Journal of Computing in Higher Education
Fall 2000, Vol. 12(1), 57-83.

EVALUATING THE EFFECT OF A COURSE WEB SITE ON STUDENT PERFORMANCE

Jesse M. Heines
Department of Computer Science
University of Massachusetts Lowell

ABSTRACT

THE AUTHOR CREATED A NUMBER OF WEB SITES to enhance traditional classroom instruction in computer science courses. These Web sites included lecture notes, assignments, downloadable programs, links to sites related to the course subject matter, and a program allowing students to see their grades on all assignments and tests and thus determine exactly where they stood in the course at any time.

This study reports on two types of data analyzed to gain insight into students' use of the site: responses to an author-created survey and students' final grades. Students demonstrated strong positive reactions to the course Web site on the survey and showed statistically significant final grade improvement after the Web site was introduced. While these results should be interpreted conservatively due to the large number of uncontrolled variables that affect student performance, they are nonetheless encouraging enough to warrant continued effort to develop and evaluate course Web sites. *(Keywords: course Web site, evaluation, analysis, student performance, course development)*

PURPOSE OF THE STUDY

THE CREATION OF A COMPREHENSIVE COURSE WEB SITE is a great deal of work. Various tools exist to help educators "put their courses on the Web," from specific tools for education such as Web Course in a Box[1] and TopClass[2], to Web page publishing wizards for existing materials in formats such as PowerPoint and Word[3], to general WYSIWYG Web page editors such as FrontPage[4], PageMill[5], and HomeSite[6]. Regardless of the tools used, an educator who creates a Web site and makes it an integral part of his or her course soon discovers that designing and producing the initial Web site are only small parts of the entire task. *Maintaining* that Web site throughout an entire semester—keeping its content current and its links up-to-date—is a much more time-consuming activity, by far.

The evaluation of a course Web site's effectiveness is difficult. There are numerous dimensions on which the site can be evaluated, and there are many variables that affect the usefulness of the site to instructors as well as to students. Nonetheless, without meaningful evaluation it is impossible to justify the effort and expense involved in creating comprehensive course Web sites. While both students and instructors may *enjoy* working with course materials on the Web and *believe* that that enjoyment translates into increased learning or at least time on task, such beliefs are unfounded without empirical evidence. Almstrum et al. (1996) have stated:

> Technology is often introduced into education to attract and excite, without any more than an assumption that it might be useful. But, if applied without deliberate study of its use in context and without the evaluation of the technology's impact on this use, "educational" technology *remains* a toy. (p. 201, emphasis in original)

CREATING AND EVALUATING COURSE WEB SITES

DAVID JORDAN (1997) attempted to "create a class Web site that would have a measurable effect on the students' education and/or morale" in a course entitled *Making of the Modern World.* Based on student responses to an instructor-designed survey and examination results in six semesters in which the course was offered (only the last of which used a course Web site), Jordan concluded that "the course *seems* to be better, but it is not knowable quite how much better, and the difference probably is not much." Given his finding that "creating and maintaining the site roughly doubles the work involved in teaching the course," Jordan was clearly disappointed that this effort did not have any "demonstrable educational effect."

There appears to be widespread agreement that the creation of a course Web site is a significant undertaking. Gilbert (1998) has remarked that many faculty, though excited about new ideas and new uses of technology to improve teaching and learning, are also "frustrated and worn out with the amount of time and effort it takes to use technology responsibly and effectively." But all efforts to improve education, particularly those related to the integration of technology, have historically required huge time and energy commitments on the part of a small number of dedicated faculty, often at the expense of their academic careers. Almstrum et al. (1996) have advised faculty facing promotion and tenure decisions that "time spent on educational issues is rarely valued for career advancement."

"The astonishing observation," Gilbert (1998) concludes, "is that so many people are willing and able to make the extra effort to try new approaches, modify them, and integrate the ones that really work." It is the determination of what really works that is so difficult. Despite that difficulty, evaluation efforts continue, because both faculty and students indicate that "they do not have the time or the desire to use

any technology unless it results in a greater understanding of the course content" (Williams Glaser, 1998).

William Trochim (1996) and his students have investigated "a variety of methods for evaluating Web site development and use and the effects of the World Wide Web on the people who use it." Their findings, however, are less than conclusive. Working with Trochim, George Fitzelle (1996) used a student survey to try to determine "whether the Web site enhanced student *perceptions* of learning" [emphasis added] by asking three Likert-scale questions:

1. The Trochim Web site significantly enhanced my learning of research methods.
2. Using the knowledge base significantly helped me learn the course content.
3. The Trochim Web site helped me to do well in the course.

Using a scale of 1 for Strongly Agree, 2 for Agree, 3 for Disagree, etc., the mean responses to these questions were 1.97, 1.87, and 1.88, respectively. Fitzelle and Trochim reported that these results indicate "students thought that the Web site *significantly* enhanced their learning of course content" [emphasis added]. The important caveats in this statement are:

1. There was no empirical measure of student performance using test or course grades as Jordan (1997) had done.
2. There was no precise definition of what constitutes "significantly enhanced learning."

Thus, while Fitzelle and Trochim's survey approach may accurately measure students' *perceptions* of their own learning, one cannot attach statistical significance to the results in terms of the Web site's effect on actual student performance. It is important to reiterate that Fitzelle and Trochim do not claim to have measured statistical significance. Rather, the fact that they did not attempt to do so may be an indication of the difficulty of doing so.

The task of evaluating whether the technology improves learning is so difficult, in fact, that some have taken a markedly pessimistic view

toward the value of such research and the types of questions that it can answer. For example, in commenting on questions such as:

- What do computers teach best?
- Does video encourage passive learning?
- Is it cheaper to teach with telecommunications?

Stephen Ehrmann (1995), who must certainly have seen his share of educational research as a program officer for the Fund for the Improvement of Postsecondary Education (FIPSE), has stated, "I don't think [those questions] *can* be answered in any reliable, valid way" [emphasis in original]. Dehoney and Reeves (1999) have reported qualitative measures of various dimensions of course Web sites, but none of their measures included assessment of actual student performance.

Others are more optimistic. In a widely quoted analysis of numerous studies on the effectiveness of computer-based instruction, Kulik and Kulik (1991) found that this technology yields, on average, a 20% improvement in learning outcomes and speed. Many feel that course Web sites will eventually show the same results. Even the more pessimistic Ehrmann (1995) believes that while the large global questions may never be answered, much is to be gained from evaluating

- not the technology per se but how it is used;
- not so much what happens in the moments when the student is using the technology, but more how those uses promote larger improvements in the fabric of the students' education; and
- not so much what we can discover about the average truth for education at all institutions, but more what we can learn about our own degree programs and our own students.

