Using Active Learning to Increase Student Retention in Introductory Computing Courses

Bruce Herring, Rochester Institute of Technology (GCCIS)

Bruce Herring is a lecturer in the Rochester Institute of Technology's Department of Software Engineering. He holds a MS in Computer Science from Florida State University. At RIT he teaches several introductory courses to incoming freshmen and sophomores in computer science and software engineering. Prior to RIT he worked in industry for 16 years for Harris Corporation, a communications company. His work was predominantly in an embedded environment dealing with the interactions between software and hardware. He has worked to integrate his extensive industry experience into the classroom to provide students with a fuller understanding of how concepts are applied after they leave the university environment.

Robert St Jacques, Rochester Institute of Technology

Robert St. Jacques is a lecturer in the Software Engineering Department within RIT's Golisano College of Computing and Information Sciences. He holds an MS in Computer Science from RIT. Robert spent nearly 20 years as a professional software engineer, with the bulk of that time spent as a member of the research and technical staff at the Xerox Research Center in Webster, NY. He is a contributing author to over 20 US Patents and served as the Director of the Advanced Software Innovation Lab at PARC before joining RIT. Over the last four years, Robert has designed and built courses for both RIT and the University of Rochester covering topics ranging from introductory computing to advanced web and mobile development. Currently, he focuses on teaching courses in the first two years of the core Software Engineering sequence.

Using Active Learning to Increase Student Retention in Introductory Computing Courses

Executive Summary

The introductory computer science sequence in the Golisano College of Computing and Information Sciences (GCCIS) at RIT is required for most incoming freshmen each year. Students that fail to successfully complete the first course in the introductory sequence (i.e. withdraw or receive a grade of D or lower) are at risk for late graduation, and a high percentage of such students end up leaving the college or university. During the Fall 2018 semester, an experiment was conducted involving 4 of the 13 sections of the course (34% of the enrolled students). For this experimental group, active learning was incorporated into what was traditionally a passive learning lecture approach. Active learning was chosen based on its success in increasing outcomes for students in STEM fields [1]. The goal of the experiment was to reduce what has historically been one of the highest failure rates in the university so that a higher percentage of students continue to the next course in the sequence, stay within the college, and ultimately graduate on time.

Changes to the traditional course format included incorporating 10-20 short, pencil-and-paper partnered active learning activities into each lecture, as well as adding an hour-long session of short, instructor-guided hands-on programming exercises in a lab environment.

The results after the first semester showed that students in the experimental sections of the course performed an average of 2.6% higher on assignments, labs, and exams than the general population of students, and students in the pilot sections received final grades that were on average 3.6% higher than the general population. Furthermore, the results of two surveys conducted to assess student satisfaction with the course showed that students in the experimental sections were significantly happier with nearly every aspect of the course.

Rationale

Each year more than 550 incoming freshmen are required to enroll in the introductory computer science sequence within RIT's Golisano College of Computing and Information Services (GCCIS). The student population is divided into many sections, with class sizes ranging from 25 to 50 students. The first course in the sequence, Computer Science I (CS1), is a study of procedural programming, algorithmic design, and data structures using Python. Students are required to pass with a grade of C- or better to continue on to Computer Science II (CS2). Students that withdraw, or receive a grade of D or lower, are required to retake the course. Many instead choose to change programs or leave the university.

The incoming freshman population each year has a very diverse set of skills and experience ranging from students with portfolios of programming projects in various languages to those that have little or no experience with computer science at all. The challenges of designing an introductory computing sequence that accommodates such a diverse population include pacing the course such that students with little or no experience may succeed while at the same time providing learning and growth opportunities to those students that do have experience.

Historically the rate of students withdrawing from or failing the course with a grade of D or lower (WDF) has averaged about 25%, which is among the highest in the university. The primary goal of the pilot conducted during the Fall 2018 semester was to reduce the WDF rate as much as possible with a secondary goal of improving student satisfaction with the sequence. The nature of the course, with tightly coordinated material introduced to students across all sections of the course, provided us with a rare opportunity to apply new teaching techniques to a subset of the sections and then compare the outcomes for those sections against the remaining population.

