2006-1164: HELPING STUDENTS VISUALIZE THEIR GRADE PERFORMANCE

Will Humphries, Virginia Tech

Justin Gawrilow, Virginia Tech

Scott Turner, Virginia Tech

Manuel Perez-Quinones, Virginia Tech

Stephen Edwards, Virginia Tech

Helping Students Visualize Their Grade Performance

Abstract

This paper discusses the design of visualizations of student grade performance. A needs analysis identified five qualities of performance that students and faculty believe to be important. These were used as requirements for the design of several visualizations. These visualizations were implemented as an integrated feature of the electronic grade book of a popular Course Management System (Moodle). This paper discusses the motivations for our work, the qualities of performance identified, the implemented visualizations, and an initial evaluation of their utility.

1. Introduction

Many professors frequently hear questions from their students regarding their performance in a course. Students often want to know what grade they have so far, what their likely final grade will be in the course, how they are doing in comparison with their peers, etc. It is important for students to self-assess and to have a grounded understanding of their actual performance in a class. Yet students rarely seem to get such an understanding just by looking at their assignment scores. Some students bring such questions to the professor, but providing answers requires making educated guesses as to how a student will perform on future work. A series of good visualizations based on student data in the course can help answer these questions for students.

In this paper, we present an analysis of student needs when self-assessing their performance in class. The analysis was intended to identify what student performance attributes, or *qualities of performance*, were most used by students. We also present visualizations designed and evaluated based on their ability to help students meet those needs. We integrated these visualizations into a course management system called Moodle[4] to help address student and faculty needs.

A common feature of course management systems (CMS) is the ability of students to view their grades in a current class. This data is traditionally presented in a tabular format showing the student's score on each assignment. Moodle, an open source CMS, follows this trend when displaying its grade book data to students[10]. The benefits of visualizing numerical data in a graphical format include a more efficient and clearer interpretation of the data[9].

The goal for this project is to inform students more effectively of their performance in a course using visualizations of current assessments, comparisons with the class as a whole, and projections of potential and expected future outcomes. In turn, this might allow students to change their behavior for the better, based on their observations of the graphs.

2. Prior Work

Edward Tufte's Visual Display of Quantitative Information[11] is a helpful reference on the topic of information visualization. His principles of graphical excellence enable us to create graphics that give the viewer the information they seek in minimal time.

Naps et al. in their paper "Evaluating the Educational Impact of Visualizations[5]" support the assertion that various visualizations enhance understanding and recognition in the classroom setting. For example, they discuss how a bar graph can be used to supplement numerical statistics about the geographic distribution of poverty. Naps indicates that the time and effort required to develop these visualizations is largely responsible for why they are not used more often. Our work is based on similar principles; we want to provide students a quick and easy way to absorb numerical data (in our case, student grades). Furthermore, because our visualizations are integrated into a CMS, they are automatically generated and always contain the most up to date data from the class grade book.

Bouyssou's work[1] is also relevant during the development of graphical models for student grades. Visualizations act as "lenses" through which grades can be evaluated. Under different lenses, the same set of scores could appear to be successful or unsatisfactory. In this way, the bias that can be introduced through the different systems for aggregating grades is made clear[1]. There is some similarity between Bouyssou's idea of lenses and our qualities of performance, as discussed later in Section 3.

Saraiya, North and Duca[8] created an innovative way to evaluate visualizations. They state that visualizations exist to provide greater insight and that the same methodologies exist across domains and may, therefore, be applied to scenarios involving different kinds of data. They proposed a methodology of evaluating visualizations based on "insights" gained about the problem domain while assessing visualizations. Insights are defined as "individual observations about the data by the participant." Part of our evaluation of the visualizations was based on the idea of insights (see Section 3).

There are other areas, such as student access logs, where visualizations might be helpful. Moodle currently provides some line graphs to display access log data (see Figure 1). GISMO[2] is an applet-based Moodle add-on that provides more features for examining Moodle access logs.