This is the more localized view taken in this study. The University of Massachusetts Lowell Computer Science Department offers a four-year B.S. program in which approximately 85% of the students are commuters. Most of the students also work at least 15 hours per week, and some as many as 30. About half of those students who work are employed in computing-related jobs.

THE UMASS LOWELL 91.353 COURSE WEB SITE

THIS RESEARCH FOCUSES ON A WEB SITE developed for course 91.353, GUI Programming I. Its URL is http://www.cs.uml.edu/~heines/91.353, and it is completely public, so readers are welcome to visit it. The discussion and figures in this section attempt to give a feel for the site and its contents.

THE COURSE HOME PAGE

The course home page is shown in Figure 1. All features of the site are accessible from the table of contents in the frame at the left:

- the course home page
- a grade display program that shows students their statuses in the course
- an on-line discussion area for posting questions to the instructors and other students
- a list of links to all course assignments
- a list of links to all class lecture notes
- a list of links to related references
- the course syllabus
- an anonymous suggestion box
- interesting downloads
- assorted utility programs

The sections that follow discuss the site's three main features: the lecture notes, the anonymous suggestion box, and the grade display program.

LECTURE NOTES

Students indicated in a survey (discussed later in this paper) that the availability of detailed lecture notes for each class was the most valuable aspect of the course Web site. However, the maintenance of these notes proved to be the most time-consuming facet of maintaining the site.

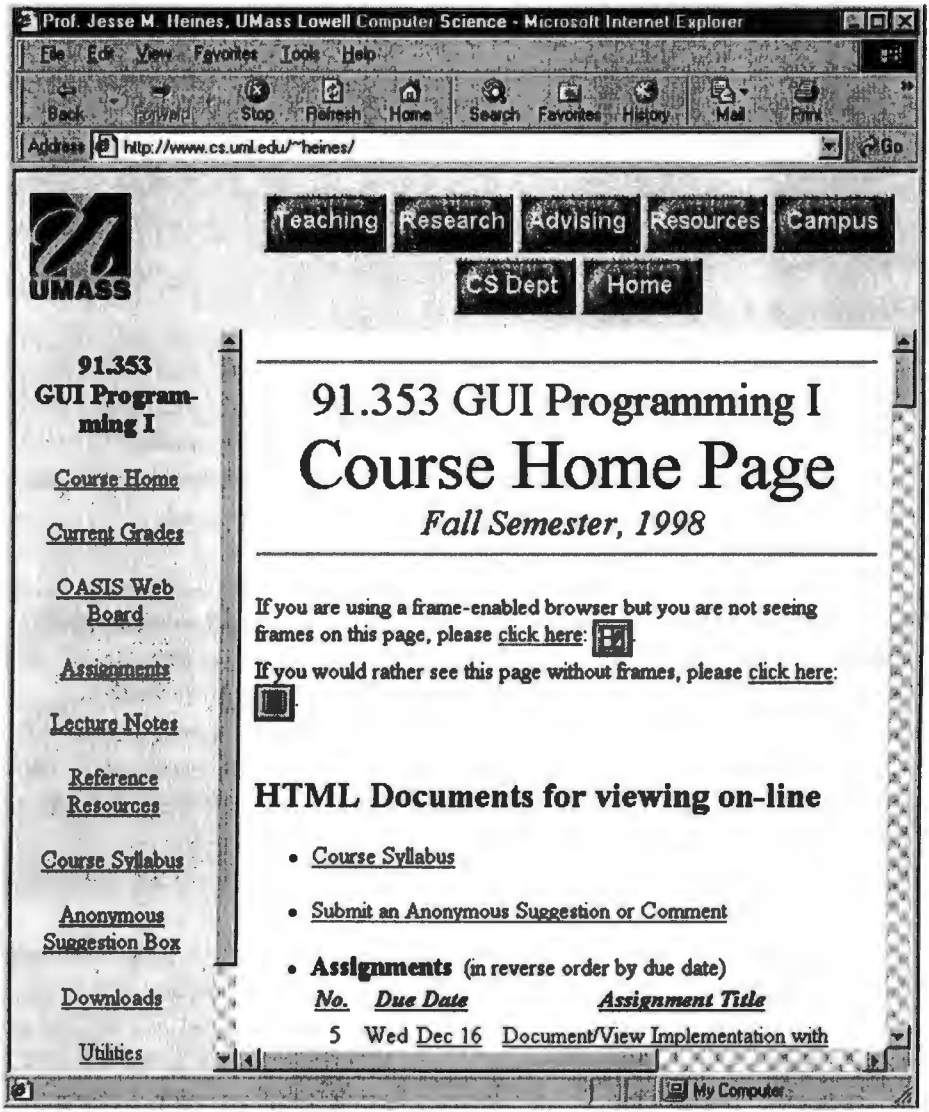

Figure 1. UMass Lowell 91.353 course home page

Notes were typically posted the day before each lecture and students often came to class with a printout of the notes. The actual lecture seldom followed the notes exactly due to shifts in response to student questions and other unforeseen factors. After each lecture, the professor updated that day's notes to accurately reflect exactly what was covered. Thus, the notes became a sort of journal of the course's dynamic flow.

ANONYMOUS SUGGESTION BOX

When the course Web site first "went public" in fall 1996, students didn't "hit" it as often as expected. The professor therefore tried to think of "carrots" that he could add to the site to encourage students to hit it. His first idea was an Anonymous Suggestion Box that allowed students to send e-mail to the professor without identifying themselves (see Figure 2).

This feature was not heavily used. The professor received no more than a half-dozen anonymous e-mail messages in any one semester. Students' main reason for e-mailing the professor was usually to get assistance, and he was unable to respond if they sent e-mail anonymously. Indeed, the professor received and replied to more than 400 student e-mail messages during the spring 1998 semester, and students ranked the Anonymous Suggestion Box number 7 out of 8 when asked which Web site components were most valuable.

GRADE DISPLAY PROGRAM

The professor's second "carrot" idea was to create a program that allowed students to view their status in the course on-line. This feature turned out to be extremely popular, and students ranked its value second only to the lecture notes.

Figure 3 shows the form students completed to gain access to their grades. Students entered their ID numbers, indicated the course in which they were enrolled, and clicked the "Show My Grades" button. Student ID numbers certainly didn't provide much security, but no students indicated that this caused any concern about other students seeing their grades.

