

Unveiling Temporal Trends in 19th Century Literature: An Information Retrieval Approach

Suchana Datta
University College Dublin, Ireland
suchana.datta@ucd.ie

Derek Greene
University College Dublin, Ireland
derek.greene@ucd.ie

Dwaipayan Roy
IISER Kolkata, India
dwaipayan.roy@iiserkol.ac.in

Gerardine Meaney
University College Dublin, Ireland
gerardine.meaney@ucd.ie

Abstract

In English literature, the 19th century witnessed a significant transition in styles, themes, and genres. Consequently, the novels from this period display remarkable diversity. This paper explores these variations by examining the evolution of term usage in 19th century English novels through the lens of information retrieval. By applying a query expansion-based approach to a decade-segmented collection of fiction from the British Library, we examine how related terms vary over time. Our analysis employs multiple standard metrics including Kendall's tau, Jaccard similarity, and Jensen-Shannon divergence to assess overlaps and shifts in expanded query term sets. Our results indicate a significant degree of divergence in the related terms across decades as selected by the query expansion technique, suggesting substantial linguistic and conceptual changes throughout the 19th century novels.

CCS Concepts

• **Information systems** → **Information extraction; Digital libraries and archives.**

Keywords

19th Century Fiction, Information Retrieval, Relevance Feedback

ACM Reference Format:

Suchana Datta, Dwaipayan Roy, Derek Greene, and Gerardine Meaney. 2024. Unveiling Temporal Trends in 19th Century Literature: An Information Retrieval Approach. In *The 2024 ACM/IEEE Joint Conference on Digital Libraries (JCDL '24)*, December 16–20, 2024, Hong Kong, China. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3677389.3702593>

1 Introduction and background

The 19th century was a period covering a profound transition in literature markets with the emergence of new literary forms, styles and thematic varieties. From the *Romanticism* and *Fictions* of the early 1800s to the *Realism* and *Naturalism* of the later half of the century, the novels written in this era mirror the complex and fascinating socio-political and cultural changes that took place.

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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

JCDL '24, December 16–20, 2024, Hong Kong, China

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1093-3/24/12

<https://doi.org/10.1145/3677389.3702593>

Influenced by several historical contexts, literature movements and the changing tastes of readers, the language and vocabulary used in the literature also evolved over this period.

In the field of information retrieval (IR), understanding how terms evolved over time is critical when developing search systems that can effectively retrieve relevant and useful information from historical collections. For traditional search systems, based on term overlap, one of the most challenging issues faced by researchers is the *vocabulary mismatch*. This arises when the terms used in a search session differs from those terms present in the relevant documents, which leads to poor retrieval performance. This issue can occur because different individuals may use various words to convey the same or similar concept, leading to low term overlap during retrieval. Query expansion, a widely-used technique in IR, aims to mitigate this problem by adding semantically similar terms to the search query [1]. This is an active research area in IR, where various models have been proposed to address the vocabulary mismatch problem [2, 11, 14, 19]. A substantial body of research has investigated the impact of temporal signals on search [4–7, 15]. From these studies, it is well established that the recency of events, in general, improve modelling of the temporal aspects of queries and documents can enhance retrieval effectiveness. However, the problem remains a concern when applied to text from different time periods. This is due to the ever-evolving landscape of literature, where the usage of words changes significantly over time. This means that a term relevant in one decade may not carry the same importance in another, and may eventually become extinct.

Previously, a range of approaches have been applied to examine trends in word usage in large text corpora. Simple methods include word frequency analysis [3] and examining the appearance of Ngrams over time (e.g. the Google Books Ngram Viewer). However, such tools need to be treated carefully, depending on the composition of the corpus being studied and due to the shifts in word meanings over time [16]. To identify such shifts as a means of investigating cultural phenomena, authors have examined distributional similarity models [8]. Considerable work has involved the generation of diachronic word embeddings [9, 12]. Given that word embeddings can capture word meaning in a static corpus, training embeddings on different time periods can allow meaning representations at different time periods to be compared, supporting the identification of changing contexts around particular words of interest. However, this requires sufficient data from each time period to construct a robust embedding.

