

Mastery Grids: An Open Source Social Educational Progress Visualization

Tomasz D. Loboda, Julio Guerra, Roya Hosseini, and Peter Brusilovsky

School of Information Sciences, University of Pittsburgh,
Pittsburgh, PA 15232, USA
{to17, jdg60, roh38, peterb}@pitt.edu

Abstract. While many pieces of educational software used in the classroom have been found to positively affect learning, they often are underused by students. Open learning model and social visualization are two approaches which have been helpful in ameliorating that low usage problem. This article introduces a fusion of these two ideas in a form of social progress visualization. A classroom evaluation indicates that this combination may be effective in engaging students, guiding them to suitable content, and enabling faster content access.

1 Introduction

Over the last 30 years researchers and practitioners have developed a range of advanced educational tools for computer science education (CSE) such as animations, simulations, programming problems, and self-assessment questions. While a number of these tools have been evaluated through user studies and proven to be useful, a limited success of many others demonstrated that educational effectiveness observed in the lab does not necessarily translate into a broad educational impact [1]. One of the reasons for this seems to be low usage by students. That is, even in a favorable situation when a good tool is adopted by a motivated instructor and made easily available to students its use may be low. That plague of low usage afflicts advanced CSE tools the most because the majority of these tools are crafted to support knowledge enhancement through self-study which makes it hard to mandate and control their use. Consequently, unlike some assessment-focused tools which students are required to use in order to complete their assignments (e.g., program grading tools or program-construction questions), self-study tools are typically provided as a recommended content and, as such, their use rarely counts towards the final grade. Moreover, the motivation to work with self-study content might be decreased when multiple content items are available (e.g., many self-assessment questions or animated code examples) because then it is even harder to choose the most appropriate item to work on. Technologies aiming at increasing student motivation or guiding them to the most appropriate self-study content could help to ameliorate this low usage problem.

Past research has identified several such technologies. For example, *learner knowledge visualization* (known as *open learner modeling*) which acknowledges

students progress and highlights gaps in their knowledge has been demonstrated to motivate student learning and guide students to the appropriate content [2]. The study of social visualization approaches based on the ideas of *social comparison* [3] has shown that the ability to compare oneself to one's peers or the entire class could increase a student's motivation to learn and improve their participation in educational activities [4, 5].

The work presented in this article proposes a progress visualization approach that combines the prospects of open student modeling and social visualization in the form of social progress visualization. To that effect, we have developed Mastery Grids, a visualization tool that combines personal and comparative progress visualization. Mastery Grids has been designed as a tunable open source educational interface and can work with several kinds of self-study content and can be a part of semester-long classes. The remainder of this article presents the idea and the implementation of Mastery Grids and reports the results of its evaluation in two different domains and two kinds of self-study content.

2 Relevant Work

The idea of an open learner model was originally explored in the area of personalized learning systems. While in traditional personalized systems, learner models were hidden “under the hood” and used to make the education process personalized, the pioneers of open modeling argued that the ability to view and modify the state of one's knowledge could be beneficial for students. Open learner models for a set of skills or topics are most frequently displayed as a set of progress bars (also known as *skillometers*) that visualize learner knowledge state understood either as the probability that a learner has mastered a skill [6] or as a fraction of total knowledge to be gained for a topic [7] (which is sometimes estimated as a fraction of material covered by that topic). Other forms of knowledge visualization (e.g., graphical or color-coded) have been explored as well [8, 9]. More recently, open learner models emerged as a feature of online learning systems which use some elements of progress modeling and personalization (e.g., Khan Academy Dashboard).

A range of benefits have been reported upon opening models to learners, such as increasing the learner's awareness of their developing knowledge, difficulties, and the learning process in general, as well as increasing students' engagement, motivation, and knowledge reflection [2, 7, 8]. Studies of both individual and group open learner models have shown an increase of reflection and helpful interactions among teammates. For more information, see Bull & Kay [2] who introduced a framework to apply open user models in adaptive learning environments and provided many in-depth examples.

