ4 in addition to the following research question: **(RQ5)** Is it sufficient to use the lead text of an entity's Wikipedia page as the entity's description?

2 Related Work

Entity Retrieval. The Sequential Dependence Model [23] assigns different weights to matching unigrams and bigarams of different types. Zhiltsov et al. [36] propose the Fielded Sequential Dependence Model (FSDM) that uses field-specific background models across multiple fields. Nikolaev et al. [25] estimate the probability of unigrams and bigrams being mapped onto a field in the FSDM dynamically. Hasibi et al. [16] leverage entity links in queries and propose a parameter-free estimation of the field weights in FSDM. Several models utilize information from Knowledge Bases. For example, Kaptein et al. [19] utilize the types of entities in the query, and Balog et al. [2] utilize category information about an entity obtained from a user. Learning-To-Rank (LTR) is another common approach. Schuhmacher et al. [30] utilize several features to re-rank entities in a LTR setting. Dietz [11] proposed ENT Rank, a LTR model that combines information about an entity, the entity's neighbors, and context using a hypergraph.

Entity Aspect Linking. Several works treat Wikipedia sections as entity aspects. For example, Fetahu et al. [14] enrich Wikipedia sections with news-article references. Banerjee et al. [3] seek to improve Wikipedia stubs by generating content for each section automatically. Nanni et al. [24] address the entity aspect linking task using the top-level sections from an entity's Wikipedia article as the entity's aspects. Their approach is based on LTR with lexical and semantic features derived from various contexts (e.g., sentence, paragraph, section) where the entity is mentioned in text. Ramsdell et al. [29] released a large dataset for entity aspect linking using the definition of aspects from Nanni et al.

Entity Support Passage Retrieval. Blanco et al. [5] rank entity support sentences using LTR with features based on named entity recognition, and term-based retrieval. Kadry et al. [18] study the importance of relation extraction for entity support passage retrieval. A related task is entity relationship explanation that aims to explain the relationship between two entities in a Knowledge Graph using a text passage. Pirro et al. [26] address the problem from a graph perspective by finding the sub-graph consisting of nodes and edges in the set of paths between the two input entities. Voskarides et al. [31,32] use textual, entity and relationship features within a LTR framework, whereas Bhatia et al. [4] address the problem from a probabilistic perspective.

3 Methodology

Datasets. The TREC Complex Answer Retrieval (CAR) [12] provides large and suitable benchmarks consisting of topical keyword queries. We use the CAR benchmarks for experiments on entity retrieval and entity support passage retrieval. Additionally, we also use the DBpedia-Entity v2 [17] dataset for our entity retrieval experiments. For experiments on entity aspect linking, we use the dataset from Ramsdell et al. [29]. We also use the aspect catalog and aspect linker implementation from Ramsdell et al. to entity aspect link the CAR corpus of English Wikipedia paragraphs.

Completed: Entity Support Passage Retrieval. We refer to the entity for which we want to find a support passage as *target entity*. Our approach [6] to entity support passage retrieval uses Learning-To-Rank with a rich set of features. The features are based on: (1) Modelling the joint relevance of a passage to the query and target entity, and (2) Salience of the target entity in a passage. Below, we describe our approach to study our research question **RQ3**. Our approach assumes that a high-precision entity ranking is available as input.

To model the relevance of a passage to a query, we retrieve a candidate set of passages $\mathcal{D}$ using BM25. We assume that a passage is relevant to the target entity if the passage contains many entities that are related to the target entity in the context of the query. We derive the query-relevant context $\mathcal{D}_e$ of the target entity from the candidate set $\mathcal{D}$ by retaining passages in $\mathcal{D}$ that contain a link to the target entity. We consider each passage $p \in \mathcal{D}_e$ as a candidate support passage. To identify entities from $\mathcal{D}_e$ which are related to the target entity, we treat $\mathcal{D}_e$ as a bag-of-entities [16,34]. We assume that entities $e_x \in \mathcal{D}_e$ which frequently co-occur with the target entity within $\mathcal{D}_e$ are related to the target entity. We then score a candidate support passage $p \in \mathcal{D}_e$ by the number of frequently co-occurring entities e_x linked to the passage. We also derive several features based on the salience of the target entity in the support passage.[1]

We outperform the state-of-the-art method from Blanco et al. [5] by a large margin on various benchmarks from TREC CAR. We find that salience is a strong indicator of support passages for entities that have a passage with a salient mention in the candidate set.

