Facilitating Peer Learning and Knowledge Sharing in STEM Courses via Pattern Based Graph Visualization

Emilio Zegarra, Shi-Kuo Chang, Jingtao Wang

Department of Computer Science
University of Pittsburgh
Pittsburgh, PA United States
{ezegarra, chang, jingtao}@cs.pitt.edu

Abstract—**High quality education in Science, Technology, Engineering and Math (STEM) majors expects not only the acquisition of comprehensive domain knowledge, but also the mastery of skills to solve open ended and even ill-defined problems. Problem-based Learning (PBL) is usually adopted to achieve such goals. However, PBL requires sustained and in-depth faculty involvement, hence making PBL not scalable. Also, existing knowledge discovery techniques do not facilitate the capture and reuse of solutions to recurring problems. To address these challenges, we present MicroBrowser, an interactive Q&A system augmented with 2D discussion visualizations and pattern based expertise sharing interfaces. MicroBrowser allows learners to browse and navigate important discussions in PBL based on topic similarity encoded by node proximity in a knowledge graph. MicroBrowser also provides a set of pattern based expertise sharing interfaces to allow both learners and instructors to highlight, share and reuse major findings in PBL. Through a 32-subject study, we found MicroBrowser to be effective at facilitating knowledge discovery. Moreover, students understood and were able to use design patterns to complete open ended tasks.**

Education, knowledge, sharing, discovery, pattern, graph, visualization

I. INTRODUCTION

Work on improving education to students in STEM majors has allowed students to acquire important domain knowledge. Yet, upon graduation, students are faced with open ended and ill-defined problems for which they might not be prepared. Innovative techniques such as PBL aim to address these limitations by encouraging students to learn by addressing everyday problems [13].

An important aspect of PBL is the in-depth participation of faculty advisors [16]. However, for large faculty-to-student ratios it becomes difficult for faculty advisors to get involved with all students resulting in a reduced time and effort involvement and a longer feedback loop A technique that has been used in the context of online education via Massive Open Online Courses (MOOCs) to address this challenge is the use of discussion forums and Q&A systems [5][12]. These systems have been found to be a common source of knowledge for students when completing homework tasks [18][20] and for interacting with instructors[12]. While these systems can address the scalability problem of faculty involvement, they introduce new problems. First of all, as knowledge bases grow in size, the sheer size of the accumulated knowledge makes it
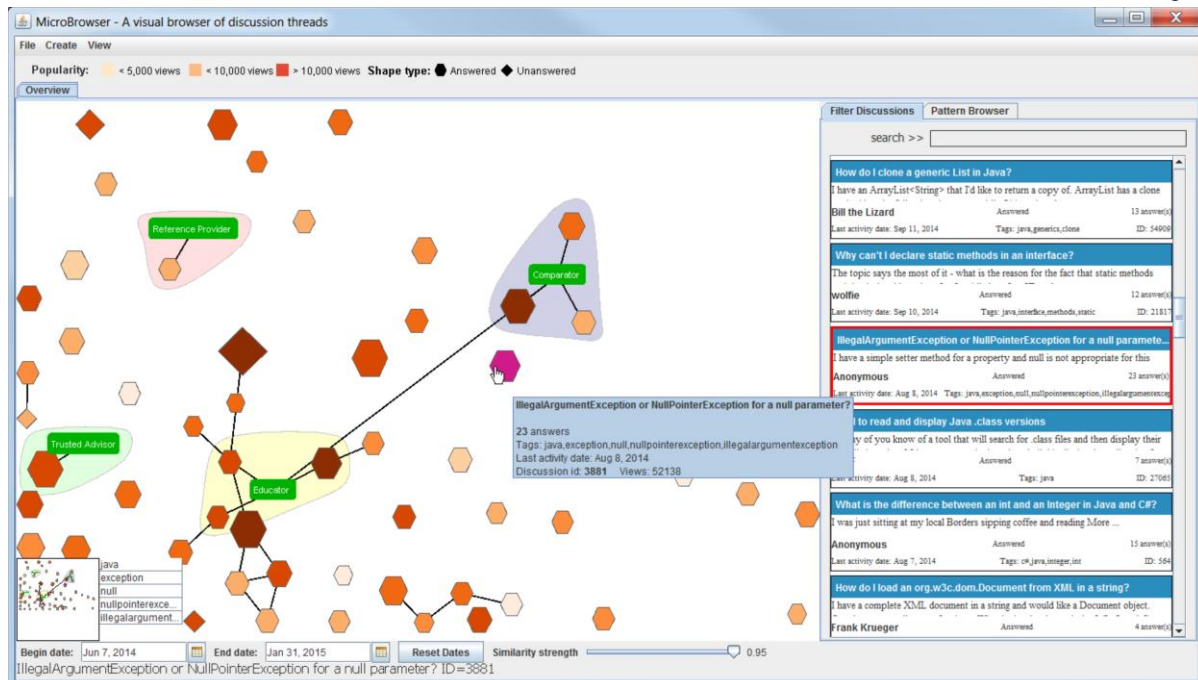

