

Unsupervised Context Retrieval for Long-tail Entities

Dario Garigliotti*

University of Stavanger, Norway
 dario.garigliotti@uis.no

Miguel Martinez

Signal, 145 City Road, London EC1V 1AZ, UK
 miguel.martinez@signal-ai.com

Dyaa Albakour

Signal, 145 City Road, London EC1V 1AZ, UK
 dyaa.albakour@signal-ai.com

Krisztian Balog

University of Stavanger, Norway
 krisztian.balog@uis.no

ACM Reference Format:

Dario Garigliotti, Dyaa Albakour, Miguel Martinez, and Krisztian Balog. 2019. Unsupervised Context Retrieval for Long-tail Entities. In *The 2019 ACM SIGIR International Conference on the Theory of Information Retrieval (ICTIR'19)*, October 2–5, 2019, Santa Clara, CA, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3341981.3344244>

1 INTRODUCTION

Entities, like people and organizations, are at the core of multiple techniques and applications devised to automatically process text documents [6, 9, 18]. One of these applications is reputation analysis, for example, in monitoring through news streams the stock performance of a company or a scandal around a celebrity. This task, termed *online reputation management* [1], is concerned with identifying and mining textual contexts from media streams, e.g. sentences in news articles, in which a given entity occurs. In order to identify such contexts, it is necessary to detect mentions of entities and disambiguate them, making this problem closely related to *entity linking* (EL) [5, 12, 16, 19]. EL is a prime example of entity-centered problems that are naturally addressed with methods that rely on knowledge bases (KBs)—large-scale repositories containing properties and facts about entities in a structured format. The suitability of such semantics-aware methods, nevertheless, is limited by the amount of information available for a given entity. EL benchmarks, for example, usually report good performance, however, the evaluation datasets are focused on the head of the entities distribution (i.e., popular entities), and systems are trained to perform well on those entities [12]. Conversely, entities in the *long tail* are out of focus for most entity-oriented approaches [19].

In this paper, we address the problem of identifying contexts for monitoring long-tail entities in news articles. Specifically, we propose a generative probabilistic framework to rank contexts where an alias of a long-tail entity occurs, by leveraging support information. As shown experimentally, our approach is in particular able to cope with out-of-KB entities in a robust manner.

* Work done while Dario Garigliotti was visiting Signal.

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 ACM 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.
 ICTIR'19, , October 2-5, 2019, Santa Clara, California.
 © 2019 Association for Computing Machinery.
 ACM ISBN 123-4567-24-567/08/06...\$15.00
<https://doi.org/10.1145/3341981.3344244>




Figure 1: Overview of our proposed framework.

To illustrate the main ideas behind our framework, consider Fig. 1, where we aim to monitor the long-tail input entity $e = \text{ISAI}$, an entrepreneurs’ fund. Then, we wish to identify contexts (here: sentences) that mention one of the known aliases of e (e.g., “ISAI” or “Isai”). It is often the case that two or more entities share the same alias. For example, in the context “Capital firm **ISAI** just raised a new \$175 million fund,” the alias indeed refers to ISAI, meanwhile in “S.J. Surya’s **Isai** begins with a curious disclaimer,” it refers to an Indian movie. Deciding which entity a given mention refers to is known as entity disambiguation, a well-studied subtask in the context of EL. Modern entity disambiguation approaches leverage rich semantic information associated with entities in a KB [2]—information which is lacking, or missing altogether, for long-tail entities. This renders supervised learning methods unsuitable and motivates us to develop an unsupervised solution. Our main underlying idea is to utilize established entities (referred to as *support entities*, with rich KB entries) similar to the input entity, and their contexts (*support contexts*), to rank the contexts in which the input entity is mentioned. Therefore, we first rank support entities (SER step in Fig. 1) according to how similar they are to e , e.g., $\tilde{e} = \text{Fundica}$. Next, we identify support contexts for each support entity via traditional EL (SCR step in Fig. 1), e.g., $\tilde{c} = \text{"Fundica held the finals of their Roadshow."}$ Finally, we rank contexts for our input entity by considering their similarity to the support contexts (CCR step in Fig. 1).

