CSCLRec: Personalized Recommendation of Forum Posts to Support Socio-collaborative Learning

Zhaorui Chen and Carrie DEMMANS EPP
EdTeKLA Research Group, Dept. of Computing Science, University of Alberta
{zhaorui, cdemmansepp}@ualberta.ca

ABSTRACT

Discussion forums are used to support socio-collaborative learning processes among students in online courses. However, complex forum structures and lengthy discourse require that students spend their limited time searching and filtering through posts to find those that are relevant to them rather than spending that time engaged in other meaningful learning activities (i.e., discussion). Moreover, existing adaptive systems do not accommodate individual learner needs in these contexts. In this work, we propose a multi-relational graph-based recommendation approach that mines student interaction logs to address the above problems within discussion-based socio-collaborative online courses. To account for the social aspects of learning, our approach incorporates learner modeling, social network analysis, and natural language processing techniques; it offers tailored recommendations of forum posts for learners with different types of interaction behaviors. In our experiments with small online courses, our approach outperformed competitor approaches in terms of recommendation precision while meeting expectations with respect to diversity and novelty. The results illustrate the proposed algorithm's effectiveness in predicting student preferences, suggesting its potential to increase student participation in discussion-related learning activities.

Keywords

Recommender systems, Discussion forums, Computer-supported collaborative learning, Online learning.

1. INTRODUCTION

Asynchronous online discussion forums are widely used to support online courses in higher education [5, 23, 25]. In these forums, many instructors post discussion topics and encourage students to expand so that knowledge can be co-created and developed through progressive discussion. In such socio-collaborative learning contexts, students' active participation and production of learning resources is essential, as less discussion could result in less sharable knowledge and thus less learning within a course [37, 81]. However, forums' complex thread structure and information-heavy posts tend to have a negative impact on student engagement, because much of their time is spent locating relevant forum posts, rather than focusing on core tasks such as debating, reflecting, and learning from each other [1, 37]. To alleviate this type of information overload problem, deploying recommender systems to recommend posts of interest or content generated by others could be beneficial.

Zhaorui Chen and Carrie Demmans Epp "CSCLRec: Personalized Recommendation of Forum Posts to Support Socio-collaborative Learning" In: *Proceedings of The 13th International Conference on Educational Data Mining (EDM 2020)*, Anna N. Rafferty, Jacob Whitehill, Violetta Cavalli-Sforza, and Cristobal Romero (eds.) 2020, pp. 364 - 373

Many recommender systems have been used to support learning across varied domains and contexts. For example, data mining approaches were used to suggest course improvements in learning management systems [33], and a workplace learning support system paired users with knowledgeable peers to enable knowledge sharing processes [8]. More recently, other systems have recommended courses to university students [7, 29, 65].

While these examples show the prior success of recommender systems in educational contexts, few have solved the problem of recommending socio-collaborative learning materials in discussion forums for smaller online courses. To fill this gap, we present a novel graph-based recommender system approach. This approach mines learner interaction data using both modelled learner types and natural language processing techniques that were specifically designed for this application domain of smaller discussion-based socio-collaborative learning environments. In our research, we posed the following question: ***How do traditional recommender algorithms and those that incorporate principles from socio-collaborative learning perform when suggesting posts in small online socio-collaborative learning contexts?***

2. Related Work

2.1 Socio-collaborative Learning

Socio-collaborative learning, also known as collaborative learning, refers to a class of learning methods in which learners cooperate in a group, relying on each other, being responsible for each other, and accomplishing a common task together [75]. This approach can be traced back to Vygotsky who pointed out that those who are more able can help others perform better [83]. Piaget claimed that the cognitive conflicts generated during social interaction could help the learner reflect on their original point of view, thus enhancing their understanding [44]. Subsequently, collaborative learning has become a widely used pedagogical theory that is also a target of many online learning environments, where it is called computer-supported collaborative learning (CSCL).

CSCL often occurs through online discussion-based forums [17, 26, 39] where information is transmitted through posts to enable knowledge sharing or co-construction among learners. The systems and mechanisms used to support CSCL are grounded in theories such as knowledge building: a specific knowledge co-construction process that emphasizes the creation of ideas through discussion [70, 71]. Many of the proposed knowledge building principles (e.g., diverse and improvable ideas or symmetric knowledge advancement [69]) provide theoretical support for our research.

2.2 Recommenders for Educational Forums

Most work has focused on supporting question and answer (Q&A) forums in university courses [34], MOOCs [49, 53, 87] or other online educational platforms [41, 80] when recommending forum posts. These systems typically aim to reduce the number of unanswered questions by recommending 1) unanswered questions to students who are able to answer them, and 2) similar questions that have already been answered to users who are about to ask one.

Using a similar recommender system design in smaller-scale socio-collaborative settings is inappropriate because the contexts differ in terms of size and pedagogical purpose. In contrast to MOOCs, these contexts suffer from both a lack of data and the cold start problem. Different from Q&A forums, developing knowledge sharing processes in discussion forums requires the algorithm support increased connectivity among users to facilitate communication [45]. It is also necessary to include posts containing diverse and novel ideas from students who express different points of view so that they might learn from each other [70].

Few studies have investigated how to deal with these challenges. Those that have depend on a priori domain knowledge (e.g., rules [2] or ontologies [18]) which is time-consuming to obtain and has limited generalizability to unseen cases [43]. Given increases in online course delivery and a desire to support students' socio-emotional development and collaborative learning [3, 24], we need approaches that can be used in the absence of domain expertise.

Many have also argued that CSCL personalization technologies should consider the social [42, 66] and other needs of learners [12, 51, 68, 79]. One study investigated learners' knowledge sharing behaviors in closely-knit communities to generate tailored notifications [45]. The notifications aimed to foster knowledge sharing processes within a learning community composed of different learner types. To extend this idea to the context of a forum post recommender system, we set out to develop recommendation algorithms that also consider learner socio-behavioral patterns and created customized strategies for each behavior pattern.

3. Recommender Algorithm: CSCLRec[1]

The target users of this system are students or learners. We will use these terms interchangeably. Our proposed algorithm, CSCLRec, relies on 3 types of data that are available in any educational forum: user interactions with forum posts which we call user-to-post (U2P) interactions; communication between users, such as reading, that we call user-to-user (U2U) interactions; and the textual content of forum posts. Using this data, it recommends posts to learners.

CSCLRec has four modules (see Figure 1): a personalized PageRank graph, a learner interaction profiler that analyzes U2U interactions, a content analyzer, and a post filtering module.

3.1 Personalized PageRank Graph

The core of the system is a modified personalized PageRank (PPR) graph [35]. As shown in Figure 2, the PPR has nodes for users, posts, and hypernyms. A hypernym is a superordinate word whose semantic meaning includes a set of other words. For example, "flower" is the hypernym for "rose" or "daisy". Multiple types of relationships including U2P interactions, inter-user relationships, and posts' relationship with hypernyms are computed by other modules and represented as edges in the graph. The weight of user-to-post edges in the graph is biased by a temporal decay rate. Edges representing U2P interactions in the past have lower weights so that the algorithm can focus on the user's recent interests.

