

# Towards Proactive Information Retrieval in Noisy Text with Wikipedia Concepts

Tabish Ahmed, Sahan Bulathwela

Centre for Artificial Intelligence, University College London, UK.

## Abstract

Extracting useful information from the user history to clearly understand informational needs is a crucial feature of a proactive information retrieval system. Regarding understanding information and relevance, Wikipedia can provide the background knowledge that an intelligent system needs. This work explores how exploiting the context of a query using Wikipedia concepts can improve proactive information retrieval on noisy text. We formulate two models that use entity linking to associate Wikipedia topics with the relevance model. Our experiments around a podcast segment retrieval task demonstrate that there is a clear signal of relevance in Wikipedia concepts while a ranking model can improve precision by incorporating them. We also find Wikifying the background context of a query can help disambiguate the meaning of the query, further helping proactive information retrieval.

## 1. Introduction

The informational needs of people are highly contextual and can depend on many different factors such as their current knowledge state, interests and goals [1, 2, 3]. However, an effective information retrieval companion should minimise the human effort required in i) expressing a human information need and ii) navigating a lengthy result set. Using topical representations of the user history (e.g. [4]) can immensely help formulating zero shot queries and refining short user queries that enable proactive information retrieval (IR). While the world has digital textual information in abundance, it can often be noisy (e.g. extracted through Automatic Speech Recognition (ASR), PDF text extraction etc.), leading to state-of-the-art neural models being highly sensitive to the noise producing sub-optimal results [5]. This demands denoising steps to refine both query and document representation.

In this paper, we argue that Wikipedia, an openly available encyclopedia, can be a humanly intuitive knowledge base [6] that has the potential to provide the world view many *noisy* information Retrieval systems need. In the midst of noisy text, we point out that it can help refine user queries (both explicit and implicit) while compiling large result sets into structured narratives that human users can navigate with optimal effort. We use a dataset with noisy text where the query description is used as a proxy of the user's context used in a proactive information retrieval system. We demonstrate that positive results can be obtained in comparison to relevant baselines when ranking documents.

---

PASIR'22: First Workshop on Proactive and Agent-Supported Information Retrieval at CIKM 2022, October 21, 2022, Atlanta, GA

✉ tabish.ahmed.21@ucl.ac.uk (T. Ahmed)



© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

## 2. Related Work

Retrieving relevant documents to a personal information need involves mining textual features that can be used to match them to user queries. While previous works have explored Wikipedia-based concept modelling for query expansion [7, 8] exclusively, our work deviates from these works as we focus on i) simultaneously adding Wiki features to the documents, ii) ranking documents with noisy texts and iii) using the query description (proxy for context extracted from user history) for query disambiguation. Latest neural ranking models dominate the state-of-the-art in IR albeit neural models struggle in the presence of noise in text [5] making much simpler, computationally efficient IR models such as BM25, DPH and PL2 much more suitable for noisy data [9]. Prior evidence has demonstrated the success of simpler models such as BM25 with the exception of failure on long documents [10]. However, more sophisticated probabilistic models such as DPH [11], motivated by the Divergence from Randomness (DFR) framework based on a hypergeometric distribution and Popper normalisation instead of Laplace normalization, have outperformed the BM25 model when it comes to noisy text [9].

This emerging need for retrieving from noisy-text has led to recent competitions such as the podcast segment retrieval challenge [12] that has paved way to improving the state-of-the-art. This competition is relevant to our work as we also tackle proactive IR in noisy text. While the best performing model within the competition uses a neural approach, it uses token embeddings rather than document embeddings [13]. The leading model uses a re-ranking model trained on an orthogonal dataset adding to computational costs significantly. The next best, Dublin City University model (DCU) uses a variety of features ranging from text, WordNet based synonyms and entities for expanding the query [14]. As the DCU model uses automatically extracted entities (via Named Entity Recognition) to expand the query, we consider this work as the most relevant model to our proposal. But, we also annotate the segments with Wikipedia concepts/entities further differentiating our work from theirs. User queries, inherently are short, averaging two to three key words in general [15]. Exploiting auxiliary information available to recover the context of the search is essential in such scenarios. In recent competitions such as the podcast segment retrieval challenge, the queries used were accompanied by a small description that provided the contextual information about the query itself. It was evidenced that the leading models in this challenge had to use the description to extract features and sometimes use the identified keywords to make queries to external search engines to further expand the queries [9]. However, in a real world system, such information is not available. Hence a proactive IR system has to rely on the user interactions from the past to build the context [16]. In this work, we use the query context as a proxy representation extracted by a proactive IR system.

