Entity Retrieval Using Fine-Grained Entity Aspects

Shubham Chatterjee
sc1242@wildcats.unh.edu
University of New Hampshire
Durham, New Hampshire, USA

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

ABSTRACT

Using entity aspect links, we improve upon the current state-of-the-art in entity retrieval. Entity retrieval is the task of retrieving relevant entities for search queries, such as "Antibiotic Use In Livestock". Entity aspect linking is a new technique to refine the semantic information of entity links. For example, while passages relevant to the query above may mention the entity "USA", there are many aspects of the USA of which only few, such as "USA/Agriculture", are relevant for this query. By using entity aspect links that indicate which aspect of an entity is being referred to in the context of the query, we obtain more specific relevance indicators for entities. We show that our approach improves upon all baseline methods, including the current state-of-the-art using a standard entity retrieval test collection. With this work, we release a large collection of entity-aspect-links for a large TREC corpus.

ACM Reference Format:
Shubham Chatterjee and Laura Dietz. 2021. Entity Retrieval Using Fine-Grained Entity Aspects. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '21), July 11–15, 2021, Virtual Event, Canada.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3404835.3463035

1 INTRODUCTION

Entity-oriented search has become ubiquitous, with 40-70% of all web searches targeting entities [12, 18, 24]. In this work, we address the task of *topical entity retrieval* for automatic article generation: Given a short topical keyword query such as *Antibiotic Use In Livestock*, return a ranked list of entities based on whether the entity must, should, or could be mentioned in an article on this topic.

In this work, we leverage a refined version of entity links, called *entity aspect links*. Analyzing entity aspect links present in a set of candidate documents allows us to significantly improve upon the current state-of-the-art in topical entity retrieval. Entity Aspect Linking [22, 25] is a recent information extraction task: Given a mention of an entity in a sentence, entity aspect linking refines the entity link to an entity aspect link that provides information on the context in which the entity is mentioned by indicating which aspect of the linked entity is referenced in this context. For example (see Figure 1), the entity "Food and Drug Administration" (FDA) may be mentioned in the context of aspects "Science and Research Programs", "History", etc. – each is called an *aspect* of the entity "Food and Drug Administration". Nanni et al. [22] suggest to derive a catalog of aspects[1] from the top-level sections of the entity's Wikipedia article, but other sources of aspects can also be used. We note that entity aspects are different from entity *types*. Aspects refer to the topics in which an entity is referenced, for example, FDA in the context of its history versus FDA as a regulator; types resolve which of many roles the entity can take on, for example, US federal agency or food safety organization. While the utility of entity types for entity retrieval is well-studied, to the best of our knowledge, we are the first to study the usefulness of entity aspects for IR tasks, such as entity retrieval.

Approaches to address the entity retrieval task often derive features for entities using entity links from a candidate set of documents retrieved with the query [3, 4, 19]. However, as in the example above, an entity link is only a unique identifier of an entity and does not preserve any further topical information about the context in which the entity is mentioned. Using a unique aspect id, entity aspect links can remedy this by resolving the context (aspect) in which the entity has been mentioned in text. Hence, entity aspect linking refines an entity link with the topical semantics of the entity's referenced aspect. In this work, we explore the extent to which such fine-grained aspects of entities can provide additional, and perhaps better signals of the relevance of an entity for a query.

Contributions. We make the following contributions:

- We propose a novel entity retrieval approach which uses entity aspect links to leverage the topical context in which the entity is relevant.[2] We outperform the previous state-of-the-art by 41%.
- We develop novel features derived from entity aspects and entity aspect linking and show that these guide our approach to more relevant and fewer non-relevant entities.
- We demonstrate that using a candidate set derived from entity support passages (passages that are suitable to explain the relevance of an entity for a query) instead of BM25 leads to further improvements.