We refer to the user who is receiving the recommendations as the active user. To recommend posts to an active user, the algorithm performs a random walk starting from their user node and its

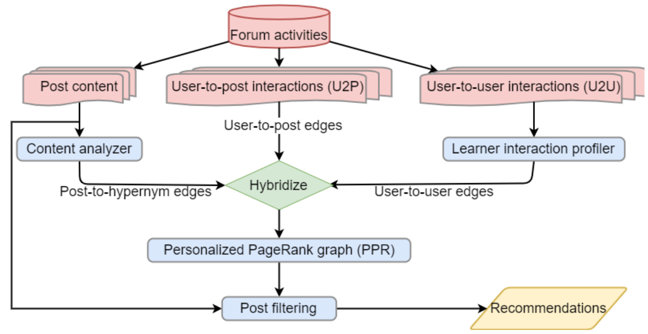

Figure 1. Overall workflow of CSCLRec

connected post nodes. When sufficient iterations have been completed, the nodes' probabilities of being visited by the random walk agent will converge to a steady state. Posts with the highest probability of being visited are presented as the recommendations. We used power iteration [60] to approximate the stationary probabilities and avoid poor computational performance.

3.2 Learner Interaction Profiler

The learner interaction profiler uses a bidirectional social network graph, which consists of different types of U2U interactions (e.g., replies and reads). Each user is a node in this graph and the interactions among users are edges. In Figure 3, the thin grey link from user U1 to user U2 indicates that U1 has read U2's post.

Students who have many interactions with the active user are their peer learners. The rich interaction history, whether in discussion or debate, indicates the active user's interest in interacting with those peers. This group of users share many outward edges with the active user in the social network graph. The inclusion threshold for number of edges required between users is set via grid search. As a result, the module generates links connecting the active user to those peer learners (the green edges in Figure 2) in the PPR graph.

The analyses over the graph also output a participation level (i.e., number of outgoing edges of reply, like, and link types from its user node) and a degree of centrality (i.e., the in-degree of a node) for each student. The more frequently other students interact with the active user's posts, the more they can increase the active user's degree of centrality. The participation level indicates the extent to which the student is actively engaging in the discussion. We used the two measures to identify four types of learners (new user, listener, single-pass user, and peripheral user) that may need differentiated recommendation strategies. These user types are identified using simple heuristics based on the literature.

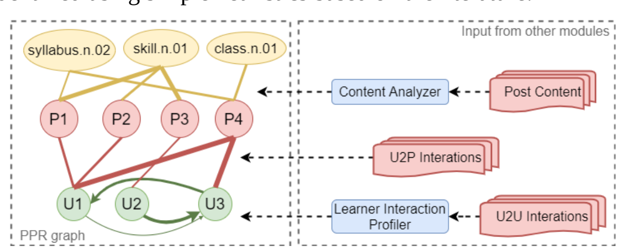

Figure 2. The modified PPR graph has 3 node types (user - green, post - red, and hypernym - yellow) and 3 edge types (user-to-user - green, user-to-post - red, and post-to-hypernym - yellow). Edges without arrows are bidirectional and edge width indicates number of occurrences.

[1] Code is available at https://github.com/EdTeKLA/CSCLRec

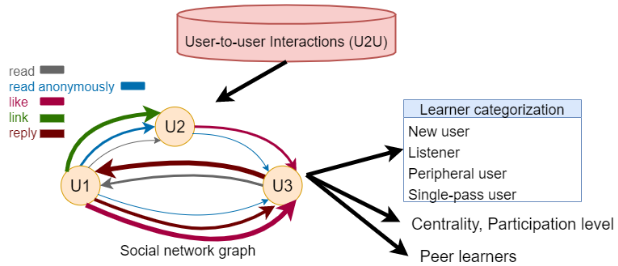

Figure 3. The workflow of the learner interaction profiler. Edge widths in the graph indicate the number of interactions.

New users are learners who have just joined the discussion. They have not created any resources nor do they have any other logged interactions. Consequently, these users are subject to the cold-start problem, which makes it difficult for the algorithms to provide suggestions because of the lack of data [10]. New users may experience greater information overload because they face many posts at once and may need tailored recommendations to help them filter information, identify their interests early, and contribute their own voices. To prevent narrow recommendations, new users are connected to every other user in the PPR graph.

Listeners read many posts but rarely post themselves [86]. The knowledge building principle of collective responsibility and symmetric knowledge advancement suggests that encouraging posting is critical to fostering activity and promoting knowledge co-construction [69]. Listeners are identified as those who have not created posts. To reduce the number of persistent listeners, we adopt the same recommendation strategy as that employed for new users since exposing these learners to different topics may increase the possibility of their expressing opinions [46].

Peripheral users are those whose centrality score is decreasing due to lost interactivity in their readership. The module aims to recover peripheral users and listeners to promote the knowledge-sharing process [47, 48] and prevent the loss of these readers' activity and interest. The learner profiler monitors the number of interactive readers for each learner: those who reply, like, or link. When the profiler detects the user's interactive reader count has dropped by half from one week to the next, that user is marked as a peripheral user. This value was tuned during the evaluation. The algorithm takes note of the lost readers and introduces connections between the peripheral user and the lost readers in the PPR graph to strengthen their connections.

Single-pass users only read new posts and ignore older posts [38]. Their widespread presence undermines socio-collaborative learning approaches because these learning processes require topics to be progressively discussed and deepened [69]. To alleviate this behavior, some have suggested encouraging students to revisit earlier posts [38]. Inspired by this idea, the learner interaction profiler identifies students who have only read posts from the previous week. For example, those who have not read posts created before week 7 are marked as single-pass users in week 8. The modified PPR graph decreases the temporal decay exerted on older posts for single-pass users so earlier posts are down-weighted less, increasing the likelihood of their recommendation to these learners.

3.3 Content Analyzer

Forum posts are hierarchically structured. Posts on the same topic have similar interaction records because users are accustomed to browsing the entire topic structure when reading a post. Therefore, algorithms based on interaction records (i.e., collaborative filtering, ordinary personalized PageRank bipartite graph) may only recommend posts that are locationally similar to those that users often interact with. Consequently, students may lose the opportunity to read posts that match their current interests because they are located elsewhere. These algorithms also bias towards post popularity [77] causing the "long-tail" problem: unpopular posts are not considered for recommendation [21, 62], which could decrease student exposure to diverse perspectives. To overcome these challenges, the content analyzer module applies natural language processing (NLP) techniques to the content of forum posts and enables links to be created between posts based on the concepts discussed rather than user interactions (as shown in Figure 2). Its workflow is shown in Figure 4. The preprocessing stage removes all html mark-up and punctuation. It also tokenizes sentences into individual words. Lemmas are extracted for nouns and verbs, and stop words are removed. To protect user privacy, person names, usernames, web URLs, and email addresses are also removed. Each post is then organized as a bag-of-words (BoW).