*Proactive information retrieval* is a field in IR that gained a lot of attention recently. Supporting user information queries with minimal/zero effort will enable users to use IR systems more effectively. However, a key part of facilitating information retrieval tasks proactively is building systems that can capture the relevant signal from the user that indicate the informational needs of them. Prior works in this domain have identified several ways to harvest user intent. The simplest approach in this direction is to explicitly ask questions [17] or use the demographic information of the user [18] to profile them. Rather than explicitly disrupting the user experience, many works harvest implicit user actions such as clicks, impressions and previous searches

[19, 18]. Especially in the social media domain, the interactions of users are used to create concept-based models [20] that can be used to infer user preferences and create zero effort queries for new content. The extracted preference features are usually used to expand information queries [8, 7], synthesise zero-effort queries [20] or personalise search results [21]. In this work, we argue in support of using concept-based user models as a form of zero-effort query for informational recommendation and learning scenarios, especially in the presence of noisy document representations.

## 2.1. Concept-based User Modelling and Wikipedia

In the field of document retrieval and personalisation, keyword extraction is a heavily used form of identifying *concepts* from textual documents. When working with short text such as social media posts, the common method is to use the words in the posts as words and build a user profile over these words [22]. While this approach provides fine grained features, the feature space can be too vast and pose challenges in increasing recall. To address this issue, other systems use topic detection using unsupervised learning approaches such as LDA [23]. Such unsupervised techniques are complex to tune and are not guaranteed to give humanly intuitive topics. While expert annotation (e.g. in education [24]) has been one of the approaches to obtain humanly intuitive, representative topics from textual content, this approach is not scalable. **Wikification**, a more recent approach, looks promising towards automatically extracting explainable concepts from text. Wikification identifies Wikipedia concepts present in a document by connecting natural text to Wikipedia articles via entity linking [25]. In this work, we argue for the usefulness of Wikification in creating precise, humanly intuitive concept representations at scale. Our experiments also demonstrate how Wikipedia concepts can be used to i) disambiguate query keyword and ii) at the absence of them, create entity features that can find relevant documents for users proactively.

## 3. Wikipedia to Support Proactive Search-based Systems

Through this work, we aim to explore how the knowledge contained in Wikipedia can be utilised to better a proactive search-based system. In order to tread towards this direction, we aim to answer multiple research questions in this work as a promising step. We hypothesise that associating Wikipedia concepts to a natural language based IR system can improve a proactive IR system by i) bettering disambiguation of the meaning of user queries, ii) by exploiting the semantic meaning of document words while being less sensitive to noisy words in the document transcripts and iii) by creating opportunities for the system to understand the learner interactions more meaningfully.

**RQ1** Do Wikipedia concepts carry a signal that indicates relevance of documents to queries?

**RQ2** Can Wikipedia annotations improve noisy information retrieval?

Through this work, we conduct a few preliminary experiments to validate these research questions. We choose the segment retrieval task from the Podcast Track at TREC 2020 [12]

for these experiments as "search" is a key component of proactive information retrieval. We use this task specifically as the query description is provided as part of the query, providing the context of the query. In many proactive IR systems, the query context is mined using the user interaction history. In this case, the description can be treated as a context representation mined from the users' prior interactions.

### 3.1. Information Retrieval Models

We restrict our model choices to a probabilistic retrieval method BM25 and a hypergeometric weighing model DPH. The reasons for restricting to probabilistic models are i) its superior performance in noisy IR [9], ii) the simplicity to introducing Wikipedia concept features, iii) the data and computational efficiency in training models and iv) the practical feasibility to real world systems. The models we utilise in this these experiments use the query  $q$ , query description  $d$  and the segment  $s$ .