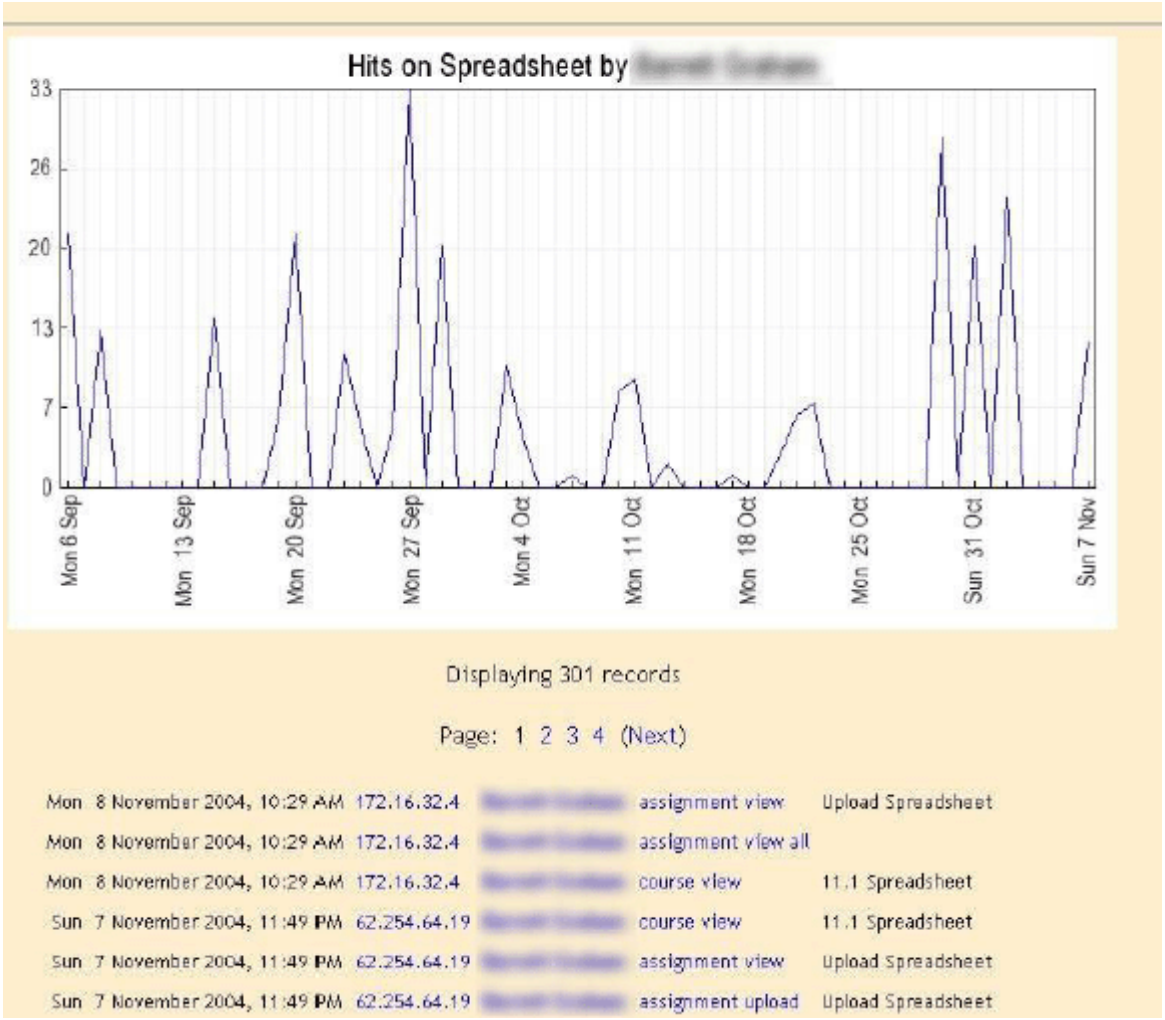

Figure 1. Instructor View of Student Access Log

CourseVis[3] is another system that uses a CMS access log to produce visualizations. In this system, different visualizations are created to help the instructor assess qualitities of a distance-learning course. The emphasis is on visualizations for the instructor instead of the student.

In general, there is a lack of literature asserting the usefulness of grade visualizations to help students assess their progress, and a lack of literature discussing which visualizations will be most helpful to students.