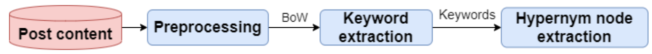

Figure 4. The workflow of the content analyzer module

TF-IDF was chosen for keyword extraction following a preliminary evaluation that compared several potential methods (i.e., RAKE [67], TextRank [54]) on independent data from the same system. TF-IDF scores are computed for each lemma to choose keywords that best differentiate the current post from others. The keywords with the top 1/5 TF-IDF scores are used to represent the post.

The extracted keywords are used to measure thematic similarity across posts. Instead of matching keywords using text similarity approaches (i.e., sentence embeddings or topic distribution vectors in vector space models), we consider two posts thematically similar provided they mentioned similar concepts regardless of student opinion towards a topic. The tools used to measure similarity included the WordNet semantic network and its collection of hypernyms [55]. We query each post in WordNet and use Lesk [50] to disambiguate hypernyms. The hypernyms are added as nodes in the PPR graph - see the yellow nodes in Figure 2. When a post contains a keyword that belongs to this hypernym, a link from the post node to the hypernym node is constructed. As a result, posts that share more concepts will share more hypernym nodes.

3.4 Post Filtering

This module analyzes, sorts, filters, and re-ranks the results produced by the recommender which may otherwise include less-informative posts that will not advance student knowledge. Posts like, "Thank you for the clarification, [name]" may be output by the algorithm if this filtering is not performed. The post filtering module refines the recommendations using two filters: one extracts verb and noun phrases as trigram models and excludes posts with fewer than 3 phrases, and the other compares post content with the Academic Word List (AWL) [20]. Posts with fewer than 3 AWL words are removed.

4. METHODS

We evaluated the performance of CSCLRec, its precursor, and other widely-used algorithms using a similar protocol to that advised by recommender system researchers [28, 73]. In each week, we recommend 10 posts to each user. Posts were selected from a candidate list consisting of those the active user has not yet read and all posts created by others in the current (evaluation) week. We hide this user 's activities from the evaluation week and use forum

Table 1. Student and instructor interactions through the course forum as a raw count (#) or *M (SD)*.

Course	Weeks (#)	Students (#)	Instructors (#)	Posts (#)	Interacted posts/ student	Interactions/ student	Reads/ student	Likes/ student	Links/ student
LA	13	26	1	1751	1176 (550.18)	1628 (1010.30)	1314 (719.23)	76 (61.16)	1.19 (2.98)
LB	13	19	4	809	358 (245.96)	441 (298.75)	365 (247.00)	21 (24.96)	0.05 (0.23)
LC	13	30	4	2090	1212 (686.41)	1373 (732.37)	1226 (698.94)	29 (23.88)	10.06 (15.20)
SA	6	23	1	627	362 (219.05)	505 (417.60)	405 (290.29)	15 (19.45)	0.26 (1.25)
SB	6	24	1	1142	616 (269.83)	731 (281.87)	635 (270.13)	8 (9.97)	0.25 (1.03)
SC	6	20	1	869	507 (223.65)	631 (269.11)	521 (223.42)	44 (44.63)	0.55 (1.57)

activities from prior weeks to train the recommenders. We start from week 2 since there are no learner posts prior to week 1.

4.1 Dataset

The evaluation used historical data from six postgraduate courses offered through an asynchronous discussion platform (PeppeR) at the University of Toronto. PeppeR provides a collaborative learning space to discuss and share ideas, making this system an ideal testbed to evaluate the proposed recommender system.

Archival data from fully online courses were used. Of the six test courses, three were regular-length courses (13-weeks long) and the others were short courses (6-weeks long). User activity statistics for each course are summarized in Table 1. The data includes forum posts and all kinds of user interactions with posts (i.e., posting, replying to others' posts, inserting hyperlinks to other posts, liking posts, and reading posts). The interactions were categorized into 7 types: create, reply, like, link, revisit, read, and anonymously read.

The large variability in student interactions (see Table 1) is consistent with the different types of users identified: Some had many forum activities, while others seldom interacted with posts. This suggests the necessity of distinguishing different learner types and employing user-specific recommendation strategies.

4.2 Recommender Algorithms

CSCLRec's performance was compared against that of 7 other recommenders. Due to the limited number of students, the diversity of student interaction behaviors, and the inter-dependence of time-series data, random cross validation was not appropriate. We tuned the hyperparameters using last block validation [9]: For each weekly evaluation, we used the prior week to validate the current weeks' recommendations. We used grid-search on the hyperparameters and trained the recommenders using data from before the validation week. Testing used data from the validation week. Using precision, the best performing hyperparameters were selected to build the recommenders for subsequent evaluations. For CSCLRec, we tuned temporal decay and the number of peer learners. All PPR-based algorithms had their damping factor tuned.

The algorithms we tested CSCLRec against are listed below. Hyper-parameter values are reported in the repository[1].

- Co-occurrence graph-based personalized PageRank (**CoPPR**) is another original method we developed. It uses the same learner profiler and post filtering modules as CSCLRec. Different from CSCLRec, Co-PPR uses the extracted keywords as nodes. Two keyword nodes are connected if they co-occurred at least once in a post. A post is connected to a keyword node if that post contains the keyword at least once. Edge weights are determined using the posts' keyword occurrence count. CoPPR helps identify the contribution of the content analyzer to CSCLRec. We tuned temporal decay, the damping factor, and number of peer learners.
- Personalized PageRank (**PPR**) is a widely used graph-based recommender [15, 58]. It uses a bipartite graph with user-to-post interactions as the only input.
- Matrix factorization collaborative filtering (**MCF**) represents a family of model-based collaborative filtering algorithms, which are commonly used in educational recommender systems [27, 78]. We used the version proposed by Hu and colleagues [40]. We tuned its confidence factor which specifies the negative weight attributed to unseen interactions.
- Keyword-based content-based recommender system (**KCB**) is frequently used to personalize discussion forums [4] and help-seeking platforms [52]. KCB relies on latent semantic indexing to create vectors from posts. Users are represented as the average of the post vectors they have interacted with before. It recommends candidate posts which are nearest to the active user in the vector space. The hyperparameters include the dimension of post vectors and the ratio of content words as the keywords (i.e., 1/7 of the content words are treated as keywords).
- Sentence embedding-based content-based recommender (**SCB**) relies on the semantics of post content [16].
- Popularity-based recommender (**PPL**) recommends popular posts. Every user receives the same recommendations. This unpersonalized algorithm is used as a baseline.
- The random recommender (**RND**) randomly draws posts from the candidate list. This algorithm is also used as a baseline.

We did not test all well-known recommendation algorithms as some structural aspects and requirements of the algorithms make them a poor fit given the nature of our dataset. For example, deep learning-based methods (i.e., autoencoders) are data-hungry and can easily overfit due to the size of our dataset [88].

4.3 Measures

Since accuracy is insufficient for determining the quality of educational recommender systems [30], we measured 3 dimensions of performance: accuracy, diversity, and novelty.

