IN SEARCH OF THE "PERFECT" BLEND BETWEEN AN INSTRUCTOR AND AN ONLINE COURSE FOR TEACHING INTRODUCTORY STATISTICS

Marsha Lovett, Oded Meyer and Candace Thille
Carnegie Mellon University, United States of America
cthille@andrew.cmu.edu

As part of the Open Learning Initiative (OLI) project, Carnegie Mellon University was funded to develop a web-based introductory statistics course, openly and freely available to individual learners online and designed so that students can learn effectively without an instructor. In addition, the course is often used by instructors in the hybrid form, to support and complement face-to-face classroom instruction. This paper documents two studies where we investigated the OLI-statistics courses' effectiveness in the hybrid instructional model. We describe the design, results and limitations of the studies and discuss the implication of the results for finding the "perfect" blend between an instructor and an online course for teaching introductory statistics.

INTRODUCTION

As part of the Open Learning Initiative (OLI) project, Carnegie Mellon University has been funded by The William and Flora Hewlett Foundation to develop an online introductory statistics course. The Foundation's interest is in providing open access to high-quality post-secondary education and educational materials to those who otherwise would be excluded due to geographic, economic or time constraints (Smith & Thille, 2004), as well as for those who due to social barriers are not encouraged to pursue higher education.

Our general approach to the task of developing a stand-alone web-based course was to create a course that would be as close to a *fully online enactment of instruction* as possible. The course was developed by a team composed of learning scientists, statistics faculty members, human-computer interaction experts, and software engineers in order to make best use of multidisciplinary knowledge for designing effective instruction. Meyer and Thille (2006) discuss in detail the learning science principles that motivated the course's design, all of which would be predicted to foster better, deeper learning. For example, the OLI-Statistics course was designed to make clear the structure of statistical knowledge, include multiple practice opportunities for each of the skills students needed to learn, to give students tailored and targeted feedback on their performance, and to effectively manage the cognitive load students must maintain while learning. (see *https://oli.web.cmu.edu/openlearning/forstudents/freecourses/statistics*). Moreover, as students work through the OLI course, we collect real-time, interaction-level data on how they are learning, and we use these data to inform further course revisions and improvements. In addition to this ongoing formative evaluation, we conduct formal learning studies on a regular basis.

The primary goal of the first two studies, conducted in Spring 2005 and Spring 2006, was to test the hypothesis that students would learn at least as much from the OLI-Statistics course in stand-alone mode as they would from traditional, instructor-led instruction. This goal represents a fairly simplified "do no harm" test of the stand-alone version of OLI-Statistics (i.e., students' learning would not be harmed relative to taking Statistics in a traditional face-to-face setting). Based on the results of these studies (see Lovett, Meyer & Thille, 2008) we were very encouraged to discover that when the OLI statistics course was used in the way it was designed to be used (as a stand-alone course), the learning gains of students were at least as good as in a traditional, instructor-led course.

ACCELERATED LEARNING STUDIES

Given that students taking the OLI-Statistics course with little or no instructor support, showed learning gains similar to students taking a traditional course covering the same materials, we sought to explore whether (and how) the OLI-Statistics course could lead students to *better* learning outcomes relative to the traditional course. First, we decided to study the OLI-Statistics course in the context of a hybrid approach in which students would receive both the online OLI materials as well as some face-to-face instruction. The idea here was to first have students work through a specified piece of the online materials on their own and then have the instructor provide

In C. Reading (Ed.), *Data and context in statistics education: Towards an evidence-based society. Proceedings of the Eighth International Conference on Teaching Statistics (ICOTS8, July, 2010), Ljubljana, Slovenia*. Voorburg, The Netherlands: International Statistical Institute. www.stat.auckland.ac.nz/~iase/publications.php

additional face-to-face instruction to address any relevant areas of difficulty. It is important to mention that within the OLI statistics course, not only do students receive tailored feedback based on their actions in the course but *instructors* receive detailed reports on the progress and performance of their students which enables them to target their instruction accordingly. Second, in studying this hybrid version of OLI-Statistics, we decided to assess effectiveness not just in terms of students' learning gains but in terms of an "accelerated learning hypothesis." That is, we wanted to see whether students taking eight weeks of OLI-Statistics in a hybrid mode could learn as much or more than students taking fifteen weeks of a traditional Statistics course. Looking on past research, it is often the case that efficiency is a more frequently used and more sensitive measure of learning (e.g., Rothkopf, 2008). Also, given the future possibility that OLI courses could help students who are at an academic disadvantage "catch up" faster, we were encouraged to test our accelerated learning hypothesis with OLI-Statistics.

