MBIPOM, B., MASSIE, S. and CRAW, S. 2018. An e-learning recommender that helps learners find the right materials. In Zilberstein, S., McIlraith, S., Weinberger, K., Youngblood, G.M., Myers, K., Eaton, E. and Wollowski, M. (eds.) *Proceedings of the 32nd Association for the Advancement of Artificial Intelligence (AAAI) Artificial intelligence conference (AAAI18), co-located with the 30th Innovative applications of artificial intelligence conference (IAAI18) and the 8th AAAI Educational advances in artificial intelligence (EAAI-18)*, 2-7 February 2018, New Orleans, Louisiana, USA. Palo Alto: AAAI Press [online], pages 7928-7933. Available from: https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16253

An e-learning recommender that helps learners find the right materials.

MBIPOM, B., MASSIE, S., CRAW, S.

2018

An e-Learning Recommender that Helps Learners Find the Right Materials

Blessing Mbipom and **Stewart Massie** and **Susan Craw**
School of Computing Science & Digital Media
Robert Gordon University
Aberdeen, UK

Abstract

Learning materials are increasingly available on the Web making them an excellent source of information for building e-Learning recommendation systems. However, learners often have difficulty finding the right materials to support their learning goals because they lack sufficient domain knowledge to craft effective queries that convey what they wish to learn. The unfamiliar vocabulary often used by domain experts creates a semantic gap between learners and experts, and also makes it difficult to map a learner's query to relevant learning materials. We build an e-Learning recommender system that uses background knowledge extracted from a collection of teaching materials and encyclopedia sources to support the refinement of learners' queries. Our approach allows us to bridge the gap between learners and teaching experts. We evaluate our method using a collection of realistic learner queries and a dataset of Machine Learning and Data Mining documents. Evaluation results show our method to outperform benchmark approaches and demonstrates its effectiveness in assisting learners to find the right materials.

Introduction

Learners often have difficulty asking an effective query of a search engine for two reasons. First, they lack sufficient knowledge about the domain they are researching, so are unable to assemble effective keywords that identify what they wish to learn (Liu, Kim, and Creel 2013). This problem results in an intent gap. Second, the vocabulary used by teaching experts is often different from that used by learners, as learners often describe their problems in different terms to how experts present the solutions (Millard et al. 2005). This presents a semantic gap.

Artificial intelligence (AI) methods have been applied to assist the teaching process in the design of an online course, by using AI agents to provide feedback to learners (Goel and Joyner 2016). AI techniques are also used in our method to assist the learning process by creating an e-Learning recommender system that provides learners with relevant documents. The developed method allows us to bridge both the intent and semantic gap between learners and domain experts. We address the intent gap by placing a learner's query within the space of learning concepts, and identifying the most similar concepts to use for refining the query. The semantic gap is addressed by leveraging the vocabulary associated with the domain concepts to support the refinement of queries. This allows us to refine the learner's query using the vocabulary of the domain. The effect is to focus the search on relevant documents and improve the recommendations.

We use background knowledge extracted automatically from a structured collection of teaching materials, to influence query refinement by providing a vocabulary for refining the queries. An e-Learning recommender system is built to evaluate the effectiveness of our approach using a collection of realistic learner queries and a dataset of Machine Learning and Data Mining resources. Evaluation results show our method to outperform a standard Bag of Words approach. The results demonstrate that using background knowledge to refine learners' queries supports the learning process by helping students to find relevant documents.

There are two key contributions from this work. First, an effective method for refining a learner's query to support the retrieval of relevant documents. Second, an e-Learning recommender system that assists the learning process and helps learners to find relevant learning materials.

Related Work

A large amount of e-Learning materials is available to learners on the Web. However, learners are often discouraged by the time spent in finding and assembling relevant resources to support their learning goals (Chen et al. 2014). Often, learners are new to the topic they are researching, so they can have difficulty asking effective queries in a search engine. The unfamiliar vocabulary often used in teaching materials poses a challenge to learners trying to find relevant materials. One way of addressing these challenges is by refining queries to improve the recommendation made to learners.

One approach to query refinement is by using internal knowledge from a document collection as a feedback method (Wu and Fang 2013). This approach is similar to pseudo relevance feedback. In this method, an initial set of documents considered to be relevant are found, then terms from these documents are used to refine the query to improve retrieval performance. A drawback of this approach is that search results may be directed towards a few documents, and this can be harmful if the documents are only about specific topics. Further, the retrieval performance for

difficult queries can be affected if the initial retrieval set does not contain relevant documents (Li et al. 2007).