For accuracy, we report both Precision at K (**P@K**) and Recall at K (**R@K**), where k is the number of recommendations. The R@K measure is affected by the number of available relevant items [73] so we report the maximum (max) R@10 to aid interpretation. Max R@10 is the average of the largest possible R@10 in each user's recommendations. We adopted the commonly used intra-list diversity (**ILD**) indicator which measures the average pairwise distance between recommended items [14, 76]. We used pre-trained Universal Sentence Encoder [16] embeddings to represent the posts and the cosine distance to compute ILDs. The mean inverse user frequency (**MIUF**) indicator is used to measure recommendation novelty [11]. The fewer people who have interacted with the post, the higher the novelty and IUF of that post. To reflect the consistency of algorithm performance, we report the

Table 2. Summary of evaluation results as *M (SD)*

	Long courses (LA, LB, LC)				Short courses (SA, SB, SC)			
Algorithm	**P@10**	**R@10**	**ILD**	**MIUF**	**P@10**	**R@10**	**ILD**	**MIUF**
CSCLRec	**0.729 (0.319)**	0.219 (0.305)	0.274 (0.125)	0.612 (0.360)	**0.751 (0.310)**	0.188 (0.254)	0.191 (0.059)	0.482 (0.140)
CoPPR	**0.718 (0.324)**	0.221 (0.304)	0.222 (0.110)	0.638 (0.380)	**0.731 (0.315)**	0.177 (0.243)	0.156 (0.048)	0.502 (0.137)
PPR	0.537 (0.408)	0.178 (0.310)	0.390 (0.162)	0.466 (0.251)	0.566 (0.383)	0.142 (0.248)	0.244 (0.106)	0.407 (0.144)
MCF	0.484 (0.391)	0.180 (0.313)	0.449 (0.192)	0.837 (0.532)	0.449 (0.406)	0.130 (0.265)	**0.357 (0.151)**	0.801 (0.485)
SCB	0.294 (0.355)	0.158 (0.313)	0.075 (0.047)	**1.216 (0.453)**	0.400 (0.378)	0.117 (0.247)	0.079 (0.019)	0.927 (0.251)
KCB	0.289 (0.359)	0.150 (0.315)	0.221 (0.105)	1.053 (0.490)	0.397 (0.369)	0.115 (0.247)	0.188 (0.072)	**1.038 (0.311)**
RND	0.307 (0.335)	0.157 (0.312)	0.406 (0.174)	**1.174 (0.404)**	0.350 (0.336)	0.113 (0.248)	**0.350 (0.130)**	**1.197 (0.385)**
PPL	0.407 (0.407)	0.177 (0.310)	0.417 (0.164)	0.420 (0.243)	0.480 (0.402)	0.140 (0.249)	**0.311 (0.136)**	0.353 (0.135)

1.The best performing algorithms are bolded as determined via a 2-Way ANOVA and post-hoc Tukey HSD tests ($p < .05$). No interactions between week and algorithm were found. Full results of statistical testing are available in the repository[1].

2.Max R@10 as *M (SD)*: long courses - 0.351 (0.343), short courses - 0.303 (0.317) · Sample size: long courses - 825, short courses - 268.

mean and standard deviation of measures. The results were averaged over those of each student in each week.

5. RESULTS

Table 2 shows that both CSCLRec and CoPPR achieve high prediction accuracy, while maintaining acceptable diversity and novelty. They outperform their competitors according to precision for both 13-week and 6-week courses. Except for the very similar CoPPR algorithm, CSCLRec's precision is more than 18% higher than that of other recommenders. According to R@10, there is no measurable performance difference among the various algorithms. Considering the maximum possible recall is capped at 35.1% and 30.3%, CSCLRec's R@10 performance (21.9% and 18.8%) suggests it successfully identifies most of the relevant items.

In short courses, two of the best performers according to ILD are the unpersonalized baseline recommenders (RND and PPL) largely due to their introducing randomness. Apart from these random methods, those that emphasize interactions (i.e., MCF and PPR) had better diversity. As a tradeoff to accuracy, CSCLRec's diversity was acceptable - it was somewhere in the middle (3rd of 6 personalized recommenders in ILD) when baseline approaches (RND and PPL) are excluded because they have low precision.

As another tradeoff to high precision, novelty is not best achieved with CoPPR or CSCLRec. Content-based algorithms (i.e., KCB and SCB) performed well from a novelty perspective as shown through their average MIUF scores. However, they had low diversity scores (ILD); they ranked last or second last.

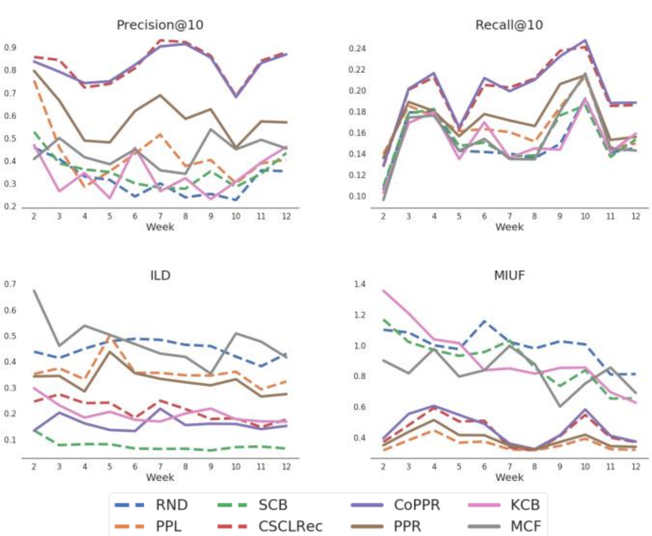

Figure 5. Weekly recommendation results for the LA course.

To illustrate differences in performance over time, we use the LA course as an example (Figure 5). Note similar patterns were present in other courses and the change at week 10 coincides with the term break. In general, CSCLRec and CoPPR remained the best performing recommenders for precision throughout the semester.

Our proposed algorithm, CSCLRec, beats its competitors in precision from weeks 2 or 3 onwards (Figure 5). In contrast, when few inputs from students were available at the beginning of courses, the performance of content-based approaches was worse than the baselines. These results suggest the inclusion of socio-collaborative elements helps address the cold-start problem.

6. DISCUSSION

6.1 Recommender Algorithm Performance

The good performance of content-based recommender algorithms (CB) such as KCB and SCB in recommendation novelty highlights their ability to discover unpopular posts. It implies these approaches are better at helping students more quickly locate difficult-to-find but conceptually related discussions when the goal is to develop narrow but deep knowledge. This class of approaches may also increase forum equity by increasing the visibility of posts made by students with minority opinions that may otherwise go unnoticed in a popularity-based recommender scheme. However, CB algorithms' poor diversity performance suggests they suffer from the over-specialization problem because they only care about content similarity. This makes them unable to recommend semantically diverse resources. Since discussions on the same thread usually have similar content, the suggestions provided by CB recommendation algorithms are likely to direct users to a few specific threads, which may prevent exposure to new ideas. This goes against the general teaching goals of learning contexts where students are expected to discuss and debate different topics.