Research Design

This paper documents two studies in which we tested our accelerated learning hypothesis in OLI-Statistics. The first was in the Spring of 2007, and the second–a replication with extension–was in the Spring of 2009. The two studies were very similar in their methods, so we describe them together. For each study, approximately 200 students were initially registered for the traditional introductory statistics course at Carnegie Mellon. One month before the semester began, we sent an email to all of these students, inviting them to participate in an accelerated learning study. In the Spring of 2007, 68 students volunteered to participate in the OLI-based course, and we randomly selected 22 of them. Of the remaining 46 volunteers, four students dropped the course before it began, so 42 students served as our primary control group, taking the course in the traditional format. The 130 non-volunteering students from the traditional course served as a secondary control group. In the Spring of 2009, we sought to "scale up" from the 2007 study by doubling the class size for the OLI-based course, so we selected 49 of the students who volunteered that semester. Because this was almost the entire set of volunteers, we did not have a special group of control students within the traditional class.

In the traditional course, besides lasting for the entire 15-week semester, class meetings occurred four times a week for 50-minutes each. Of these four class meetings, three were lectures and the fourth was a computer lab session. The OLI based course, besides lasting only eight weeks, met only twice a week. For each of these two class sessions, the instructor selected material (usually problems to solve or concepts to discuss) based on the OLI system's report of students' difficulties with the recently completed online activities.

Student Learning Measures

We used several measures to assess students' learning outcomes for the Spring 2007 and Spring 2009 studies. First, we compared students' scores on in-class exams. Students in the OLI-Statistics course and in the traditional course took three midterms and a final. All of the tests were matched for content and level of difficulty. Note, however, that the exams were not identical across the two courses. This was a practical decision based on the different timing of the exams in the 15-week versus 8-week version of the course.

Second, we administered the CAOS test as a pre- and post-test to both course sections in both Spring 2007 and Spring 2009. The CAOS test is a Statistics knowledge assessment developed by statistics education researchers (delMas, Ooms, Garfield & Chance, 2006). This test is named CAOS for "Comprehensive Assessment of Outcomes in a first Statistics course", and it is designed to measure students' basic statistical reasoning. It includes 40 multiple-choice items on statistical reasoning in general and targets several difficult concepts in statistics. The CAOS test not only represents a generally accepted measure of statistical literacy, it offers a set of national benchmarks for performance that we used to compare with our OLI-Statistics groups. We administered the CAOS test to the OLI-Statistics students and the traditional students at the beginning and end of the semester in order to assess students' learning gain. Unfortunately, due to unforeseen time constraints in the traditional course's administration of the CAOS pre-test for Spring 2009, many students did not come close to finishing, so for the Spring 2009 study we only have comparison data between the OLI-Statistics course and the traditional course for the CAOS post-test.

Results

In the Spring 2007 study, of the 22 students in the OLI-Statistics course, 21 completed the work and took the final exam. Of the 42 students in the control condition for Spring 2007, 40 took the final exam. In the Spring 2009 study, of the 49 students in the OLI-Statistics course, 46 completed the work and took the final exam. With such similar retention rates across these three contexts looking at students who had volunteered to participate in our studies, there do not appear to be any obvious drop-out rate differences that would affect a comparison of the accelerated OLI-Statistics course and the traditional course.

In-class exams showed no significant difference between the traditional and online groups. That is, students in OLI-Statistics were performing as well as traditional students on in-class exams after having spent approximately half the time learning the material. This was true for both the Spring 2007 and Spring 2009 studies. (Figure 1 below shows the final exam scores for Spring 2007.) This non-significant difference, however, is not surprising given that there are so many variables affecting students' performance on a final exam (e.g., amount of time studying, incoming prior knowledge, and study strategies, just to name a few). Moreover, even a carefully planned final exam does not have any guarantee of validity and reliability. So, we relied on the CAOS test as a more rigorous instrument.