Another approach to query refinement involves using external knowledge sources for refining queries (Meij and de Rijke 2010; Meij et al. 2011). This approach entails using terms from domain sources to refine queries (Bendersky, Metzler, and Croft 2012). A knowledge source with a good coverage is usually recommended for this task. Sources such as Wikipedia (He and Ounis 2007; Xu, Jones, and Wang 2009), and DBpedia (Meij et al. 2009) have been used to identify potentially relevant terms to use for refining queries. The effectiveness of this approach has been demonstrated in previous work (Bendersky, Metzler, and Croft 2012). WatsonPaths applies a similar approach to reason over domain knowledge sources for answering queries in the medical domain (Lally et al. 2017). One potential challenge is query drift, where a refined query deviates from the original query (Xu, Jones, and Wang 2009). So, one needs to determine how much knowledge is sufficient for refining a query.

The approach in this paper draws insight from methods that use external knowledge sources. Our challenge is addressed in an e-Learning domain. So, the knowledge source used is drawn from learning concepts from the Tables of contents (TOCs) of e-Books and enriched with descriptive text from DBpedia abstracts. An initial set of similar learning concepts are automatically identified for each query. Highly weighted terms from the identified concepts then form the vocabulary that is used to refine a learner's query.

Background Knowledge

Background knowledge refers to specialized information about a domain that can be used for general understanding and problem solving (Zhang, Liu, and Cole 2013). We attempt to capture background knowledge as a set of domain concepts, each representing an important topic in the domain. For example, in a learning domain, such as Data Mining, you would find topics such as classification, association rules, and regression. Each of these topics is represented by a concept, in the form of a concept label and a pseudo-document which describes the concept.

We use the background knowledge that we developed in earlier work (Mbipom, Craw, and Massie 2016), to underpin e-Learning recommendation for a broad learning topic such as Machine Learning and Data Mining. In this work, the background knowledge is employed to support the refinement of queries. The background knowledge creation process is shown in Figure 1. The input to this process are domain knowledge sources, such as a structured collection of teaching materials and an encyclopedia source. Two knowledge sources are used as initial inputs for discovering domain concepts. First, the TOCs of 20 e-Books are used as a structured collection of teaching materials, which provide a source for extracting important topics identified by teaching experts in the domain. Second, a domain lexicon is used to verify that the concept labels identified from the teaching materials are directly relevant. The lexicon is created from Wikipedia because it contains articles for many learning domains (Zheng et al. 2010), and the contributions of many people (Yang and Lai 2010). The domain lexicon containing 664 Wiki-phrases provides a broader but more detailed coverage of the relevant topics in the domain. Thereafter, an encyclopedia source, such as DBpedia abstracts is searched and this provides the relevant text to form a pseudo-document for each verified concept label. The final output from this process is the background knowledge containing a set of 150 domain concepts each comprising a concept label and an associated pseudo-document.

The concept vocabulary with terms t_1 to t_c, from concepts, C_1 to C_m is used to create a concept term matrix with TF-IDF weighting (Salton and Buckley 1988). TF-IDF is useful for distinguishing concept terms in the concept space, and for identifying concepts that are relevant to queries hence its use in this method. A set of potentially useful concept terms are selected from the concept vocabulary as a means of scaling up the representation. So, we represent the background knowledge by using the top 10% of concept terms that have the highest TF-IDF values. The selected concept terms are used to create a concept term matrix. The selected terms t_{c1} to t_{cn}, from the concepts, C_1 to C_m are the set of potential terms that would be used for refining a query.

Refining Queries using Domain Concepts

The background knowledge representation is used to support the refinement of queries as a step towards addressing the intent and semantic gap learners face. When a new query is received from a learner, a search is performed on all the domain concepts. A ranked list of domain concepts that are similar to the query is retrieved. The terms from the term-vectors of the most similar concepts are put together to create a potential refined query. Terms with the highest weights are selected from the potential refined query and added to the initial query to create a refined query. The refined query can then be used to search on a document collection, and documents would be retrieved and presented to the learner. We expect the retrieved documents to be relevant because the query used for the search has been generated using domain concepts related to the initial query.