3. Qualities of Performance

Students currently view their grades in a textual format, even in courses that use a CMS. This representation does not provide them with the best context to enable the student to visually track their performance over time, and/or to make a comparison with the rest of the class.

Evaluating potential visualizations requires considering which factors affect their worth. Their usefulness to the students is a primary characteristic for determining which visualizations to use, and where they should be placed within the CMS.
To identify student needs, we first identified the most typical questions that students ask of us (as professors) regarding their performance in class. This included a short list of:

- How did I do compared to the rest of the class?
- Given what I have done so far, what grade will I get at the end?
- What should I do to improve my final grade?

With this initial list, we formulated scenarios of use[6,7]. These scenarios were focused on students accessing grades on the course web site with a particular performance-related purpose. For example, one scenario detailed the thoughts and actions of a student attempting to determine whether she/he should drop a class or stick with it. In another example, a student is concerned with her/his performance, and wants to determine whether she/he is keeping up with the rest of the class.

The development, analysis, and iterative refinement of these scenarios provided unique insights into the needs of students viewing their performance data. We discussed these scenarios with other students, faculty, and as a semester long project in a graduate seminar course at Virginia Tech.

The insights gained helped us identify several *qualities of performance* that students are interested in at different times. These include:

a) Performance relative to other students
b) Perfect scores/missed scores
c) Adequacy of effort thus far
d) Impact of a particular assignment on one's overall grade
e) Current grade

The quality of '*performance relative to other students*' is often important to students. This is because they are concerned about their potential class rank, and their final grade after the professor curves class grades. Another quality identified was '*perfect scores/missed assignments*', which is not always a piece of information students actively seek out, but one they should be made aware of, especially in the case of missed assignments. This quality of performance was significant to professors. Many felt that several missed assignments indicated a poor student disposition towards the class. The third quality was '*adequacy of effort thus far*', a quality that is important when students are attempting to determine if they have been putting enough work into a particular class, or if they can afford to put in less work. The last two qualities '*impact of a particular assignment on overall grade*,' and '*current grade*' were typical requests that students made in our interviews. In the next section, we describe how different visualizations were created and assessed with regards to their ability to convey these qualities.

3.1 User Survey

We developed several prototype visualizations. These initial prototypes were generally based upon other graph types that are well known, accurate, and powerful. Some of these graphs were tailored to meet the needs of students, but for the most part the focus of this early phase was to bring forward many graphs that would be understood easily by the target user group and that were capable of quickly and accurately

conveying grade data. These graphs were then refined or removed from consideration based on their ability to help students make conclusions and decisions based on their performance.

A user survey was conducted to determine how effectively these visualizations could convey the five qualities of performance identified in the last section. The survey was provided to 19 Computer Science students at Virginia Tech (12 undergraduates and 7 graduates) during the Spring 2005 semester. The survey contained one section for each visualization under consideration. Each section was composed of an image of the graph, a brief description of the graph's content, and several Likert-type questions about the graph. The Likert-type questions were used to assess the graphs' helpfulness when assessing a particular quality of performance. The questions asked were:

a) How helpful was this graph at determining your current grade in this course?
b) How helpful was this graph at determining whether you are in the top third of this class?
c) How helpful was this graph at assessing whether your effort thus far has been adequate?
d) How helpful was this graph at highlighting your perfect scores and / or missed assignments?
e) How helpful was this graph at assessing the impact any individual assignment had on your grade?
f) On the whole, how helpful was this graph?
g) Any comments for this graph / chart?

The survey contained five graphs. The first was a line graph that made predictions about future grade outcomes (referred to here as the projection graph), such as best/worst possible grade and the most likely final grade. The second was a bar graph showing grades for every assignment and their relative weightings. The third and fourth visualizations displayed the student's grades in relation to the rest of the class. One was displayed as a line graph and the other as a histogram. The final graph used a box and whiskers plot to describe the outcome of a single assignment.