Completed: Entity Retrieval. We study **RQ1** based on the hypothesis that different mentions of an entity in a query-specific context contribute differently to determine the relevance of that entity for the query. For example, when determining the relevance of the entity “Oyster” for the query “Cholera”, the aspect “Diseases” is more important than “Cultivation”. Our approach [7] is based on Learning-To-Rank (LTR) with features derived from entity aspects: (1) Aspect Retrieval features: We rank aspects via their text[2] using the query, and (2)

[1] We use the salience detection system from Ponza et al. [27].
[2] Available from the aspect catalog from Ramsdell et al.

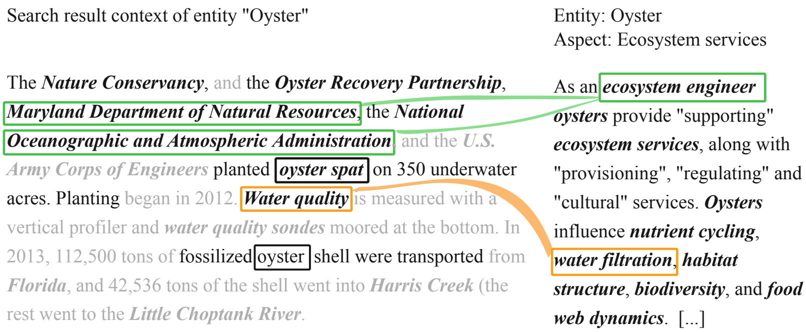

Fig. 2. Depiction of our entity aspect linking approach for the entity "Oyster". Left: context from search results. Right: Correct aspect "Ecosystem services" of the entity "Oyster". The example text, entities, and aspects taken from Wikipedia. Entity links marked in bold italics. In objective 1, we address the issue that not all words are relevant for the decision – non-relevant words are depicted in grey. As described in objective 2, it is rare that identical entities are mentioned in both context and aspect content, hence we need to identify which entities are related in this context, such as entities related to ecosystems (green frame) and regarding water quality (orange frame). In objective 3, we study how integrating the prediction of relevant words and entities is helpful for most accurate predictions of entity aspect links. (Color figure online)

Aspect Link PRF features: After retrieving an initial candidate set of passages using BM25, the frequency distribution of entity aspect links in these passages are weighted by the retrieval score of the passages to obtain a distribution of relevant entity aspects. To study **RQ4**, we also build the candidate set of passages from an entity support passage ranking instead of BM25 when deriving Aspect Link PRF features. Furthermore, we study RQ1, RQ4, and **RQ5** by fine-tuning a BERT model for the entity ranking task using entity aspects and entity support passages as query-specific entity descriptions.[3]

We find that our LTR and BERT models trained using entity aspects and entity support passages significantly outperform both neural and non-neural baselines using both TREC CAR and DBpedia-Entity v2. Moreover, significant performance improvements are obtained by replacing a query-independent entity description (e.g., lead text of an entity's Wikipedia article) with a query-specific description (e.g., entity support passage).

Proposed: Entity Aspect Linking. Below, we identify three research objectives to study **RQ2**.

1. **Identify relevant words in context and aspect.** Nanni et al. [24] find that using all words from the entity's context leads to poor results. They alleviate this by considering only the sentence mentioning the entity. However, as shown in Fig. 2, to help us make the aspect linking decision, we need to consider the whole passage which mentions the entity "Oyster", and not

[3] Paper under review.

