

Teaching a New Dog Old Tricks: Resurrecting Multilingual Retrieval Using Zero-shot Learning

Sean MacAvaney¹, Luca Soldaini², and Nazli Goharian¹

¹ IR Lab, Georgetown University, Washington DC 20057, USA
`{firstname}@ir.cs.georgetown.edu`

² Amazon Alexa Search, Manhattan Beach CA, 90266, USA
`lssoldai@amazon.com`

Abstract. While billions of non-English speaking users rely on search engines every day, the problem of ad-hoc information retrieval is rarely studied for non-English languages. This is primarily due to a lack of data set that are suitable to train ranking algorithms. In this paper, we tackle the lack of data by leveraging pre-trained multilingual language models to transfer a retrieval system trained on English collections to non-English queries and documents. Our model is evaluated in a zero-shot setting, meaning that we use them to predict relevance scores for query-document pairs in languages never seen during training. Our results show that the proposed approach can significantly outperform unsupervised retrieval techniques for Arabic, Chinese Mandarin, and Spanish. We also show that augmenting the English training collection with some examples from the target language can sometimes improve performance.

1 Introduction

Every day, billions of non-English speaking users [22] interact with search engines; however, commercial retrieval systems have been traditionally tailored to English queries, causing an information access divide between those who can and those who cannot speak this language [39]. Non-English search applications have been equally under-studied by most information retrieval researchers. Historically, ad-hoc retrieval systems have been primarily designed, trained, and evaluated on English corpora (e.g., [1,5,6,23]). More recently, a new wave of supervised state-of-the-art ranking models have been proposed by researchers [11,14,21,24,26,35,37]; these models rely on neural architectures to rerank the head of search results retrieved using a traditional unsupervised ranking algorithm, such as BM25. Like previous ad-hoc ranking algorithms, these methods are almost exclusively trained and evaluated on English queries and documents.

The absence of rankers designed to operate on languages other than English can largely be attributed to a lack of suitable publicly available data sets. This aspect particularly limits supervised ranking methods, as they require samples for training and validation. For English, previous research relied on English collections such as TREC Robust 2004 [32], the 2009-2014 TREC Web Track [7], and MS MARCO [2]. No datasets of similar size exist for other languages.

While most of recent approaches have focused on ad hoc retrieval for English, some researchers have studied the problem of cross-lingual information retrieval. Under this setting, document collections are typically in English, while queries get translated to several languages; sometimes, the opposite setup is used. Throughout the years, several cross lingual tracks were included as part of TREC. TREC 6, 7, 8 [4] offer queries in English, German, Dutch, Spanish, French, and Italian. For all three years, the document collection was kept in English. CLEF also hosted multiple cross-lingual ad-hoc retrieval tasks from 2000 to 2009 [3]. Early systems for these tasks leveraged dictionary and statistical translation approaches, as well as other indexing optimizations [27]. More recently, approaches that rely on cross-lingual semantic representations (such as multilingual word embeddings) have been explored. For example, Vulic and Moens [34] proposed BWESG, an algorithm to learn word embeddings on aligned documents that can be used to calculate document-query similarity. Sasaki et al [28] leveraged a data set of Wikipedia pages in 25 languages to train a learning to rank algorithm for Japanese-English and Swahili-English cross-language retrieval. Litschko et al [20] proposed an unsupervised framework that relies on aligned word embeddings. Ultimately, while related, these approaches are only beneficial to users who can understand documents in two or more languages instead of directly tackling non-English document retrieval.

A few monolingual ad-hoc data sets exist, but most are too small to train a supervised ranking method. For example, TREC produced several non-English test collections: Spanish [12], Chinese Mandarin [31], and Arabic [25]. Other languages were explored, but the document collections are no longer available. The CLEF initiative includes some non-English monolingual datasets, though these are primarily focused on European languages [3]. Recently, Zheng et al. [40] introduced Sogou-QCL, a large query log dataset in Mandarin. Such datasets are only available for languages that already have large, established search engines.