#### 3.1.1. Baseline Models

We use BM25 and DPH models as probabilistic baseline models for this work. These two models use the textual tokens of the query  $q$  and the segment  $s$  to compute the relevance score  $rel$  as per equation 1 where  $txt$  identifies the text token features..

$$rel(q, s) = f(q_{txt}, s_{txt}) \quad (1)$$

As described in section 2, we implement and test the DCU run 2 model [14] (referred to as DCU model hereafter) that expands the query with entities extracted from the description as per equation 2 that expands equation 1. The entities are extracted using Named Entity Recognition (NER).

$$rel(q, d, s) = f(q_{txt} + d_{ent}, s_{txt}) \quad (2)$$

where  $ent$  represents the entities extracted from the query description  $d$ .

### 3.2. Proposed Models

The proposed models in this work uses Wikipedia concepts to enrich the representations used in the IR model. We hypothesise that adding Wikipedia concepts can improve precision of the model as exact entities can be matched between the query + context and the segments. We propose two models that use Wikipedia topics both in query and segment expansion.

**Wiki\_rel** , This model expands the query and the document by adding the Wikipedia concepts extracted using both the query  $q$  and the query description  $d$  as per equation 4. This model replaces the entities in the DCU model by Wikipedia concepts.

$$rel(q, d, s) = f(q_{txt} + q_{wiki} + d_{wiki}, s_{txt} + s_{wiki}) \quad (3)$$

where  $wiki$  represents the concepts extracted from the query  $q$ , description  $d$  and segment  $s$ .

**Ent\_Wiki\_rel** , This model expands the DCU model by enriching the query, description and the segment with Wikipedia concepts. This formulation is presented in equation 4.

$$rel(q, d, s) = f(q_{\text{txt}} + d_{\text{ent}} + q_{\text{wiki}} + d_{\text{wiki}}, s_{\text{txt}} + s_{\text{wiki}}) \quad (4)$$

The two proposed models help us validate RQ2 by i) Wiki\_rel model validating if *replacing* named-entities with Wikipedia features improves the ranking models and ii) Wiki\_Ent\_rel model validating if *adding* Wikipedia features improves the model.

### 3.3. Data

As per section 3, we use the Podcast segment retrieval task to demonstrate usefulness of Wikipedia features in proactive information retrieval. The Spotify podcast Dataset contains approx 100,000 transcripts from around 60,000 hours of audio data culminating in the largest corpus of transcribed speech data. The episodes in the dataset were randomly sampled from 105,360 English podcast episodes published between January 1, 2019 to March 1, 2020 on Spotify with 10% of the podcasts from professional creators with high production values and the other 90% coming from amateurs [26]. The data is also provided with training and testing topics (queries) along with human relevance judgements. There are 8 topics in the training dataset and another 50 topics in the test set.

To keep the computational costs low, we run the experiments outlined in section 3.4 exclusively using the training data in the podcast dataset. We first transform the dataset into two minute overlapping segments as described in [9]. Then, we take all the relevant segments as positive documents for the different queries. As there is a substantial number of segments that are irrelevant (as the whole dataset contains 3.5 Mn segments), we down sample this data and get a negative segment set of  $\approx 14000$  segments that are not relevant to any of the labelled queries. This set of positive and negative segments make up the smaller scale dataset that we use in our empirical experiments. We refer to this dataset as Podcast Small dataset. We use the Wikifier [25] to associate Wikipedia concepts to both the queries (query + description) and the podcast segments as the models in section 3.1 require.

### 3.4. Experiments

The empirical experiments we run in this work aim to answer RQ 1 and 2. We use the Podcast Small dataset for these experiments. In a proactive information retrieval setting, the user queries are inferred from users’ historical interactions. We treat the query description provided in the dataset in the place of the features extracted from the user history. We hypothesise that the query description, can be used in two ways, i) as the user context that can be used in a zero effort query to make up the context and rank relevant documents from the corpus and/or ii) as a mean to disambiguate the true meaning of a short query provided by the user without prompting the user to provide further clarifications.