2 RELATED WORK

Term-based Models. Often, the first step in entity retrieval is to create a term-based representation of entities called *entity descriptions* using either unstructured sources such as document corpora, or structured knowledge bases such as DBpedia [17] and Freebase [2]. The second step is to rank these descriptions using traditional document retrieval models such as BM25 [27]. For example, Meij et al. [20] retrieve an initial candidate set of entities using the entity descriptions. Then, a supervised machine learning algorithm

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGIR '21, July 11–15, 2021, Virtual Event, Canada
© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8037-9/21/07…$15.00
https://doi.org/10.1145/3404835.3463035

[1] We use *entity aspect* and *aspect* interchangeably in this work.
[2] Data and code available https://github.com/shubham526/SIGIR2021-Entity-Retrieval

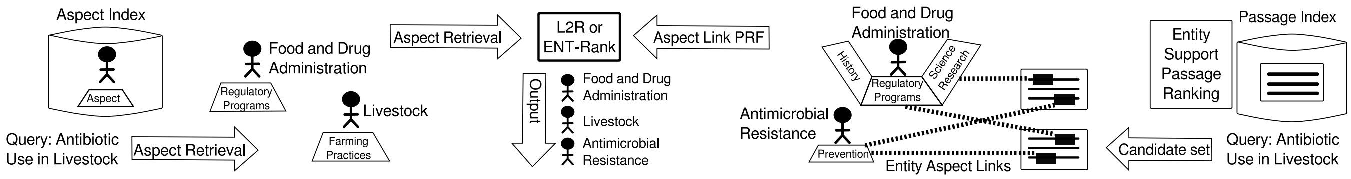

Figure 1: For the example query "Antibiotic Use in Livestock", we identify the relevant entity "Food and Drug Administration" as its relevant aspect "Regulatory Programs" is both retrieved from an aspect index and linked in the candidate passage set.

is used to classify each candidate entity as relevant or not for the query.

Probabilistic Graphical Models. Models based on Markov Random Fields (MRF) [21] have also been employed to represent a joint distribution over the terms from an entity's description, and the information from a semi-structured data about the entity. Such models include the Sequential Dependence Model [21] and its variants [23, 31]. Hasibi et al. [13] leverage the entity annotations in the queries using MRFs, wherein they introduce a new component for matching the linked entities from the query. Raviv et al. [26] model the different representations of an entity (description, type, and name) jointly with the query terms using MRFs.

Semantically Enriched Models. With the availability of large-scale knowledge bases, semantically enriched models that utilize the types of entities have also been proposed. For example, Kaptein et al. [14] utilize the types of the entities in the query to rank Wikipedia pages according to the similarity with the query and the entity types. While Balog et al. [1] utilize category information about an entity obtained from a user, Garigliotti et al. [9] utilize type information about entities available in type taxonomies such as DBpedia and Freebase. Similarly, relationships between entities in a knowledge base have also been leveraged for entity ranking. For example, Tonon et al. [29] first identify a set of seed entities using term-based retrieval, and then traverse the edges of these seed entities in a knowledge graph to identify additional entities. Recently, Gerritse et al. [10] have shown that graph embeddings obtained using Wikipedia2Vec [30] are useful for entity ranking [10]. They re-rank each entity in an initial entity ranking by the cosine similarity between the entity's embedding, and the embedding of entities in the query (identified with TagMe [8]). We add this as a reference baseline in this work.

Learning-to-Rank Models. Another class of entity retrieval methods are based on Learning-To-Rank (L2R). Schuhmacher et al. [28] utilize several features to re-rank entities in a L2R setting. Graus et al. [11] represent the entity as a fielded document and then combine the content from each field using L2R. More recently, Dietz [6] proposed ENT Rank, a L2R model which utilizes the information about text for entity retrieval. Neighbor relations between entities are defined using the context of the entity such as the passage in which the entity appears. We include ENT Rank as a reference baseline in this work.