The survey results are shown in Figure 2. All the graphs, except the box and whiskers, were ranked about the same on how well they displayed adequacy of effort. There was more variance with the other questions. The projection graph was rated the highest overall, despite the fact that some of the additional comments indicated that people thought the graph was somewhat confusing and complicated. In general, they felt it was useful in determining the current grade and impact of assignments on the overall grade. The bar graph was, by far, the best at showing perfect scores/missed scores and was helpful in assessing the impact of assignments. However, several students noted that it lacked any information about the performance of the rest of the class or about the current grade. The line graph and the histogram proved to have strengths and weaknesses opposite to those of the bar graph. While they displayed the current score and the position relative to the rest of the class, they said very little about the weight of the assignment. Compared to the line graph, the histogram responses were more moderate. It was not ranked as high on its strong points nor was it ranked as low on its weak points. The box and whiskers plot scored poorly across all of the questions. This may be due, at least in part, to the students' unfamiliarity with this type of graph.

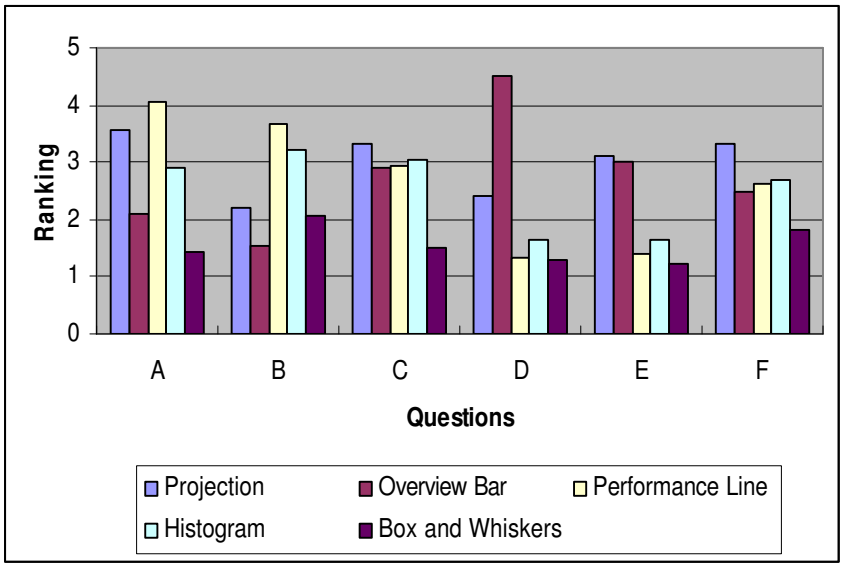

Figure 2. Graph Ranking by Question

4. Visualization Implementation

Based on the findings of our survey, we selected four visualizations for implementation. The high acceptance of the bar graph led us to divide it into two separate graphs: one that shows an overview of all the grades and one that shows a more detailed view of the grades within a particular category of assignments. These two separate views should give the students the information they want without becoming needlessly noisy. The histogram was also implemented to show the details of a single assignment. We had hoped to implement the prediction graph, which was rated as valuable by students, but significant implementation issues in accessing the appropriate data to generate the graph caused us to defer it for future work.

In addition, we implemented another visualization that provides an alternative overview of the grade information. While not covered by the first survey, this graph, called a star graph, was motivated by the need to address quality of performance (a) when examining overall grades. The star graph appeared to have enough potential usefulness for this goal that it was added to the interface. It is discussed in Section 4.1.4.

The locations of the graphs in the interface were based on what the authors felt users would be attempting to determine, as inferred from their position. For example, a user looking at her/his grades on all tests would probably have an interest in whether their effort thus far has been adequate.

4.1.1 Grade Overview Bar Graph

Figure **3** shows an example of a bar graph that gives an overview across all assignments. The bar graph has one major column for each grade category assigned by the instructor (e.g. exams, homework assignments, projects, etc.). Each of these category columns is further subdivided into smaller columns for each assignment within the category (e.g. exam 1, exam 2). The columns are then 'filled' vertically with the percentage the student scored on the assignment. The percentage of credit lost is then colored

black. The width of each column is proportional to the weight of that assignment in computing the student's cumulative score. This visualization addresses performance qualities (b), (c), and (d). We experimented with adding class averages to the graph, but that increased clutter.