3 Social Educational Progress Visualization (Mastery Grids)

Mastery Grid is out attempt to implement the ideas of social progress visualization in a generic, usable, and re-usable form. The most important component around which Mastery Grids is build is a *grid* which is a three-dimensional visualization component. The two geometric dimensions are topics defined in a domain (e.g., “Variables” or “Arrays”; the horizontal dimension) and resources which learners can access (e.g., questions, examples, readings, lecture recordings, etc.; the vertical dimension). The current classroom study used only two types of resources: Questions and examples. The third dimension is given by the intensity of the color in each cell of a grid and denotes the level of completeness, level of mastery, or level of progress a learner has for the given combination of topic and resource type. For example, the more saturated the color the more mastery the learner has achieved. This design allows learners to quickly get an idea about how they are doing in a course and where to allocate their efforts next.

Figure 1 shows the first of the two modes Mastery Grids can work in, the *All Resources* mode. In that mode, the “Me” grid shows progress the learner has made, the “Group” grid shows progress that the currently selected group has made, and the “Me vs group” grid shows the difference between the two grids. By manipulating the group the learner decides what they want to compare themselves to. For example, they could choose “class average” or “top 10 learners,” both of which were used in the current classroom study. Below the three main grids are individual learner grids. Each learner in the currently selected group is represented by one grid and these grids are ordered according to the average progress.

In the All Resources mode described above a learner sees the entirety of the material which is available to them (i.e., all topics and all resource types). However, a learner can choose to focus on one resource in particular (e.g., questions) and enter the *Resource Focus* mode (Figure 2). In that mode, the three main grids are fused into a single main grid. That single grid displays only the selected resource rows from the three main grids, i.e., the first row shows the learner, the third row shows the group, and the middle row shows the difference between the two. The Resource Focus mode makes the me-group types of comparisons easier by displaying the relevant information closer together. As a result of focusing on one resource type, individual learners from the current group are shown in a single grid as well. A learner can control the height of rows of that grid to fit groups of even several hundred learners on one screen. The Resource Focus mode features a timeline which makes it easier for learners to identify which topic is currently being covered in class (the big circle). Furthermore, green and red circles respectively denote already-covered and to-be-covered topics. The timeline was not available to subjects in the current classroom study.

Each cell from either of the three main grids in the All Resources mode or the single main grid in the Resource Focus mode grants access to the list of activities. When a cell is clicked, topic grid (or grids) are replaced by activity grid (or grids) as shown in Figure 3. Each cell of an activity grid corresponds to a




Fig. 1. Mastery Grids: All Resources mode. In this mode, the student sees all resources available to them. Each resource is a collection of activities of a different kind (here, Questions and Examples) and these resources are represented by rows of a grid. The columns of a grid show topic into which the course is divided. “Resources \times topics” create a two-dimensional representation of the course material available through Mastery Grid. The third dimension is given by the intensity of the color of a grid cell; the more deep the color the more progress. The student can see their own progress on the first grid, the average progress of the current group (e.g., all students in the class) on the third grid, and the difference between the two on the second grid. A list of individual progress of all members of the current group is shown below and in the case of large groups the students would scroll down to inspect it (only two individual progress grids are shown here). The student’s own progress is shown in that list as well so that they know where they are relative to other students. As indicated in the header of that section of the interface, the example student is 18th out of the total of 55 students.