Figure 1. The primary UI of MicroBrowser displaying discussion threads and design patterns. MicroBrowser consists of three components: a knowledge exploration graph, a discussion listing view and a knowledge timeline

harder to locate the desired information. Secondly, existing knowledge discovery techniques do not provide effective means for the capture and reuse of solutions to recurring problems.

To address these challenges, we designed MicroBrowser (Fig. 1). MicroBrowser is a pattern based visualization system designed for educational settings that aims to improve peer-learning by facilitating knowledge discovery and reuse. Its contributions are the ability to discover knowledge from large online discussion boards by using visualization techniques as well as the introduction of design patterns as a method to not only discover knowledge but also to improve peer-learning by facilitating knowledge creation and reuse. In a 32-subject user study, we found that MicroBrowser facilitates the completion of knowledge discovery and reuse tasks when compared to traditional Q&A systems. We also found that for open ended tasks, students made use of design patterns to complete the tasks in over 50% of the time. Also, results from our subjective evaluation indicate students understood the concept of design patterns and found the features in MicroBrowser to be effective and useful. Finally, students found MicroBrowser to be easy to use.

II. RELATED WORK

A. Visualizing Discussion Forums

Discovering of information thru visualization is not limited to finding relevance between posts, but also about finding structure[19]. Several approaches have been used to address forum visualization. ForAvis[23] uses the "overview, zoom and filter" visualization technique as well as color coding to provide a layered approach to displaying information. In [7], the authors found that using treemaps, with color and size encoding, made finding largest and most active discussions faster than traditional text based interfaces. VIDI Toolbar [21] displayed Political Science discussion on a graph with node proximity encoding similarity and cluster of nodes representing topics. Anagora[10] visualized discussion activity along a horizontal chronogram axis. The length of the discussions represented activity and high activity in a forum was represented thru discussion overlap.
In MicroBrowser, we learned from these works on discussion forums and extended them to a Q&A system in an educational domain.

B. Using Design Pattern

Design patterns are a shared language used to communicate proven solutions to recurring problems. Christopher Alexander [1] crafted the notion of patterns to facilitate design and construction of towns and buildings.
The adaptability benefits of design patterns have allowed their extension to other contexts. Gamma et.al used patterns to document recurring problems and their solutions in object oriented programming[8]. Pedagogical patterns assist instructors in preparing effecting instructional material based on learned experience from other instructors[3][14]. Also, human computer interaction and user interface design patterns have been defined to assist with valuable UI design solutions [15][22].

MicroBrowser takes advantages of the benefits of design patterns and applies them to discussion threads in an educational domain. Different from pedagogical patterns, design patterns in MicroBrowser are intended to address the problem of answering recurring questions posted in Q&A systems. MicroBrowser also provides pattern browsing capabilities which allow students to discover patterns and reuse them by either associating them to discussion threads or using their solutions to answer similar questions.