This study aims to explore trends in over 10,000 19th century novels from the British Library Digital Collection (BL19)¹. By splitting the collection of novels into decades, we investigate how the expansion terms selected through the Relevance-based Language Model (RLM) [10, 13] for the same query differ from one decade to another. Specifically, we analyse how the selected expansion terms and their associated weights vary when RLM is applied to novels published in different decades, as well as when applied to the entire collection of novels. Our hypothesis is that the expansion terms with the highest weights will exhibit significant variations across different decades, reflecting the evolving nature of the language. Through this study, we uncover nuanced differences in the semantic characteristics of the language by comparing the expansion terms generated from each decade with the same produced from the entire collection taken together. The primary contributions of this work are as follows:

- We conduct an in-depth analysis of how specific concepts evolved in the BL19 fiction collection, based on semantically similar terms identified by the pseudo-relevance feedback. Specifically, we measure the correlation between the expansion terms estimated via the relevance feedback model for a given decade in the 19th century and those from the full collection. This allows us to gain an insight into how the terms used in novels evolved over time.
- We demonstrate that applying relevance feedback on a given decade from the 19th century exhibits different relevance behaviour in comparison to feedback from the entire time span of the collection.

2 Quantifying the evolution of concepts

In this paper, we aim to capture how the relevance and meaning of key concepts have changed over different decades by analyzing term overlaps, rank correlations, and semantic shifts for understanding the temporal dynamics of language and themes in literary texts. The comparison is structured around several key analyses that collectively offer insights into understanding how the usage of words has evolved resulting in varying term selection, retrieval outcomes, and semantic meanings across different periods in the collection. First, the analysis examines the overlap of the most significant expansion terms selected by RLM from different sub-collections. We can identify the degree of consistency in term selection across periods by comparing the top terms generated for each decade. Additionally, the overlap in terms selected across decades is quantified by calculating the Jaccard similarity between feedback queries generated from different sub-collections. This serves to highlight which terms maintain their relevance over multiple decades and which ones are unique to specific periods. Another approach is to measure the similarity between ranked lists produced by queries expanded from the entire collection and those expanded from individual sub-collections (decades). Using Kendall's τ to quantify rank correlation, this comparison provides insights into how closely the retrieval outcomes from sub-collections align with those from the full collection. The differences in term weight distributions across decades are further explored using Jensen-Shannon (JS) divergence. Through these analyses, we aim to provide a comprehensive view of the temporal evolution of language and themes in literary texts,

¹See <https://labs.biblios.tech/>

Table 1: Number of novels in each decade in BL19 collection.

Decades	1830s	1840s	1850s	1860s	1870s	1880s	1890s
#Novels	20	144	746	1139	1750	2034	4377

Table 2: Query terms of three different categories that are used to study the temporal changes in the BL19 collection.

Thematic	immigrant, emigrant, foreign, newcomer, alien, enslaved, colony, vampire
Plot	engagement, proposal, wedding, suitor, lover, betrothal, eligible, consent, love, mesalliance, heiress, eviction
Genre	crime, murder, mystery, villain, adventure

highlighting some valuable insights into how the concepts have evolved, across different historical decades in the collection.

3 Experimental setup

Collection. In our experiments, we focus on the set of 10,210 English-language fiction texts from the BL19 collection. Their publication dates span from the 1830s to the end of the 19th century, representing a key period both historically and in terms of the development of literary fiction. These texts include well-known novels by prominent authors such as Charles Dickens and Jane Austen and a significant number of little-known and infrequently studied works. As such, this potentially provides a mirror on attitudes and views towards important societal issues during this time period, such as gender, migration, and public health. Furthermore, it presents an opportunity to understand how word usage and meanings have changed around these topics during the course of the 19th century. Studying semantic change and word usage in this way allows humanities researchers to understand the evolution of language and how meanings of words have shifted over time, reflecting cultural and societal changes. These texts were grouped into seven decades, spanning from the 1830s to the 1890s. The distribution of texts across these decades is uneven, with a notable increase in the late 19th century (see Table 1). This pattern mirrors both the composition of the BL19 collection and the broader rise in British fiction publication during that period. Henceforth, we will refer to the entire collection by C . Further, the sub-collections consisting of novels published in the 1830s, 1840s, and subsequent decades will be denoted by C^{30} , C^{40} , and so on. More generally, let C^i denote any sub-collection corresponding to a specific decade i . For our experiments, we split each sub-collection C^i in the unit of paragraphs prior to retrieval.