The family of collaborative filtering algorithms (CF), represented by MCF, showed relatively poor novelty when compared with the CB algorithms. This lack of novelty may discourage the participation of students who hold minority opinions, as has been seen in other investigations [64]. When comparing with PPR, CSCLRec's considerable enhancement in precision demonstrates the effectiveness of its three add-on modules. CoPPR also performs well, but its recommendation diversity appears to be lower. This finding indicates that the design of CSCLRec's content analyzer module benefits recommendation diversity as it is the only difference between CSCLRec and CoPPR. The outstanding performance of CSCLRec makes it the most appropriate choice to provide personalized suggestions in small-scale socio-collaborative learning contexts. However, we expect CoPPR to work better than CSCLRec in environments where the domain of discussion topics is narrower than those in our dataset as the use of hypernyms in CSCLRec is prone to mis-classification due to the granularity of the WordNet ontology. For example, in Chemistry, both “Sodium chloride” and “Copper(II) sulfate” are a “chemical compound”, but it makes little sense to link these two terms together as students might be talking about different things.

In contrast to traditional algorithms, CSCLRec and CoPPR integrate pedagogical considerations. The learner interaction profiler module is an obvious case. Unlike the one-size-fits-all recommendation strategy in other algorithms, it employs different strategies (i.e., adding more user-to-user edges) depending on learner type. Compared with the results of other approaches, the user-centered recommendation algorithm design of CSCLRec provided better prediction results by taking advantage of socio-collaborative learning principles.

While CSCLRec tended to perform well in recommendation accuracy, we should acknowledge that such support is not always what is needed for some learner types. For example, listeners not actively engaging in the discussions could be attributed to the recommendation lacking diversity. In this case, collaborative filtering approaches such as MCF might be a better choice. Moreover, new users may benefit from unpersonalized recommenders. For example, PPL could be used when we lack information about that learner because popular discussions may pique newcomer’s interest and encourage them to participate.

6.2 Recommender Support for Learning

The evaluations confirmed that our recommender system can forecast student behavior and give recommendations that match students’ preferences, as represented through their behaviors, in an e-learning discussion forum. Here we discuss the system’s potential to enhance students’ learning processes and outcomes in socio-collaborative learning spaces.

Rooted in learner interest, the generated recommendations can help reduce the time students spend searching for useful resources, thereby increasing the proportion of time dedicated to learning activities (i.e., discussing and sharing). The increased interaction should enable more knowledge-construction within the forum [37, 72], benefiting every learner with more opportunities to review and increase their understanding of the knowledge they have learned [85]. Many empirical studies have also found that student’s active participation in sharing can develop their critical thinking abilities [13] and benefit their overall course performance [19, 61, 84].

At the same time, pedagogical research shows that the diversity and novelty of ideas are critical to learning outcomes, especially during the process of knowledge co-construction. According to the theory of social constructivism, learner exposure to diverse perspectives can help them experience the types of cognitive conflict that lead to knowledge gain [32, 44]. Knowledge building principles also emphasize the importance of diversity and novelty of ideas to the knowledge scaffolding process [69]. Fortunately, CSCLRec’s novelty and diversity performance demonstrated the algorithm’s potential to support various collaborative learning activities in small discussion-based e-learning forums.

6.3 Potential Expansions

There are many ways to further improve the system’s performance when it is deployed online. First, real-time feedback from students can be collected and used to steer the strategies for the next round of recommendations. Second, the system could allow instructor and student configuration. This would allow users to refine the quality of recommendations and offer increased transparency to improve user satisfaction and trust in the recommendation mechanism [82]. In the future, we may adopt a human-in-the-loop approach and let course instructors adjust recommender parameters so they are more consistent with desired teaching plans.

More advanced NLP methods could also be used. For example, using knowledge graphs could benefit graph-based recommenders [56, 57, 59, 63]. Using such approaches could extend the semantic network in the content analyzer. Knowledge graphs relying on Linked Open Data usually have a wider coverage of entities which may allow them to overcome the current algorithms’ lack of phrases for representing key domain-specific concepts [74]. We had tried to use entity linking tools (e.g., DBpedia spotlight [22] and TagMe [31]), to query post content so that key phrases could be linked to entities in the knowledge base which would have replaced the hypernym portion of the PPR graph. However, their performance seemed poor in our context: many key phrases were not linked to the correct knowledge graph entities. The main reason may be that forum posts present disambiguation challenges to entity linking tools [36]. Moreover, some knowledge bases, such as DBpedia [6], have a limited number of verb entities because most verbs are treated as relations. Thus, building a knowledge graph specifically for an individual course seems to be the only realistic approach even though it would require considerable effort.

Lastly, while the proposed recommender performed relatively well, the ability of this recommender to support socio-collaborative learning processes within discussion forums still needs to be validated through in-vivo studies. Due to the limitations of using historical data, the present evaluation does not allow the direct observation of how learners will respond to the recommendations nor does it allow the measurement of the recommendations’ effect on learning processes [28, 30, 51].

7. CONCLUSIONS

In this paper, a novel recommendation approach that accounts for socio-collaborative learning principles in small discussion forums was proposed. This multi-relational graph-based recommendation scheme, CSCLRec, incorporates social network analysis, learner categorization, and natural language processing techniques. A similarly structured recommender, CoPPR, was also introduced for potential use in socio-collaborative learning contexts.

The performance of these proposed algorithms was evaluated in an offline experiment where they were compared against six other recommendation algorithms. The results from this evaluation show our approaches outperform others. Going beyond these measures, we discussed CSCLRec’s potential to help socio-collaborative learning processes, as well as its use cases and potential expansions from the perspective of a variety of measures (e.g., precision, diversity) and learning goals. As future work, we plan to deploy the system to examine its influence on student behaviors and learning.

8. ACKNOWLEDGMENTS

We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada [RGPIN-2018-03834].

9. REFERENCES

[1] Abel, F., Bittencourt, I.I., Costa, E., Henze, N., Krause, D. and Vassileva, J. 2010. Recommendations in online discussion forums for e-learning systems. *IEEE Transactions on Learning Technologies*. 3, 2 (2010), 165–176. DOI:https://doi.org/10.1109/TLT.2009.40.

[2] Abel, F., Bittencourt, I.I., Henze, N., Krause, D. and Vassileva, J. 2008. A Rule-Based Recommender System for Online Discussion Forums. *Adaptive Hypermedia and Adaptive Web-Based Systems* (Berlin, Heidelberg, 2008), 12–21.

[3] Akyol, Z. and Garrison, D.R. 2008. The development of a community of inquiry over time in an online course: Understanding the progression and integration of social, cognitive and teaching presence. *Journal of Asynchronous Learning Networks*. 12, (2008), 3–22.

[4] Albatayneh, N.A., Ghauth, K.I. and Chua, F.-F. 2018. Utilizing Learners' Negative Ratings in Semantic Content-based Recommender System for e-Learning Forum. *Journal of Educational Technology & Society*. 21, 1 (2018), 112–125.