Traditional Format

Each section of CS1 meets for a two-hour lecture, typically on the first day of each week. The lectures are traditionally a passive learning experience during which most instructors deliver content as some combination of slides, live coding, and drawing or writing on a whiteboard while students observe and take notes. End-of-semester course evaluations have consistently demonstrated that many students find it difficult to maintain focus for the entire lecture.

On the second day of each week, students meet with their instructor for a two-hour lab. During the first hour, referred to as Problem Solving, students are divided into groups of 3 or 4 and given a series of problems to solve together on paper or a whiteboard. At the end of the first hour, the instructor goes over the solution with the class before handing out the lab assignment, which is meant to be completed by each student individually for submission the following week. Students spend the second hour working individually on a few small coding tasks meant to get them started on the lab.

On the third day, students meet for a 1-hour recitation, at the start of which they are given a short quiz to refresh the material covered that week during lecture. This is followed by a brief presentation meant to summarize that week’s lessons, and a 20-25 minute pair programming exercise. Finally, before leaving, students are given another short quiz.

In a typical week of the course, students are given only two opportunities to practice the material that they have learned: a homework assignment and a lab assignment.

Experimental Setup

In the fall of 2018, a total of 574 students enrolled in Computer Science I. The students were distributed across 13 sections of the course. The experiment comprised 4 of the 13 sections, including 197 students (34.3% of the total population). While the content delivered was tightly coordinated and consistent across all sections of the course, the experimental sections included several significant changes to the ways in which the content was delivered to the students, each of which is discussed in detail below.

During the two-hour lecture each week, material was presented using slides, but incorporated 10-20 activities, or about 1 for every 2-3 slides of content. This had the effect of dividing the two-hour period into a series of 5-15 minute segments, each of which was punctuated with one or more short active learning activities meant to be completed by students in pairs. This approach was chosen based on existing research regarding the positive impact of active learning activities on a student's performance [2][3][4]. The primary goal of the activities was to allow students to immediately think about and attempt to apply the information they had just learned to a small problem.

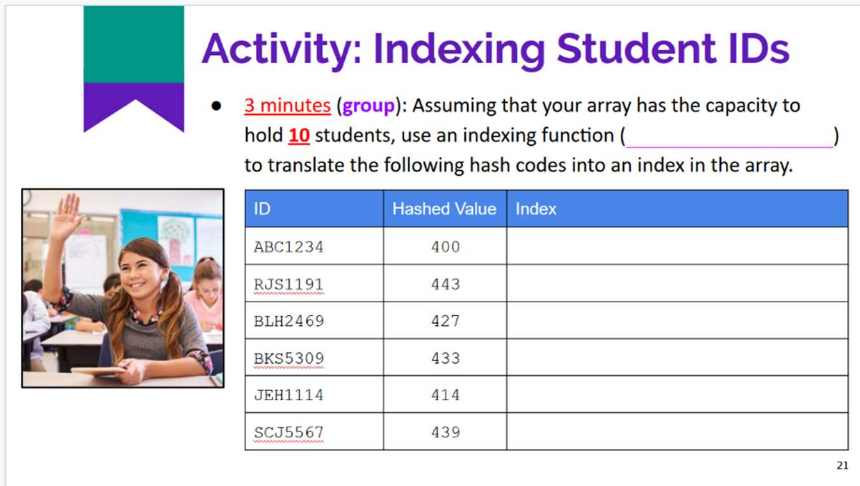

Figure 1: Example in-class activity designed to be completed by a pair of students

Because not all students are equipped with a computer during the lecture, all activities were necessarily designed to be completed using pencil and paper in 1-5 minutes. For example, a 1-minute activity might involve the groups discussing the answer to a question, matching vocabulary words and their definitions, or writing a single line of code. A 2-minute activity might task the students to draw a diagram or write a small block of code. 5-minute activities were reserved for writing larger blocks comprising as many as 10 lines of code. The lectures were designed such that each activity contributed a small part of the solution to a larger problem, often culminating in the development of a small application by the end of the lecture.