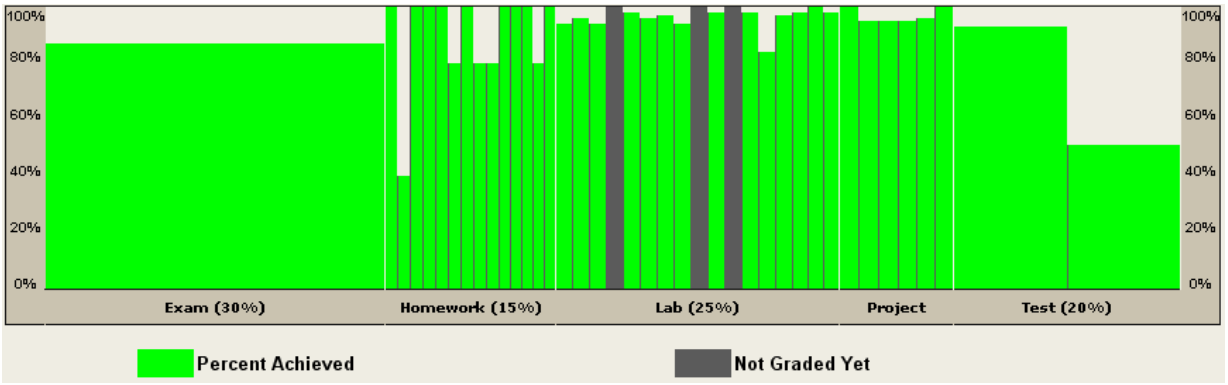

Figure 3. Grade Overview Bar Graph

This graph is located on the main grades page that students visit to access any of their grade information. Based on the strengths identified in the survey, its primary purpose is to give students a general feeling about their effort thus far and to highlight any perfect scores and/or missed assignments. It also provides insight into the impact of individual assignments on the total grade, because students can see what grading categories exist in the course and the relationships between their weightings.

4.1.2 Category Bar Graph

When students drill down to see their scores in a particular assignment category (e.g., all program assignments), they see a different graph, as shown in Figure 4. This graph is very similar to the grade overview graph, with more detail on a smaller data set. The primary differences are that assignments are now only shown within a single category and the visualizations are turned sideways. This visualization was rendered horizontally to provide more flexibility in displaying assignment names. Unlike the overview graph, all bars have the same thickness. This ensures that assignments that are worth only a small percentage of the final grade still can be clearly seen, which is not the case with the overview graph. This graph addresses performance qualities (a), (b), and (c).

This graph can be found on the grade page for each assignment category. Students honing in on a particular category are more likely to be concerned with whether their efforts have been adequate relative to others in the class. This is the reason that the average line was added (as shown in Figure 4). We feel it will allow students to note whether they did enough to remain above class average on a particular assignment. At the same time, the graph will still highlight any assignments that have been missed.