Topics. We consider a diverse list of 25 query keywords from three distinct categories (thematic, plot, and genre), as identified by experts of 19th century British and Irish literature (see Table 2). Many of these keywords broadly reflect key areas of societal change during the 19th century, including gender and migration. Thematic terms associated with 19th century migration could be expected to occur more frequently as both inward and outward migration increased exponentially from Victorian Britain². The frequency

²In British history, the Victorian era is the period between approximately 1820 and 1914, corresponding roughly but not exactly to the period of Queen Victoria's reign.

Table 3: Top 15 feedback terms chosen by RLM for two sample queries ('immigrant' and 'murder') from Table 2. Feedback terms are obtained from different decades and also from the entire 19th century fictions (see the last row of each group). Term overlaps between two decades are in bold and the decade in which the corresponding query term does not occur is grayed out.

Query	Time Period	Feedback terms
Immigrant	1830-40	
	1841-50	puident, avez, censu, raison, soulless, asop, axl, aptli, diadem, altitud, radic, fry, shred, wight, cape
	1851-60	trade, countri , emigr, soil, saxon , great , flourish, melbourn, cite, irish, gener, furrugn, nativ , anccstii, intrus
	1861-70	britain, land , foreign, govern , race, distdleri, saxon , countri , law, influx, great , thed, tax, compani, natur
	1871-80	coloni , govern , land , countri , peopl, home , durban, australia , popul, state, settler, dimsdal, good , number, nativ
	1881-90	land , countri , queensland, room, nativ , work, men, australia , home , great , year , place, chines, trial, time
	1891-99	alien , favish , land , time , young, turn, yenta, chananya , year , man, mendel, dai, america, good , found
	1831-99	alien , coloni , favish , countri , land , turn , emigr, chananya , year , found , time , good , make, australia , thing
	Murder	1831-40
1841-50		man , commit , bodi , cri , crime , guilti , found , kill , evid , case , accus , person, heard , head , deed
1851-60		commit , man , deed , crime , cri , accus , death , word , blood , answer, exclaim , case , bodi , head , evid
1861-70		man , commit , crime , bodi , dead , heard , cri , word , mr , blood , men, horror , prove, found , hear
1871-80		commit , man , blood , death , bodi , heard , finn, crime , found , night, mr , bonteen, guilti , evid , phinea
1881-90		commit , man , crime , victim, polic , kill , found , case , mysteri, arrest, hand , mr , dai, blood , time
1891-99		man , commit , crime , dead , polic , kill , horror , face, bodi , peopl, word , blood , found , hand , cri
1831-99		commit , man , horror , kill , dead , crime , cri , victim , bodi , blood , hand , death , call, life, hear

of terms relating to it offers insight into the relationship between the fiction in BL19 and historical developments. However, the fact that the lexical search indicates a strong association between the terms 'immigrant' and 'vampire' would indicate that this topic is not treated with sober realism in this body of fiction and that fear and suspicion dominate. 'Engagements' and 'proposals' are key elements in both the domestic, realist fiction which came to dominate fiction from the early to the mid nineteenth century. The relative frequency of terms associated with this marriage plot as the century progressed offers an insight into the extent to which this form of fiction dominates the corpus. It also indicates the extent to which plots and issues primarily concerned with the life of young women are a central focus of that fiction over the period of analysis.