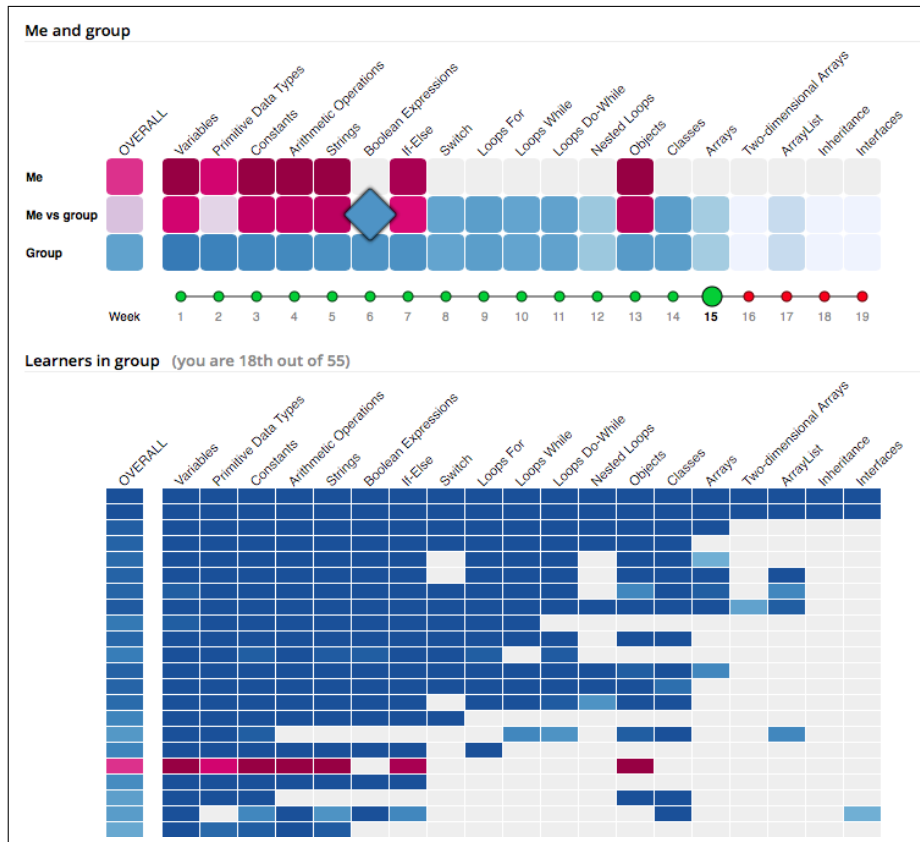


Fig. 2. Mastery Grids: Resource Focus mode. Although many resources can be available to students, they may choose to focus on only one of them at any given point. For example, when they do not yet have a good understanding of the material, they may focus on readings or example source code snippets. Later on, however, especially before an upcoming in-class exam, they may choose to focus on evaluation-type resources to see how well they can answer questions related to the material. The Resource Focus mode gives students the ability to work with one resource only. The three main grids from the All Resources mode are summarized by only one grid. That grid shows the student’s progress in the first row, the average progress of the current group (e.g., top 10 students in the class) in the third row, and the difference between the two in the second row. Below, a list of individual progress of all members of the current grid is shown. As shown on this screenshot, the student’s own progress is displayed in that list as well; that gives them a visual clue about where they are relative to other students. This interface is more compact compared to the All Resources mode. Switching between the two modes happens through a simple always-on-the-screen toolbar which is not shown here.




Fig. 3. Mastery Grids: List of questions for the Variables topic (Resource Focus mode). Irrespective of the mode the interface is in (i.e, All Resources of Resource Focus), when the student clicks on one of the cells of any of the main grids, a list of activities (see Figure 4) associated with that cell is shown. In this example, the interface is in the Resource Focus mode and therefore there is only one activities grid (vs. three for the All Resource mode) and it shows progress that has been made so far towards all activities for the selected resource (here, Questions) and topic (here, Variables). To invoke an activity, the student clicks on one of the cells of this activities grid.

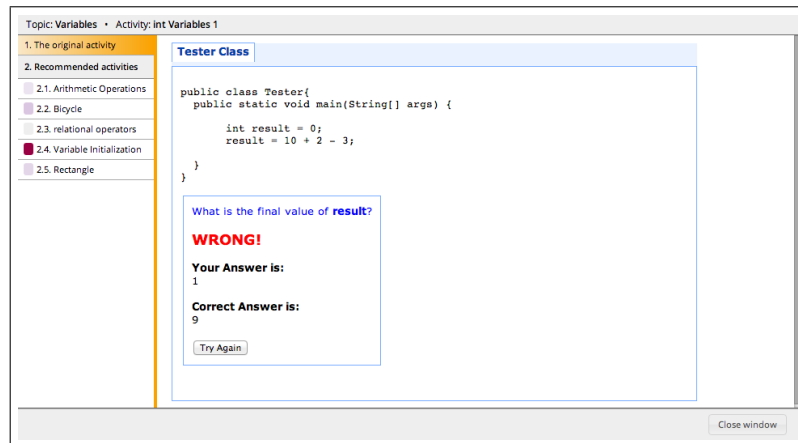


Fig. 4. Mastery Grids: A sample activity (here, a question) shown after the student clicks on one of the activity grid cells. In this example, the question has been answered incorrectly and as a result a list of recommended activities is shown on the left.

single activity the learner can access and clicking on that cell opens the activity window (Figure 4). To enhance the learning experience, when a learner answers a question incorrectly, a set of examples is recommended to them as remediation. These examples are chosen based on the similarity between the question and all the examples available.