III. THE DESIGN OF MICROBROWSER

The design goals of MicroBrowser are to: a) Identify, create and reuse design patterns in discussions and b) Make it easier to visualize and browse discussion threads and associated patterns.
MicroBrowser was implemented on Java using the Prefuse[11] library.

A. Data Processing and Graph Generation

Topic Modeling facilitates the analysis of large volumes of unstructured text. We use the topic models information to quantify similarities between discussions. For our corpus, we created a text document for each discussion by combining the question title, question text, tags and the answers' text. Stop words were removed from the documents. We used the Mallet [17] topic modeling toolkit to perform topic analysis and build our topic model. The initial model parameters were set to the default values except for the maximum number of topics which was set to 100 and the maximum iteration was set to 2000. Using the topic model, we created a dissimilarity matrix between each discussion. We first used (1) to compute the similarity between two discussions and (2) for their dissimilarity and to populate our dissimilarity matrix.

$$similarity(i,j) = \frac{similar\_topics[i][j]}{count\_topics[i]} \qquad (1)$$

$$dissimilarity(i,j) = 1 - similarity(i,j) \qquad (2)$$

As in [21], we then applied Multi-Dimensional Scaling (MDS) to the dissimilarity matrix to reduce the discussions to a 2-dimension space for plotting. We stored the resulting similarity value between discussions separately to be used when retrieving related discussions. Finally, we used the GraphML[1] format to describe the structural properties of the knowledge base with nodes representing discussions and edges representing relationship. For each discussion, we created a connection to each of its top 20 most similar discussions. For the case of design patterns, we build a connection to each associated discussion. For each edge created between two nodes, we stored the calculated similarity value, except for design pattern edges where we set the value to 1.

---

[1] http://graphml.graphdrawing.org

B. Knowledge Graph View

Similar to [23], we displayed discussions in an overview graph. Thru visualization techniques, students can use their visual abilities to discover relationships, structures or patterns in discussion threads.

We used the shape of a discussion node to encode whether a discussion was either answered or unanswered, with answered ones displayed as a hexagon and unanswered ones displayed as diamond. Having the name of the pattern available makes it easier to identify them. Thus, we used a green label node to display design patterns. To highlight associations to design patterns, we build a convey circle around those discussion nodes belonging to the design pattern and centered the design pattern node among them.

We encoded popularity using a gradient color scheme similar to [2][23] with darker shades denoting more popular items.

A connection between a discussion node and a pattern node indicates that a discussion node is an instance of the pattern. A connection between two discussion nodes indicates that the two discussions are related to each other. A challenge we faced with the number of connections between nodes was that it was cluttering the drawing area making it unreadable. We address this by using progressive disclosure [4] and used the similarity strength value assigned to each connection or edge between two discussions to determine the color intensity of the similarity with more similar items darker. We also added a similarity slider to allow students to control the visibility of the nodes based on the similarity strength.

C. Discussion Listing View

Traditional text based browsing schemes provide simple and effective means for displaying textual information that cannot be represented in a graph node. The discussion listing view provides a text based browsing scheme to allow students to access more information about the discussions that would have been difficult to include in the graph. Since we provided two
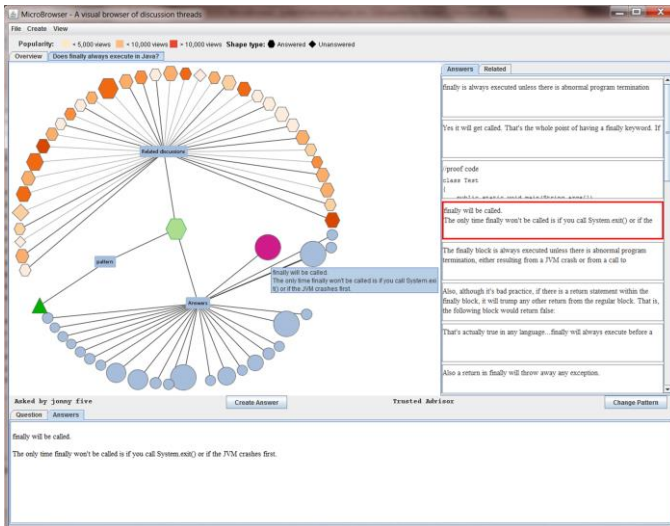