In contrast, terms associated with 'mystery' and 'crime' in BL19 indicate the extent to which the relatively new genre of crime fiction is represented. Modern accounts identify the publication of short stories of Edgar Allan Poe in the 1840s and the novels of Arthur Conan Doyle in the 1880s as key points in the development of the genre. The frequency of terms characteristic of this genre offers an opportunity to test these dates against this corpus, though the fact that the lexical search turned up adventure as a closely associated term indicates that genre boundaries are not rigid in this period.

4 Results and analysis

We explore whether the demarcation of the collection by decade leads to different outcomes by comparing the two expanded queries generated from the individual sub-collections based on these time periods. Specifically, we form expanded queries EQ_C and EQ_C^i using RLM respectively from C and C^i and perform a retrieval on

the entire collection C . We apply BM25 [17, 18] both for the initial and second stage retrieval. Specifically, we choose top-scored 100 expansion terms that the RLM estimates from the initial set of 100 documents retrieved by BM25. In Table 3, we present the top 15 terms with the highest weights as determined by RLM for the queries 'immigrant' and 'murder', generated from both the individual sub-collections based on decades and the entire fiction collection. The terms that appear consistently across different decades are highlighted in bold. From Table 3, we observe the terms contextually similar to 'immigrant' exhibit less overlap across decades in comparison to those related to 'murder'. This suggests that the semantic context of the term 'immigrant' may have varied more significantly over time, while the context for 'murder' remained more stable. One interesting observation is the presence of 'Australia' among the top terms. Note that, the late 18th and early 19th centuries marked a pivotal turning point in Australian history, as this period witnessed the establishment of British colonial rule in Australia. The arrival of the First Fleet in 1787-1788 ushered in a new era, transforming the landscape, society, and culture of this continent through the introduction of European systems, values, and languages³. From the results, we can see that this historical backdrop has also been reflected in British novels of the time. The themes, vocabulary, and narratives found in these literary works echo the broader changes occurring within the British Empire, including its colonial expansion and the resulting cultural exchanges. This suggests a literary consciousness in British novels that parallels the historical events of the period, capturing the impact of colonialism on both the colonized and the colonizers. Another interesting term is 'alien', rooted in the Latin word *alienus* meaning

³https://en.wikipedia.org/wiki/Immigration_history_of_Australia and <https://www.homeaffairs.gov.au/about-us-subsite/files/immigration-history.pdf>.

Table 4: Comparisons of rank correlation values (measured with Kendall’s τ) between two re-ranked lists retrieved by RLM augmented queries, where one query is obtained from the entire century collection and the other is estimated from an individual decade (see the leftmost column). For each query, the correlations are to be compared between decades, i.e. along the columns, and the highest and lowest correlations are respectively bold-faced and underlined. Grey cells indicate the absence of the query term in the corresponding decade’s collection.

	Immigrant	Emigrant	Foreign	Newcomer	Alien	Enslaved	Colony	Vampire	Engagement	Proposal	Wedding	Suitor	Lower	Retribution	Eligible	Consent	Love	Misalliance	Heiress	Crime	Murder	Mystery	Willain	Adventure	Evilution
1830s		0.069	0.010	-0.115	-0.041	-0.049	-0.021		0.013	-0.004	<u>-0.031</u>	0.011	0.007	-0.038	-0.004	0.042	0.019		-0.009	0.122	0.103	-0.005	<u>-0.022</u>	-0.031	
1840s	<u>-0.024</u>	-0.059	0.076	-0.017	0.008	-0.047	0.008	0.027	-0.017	0.032	0.009	0.043	0.012	-0.020	0.031	-0.013	-0.007	0.024	0.021	0.042	0.018	<u>-0.022</u>	0.019	0.039	-0.021
1850s	-0.021	-0.082	-0.076	<u>-0.033</u>	-0.139	0.031	0.007	0.015	-0.029	0.021	0.019	-0.009	0.073	-0.037	<u>-0.028</u>	-0.011	0.028	0.041	0.028	<u>0.009</u>	0.037	0.028	0.049	<u>-0.052</u>	0.019
1860s	0.059	-0.002	0.041	-0.095	0.113	0.076	-0.013	0.023	0.034	<u>-0.021</u>	0.045	0.038	-0.005	0.014	0.056	0.011	0.078	<u>-0.023</u>	<u>-0.075</u>	0.089	0.023	0.063	-0.008	<u>-0.026</u>	0.042
1870s	-0.002	0.091	-0.078	-0.003	<u>-0.084</u>	<u>-0.061</u>	<u>-0.043</u>	0.029	0.052	0.090	-0.007	-0.024	0.038	0.051	0.032	<u>-0.023</u>	-0.006	0.013	-0.054	0.040	0.022	0.018	0.021	0.009	0.065
1880s	-0.017	0.130	0.043	0.059	0.028	0.065	0.028	-0.021	-0.030	0.041	0.033	0.013	<u>-0.014</u>	<u>-0.047</u>	0.029	0.046	0.038	-0.009	0.011	0.015	0.039	0.048	0.081	-0.039	0.041
1890s	0.025	-0.013	0.031	0.107	0.087	0.039	0.081	<u>-0.023</u>	<u>-0.041</u>	0.007	0.056	<u>-0.042</u>	0.072	-0.029	0.081	-0.021	<u>-0.049</u>	0.032	0.051	0.027	0.009	0.054	0.021	-0.042	<u>-0.037</u>