For evaluation, we build a curated test collection of 165 long-tail entities (comprised of 92 in-KB entities with Wikipedia pages and 73 out-of-KB entities) and collect relevance assessments for contexts retrieved from a large news corpus. We show experimentally that our method substantially outperforms both an EL and a retrieval baseline. Moreover, it performs just as well for entities without any representation in Wikipedia, a setting where the EL baseline fails inevitably. The contributions that we make in this work are: (i) an unsupervised context retrieval approach for long-tail entities, (ii) a purpose-built test collection, and (iii) a detailed component-level evaluation and an analysis of performance for in-KB vs. out-of-KB entities.

2 PROBLEM STATEMENT

We first introduce some terminology and clarify our assumptions.

Each entity e in a knowledge base (KB) has a unique identifier. A *surface form*, or *alias*, of e is a textual phrase that is used to refer to e . For example, “Obama” or “US President” are aliases of the entity whose identifier is Barack_Obama. An entity *mention* is an alias occurring in some context (span of text). For simplicity, we use sentences as contexts, i.e., the unit of retrieval is a sentence.

In the rest of this work, e is a long-tail entity, and C is a set of contexts each with at least one alias of e being mentioned. We assume that we know (i) a brief textual description e_{desc} , (ii) the entity type (among Person, Location, or Organization), and (iii) all the entity surface forms. These assumptions are intuitive in our scenario of entity monitoring in media, where all this information can be easily traced and provided by an external component. We also assume that we have access to a large collection of entities \mathcal{E} from an entity catalog or a KB, e.g. Wikipedia. Another assumption is that all entity mentions in the contexts are detected.

Given e and C , we aim to rank each context $c \in C$ according to the likelihood that the (potentially ambiguous) entity mention in c refers to e . The context retrieval problem reduces to assigning a score $s(c, e)$ to each pair (c, e) indicating our confidence that e is the entity mentioned in c .

3 GENERATIVE MODELING FRAMEWORK

We introduce a generative probabilistic model for scoring a context $c \in C$ according to $P(c|e)$, i.e., the probability that the surface form mentioned in c refers to the long-tail entity e . Formally:

$$\begin{aligned} P(c|e) &= \sum_{\tilde{e}} P(c|e, \tilde{e})P(\tilde{e}|e) \\ &= \sum_{\tilde{e}} \left(\sum_{\tilde{c}} P(c|e, \tilde{e}, \tilde{c})P(\tilde{c}|e, \tilde{e}) \right) P(\tilde{e}|e) \\ &= \sum_{\tilde{e}} \left(\sum_{\tilde{c}} P(c|e, \tilde{c})P(\tilde{c}|\tilde{e}) \right) P(\tilde{e}|e) \end{aligned} \quad (1)$$

This model has three components (illustrated in Fig. 1):

- (i) $P(\tilde{e}|e)$ (*Support Entity Ranking*, or SER) expresses the importance of a *support* entity \tilde{e} for the given long-tail entity e ;
- (ii) $P(\tilde{c}|\tilde{e})$ (*Support Context Ranking*, or SCR) represents the importance of a *support* context \tilde{c} for a support entity \tilde{e} ; and
- (iii) $P(c|e, \tilde{c})$ (*Context-to-Context Ranking*, or CCR) is the importance of a support context \tilde{c} for a context c , given that an alias of the long-tail entity e is mentioned in c .

The last expression in Eq. (1) is obtained after assuming independence of c w.r.t. \tilde{e} given \tilde{c} , and of \tilde{c} w.r.t. e given \tilde{e} .

3.1 Support Entity Ranking

Since e has a limited representation in KBs and in contexts where its known aliases might be mentioned, the textual description e_{desc} we assume known (cf. Sect. 2) is key to find relevant entities. It is appropriate then to consider a text matching method that ranks support entities for e . This method can be, for example, a retrieval model that uses e_{desc} to query the entity collection \mathcal{E} . This collection should contain mostly entities with good data representation, e.g. by being taken from an actual KB.

Given an entity ranking function ser , we denote the importance score of \tilde{e} for e by $ser(e, \tilde{e})$. We can then estimate $P(\tilde{e}|e)$ by $ser(e, \tilde{e})/\sum_{e' \in \tilde{\mathcal{E}}} ser(e, e')$.