To answer RQ1, we aim to see if there is a significant alignment between the Wikipedia concepts present in the query and a relevant document in contrast to a non-relevant document. To measure this quantitatively, we measure the Jaccard similarity coefficient  $\rho$  between the query and the document using the sets of Wikipedia concepts. Then, for each query, we take

the sets of values  $\rho_{\text{rel}}$  and  $\rho_{\text{non rel}}$  and compare the difference of medians using a one-tailed Mann–Whitney U test ( $H_{\text{alt.}} : \rho_{\text{rel}} > \rho_{\text{non rel}}$ ).

To answer RQ2, we take the annotated Wikipedia concepts of the query and the podcast segments as additional features. Then we develop a new set of information retrieval algorithms outlined in section 3.1 to account for the alignment of Wikipedia concepts present. Then we use the same Podcast Small dataset to evaluate if the ranking of relevant documents is different between the baselines and the proposed models. To validate how Wikifying the context can be used to disambiguate a query, we ran Wikification on each query using i) the query words exclusively and ii) the query words + the description words, to observe if more precise identification of Wikipedia concepts is carried out.

In order to run the evaluation experiments needed for RQ2, we utilise Python Terrier library [27]. The text processing and NER enrichment needed for the DCU model is done using the Spacy library. The Wikipedia annotations are sourced using the Wikifier service [25].

### 3.4.1. Evaluation

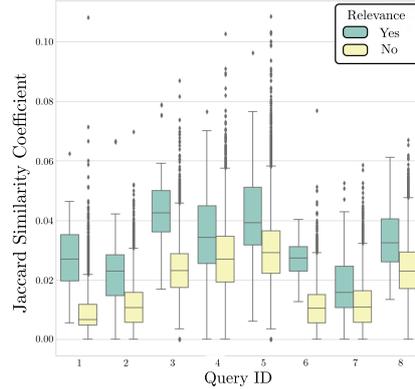
To evaluate if the Jaccard similarity coefficients in relation to RQ1 are significantly different, we use the test statistic of the hypothesis test. We present the performance of the ranking models related to validating RQ2 using NDCG, NDCG@30 and Precision@10 as these are the same evaluation metrics used in prior work [9]. NDCG uses graded relevance instead of binary relevance/non-relevance to focus on highly relevant documents, implying that retrieving a highly relevant document is much more important than retrieving a less important document. Whereas precision measures the number of relevant documents retrieved over the total relevant documents. Similar to the prior work, we calculate NDCG for the entire set of relevant segments available and up-till rank 30. Precision is computed by using a cut-off of top-10 ranked segments.

## 4. Experimental Results

The results of the experiments are reported in this section. Figure 1 shows the difference of medians of Jaccard similarity coefficients, while Table 1 summarises the results obtained from the hypothesis test conducted to verify RQ1. Table 2 outlines the ranking performance obtained in the experiments run in order to verify RQ2.

### 4.1. Query Keyword Disambiguation using the Context

We hypothesise that the context gathered by the user can also be used to disambiguate the meaning of keywords in a search query. To test this, we take 7 queries from the training dataset that belongs to categories *topical* and *known item* (“story about riding a bird” is ignored as it is a query that belongs to the *refinding* category). We enrich these queries using Wikification in two settings, i) query only vs. ii) query + description, to investigate if the identified Wiki concepts are different. Table 3 reports how the salient entities in the query keywords are detected differently in the two settings.



**Figure 1:** Jaccard similarity distribution of relevant vs. non-relevant segments for each query in the podcast dataset.

**Table 1**

The difference of medians test for each query reported with the statistical significance.

Query ID	Number of Documents		Median Jaccard Similarity			MannWhitney U test	
	Relevant	Non-Relevant	Relevant	Non-Relevant	Difference	Test Statistic	p value
1	70	14109	0.027	0.007	0.020	25680	1.03E-49
2	63	14116	0.023	0.011	0.012	47466	8.84E-46
3	78	14101	0.043	0.023	0.019	264688	1.22E-16
4	78	14101	0.034	0.027	0.007	393033	1.41E-06
5	80	14099	0.039	0.029	0.010	455015	1.42E-03
6	37	14142	0.027	0.010	0.017	35748	8.37E-48
7	77	14102	0.016	0.011	0.005	56322	2.80E-44
8	80	14099	0.032	0.023	0.010	271961	6.31E-16

## 5. Discussion

The results obtained in the preliminary experiments conducted are very promising. Majority of the results suggest that Wikipedia concepts can help the IR system.