Figure 2. Discussion details view providing question details, answers, related discussions and design patterns. In this figure, a student has selected an answer. The selected answer is brought into focus on the answers list to the right and its content is displayed in the *Answers* tab below.
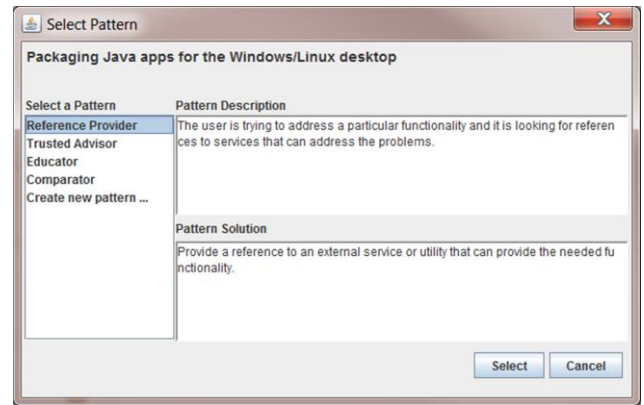

Figure 3. Pattern selection dialog allows students to select existing or create new design patterns.

different browsing approaches (graph and text-based), it was important to ensure proper integration between the two. We enabled bi-directional synchronization between the graph and the listing view. As students moved the mouse over a node in the graph, the corresponding entry in the discussion listing view was brought into focus. Likewise, if a student was more comfortable browsing discussions in textual form, as they moved the mouse over the discussions the corresponding node in the graph got highlighted.

The discussion view also provides keyword filtering to allow users to filter out results.

D. Knowledge Timeline

Only providing keyword filtering might not be enough. In [10][9], the authors allowed users to browse discussions thru their temporal attributes. In an educational setting, time information is very important as discussions can be associated around quizzes or deadlines. To support such scenarios, the knowledge timeline allows students to filter discussions by specifying starting and ending dates. With the timeline in combination with the keyword filters, students can further narrow the available discussions.

E. Discussion Details View

For our discussions, we wanted to represent an organization of answers, related discussions and associated design patterns. In [2], the authors found that using a tree layout provides a simple way to represent perceptual organization. We extended the approach and built a hierarchy for answers, related discussions and design patterns (Fig. 2). The discussion detail view extends the color and size encodings from the knowledge graph view. Textual content for the question and answers was available from the corresponding tab.

F. Associating Design Patterns

When students open up the pattern selection dialog, Fig. 3, they are presented with a list of available patterns they can select from. Selecting a pattern displays its description and the solution. The description helps the student to identify the appropriate question type being evaluated. The solution provides the student with a recommendation for how to answer the type of question.

Once a pattern is selected, the discussion becomes an instance of that pattern. If students need to find samples of how similar questions were answered, they can look at other instances of discussions associated to the selected pattern.

Finally, [6] found that tools that manage pattern collections should be able store and organize the patterns for easy access and exploration. The Pattern Browser view provides such functionality and it is available from the main navigation window along with the discussion listing view.

IV. USER STUDY

We conducted a 32-subject user study of MicroBrowser to find out:

- Does MicroBrowser promote peer-learning and knowledge discovery? Were students able to discover and reuse knowledge efficiently?

- Did students understand the concept of design patterns? Were design patterns helpful for completing knowledge discovery and reuse?

- Were the design features of color coding, node sizes, keyword and timeline search and discussion overview effective and easy to understand?

- Is MicroBrowser easy to use?