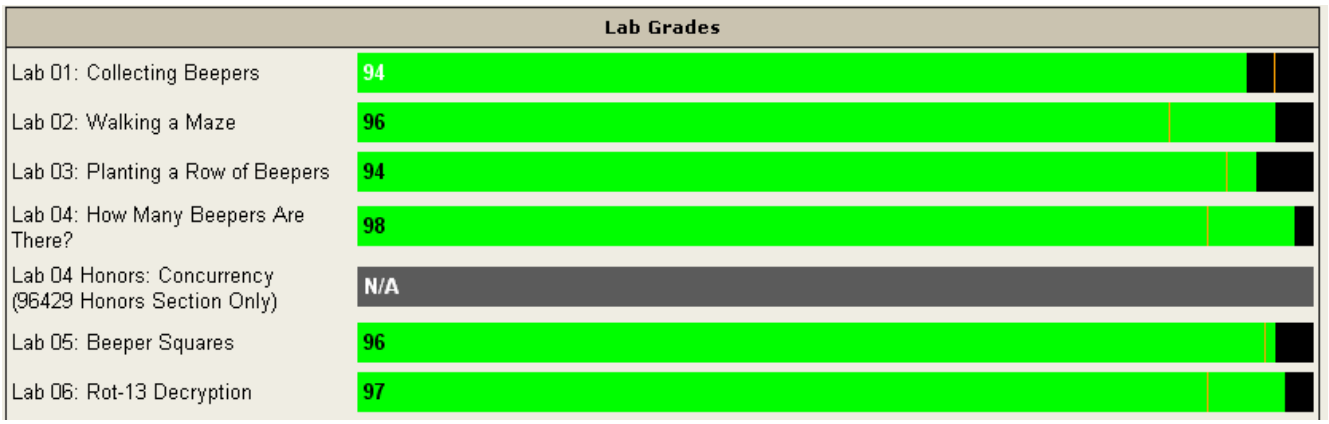

Figure 4. Category Bar Graph

4.1.3 Assignment Histogram

Figure 5 shows the histogram that students see when viewing the results of a single assignment. The student grades are divided into approximately 10 ranges each of which is labeled on the x-axis. The bar that contains the student's score is displayed in a different color to highlight it. This graph is currently displayed when users request more information about their grades on a single assignment. The graph was positioned here based on the high ranking it received from students at determining current score and determining position relative to other students. This histogram addresses performance qualities (a) and (b).

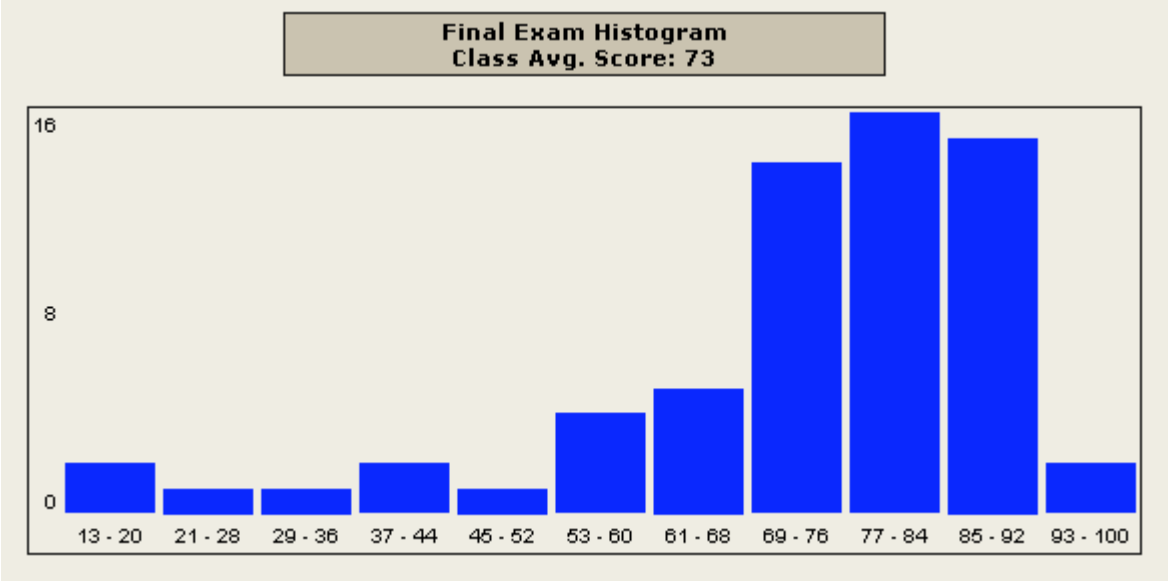

Figure 5. Assignment Histogram

4.1.4 Star Graph

The star graph illustrated in Figure 6 is intended to provide a high level overview of grade information and to emphasize the areas where a student is strong or weak. It targets performance qualities (a) and

(c). Essentially, the graph is a regular polygon with a vertex for each category in the grade book. The colored bands mark off evenly spaced ranges in the grading scale. In our implementation, they divide each category into quarters. The student's score for each category is plotted along the corresponding radius and is used to form another, usually irregular, polygon inside the first. The graph on the left shows a student's individual performance, while the graph on the right shows the class averages for the categories.

This graph is intended to be an alternative view of overall performance, so it was also placed on the main grades page. A simple dropdown box allows students to switch between the overview bar graph and the star graph. The overview bar graph is displayed by default.