Inspired by the success of neural retrieval methods, this work focuses on studying the problem of monolingual ad-hoc retrieval on non English languages using supervised neural approaches. In particular, to circumvent the lack of training data, we leverage transfer learning techniques to train Arabic, Mandarin, and Spanish retrieval models using English training data. In the past few years, transfer learning between languages has been proven to be a remarkably effective approach for low-resource multilingual tasks (e.g. [16,17,29,38]). Our model leverages a pre-trained multi-language transformer model to obtain an encoding for queries and documents in different languages; at train time, this encoding is used to predict relevance of query document pairs in English. We evaluate our models in a zero-shot setting; that is, we use them to predict relevance scores for query document pairs in languages never seen during training. By leveraging a pre-trained multilingual language model, which can be easily trained from abundant aligned [19] or unaligned [8] web text, we achieve competitive retrieval performance without having to rely on language specific relevance judgements. During the peer review of this article, a preprint [30] was published with similar observations as ours. In summary, our contributions are:

- We study zero shot transfer learning for IR in non-English languages.
- We propose a simple yet effective technique that leverages contextualized word embedding as multilingual encoder for query and document terms. Our approach outperforms several baselines on multiple non-English collections.
- We show that including additional in-language training samples may help further improve ranking performance.
- We release our code for pre-processing, initial retrieval, training, and evaluation of non-English datasets.³ We hope that this encourages others to consider cross-lingual modeling implications in future work.

2 Methodology

Zero-shot Multi-Lingual Ranking. Because large-scale relevance judgments are largely absent in languages other than English, we propose a new setting to evaluate learning-to-rank approaches: zero-shot cross-lingual ranking. This setting makes use of relevance data from one language that has a considerable amount of training data (e.g., English) for model training and validation, and applies the trained model to a different language for testing.

More formally, let \mathcal{S} be a collection of relevance tuples in the source language, and \mathcal{T} be a collection of relevance judgments from another language. Each relevance tuple $\langle \mathbf{q}, \mathbf{d}, r \rangle$ consists of a query, document, and relevance score, respectively. In typical evaluation environments, \mathcal{S} is segmented into multiple splits for training (\mathcal{S}_{train}) and testing (\mathcal{S}_{test}), such that there is no overlap of queries between the two splits. A ranking algorithm is tuned on \mathcal{S}_{train} to define the ranking function $R_{\mathcal{S}_{train}}(\mathbf{q}, \mathbf{d}) \in \mathbb{R}$, which is subsequently tested on \mathcal{S}_{test} . We propose instead tuning a model on all data from the source language (i.e., training $R_{\mathcal{S}}(\cdot)$), and testing on a collection from the second language (\mathcal{T}).

Datasets. We evaluate on monolingual newswire datasets from three languages: Arabic, Mandarin, and Spanish. The Arabic document collection contains 384k documents (LDC2001T55), and we use topics/relevance information from the 2001–02 TREC Multilingual track (25 and 50 topics, respectively). For Mandarin, we use 130k news articles from LDC2000T52. Mandarin topics and relevance judgments are utilized from TREC 5 and 6 (26 and 28 topics, respectively). Finally, the Spanish collection contains 58k articles from LDC2000T51, and we use topics from TREC 3 and 4 (25 topics each). We use the topics, rather than the query descriptions, in all cases except TREC Spanish 4, in which only descriptions are provided. The topics more closely resemble real user queries than descriptions.⁴ We test on these collections because they are the only document collections available from TREC at this time.⁵

We index the text content of each document using a modified version of Anserini with support for the languages we investigate [36]. Specifically, we add

³ <https://github.com/Georgetown-IR-Lab/multilingual-neural-ir>

⁴ Some have observed that the context provided by query descriptions are valuable for neural ranking, particularly when using contextualized language models [9].

⁵ https://trec.nist.gov/data/docs_noneng.html

Anserini support for Lucene’s Arabic and Spanish light stemming and stop word list (via `SpanishAnalyzer` and `ArabicAnalyzer`). We treat each character in Mandarin text as a single token.