A. Experimental design

To validate the effectiveness of MicroBrowser (MB), we compared it against a state-of-the-art Q&A system (Q&A). Subjects were randomly assigned to start with either MB or Q&A. The user study consisted of three parts. First, students completed a brief tutorial and were asked to complete a total of 12 tasks. After each task, students rated the perceived difficulty of completing the task. Then students completed 12 similar tasks using the other system. There was no time limit for completing these tasks. Finally, students completed an exit questionnaire.

B. Task Descriptions

We defined 12 tasks around 4 scenarios:

1) Discovering knowledge using keywords
Task 1 asked students to find a question about a given keyword. Task 2 asked students to find the most popular discussion for a given keyword. Task 3 asked students to first find an unanswered question about a given keyword and then find related discussion suggested by the system.

2) Discovering knowledge using timeline
Task 4 asked students to find any discussion that occurred during a given timeframe. Task 5 asked students to find discussions about given keyword that occurred during a given timeframe. Task 6 asked students to find the most recent unanswered question.

3) Using Patterns for knowledge reuse
Task 7 asked students to find examples of discussions whose answer either provide recommendations or advise or suggested the use of references. Task 8 asked students to modify a given discussion such that other students looking for similar approaches to answering the question could reference it. Task

9 asked students to submit a given answer to an unanswered discussion. Then they were asked to modify the discussion such that other students could reuse how the answer to the question was presented.

4) Using Patterns for knowledge discovery
Task 10 asked students to find discussions based on the approach taken to answer them. Task 11 asked students to refer to a particular discussion and then find other discussions whose approach to answering it was similar. Task 12 asked students how they can find recommended ways to answer discussions.

C. Participants and Apparatus

We recruited 32 subjects (6 female) over 18 years (18-21: 38%, 22-25: 41%, 26+: 22%) among STEM major students from a local university. Each session lasted 1 hour 30 minutes and participants were compensated with a $10 Amazon.com gift card.

Participants used a Dell Optiplex 745 (Intel Core2 Duo T6300 1.86GHz, 2GM RAM) using Windows 8 and a 19" display at a screen resolution of 1280x1024. MB ran on top of Eclipse V4.2.2 and Java SE 7u51.

As our baseline Q&A system, we configured an instance of the open source, PHP based Question2Answer[2] platform and access it via Internet Explorer V10.

D. Data

We loaded the system with the 500 most recent discussion threads and their answers retrieved from StackOverflow[3] associated to the tag 'Java'. This data consisted of 4041 records (500 questions, 3541 answers) from 2091 different users. We created 4 initial design patterns and assigned them sample discussions.

E. Usage Behavior

Table 1 summarizes the activities participants performed the most while completing the tasks using MB. We noted that participants visited a median of 182 nodes using the knowledge graph but only 12 nodes using the discussion listing view suggesting a preference for browsing knowledge using the knowledge graph over the discussion listing view. Participants opened an average of 15 discussions and 8 design patterns suggesting that integrating discussion details in the knowledge graph reduced the need to open discussions for information. When students opened the details of a

TABLE 1. DESCRIPTIVE STATISTICS SUMMARIZING PARTICIPANTS' ACTIVITIES USING MICROBROWSER (SD=STD DEV).

Activities	Median	Mean	SD
Open discussion details	15.0	15.37	7.66
Open design pattern details	8.0	8.12	3.88
Open answer details	56.5	81.50	76.49
Perform keyword searches	13.0	13.81	5.14
Visit node on knowledge graph	179.0	182.56	72.20
Visit node on discussion listing	8.0	12.28	13.77

discussion, they could explore related discussions and answers associated to the discussion Data shows that on average participants viewed the details of about 50 answers (SD = 76.49) suggesting that MB facilitated knowledge exploration not only of discussion but also of their answers.   Finally, participants performed an average of 13 search events.

F. Perception of Difficulty

1) Discovering knowledge using keywords

Results show no statistically significant differences between Q&A and MB for Task 1 (p = 0.8514). For Tasks 2 and 3, results show statistically significant differences between the two methods. Data shows that MB is easier to use for completing those tasks (Table 2).