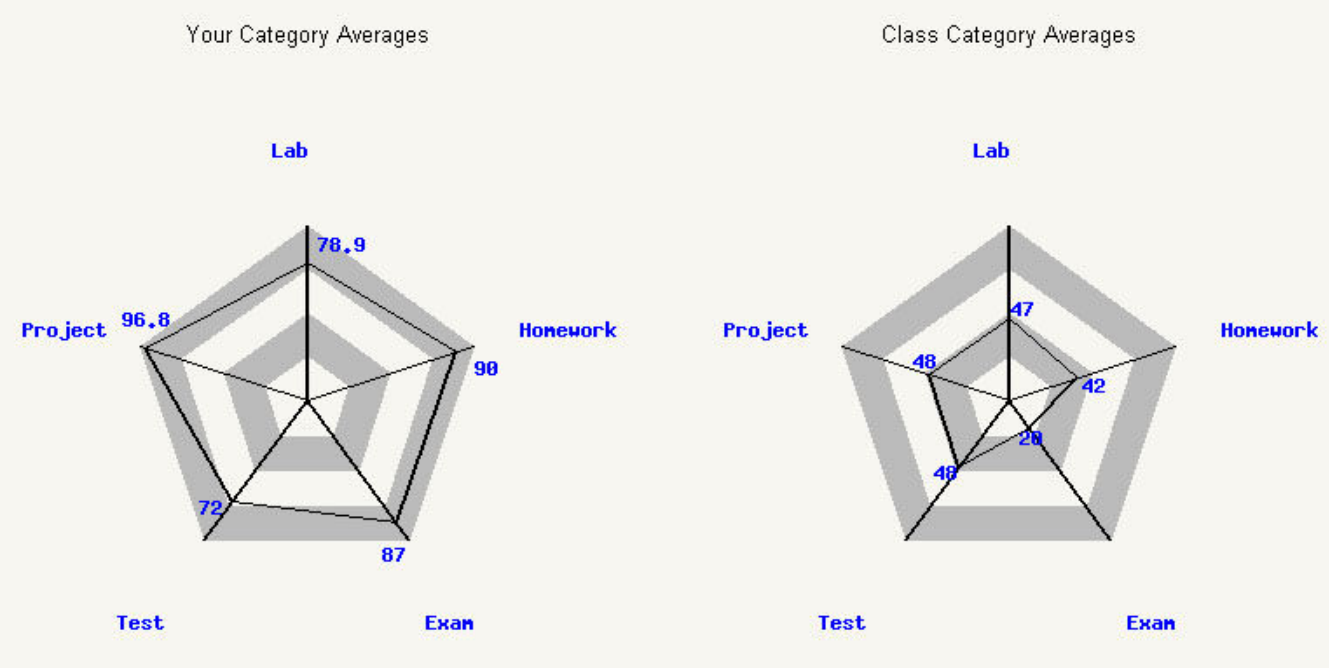

Figure 6. Star Graph

5. Evaluation

Once a set of graphs was decided upon, we integrated them into Moodle for use during summer session classes. With visualizations in place, we surveyed a few students who were using the enhanced grade book in a summer course. This provided us with early feedback about our implementation. A second, larger survey was also performed at the end of the Fall 2005 semester. Even though the graphs were still going through modifications, they were made available to students in three computer science courses for the entire term. At the end of the semester, all students in these courses were given an optional on-line survey that covered how they used their grade information and how they responded to the visualizations. 56 of the students completed the survey.

The grade overview bar graph, the category bar graph, and the assignment histogram seemed to be accepted well by the students. For both the overview bar graph and the assignment histogram, approximately two-thirds of the students rated them as easy to understand and use. The category graph

did better with 82% ranking it positively. The star graph was not rated as highly. Only 27% of students thought it was easy to use, with an additional 30% rating it neutral. From some of the comments about the star graph received on the survey, it seems its purpose may be unclear to some students. Currently, a short written explanation in the form of a help box is presented to the students. A more detailed or more prominent description might be required.

Since the graphs were located in different places within Moodle's grade book, the students were also asked about the visibility of the graphs and if they saw them during the semester. Nearly everyone saw the overview bar graph. This is not surprising, given that it is the default graph on the main grade page. The star graph and the category bar graph were not as visible, with 18% and 21% of students, respectively, commenting that they had never seen them before. It is interesting to note that 20% of the class never looked in detail at the grades within the individual categories but they found the category bar graph the most useful of all the graphs. Perhaps the category grade pages should be made more prominent by creating more ways to reach them. Nearly 1/3rd of the students missed the assignment histogram. Likely, this is due to the fact that the link to the graph is an icon under each assignment name in the category grade page and it requires the most navigation to get to. Again, providing more ways to find the graph should increase its visibility.