Modeling. We explore the following ranking models:

- **Unsupervised baselines.** We use the Anserini [36] implementation of BM25, RM3 query expansion, and the Sequential Dependency Model (SDM) as unsupervised baselines. In the spirit of the zero-shot setting, we use the default parameters from Anserini (i.e., assuming no data of the target language).
- **PACRR** [14] models n-gram relationships in the text using learned 2D convolutions and max pooling atop a query-document similarity matrix.
- **KNRM** [35] uses learned Gaussian kernel pooling functions over the query-document similarity matrix to rank documents.
- **Vanilla BERT** [21] uses the BERT [10] transformer model, with a dense layer atop the classification token to compute a ranking score. To support multiple languages, we use the `base-multilingual-cased` pretrained weights. These weights were trained on Wikipedia text from 104 languages.

We use the embedding layer output from `base-multilingual-cased` model for PACRR and KNRM. In pilot studies, we investigated using cross-lingual MUSE vectors [8] and the output representations from BERT, but found the BERT embeddings to be more effective.

Experimental Setup. We train and validate models using TREC Robust 2004 collection [32]. TREC Robust 2004 contains 249 topics, 528k documents, and 311k relevance judgments in English (folds 1-4 from [15] for training, fold 5 for validation). Thus, the model is only exposed to English text in the training and validation stages (though the embedding and contextualized language models *are* trained on large amounts of unlabeled text in the languages). The validation dataset is used for parameter tuning and for the selection of the optimal training epoch (via nDCG@20). We train using pairwise softmax loss with Adam [18].

We evaluate the performance of the trained models by re-ranking the top 100 documents retrieved with BM25. We report MAP, Precision@20, and nDCG@20 to gauge the overall performance of our approach, and the percentage of judged documents in the top 20 ranked documents (judged@20) to evaluate how suitable the datasets are to approaches that did not contribute to the original judgments.

3 Results

We present the ranking results in Table 1. We first point out that there is considerable variability in the performance of the unsupervised baselines; in some cases, RM3 and SDM outperform BM25, whereas in other cases they underperform. Similarly, the PACRR and KNRM neural models also vary in effectiveness, though more frequently perform much worse than BM25. This makes sense because these models capture matching characteristics that are specific to English. For instance, n-gram patterns captured by PACRR for English do not necessarily transfer well to languages with different constituent order, such as

Table 1. Zero-shot multi-lingual results for various baseline and neural methods. Significant improvements and reductions in performance compared with BM25 are indicated with \uparrow and \downarrow , respectively (paired t-test by query, $p < 0.05$).