Technically, Mastery Grids is a Web application with the front-end written in HTML, CSS3, and JavaScript and developed by the first author. At the back-end, it is supported by Java servlets logic with a relational database persistence developed by the second and third authors. The front-end does not hard-code anything content-related and expects to receive everything from the server. In that respect, it is a flexible interface which can be deployed in many progress visualization scenarios.

4 Evaluation

4.1 Classroom Study

To examine the role of Mastery Grids as both a visualization and a content access interface we gave students an alternative way of accessing content through a simple two-level hierarchy of HTTP links which from now on we refer to as *Links interface* or simply *Links*. The first level links listed topics and the second level links listed activities (i.e., questions and examples).

A combination of Links and Mastery Grids was used in three courses in the Fall 2013 term at the School of Information Sciences, University of Pittsburgh. The three courses were: Introduction to Object Oriented Programming (undergraduate) and Database Management (graduate and undergraduate). Both tools were introduced to students by the first and second author on the second week of classes in the Java course and on the fourth week of classes in both database courses. Students were informed that the use of these tools was non-mandatory and that there was no penalty for not using them. To engage students, the Java course instructor offered extra points (5 out of 100) towards class participation for solving at least 15 questions using either Links or Mastery Grids¹. Table 1 reports the number of students in each course along with the number of student sessions recorded.

4.2 Educational Material

The Links and Mastery Grids interfaces were used to provide access to and visualize progress in two kinds of self-study learning activities from two domains, Java and SQL. The first type of activities were examples delivered by WebEx system [10]. Each example was a snippet of Java or SQL code with annotated

¹ None of the students did exactly 15 questions though.

Table 1. Classroom sizes and numbers of sessions in the three courses used in the current study.

Course	Students	Sessions		
		0	1-3	4+
O.O. Programming (Java)	35	4	18	13
Databases (grad)	83	54	26	3
Databases (undergrad)	35	24	9	2
TOTAL	153	82	53	18

lines. All annotation were initially hidden and to see them a student needed to click on a line. The second type of activities were questions. Each question was a parameterized snippet of code that required students to mentally execute it and provide an answer. Java questions were delivered by QuizJet system [11] and SQL questions were delivered by SQL-Knot system [12]. Activities in both domains were organized into topics to reflect the structure of the three courses used in the current study. The Java corpus contained 19 topics with a total of 75 WebEx examples and 94 QuizJet questions while the SQL corpus contained 19 topics with a total of 64 WebEx examples and 46 SQL-Knot questions.

4.3 Usage Pattern Analysis

We used log analysis to understand how Links and Mastery Grids and their features were used by the students. Only Java class logs were used because they offered substantially more data than either of the other two courses².

Based on our analysis, students could be split into six groups based on how much they used both tools: (Z) Four students did not use either of the tools; (L1) Two students used Links only and never used visualization; (L2) Three students used Links for content access and loaded the visualization but did not use its interface; (L.MG1) 16 students used Links for content access and interactively explored visualization but did not use it for content access; (L.MG2) Six students used both Links and Mastery Grids for content access; (MG) Five students used exclusively Mastery Grids for content access. This data demonstrates a relatively high engagement levels in that while the use of Mastery Grids was not required, the majority of students used it as an information source and a considerable amount of them used it for content access despite the availability of a faster content access tool, Links.

To analyze the impact of Mastery Grids on student behavior, we combined the five groups of students who used either tool into two larger groups depending on their use of visualization (Table 2): Those who had little to no use of Mastery Grids (L1+L2; the two first columns) and those who used it considerably (L.MG1+L.MG2+MG; the two last columns). 26 students used Mastery

² We excluded student sessions shorter than 30 seconds. Additionally, we excluded two students which accessed a total number of questions larger than 3*SDs than the mean.