Figure 2 contains an illustration of how a refined query is generated. In this example, Cq_1, Cq_2, and Cq_k are the k most similar concepts to the query, while t_{c1} to t_{cn}, are the selected concept terms. The entries into the matrix are the tf-idf weights of the terms in the respective concepts. While, $SimScore_1$, $SimScore_2$ and $SimScore_k$ are the similarity scores between the query and concepts Cq_1, Cq_2, and Cq_k respectively. The weight of a concept term such as t_{c1} in the potential refined query is generated by computing the weighted sum for that term.

Equation 1 shows how the weight of term t_{c1} is computed. The weight t_{c1} is achieved by multiplying the weight $SimScore1$ with the tf-idf scores of terms that appear in concept Cq_1. This is also done for terms that appear in Cq_2 and Cq_k respectively. The column sum for t_{c1} is then computed. Altering the tf-idf weights of concept terms with the respective similarity scores allows terms from concepts that are more similar to the query to have more influence in the refined query. The output is a potential refined query with concept terms, t_{c1} to t_{cn} and their respective weights.

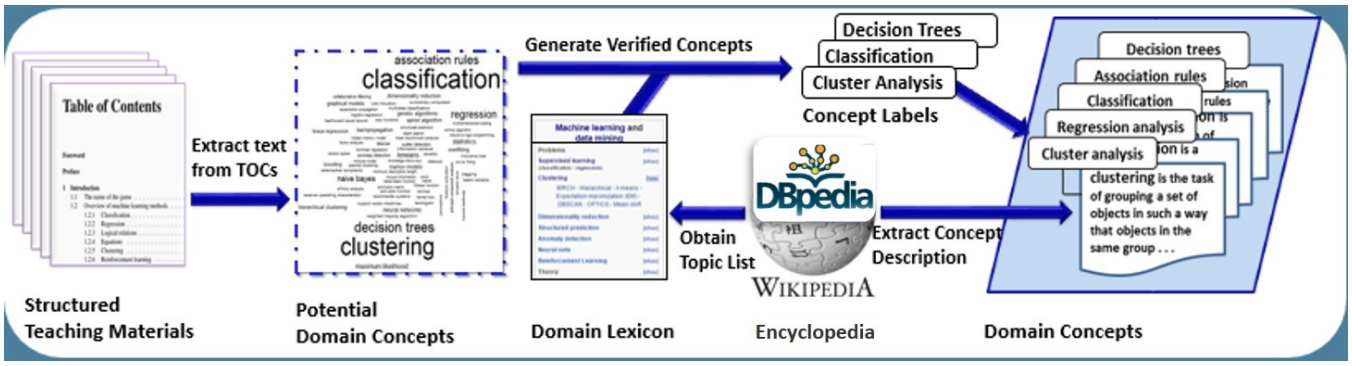

Figure 1: An overview of the background knowledge creation process

Concepts \ Terms	t_{c1}	t_{c2}	...	t_{cn}	SimScore
C_{q1}	tf-idf(t_{c1}, C_{q1})				$SimScore_1$
C_{q2}	tf-idf(t_{c1}, C_{q2})				$SimScore_2$
C_{qk}	tf-idf(t_{c1}, C_{qk})				$SimScore_k$
Potential refined query	t_{c1}	t_{c2}	...	t_{cn}	

Figure 2: Generating a refined query

$$Weight(t_{c1}) = \sum_{i=1}^{k} (\text{tf-idf}(t_{c1}, C_{qi}) \times SimScore_i) \quad (1)$$

where t_{c1} is a concept term, and $\text{tf-idf}(t_{c1}, C_{qi})$ is the tf-idf score of term, t_{c1} in the i-th concept, C_{qi}, and $SimScore_i$ is the similarity between the query, q and the i-th concept C_{qi}.

The weight of a term in the potential refined query gives an indication of the importance of the term within the concept space in relation to the given query. We take advantage of this weight by selecting the highly weighted terms from the potential refined query. These terms are used for generating a refined query. We adopt this approach so that noisy terms would not be included during query refinement (Xu, Jones, and Wang 2009). We include the initial query as part of the refined query to maintain the context of the query.

For example, given an initial query: *"How do you implement gradient descent algorithm?"*. A search is performed on the set of concepts and the 3 most similar concepts to this query are: *stochastic gradient descent*, *backpropagation*, and *winnow algorithm*. The terms from these concepts are put together as described above. We use the top 25 highest weighted terms from an amalgamation of these concepts as the generated concept terms. In this example, the first 10 terms are: gradient, descent, stochastic, formula, update, momentum, delta, rate, derivative, backpropagation. The top 25 terms are then added to the initial query. So, the refined query becomes: *initial query + generated concept terms*. The refined query is then used to search for documents.