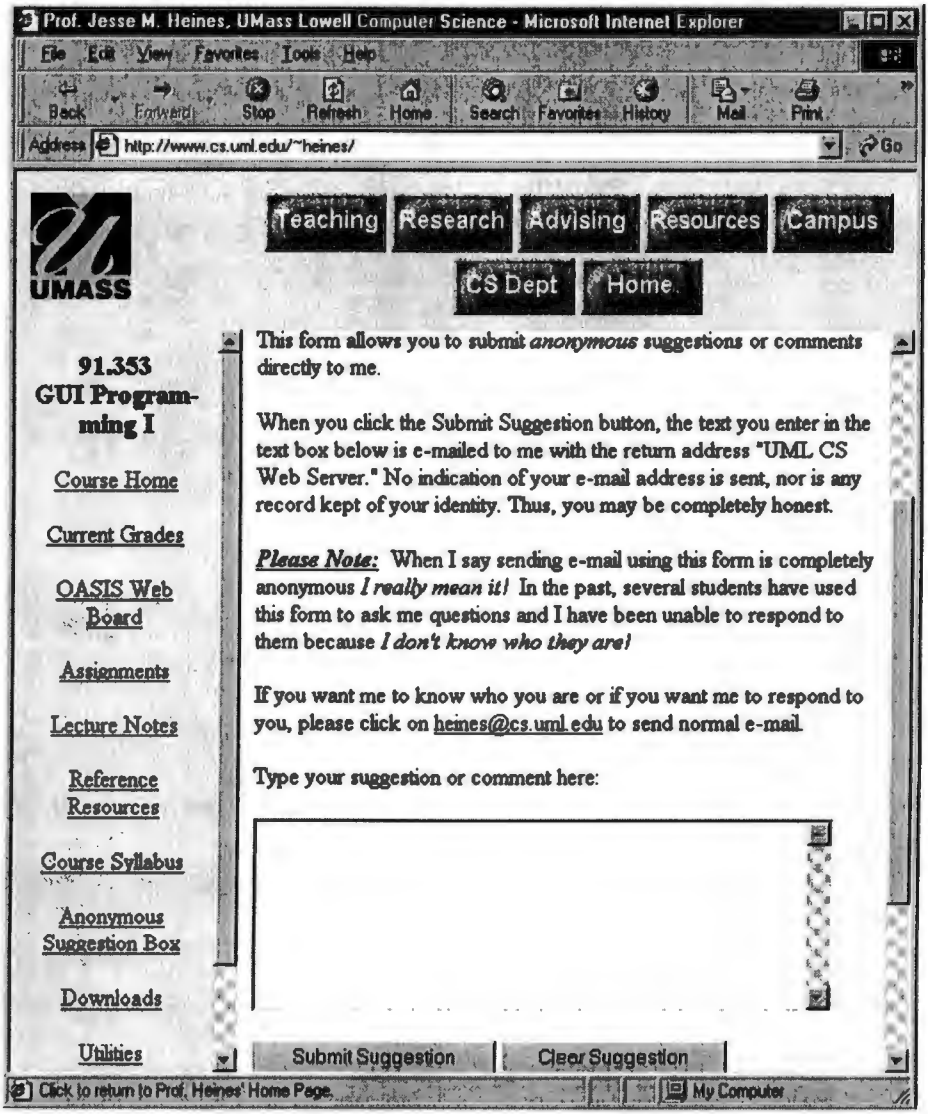

Figure 2. Anonymous suggestion box

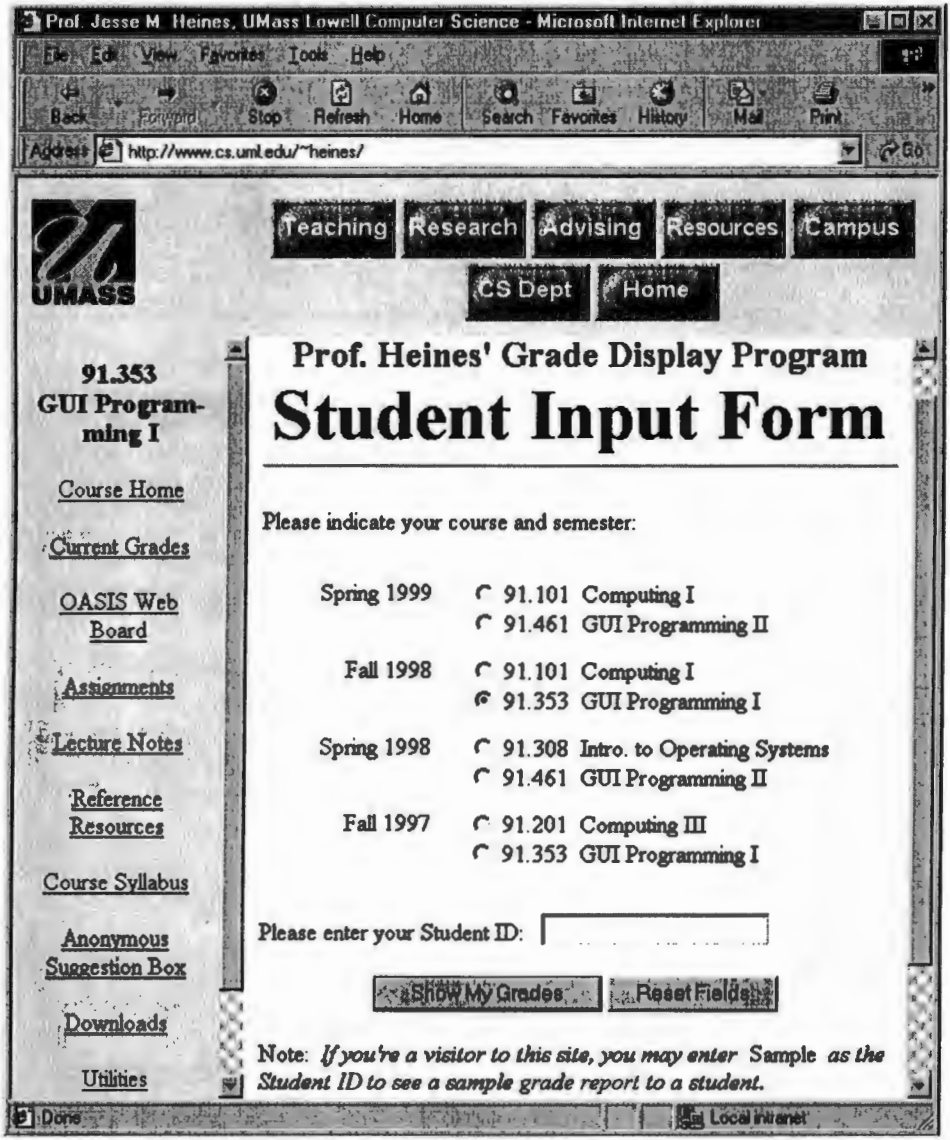

Figure 3. Student input form for the grade display program

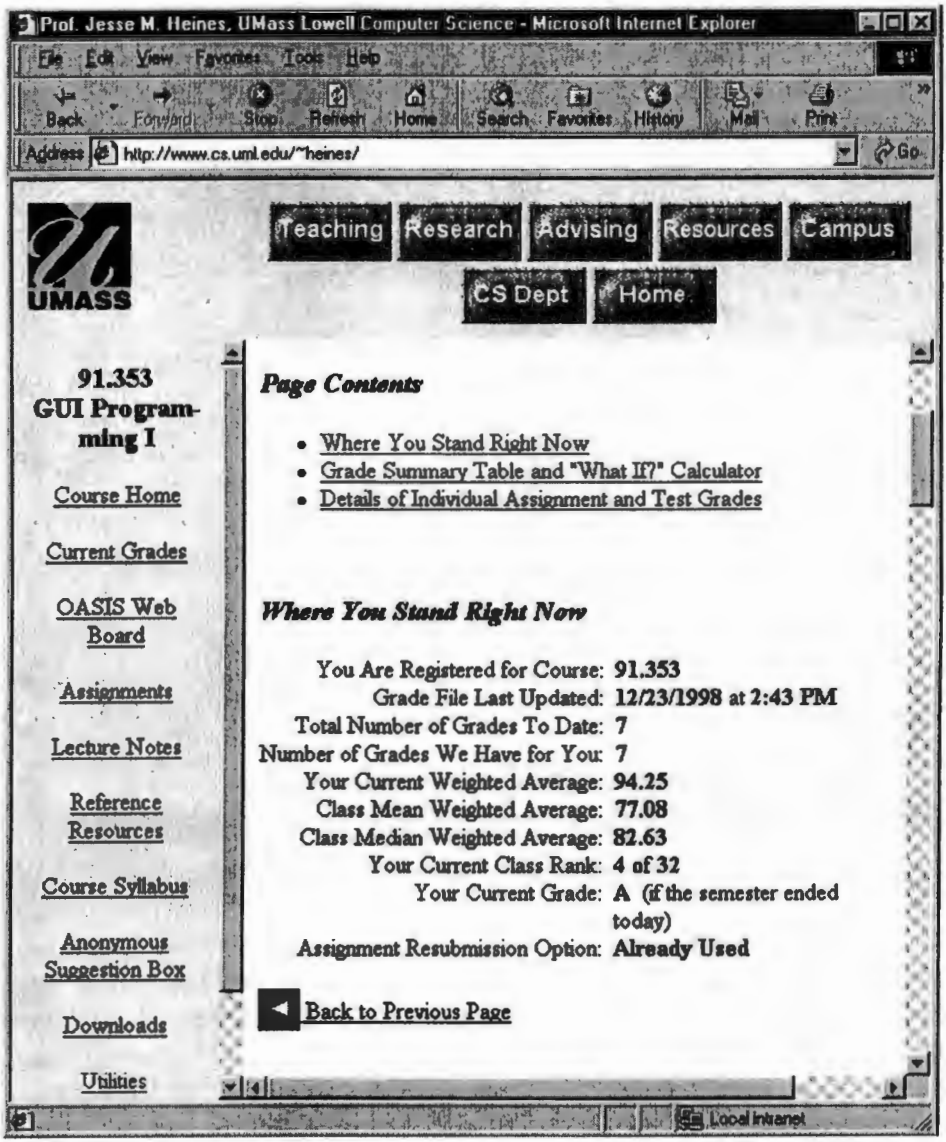

Figure 4. Grade display program: "Where You Stand" section

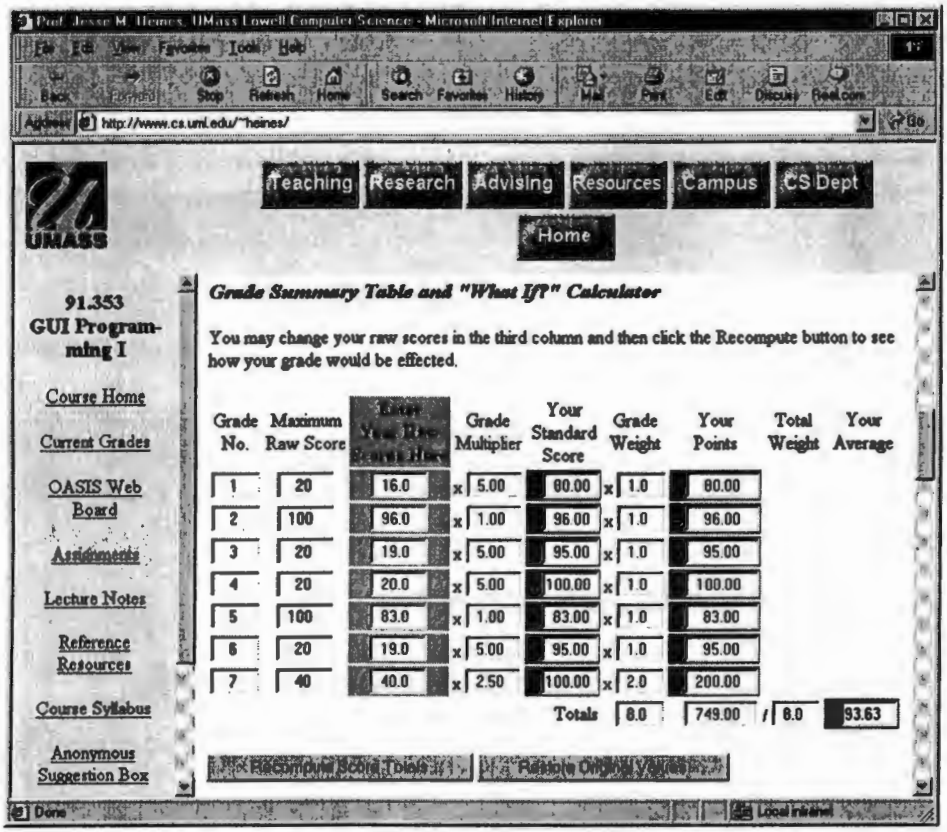

Figure 5. Grade display program: "What If" calculator section

Figure 4 shows the first part of the grade display for a sample student. These data show the student's overall status and where he or she stands relative to the class as a whole. These data were read from the same Excel spreadsheet the professor used to store grades, so the report that students saw was always up to date. The program ran on a system running the Microsoft Personal Web Server, and the spreadsheet could not be accessed directly by students.