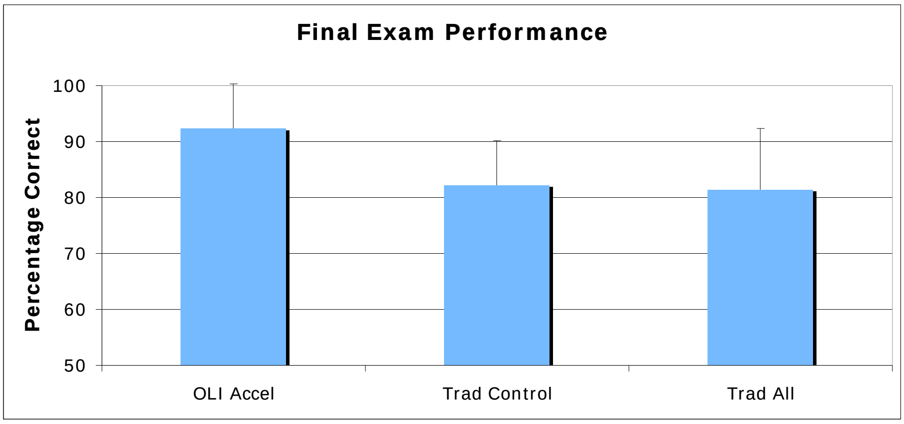

Figure 1. Final exam performance of accelerated OLI-Statistics compared to traditional

For the CAOS test scores in the Spring 2007 study, we assessed not only whether OLI-Statistics students showed significant learning gains across their 8-week course but whether those gains were different in size compared to our traditional control group (see Figure 2). The OLI-Statistics students gained, on average, 18 percentage points from the beginning to the end of the semester, a significant increase, $t(20) = 6.9$, $p < .01$. The control students from the traditional course gained on average only 3 percentage points from the beginning to the end of the semester, $t(39) = 1$, an increase that was not significantly different from zero. Moreover, as these numbers suggest, the size of the learning gain was significantly larger for the OLI-Statistics students compared to the traditional controls, $t(46) = 4.0$, $p < .01$. Similar results were obtained when this analysis is done with the raw pretest and posttest scores submitted to an analysis of covariance (ANCOVA), with pretest as the covariate and group (OLI-Statistics vs. control) as the factor.

For the Spring 2009 study, the OLI-Statistics students gained, on average, 18 percentage points from the beginning to the end of the semester, a significant increase, $t(45) = 13.9$, $p < .01$. They started the semester with 54% correct, on average, and ended with 72% correct. Although, we do not have a pre-to-posttest learning gains comparison with the traditional students for Spring 2009, we can compare the posttest results for the OLI-Statistics and traditional students. Compared to 72% correct at post-test among the OLI-Statistics students, the traditional students averaged only 59% correct, a significant difference, $t(105) = 5.3$, $p < .01$. Moreover, the learning gain results for the OLI-Statistics students in Spring 2009 are not significantly different from the OLI-Statistics students' learning gains in Spring 2007. In fact, the 2009 learning gains are strikingly similar to those from 2007 even though the class size was more than double.

OLI Accelerated	n	Average % Correct
Pre	21	55
Post	21	73
Increase: 18 percentage points		
$t(20) = 6.9$, $p < .001$		

Traditional Control	n	Average % Correct
Pre	40	50
Post	40	53
Increase: 3 percentage points		
$t(39) = 1$, not significant		

Figure 2. CAOS results for accelerated OLI-Statistics versus traditional control

The above results from the Spring 2007 and Spring 2009 accelerated learning studies have shown that the OLI-Statistics students obtained learning outcomes that were as great or greater than those of the traditional course students. In this sense, our accelerated learning hypothesis was supported: students in OLI-Statistics learned 15 weeks' worth of material as well or better than traditional students in a mere 8 weeks. However, it is still possible that students in the OLI-Statistics course were actually making up for lost time by spending twice (or more) study time per week compared to the traditional students. While there is no particular reason to suspect this, we wanted to verify that it was not the case.