3 ENTITY ASPECTS FOR ENTITY RETRIEVAL

Our work is 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 "Food and Drug Administration" for the query "Antibiotic Use In Livestock" (see Figure 1), the aspect "Regulatory Programs" is more important than "History". Our approach is to identify relevant entity-aspects, based on the assumption that only entities with relevant aspects are actually relevant for the query. Entity aspects allow us to be more specific about the topical context in which the entity is relevant. The main effect is that our predicted entity rankings contain fewer mistakes in the top of the ranking, hence obtaining significant performance improvements overall.

Our approach relies on passages with entity aspect links. In this work, we utilize the aspect catalog and aspect linker implementation from Ramsdell et al. [22, 25][3] to aspect link[4] the corpus from TREC Complex Answer Retrieval benchmark [7] which is derived from the English Wikipedia.

Aspect retrieval features. A common approach to rank entities is to use a document retrieval model such as BM25 to rank the entity's term-based representation. We transfer this idea to entity aspects by creating a search index of entity aspect descriptions, then rank aspects by retrieving from this index using several retrieval models (described in Section 4). This index contains the name of the Wikipedia page corresponding to the entity (such as "Food and Drug Administration"), name of the aspect (such as "Regulatory Process"), content of the aspect (the text of the section on "Regulatory Process" in the entity's Wikipedia article), and other entities mentioned in the aspect's content. This information is obtained from the aspect catalog. By using multiple retrieval models, we obtain multiple aspect retrieval features which are combined with learning-to-rank.

One may argue that since aspects are derived from sections in the entity's Wikipedia article, the same information is also available when retrieving entities from a full text index of Wikipedia articles. Creating a search entry for each aspect (section) avoids the common pitfall where entities with many aspects are penalized for their diversity by the retrieval model, via document length components or L2-normalization of term vectors. The downside is that Wikipedia articles do not always contain all relevant information about an entity due to reasons such as articles being out-of-date or some (negative) information being removed. Hence, we incorporate fall-back strategies to maximize the recall (detailed below).

Aspect link PRF features. An idea akin to Pseudo Relevance Feedback (PRF) [16] is often used in entity-oriented search: After retrieving an initial candidate set of text passages, the frequency distribution of entity links in these passages are weighted by the retrieval score of the passages to obtain a distribution of relevant

[3]https://www.cs.unh.edu/~dietz/eal-dataset-2020/
[4]We use *entity-aspect link* and *aspect link* interchangeably in this work.

entities. This is often a strong relevance indicator [4, 6, 26]. We translate this idea to entity aspect links. Using a candidate set of passages for the query (see below), we obtain a distribution over relevant aspects a for query q using candidate passages D:

$$\text{score}(a|q) = \sum_{d \in D} \text{score}(d|q) \cdot \frac{\text{number of aspect links to } a \text{ in } d}{\text{total number of aspect links in } d}$$

Using entity aspect links instead of entity links offers acess to more fine-grained topical information ("FDA/Regulatory Program" versus "FDA/History"). This helps to promote entities that are mentioned in the context of the same aspect across multiple candidate passages. Using several retrieval models (detailed in Section 4) and different candidate sets (detailed below), we obtain multiple aspect link PRF features which are combined with learning-to-rank.

Candidate set. Our aspect-based features described above use a candidate set of passages D for the query. A common approach is to create the candidate set using the top-K documents of a BM25 ranking. However, non-relevant entities can often dominate such candidate passages, which can negatively affect the identification of relevant aspects. We avoid this by leveraging work on entity support passage retrieval [3] where the task is, given a query and a target entity, retrieve passages which best explain why the entity is relevant to the query.

We use the method from Chatterjee et al. [3]: We build an entity description consisting of passages from a query-relevant candidate set that mention the entity. These description passages are then re-ranked using query words, expansion words, and expansion entities.[5] By using entity support passages instead of a direct BM25 ranking as candidates, we maximize the query-relevant information about each entity. Finally, we merge all entity support passage rankings across entities from a high-precision entity ranking.[6] We merge multiple support passage rankings by marginalizing over these entities: $\text{Score}(p|q) = \sum_{e_i} \text{Score}(p|e_i, q)$, where p is a support passage for the entity e_i given the query q. The top-K of this ranking is used to build the candidate set of passages D for the query when deriving Entity Aspect Link PRF features. Such a candidate set promotes passages that are good explanations for multiple entities and avoids that the candidate ranking is dominated by a single frequently occurring entity.