### 5.1. The Usefulness of Wikipedia Features (RQ1)

Figure 1 clearly shows that the Jaccard similarity coefficient between the relevant segments  $\rho_{rel}$  and the non-relevant documents  $\rho_{non\ rel}$  are different from each other. It is even evident that the confidence intervals don't even overlap in some of the queries in the training data. The test statistics presented in table 1 further confirms this observation. In a one-tailed hypothesis test, all the queries demonstrate statistically significant differences of medians where the median Jaccard similarity of the Wikipedia concepts of the relevant segments and the query is greater than that of the non-relevant segments. These results enable us to conclude that the the overlap of Wikipedia concepts between queries and relevant documents is an influential signal that can be incorporated in relevance computation. The results indicate that there is promise in using

**Table 2**

Ranking performance of the models. The best performance in **bold** and the second best in *italic* face.

Model	Features			Metrics		
	Query Text	NER Entities	Wiki Concepts	NDCG	NDCG at 30	Precision at 10
<i>Baselines</i>						
DPH	×	o	o	0.48	<i>0.30</i>	<i>0.32</i>
BM25	×	o	o	0.48	0.28	0.31
DCU	×	×	o	<b>0.51</b>	<b>0.32</b>	0.30
<i>New Proposals</i>						
Wiki_rel	×	o	×	<i>0.49</i>	0.29	0.31
Ent_Wiki_rel	×	×	×	<b>0.51</b>	<i>0.30</i>	<b>0.36</b>

**Table 3**

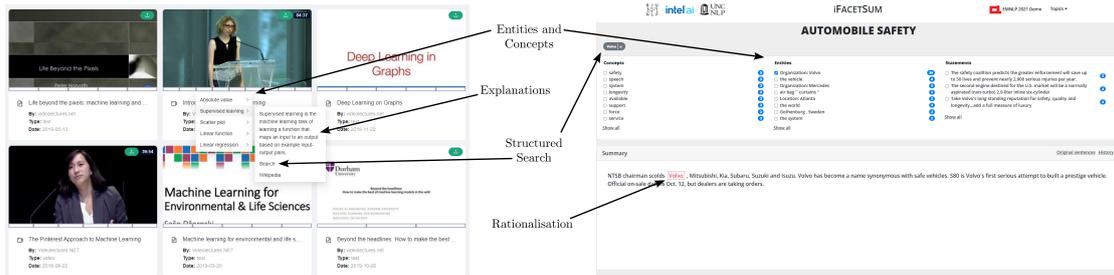
The Quality of Query Word Disambiguation/Refinement when using the query exclusively vs. when using the description.

Query	Query Only	Query + Description
<u>coronavirus</u> spread	wiki/Coronavirus	wiki/Novel_coronavirus
<u>greta thunberg</u> cross atlantic	-	wiki/Greta_Thunberg
<u>black hole</u> image	wiki/Black_hole	wiki/Black_hole
<u>daniel ek</u> interview	-	wiki/Daniel_Ek
<u>michelle obama</u> <u>becoming</u>	-	wiki/Michelle_Obama
	wiki/Becoming_(philosophy)	wiki/Becoming_(book)
<u>anna delvey</u>	wiki/Indian_anna	wiki/Anna_Sorokin
<u>facebook</u> <u>stock</u> prediction	wiki/Facebook	wiki/Facebook
	wiki/Stock	wiki/Stock

Wikipedia concepts as a feature in IR systems where noise is present.