User Evaluation

An e-Learning recommender system is developed to demonstrate how refined queries can be used for e-Learning recommendation. Three methods are implemented. First our CONCEPTBASED query refinement method which uses the most similar domain concepts to create a concept based representation of a query. Second, the Bag-Of-Words (BOW) method, which is a standard Information Retrieval method, where a learner's query is represented using the terms in the query only. Finally, a HYBRID method which takes advantage of the features in a query to make a dynamic choice in determining when to apply CONCEPTBASED or BOW to refine a query. The evaluation aims to compare the relevance of documents retrieved using CONCEPTBASED and HYBRID methods against the standard BOW method.

Data

The data used for evaluation is drawn from 2 collections. First, a document collection for recommendation and second, a collection of learner queries. The document collection contains 504 chapters from 32 Machine Learning and Data Mining (ML/DM) e-Books. The collection used is fairly distributed across the domain concepts. The query collection contains realistic learner-focused queries which we use for evaluating the system. We used 2 sources to generate our queries. First, postgraduate students in the School of Computing Science and Digital Media took part as learners in generating queries. An e-mail specifying the task was sent to them. In order to allow learners to send anonymous responses, and return their queries without seeing what others had asked, a Google form was created to capture the queries. Second, online sources such as Coursera's Machine Learning MOOC and Quora were used to generate queries. Course specific questions were accessed from Coursera's MOOC, while the open questions in Quora from the Machine Learning and Data Mining topics were chosen.

For the query sources, our aim was to have realistic learner queries, so we used queries where the user wanted to learn about a technique, for example: "what are the various data mining techniques for fraud detection". We did not use generic or career-related queries such as "What is it like to be a data scientist at Amazon?", or queries that were

out of scope such as "is there any course on ML?". Overall, 11 queries are from learners and 59 queries from online sources, resulting in 70 queries.

The evaluation system was deployed using Microsoft Azure (Copeland et al. 2015), so the system could be accessible to users online. The evaluation system was available online for 8 weeks. To generate the recommendations for evaluation, the 70 queries were run using the methods, and the top 3 recommendations from each method was stored. A link to the system was shared with users working in the ML/DM field. Each user completed a questionnaire at the start to provide data about their background in the ML/DM domain. The data users provided gave us an idea of the experience and expertise of the users. An analysis of the evaluation results using the questionnaire data allowed us to gain valuable insights into the way different users judged the recommendations made by the system.

There were responses from 22 users. There were 16 PhD students, 3 Researchers, and 3 Lecturers/Professors. All users had at least an MSc degree or higher. There were 3 users with over 10 years experience in ML/DM, 3 users had over 5 years experience, 10 users had between 3-5 years experience, 5 users had 1-2 years experience, and only 1 user had less than a year's experience in ML/DM. This level of experience in the ML/DM topic is useful, because the judgements should be from people who know the domain. For expertise, there were 2 experts, 16 competent users and only 4 beginners. So most users are competent or expert.

Design of the User Evaluation

At the start, each user was shown a briefing containing a guide on the evaluation study, the task, a note on confidentiality, and the researcher's contact information. During the study, a user was shown one query each time to evaluate. For each query, the user could choose to skip, if the user had no idea about the query, or proceed to evaluate because the user had some understanding of the query. This allowed each user to evaluate recommendations for queries they were knowledgeable about. When evaluating a query, the user was shown up to 6 retrieved documents in random order. The set of documents were the top 3 documents from the CONCEPTBASED and BOW methods. Since HYBRID applies either CONCEPTBASED or BOW, the documents for HYBRID are already included in the retrieval set shown to users.

It is important that the retrieval set of documents shown to users is presented in a way that avoids any potential bias. Three issues of bias are considered and addressed. First, the users do not know which method produced the recommendations they are evaluating, this prevents a user from favouring one method over the others. Second, the order of documents presented to users is randomized. This rules out the bias of documents shown at the top being considered to be relevant over those lower down the list. Third, the same user evaluates recommendations from both BOW and CONCEPTBASED methods for the same query. This ensures that the same user gives an evaluation for all methods at the same time for a given query. This prevents the possibility of receiving ratings from a positive user for one method, and ratings from a generally negative user for another method.