The most common complaint made by the students is that the graphs were not always up to date. That is, the graphs would show averages that were lower than what the student expected because the averages contained ungraded work. This occurred because the grades were not entered into Moodle as quickly as the students might want. While this is not actually a failing of the graphs per se, it is important to note this issue, how it is perceived by students, and how it might be addressed.

6. Implementation

The grade visualizations are implemented as a patch to the standard Moodle 1.5 grade book. Written in PHP, most of the visualizations were created with standard HTML, with the star graph using PHP's GD graphics library to generate images. Although Moodle supports multiple languages, currently our visualizations only support English.

Our future plans include releasing these visualizations as an extension to the Moodle grade book so other users of Moodle can benefit from our work.

7. Future Work

We are currently using these visualizations in our Moodle production server. This server hosts eight computer science courses covering all four years of our curriculum. We will improve and evaluate these visualizations through upcoming semesters. This will help determine how useful the graphs are to the students, in what situations they are useful, and if they cover all the students' needs.

Other possible work includes the implementation of additional features to the visualizations. One idea is to build links into the visualizations as a way to allow students to get more detailed information on a particular portion of the graph, and to 'drill down' for a more detailed view.

We are also considering creating more graphs specifically targeted for instructors. Visualizations have been provided for the students to allow them to more clearly understand where they stand in a class but there is a general lack of visualizations for instructors to track student progress and overall class performance.

8. Conclusions

In this paper we presented a series of performance qualities that students are interested in observing from their class grades. Students naturally ask these questions about their own performance, and informed self-assessment is critical to making decisions about where to spend effort, whether or not to drop a course, and the preconditions for desired outcomes. The performance qualities were used as requirements for the development of visualizations of grade book data. These visualizations were evaluated through two surveys with students and are currently implemented as an enhancement to Moodle. The goal of our effort is to enable students to more rapidly make conclusions and decisions about their performance in classes. Further evaluation will reveal the extent to which these visualizations are viewed, and the effect they have on the students using them.

1 Bouyssou, Denis. Evaluation and Decision Models: a critical perspective. Kluwer Academic, Boston, MA, 2000.

2 GISMO http://gismo.sourceforge.net/ Aug 2005.

3 Mazza, R., Dimitrova, V. (2004) Visualising student tracking data to support instructors in web-based distance education. *Proceedings of the 13th International World Wide Web Conference*. ACM Press, NY, May 19-21, 154-161.

4 Moodle 1.5. http://moodle.org. Aug 2005.

5 Naps, Thomas, et al. Evaluating the educational impact of visualization. In *Working group reports from ITiCSE on Innovation and technology in computer science education.* ACM Press, Thessaloniki, Greece, ACM Press 2003, 124-136.

6 Rosson, M.B., Carroll, J.M. (2002) Usability Engineering: Scenario-Based Development of Human-Computer Interaction. Academic Press, San Diego, CA.

7 Rosson, M.B., Carroll, J.M., and Rodi, C. (2004) Case Studies for Teaching Usability Engineering. *Proceedings of the 35th SIGCSE Technical* Symposium on Computer Science Education. ACM Press, Norforlk, Virginia, March 3-7, pp. 36-40.

8 Saraiya, P., North, C., Duca, K. An Insight Based Methodology for Evaluating Bioinformatics Visualization. In *IEEE Transactions on Visualizations and Computer Graphics*, Vol. 11, No. 4, 2005 July/August.

9 Shneiderman, B. (1996). The eyes have it: A task by data type taxonomy of information visualizations, Proc. In *Proceedings of the IEEE Symposium on Visual Languages '96*. IEEE, Los Alamitos, CA, September 1996, 336-343.

10 Smith, D. Using your Moodle. In *Teaching ICT Journal*, Issue 5, Winter 2004.

11 Tufte, Edward. *The Visual Display of Quantitative Information.* Graphics Press, Cheshire, CT, 1983.