Conversion of relevant aspects to entities. To convert rankings of entity aspects to a ranking of entities, we consider the top-K aspects, then aggregate multiple aspects of the same entity either by sum or max. Empirically we choose $K = 100$. A separate entity ranking is obtained per aspect feature.

Entity features. We include various entity relevance features used in previous work. We use an entity index of full Wikipedia pages, and description field with meta data to obtain entity rankings using several retrieval models (detailed in Section 4). We also include features based on the Entity Context Model [4].

Combinations and Learning-to-rank. We find that a strong method is a combination of aspect retrieval and aspect link PRF. Hence, we filter the aspect ranking obtained using aspect retrieval features, and only retain aspects that are linked in passages from the candidate set (either using BM25 or the support passage ranking).

[5] In this work, this candidate set is retrieved with BM25, but the method can be adjusted to other methods as well.

[6] We retrieve entities via the lead text of their Wikipedia articles using BM25.

We also include an aggregate feature via reciprocal rank aggregation on our aspect link features into our learning-to-rank system. Reciprocal rank aggregation is an unsupervised rank aggregation method that has been found to be a strong relevance indicator [6]: All distinct items d across all rankings R are assigned a new aggregated rank score from reciprocal ranks $\frac{1}{rank(d)}$.

We use a Learning-to-Rank[7] approach to train an ideal weighed combination of all features. We also extend the features of the original ENT-Rank approach with our entity-aspect features to demonstrate the potential for further performance improvements to both the previous state-of-the-art results, as well as an improvement over learning-to-rank. In both cases, we optimize for Mean Average Precision using coordinate ascent with Z-score normalization.

4 EVALUATION

We evaluate our approach using the dataset from TREC Complex Answer Retrieval (CAR) [7].[8] The dataset contains an entity linked corpus consisting of passages from the entire English Wikipedia. We use two subsets from the CAR dataset for our experiments:

- **BenchmarkY1-Train.** We conduct **page-level experiments** using this benchmark and evaluate our results using the automatic entity ground truth provided with the dataset. This ground truth is created synthetically by defining all entities on the Wikipedia page corresponding to the query as relevant for the query. This benchmark is based on a Wikipedia dump from 2016. It contains 117 page-level queries with a total of 13,031 positive entity assessments.
- **BenchmarkY2-Test.** As the official CAR results are on section-level queries, we conduct **section-level experiments** using this benchmark and evaluate our results using the manual entity ground truth provided. This ground truth is created after manual assessments by NIST using pool-based evaluation. This benchmark is based partly on a Wikipedia dump from 2018 and partly on the Textbook Question Answering [15] dataset. It contains 271 section-level queries with a total of 1356 positive assessments.

Training and Evaluation. We train and test using 5-fold cross-validation. We use Mean Average Precision (MAP), Precision at R (P@R) and Mean Reciprocal Rank (MRR) as our evaluation metrics and conduct significance testing using paired-t-tests.

Feature generating retrieval models. We use the following retrieval models with our entity features and entity-aspect link features to generate features per retrieval model.: (1) BM25, and (2) Query Likelihood with Dirichlet Smoothing (QL-DS) ($\mu = 1500$) using RM3-style query expansion.

Baselines. We compare our approach to the following methods:

- **CatalogRetrieval.** We index the aspect catalog, and directly retrieve aspects from this index with the query using BM25 without any other components of our approach.
- **Wiki2Vec-ReRank.** Entity re-ranking method from Gerritse et al. [10] which uses Wikipedia2Vec [30].
- **BERT-ReRank.** Similar entity re-ranking method as Gerritse et al. [10] but using BERT [5]. We embed the query and candidate entity[9] using BERT's contextual sentence embedding.