Figure 5 shows the second part of the Grade Display Program, the "What If?" Grade Calculator. This is a JavaScript program that clearly shows students the effect of missed homework assignments

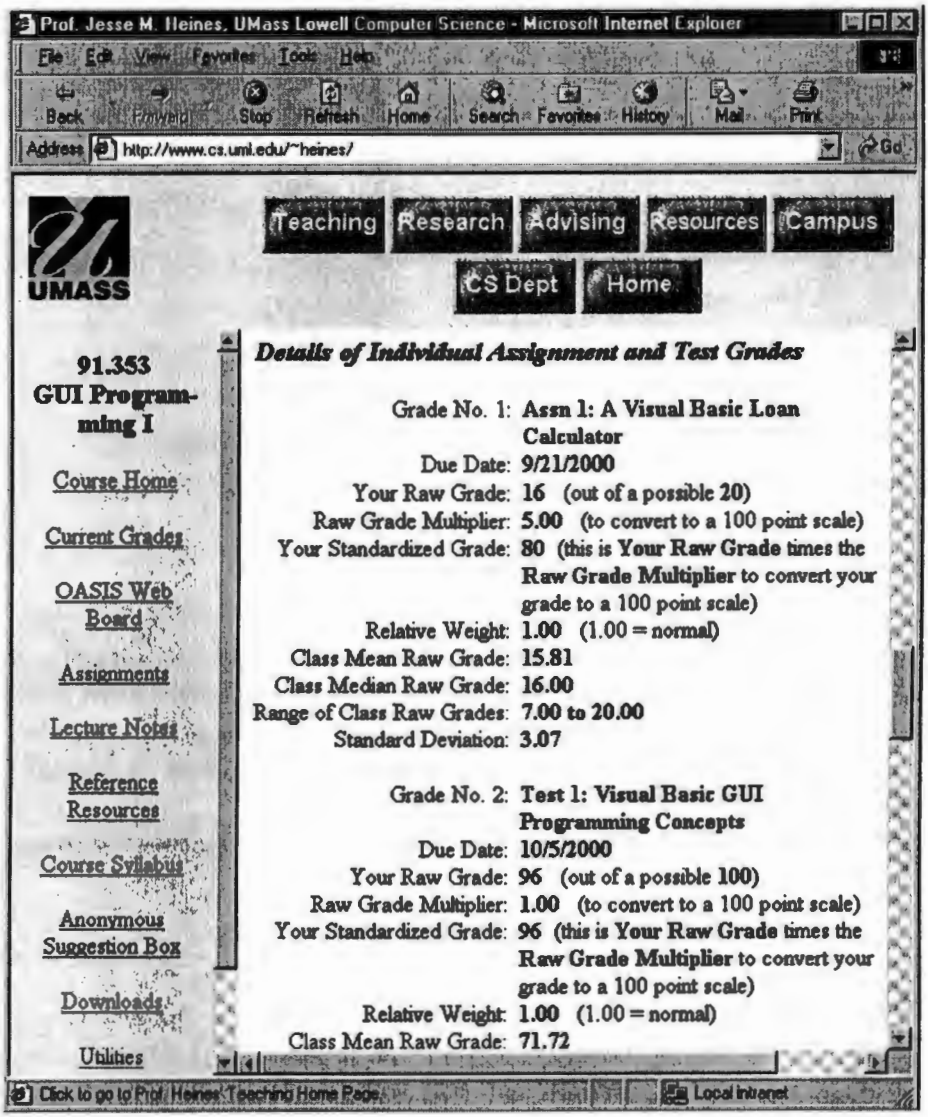

Figure 6. Grade display program: Details of individual grades section

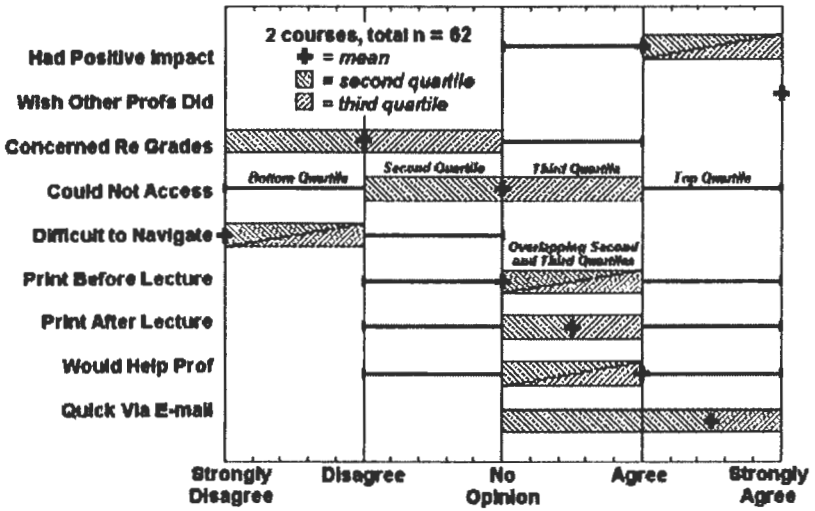

Figure 7. Likert scale survey responses

that were scored as 0. The calculator allowed students to determine how their grade would be effected if they submitted a missed assignment late or resubmitted an assignment to improve its grade. The Details of Individual Assignment and Test Grades section shown in Figure 6 provided all of the data stored for each grade awarded in the course so that students knew exactly how their grades were computed.

STUDENT SURVEY ANALYSIS

AS IN MANY OTHER STUDIES, evaluation of the site began with a student survey. Sixty-two of the 84 students enrolled in the two courses taught by the professor during the 1997 fall semester voluntarily completed an on-line survey.

The survey first presented Likert scale questions to learn about students' attitudes toward the Web site. The texts of these questions and a tabulation of student responses are shown in Table 1. These data are represented graphically in the box-and-whisker quartile chart in Figure 7.

Table 1.
Student Responses to Likert Scale Survey Questions 4-12

	Question	Responses				
		SA	A	NO	D	SD
4.	The presence of course materials on the Web had a significant impact on my ability to keep up with this course.	1	2	4	30	25
5.	I wish other professors maintained course Web sites like this one.	0	1	0	10	51
6.	I was concerned that others could see my grades on the Web site.	20	21	20	1	0
7.	There were times when I could not access the Web when I wanted to.	2	27	5	25	3
8.	I found the Web site difficult to navigate.	40	18	3	1	0
9.	I like to print out the lecture notes before attending the lecture.	1	14	23	15	9
10.	I like to print out the lecture notes after attending the lecture.	2	7	22	22	9
11.	I would be willing to help the professor develop the course Web site for other students.	0	5	24	18	13
12.	The professor was quick to answer questions asked via e-mail.	0	0	15	13	33

Note: SA = Strongly Agree, A = Agree, NO = No Opinion, D = Disagree, and SD = Strongly Disagree

One can see that students overwhelmingly agreed or strongly agreed with the statement, "The presence of course materials on the Web had a significant impact on my ability to keep up with course" (the first item in the chart), and virtually all students strongly agreed with the statement, "I wish other professors maintained course Web sites like this one." Additional questions on the survey revealed that most students (31 of 58 responding) accessed the Web site 3-5 times per week and that about half (29 of 60 responding) printed out materials from the site greater than 10 times.