Table 2. Usage statistics by group (means; SDs in brackets)

Number of students in group	L1+L2		L.MG1+L.MG2		MG		L.MG1+L.MG2+MG	
	5	-	20	-	6	-	26	-
Session info								
Number of sessions	4.00	(1.58)	4.05	(2.84)	2.67	(2.07)	3.73	(2.71)
Number of MG sessions	0.80	(0.84)	3.00	(2.00)	2.67	(2.07)	2.92	(1.98)
Number of MG UI actions	0.00	(0.00)	25.30	(57.48)	111.83	(82.99)	45.27	(72.60)
Sessions with 1+ MG UI actions	0.00	(0.00)	0.46	(0.50)	0.94	(0.25)	0.54	(0.50)
Sessions with 5+ MG UI actions	0.00	(0.00)	0.14	(0.34)	0.69	(0.48)	0.23	(0.42)
Sessions with 10+ MG UI actions	0.00	(0.00)	0.10	(0.30)	0.62	(0.50)	0.19	(0.39)
Duration (h:mm)	1:00	(1:46)	1:13	(1:53)	0:24	(0:30)	1:05	(1:46)
Question answering performance								
Correct answer ratio	0.58	(0.11)	0.69	(0.13)	0.73	(0.12)	0.70	(0.13)
Material exposition (total number of accesses)								
Activities	86.60	(53.94)	111.80	(96.70)	68.17	(46.41)	101.73	(88.82)
Questions	66.20	(37.37)	85.65	(81.83)	61.00	(44.14)	79.96	(74.77)
Examples	20.40	(20.51)	26.15	(28.02)	7.17	(6.94)	21.77	(25.94)
Unique activities	58.00	(33.90)	69.75	(51.90)	45.00	(26.84)	64.04	(48.00)
Unique questions	39.80	(22.10)	45.30	(29.54)	38.00	(23.19)	43.62	(27.94)
Unique examples	18.20	(15.79)	24.45	(25.43)	7.00	(6.63)	20.42	(23.59)
Example lines	55.00	(36.10)	106.70	(133.99)	29.67	(44.44)	88.92	(123.02)
Unique example lines	47.40	(30.28)	93.75	(115.42)	22.33	(32.30)	77.27	(106.19)
Material exposition (rate, i.e., total number of accesses per minute)								
Activities	0.84	(0.82)	0.66	(0.47)	1.27	(0.61)	0.80	(0.56)
Questions	0.71	(0.74)	0.53	(0.43)	1.15	(0.66)	0.68	(0.55)
Examples	0.13	(0.09)	0.12	(0.09)	0.12	(0.15)	0.12	(0.10)
Unique activities	0.58	(0.58)	0.44	(0.30)	0.89	(0.40)	0.54	(0.37)
Unique questions	0.45	(0.49)	0.32	(0.26)	0.77	(0.43)	0.43	(0.35)
Unique examples	0.13	(0.10)	0.12	(0.08)	0.12	(0.15)	0.12	(0.10)

Grids on about half of their sessions while only six did not use the visualization interactively. As the data shows further, students who used the visualization seemed to be more engaged with self-study content in that they answered more questions, tried more examples, inspected more example line comments, and got a higher correct question answer ratio. The groups did not differ with respect to how quickly they explored the material space (measured by activity access per minute).

To examine the difference between students who used Mastery Grids exclusively and those who used it only some of the time, we also contrasted the MG group with groups L.MG1 and L.MG2 combined. According to data reported in Table 2 (3rd through 8th columns of numbers), students from the MG group worked with the content more productively by accessing questions (both those they had already seen and those they had not) at a higher rate than students from the two other groups. We return to this observation in the discussion.

While these results seem encouraging, they are based on a small number of students which incidentally prevented us from performing statistical inferences. Because of that, the above analysis needs another evaluation done in light of better data. In the future, we hope to address this concern and thus get a better insight into the different use patterns.

4.4 Material Exposition and Final Grades

While data from the current study cannot be used to draw causal links between material exposition and final grade (or any other indicators of educational performance) we wanted to discover if more student activity was associated with worse or better final grade. To that effect, we fitted four linear mixed models, each having the final grade as the response variable. To use all data we had at our disposal, we pooled data from all three courses described earlier in this section. To properly model the structure of the combined data, we included class random effect as a higher level unit. Student random effect was nested in the class random effect. Because relatively few students from the two database courses used the supplementary tools at all, we considered only students who logged in at least once (we required no activity beyond that though).³

First, we wanted to see if more interaction with the material as a whole could have affected the grade. To do that, we checked if the total number of questions and the total number of examples accessed (through folders and visualization together) were significant predictors of the final grade. After discovering that they were not, we broke the material apart into two parts and asked if perhaps the amount of content accessed only through folders or only through the visualization had a significant association with the grade. That, however, was not the case either. To summarize, we did not establish any link between the sheer amount of interaction with the material and the grade.