3.2 Support Context Ranking

Once N support entities $\tilde{\mathcal{E}} = (\tilde{e}_1, \dots, \tilde{e}_N)$ are ranked for e , we aim to exploit the contexts where it is known that entities from $\tilde{\mathcal{E}}$ are mentioned. Specifically, for each $\tilde{e}_i, i \in 1, \dots, N$, let $C_{\tilde{e}_i}$ be a collection of contexts, such that for each $\tilde{c}_x \in C_{\tilde{e}_i}$, \tilde{c}_x contains a mention *linked* to \tilde{e}_i . We assume that each of these collections is previously built, by applying an entity linking system over a universe of context candidates. Then, we take M_i support contexts $\tilde{C}_i = \{\tilde{c}_1, \dots, \tilde{c}_{M_i}\}$ among all the contexts in $C_{\tilde{e}_i}$. For each \tilde{C}_i , its size M_i would be informative enough, since the “support” entity collection \mathcal{E} is mostly head-oriented, and there the EL system should link each \tilde{e}_i from a sufficiently large number of contexts. We refer to the contexts of any \tilde{C} as *support contexts*, and say that the previously applied EL method links an entity mention in \tilde{c} to \tilde{e} with confidence $conf(\tilde{c}, \tilde{e})$. The set of all support contexts for the long-tail entity e is defined as $\tilde{C} = \bigcup_i \tilde{C}_i$. We can then estimate $P(\tilde{c}|\tilde{e})$ with $conf(\tilde{c}, \tilde{e})/\sum_{c' \in \tilde{C}} conf(c', \tilde{e})$.

3.3 Context-to-Context Ranking

Given e and a context $c \in C$, $ccr(c, \tilde{c})$ denotes the relevance score of c for each support context $\tilde{c} \in \tilde{C}$. In order to estimate $P(c|e, \tilde{c})$, we define the following estimator: $ccr(c, \tilde{c})/\sum_{c' \in C} ccr(c', \tilde{c})$.

4 EXPERIMENTAL SETUP

This section describes the estimators, dataset, evaluation metrics and baselines that we use in our experiments.

Component estimators. For estimating the SER component (cf. Sect. 3.1), we take articles from Wikipedia (2018-10-01 dump) to be our entity collection \mathcal{E} . We then approach entity ranking via a standard unstructured retrieval model, specifically, BM25 [17] with parameters set as $k1 = 1.2$ and $b = 0.8$, as recommended in [11]. The entity description e_{desc} is used as a query over the opening_text field to rank at most 200 support entities. We refer to this SER definition as *basic*. We also consider two variants to this method: (i) *pop*, in which the basic retrieval score is multiplied by the popularity of the support entity estimated by the number of incoming links of its Wikipedia article; and (ii) *types*, where we filter out, from the ranking, entities which do not have common types (assigned from the DBpedia Ontology) with the type of e ; *types* corresponds to the strict filtering method in [7]. Type information has been shown to significantly improve retrieval performance, e.g., in the scenario of ad hoc entity retrieval [8]. A parameter $N \in \{50, 100\}$ controls the cut-off of the support entity ranking.

Secondly, we describe the SCR estimation (cf. Sect. 3.2). We work with a sample of 5 million news articles, published between Oct. 2017 and Sept. 2018, provided by commercial news aggregators. We perform entity linking on the content of each article, using Spotlight [4] version 1.0.0 (with the latest model released for English language: 2016-10)¹. For each entity mention in an article, we retain an entity linking candidate only if its final confidence score is above a threshold θ of 0.9, to ensure high-quality support information. Then, from each article, we pick a single context (sentence) with an entity mention that has at least one retained candidate, regardless

¹<https://sourceforge.net/projects/dbpedia-spotlight/files/2016-10/en/model/en.tar.gz/download>

which entity it is. Given a support entity \tilde{e} , $C_{\tilde{e}}$ is the set of all contexts where \tilde{e} was linked, and so we take \tilde{C} for \tilde{e} as any non-empty subset of $C_{\tilde{e}}$. A parameter $M \in \{50, 100\}$ controls the size of \tilde{C} .

Finally, regarding the estimation of the CCR component (cf. Sect. 3.3), we consider two approaches to model ccr . Firstly, we define $ccr(c, \tilde{c})$ to be the retrieval score (BM25, $k1 = 1.2, b = 0.8$) of c when using \tilde{c} as a query over an index of contexts. We refer to this method as *retrieval*. Alternatively, we model $ccr(c, \tilde{c})$ as a semantic similarity between the long-tail entity context c and a support context \tilde{c} . Specifically, we define a *semantic* method to be the cosine similarity between the vectors \mathbf{v}_c and $\mathbf{v}_{\tilde{c}}$, where \mathbf{v}_s is the average, for all the terms in a context s , of term vectors in the pre-trained *word2vec* word embeddings [14].