Given the popularity of the Grade Display Program and the increased number of "hits" that it generated, the professor expected this feature to be mentioned most often. He was pleasantly surprised to see that students actually valued more highly having the lecture notes available. Not only did students mention the lecture notes more often as the best thing about the Web site, but they also claimed to access the Web site more often to see the lecture notes than to use the Grade Display Program.

The survey also asked students to rank the eight major site components in order of their value, with 1 indicating the most valuable and 8 the least. Data on student responses to this question are presented in Table 2. These data corroborate those presented in the previous two sections: that students valued most having the course lecture notes on-line, followed by the availability of the Grade Display Program and the assignments (note that ranks were determined by the mode, not the mean).

STUDENT GRADE ANALYSIS

IMPORTANT ANALYSIS ISSUES

ONE MUST RECOGNIZE that favorable student reaction to a course Web site is not evidence that they actually learn better because of it. It is virtually impossible to prove conclusively that students learn better as a result of any application of technology

Table 2.
Student Rankings of the Value of Web Site Components in Mode Order

Rank	Component	Mean	Median	Mode
1.	Lecture notes	1.9	1	1
2.	Grade display program	3.1	3	2
3.	Assignments	2.6	3	3
4.	Course syllabus and calendar	4.0	4	5
5.	File downloads	4.6	5	5
6.	Industry news items	6.2	6	6
7.	Anonymous suggestion box	6.6	7	8
8.	Links to external sites	6.7	7	8

due to the large number of uncontrollable variables in such studies, not to mention the numerous extraneous conditions that influence results such as the Hawthorne (novelty) and Experimenter Bias Effects (as cited in Biehler, 1971, p. 48), which were also recognized in the study by Dehoney and Reeves (1999).

In addition to these effects, it is difficult in today's academic environment to give two sets of students different versions of a course and compare their results in a tightly controlled manner. One can get approval to conduct small studies on single lessons unrelated to core curricula, but campus ethics and students rights groups are loathe to approve studies in core subjects in which one class is taught in a manner that most believe will enhance their learning while another is deprived the benefit of that approach so that they can serve as a control group. It is even more difficult to get approval to compare one professor's class against that of another professor, for fear that such analysis will reveal the teaching deficiencies of one professor or the other.

Thus, one has to look for situations that already exist that one can use to compare performance. Such situations will virtually always be flawed from a purely statistical point of view, but if one interprets results in a limited scope as Ehrmann (1995) suggests, much can be learned about improving teaching and learning in a local environment

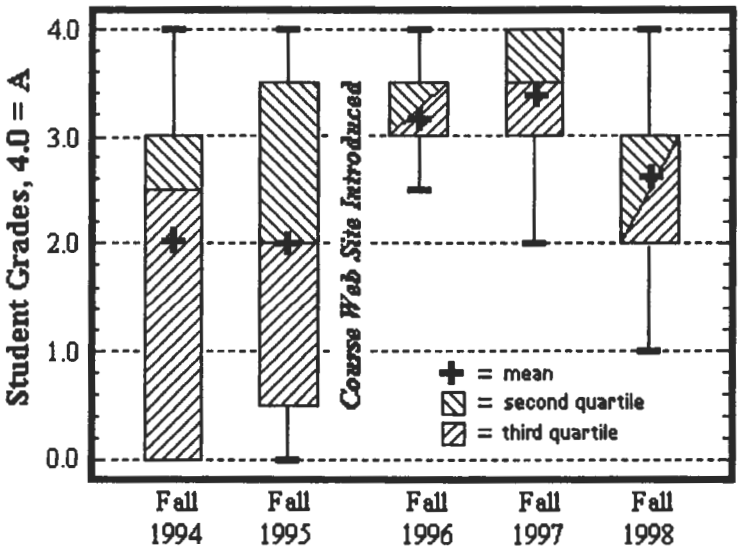

Figure 8. Final student grades for the same course over five years, before and after the introduction of the course Web site

statistically significant difference in student performance after the Web site was introduced into the course. Fortunately, such a pre-existing situation existed at University of Massachusetts Lowell when the current study was undertaken. That situation was carefully analyzed to yield the results presented in the next section.

ANALYSIS OF FINAL STUDENT GRADES

Beginning in 1994, the same computer science professor had taught the same course, 91.353 GUI Programming I, during five successive fall semesters. The course was taught without a Web site in 1994 and 1995, and with a Web site from 1996 on. Except for the addition of the Web site and the updating of some of the technical material, the course was essentially the same in each of the five semesters. A box-and-whisker quartile chart showing the distribution of final student grades in each of the five years is shown in Figure 8.

Table 3.
ANOVA Results of Final Student Grade Data Presented in Figure 8

Source	Sum of Squares	Degrees of Freedom	Mean Square	F Ratio	P Value
Between groups	40.71	4	10.18	7.65	1.0E-4
Within groups	163.65	123	1.33		
Total	204.36	127			

Note: Since the P-value of the F-test is less than 0.05, there is a statistically significant difference between the means of the various groups at the 5.0% significance level (Statlets interpretation).

Analysis of variance results on the data in Figure 8 are presented in Table 3, and analyses of their ranges are shown in Table 4. These tables confirm what the professor had hoped: that there was a statistically significant difference in student performance after the Web site was introduced into the course.

ANALYSIS OF THE PROBABILITY OF SUCCESS

On April 16, 1997, the author had the opportunity to discuss this research with Dr. Ben Shneiderman, the University of Maryland's renowned Human Factors researcher. While lending his insights to the topic and his advice on data analysis, Shneiderman suggested that the author research analytical techniques for focusing on the performance of the lower half of the class, that is, the poorer students. Shneiderman's reasoning was that good students will learn regardless of the instructional techniques employed. Poor students, however, Shneiderman felt might show marked differences in performance when something like a course Web site is made available as an enhancement to traditional classroom instruction.