## 5.2. Utility of Wikification in Information Retrieval (RQ2)

The usefulness of Wikipedia concepts in relevance prediction is validated in the results presented in table 2. The results in the table show that replacing the NER-based entities with Wikipedia concepts is not as fruitful as augmenting them while the entities are retained. It is observed that both the proposed models outperform the entity based DCU model in precision at 10. This aligns with our understanding of the effect of Wikipedia concepts. NER, while detecting specific persons, organisations and other entities in the text, only adds the textual representations of these entities to the query to expand it. Therefore, the newly added queries are not capable of entity disambiguation (when the same term occurs in a segment meaning something else. e.g. Apple, the fruit vs. the company) and coreference resolution (when a different textual form is used to mean the same entity. e.g. United Nations vs. UN). On the contrary, enriching both the query and the document with Wikipedia entities enable accounting for both of these effects, making the relevance calculation more precise. The *Ent\_Wiki\_rel* model that contains both the NER-based entities and Wikipedia concepts, manages to improve the precision at 10 while retaining the NDCG performance. This shows that keeping the entities in the query is



**Figure 2:** (i) X5Learn [28] and IFacetSum [29] showing that using Wikipedia concepts to represent documents can support proactive information retrieval by summarising the key concepts/entities in the result collection, helping the user to make more informed decisions on their next steps.

important to retain the NDCG gains found in the DCU model.

In this scenario, we use the Wikipedia concepts and entities extracted from the description to improve the search ranking. While we treat the description in this scenario as a representation extracted from the user history, many concept-based user models described in section 2 support extracting features such as keywords [21] and Wikipedia concepts [1] from content.

**Query Keyword Disambiguation** The results in table 3 also give us insights into the potential of using query context for keyword disambiguation. It is seen that in multiple queries, certain concepts were not even detected when the query was used exclusively for Wikification (e.g. Greta Thunberg, Michelle Obama etc.). However, with the presence of context (that can be extracted from historic encounter of concepts of a user [4]), the algorithm managed to identify these concepts in the query increasing recall. Table 3 also shows instances where the precision of identification is improved in the presence of the context. The more precise *Novel Coronavirus* was detected in the first query while the system understood that the query was about Michelle Obama’s book, *Becoming* in the fifth query, not the philosophical concept. These results show how many scenarios where short query keywords can be clarified automatically by referring to the context rather than needlessly prompting the human user.

### 5.3. Wikipedia Concepts and the User Experience

The previous section discussed one opportunity of Wikification, automatic query word disambiguation. Such features will substantially decrease the number of user prompts in a system leading to a smoother experience. Wikipedia concepts also bring the advantage of being grounded in a taxonomy that is intuitive to humans. The human intuitive nature of Wikipedia concepts allows the system to provide more transparent result sets that the user can navigate more efficiently. Figure 2 shows two recent systems demonstrating the potential of using Wikipedia based representations to improve human-in-the-loop information retrieval.

The left system, X5Learn [28] shows how different Wikipedia concepts are related to different segments of videos. The Jaccard-based relevance can be used to visualise the degree of concept overlap between the segments and the query indicating relevance (connecting to the hypothesis of RQ1). Also, using Wikipedia concepts rather than free-form keywords or NER-based entities

allows the system to provide explanations of the concepts users encounter. This can be done using the knowledge base behind the taxonomy (in this case, Wikipedia as per figure 2 (left)) Many users seeking information do not fully understand the knowledge landscape they navigate and can use learning opportunities to get more familiar with the topics.

IFacetSum, a multi-faceted mix-initiative interface shown in figure 2(right) is an example of how the result sets can be summarised using the entities and concepts in the document. As shown in the demonstration, the detected entities and concepts allow the user to understand the contents of the relevant documents quickly. Their user study also found that the users didn't utilise the statements and summaries as much as they used the concepts and entities [29]. However, this system lacks the capabilities of explaining the topics as it is not backed by a knowledge base pointing to the advantages of using a taxonomy that has auxiliary information that can be used. In a nutshell, many recent visualisation approaches use concepts to represent documents. Having Wikipedia as a topic taxonomy enables these systems to make use of the advancements of Wikification that is advancing rapidly at present. It also allows taking use of other information in Wikipedia (e.g. semantic relatedness between topics, hierarchy of concepts etc.) to model the concepts in a much more sophisticated manner. This in turn, helps bridging the system seamlessly to the end user, creating cleaner interaction footprints that transform into the input for proactive search.