[7] https://www.cs.unh.edu/~dietz/rank-lips/

[8] http://trec-car.cs.unh.edu

[9] We use the name of the Wikipedia page corresponding to the entity.

Table 1: Results on BenchmarkY1-Train page-level (using automatic ground truth) and BenchmarkY2-Test section-level (using manual ground truth). $^\blacktriangle$ denotes significant improvement and $^\blacktriangledown$ denotes significant deterioration compared to ENT-Rank [6] using a paired-t-test at $p < 0.05$.

	CAR Y1-Train Page			CAR Y2-Test Section		
	MAP	P@R	MRR	MAP	P@R	MRR
CatalogRetrieval	0.04$^\blacktriangledown$	0.09$^\blacktriangledown$	0.55$^\blacktriangledown$	0.14$^\blacktriangledown$	0.18$^\blacktriangledown$	0.55$^\blacktriangledown$
Wiki2Vec-ReRank[10]	0.09$^\blacktriangledown$	0.16$^\blacktriangledown$	0.53$^\blacktriangledown$	0.19$^\blacktriangledown$	0.20$^\blacktriangledown$	0.62$^\blacktriangledown$
BERT-ReRank	0.21$^\blacktriangledown$	0.28$^\blacktriangledown$	0.61$^\blacktriangledown$	0.15$^\blacktriangledown$	0.17$^\blacktriangledown$	0.61$^\blacktriangledown$
CAR Rank 1: UNH-e-L2R	–	–	–	0.31$^\blacktriangledown$	0.31$^\blacktriangledown$	0.75$^\blacktriangledown$
ENT-Rank* [6]	0.32	0.36	0.67	0.32	0.32	0.74
Our Approach (L2R)	0.50$^\blacktriangle$	0.50$^\blacktriangle$	0.77$^\blacktriangle$	0.45$^\blacktriangle$	0.45$^\blacktriangle$	0.83$^\blacktriangle$
Ours + ENT-Rank	0.53$^\blacktriangle$	0.55$^\blacktriangle$	0.93$^\blacktriangle$	0.47$^\blacktriangle$	0.48$^\blacktriangle$	0.84$^\blacktriangle$

Table 2: Results of ablation study on BenchmarkY1-Train for page-level experiments.

L2R:	Aspect		Entity			
Feedback:	BM25	Entity-Support-Psg		MAP	P@R	MRR
			X	0.37	0.39	0.71
	X	X		0.42	0.47	0.77
		X	X	0.48	0.49	0.77
	X		X	0.41	0.44	0.78
	X	X	X	0.50	0.50	0.77

- **ENT-Rank.** ENT-Rank method using entity, neighbors, and text features [6]. This is the current state-of-the-art entity retrieval method on our benchmark.
- **UNH-e-L2R.** This is the best performing official submission (using BenchmarkY2-Test) to the entity retrieval track in TREC CAR Year 2. Results taken from the TREC CAR submission, which are not available for BenchmarkY1-Train.

4.1 Results

Overall performance. We observe from Table 1 that on both benchmarks, our proposed approach outperforms all baselines. On BenchmarkY1-Train, we obtain $MAP = 0.50$ using L2R with our features. This is an improvement of 56% over the current state-of-the-art method ENT-Rank. On BenchmarkY2-Test, we achieve $MAP = 0.45$ using L2R. This is an improvement of 41% over ENT-Rank and UNH-e-L2R, the best performing entity retrieval system from TREC CAR year 2. We also outperform the neural re-ranking methods based on BERT and Wikipedia2Vec by a large margin. Using our features within the ENT-Rank framework gives further performance improvements on both benchmarks.