2) Discovering knowledge using timeline

For Tasks 4-6 results show statistically significant differences between Q&A and MB. Data show Task 4 and 5 were easier to complete using MB while for Task 6 Q&A was easier (Table 2).

TABLE 2. STUDENTS' PERCEPTION OF DIFFICULTY FOR TASKS 1-6. SCALE OF 1-EASIER TO 5-MORE DIFFICULT. (GRAY BACKGROUND = P < 0.05, WILCOXON SIGNED-RANK TEST, IQR=INTERQUARTILE RANGE, SD=STD DEV)

Task	Q&A		MB	
	Median (IQR)	Mean (SD)	Median (IQR)	Mean (SD)
1	1 (0)	1.25 (0.56)	1 (0.25)	1.28 (0.52)
2	2 (1.15)	2.51 (1.15)	**2 (0.59)**	**1.68 (0.59)**
3	3 (1.09)	3.00 (0.83)	**1 (0.83)**	**1.67 (0.83)**
4	3 (1.14)	3.03 (1.14)	**1 (0.67)**	**1.50 (0.67)**
5	2 (1.10)	2.25 (1.10)	**1 (0.56)**	**1.46 (0.56)**
6	**1 (1.00)**	**1.65 (1.00)**	3 (0.98)	2.93 (0.98)

3) Using Patterns for knowledge reuse

No statistically significant differences were found for Task 8 (p=0.0675) or Task 9 (p=0.1164) between the two systems. However, data shows a trend towards MB being easier to complete the tasks. Results show MB was easier to complete Task 7.

4) Using Patterns for knowledge discovery

We found statistically significant differences between Q&A and MB for Tasks 10-12. Students found using MB to be less difficult to complete these tasks.

TABLE 3. STUDENTS' PERCEPTION OF DIFFICULTY FOR TASKS 7-12. SCALE OF 1-EASIER TO 5-MORE DIFFICULT. (GRAY BACKGROUND = P < 0.05, WILCOXON SIGNED-RANK TEST, IQR=INTERQUARTILE RANGE, SD=STD DEV)

Task	Q&A		MB	
	Median (IQR)	Mean (SD)	Median (IQR)	Mean (SD)
7	3 (1)	2.83 (1.00)	**2 (2)**	**2.09 (0.92)**
8	3 (2.5)	2.90 (1.37)	2 (1.25)	2.37 (1.23)
9	2.5 (2)	2.81 (1.40)	2 (1)	2.40 (1.04)
10	2 (1)	2.60 (1.16)	**2 (2)**	**2.00 (1.09)**
11	2 (1.5)	2.50 (1.29)	**2 (1)**	**1.76 (0.81)**
12	3 (1.37)	3.37 (1.37)	**2 (1.14)**	**2.30 (1.14)**

G. Qualitative Results

The majority of students found the use of design patterns effective and most importantly easy to learn (Fig. 4). This was an important finding as we wanted to validate students could understand their value. Students were ecstatic about their capability and usefulness as noted by the following comments:
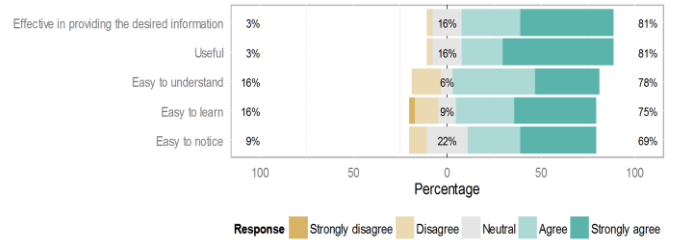

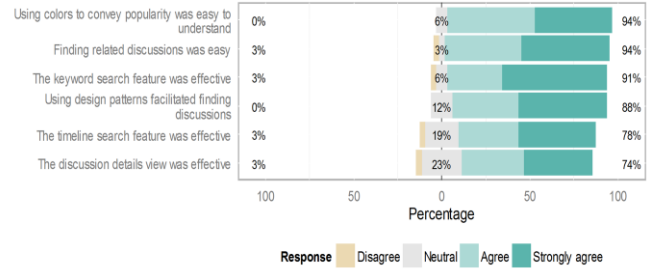Figure 4. Students' opinion about design patterns. 5-point Likert-scale