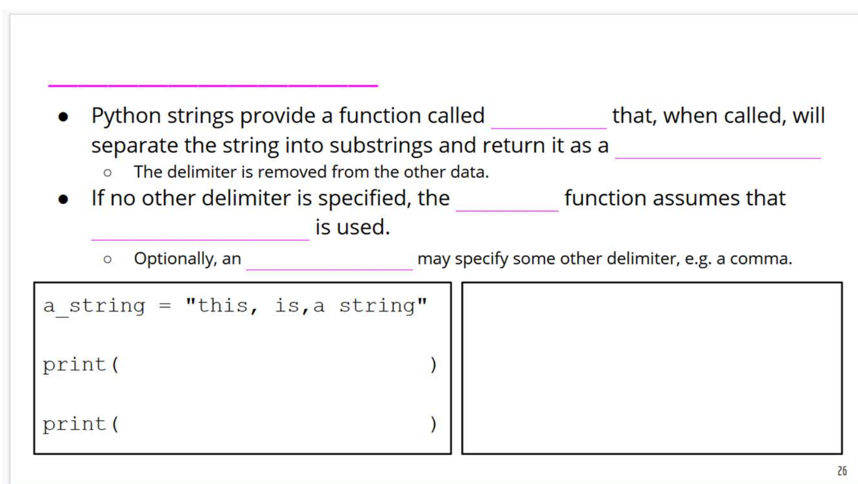

Figure 2: Example content slide with blanks to be filled in by students

While active learning was the primary goal of the for the interspersed activities, a secondary goal was to help students remain engaged throughout the lecture. An additional step was taken as well: each content slide included a series of blanks that were filled in by the instructor as the material was presented. Students were provided with a hard copy of the presentation materials at the start of each lecture, and each student followed along with the instructor, filling each blank with new terms or salient phrases. While a digital copy of each presentation (including blanks) was made available as a study aid, students were responsible for filling in the blanks themselves; at no point was a completed version of any presentation made available. This was purposefully done to discourage students from skipping or disengaging during lecture and waiting for a completed presentation to be published after the lecture.

On the second day each week, students in the experimental sections met with their instructor for a two-hour lab. In the first hour, students were presented with a series of short coding challenges, each of which was designed to be completed in under 5 minutes. In the lab, every student was provided with a computer, and the expectation was that students complete the coding challenges individually by writing and running real code. For many students, this was a first hands-on opportunity to practice the concepts and syntax learned in that week's lecture. In the second hour, students are divided into teams of 3 or 4 to participate in the same problem solving activities as students in the control group. No time is devoted to students beginning their individual lab assignment; they are expected to begin on their own time after class.

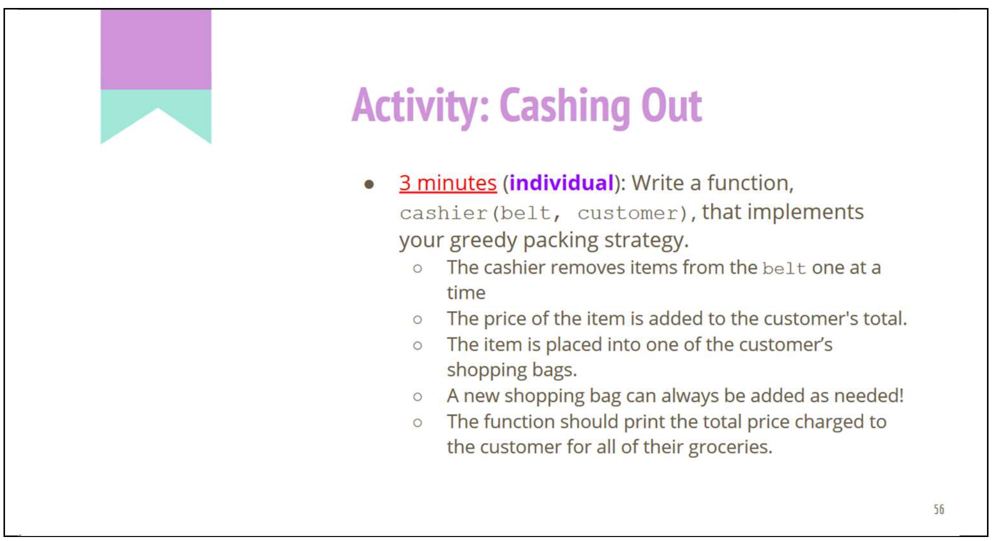

Figure 3:Example individual coding challenge given to students to complete during lab