Given this direction, the author enlisted the help of Dr. Shelley Rasmussen, of the University of Massachusetts Lowell, to devise such an analysis. Rasmussen (1992) suggested computing an estimated "probability of success," defined as the ratio of the number of "suc-

Table 4.
Analysis of Ranges of Final Student Grade Data Presented in Figure 8

Contrast	Difference	+/− Limits	Significant
Fall 1994 - Fall 1995	0.038	0.646	
Fall 1994 - Fall 1996	− 1.128	0.740	*
Fall 1994 - Fall 1997	− 1.348	0.607	*
Fall 1994 - Fall 1998	− 0.587	0.603	
Fall 1995 - Fall 1996	− 1.166	0.751	*
Fall 1995 - Fall 1997	− 1.387	0.620	*
Fall 1995 - Fall 1998	− 0.625	0.617	*
Fall 1996 - Fall 1997	− 0.220	0.718	
Fall 1996 - Fall 1998	0.542	0.714	
Fall 1997 - Fall 1998	0.762	0.575	*

Note: Course Web site introduced for the Fall 1998 class. * denotes statistically significant difference at the 5.0% significance level using Fisher's least significant difference procedure (Statlets interpretation).

cessful" students to the total number of students. Assuming random sampling, the standard deviation of this ratio could be approximated by its standard error. Thus, if the estimated probability of success is

$$\hat{p} = \frac{number\ of\ successes}{total\ number\ of\ cases}$$

then the standard error of $\hat{p}$ is

$$SE(\hat{p}) = \sqrt{\frac{\hat{p}\,(1-\hat{p})}{n}}$$

A traditional rule of thumb cautions that these formulae are "reasonable to use if the numbers of observed successes and failures is at least five." The problem, then, is to define what constitutes "success." Unfortunately, there was no single final letter grade at which the data could be cut so that at least five students fell into the "success" and "failure" categories in all four years. Setting the success criterion as "BC or Better" resulted in only 1 student in the "C or Worse" category for 1997. Setting the criterion as "B or better" re-

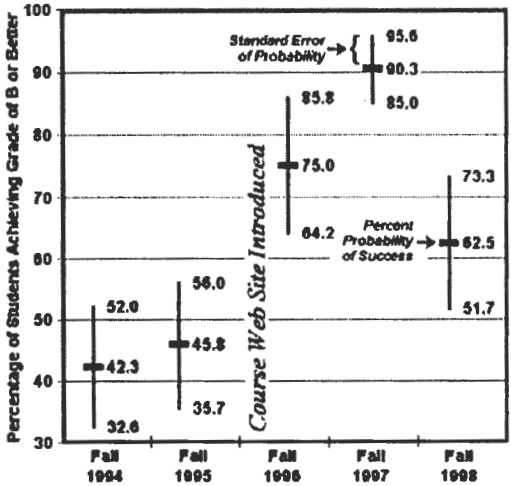

Figure 9. Probability of B or better grade success (same course over five years)

sulted in only 3 students in the "BC or Worse" category for 1997. And setting the criterion as "AB or Better" resulted in only 3 students in the "AB or Worse" category for 1994. Thus, statistical significance on this measure cannot be established on this dimension for the data used in this study. Nonetheless, Figure 9 shows these data presented using a success criterion of "B or Better" to provide a feel for the type of analysis that Shneiderman suggested, and one that would be reasonable to use on larger samples.

USING LARGER SAMPLES

To address the need for a larger sample to apply Rasmussen's techniques, the author analyzed final grades over 10 semesters in Computing I, which enrolled a total of 1,451 students over that time. These courses historically experience high dropout and failure rates and provide an excellent platform for testing Shneiderman's hypothesis. Prior to the fall 1998 semester, Computing I was taught without a course Web site. The author implemented a course Web site modeled after the 91.353 Web site and used it to enhance his teaching of this course in the fall 1998 and spring 1999 semesters. He then obtained historical grade data from the university registrar to use in the analysis that follows.

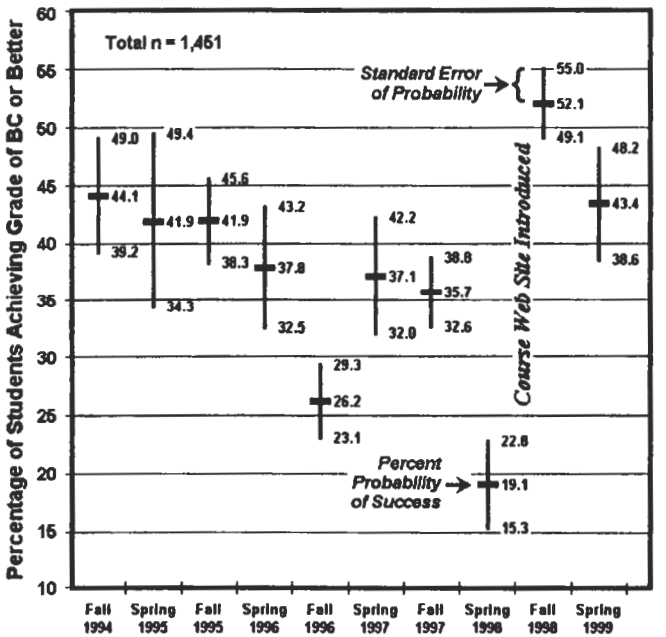

Figure 10. Probability of BC or better grade success in Computing I over ten semesters

Analysis using Rasmussen's "probability of success" technique for this data is shown in Figure 10. While one would like to attribute the huge increase in success from the spring 1998 to fall 1998 semesters to the introduction of a course Web site, this would be irresponsible even with the large number of subjects. The professor changed between these semesters, and indeed the professor who taught after the course Web site was introduced was the same professor who taught the fall 1994 semester. However, it is clear that there is *some* factor that affects students' probability of success in different semesters.

While it is not possible to identify a single cause for the grade differences from year to year, it is possible to establish their significance. Figure 11 shows a box-and-whisker plot of the mean grades over the ten semesters studied. The intervals surrounding the mean are based on Fisher's least significant difference (LSD) procedure. According to the

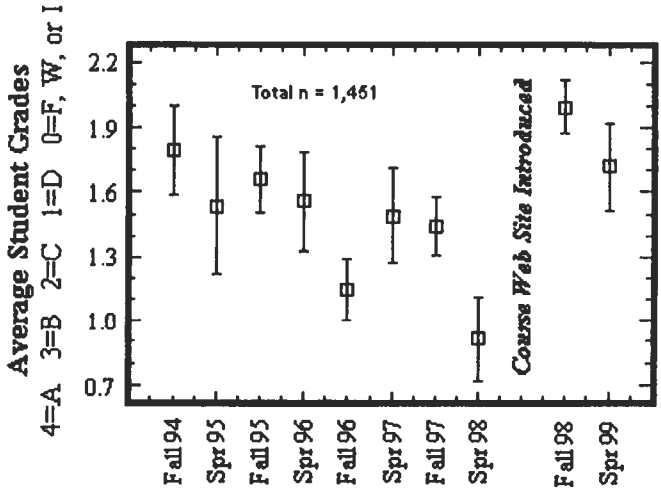

Figure 11. Means plot for 10 semesters of Computing 1

interpretation provided in the Statlets[7] statistical program, the intervals "are constructed in such a way that if two means are the same, their intervals will overlap 95.0% of the time."