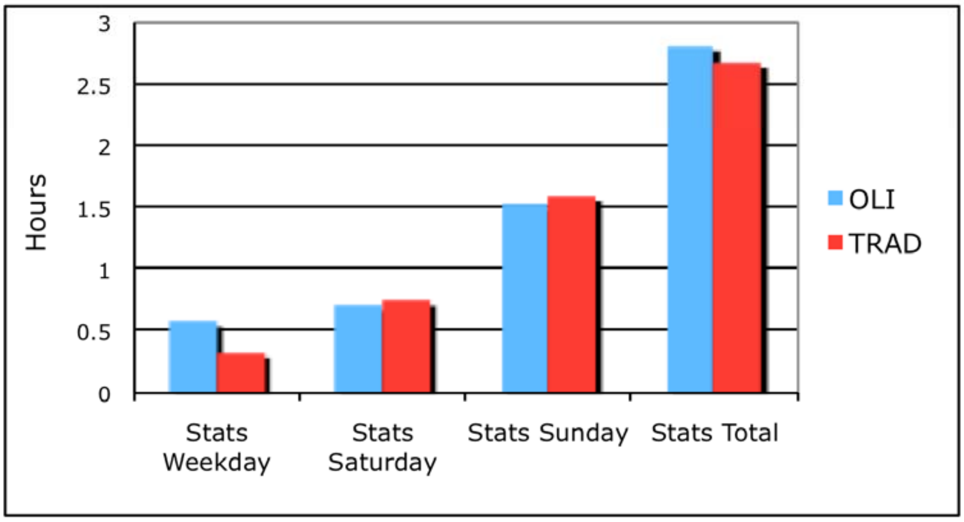

Figure 3. Outside-of-class time data from both groups of students

Figure 3 shows the average self-reported amount of time that students spent on Statistics outside of class in both the OLI-Statistics group and the traditional group in Spring 2007. The first three pairs of columns show students' time broken down by "weekday" and each weekend day, and the rightmost pair of columns gives the total time for the six days students were surveyed. Several things are worth noting about these data. First, there is almost no difference between the two groups in their total time spent per week. This suggests that even though OLI-Statistics students were covering approximately twice the material in a given week, they were *not* spending twice the time learning it. Thus, the learning outcomes results presented above document a significantly more efficient learning experience among the OLI-Statistics students, confirming our accelerated learning results. (Note that OLI-Statistics students' in-class time was exactly half that of the traditional students, with two instead of four 50-minute class meetings per week.) Second, although it is not statistically significant, the case where OLI-Statistics students spent more time studying statistics is during the week (more than one hour per weekday compared to about a half hour per weekday). This result suggests that the OLI-Statistics course (at least as it was conducted in this study) may lead students to spread their study time more evenly rather than cramming study time into long weekend sessions.

RETENTION COMPONENT OF THE ACCELERATED LEARNING STUDY

Retention Study motivation and method

Because the results of the Spring 2007 study were so encouraging–namely, students in OLI-Statistics took half the time to learn as much or more than their traditional counterparts–we sought to extend the study by conducting a retention follow-up study that would test students' abilities to retain and use what they learned during Spring 2007 at a considerable delay. This retention study was also designed as an authentic assessment of students' learning by testing what they retained by the beginning of the following semester, i.e., precisely when they would be expected to build on their previous knowledge. So, at the beginning of the following semester (Fall of 2007), we invited students from both groups (the OLI-Statistics students and the traditional control) to participate in an additional study for pay. This additional study included three activities: taking the CAOS test again, solving open-ended problems from introductory statistics, and learning a new topic (and answering questions about it). It is worth noting that the OLI-Statistics students, who had finished their statistics course at the beginning of March 2007, completed the retention study at a 7-month delay whereas the traditional students, who finished their statistics course in the middle of May 2007, completed the retention study at a 5-month delay. So, even if students' memory decay rate was equivalent during this time period, we might expect somewhat lower performance among the OLI-Statistics students.

Before presenting the results for the three activities in this retention study, we should note an important practical challenge we encountered. Out of the 60 students we emailed to invite to participate in the study, only eleven students responded and completed the retention activities. Conveniently, they were almost evenly balanced between the two groups, with six OLI-Statistics students and five traditional students. Nevertheless, we must take the following results as merely suggestive because of the small sample size.

Retention Study results

For the CAOS test, we found no significant difference between the two groups (Accelerated OLI-Statistics group averaged 72% correct; traditional controls averaged 67% correct). Even without finding a difference between groups, it is interesting to note that students' retention scores tracked their Spring 2007 posttest scores rather well (70% and 66% for the corresponding students from the two groups). Such a result is consistent with previous research showing that students who learn more retain more. It also encourages us to expect that with a larger retention sample, we might have been able to show a significant difference in CAOS scores between the OLI-Statistics students and traditional students.