Test collection. In order to ensure quality in the earliest step, we manually identify a set E of 165 long-tail entities. These correspond to either entities that have recently emerged, or entities with certain lifespan yet presenting difficulties to be linked with reasonable accuracy. 162 out of the 165 entities (i.e., 98.18%) are of type Organization, and 3 are people. After requesting for responses to the Wikipedia RESTful API², we find that 73 out of the 165 entities (i.e., 44.24%)—all organizations—do not have a corresponding Wikipedia page by October 1st, 2018. For this subset of entities, the average number of known surface forms is 1.07 per entity; for the remainder 92 in Wikipedia, the average is 1.11 per entity.

For each long-tail entity $e \in E$, we search for contexts where at least one of its surface forms occurs. Specifically, for a given e , we rank a set C of at most 5,000 contexts from a proprietary collection using each of the possible *configurations* (i.e., combinations of estimators and parameter settings for SER, SCR and CCR components; cf. Sect. 3). We use the well-established top- k pooling technique [20] for our relevance assessment. We take the top 20 contexts from each of these rankings, and build a pool. Each of these 4,536 pooled contexts is assessed by an expert media analyst regarding whether it is relevant or not for its long-tail entity.

Evaluation Metrics. We use two evaluation metrics: Mean Average Precision (MAP) and Mean Reciprocal Ranking (MRR), both very sensitive to highly ranked relevant contexts.

Baselines. Our baseline, a strong sentence retrieval model [3], ranks the contexts in C for a long-tail entity e by the standard BM25 retrieval score of querying with e_{desc} an index of C .³ (Note that this is different from the *retrieval* CCR component described above.) We also apply the Spotlight entity linker, previously used for estimating SCR (cf. Sect. 3.2), over each context in the universe for all long-tail entities, $\bigcup_e C_e$. We experiment with two values for the Spotlight final confidence threshold θ : 0.6, as the best performing Spotlight setting [13], and 0.9 as before.

5 EXPERIMENTAL RESULTS

Our experiments address three research questions: (RQ1) What is the best way to estimate each component?, (RQ2) How does our approach perform for context retrieval?, and (RQ3) How does it perform for entities with and without representation in a KB?

²<https://en.wikipedia.org/w/api.php>

³We also experimented with a combined method that mixes entity frequency and rarity [3], but it performed worse than our BM25 baseline. It appears that noise is introduced when counting mentions of ambiguous aliases in contexts.

Table 1: Context retrieval performance, measured in terms of MAP and MRR metrics. Best values per block are typeset in bold. Statistical significance is tested using a two-tailed paired t-test at $p < 0.05$ ([†]) and $p < 0.001$ ([‡]). For each line in the *pop* and *types* blocks, significance is tested against the corresponding row of the *basic* block.

<i>ser</i>	<i>N</i>	<i>M</i>	<i>ccr</i>			
			semantic		retrieval	
			MAP	MRR	MAP	MRR
basic	50	50	0.5195	0.6133	0.2758	0.4440
	50	100	0.5076	0.6138	0.2796	0.4453
	100	50	0.5133	0.6357	0.2927	0.4449
	100	100	0.5081	0.6341	0.2954	0.4489
pop	50	50	0.4441 [‡]	0.5738 [†]	0.2636	0.4030
	50	100	0.4499 [‡]	0.5738	0.2695	0.3826 [†]
	100	50	0.4564 [‡]	0.6061 [†]	0.2683 [†]	0.3988 [†]
	100	100	0.4633[‡]	0.6200	0.2706 [†]	0.4067 [†]
types	50	50	0.4512 [‡]	0.5818	0.2537	0.3942
	50	100	0.4475 [‡]	0.5818	0.2558	0.3964
	100	50	0.4544[‡]	0.5866 [†]	0.2504 [†]	0.3828 [†]
	100	100	0.4477 [‡]	0.5881[†]	0.2572 [†]	0.3913 [†]

5.1 Framework Components

Our probabilistic model (cf. Sect. 3) consists of three components: (i) Support Entity Ranking (SER), estimated by defining the function *ser* using either basic, pop or types, and with cut-off parameter N ; (ii) Support Context Ranking (SCR), estimated via entity linking confidence, with parameter M for the size of the support contexts set; and (iii) Context-to-Context Ranking (CCR), modeled with retrieval or semantic estimator to define the function *ccr*.