“belonging to another”. Although the *Aliens Act* was imposed by the British Parliament in 1793, the use of the term associated with immigrants in novels appeared in the late 1800s.

The impact of dividing the collection based on publishing decade is further investigated by the correlation between the ranked lists produced by the expanded queries formed from the individual decades. We compute the rank correlation between the two ranked lists produced by EQ_C and EQ_C^i in different settings. In Table 4, we report the rank correlation (Kendall’s τ) between the ranked lists with the two expanded queries. The results indicate that the correlation between the ranked lists is generally very low, with the highest observed correlation being a trivial 0.130 for the query *Emigrant*, suggesting that the expanded queries derived from different sub-collections lead to significantly different retrieval results.

To further quantify the similarity in term selection by RLM, we compute the Jaccard similarity between the terms selected by RLM for the considered collections. Specifically, we evaluate the Jaccard similarity between any pair of sub-collections C_i and the entire collection C , as well as between pairs of sub-collections themselves. For each pair of decades, C_i and C_j , we calculate the average Jaccard similarity based on the expanded queries generated by RLM for all 25 queries. To quantify the variability in term overlap across different topic pairs, the standard deviation of these Jaccard similarities is also reported. The average Jaccard similarity between the expanded query term sets for all 25 topics is presented in the upper triangle of Table 5. Each cell (i, j) within this triangle displays the mean Jaccard similarity between collections C^i and C^j . To provide a measure of dispersion, the standard deviation of these Jaccard similarities is included in parentheses within each cell.

The JS divergence is also computed and presented in the lower half of Table 5 to complement the Jaccard similarity analysis. While the Jaccard measure focuses solely on the overlap of terms in two expanded query sets, disregarding term weights, JS divergence takes these weights into account and calculates the divergence between a pair of expanded queries from two different sub-collections. Interestingly, despite variations in specific term compositions (indicated by lower Jaccard values), the average JS divergence across all decade pairs remains relatively constant around 0.5 with lowest divergence value of 0.48. A relatively lower standard deviation across all queries indicates minimal variation among the values, suggesting both the Jaccard similarity and JS divergence for all queries within a decade exhibit consistent patterns. This suggests a progressive divergence in the semantic space of query expansions,

Table 5: Average variations of feedback term distributions obtained by RLM comparing decades of 19th century. Upper triangle (green cells) shows Jaccard similarities of feedback queries generated from the corresponding row and column decades, whereas lower part (red cells) depicts JS divergence between two feedback term distributions. Each cell reports average variation and standard deviation (in brackets) of 25 keyword queries as in Table 2. Highest Jaccard similarity and lowest JS divergence are bold-faced.