The open-ended problem solving portion of the retention study was scored by a rater who was trained to use a scoring rubric that gave up to a total of 9 points for (1) the accuracy of the solution, (2) the appropriateness of statistical tools used, and (3) the clarity and accuracy of the written interpretation of the statistical results. The rater was blind to participants' condition. With such a small sample, it is not surprising that these scores did not reach statistical significance, $t(11) = 1.6$, $p < .13$. Nevertheless, the OLI-Statistics group scored numerically quite a bit higher: 6.3 versus 3.9. Moreover, it is interesting to note that none of the six OLI students made an egregious error in their answers, whereas two of the five students in the traditional group made a serious interpretive error (for example, reporting that $p < .05$ means that null hypothesis is accepted).

Finally, the third activity in the retention study asked students to read a short passage explaining a new statistical tool, Analysis of Covariance, and then to answer a few conceptual questions about this tool. Accuracy scores on these questions were again scored on a scale from 0 to 9. Results showed no difference between the two groups, with both groups averaging 7 points.

DISCUSSION

The most striking finding in this set of studies is that students in the accelerated OLI-Statistics course were able to learn better and in half the time as compared to students with traditional instruction. Usually, that kind of effectiveness or efficiency effect would be the result of individualized, human tutored instruction (e.g., Bloom, 1984). And yet, we had close to 50 students in a class that met for less than two hours per week, showing such results. The mechanism we posit for this striking result is that the accelerated OLI-Statistics students actually attended their class meetings in a much better prepared state than students usually do. As opposed to skimming (or skipping) the reading before a traditional lecture, our accelerated students prepared for class by actively engaging with the material in numerous ways by completing comprehension checks of their understanding as they read, applying their new skills to problems for practice, receiving tailored feedback on their answers, and reflecting on their own understanding and questions as they proceeded. In this way students came to class ready to make best use of their time with the instructor. And, the instructor came to class better prepared to teach. Thanks to OLI's automatically generated instructor reports, the instructor was able to see reports on student progress, review summaries of students' quiz performance, and read students' reflections and questions about the previous week's material. With this information in hand, he was able to select discussion topics and example problems that targeted the topics with which the students were struggling. Then, class time was spent with students actively engaged on the material that was most likely to need more supported practice or a novel explanation from the instructor.

It is this combination of preparedness of both the students and the instructor, facilitated by the OLI-Statistics course, that we believe is the key to the success of using this course in accelerated hybrid mode. Ironically, the fact that the OLI statistics course was designed as a stand-alone course–making knowledge structures explicit, following as many principles of learning as possible– s likely the reason that it was so successful when used in hybrid form.

Finally, one of the challenges that academic institutions are facing and are hoping to solve by using online education is how to provide effective instruction under limited resources. The more a course is web-based and relies less on an instructor, the more resources are saved. In addition, some colleges do not have statistics experts to teach their introductory statistics courses and instead rely on mathematicians to teach such courses. In such cases, using online instructional support such as OLI-Statistics could provide "pedagogical scaffolding" so that the overall quality of instruction is improved. So, although our main findings involve not just stand-alone online instruction but document the effectiveness of a pedagogically active instructor working *with* the OLI-Statistics course, there are still a lot of resources saved in comparison to a traditional course (e.g., two course meetings per week instead of four). In addition, resources could be saved since the course can be taught in half a semester with no extra time cost to the students and impressive benefits in the form of solid learning gains and substantial retention of the material.

REFERENCES

Bloom, B. (1984). The two sigma problem: The search for methods of group instruction as effective one-to-one tutoring. *Educational Researcher, July,* 4-15.

DelMas, R., Ooms, A., Garfield, J., & Chance, B. (2006). Assessing students' statistical reasoning. *Proceedings of the seventh international conference on teaching statistics*. Voorburg, The Netherlands: International Statistical Institute.

Lovett, M., Meyer, O., & Thille, C. (2008). The open learning initiative: Measuring the effectiveness of the OLI statistics course in accelerating student learning. *Journal of Interactive Media in Educatio.* Online: http://jime.open.ac.uk/2008/14/.

Meyer, O., & Thille, C. (2006). Developing statistical literacy across social, economic and geographical barriers using a "stand-alone" online course. *Proceedings of the seventh international conference on teaching statistics*. Voorburg, The Netherlands: International Statistical Institute.

Rothkopf, E. Z. (2008). Reflections on the field: Aspirations of learning science and the practical logic of instructional enterprises. *Educational Psychology Review, 20*(3), 351-368.

Smith, J. M., & Thille, C. (2004). The open learning initiative: cognitively informed E-learning, *The Observatory on borderless higher education*, October 2004.