Table 1 presents the performance of every possible configuration for context retrieval. In order to answer (RQ1: what are the best component estimators?), we first notice that SER defined as *basic* outperforms its *pop* and *types* variants. It does that for each of both CCR settings, and for any pair of values for the (N, M) parameters. The differences in favor of SER are highly significant in terms of MAP, and less or not significant according to MRR. Filtering by common type should help retain relevant support entities, but the results suggest that, even with slight differences in terms of MRR, the performance of the whole ranking in terms of MAP is degraded. Secondly, there is not a clear pattern in the performance when comparing settings for the M parameter. Regarding the last component, the semantic definition of CCR leads to the best estimator. This shows that the vocabulary mismatch affects CCR, a gap against which the semantic method is more robust.

5.2 Unsupervised Context Retrieval

Moving to address RQ2, we compare, against the baseline, the best performing configuration of our approach, found in Sect. 5.1 (i.e., using basic SER, $N = 50, M = 50$, and semantic CCR). The results in Table 2 corresponding to the entire set E of long-tail entities show that our method outperforms the baseline with high significance. This validates the suitability of our framework, that leverages support information to rank contexts, for long-tail entities. The very low performance of Spotlight 2016-10 confirms that most

Table 2: Context retrieval performance, comparing each of the baselines and the best configuration for our approach (according to MAP; cf. Table 1). For a given metric, the symbol \blacktriangle denotes statistical significance against the baseline, tested using a two-tailed paired t-test at $p < 0.001$.

Entities	All		In Wikipedia		Not in Wikipedia	
Method	MAP	MRR	MAP	MRR	MAP	MRR
Spotlight, $\theta = 0.6$	0.0406	0.0826	0.0726	0.1475	0.0004	0.0008
Spotlight, $\theta = 0.9$	0.0393	0.0811	0.0701	0.1447	0.0004	0.0008
Baseline	0.2248	0.3357	0.1732	0.2787	0.2899	0.4077
Our method	0.5195\blacktriangle	0.6133\blacktriangle	0.4900\blacktriangle	0.6032\blacktriangle	0.5565\blacktriangle	0.6261\blacktriangle

of the long-tail entities that happen to exist by October 2018 were indeed recently added between those two dates.

5.3 Out-of-KB Entities

To answer RQ3, we add the two last blocks in Table 2 that present the retrieval performances in the two disjoint subsets of E , namely with and without a corresponding Wikipedia article. We observe that our approach not only outperforms the baseline highly significantly in each of both subsets, but is also robust for the subset of entities that are not present in Wikipedia. It is in this subset of entities, clearly long-tail, that our approach performs best.

6 RELATED WORK

Blanco and Zaragoza [3] are the first to formalize and evaluate the task of finding sentences to support the relationship between an entity and an ad-hoc query. In our task, we relax the query-aware constraint by only ranking contexts for an entity. Another difference is that we require for a context to contain an entity mention. Closely developed with production-scale media monitoring, recent related work has addressed a number of applications like tracking entity popularity in media [9], and entity-centric approaches for topic modeling [18]. Färber et al. [6] formalize a series of challenges on missing entities and/or surface forms in KBs for entity linking, focusing on the detection of emerging entities in news articles. Our work faces all these long-tail entity challenges, i.e., whether or not a surface form matches any known surface form for the entities in the KB, and whether or not the correct entity actually exists in the KB. Reinanda et al. [15] introduce EIDF, an entity-independent, supervised learning model to retain documents in news that are vital to enhance a textual entity profile. Our task concerns with much smaller semantic units (sentences) as contexts for entities. Then, in our setting, an unsupervised approach appears more suitable, since structural document features they use, like informativeness or timeliness, are no longer able to be captured.

Supervised learning models for EL highly depend on fine-tuned training data from KBs during the candidate selection stage [16]. Related work also finds that EL datasets are skewed towards popular and frequent entities of general-purpose KBs, and emphasizes the need for datasets that focus on the long tail [19]. In this line, a very large proportion of entities lack of Wikipedia coverage, leading to a long tail of entities unlinked by state-of-the-art systems [5]. Very recently, a systematic study verifies that EL systems are heavily optimized on head entities, yet perform worst for infrequent, highly ambiguous entities [12].

7 CONCLUSIONS

We have presented an unsupervised approach to retrieve contexts for long-tail entities. We built a test collection, consisting in a set of long-tail entities manually identified, and relevance assessments for contexts. The experimental results show that our approach outperforms a text matching baseline, and also is robust for entities without representation in Wikipedia.