[5] Andresen, M.A. 2009. Asynchronous discussion forums: Success factors, outcomes, assessments, and limitations. *Journal of Educational Technology & Society*. 12, 1 (2009).

[6] Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R. and Ives, Z. 2007. Dbpedia: A nucleus for a web of open data. *The semantic web*. Springer. 722–735.

[7] Backenköhler, M., Scherzinger, F., Singla, A. and Wolf, V. 2018. Data-Driven Approach towards a Personalized Curriculum. *Proceedings of the 11th International Conference on Educational Data Mining, EDM 2018* (2018), 246–251.

[8] Beham, G., Kump, B., Ley, T. and Lindstaedt, S. 2010. Recommending knowledgeable people in a work-integrated learning system. *Procedia Computer Science*. 1, 2 (2010), 2783–2792. DOI:https://doi.org/10.1016/j.procs.2010.08.003.

[9] Bergmeir, C. and Benítez, J.M. 2012. On the use of cross-validation for time series predictor evaluation. *Information Sciences*. 191, (2012), 192–213. DOI:https://doi.org/10.1016/j.ins.2011.12.028.

[10] Bobadilla, J., Ortega, F., Hernando, A. and Bernal, J. 2012. A collaborative filtering approach to mitigate the new user cold start problem. *Knowledge-based systems*. 26, (2012), 225–238.

[11] Breese, J.S., Heckerman, D. and Kadie, C. 1998. Empirical Analysis of Predictive Algorithms for Collaborative Filtering. *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence* (San Francisco, CA, USA, 1998), 43–52.

[12] Buder, J. and Schwind, C. 2012. Learning with personalized recommender systems: A psychological view. *Computers in Human Behavior*. 28, 1 (Jan. 2012), 207–216. DOI:https://doi.org/10.1016/j.chb.2011.09.002.

[13] Carini, R.M., Kuh, G.D. and Klein, S.P. 2006. Student engagement and student learning: Testing the linkages. *Research in Higher Education*. 47, 1 (Feb. 2006), 1–32. DOI:https://doi.org/10.1007/s11162-005-8150-9.

[14] Castells, P., Hurley, N.J. and Vargas, S. 2015. Novelty and diversity in recommender systems. *Recommender systems handbook*. Springer. 881–918.

[15] Catherine, R. and Cohen, W. 2016. Personalized Recommendations Using Knowledge Graphs: A Probabilistic Logic Programming Approach. *Proceedings of the 10th ACM Conference on Recommender Systems* (New York, NY, USA, 2016), 325–332.

[16] Cer, D., Yang, Y., Kong, S., Hua, N., Limtiaco, N., John, R.S., Constant, N., Guajardo-Cespedes, M., Yuan, S., Tar, C. and others 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*. (2018).

[17] Chen, B., Chang, Y.H., Ouyang, F. and Zhou, W. 2018. Fostering student engagement in online discussion through social learning analytics. *Internet and Higher Education*. 37, (2018), 21–30. DOI:https://doi.org/10.1016/j.iheduc.2017.12.002.

[18] Chen, W. and Persen, R. 2009. A recommender system for collaborative knowledge. *Proceedings of the 2009 conference on Artificial Intelligence in Education: Building Learning Systems that Care: From Knowledge Representation to Affective Modelling* (2009), 309–316.

[19] Cheng, C.K., Paré, D.E., Collimore, L.-M. and Joordens, S. 2011. Assessing the effectiveness of a voluntary online discussion forum on improving students' course performance. *Computers & Education*. 56, 1 (2011), 253–261. DOI:https://doi.org/https://doi.org/10.1016/j.compedu.2010.07.024.

[20] Coxhead, A. 2000. A New Academic Word List. *TESOL Quarterly*. 34, 2 (2000), 213. DOI:https://doi.org/10.2307/3587951.

[21] Cremonesi, P., Koren, Y. and Turrin, R. 2010. Performance of recommender algorithms on top-N recommendation tasks. *RecSys'10 - Proceedings of the 4th ACM Conference on Recommender Systems* (New York, New York, USA, 2010), 39–46.

[22] Daiber, J., Jakob, M., Hokamp, C. and Mendes, P.N. 2013. Improving Efficiency and Accuracy in Multilingual Entity Extraction. *Proceedings of the 9th International Conference on Semantic Systems (I-Semantics)* (2013).

[23] Demmans Epp, C., Phirangee, K. and Hewitt, J. 2017. Student actions and community in online courses: The roles played by course length and facilitation method. *Online Learning Journal*. 21, 4 (2017).

[24] Demmans Epp, C., Phirangee, K. and Hewitt, J. 2017. Talk with Me: Student Behaviours and Pronoun Use as Indicators of Discourse Health across Facilitation Methods. *Journal of Learning Analytics*. 4, 3 (Dec. 2017), 47–75. DOI:https://doi.org/10.18608/jla.2017.43.4.

[25] Dowell, N.M.M., Brooks, C., Kovanović, V., Joksimović, S. and Gašević, D. 2017. The changing patterns of MOOC discourse. *Proceedings of the Fourth (2017) ACM Conference on Learning@ Scale* (2017), 283–286.

[26] Dowell, N.M.M., Skrypnyk, S., Joksimović, S., Graesser, A., Dawson, S., Gašević, D., Hennis, T.A., Vries, P. de

and Kovanović, V. 2015. Modeling Learners ’ Social Centrality and Performance through Language and Discourse. *Proceedings of the 8th International Conference on Educational Data Mining, EDM 2015* (2015), 250–257.

[27] Drachsler, H., Verbert, K., Santos, O.C. and Manouselis, N. 2015. Panorama of Recommender Systems to Support Learning. *Recommender Systems Handbook*. F. Ricci, L. Rokach, and B. Shapira, eds. Springer US. 421–451.

[28] Erdt, M., Fernández, A. and Rensing, C. 2015. Evaluating Recommender Systems for Technology Enhanced Learning: A Quantitative Survey. *IEEE Transactions on Learning Technologies*. 8, 4 (2015), 326–344. DOI:https://doi.org/10.1109/TLT.2015.2438867.

[29] Esteban, A., Zafra, A. and Romero, C. 2018. A Hybrid Multi-Criteria Approach Using a Genetic Algorithm for Recommending Courses to University Students. *Proceedings of the 11th International Conference on Educational Data Mining, EDM 2018* (2018), 273–279.

[30] Fazeli, S., Drachsler, H., Bitter-Rijpkema, M., Brouns, F., Van Der Vegt, W. and Sloep, P.B. 2018. User-Centric Evaluation of Recommender Systems in Social Learning Platforms: Accuracy is Just the Tip of the Iceberg. *IEEE Transactions on Learning Technologies*. 11, 3 (2018), 294–306. DOI:https://doi.org/10.1109/TLT.2017.2732349.

[31] Ferragina, P. and Scaiella, U. 2010. Tagme: on-the-fly annotation of short text fragments (by wikipedia entities). *Proceedings of the 19th ACM international conference on Information and knowledge management* (2010), 1625–1628.