Research giving students more opportunities to practice improves student learning (provided that the practice activities are related to tasks that the students are expected to perform)[5]. Between the pencil-and-paper activities during the lecture and the coding challenges during lab, students participating in the experimental format of the course are given 20-30 opportunities each week to practice what they have learned. This is in addition to the homework assignment and lab assignment that all students must complete each week.

Collection Techniques

Students were given two surveys, one at mid-semester and one at the end of the semester, designed to gauge their satisfaction with the course. Each of the two surveys included approximately 30 questions. Most of the questions asked students to choose one of five available options or to rank order a series of options.
For example:

> On the average, I am able to maintain my focus on the lecture this much of the time:
> 1) hardly at all
> 2) one quarter of the time
> 3) half the time
> 4) three quarters of the time
> 5) almost the whole time

CSI assesses students using both individual assignments (10 homework assignments, 12 labs, and 1 project) and exams (2 midterms and a final). While each instructor is responsible for making their own midterm exams, all students in the course, regardless of in which section they are enrolled, are given the same homework and lab assignments, and the same cumulative final. The grades from each of these artifacts was collected along with the number of students that withdrew from a course and students' final cumulative grade.

Data Analysis

Data was collected from all 13 sections of the course. The student populations for each section are listed in Table 1. Sections 01, 02, 04, and 06 were the experimental sections.

Section	**01**	**02**	**03**	**04**	**05**	**06**	**08**	**09**	**10**	**11**	**12**	**13**	**14**
#Students	50	50	50	48	47	49	31	42	39	50	44	26	48

Table 1: Student Population Distribution

The results from the two surveys are provided in Table 1 and Table 2. Each row includes a different question with the columns indicating the experimental score (and sample size), the traditional course score (and sample size), and the difference between the two.

The experimental sections were rated higher than the traditional sections in all areas except for how challenges the course was, in which all sections were rated similar. The three survey questions that had the largest deviation will be discussed individually.

Lecture Ranking, in which the experiment showed a 30% positive difference, is a ranked order of how valuable students found lecture in comparison to four other areas in the course. The other ranked items included Lab Session, Lab Implementation, Homework, and Recitation. For reference, while not shown in the above data, lecture was ranked first over twice as often as the lab session which had the second most number of first rank (73 Lecture vs 31 Lab Session). In contrast, the traditional sections ranked the lab session first more than twice as often as lecture was ranked first (131 Lab Session vs 55 Lecture). Since the same lab was assigned to all students the vast increase in perceived value of the lecture by the experimental sections points towards a strong preference towards active learning.

	Control	Variable	Diff	Diff %
Lecture Ranking (1 = Ranked Highest, 5 = Ranked Lowest)	3.065 SS 309	1.732 SS 179	1.333	26.66%
Lecture Materials Adequately Covered Topics (1 = Strongly Agree, 5 = Strongly Disagree)	2.152 SS 309	1.726 SS 179	0.426	8.52%
Ability to Maintain Focus (1 = Almost Entire Time, 5 = Hardly at All)	2.592 SS 309	1.631 SS 179	0.961	19.22%
Lecture Materials Have Been (1 = Extremely Useful, 4 = Neither Useful nor Useless)	2.246 SS 309	1.737 SS 179	0.509	12.71%
Overall Course Satisfaction (1 = Extremely Satisfied, 4 = Neither Satisfied nor Dissatisfied)	2.143 SS 308	1.855 SS 179	0.288	7.20%
How Interesting is the Course (1 = Extremely Interesting, 5 = Not Interesting at All)	2.460 SS 309	2.084 SS 178	0.375	7.51%
How Challenging is the Course (1 = Extremely Challenging, 5 = Not Challenging at All)	2.566 SS 309	2.553 SS 179	0.013	0.27%
How Much Are You Learning (1 = A Great Deal, 5 = Nothing at All)	2.474 SS 304	2.113 SS 177	0.361	7.21%
How Much Are You Enjoying the Class (1 = A Great Deal, 5 = Nothing at All)	2.588 SS 306	1.732 SS 176	0.856	17.13%