Ablation study. To gain insights into the performance improvements, we analyze our results on BenchmarkY1-Train. We perform an ablation study where we divide our features into three subsets: aspect-based features using BM25 candidate passages, aspect-based features using entity support passages, and other entity features. From Table 2, we observe that aspect features alone achieve $MAP = 0.42$ whereas entity features alone achieve $MAP = 0.37$. However, a combination of the three achieves the maximum improvement. The aspect-based features improve performance by 35% over the entity features.

We also find that a combination of aspect retrieval and aspect link PRF is a strong feature. Our aspect retrieval feature with BM25 achieves $MAP = 0.04$ whereas aspect link PRF with a BM25 candidate set achieves $MAP = 0.17$. However, the combination of the two achieves $MAP = 0.20$. Replacing the BM25 candidate set with a support passage candidate set for aspect link PRF in the combination achieves further improvement with $MAP = 0.40$.

Entity-level analysis. For each query in BenchmarkY1-Train, we inspect the top-100 entities in the ranking obtained by aspect-based features and entity features. We then find the number of relevant entities found by: (1) Both, (2) Only aspect-based features, and (3) Only entity features. We use the query "Antibiotic Use In Livestock" as an illustrative example here. We find that both aspect-based features and entity features retrieve seven relevant entities; however, aspect-based features always places these entities higher in the ranking. For example, the entity "Food and Drug Administration" is placed at rank 3 by aspect-based features but at rank 55 by entity features. Moreover, aspect-based features also retrieve more relevant entities than entity features such as "Animal Feed", "Antibiotic Misuse", and "Antifungal".

Entity support passages for deriving aspect-based features. From Table 2, we observe that there is a huge performance drop when we remove aspect-based features derived from the candidate set obtained using entity support passages ($MAP = 0.50$ to $MAP = 0.41$). However, the performance drops only slightly when we remove the aspect-based features obtained using BM25 candidate set ($MAP = 0.50$ to $MAP = 0.48$). On further investigation, we find that the top passages of the support passage candidate set contain more relevant entities than that obtained using BM25 candidate set, hence yielding better results.

5 CONCLUSION

We demonstrate the benefits of integrating entity aspects into an entity retrieval system. Our results show that aspect-based features are more informative than traditional entity ranking features, and outperform several strong baselines, including BERT-based re-ranking method. We obtain this performance boost because previously missing relevant entities are identified by aspect retrieval and aspect link PRF features when combined with L2R. Hence, relevant entities are promoted to the top of the ranking. We further find positive effects of carefully choosing the candidate set of passages for a query: significant performance improvements are obtained when replacing a BM25 candidate set with a candidate set derived from entity-support passages. While entity aspect linking is not widely studied in the IR community, we hope that the demonstration of its merits leads to further research in this area.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1846017. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

[1] Krisztian Balog, Marc Bron, and Maarten De Rijke. 2011. Query Modeling for Entity Search Based on Terms, Categories, and Examples. *ACM Trans. Inf. Syst.* 29, 4, Article 22 (Dec. 2011), 31 pages. https://doi.org/10.1145/2037661.2037667

[2] Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: A Collaboratively Created Graph Database for Structuring Human Knowledge. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data* (Vancouver, Canada) *(SIGMOD '08)*. Association for Computing Machinery, New York, NY, USA, 1247–1250. https://doi.org/10.1145/1376616.1376746

[3] Shubham Chatterjee and Laura Dietz. 2019. 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* (Santa Clara, CA, USA) *(ICTIR '19)*. Association for Computing Machinery, New York, NY, USA, 221–224. https://doi.org/10.1145/3341981.3344243

[4] Jeffrey Dalton, Laura Dietz, and James Allan. 2014. Entity Query Feature Expansion Using Knowledge Base Links. In *Proceedings of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Gold Coast, Queensland, Australia) *(SIGIR '14)*. Association for Computing Machinery, New York, NY, USA, 365–374. https://doi.org/10.1145/2600428.2609628

[5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. https://doi.org/10.18653/v1/N19-1423

[6] Laura Dietz. 2019. 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* (Paris, France) *(SIGIR'19)*. Association for Computing Machinery, New York, NY, USA, 215–224. https://doi.org/10.1145/3331184.3331257

[7] Laura Dietz and John Foley. 2019. TREC CAR Y3: Complex Answer Retrieval Overview.. In *Proceedings of Text REtrieval Conference (TREC)*.