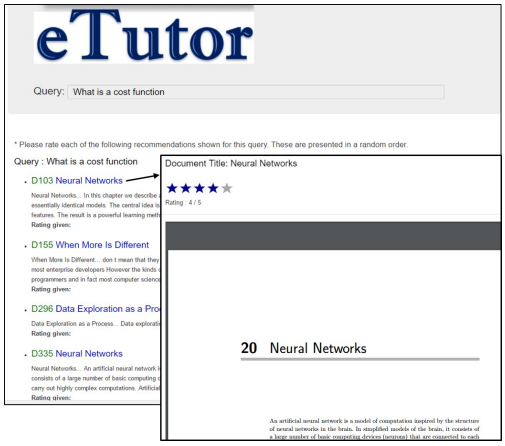

Figure 3: Recommendation screen

Figure 3 shows a page for displaying recommendations, and the screen for a document the learner has selected. The relevance of each document to the query is captured using a rating scale of 1 to 5 stars where 1 is least likely to be relevant and 5 stars is very relevant. The rating stars were included in the page that contained the document, so that each user had an opportunity of reading the document before rating it.

Evaluation Metrics

The evaluation uses the ratings given by users for query-recommendation pairs. We compute the $rating$ for a query-recommendation (q, r) pair as:

$$rating\,(q, r) = \sum_{u \in U_q} \frac{R_u(q, r)}{|U_q|} \tag{2}$$

where U_q are users that evaluated query q, and R_u is the rating user u has given to a (q, r) pair. Performance of a method is computed by taking the average rating for the queries.

Recommendation Results

Users evaluated 105 queries and provided ratings for 521 query-recommendation (q, r) pairs. There were 6 of the total 70 queries that were not evaluated. Figure 4 shows the spread of ratings for (q, r) pairs that were evaluated as heat maps for the CONCEPTBASED, HYBRID and BOW methods respectively. The ratings range from 5 to 1, for the highest rating of 5, to the lowest rating, 1. In plotting the heat map, the average rating values per (q, r) pair are sorted in descending order. The 3 heat maps are plotted in the same way using the actual average rating value given for each (q, r) pair. Lines are included to show a change in ratings.

In considering high ratings, HYBRID does best in producing documents with high quality ratings followed by CONCEPTBASED and then BOW. HYBRID is able to correctly identify when to use either BOW or CONCEPTBASED for refining a query in order to produce such good quality documents. For (q, r) pairs with the lowest ratings, the standard BOW method produces the highest number of documents with very poor ratings, BOW has difficulty pre-

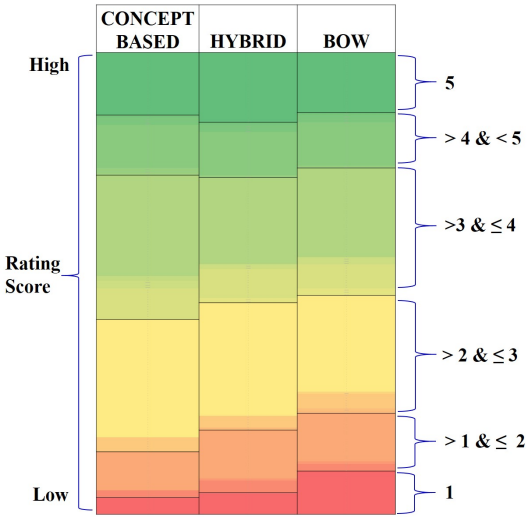

Figure 4: Spread of ratings for query-recommendation pairs

venting poor retrievals from being shown. CONCEPTBASED has the fewest number of (q, r) pairs with poor ratings. In particular CONCEPTBASED is very good at not presenting poor retrievals to users. These results for all the marked areas show that users gave higher ratings to the recommendations made using the HYBRID and CONCEPTBASED methods than those made using the standard BOW method.

We wish to know if a user's expertise affects the ratings they provided. This would allow us to confirm if there is some agreement among the users irrespective of their level of expertise. The rating for each (q, r) pair is computed based on Equation 2. Table 1 contains the average ratings of all users as well as ratings based on expertise of users for each of the 3 methods. For all the users, we find that CONCEPTBASED (CB) $>$ HYBRID $>$ BOW. So, using a CONCEPTBASED representation of a query to find learning materials is better than when HYBRID or BOW is used. For experts, the average rating scores for all methods are lower, nonetheless the experts still agree that the best performance is from the CONCEPTBASED method. We are confident in the results received from experts because they know what topics learners should be interested in. The competent users have higher rating scores across all methods, and they also agree that the CONCEPTBASED method performs better. Although the ratings by the beginners for all the methods are very similar, their rating scores also agree with the other users that the CONCEPTBASED method performs best.