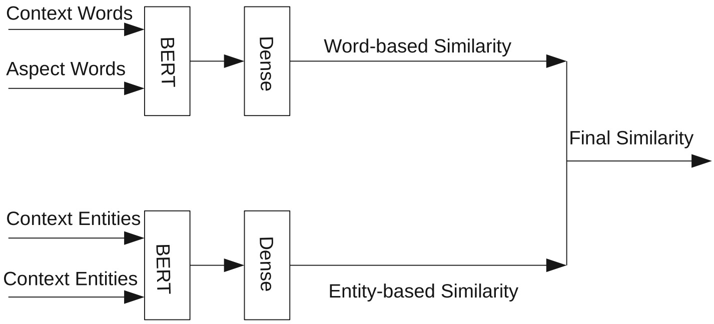

Fig. 3. Depiction of our proposed end-to-end entity aspect linking system. The top model takes the words from the aspect and context and learns a representation that pays attention to the words from the context that are important for the aspect linking decision. Similarly, the bottom model learns a context-dependent similarity between the entities from the context and aspect. The final similarity is obtained by combining the word-based and entity-based similarity.

just the sentence. Since the majority of words in the larger context are not relevant, we propose to use the attention mechanism in deep learning to select words from the context which are most beneficial for the aspect linking decision.

2. **Identify contextually related entities.** Nanni et al. base the aspect linking decisions on whether a direct relationship exists between an aspect-entity and a context-entity. However, as shown in Fig. 2, otherwise unrelated entities are related in the given context. Hence, we propose to base the aspect linking decisions on whether two entities are related in context, by learning embeddings of these entities using BERT, taking the context into account. Our preliminary work (See Footnote 3) in learning query-specific entity embeddings has shown promising results.
3. **Integrate information from words and entities.** Previous works which leverage entities for retrieving text [13,20,33–35] have found that combining indicators of relevance obtained using words and entities leads to better performance for distinguishing relevant from non-relevant text. In this light, I propose to integrate the information from relevant words and entities (from 1 and 2 above) by learning the similarity between the context and the aspect end-to-end using a Siamese Neural Network (Fig. 3).

4 Conclusion

In this research statement, we describe our entity-oriented approach to identify relevant text and entities for answering a topical keyword query such as "Cholera". We describe three related tasks that we address for this purpose, and identify several overarching research questions in entity-oriented research that we aim to answer with this work. We envision this work to serve as a stepping stone towards building more intelligent systems. Such systems would one day respond to a user's open-ended and complex information needs with a complete answer instead of a ranked list of results, thus transforming the "search" engine into an "answering" engine.

References

1. Balog, K., Azzopardi, L., de Rijke, M.: Formal models for expert finding in enterprise corpora. In: Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2006, pp. 43–50. Association for Computing Machinery, New York (2006). https://doi.org/10.1145/1148170.1148181
2. Balog, K., Bron, M., De Rijke, M.: Query modeling for entity search based on terms, categories, and examples. ACM Trans. Inf. Syst. **29**(4) (2011). https://doi.org/10.1145/2037661.2037667
3. Banerjee, S., Mitra, P.: WikiKreator: improving Wikipedia stubs automatically. 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, pp. 867–877. Association for Computational Linguistics, July 2015. https://doi.org/10.3115/v1/P15-1084. https://aclanthology.org/P15-1084
4. Bhatia, S., Dwivedi, P., Kaur, A.: That's interesting, tell me more! Finding descriptive support passages for knowledge graph relationships. In: Vrandečić, D., et al. (eds.) ISWC 2018. LNCS, vol. 11136, pp. 250–267. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00671-6_15
5. Blanco, R., Zaragoza, H.: Finding support sentences for entities. In: Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2010, pp. 339–346. Association for Computing Machinery, New York (2010). https://doi.org/10.1145/1835449.1835507
6. Chatterjee, S., Dietz, L.: Why does this entity matter? Support passage retrieval for entity retrieval. In: Proceedings of the 2019 ACM SIGIR International Conference on Theory of Information Retrieval, ICTIR 2019, pp. 221–224. Association for Computing Machinery, New York (2019). https://doi.org/10.1145/3341981.3344243
7. Chatterjee, S., Dietz, L.: Entity retrieval using fine-grained entity aspects. In: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2021. Association for Computing Machinery, New York (2021). https://doi.org/10.1145/3404835.3463035
8. Chen, J., Xiong, C., Callan, J.: An empirical study of learning to rank for entity search. In: Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2016, pp. 737–740. Association for Computing Machinery, New York (2016). https://doi.org/10.1145/2911451.2914725
9. Conrad, J.G., Utt, M.H.: A system for discovering relationships by feature extraction from text databases. In: Croft, B.W., van Rijsbergen, C.J. (eds.) SIGIR 1994, pp. 260–270. Springer, London (1994). https://doi.org/10.1007/978-1-4471-2099-5_27