Table 2: Midterm Survey

	Control	Variable	Diff	Diff %
Lecture Ranking (1 = Ranked Highest, 5 = Ranked Lowest)	3.198 SS 263	1.691 SS 123	1.507	30.13%
Lecture Materials Adequately Covered Topics (1 = Strongly Agree, 5 = Strongly Disagree)	2.120 SS 274	1.625 SS 144	0.495	9.91%
Ability to Maintain Focus (1 = Almost Entire Time, 5 = Not at All)	2.821 SS 273	1.681 SS 144	1.140	22.80%
Usefulness of Lecture (1 = Extremely Useful, 5 = Not at All Useful)	2.894 SS 273	1.650 SS 143	1.243	24.87%
Usefulness of Lecture Notes (1 = Extremely Useful, 5 = Not at All Useful)	2.777 SS 273	2.105 SS 143	0.672	13.43%
Overall Course Satisfaction (1 = Extremely Satisfied, 4 = Neither Satisfied nor Dissatisfied)	2.106 SS 273	1.699 SS 143	0.407	10.17%
How Interesting is the Course (1 = Extremely Interesting, 5 = Not Interesting at All)	2.352 SS 273	1.881 SS 143	0.471	9.41%
How Challenging is the Course (1 = Extremely Challenging, 5 = Not Challenging at All)	2.304 SS 273	2.427 SS 143	-0.123	-2.45%
How Much Are You Learning (1 = A Great Deal, 5 = Nothing at All)	2.165 SS 273	1.810 SS 142	0.355	7.10%
How Much Are You Enjoying the Course (1 = A Great Deal, 5 = Not at All)	2.571 SS 273	2.106 SS 142	0.466	9.32%

Table 3: End Term Survey

The ability to maintain focus for the duration of a two-hour lecture was a major impetuous for choosing the active learning approach. The results of the Ability to Maintain Focus survey question indicate active learning did substantially increase a student's ability to remain engaged for more of the course. The possible responses were: Almost the Entire Time, Three Quarters, Half, One Quarter, Not at All. Based on the survey results the experimental section saw an increase of approximately 25 minutes over the traditional sections (~95 minutes vs ~60 minutes).

The Usefulness of Lecture survey question, in which the experimental sections scored substantially better, is closely aligned to the Lecture Ranking. As with the Lecture Ranking question, the large difference in perceived value/usefulness of the lecture sessions indicates student preference towards an active learning environment.

Since the homework and lab assignments are uniform across all sections individual assignments were not analyzed. The results from the each of the graded components are detailed in Table 3. Each column is a graded item and each row indicates the experimental or traditional score.

	Control	**Variable**	**Diff**
Homework	80.3%	82.5%	2.18%
Labs	85.9%	87.5%	1.59%
Projects	69.67%	73.24%	3.57%
Final Exam	83.4%	86.4%	3.01%
Final Grade	81.6%	85.2%	3.63%

Table 4: Course Grades

The midterm exams are created by each section's instructor, and while they all cover the same content, the format, specific questions, and grading schemes are determined by the individual instructor. For this reason, midterm exam grades were not included in the analysis. The final exam is a standardized exam that is given to all students in the course. Due to its uniform nature it is used as one of the primary indicators when determining the experiment's success. Since retention is the ultimate goal, the courses WDF rate is also considered a predominate indicator of the experiments success.