Table 1: Average rating

Method	All	Expert	Competent	Beginner
CB	3.54	3.13	3.66	3.29
HYBRID	3.45	2.71	3.58	3.27
BOW	3.33	2.58	3.46	3.25

Relevance judgement is subjective and depends on the users who are providing ratings for the documents seen. We wanted to know how many users preferred the recommendations produced using either our CONCEPTBASED method (CB) or the standard BOW method. We use the rating as given in Equation 2 for BOW and CB and we count how many users rated documents from one method higher than the other. Table 2 contains the results for the preference users had for either CB or BOW. Half of the experts preferred CONCEPTBASED, while the other half of the experts thought both methods were the same. None of the experts thought that BOW was better than CONCEPTBASED. 14 users preferred the recommendations produced using the CONCEPTBASED method over those of the standard BOW method.

We trust the judgement of experts because they are more knowledgeable about the domain, and they know what documents should be relevant to learners. We can recall from Table 1 that the scores provided by the beginners for all methods were very similar, so we cannot rely totally on the judgements provided by beginners. Four users rated documents seen from both methods equally. This could come from users that gave equal scores to all documents seen for a query. Overall, majority of competent users preferred the documents from CONCEPTBASED over BOW.

Table 2: Preferences of methods

Preference	#	Demographics
CB $>$ BOW	14	1expert,12competent,1 beginner
BOW $>$ CB	4	1 competent, 3 beginners
CB $=$ BOW	4	1 expert, 3 competent

Coverage of Relevant Topics

Having recommendations with high ratings is good, but recommending documents that cover topics relevant to the query is important. We capture the coverage by asking users to provide feedback after they evaluate each query. This feedback is optional for users. Users were asked what extent they thought the documents they were shown covered the topics relevant to the query. The user could make a selection from 4 options: Complete, Good, Partial and Limited coverage. 50% of entries from users said the documents had good coverage. An additional 19% of entries said the documents had complete coverage, while 21% of entries said the documents had partial coverage. Only 10% of entries said the documents had limited coverage. So, most recommendations covered topics relevant to the queries.

Figure 5 contains a heat map which shows a broader view of the spread of coverage scores for the queries evaluated. The heat map is plotted by converting the coverage options to numeric values where Complete is 4, Good is 3, Partial is 2, and Limited is 1. The heat map is sorted twice. First based on the average coverage scores per query and second based on the average coverage scores per user. Figure 5 captures queries with the best coverage at the top, and those with least coverage at the bottom of the diagram.

The queries rated had consistent coverage scores from more than one user, so this allowed us to gain more insight to the queries. For example, a query such as: *"How does cluster analysis work?"* had "complete coverage" from 3 users

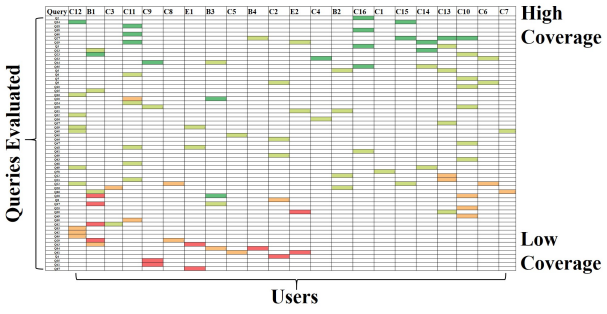

Figure 5: Spread of scores for the coverage

and "good coverage" from 1 user. While the query: *"Is it possible to use reinforcement learning to solve any supervised or unsupervised problem?"* had partial coverage from 1 user and limited coverage from 1 user. We draw the following insights from the coverage. First, some queries may not be well written and so can be hard to understand. Second, the topic contained in some queries may be very specialized. The results show potential in exploring the features of queries when refining them for e-Learning recommendation.

Conclusions

There are large amounts of e-Learning materials available to learners on the Web. However, learners often have difficulty finding relevant materials. Learners are often new to the topic they are researching, and so are unable to create effective queries in a search engine. We have created an e-Learning recommender system that uses background knowledge extracted from a collection of teaching materials to support the refinement of learners' queries. The rich vocabulary from the background knowledge is used to refine the queries and this allows us to focus the search on documents that are relevant to learners.