10. Dalton, J., Dietz, L., Allan, J.: Entity query feature expansion using knowledge base links. In: Proceedings of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2014, pp. 365–374. Association for Computing Machinery, New York (2014). https://doi.org/10.1145/2600428.2609628
11. Dietz, L.: ENT rank: retrieving entities for topical information needs through entity-neighbor-text relations. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2019, pp. 215–224. Association for Computing Machinery, New York (2019). https://doi.org/10.1145/3331184.3331257
12. Dietz, L., Foley, J.: TREC CAR Y3: complex answer retrieval overview. In: Proceedings of Text REtrieval Conference (TREC) (2019)
13. Ensan, F., Bagheri, E.: Document retrieval model through semantic linking. In: Proceedings of the 10th ACM International Conference on Web Search and Data Mining, WSDM 2017, pp. 181–190. Association for Computing Machinery, New York (2017). https://doi.org/10.1145/3018661.3018692
14. Fetahu, B., Markert, K., Anand, A.: Automated news suggestions for populating wikipedia entity pages. In: Proceedings of the 24th ACM International on Conference on Information and Knowledge Management, CIKM 2015, pp. 323–332. Association for Computing Machinery, New York (2015). https://doi.org/10.1145/2806416.2806531
15. Graus, D., Tsagkias, M., Weerkamp, W., Meij, E., de Rijke, M.: Dynamic collective entity representations for entity ranking. In: Proceedings of the Ninth ACM International Conference on Web Search and Data Mining, WSDM 2016, pp. 595–604. Association for Computing Machinery, New York (2016). https://doi.org/10.1145/2835776.2835819
16. Hasibi, F., Balog, K., Bratsberg, S.E.: Exploiting entity linking in queries for entity retrieval. In: Proceedings of the 2016 ACM International Conference on the Theory of Information Retrieval, ICTIR 2016, pp. 209–218. Association for Computing Machinery, New York (2016). https://doi.org/10.1145/2970398.2970406
17. Hasibi, F., et al.: DBpedia-entity V2: a test collection for entity search. In: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2017, pp. 1265–1268. Association for Computing Machinery, New York (2017). https://doi.org/10.1145/3077136.3080751
18. Kadry, A., Dietz, L.: Open relation extraction for support passage retrieval: merit and open issues. In: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2017, pp. 1149–1152. Association for Computing Machinery, New York (2017). https://doi.org/10.1145/3077136.3080744
19. Kaptein, R., Serdyukov, P., De Vries, A., Kamps, J.: Entity ranking using wikipedia as a pivot. In: Proceedings of the 19th ACM International Conference on Information and Knowledge Management, CIKM 2010, pp. 69–78. Association for Computing Machinery, New York (2010). https://doi.org/10.1145/1871437.1871451
20. Liu, X., Fang, H.: Latent entity space: a novel retrieval approach for entity-bearing queries. Inf. Retr. J. **18**(6), 473–503 (2015)
21. Manotumruksa, J., Dalton, J., Meij, E., Yilmaz, E.: CrossBERT: a triplet neural architecture for ranking entity properties. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2020, pp. 2049–2052. Association for Computing Machinery, New York (2020). https://doi.org/10.1145/3397271.3401265