[32] Fosnot, C.T. and Perry, R.S. 1996. Constructivism: A psychological theory of learning. *Constructivism: Theory, perspectives, and practice*. 2, (1996), 8–33.

[33] García, E., Romero, C., Ventura, S. and Castro, C. De 2009. An architecture for making recommendations to courseware authors using association rule mining and collaborative filtering. *User Modeling and User-Adapted Interaction*. 19, 1-2 SPEC. ISS. (2009), 99–132. DOI:https://doi.org/10.1007/s11257-008-9047-z.

[34] Greer, J., McCalla, G., Cooke, J., Collins, J., Kumar, V., Bishop, A. and Vassileva, J. 1998. The Intelligent Helpdesk: Supporting Peer-Help in a University Course. *International Conference on Intelligent Tutoring Systems*. Springer Verlag. 494–503.

[35] Haveliwala, T.H. 2002. Topic-sensitive PageRank. *Proceedings of the 11th International Conference on World Wide Web, WWW ’02* (New York, New York, USA, 2002), 517–526.

[36] Heitmann, B. and Hayes, C. 2010. Using linked data to build open, collaborative recommender systems. *AAAI Spring Symposium - Technical Report*. SS-10-07, October 2010 (2010), 76–81.

[37] Hew, K.F. and Cheung, W.S. 2012. *Student participation in online discussions: Challenges, solutions, and future research*. Springer Science & Business Media.

[38] Hewitt, J. 2005. Toward an understanding of how threads die in asynchronous computer conferences. *Journal of the Learning Sciences*. 14, 4 (2005), 567–589. DOI:https://doi.org/10.1207/s15327809jls1404_4.

[39] Hmelo-Silver, C. 2013. *The International Handbook of Collaborative Learning*. Routledge.

[40] Hu, Y., Koren, Y. and Volinsky, C. 2008. Collaborative Filtering for Implicit Feedback Datasets. *2008 Eighth IEEE International Conference on Data Mining* (Dec. 2008), 263–272.

[41] Ishola, O.M. and McCalla, G. 2017. Predicting prospective peer helpers to provide just-in-time help to users in question and answer forums. *Proceedings of the 10th International Conference on Educational Data Mining, EDM 2017* (2017), 238–243.

[42] Jeong, H. and Hmelo-Silver, C.E. 2016. Seven Affordances of Computer-Supported Collaborative Learning: How to Support Collaborative Learning? How Can Technologies Help? *Educational Psychologist*. 51, 2 (Apr. 2016), 247–265. DOI:https://doi.org/10.1080/00461520.2016.1158654.

[43] Jones, D., Bench-Capon, T. and Visser, P. 1998. Methodologies for Ontology Development. *Proc. IT&KNOWS Conference of the 15th IFIP World Computer Congress*. (1998).

[44] Kamii, C. 1986. The Equilibration of Cognitive Structures: The Central Problem of Intellectual Development . Jean Piaget , Terrance Brown , Kishore Julian Thampy . *American Journal of Education*. 94, 4 (Aug. 1986), 574–577. DOI:https://doi.org/10.1086/443876.

[45] Kleanthous Loizou, S. and Dimitrova, V. 2013. Adaptive notifications to support knowledge sharing in close-knit virtual communities. *User Modeling and User-Adapted Interaction*. 23, 2–3 (2013), 287–343. DOI:https://doi.org/10.1007/s11257-012-9127-y.

[46] Kleanthous, S. and Dimitrova, V. 2010. Analyzing community knowledge sharing behavior. *International Conference on User Modeling, Adaptation, and Personalization* (2010), 231–242.

[47] Kleanthous, S. and Dimitrova, V. 2009. Detecting changes over time in a knowledge sharing community. *Proceedings - 2009 IEEE/WIC/ACM International Conference on Web Intelligence, WI 2009*. 1, (2009), 100–107. DOI:https://doi.org/10.1109/WI-IAT.2009.21.

[48] Kleanthous, S. and Dimitrova, V. 2008. Modelling Semantic Relationships and Centrality to Facilitate Community Knowledge Sharing. *Adaptive Hypermedia and Adaptive Web-Based Systems*. Springer Berlin Heidelberg. 123–132.

[49] Lan, A.S., Spencer, J.C., Chen, Z., Brinton, C.G. and Chiang, M. 2019. Personalized Thread Recommendation for MOOC Discussion Forums. *Machine Learning and Knowledge Discovery in Databases* (Cham, 2019), 725–740.

[50] Lesk, M. 1986. Automatic sense disambiguation using machine readable dictionaries: How to tell a pine cone from an ice cream cone. *Proceedings of the 5th Annual International Conference on Systems Documentation, SIGDOC 1986* (New York, New York, USA, Jun. 1986), 24–26.

[51] Manouselis, N., Drachsler, H., Verbert, K. and Duval, E. 2012. *Recommender Systems for Learning*. Springer Publishing Company, Incorporated.

[52] McMillan, C., Poshyvanyk, D. and Grechanik, M. 2010. Recommending source code examples via api call usages and documentation. *Proceedings of the 2nd International Workshop on Recommendation Systems for Software Engineering* (2010), 21–25.

[53] Mi, F. and Faltings, B. 2017. Adaptive Sequential Recommendation for Discussion Forums on MOOCs using Context Trees. *Proceedings of the 10th International Conference on Educational Data Mining*. (2017), 24–31.

[54] Mihalcea, R. and Tarau, P. 2004. TextRank: Bringing Order into Text. *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing* (Barcelona, Spain, 2004), 404–411.

[55] Miller, G.A. 1995. WordNet: A Lexical Database for English. *Communications of the ACM*. 38, 11 (Nov. 1995), 39–41. DOI:https://doi.org/10.1145/219717.219748.

[56] Musto, C., Basile, P., Lops, P., de Gemmis, M. and Semeraro, G. 2017. Introducing linked open data in graph-based recommender systems. *Information Processing & Management*. 53, 2 (2017), 405–435. DOI:https://doi.org/https://doi.org/10.1016/j.ipm.2016.12.003.

[57] Musto, C., Lops, P., Basile, P., de Gemmis, M. and Semeraro, G. 2016. Semantics-Aware Graph-Based Recommender Systems Exploiting Linked Open Data. *Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization* (New York, NY, USA, 2016), 229–237.

[58] Musto, C., Semeraro, G., de Gemmis, M. and Lops, P. 2017. Tuning Personalized PageRank for Semantics-Aware Recommendations Based on Linked Open Data. *The Semantic Web* (Cham, 2017), 169–183.

[59] Nguyen, P.T., Tomeo, P., Di Noia, T. and Di Sciascio, E. 2015. An evaluation of simrank and personalized pagerank to build a recommender system for the web of data. *WWW 2015 Companion - Proceedings of the 24th International Conference on World Wide Web* (2015), 1477–1482.

[60] Page, L., Brin, S., Motwani, R. and Winograd, T. 1998. The PageRank Citation Ranking: Bringing Order to the Web. *World Wide Web Internet And Web Information Systems*. 54, 1999–66 (1998), 1–17. DOI:https://doi.org/10.1.1.31.1768.