	1830s	1840s	1850s	1860s	1870s	1880s	1890s	1831-99
1830s	1	0.1756 (0.1392)	0.1718 (0.1430)	0.1517 (0.1245)	0.1541 (0.1275)	0.1539 (0.1193)	0.1322 (0.1160)	0.1405 (0.1220)
1840s	0.5249 (0.2084)	1	0.2618 (0.1422)	0.2541 (0.1314)	0.2465 (0.1295)	0.2383 (0.1241)	0.2245 (0.1361)	0.2257 (0.1377)
1850s	0.5190 (0.2140)	0.5188 (0.1954)	1	0.3214 (0.1108)	0.2967 (0.1181)	0.2801 (0.1047)	0.2717 (0.1256)	0.3166 (0.1307)
1860s	0.5341 (0.1575)	0.5616 (0.1434)	0.5469 (0.1770)	1	0.3304 (0.1117)	0.3166 (0.1173)	0.2946 (0.1214)	0.3286 (0.1248)
1870s	0.5633 (0.1880)	0.5403 (0.1776)	0.5274 (0.1963)	0.5493 (0.1766)	1	0.3450 (0.0928)	0.3064 (0.1198)	0.3543 (0.1217)
1880s	0.5912 (0.2082)	0.5443 (0.1726)	0.5899 (0.1064)	0.5267 (0.1971)	0.5283 (0.1975)	1	0.3354 (0.1126)	0.3646 (0.1190)
1890s	0.5234 (0.1594)	0.4817 (0.2233)	0.5891 (0.1031)	0.5892 (0.1072)	0.5918 (0.1004)	0.5732 (0.1393)	1 (0)	0.4491 (0.1484)
1831-99	0.5483 (0.1593)	0.5835 (0.1078)	0.5055 (0.2121)	0.5912 (0.1005)	0.4861 (0.2309)	0.6133 (0.0085)	0.5505 (0.1704)	1 (0)

highlighting the dynamic nature of language and conceptual frameworks over time. This also indicates a consistent level of semantic separation between the different time periods which varies rarely.

5 Conclusion and future work

In this paper, we presented our findings on the study of evolving language usage in 19th century fiction via an information retrieval approach. We investigated how the use of words has evolved over time by applying a Relevance-based Language Model, an effective query expansion technique that leverages the statistical occurrences of words. Our approach compares the terms selected in one decade with those from another, focusing on the overlap of expansion terms. Additionally, we analysed the differences in term importance by computing the divergence in their weights.

As part of future work, we plan to conduct a comparative study of word usage in 19th century novels and non-fiction texts from the same period, as provided by the British Library Digital Collection. Additionally, we aim to develop an interactive tool that enables humanities researchers to visualise and explore word usage trends in the context of the original documents.

Acknowledgement. This publication is part of a project that has received funding from (i) the European Research Council (ERC) under the Horizon 2020 research and innovation program (Grant agreement No. 884951); (ii) Science Foundation Ireland (SFI) to the Insight Centre for Data Analytics under grant No 12/RC/2289_P2.