Ranker	P@20	nDCG@20	MAP	judged@20
Arabic (TREC 2002) [25]				
BM25	0.3470	0.3863	0.2804	99.0%
BM25 + RM3	0.3320	0.3705	\downarrow 0.2641	95.1%
SDM	0.3380	0.3775	\downarrow 0.2572	98.1%
PACRR multilingual	0.3270	0.3499	\downarrow 0.2517	96.4%
KNRM multilingual	0.3210	\downarrow 0.3415	\downarrow 0.2503	95.2%
Vanilla BERT multilingual	\uparrow 0.3790	0.4205	0.2876	97.4%
Arabic (TREC 2001) [25]				
BM25	0.5420	0.5933	0.3462	97.2%
BM25 + RM3	\downarrow 0.4700	0.5458	\downarrow 0.2903	85.6%
SDM	0.5140	0.5843	0.3213	96.2%
PACRR multilingual	\downarrow 0.3880	\downarrow 0.3933	\downarrow 0.2724	90.6%
KNRM multilingual	\downarrow 0.4140	\downarrow 0.4327	\downarrow 0.2742	91.0%
Vanilla BERT multilingual	0.5240	0.5628	0.3432	91.0%
Mandarin (TREC 6) [31]				
BM25	0.5962	0.6409	0.3316	89.6%
BM25 + RM3	\downarrow 0.5019	\downarrow 0.5571	0.2696	75.6%
SDM	0.5942	0.6320	0.3472	92.1%
PACRR multilingual	\downarrow 0.4923	\downarrow 0.5238	0.2856	79.0%
KNRM multilingual	\downarrow 0.5308	\downarrow 0.5497	\downarrow 0.3107	80.8%
Vanilla BERT multilingual	\uparrow 0.6615	\uparrow 0.6959	\uparrow 0.3589	92.7%
Mandarin (TREC 5) [33]				
BM25	0.3893	0.4113	0.2548	85.4%
BM25 + RM3	\downarrow 0.2768	\downarrow 0.3021	\downarrow 0.1698	64.6%
SDM	\uparrow 0.4536	\uparrow 0.4744	\uparrow 0.2855	94.1%
PACRR multilingual	0.3786	0.3998	0.2331	83.2%
KNRM multilingual	\downarrow 0.3232	\downarrow 0.3449	\downarrow 0.2223	77.5%
Vanilla BERT multilingual	\uparrow 0.4589	\uparrow 0.5196	\uparrow 0.2906	92.0%
Spanish (TREC 4) [12]				
BM25	0.3080	0.3314	0.1459	83.8%
BM25 + RM3	0.3360	0.3358	\uparrow 0.2024	85.2%
SDM	0.2780	0.3061	0.1377	78.6%
PACRR multilingual	0.2440	0.2494	0.1294	69.4%
KNRM multilingual	0.3120	0.3402	0.1444	79.2%
Vanilla BERT multilingual	\uparrow 0.4400	\uparrow 0.4898	\uparrow 0.1800	85.6%
Spanish (TREC 3) [13]				
BM25	0.5220	0.5536	0.2420	84.8%
BM25 + RM3	\uparrow 0.6100	0.6236	\uparrow 0.3887	93.0%
SDM	0.4920	0.5178	0.2258	83.8%
PACRR multilingual	\downarrow 0.4140	\downarrow 0.4092	0.2260	76.0%
KNRM multilingual	0.5560	0.5700	0.2449	85.2%
Vanilla BERT multilingual	\uparrow 0.6400	\uparrow 0.6672	\uparrow 0.2623	90.8%

Table 2. Zero-Shot (ZS) and Few-Shot (FS) comparison for Vanilla BERT (multilingual) on each dataset. Within each metric and dataset, the top result is listed in bold. Significant increases from using FS are indicated with \uparrow (paired t-test, $p < 0.05$).

Dataset	P@20		nDCG@20		MAP	
	ZS	FS	ZS	FS	ZS	FS
Arabic 2002	0.3790	0.3690	0.4205	0.3905	0.2876	0.2822
Arabic 2001	0.5240	\uparrow 0.6020	0.5628	\uparrow 0.6405	0.3432	0.3529
Mandarin 6	0.6615	0.6808	0.6959	0.7099	0.3589	0.3537
Mandarin 5	0.4589	0.4643	0.5196	0.5014	0.2906	0.2895
Spanish 4	0.4400	\uparrow 0.5060	0.4898	\uparrow 0.5636	0.1800	\uparrow 0.2020
Spanish 3	0.6400	0.6560	0.6672	0.6825	0.2623	0.2684

Arabic (VSO instead of SVO). An interesting observation is that the Vanilla BERT model (which recall is only tuned on English text) generally outperforms a variety of approaches across three test languages. This is particularly remarkable because it is a single trained model that is effective across all three languages, without any difference in parameters. The exceptions are the Arabic 2001 dataset, in which it performs only comparably to BM25 and the MAP results for Spanish. For Spanish, RM3 is able to substantially improve recall (as evidenced by MAP), and since Vanilla BERT acts as a re-ranker atop BM25, it is unable to take advantage of this improved recall, despite significantly improving the precision-focused metrics. In all cases, Vanilla BERT exhibits judged@20 above 85%, indicating that these test collections are still valuable for evaluation.

To test whether a small amount of in-language training data can further improve BERT ranking performance, we conduct an experiment that uses the other collection for each language as additional training data. The in-language samples are interleaved into the English training samples. Results for this few-shot setting are shown in Table 2. We find that the added topics for Arabic 2001 (+50) and Spanish 4 (+25) significantly improve the performance. This results in a model significantly better than BM25 for Arabic 2001, which suggests that there may be substantial distributional differences in the English TREC 2004 training and Arabic 2001 test collections. We further back this up by training an “oracle” BERT model (training on the test data) for Arabic 2001, which yields a model substantially better (P@20=0.7340, nDCG@20=0.8093, MAP=0.4250).