Given these data, we can use Dunnett's procedure to compare the means of different groups against a control. Table 5 shows the number of students in each semester and compares the mean of each semester to that of the fall 1998 and spring 1999 classes. While these data cannot conclusively establish that the presence of a course Web site invariably increases student performance, it is at least encouraging to note that comparisons that do yield statistically significant results are between classes with very large enrollments. That is, these classes had enrollments of 202, 241, and 290 students, respectively, and the fall 1996 and 1997 classes yielded statistically significant means from the fall 1998 class, while the fall 1996 class also yielded a statistically significant mean from the spring 1999 class.

Table 5.
Comparison of Final Grade Means to that of the Fakkl 1998 Class as a Control Using Dunnett's Procedure

Contrast This Semester	N	to Fall 1998 Semester			to Spring 1999 Semester		
		Diff.	Limits	Signif.	Diff.	Limits	Signif.
Fall 1994	102	- 0.202	0.494		- 0.202	0.494	
Spring 1995	43	- 0.456	0.702		- 0.456	0.702	
Fall 1995	186	- 0.330	0.403		- 0.330	0.403	
Spring 1996	82	- 0.436	0.537		- 0.436	0.537	
Fall 1996	202	- 0.845	0.393	*	- 0.845	0.393	*
Spring 1997	89	- 0.502	0.520		- 0.502	0.520	
Fall 1997	241	- 0.545	0.374	*	- 0.545	0.374	
Spring 1998	110	- 1.077	0.481	*	- 1.077	0.481	*
Fall 1998	290	n/a	n/a		0.274	0.487	
Spring 1999	106	0.274	0.487		n/a	n/a	
Total	1,451						

Note: Course Web site introduced for the Fall 1998 class. * denotes statistically significant difference at the 5.0% significance level using Dunnett's procedure (Statlets interpretation).

CONCLUSIONS

EVALUATION OF A COURSE WEB SITE is clearly an inexact science, but it can be seen that analysis is possible within limited parameters and when the scope of interpretation of results is well defined. Work such as that by Trochim (1996) may shed light on which analysis techniques are better than others for specific purposes, but it will remain difficult to control all variables that can influence student performance, particularly in university core courses.

Nonetheless, the author is encouraged not only by his students' positive feedback on the value of course Web sites, but also by the cautiously positive effects on their performance indicated by the results of statistical analysis. While one cannot generalize these results to all courses taught by all professors, one can surely state that contrary to the conclusion drawn by Jordan (1997), creation and maintenance of the course Web site is indeed worth the extra time and effort that it entails.

ACKNOWLEDGMENTS

The author is deeply indebted to Dr. Shelley Rasmussen, Associate Professor, University of Massachusetts Lowell Department of Mathematics, for providing critical assistance and advice on analyzing, interpreting, and displaying the data presented in this study. He also wishes to thank Mr. Brian O'Neill, former Director of Computing, University of Massachusetts Lowell Department of Computer Science, for providing invaluable assistance and advice in so many phases of the Web site implementation.

The graphs included in this paper were enhanced from ones generated by Statlets (http://www.statlets.com), a statistical package available from NWP Associates, Inc., of Princeton, NJ, and provided free of charge to academic institutions. The Microsoft software used in this study was supplied in part by a grant from the Microsoft Academic Cooperative (http://www.academicoop.com/).

ENDNOTES

[1] MadDuck Technologies, http://www.madduck.com/wcbinfo/wcb.html.

[2] WBT Systems, http://www.wbtsystems.com/products/products.html.

[3] Microsoft Corporation, http://www.microsoft.com/office/powerpoint/ and http://www.microsoft.com/office/word/.

[4] Microsoft Corporation, http://www.microsoft.com/frontpage.

[5] Adobe Systems, http://www.adobe.com/products/pagemill/main.html.

[6] Allaire Corporation, http://www.allaire.com/Products/HomeSite/.

[7] See http://www.statlets.com/.

[8] See http://www.tltgroup.org/.

[9] See http://www.aahe.org/.

REFERENCES

Alstrum, V. L., Dale, N., Berglund, A., Granger, M., Little, J. C., Miller, D.M., Petre, M., Schragger, P., & Springsteel, F. (1996). Evaluation: Turning technology from toy to tool. *ACM SIGCSE Bulletin*, 28, 201-217.

Biehler, R. F. (1971). *Psychology applied to teaching*. Boston, MA: Houghton-Mifflin Company.

Dehoney, J., & Reeves, T.C. (1999). Instructional and social dimensions of class Web pages. *Journal of Computing in Higher Education*, 10(2), 19-41.

Ehrmann, S. C. (1995). Asking the right questions: what does research tell us about technology and higher learning? *Change*, 27(2), 20-27.

Fitzelle, G. T., Jr., & Trochim, W. M. K. (1996, November). *Survey evaluation of Web site instructional technology: does it increase student learning*? Paper presented at the Annual Conference of the American Evaluation Association, Atlanta, GA. [Online] Available: http://trochim.human.cornell.edu/webeval/webques/webques.htm.)

Gilbert, S.W. (1998). *Lazy students & faculty*? Posting to the Teaching, Learning, and Technology[2] ListServ of the American Association for Higher Education[9] dated February 5, 1998, message number AAHESGIT 2.

Jordan, D. (1997). *Evaluation of a Web site as a teaching tool in MMW-1* (track C). Unpublished manuscript, University of California, San Diego.

Kulik, C.C., & Kulik, J.A. (1991). Effectiveness of computer-based instruction: An updated analysis. *Computers in Human Behavior*, 7(1-2), 75-94.

Rasmussen, S. (1992). An introduction to statistics with data analysis. Brooks/Cole Publishing Company, Pacific Grove, CA. Also personal conversations, Spring 1998.

Trochim, W. (1996, November). *Evaluating Web sites: Introduction.* Annual Conference of the American Evaluation Association, Atlanta, GA. [ONLINE] Available: http://trochim.human.cornell.edu/webeval/intro.htm.)

Williams Glaser, C. (1998). *Creating New Standards for Higher Education: Effecting Pedagogical Change in the Undergraduate Curriculum through the Integration of Technology*. Unpublished doctoral dissertation, Peabody College of Education and Human Development, Vanderbilt University.

ABOUT THE AUTHOR

Jesse M. Heines is an Associate Professor in the Department of Computer Science at University of Massachusetts Lowell. He has been on sabbatical for the 1999-2000 academic year, serving as a Visiting Scholar at the Massachusetts Institute of Technology's Center for Educational Computing Initiatives. Heines specializes in the implementation of graphical user interfaces (GUIs) in interactive programs. He is the owner of CBT Artisan, Inc., a consulting firm founded in 1985 that provides training and programming services for GUI-based applications. Author's present address: University of Massachusetts Lowell, Dept. of Computer Science, One University Avenue, Lowell, MA 01854. E-mail: heines@cs.uml.edu. Home page: http://www.cs.uml.edu/~heines.