References

- [1] Hiteshwar Kumar Azad and Akshay Deepak. 2019. Query expansion techniques for information retrieval: A survey. *Information Processing & Management* 56, 5 (2019), 1698–1735.
- [2] Elias Bassani, Nicola Tonello, and Gabriella Pasi. 2023. Personalized Query Expansion with Contextual Word Embeddings. *ACM Trans. Inf. Syst.* 42, 2, Article 61 (dec 2023), 35 pages.
- [3] Joan Bybee. 2006. *Frequency of use and the organization of language*. Oxford University Press.
- [4] Olga Craveiro, Joaquim Macedo, and Henrique Madeira. 2014. Query Expansion with Temporal Segmented Texts. In *Proceedings of the 36th European Conference on IR Research on Advances in Information Retrieval - Volume 8416* (Amsterdam, The Netherlands) (*ECIR 2014*). Springer-Verlag, Berlin, Heidelberg, 612–617.
- [5] Wisam Dakka, Luis Gravano, and Panagiotis G. Ipeirotis. 2008. Answering general time sensitive queries. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management* (Napa Valley, California, USA) (*CIKM '08*). Association for Computing Machinery, New York, NY, USA, 1437–1438.
- [6] Miles Efron and Gene Golovchinsky. 2011. Estimation methods for ranking recent information. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Beijing, China) (*SIGIR '11*). Association for Computing Machinery, New York, NY, USA, 495–504.
- [7] Miles Efron, Jimmy Lin, Jiyin He, and Arjen de Vries. 2014. Temporal feedback for tweet search with non-parametric density estimation. In *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval* (Gold Coast, Queensland, Australia) (*SIGIR '14*). Association for Computing Machinery, New York, NY, USA, 33–42.
- [8] Kristina Gulordava and Marco Baroni. 2011. A distributional similarity approach to the detection of semantic change in the Google Books Ngram corpus. In *Proceedings of the GEMS 2011 workshop on geometrical models of natural language semantics*. 67–71.
- [9] William L. Hamilton, Jure Leskovec, and Dan Jurafsky. 2016. Diachronic Word Embeddings Reveal Statistical Laws of Semantic Change. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*. The Association for Computing Machinery, New York, NY, USA, 515–524.
- [10] Nasreen AbdulJaleel, James Allan, W. Bruce Croft, Fernando Diaz, Leah S. Larkey, Xiaoyan Li, Mark D. Smucker, and Courtney Wade. 2004. UMass at TREC 2004: Novelty and HARD. In *Proceedings of the Thirteenth Text REtrieval Conference, TREC 2004, Gaithersburg, Maryland, USA, November 16-19, 2004 (NIST Special Publication, Vol. 500-261)*. National Institute of Standards and Technology (NIST).
- [11] Pau Perng-Hwa Kung, Zihao Fan, Tong Zhao, Yozen Liu, Zhixin Lai, Jiahui Shi, Yan Wu, Jun Yu, Neil Shah, and Ganesh Venkataraman. 2024. Improving Embedding-Based Retrieval in Friend Recommendation with ANN Query Expansion. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Washington DC, USA) (*SIGIR '24*). Association for Computing Machinery, New York, NY, USA, 2930–2934.
- [12] Andrey Kutuzov, Lilja Øvrelid, Terrence Szymanski, and Erik Velldal. 2018. Diachronic word embeddings and semantic shifts: A survey. *arXiv preprint arXiv:1806.03537* (2018).
- [13] Victor Lavrenko and W. Bruce Croft. 2001. Relevance based language models. In *Proceedings of the 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (New Orleans, Louisiana, USA) (*SIGIR '01*). Association for Computing Machinery, New York, NY, USA, 120–127.
- [14] Minghan Li, Honglei Zhuang, Kai Hui, Zhen Qin, Jimmy Lin, Rolf Jagerman, Xuanhui Wang, and Michael Bendersky. 2024. Can Query Expansion Improve Generalization of Strong Cross-Encoder Rankers?. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Washington DC, USA) (*SIGIR '24*). Association for Computing Machinery, New York, NY, USA, 2321–2326.
- [15] Xiaoyan Li and W. Bruce Croft. 2003. Time-based language models. In *Proceedings of the Twelfth International Conference on Information and Knowledge Management* (New Orleans, LA, USA) (*CIKM '03*). Association for Computing Machinery, New York, NY, USA, 469–475.
- [16] Michael Pettit. 2016. Historical time in the age of big data: Cultural psychology, historical change, and the Google Books Ngram Viewer. *History of Psychology* 19, 2 (2016), 141.
- [17] Stephen Robertson and Hugo Zaragoza. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. *Foundations and Trends in Information Retrieval* 3, 4 (apr 2009), 333–389.
- [18] S. E. Robertson and S. Walker. 1994. Some Simple Effective Approximations to the 2-Poisson Model for Probabilistic Weighted Retrieval. In *SIGIR '94*. Springer London, London, 232–241.
- [19] Le Zhao and Jamie Callan. 2012. Automatic term mismatch diagnosis for selective query expansion. 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, 515–524.