4 Conclusion

We introduced a zero-shot multilingual setting for evaluation of neural ranking methods. This is an important setting due to the lack of training data available in many languages. We found that contextualized languages models (namely, BERT) have a big upper-hand, and are generally more suitable for cross-lingual performance than prior models (which may rely more heavily on phenomena exclusive to English). We also found that additional in-language training data may improve the performance, though not necessarily. By releasing our code and models, we hope that cross-lingual evaluation will become more commonplace.

Acknowledgements. This work was supported in part by ARCS Foundation.

References

1. Amati, G., Van Rijsbergen, C.J.: Probabilistic models of information retrieval based on measuring the divergence from randomness. *ACM Transactions on Information Systems (TOIS)* **20**(4), 357–389 (2002)
2. Bajaj, P., Campos, D., Craswell, N., Deng, L., Gao, J., Liu, X., Majumder, R., McNamara, A., Mitra, B., Nguyen, T., et al.: Ms marco: A human generated machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268* (2016)
3. Braschler, M.: Clef 2003—overview of results. In: *Workshop of the Cross-Language Evaluation Forum for European Languages*. pp. 44–63. Springer (2003)
4. Braschler, M., Schäuble, P., Peters, C.: Cross-language information retrieval (clir) track overview. In: *TREC (2000)*
5. Cao, Z., Qin, T., Liu, T.Y., Tsai, M.F., Li, H.: Learning to rank: from pairwise approach to listwise approach. In: *Proceedings of the 24th international conference on Machine learning*. pp. 129–136. ACM (2007)
6. Carpineto, C., Romano, G.: A survey of automatic query expansion in information retrieval. *Acm Computing Surveys (CSUR)* **44**(1), 1 (2012)
7. Collins-Thompson, K., Macdonald, C., Bennett, P., Diaz, F., Voorhees, E.M.: Trec 2014 web track overview. *Tech. rep., MICHIGAN UNIV ANN ARBOR* (2015)
8. Conneau, A., Lample, G., Ranzato, M., Denoyer, L., Jégou, H.: Word translation without parallel data. *arXiv preprint arXiv:1710.04087* (2017)
9. Dai, Z., Callan, J.: Deeper text understanding for ir with contextual neural language modeling. In: *SIGIR* (2019)
10. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: 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)*. pp. 4171–4186. Association for Computational Linguistics, Minneapolis, Minnesota (Jun 2019)
11. Guo, J., Fan, Y., Ai, Q., Croft, W.B.: A deep relevance matching model for ad-hoc retrieval. In: *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*. pp. 55–64. ACM (2016)
12. Harman, D.: Overview of the fourth text retrieval conference (trec-4). *NIST SPECIAL PUBLICATION SP* pp. 1–24 (1996)
13. Harman, D.K.: Overview of the third text retrieval conference (TREC-3). *DIANE Publishing* (1995)
14. Hui, K., Yates, A., Berberich, K., de Melo, G.: Pacrr: A position-aware neural ir model for relevance matching. *arXiv preprint arXiv:1704.03940* (2017)
15. Huston, S., Croft, W.B.: Parameters learned in the comparison of retrieval models using term dependencies. *Technical Report* (2014)
16. Johnson, M., Schuster, M., Le, Q.V., Krikun, M., Wu, Y., Chen, Z., Thorat, N., Viégas, F., Wattenberg, M., Corrado, G., Hughes, M., Dean, J.: Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association for Computational Linguistics* **5**, 339–351 (2017)
17. Kim, J.K., Kim, Y.B., Sarikaya, R., Fosler-Lussier, E.: Cross-lingual transfer learning for pos tagging without cross-lingual resources. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. pp. 2832–2838 (2017)
18. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: *ICLR* (2015)