Because the visualization offers an elaborated user interface that presents students with a lot of information, we decided to broaden our search and see

³ Using all student data did not change the results substantially.

if increased student activity in the visualization itself could be held responsible for a change in the grade. In order to address this question, we looked at the number of *educational actions* students performed while using Mastery Grids. That number was a sum of the number of times students performed any of the following eight actions: (1) Selecting a topic; (2) Accessing a question; (3) Accessing an example; (4) Accessing a recommended example; (5) Accessing the original question again after inspecting at least one recommended example; (6) Clicking on a line of an example to reveal the comment associated with that line; (7) Assessing difficulty of a question; (8) Navigating back to topics grid from activity grids. We found a statistically significant amount of support for this relationship. More specifically, performing one educational action was associated with an increase of 0.016 in the grade (SE=0.007; $p=.0187$). That is, performing a 100 such actions would be associated with a 1.6 increase in the grade. If such an increase was to occur, then that would be a reasonable effect size.

4.5 Subjective Responses

To understand how students perceived Mastery Grids, we asked them to complete a survey consisting of two parts:⁴ (1) general evaluation containing 14 questions addressing usefulness, usability, and satisfaction, and (2) feature-specific evaluation containing 15 question about the most important components of Mastery Grids.⁵ All questions were phrased as positive statements and a 5-points Likert scale was used to collect students' responses (*strongly disagree*, *disagree*, *no opinion*, *agree*, and *strongly agree*).

A total of 39 students from the three courses completed the survey. However, we excluded three students who answered all questions giving the same score, and 11 other students who used either Links or Mastery Grids in no more than two sessions each shorter than 30 seconds. Thus, responses of 25 students were considered in the analysis.

The survey data shows that students' perception of Mastery Grid's usefulness and usability was positive (Figure 5). The majority of the students agreed or strongly agreed that the application was useful (F1, 70.8%), that they liked the tool and would like to use it in other courses (S2, S14, 83.3% and 70.8%, respectively), that the tool was easy to use (F10, 66.7%), and that it was easy to remember how to use it (U11, 79.2%). Interestingly, social comparison aspects of progress visualization were perceived more positively than personal aspects. That is, while the majority of students (F7, 54.2%) agreed or strongly agreed that viewing classmates' progress motivated their own progress, less than a third agreed or strongly agreed that progress visualization helped to identify their weak points (F5, 29%) or plan their class work (F6, 29%).

⁴ Due to space limitations we do not provide text of questions but we refer to them as much as possible in the main text.

⁵ Screenshots were used to indicate the parts of user interface questions were referring to.




Fig. 5. Subjective responses: General (F–Usefulness; U–Usability; S–Satisfaction).




Fig. 6. Subjective responses: Feature-specific.

The feature-specific responses (Figure 6) were also generally positive: 76% and 68% of the students agreed or strongly agreed that the “Me” grid was useful for viewing one’s own progress at the level of topics and activities, respectively (questions 1 and 2). Similarly, 72% and 60% of students agreed or strongly agreed that the “Group” grid was useful at the level of topics and activities, respectively (questions 3 and 4). However, a smaller fraction of students thought that it was useful to be able to compare oneself against the group using either the “Group” grid (52%) or the “Me vs group” grid (44%). The usefulness of the individual students’ progress (the grids displayed below the main grids) was evaluated positively by about half of the students (48% for topics, question 9; 52.9% for activities, question 10).

Regarding the feature-level usability, 60% of students thought that progress at both the topic level and the activities levels was clearly shown. However, only 48% of the students stated that they quickly learned how to go from topics to activities and back (question 13) which could explain why majority of students used Mastery Grids interface to view their progress and progress of others and only a third of them used it for content access.

5 Discussion and Conclusions

In this article, we have presented Mastery Grids, an open source implementation of a social progress visualization interface. We have also reported results of a classroom study that examined how students used the system as well as their feedback. In that semester-long study, students in three courses were offered access to two kinds of advanced self-study content and progress visualization through Mastery Grids. They were also offered a simpler and more traditional way of accessing content we have called Links.