Figure 5. Students' opinion about MicroBrowser features and functionality. 5-point Likert-scale

"Pattern browsing is a great concept", "I really loved the pattern browser. Just a brilliant idea" and *"Pattern browser, it is very useful to see "trusted advisor". If I want to see reliable answers and "reference" and "comparator" are very useful."*

When directly asked about their feedback, students were very positive on the usefulness and effectiveness of the main features (Fig. 5).

When openly asked to list their most liked features, 56% students cited the integration of design patterns and the pattern browser, 47% the visual representation of discussions, 13% for visualizing relationships, 25% the use of color to encode popularity, 25% for shape encoding, and 22% the knowledge timeline.

Among the features students disliked the most were the initial learning curve (13%), the knowledge timeline (13%), the slow response (9%) and the use of patterns (9%).

Finally, students found MicroBrowser to be easy to use (10-point Likert-scale, median=8).

V.   DISCUSSION

In this study, we aimed to determine if our innovative system MB was easy to use and facilitated completion of the tasks when compared to a traditional Q&A. Based on the results, students found MB to be easier to complete most of the tasks when compared to Q&A. Task 1 was very simple which can explain why no statistically significant differences were found between the systems. For Tasks 2 and 3, students benefited in MB from the use of color coding to quickly identify popular discussions and the clear identification of the related discussions in the discussion details view. For Tasks 4 and 5, students benefited from using the knowledge timeline to narrow discussions as well as the ability to filter by both keyword and timeline simultaneously.   When using Q&A, students could only sort by recent discussions but had to page thru the discussions until reaching the desired date. Surprisingly, based on results, Q&A was easier for completing

Task 6. Q&A showed unanswered discussions already sorted by update date. This explains why, under Qualitative results, students recommended adding sorting capabilities to MB.

Even though Task 7 was open-ended, we wanted to see if students could use the systems to identify instances of reusable knowledge. Using MB, students were able to use the pattern browser and from there refer to the sample discussions. In the case of Q&A, students used different methods, such as performing keyword searches and then reading the details of identified discussions. Sometimes students simply opted for selecting the discussion that matched a keyword search. For Task 8, students benefited by the use of design patterns because 66% of the students used a design pattern to complete the task. For the case of Q&A, students tended to associate a tag to create a classification of the discussions. For Task 9-12, we noted a large percentage of students benefited from design patterns by using them when completing the tasks (50%, 81%, 78% and 77%, respectively). For Q&A, students had to find techniques for identifying recommended discussions such as using keywords or using page view count information.

Finally, qualitative results showed students found the MB features to be useful and effective. Students identified some usability issues and also made recommendations for improvements to MB. Among the recommendation were to add sorting capabilities, to improve the knowledge timeline by using a slider, to add filtering by users or question state, to associate more patterns to questions, to add actions on right-click menus and to design for accessibility.

VI. CONCLUSION

In this paper, we presented MicroBrowser, a system that facilitates peer learning and knowledge discovery in classroom settings. Results from our 32-subject user study show reduced difficulty at completing tasks when compared to traditional Q&A system. More importantly, students found benefit in the use of design patterns and found the system and its features effective and easy to use.

The integration of Design Patterns in knowledge discovery and generation is a key innovation in MicroBrowser. Overall, MicroBrowser achieves this with an interactive knowledge visualization and exploration system.

REFERENCES

[1] Alexander, C. The origins of pattern theory: The future of the theory, and the generation of a living world. *Software, IEEE 16*, 5 (1999), 71–82.

[2] Cai, Q. and Sheth, N. Visualizing MeSH Dataset using Radial Tree Layout. (2003).

[3] Carle, A., Clancy, M., and Canny, J. Working with pedagogical patterns in PACT. *ACM SIGCSE Bulletin 39*, 2007, 238.