19. Lample, G., Conneau, A.: Cross-lingual language model pretraining. arXiv preprint arXiv:1901.07291 (2019)
20. Litschko, R., Glavaš, G., Ponzetto, S.P., Vulić, I.: Unsupervised cross-lingual information retrieval using monolingual data only. In: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval. pp. 1253–1256. ACM (2018)
21. MacAvaney, S., Yates, A., Cohan, A., Goharian, N.: Cedr: Contextualized embeddings for document ranking. In: Proceedings of the 42Nd International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 1101–1104. SIGIR’19, ACM, New York, NY, USA (2019)
22. Max Roser, H.R., Ortiz-Ospina, E.: Internet. <https://ourworldindata.org/internet> (2019), last accessed: 2019/09/15
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. pp. 472–479. SIGIR ’05, ACM, New York, NY, USA (2005)
24. Mitra, B., Craswell, N., et al.: An introduction to neural information retrieval. *Foundations and Trends® in Information Retrieval* **13**(1), 1–126 (2018)
25. Oard, D.W., Gey, F.C.: The TREC 2002 arabic/english clir track. In: TREC (2002)
26. Onal, K.D., Zhang, Y., Altingovde, I.S., Rahman, M.M., Karagoz, P., Braylan, A., Dang, B., Chang, H.L., Kim, H., McNamara, Q., et al.: Neural information retrieval: At the end of the early years. *Information Retrieval Journal* **21**(2-3), 111–182 (2018)
27. Peters, C., Braschler, M., Clough, P.: *Multilingual information retrieval: From research to practice*. Springer Science & Business Media (2012)
28. Sasaki, S., Sun, S., Schamoni, S., Duh, K., Inui, K.: Cross-lingual learning-to-rank with shared representations. In: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers). pp. 458–463. Association for Computational Linguistics, New Orleans, Louisiana (Jun 2018)
29. Schuster, S., Gupta, S., Shah, R., Lewis, M.: Cross-lingual transfer learning for multilingual task oriented dialog. 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). pp. 3795–3805. Association for Computational Linguistics, Minneapolis, Minnesota (Jun 2019)
30. Shi, P., Lin, J.: Cross-lingual relevance transfer for document retrieval. ArXiv **abs/1911.02989** (2019)
31. Voorhees, E., Harman, D., Wilkinson, R.: The sixth text retrieval conference (trec-6). In: The Text REtrieval Conference (TREC). vol. 500, p. 240. NIST (1998)
32. Voorhees, E.M.: Overview of the TREC 2005 Robust Retrieval Track. In: TREC (2005)
33. Voorhees, E.M., Harman, D.: Overview of the fifth text retrieval conference (trec-5). In: TREC. vol. 97, pp. 1–28 (1996)
34. Vulić, I., Moens, M.F.: Monolingual and cross-lingual information retrieval models based on (bilingual) word embeddings. In: Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 363–372. SIGIR ’15, ACM, New York, NY, USA (2015)
35. Xiong, C., Dai, Z., Callan, J., Liu, Z., Power, R.: End-to-end neural ad-hoc ranking with kernel pooling. In: Proceedings of the 40th International ACM SIGIR conference on research and development in information retrieval. pp. 55–64. ACM (2017)

36. Yang, P., Fang, H., Lin, J.: Anserini: Reproducible ranking baselines using lucene. *J. Data and Information Quality* **10**, 16:1–16:20 (2018)
37. Yang, W., Zhang, H., Lin, J.: Simple applications of BERT for ad hoc document retrieval. arXiv preprint arXiv:1903.10972 (2019)
38. Yang, Z., Salakhutdinov, R., Cohen, W.W.: Transfer learning for sequence tagging with hierarchical recurrent networks. arXiv preprint arXiv:1703.06345 (2017)
39. Young, H.: The digital language divide. <http://labs.theguardian.com/digital-language-divide/> (2015), last accessed: 2019/09/15
40. Zheng, Y., Fan, Z., Liu, Y., Luo, C., Zhang, M., Ma, S.: Sogou-qcl: A new dataset with click relevance label. In: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. pp. 1117–1120. ACM (2018)