[61] Palmer, S., Holt, D. and Bray, S. 2008. Does the discussion help? the impact of a formally assessed online discussion on final student results. *British Journal of Educational Technology*. 39, 5 (Sep. 2008), 847–858. DOI:https://doi.org/10.1111/j.1467-8535.2007.00780.x.

[62] Park, Y.J. and Tuzhilin, A. 2008. The long tail of recommender systems and how to leverage it. *RecSys'08: Proceedings of the 2008 ACM Conference on Recommender Systems* (New York, New York, USA, 2008), 11–18.

[63] Pereira, C.K., Campos, F., Ströele, V., David, J.M.N. and Braga, R. 2018. BROAD-RSI – educational recommender system using social networks interactions and linked data. *Journal of Internet Services and Applications*. 9, 1 (2018), 7. DOI:https://doi.org/10.1186/s13174-018-0076-5.

[64] Phirangee, K., Demmans Epp, C. and Hewitt, J. 2016. Exploring the Relationships between Facilitation Methods, Students' Sense of Community, and Their Online Behaviors. *Online Learning*. 20, 2 (2016), 134–154.

[65] Polyzou, A., Athanasios, N. and Karypis, G. 2019. Scholars Walk: A Markov Chain Framework for Course Recommendation. *Proceedings of the 12th International Conference on Educational Data Mining* (2019), 396–401.

[66] Rosé, C.P. and Ferschke, O. 2016. Technology Support for Discussion Based Learning: From Computer Supported Collaborative Learning to the Future of Massive Open Online Courses. *International Journal of Artificial Intelligence in Education*. 26, 2 (Jun. 2016), 660–678. DOI:https://doi.org/10.1007/s40593-016-0107-y.

[67] Rose, S., Engel, D., Cramer, N. and Cowley, W. 2010. Automatic Keyword Extraction from Individual Documents. *Text Mining: Applications and Theory*. John Wiley and Sons. 1–20.

[68] Santos, O.C., Boticario, J.G. and Pérez-Marín, D. 2014. Extending web-based educational systems with personalised support through User Centred Designed recommendations along the e-learning life cycle. *Science of Computer Programming*. 88, (Aug. 2014), 92–109. DOI:https://doi.org/10.1016/j.scico.2013.12.004.

[69] Scardamalia, M. 2002. Collective cognitive responsibility for the advancement of knowledge. *Liberal education in a knowledge society*. (2002), 67–98.

[70] Scardamalia, M. and Bereiter, C. 1994. Computer Support for Knowledge-Building Communities. *Journal of the Learning Sciences*. 3, 3 (Jul. 1994), 265–283. DOI:https://doi.org/10.1207/s15327809jls0303_3.

[71] Scardamalia, M. and Bereiter, C. 2006. Knowledge Building: Theory, Pedagogy, and Technology. *The Cambridge handbook of: The learning sciences.* Cambridge University Press. 97–115.

[72] Schellens, T. and Valcke, M. 2006. Fostering knowledge construction in university students through asynchronous discussion groups. *Computers and Education*. 46, 4 (2006), 349–370. DOI:https://doi.org/10.1016/j.compedu.2004.07.010.

[73] Shani, G. and Gunawardana, A. 2011. Evaluating Recommendation Systems. *Recommender Systems Handbook*. Springer US. 257–297.

[74] Shen, W., Wang, J. and Han, J. 2015. Entity Linking with a Knowledge Base : *IEEE Transactions on Knowledge and Data Engineering*. 27, 2 (2015), 443–460. DOI:https://doi.org/10.1109/TKDE.2014.2327028.

[75] Smith, L. and Macgregor, J.T. 1992. What is Collaborative Learning ? *Assessment*. 117, 5 (1992), 10–30.

[76] Smyth, B. and McClave, P. 2001. Similarity vs. diversity. *International conference on case-based reasoning* (2001), 347–361.

[77] Steck, H. 2011. Item popularity and recommendation accuracy. *RecSys'11 - Proceedings of the 5th ACM Conference on Recommender Systems* (New York, New York, USA, 2011), 125–132.

[78] Thai-Nghe, N., Drumond, L., Horváth, T., Krohn-Grimberghe, A., Nanopoulos, A. and Schmidt-Thieme, L. 2012. Factorization techniques for predicting student performance. *Educational recommender systems and technologies: Practices and challenges*. IGI Global. 129–153.

[79] Vassileva, J. 2008. Toward Social Learning Environments. *IEEE Transactions on Learning Technologies*. 1, 4 (2008), 199–214. DOI:https://doi.org/10.1109/TLT.2009.4.

[80] Vassileva, J., McCalla, G.I. and Greer, J.E. 2016. From Small Seeds Grow Fruitful Trees: How the PHelpS Peer Help System Stimulated a Diverse and Innovative Research Agenda over 15 Years. *International Journal of Artificial Intelligence in Education*. 26, 1 (2016), 431–447. DOI:https://doi.org/10.1007/s40593-015-0073-9.

[81] Veldhuis-Diermanse, A.E. 2002. *Participation, learning activities and knowledge construction in computer-supported collaborative learning in higher education*.

[82] Verbert, K., Manouselis, N., Ochoa, X., Wolpers, M., Drachsler, H., Bosnic, I. and Duval, E. 2012. Context-aware recommender systems for learning: A survey and future challenges. *IEEE Transactions on Learning Technologies*.

[83] Vygotsky, L.S. 1978. *Mind in society: Development of higher psychological processes.* Harvard University Press.

[84] Webb, E., Jones, A., Barker, P. and Van Schaik, P. 2004. Using e-learning dialogues in higher education. *Innovations in Education and Teaching International*. 41, 1 (2004). DOI:https://doi.org/10.1080/1470329032000172748.

[85] Weber, K., Maher, C., Powell, A. and Lee, H.S. 2008. Learning opportunities from group discussions: Warrants become the objects of debate. *Educational Studies in Mathematics*. 68, 3 (2008), 247–261.

[86] Wise, A.F., Hausknecht, S.N. and Zhao, Y. 2014. Attending to others' posts in asynchronous discussions: Learners' online "listening" and its relationship to speaking. *International Journal of Computer-Supported Collaborative Learning*. 9, 2 (May 2014), 185–209. DOI:https://doi.org/10.1007/s11412-014-9192-9.

[87] Yang, D., Piergallini, M., Howley, I. and Rose, C.P. 2014. Forum Thread Recommendation for Massive Open Online Courses. *Proceedings of the 7th International Conference on Educational Data Mining*. (2014), 257–260. DOI:https://doi.org/https://www.cs.cmu.edu/~diyiy/docs/edm14_recom.pdf.

[88] Zhang, S., Yao, L., Sun, A. and Tay, Y. 2019. Deep Learning Based Recommender System: A Survey and New Perspectives. *ACM Comput. Surv.* 52, 1 (2019). DOI:https://doi.org/10.1145/3285029.