[8] Paolo Ferragina and Ugo Scaiella. 2010. TAGME: On-the-Fly Annotation of Short Text Fragments (by Wikipedia Entities). In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management* (Toronto, ON, Canada) *(CIKM '10)*. Association for Computing Machinery, New York, NY, USA, 1625–1628. https://doi.org/10.1145/1871437.1871689

[9] Darío Garigliotti and Krisztian Balog. 2017. On Type-Aware Entity Retrieval. In *Proceedings of the ACM SIGIR International Conference on Theory of Information Retrieval* (Amsterdam, The Netherlands) *(ICTIR '17)*. Association for Computing Machinery, New York, NY, USA, 27–34. https://doi.org/10.1145/3121050.3121054

[10] Emma J Gerritse, Faegheh Hasibi, and Arjen P de Vries. 2020. Graph-Embedding Empowered Entity Retrieval. In *European Conference on Information Retrieval*. Springer, 97–110.

[11] David Graus, Manos Tsagkias, Wouter Weerkamp, Edgar Meij, and Maarten de Rijke. 2016. Dynamic Collective Entity Representations for Entity Ranking. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining* (San Francisco, California, USA) *(WSDM '16)*. Association for Computing Machinery, New York, NY, USA, 595–604. https://doi.org/10.1145/2835776.2835819

[12] Jiafeng Guo, Gu Xu, Xueqi Cheng, and Hang Li. 2009. Named Entity Recognition in Query. In *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Boston, MA, USA) *(SIGIR '09)*. Association for Computing Machinery, New York, NY, USA, 267–274. https://doi.org/10.1145/1571941.1571989

[13] Faegheh Hasibi, Krisztian Balog, and Svein Erik Bratsberg. 2016. Exploiting Entity Linking in Queries for Entity Retrieval. In *Proceedings of the 2016 ACM International Conference on the Theory of Information Retrieval* (Newark, Delaware, USA) *(ICTIR '16)*. Association for Computing Machinery, New York, NY, USA, 209–218. https://doi.org/10.1145/2970398.2970406

[14] Rianne Kaptein, Pavel Serdyukov, Arjen De Vries, and Jaap Kamps. 2010. Entity Ranking Using Wikipedia as a Pivot. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management* (Toronto, ON, Canada) *(CIKM '10)*. Association for Computing Machinery, New York, NY, USA, 69–78. https://doi.org/10.1145/1871437.1871451

[15] Aniruddha Kembhavi, Minjoon Seo, Dustin Schwenk, Jonghyun Choi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. Are You Smarter Than a Sixth Grader? Textbook Question Answering for Multimodal Machine Comprehension. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4999–5007.

[16] Victor Lavrenko and W. Bruce Croft. 2001. Relevance-Based Language Models. *SIGIR Forum* 51, 2 (Aug. 2001), 260–267. https://doi.org/10.1145/3130348.3130376

[17] Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick Van Kleef, Sören Auer, et al. 2015. Dbpedia–A Large-scale, Multilingual Knowledge Base Extracted from Wikipedia. *Semantic web* 6, 2 (2015), 167–195.

[18] Thomas Lin, Patrick Pantel, Michael Gamon, Anitha Kannan, and Ariel Fuxman. 2012. Active Objects: Actions for Entity-Centric Search. In *Proceedings of the 21st International Conference on World Wide Web* (Lyon, France) *(WWW '12)*. Association for Computing Machinery, New York, NY, USA, 589–598. https://doi.org/10.1145/2187836.2187916

[19] 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. https://doi.org/10.1007/s10791-015-9267-x

[20] Edgar Meij, Marc Bron, Laura Hollink, Bouke Huurnink, and Maarten de Rijke. 2011. Mapping Queries to the Linking Open Data cloud: A case study using DBpedia. *Journal of Web Semantics* 9, 4 (2011), 418 – 433. https://doi.org/10.1016/j.websem.2011.04.001 JWS special issue on Semantic Search.