## 6. Conclusion

Through this work, we explore the potential of using Wikipedia concepts to drive the proactive search experience of an IR system working on noisy documents. We explored the usefulness of Wikipedia concepts focusing on 2 research questions. We treated the query description in a podcast segment retrieval task dataset as the user context in a proactive IR system and ran a series of experiments. Our first experiment showed that the overlap of Wikipedia concepts between the relevant segments and the query are significantly larger than with the non-relevant segments. A subsequent experiment that added Wikipedia concept features to the search query and the document showed that the relevance model that uses Wikipedia concepts can improve precision of the ranking model while retaining the NDCG score in comparison to a counterpart that uses entities instead of Wikipedia concepts. A deeper analysis on the Wikification of the queries also showed that the meaning of the query keywords can be enhanced by using the query context. We also demonstrated using two recent works on how Wikipedia taxonomy can help an IR system to better connect with the end user and enable them to carry out structured IR tasks by proactively helping them navigate the information in the result sets. It is seen that having Wikipedia concepts allows presenting the result set in a humanly intuitive form that encourages the user to efficiently refine their query in the subsequent steps. All the observations from the analysed systems indicate that Wikipedia concepts have a key role to play in creating proactive information retrieval systems that can facilitate a higher degree human-in-the-loop operation. In the future work, we aim to use more sophisticated representations of Wikipedia concepts and advanced relevance models to run large scale experiments on the full podcast dataset with 3.5 Mn segments. Understanding how to use Wikipedia concepts effectively (e.g. by ranking them) is a top priority.

## Acknowledgments

This research was partially conducted as part of the X5GON project funded from the EU's Horizon 2020 research programme grant No 761758. This work is also supported by the European Commission-funded project "Humane AI: Toward AI Systems That Augment and Empower Humans by Understanding Us, our Society and the World Around Us" (grant 820437), EU Erasmus+ project "European Network for Catalysing Open Resources in Education" (project ref: 621586-EPP-1-2020-1-NO-EPPKA2-KA), and the AT2030 Programme. The AT2030 programme is funded by aid from the UK government and led by the Global Disability Innovation Hub.

## References

- [1] S. Bulathwela, M. Pérez-Ortiz, E. Yilmaz, J. Shawe-Taylor, Towards an integrative educational recommender for lifelong learners, in: AAAI Conference on Artificial Intelligence, AAAI 20, 2020.
- [2] W. Jiang, Z. A. Pardos, Q. Wei, Goal-based course recommendation, in: Proceedings of International Conference on Learning Analytics & Knowledge, 2019.
- [3] S. Bulathwela, M. Pérez-Ortiz, E. Yilmaz, J. Shawe-Taylor, Power to the learner: Towards human-intuitive and integrative recommendations with open educational resources, *Sustainability* 14 (2022) 11682.
- [4] S. Bulathwela, M. Pérez-Ortiz, E. Yilmaz, J. Shawe-Taylor, Truelearn: A family of bayesian algorithms to match lifelong learners to open educational resources, in: AAAI Conference on Artificial Intelligence, AAAI 20, 2020.
- [5] G. Sidiropoulos, S. Vakulenko, E. Kanoulas, On the impact of speech recognition errors in passage retrieval for spoken question answering, *arXiv preprint arXiv:2209.12944* (2022).
- [6] S. Bulathwela, M. Pérez-Ortiz, C. Halloway, J. Shawe-Taylor, Could ai democratise education? socio-technical imaginaries of an edtech revolution, in: In Proc. of the NeurIPS Workshop on Machine Learning for the Developing World (ML4D), 2021.
- [7] H. K. Azad, A. Deepak, A new approach for query expansion using wikipedia and wordnet, *Information sciences* 492 (2019) 147–163.
- [8] J. A. Nasir, I. Varlamis, S. Ishfaq, A knowledge-based semantic framework for query expansion, *Information processing & management* 56 (2019) 1605–1617.
- [9] R. Jones, B. Carterette, A. Clifton, M. Eskevich, G. J. Jones, J. Karlgren, A. Pappu, S. Reddy, Y. Yu, Trec 2020 podcasts track overview, *arXiv preprint arXiv:2103.15953* (2021).
- [10] Y. Lv, C. Zhai, When documents are very long, bm25 fails!, in: Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval, 2011, pp. 1103–1104.
- [11] G. Amati, Frequentist and bayesian approach to information retrieval, in: European Conference on Information Retrieval, Springer, 2006, pp. 13–24.
- [12] Y. Yu, J. Karlgren, H. Bonab, A. Clifton, M. I. Tanveer, R. Jones, Spotify at the trec 2020 podcasts track: Segment retrieval, in: Proceedings of the Twenty-Ninth Text REtrieval Conference (TREC 2020), 2020.