[4] Chuang, J., Ramage, D., Manning, C.D., and Heer, J. Interpretation and Trust : Designing Model-Driven Visualizations for Text Analysis. (2012), 443–452.

[5] Coetzee, D., Fox, A., Hearst, M., and Hartmann, B. Should your MOOC forum use a reputation system? *CSCW'14*, (2014), 1176–1187.

[6] Deng, J., Kemp, E., and Todd, E.G. Managing UI pattern collections. *Proceedings of the 6th ACM SIGCHI New Zealand chapter's international conference on Computer-human interaction making CHI natural - CHINZ '05*, (2005), 31–38.

[7] Engdahl, B., Koksal, M., and Marsden, G. Using treemaps to visualize threaded discussion forums on PDAs. *Proceedings of ACM CHI 2005 Conference on Human Factors in Computing Systems*, (2005), 1355–1358.

[8] Gamma, E., Helm, R., Johnson, R., and Vlissides, J. *Design Patterns: Elements of Reusable Object-Oriented Software.* 1994.

[9] Gibbs, W.J., Olexa, V., and Bernas, R.S. A Visualization Tool for Managing and Studying Online Communications. *9*, (2006), 232–243.

[10] Giguet, E. and Lucas, N. Creating discussion threads graphs with Anagora. *CSCL'09 Proceedings of the 9th international conference on Computer supported collaborative learning*, (2009), 616–620.

[11] Heer, J., Card, S., and Landay, J. Prefuse: a toolkit for interactive information visualization. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM Press (2005), 421.

[12] Huang, J., Dasgupta, A., Ghosh, A., Manning, J., and Sanders, M. Superposter behavior in MOOC forums. *Proceedings of the first ACM conference on Learning @ scale conference - L@S '14*, (2014), 117–126.

[13] Hung, Woei (University of Arizona South, Sierra Vista, A., Jonassen, David H (University of Missouri, Columbia, M., and Liu, Rude (Beijing Normal University, Beijing, C. Problem-Based Learning. *Encyclopedia of the Sciences of Learning*, (2012), 2687–2690.

[14] Köppe, C. and Utrecht, H. A pattern language for teaching design patterns. *Proceedings of the 18th Conference on Pattern Languages of Programs - PLoP '11*, Part 2 (2011), 1–16.

[15] Lin, J. and Landay, J.A. Employing patterns and layers for early-stage design and prototyping of cross-device user interfaces. *Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI '08*, (2008), 13.

[16] Massie, D. and Massie, C. Framework for organization and control of capstone design/build projects. *Journal of STEM Education 7*, 3 (2006), 36–43.

[17] McCallum, A.K. MALLET: A Machine Learning for Language Toolkit. 2002. http://mallet.cs.umass.edu/.

[18] Seaton, D.T., Bergner, Y., Chuang, I., Mitros, P., and Pritchard, D.E. Who does what in a massive open online course? *Communications of the ACM 57*, 2014, 58–65.

[19] Seo, J., Croft, W.B., and Smith, D. a. Online community search using thread structure. *Proceeding of the 18th ACM conference on Information and knowledge management - CIKM '09*, (2009), 1907.

[20] Sonnino, E. 4 Ways to Get the Most out of a MOOC. *Learning Advisor*, 2013. http://blog.studentadvisor.com/four-ways-get-mooc/.

[21] Trampuš, M. and Grobelnik, M. Visualization of online discussion forums. *Workshop on Pattern Analysis Applications 11*, (2010), 134–141.

[22] Vanderdonckt, J. and Simarro, F.M. Generative pattern-based design of user interfaces. *Proceedings of the 1st International Workshop on Pattern-Driven Engineering of Interactive Computing Systems - PEICS '10*, (2010), 12–19.

[23] Wanner, F., Ramm, T., and Keim, D. ForAVis: explorative user forum analysis. *WIMS '11: Proceedings of the International Conference on Web Intelligence, Mining and Semantics*, (2011).