	Control	**Variable**	**Diff**	**Std. Dev.**	**p-score**
Withdrawal	9.9%	7.1%	-2.8%	4.46%	0.3057
Grade of D	7.5%	3.1%	-4.4%	4.53%	0.0085
Grade of F	6.8%	4.5%	-2.3%	3.82%	0.3576
Total	24.2%	14.8%	-9.4%	8.26%	0.0065

Table 5: Course WDF Rates

The p-scores were calculated using a two-tailed Welch's t-test. The high variance in D's, and ultimately total WDF rates, aligns with their increased p-scores and indicates that the success of the experimental groups in highly unlikely to have occurred due to random sampling.

Conclusions

The primary goal of the experiment was to determine whether or not an active learning approach to lecture would provide a positive increase in student outcomes with the ultimate goal of increasing successful course completion rates and retaining students that might otherwise change programs or leave the university. Leading indicators of whether or not a student is likely to change programs include grades earned on assignments and exams, and student satisfaction with the course.

The results of the experiment indicate that the active learning techniques applied to the course had a positive impact on grades earned by students. The students in the experimental sections scored 3% better on the final exam, 2.6% better on assignments, and 3.6% better overall. Since the final exam was uniformly administered to all students, it serves as a strong indicator of the impact of the experiment on student success with respect to the graded elements of the course.

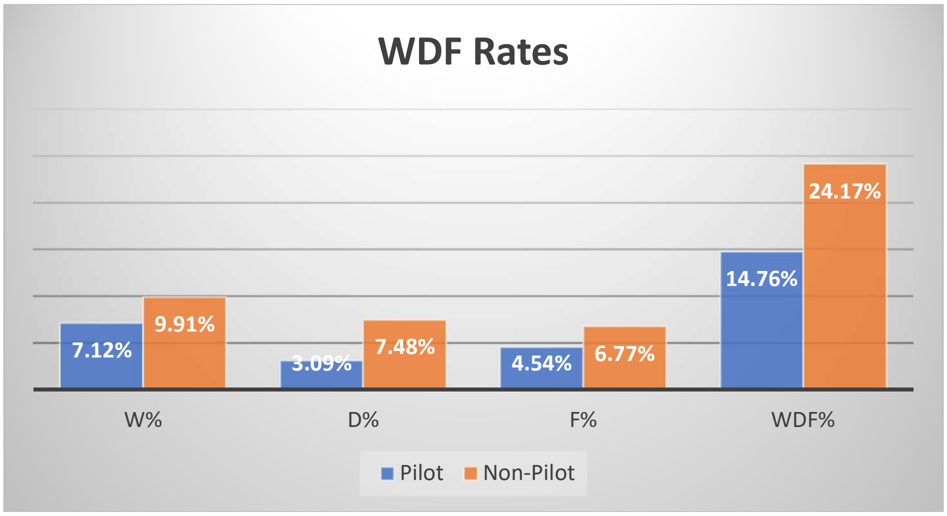

Figure 4: WDF Rates as Percentage of Sample

Perhaps most significant, however, was that the WDF rate for students in the experimental sections was 14.8%, 9.4% less than those of students enrolled in the traditionally conducted sections (Figure 4).