- [13] P. Galuščáková, S. Nair, D. W. Oard, Combine and re-rank: The university of maryland at the trec 2020 podcasts track (2020).
- [14] Y. Moriya, G. J. Jones, Dcu-adapt at the trec 2020 podcasts track., in: TREC, 2020.
- [15] A. Arampatzis, J. Kamps, A study of query length, in: Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval, 2008, pp. 811–812.
- [16] P. Sen, Proactive information retrieval, Ph.D. thesis, Dublin City University, 2021.
- [17] G. H. Torbati, A. Yates, G. Weikum, Personalized entity search by sparse and scrutable user profiles, in: Proceedings of the 2020 Conference on Human Information Interaction and Retrieval, 2020, pp. 427–431.
- [18] L. Yang, Q. Guo, Y. Song, S. Meng, M. Shokouhi, K. McDonald, W. B. Croft, Modeling user interests for zero-query ranking, in: European Conference on Information Retrieval, Springer, 2016, pp. 171–184.
- [19] S. Bulathwela, M. Perez-Ortiz, E. Novak, E. Yilmaz, J. Shawe-Taylor, Peek: A large dataset of learner engagement with educational videos, in: Proc. of RecSys Workshop on Online Recommender Systems and User Modeling (ORSUM'21), 2021. URL: <https://arxiv.org/abs/2109.03154>.
- [20] F. Zarrinkalam, S. Faralli, G. Piao, E. Bagheri, Extracting, mining and predicting users' interests from social media (2020).
- [21] R. Syed, K. Collins-Thompson, Retrieval algorithms optimized for human learning, in: Proc. of Int. Conf. on Research and Development in Information Retrieval (SIGIR), 2017.
- [22] G. Piao, J. G. Breslin, Analyzing mooc entries of professionals on linkedin for user modeling and personalized mooc recommendations, in: Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization, UMAP '16, 2016.
- [23] F. Zarrinkalam, H. Fani, E. Bagheri, M. Kahani, Predicting users' future interests on twitter, in: European Conference on Information Retrieval, Springer, 2017, pp. 464–476.
- [24] A. T. Corbett, J. R. Anderson, Knowledge tracing: Modeling the acquisition of procedural knowledge, *User Modeling and User-Adapted Interaction* 4 (1994).
- [25] J. Brank, G. Leban, M. Grobelnik, Annotating documents with relevant wikipedia concepts, in: Proc. of Slovenian KDD Conf. on Data Mining and Data Warehouses (SiKDD), 2017.
- [26] A. Clifton, A. Pappu, S. Reddy, Y. Yu, J. Karlgren, B. Carterette, R. Jones, The spotify podcast dataset, arXiv preprint arXiv:2004.04270 (2020).
- [27] C. Macdonald, N. Tonello, Declarative experimentation in information retrieval using PyTerrier, in: Proceedings of the 2020 ACM SIGIR on International Conference on Theory of Information Retrieval, ACM, 2020. doi:10.1145/3409256.3409829.
- [28] S. Bulathwela, S. Kreitmayer, M. Pérez-Ortiz, What's in it for me? augmenting recommended learning resources with navigable annotations, in: Proceedings of the 25th International Conference on Intelligent User Interfaces Companion, IUI 20, 2020.
- [29] E. Hirsch, A. Eirew, O. Shapira, A. Caciularu, A. Cattan, O. Ernst, R. Pasunuru, H. Ronen, M. Bansal, I. Dagan, iFacetSum: Coreference-based interactive faceted summarization for multi-document exploration, in: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, Association for Computational Linguistics, Online and Punta Cana, Dominican Republic, 2021, pp. 283–297. URL: <https://aclanthology.org/2021.emnlp-demo.33>. doi:10.18653/v1/2021.emnlp-demo.33.