22. Meij, E., Bron, M., Hollink, L., Huurnink, B., de Rijke, M.: Mapping queries to the linking open data cloud: a case study using DBpedia. J. Web Semant. **9**(4), 418–433 (2011). https://doi.org/10.1016/j.websem.2011.04.001. http://www.sciencedirect.com/science/article/pii/S1570826811000187. jWS special issue on Semantic Search
23. Metzler, D., Croft, W.B.: A Markov random field model for term dependencies. In: Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2005, pp. 472–479. Association for Computing Machinery, New York (2005). https://doi.org/10.1145/1076034.1076115
24. Nanni, F., Ponzetto, S.P., Dietz, L.: Entity-aspect linking: providing fine-grained semantics of entities in context. In: Proceedings of the 18th ACM/IEEE on Joint Conference on Digital Libraries, JCDL 2018, pp. 49–58. Association for Computing Machinery, New York (2018). https://doi.org/10.1145/3197026.3197047
25. Nikolaev, F., Kotov, A., Zhiltsov, N.: Parameterized fielded term dependence models for ad-hoc entity retrieval from knowledge graph. In: Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2016, pp. 435–444. Association for Computing Machinery, New York (2016). https://doi.org/10.1145/2911451.2911545
26. Pirrò, G.: Explaining and suggesting relatedness in knowledge graphs. In: Arenas, M., et al. (eds.) ISWC 2015. LNCS, vol. 9366, pp. 622–639. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-25007-6_36
27. Ponza, M., Ferragina, P., Piccinno, F.: SWAT: a system for detecting salient Wikipedia entities in texts. Comput. Intell. (2018). https://doi.org/10.1111/coin.12216
28. Raghavan, H., Allan, J., McCallum, A.: An exploration of entity models, collective classification and relation description. In: KDD Workshop on Link Analysis and Group Detection, pp. 1–10 (2004)
29. Ramsdell, J., Dietz, L.: A large test collection for entity aspect linking. In: Proceedings of the 29th ACM International Conference on Information and Knowledge Management, CIKM 2020, pp. 3109–3116. Association for Computing Machinery, New York (2020). https://doi.org/10.1145/3340531.3412875
30. Schuhmacher, M., Dietz, L., Paolo Ponzetto, S.: Ranking entities for web queries through text and knowledge. In: Proceedings of the 24th ACM International on Conference on Information and Knowledge Management, CIKM 2015, pp. 1461–1470. Association for Computing Machinery, New York (2015). https://doi.org/10.1145/2806416.2806480
31. Voskarides, N., Meij, E., de Rijke, M.: Generating descriptions of entity relationships. In: Jose, J.M., et al. (eds.) ECIR 2017. LNCS, vol. 10193, pp. 317–330. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56608-5_25
32. Voskarides, N., Meij, E., Tsagkias, M., de Rijke, M., Weerkamp, W.: Learning to explain entity relationships in knowledge graphs. 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, pp. 564–574. Association for Computational Linguistics, July 2015. https://doi.org/10.3115/v1/P15-1055. https://www.aclweb.org/anthology/P15-1055
33. Xiong, C., Callan, J.: Query expansion with freebase. In: Proceedings of the 2015 International Conference on The Theory of Information Retrieval, ICTIR 2015, pp. 111–120. Association for Computing Machinery, New York (2015). https://doi.org/10.1145/2808194.2809446

34. Xiong, C., Callan, J., Liu, T.Y.: Bag-of-entities representation for ranking. In: Proceedings of the 2016 ACM International Conference on the Theory of Information Retrieval, ICTIR 2016, pp. 181–184. Association for Computing Machinery, New York (2016). https://doi.org/10.1145/2970398.2970423
35. Xiong, C., Callan, J., Liu, T.Y.: Word-entity duet representations for document ranking. In: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2017, pp. 763–772. Association for Computing Machinery, New York (2017). https://doi.org/10.1145/3077136.3080768
36. Zhiltsov, N., Kotov, A., Nikolaev, F.: Fielded sequential dependence model for ad-hoc entity retrieval in the web of data. In: Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2015, pp. 253–262. Association for Computing Machinery, New York (2015). https://doi.org/10.1145/2766462.2767756