We use a collection of realistic queries and a dataset of Machine Learning and Data Mining documents for evaluation. Relevance judgement is subjective and depends on the opinions of the users taking part in the evaluation. However, we had a good level of consensus on the relevance judgements provided for each method by users with different levels of expertise. Results from experts, competent users and beginners all showed that using a CONCEPTBASED representation of a query to search produced documents that were consistently more relevant to learners than when the standard method was used. User evaluation results demonstrate the effectiveness of our approach to assist students to find the right learning materials. In future, the background knowledge can be developed to provide a guided view of a learning domain and support intelligent browsing of documents.

References

Bendersky, M.; Metzler, D.; and Croft, W. B. 2012. Effective query formulation with multiple information sources. In *5th ACM International Conference on WSDM*, 443–452.

Chen, W.; Niu, Z.; Zhao, X.; and Li, Y. 2014. A hybrid recommendation algorithm adapted in e-learning environments. *World Wide Web* 17(2):271–284.

Copeland, M.; Soh, J.; Puca, A.; Manning, M.; and Gollob, D. 2015. *Microsoft Azure: Planning, Deploying, and Managing Your Data Center in the Cloud*. Apress.

Goel, A. K., and Joyner, D. A. 2016. Design of an online course on Knowledge-Based AI. In *AAAI*, 4089–4094.

He, B., and Ounis, I. 2007. Combining fields for query expansion and adaptive query expansion. *Information Processing & Management* 43(5):1294–1307.

Lally, A.; Bagchi, S.; Barborak, M. A.; Buchanan, D. W.; Chu-Carroll, J.; Ferrucci, D. A.; Glass, M. R.; Kalyanpur, A.; Mueller, E. T.; Murdock, J. W.; Patwardhan, S.; Prager, J. M.; and Welty, C. A. 2017. WatsonPaths: Scenario-based question answering and inference over unstructured information. *AI Magazine* 38(2).

Li, Y.; Luk, W. P. R.; Ho, K. S. E.; and Chung, F. L. K. 2007. Improving weak ad-hoc queries using Wikipedia as external corpus. In *30th ACM SIGIR*, 797–798.

Liu, J.; Kim, C. S.; and Creel, C. 2013. Why do users feel search task difficult? *ASIS&T* 50(1):1–4.

Mbipom, B.; Craw, S.; and Massie, S. 2016. Harnessing background knowledge for e-learning recommendation. In *36th BCS SGAI*. Springer. 3–17.

Meij, E., and de Rijke, M. 2010. Supervised query modeling using Wikipedia. In *33rd ACM SIGIR*, 875–876.

Meij, E.; Bron, M.; Hollink, L.; Huurnink, B.; and De Rijke, M. 2009. Learning semantic query suggestions. In *The Semantic Web - ISWC*, 424–440. Springer.

Meij, E.; Bron, M.; Hollink, L.; Huurnink, B.; and de Rijke, M. 2011. Mapping queries to the linking open data cloud: A case study using DBpedia. *Web Semantics: Science, Services and Agents on the World Wide Web* 9(4):418–433.

Millard, D. E.; Gibbins, N. M.; Michaelides, D. T.; and Weal, M. J. 2005. Mind the semantic gap. In *16th ACM Conference on Hypertext and Hypermedia*, 54–62.

Salton, G., and Buckley, C. 1988. Term-weighting approaches in automatic text retrieval. *Information Processing & Management* 24(5):513 – 523.

Wu, H., and Fang, H. 2013. An incremental approach to efficient pseudo-relevance feedback. In *36th ACM SIGIR*, 553–562.

Xu, Y.; Jones, G. J.; and Wang, B. 2009. Query dependent pseudo-relevance feedback based on Wikipedia. In *32nd ACM SIGIR*, 59–66.

Yang, H.-L., and Lai, C.-Y. 2010. Motivations of Wikipedia content contributors. *Computers in Human Behavior* 26(6):1377–1383.

Zhang, X.; Liu, J.; and Cole, M. 2013. Task topic knowledge vs. background domain knowledge: Impact of two types of knowledge on user search performance. In *Advances in Information Systems and Technologies*. Springer. 179–191.

Zheng, Z.; Li, F.; Huang, M.; and Zhu, X. 2010. Learning to link entities with knowledge base. In *NAACL HLT*, 483–491.