When analyzing usage patterns we observed that students who used Mastery Grids had a higher ratio of questions answered correctly than those who used Links only. It is possible then that the visualization guided students to questions which were more suited for their level of understanding of the material. We also observed that the visualization directed students to new material at rates higher than the alternative Links interface. A similar finding has been reported before in the context of adaptive explanatory visualization [13]. The mechanisms responsible for elevating material exposition rates in both the case of that earlier report and the current study are likely the same and may be the result of the visualization attempting to stay in sync with students’ progress and thus being able to direct them to new content more quickly. If that was indeed the case, then our results would underline the importance of accurate student modeling.

Our finding of the number of educational actions performed in Mastery Grids being predictive of the final grade is a hopeful one. After all, a tool like that should help students to climb the mountain of scholastic achievements. However, because of the lack of manipulation, we recommend that this relationship be studied more systematically. More importantly, the current study cannot adjudicate if it is indeed the case that using the visualization more helped students

with getting a better grade or if instead students which ended up getting a better grade were also the ones more likely to be engaged with supplementary educational tools.

Finally, student feedback analysis demonstrated that students assessed the usefulness and usability of the system quite positively. At the same time, some features of the system were regarded as less positive than others indicating that there is room for improvement for the future versions of Mastery Grids.

References

- [1] T. Naps, G. Rößling, J. Anderson, S. Cooper, W. Dann, R. Fleischer, B. Koldehofe, A. Korhonen, M. Kuittinen, C. Leska, M. McNally, L. Malmi, J. Rantakokko, and R.J. Ross. Evaluating the educational impact of visualization. *ACM SIGCSE bulletin*, 35(4):124–136, 2003.
- [2] S. Bull and J. Kay. Student Models That Invite the Learner In: The SMILL() Open Learner Modeling Framework. *International Journal of Artificial Intelligence in Education*, 17(2):89–120, 2007.
- [3] L.A. Festinger. Theory of social comparison processes. *Human relations*, 7(2):117–140, 1954.
- [4] F. Linton and H.-P. Schaefer. Recommender systems for learning: Building user and expert models through long-term observation of application use. *User Modeling and User-Adapted Interaction*, 10(2-3):181–208, 2000.
- [5] J. Vassileva and L. Sun. Using community visualization to stimulate participation in online communities. *e-Service Journal*, 6(1):3–39, 2007.
- [6] A.T. Corbett and J.R. Anderson. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction*, 4(4):253–278, 1995.
- [7] A. Mitrovic and B. Martin. Evaluating the effect of open student models on self-assessment. *International Journal of Artificial Intelligence in Education*, 17(2):121–144, 2007.
- [8] J.-D. Zapata-Rivera and J.E. Greer. Interacting with inspectable Bayesian student models. *International Journal of Artificial Intelligence in Education*, 14(1):1–37, 2004.
- [9] I.H. Hsiao., F. Bakalov, P. Brusilovsky, and B. König-Ries. Progressor: Social navigation support through open social student modeling. *New Review of Hypermedia and Multimedia*, 19(2):112–131, 2013.
- [10] P. Brusilovsky, M.V. Yudelson and I.H. Hsiao. Problem solving examples as first class objects in educational digital libraries: Three obstacles to overcome. *Journal of Educational Multimedia and Hypermedia*, 18(3):267–288, 2009.
- [11] I.H. Hsiao, S. Sosnovsky, and P. Brusilovsky. Guiding students to the right questions: Adaptive navigation support in an e-learning system for Java programming. *Journal of Computer Assisted Learning*, 26(4):270–283, 2010.
- [12] P. Brusilovsky, S. Sosnovsky, D.H. Lee, M.V. Yudelson, V. Zadorozhny, and X. Zhou. An open integrated exploratorium for database courses. In *Proceedings of 13th Annual Conference on Innovation and Technology in Computer Science Education (ITiCSE)*, pages 22–26, 2008.
- [13] T.D. Loboda and P. Brusilovsky. Adaptation in the context of explanatory visualization. In *Proceedings of 3rd European Conference on Technology Enhanced Learning (ECTEL)*, pages 250–261, 2008.