[21] Donald Metzler and W. Bruce Croft. 2005. 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* (Salvador, Brazil) *(SIGIR '05)*. Association for Computing Machinery, New York, NY, USA, 472–479. https://doi.org/10.1145/1076034.1076115

[22] Federico Nanni, Simone Paolo Ponzetto, and Laura Dietz. 2018. Entity-Aspect Linking: Providing Fine-Grained Semantics of Entities in Context. In *Proceedings of the 18th ACM/IEEE on Joint Conference on Digital Libraries* (Fort Worth, Texas, USA) *(JCDL '18)*. Association for Computing Machinery, New York, NY, USA, 49–58. https://doi.org/10.1145/3197026.3197047

[23] Fedor Nikolaev, Alexander Kotov, and Nikita Zhiltsov. 2016. 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* (Pisa, Italy) *(SIGIR '16)*. Association for Computing Machinery, New York, NY, USA, 435–444. https://doi.org/10.1145/2911451.2911545

[24] Jeffrey Pound, Peter Mika, and Hugo Zaragoza. 2010. Ad-Hoc Object Retrieval in the Web of Data. In *Proceedings of the 19th International Conference on World Wide Web* (Raleigh, North Carolina, USA) *(WWW '10)*. Association for Computing Machinery, New York, NY, USA, 771–780. https://doi.org/10.1145/1772690.1772769

[25] Jordan Ramsdell and Laura Dietz. 2020. A Large Test Collection for Entity Aspect Linking. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management* (Virtual Event, Ireland) *(CIKM '20)*. Association for Computing Machinery, New York, NY, USA, 3109–3116. https://doi.org/10.1145/3340531.3412875

[26] Hadas Raviv, David Carmel, and Oren Kurland. 2012. A Ranking Framework for Entity Oriented Search Using Markov Random Fields. In *Proceedings of the 1st Joint International Workshop on Entity-Oriented and Semantic Search* (Portland, Oregon, USA) *(JIWES '12)*. Association for Computing Machinery, New York, NY, USA, Article 1, 6 pages. https://doi.org/10.1145/2379307.2379308

[27] Stephen Robertson and Hugo Zaragoza. 2009. *The probabilistic relevance framework: BM25 and beyond*. Now Publishers Inc.

[28] Michael Schuhmacher, Laura Dietz, and Simone Paolo Ponzetto. 2015. Ranking Entities for Web Queries Through Text and Knowledge. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management* (Melbourne, Australia) *(CIKM '15)*. Association for Computing Machinery, New York, NY, USA, 1461–1470. https://doi.org/10.1145/2806416.2806480

[29] Alberto Tonon, Gianluca Demartini, and Philippe Cudré-Mauroux. 2012. Combining Inverted Indices and Structured Search for Ad-Hoc Object Retrieval. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Portland, Oregon, USA) *(SIGIR '12)*. Association for Computing Machinery, New York, NY, USA, 125–134. https://doi.org/10.1145/2348283.2348304

[30] Ikuya Yamada, Akari Asai, Hiroyuki Shindo, Hideaki Takeda, and Yoshiyasu Takefuji. 2018. Wikipedia2Vec: An Optimized Tool for Learning Embeddings of Words and Entities from Wikipedia. *CoRR* abs/1812.06280 (2018). arXiv:1812.06280 http://arxiv.org/abs/1812.06280

[31] Nikita Zhiltsov, Alexander Kotov, and Fedor Nikolaev. 2015. 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* (Santiago, Chile) *(SIGIR '15)*. Association for Computing Machinery, New York, NY, USA, 253–262. https://doi.org/10.1145/2766462.2767756