In a scenario of a newly emerged entity under media management, the entity will likely lack a KB entry during a long period [9]. Our method can be used in such cases to retrieve contexts, powering entity mining in news articles. We see at least two additional applications of our approach. Firstly, the top-ranked contexts identified by our approach could be used as weakly labeled instances in training entity linking models by distant supervision. Secondly, these contexts could be used for augmenting textual entity profiles, e.g., by leveraging information from filtered news articles [15] and by incorporating media streams in the entity representation [10].

REFERENCES

- [1] Enrique Amigó, Adolfo Corujo, Julio Gonzalo, Edgar Meij, and Maarten De Rijke. 2012. Overview of RepLab 2012: Evaluating Online Reputation Management Systems. *CEUR Workshop Proceedings* 1178 (01 2012).
- [2] Krisztian Balog. 2018. *Entity-Oriented Search*. The Information Retrieval Series, Vol. 39. Springer.
- [3] Roi Blanco and Hugo Zaragoza. 2010. Finding Support Sentences for Entities. In *Proc. of SIGIR*. 339–346.
- [4] Joachim Daiber, Max Jakob, Chris Hokamp, and Pablo N. Mendes. 2013. Improving Efficiency and Accuracy in Multilingual Entity Extraction. In *Proc. of I-Semantics*. 121–124.
- [5] José Esquivel, Dyaa Albakour, Miguel Martinez, David Corney, and Samir Moussa. 2017. On the Long-Tail Entities in News. In *Proc. of ECIR*. 691–697.
- [6] Michael Färber, Achim Rettinger, and Boulos Asmar. 2016. On Emerging Entity Detection. In *Proc. of EKAW*. 223–238.
- [7] Dario Garigliotti and Krisztian Balog. 2017. On Type-Aware Entity Retrieval. In *Proc. of ICTIR*. 27–34.
- [8] Dario Garigliotti, Faegheh Hasibi, and Krisztian Balog. 2019. Identifying and exploiting target entity type information for ad hoc entity retrieval. *Information Retrieval Journal* 22, 3 (2019), 285–323.
- [9] David Graus, Daan Odijk, and Maarten de Rijke. 2018. The birth of collective memories: Analyzing emerging entities in text streams. *JASIST* 69, 6 (2018), 773–786.
- [10] David Graus, Manos Tsagkias, Wouter Weerkamp, Edgar Meij, and Maarten de Rijke. 2016. Dynamic Collective Entity Representations for Entity Ranking. In *Proc. of WSDM*. 595–604.
- [11] 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 *Proc. of SIGIR*. 1265–1268.
- [12] Filip Ilievski, Piek Vossen, and Stefan Schlobach. 2018. Systematic Study of Long Tail Phenomena in Entity Linking. In *Proc. of COLING*. 664–674.
- [13] Pablo N. Mendes, Max Jakob, Andrés García-Silva, and Christian Bizer. 2011. DBpedia Spotlight: Shedding light on the web of documents. In *Proc. of I-Semantics*. 1–8.
- [14] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed Representations of Words and Phrases and Their Compositionality. In *Proc. of NIPS*. 3111–3119.
- [15] Ridho Reinanda, Edgar Meij, and Maarten de Rijke. 2016. Document Filtering for Long-tail Entities. In *Proc. of CIKM*. 771–780.
- [16] Giuseppe Rizzo, Bianca Pereira, Andrea Varga, Marieke van Erp, and Amparo Elizabeth Cano Basave. 2017. Lessons learnt from the Named Entity rEcognition and Linking (NEEL) challenge series. *Semantic Web* 8, 5 (2017), 667–700.
- [17] Stephen Robertson and Hugo Zaragoza. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. *Found. Trends Inf. Retr.* 3, 4 (April 2009), 333–389.
- [18] Andreas Spitz and Michael Gertz. 2018. Exploring Entity-centric Networks in Entangled News Streams. In *Proc. of WWW*. 555–563.
- [19] Marieke van Erp, Pablo N. Mendes, Heiko Paulheim, Filip Ilievski, Julien Plu, Giuseppe Rizzo, and Jörg Wäitelonis. 2016. Evaluating Entity Linking: An Analysis of Current Benchmark Datasets and a Roadmap for Doing a Better Job. In *Proc. of LREC*. 4373–4379.
- [20] Ellen M. Voorhees. 2014. The Effect of Sampling Strategy on Inferred Measures. In *Proc. of SIGIR (SIGIR '14)*. 1119–1122.