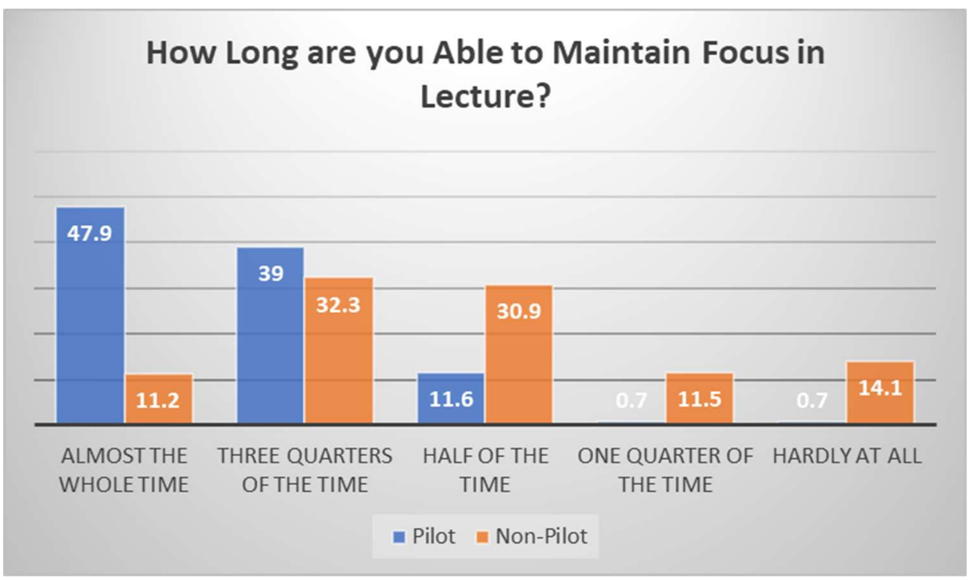

Figure 5: Results - Ability to maintain focus during lecture

In addition to the primary goal of increasing the student's immediate success rate there was a secondary goal of increasing students perceived satisfaction with the course. Previous studies have indicated that increased satisfaction in a course leads to increased retention rates [6]. In the area of perceived satisfaction and self-measured engagement rates the experimental sections were rated significantly higher. Student engagement for a two-hour lecture increased from 43.5% of students claiming that they could pay attention for most or all of the lecture in the traditional sections to 87% of students in the experimental sections. Perceived satisfaction, as measured by targeted survey questions was also markedly higher for the experimental sections.

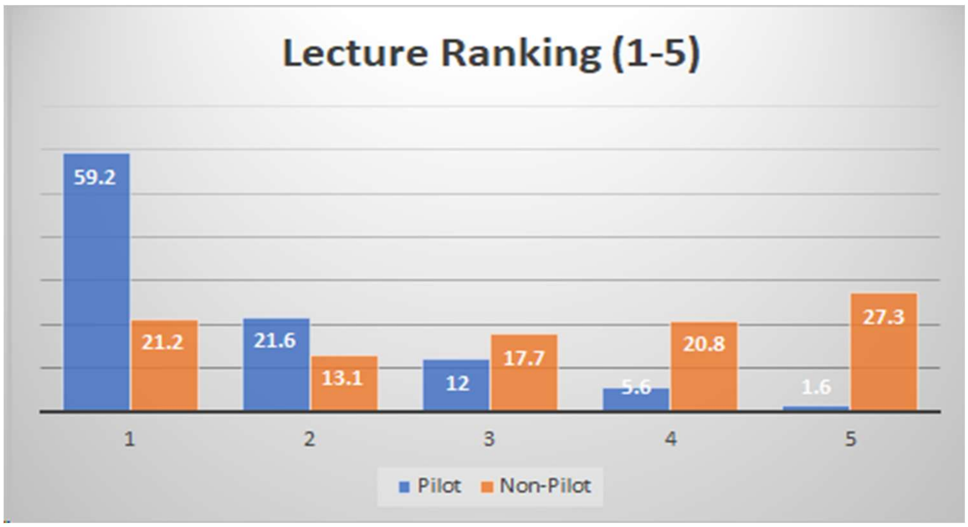

Figure 6: Results – Lecture Ranking 1 (most important) to 5 (least important)

In conclusion, both the course success rates and perceived student satisfaction were increased by a significant degree in the experimental sections in comparison to the control sections.

Future Work and Concerns

The experiment's positive results have led to continuing the experiment for the follow-on course in the introductory computing sequence, Computer Science II. Students enrolling in the experimental sections will contain a mix of students:

1. Continuing in the experimental sections (in both the CS1 pilot and the CS2 pilot)
2. Participating in the active learning lectures for first time (not in the CS1 pilot)
3. Students that previously participated in the active learning lectures but are attending the traditional lectures for the follow-on course (in the CS1 pilot but not in the CS2 pilot)

Looking even further ahead, there is an opportunity to perform a longitudinal study on students involved in the experimental sections. The goal of such a study would be to ascertain any variations in retention and graduation rates.

In addition to continuing and future experiments the completed active learning pilot has raised some concerns with attempting to mainstream it to all sections. One major concern is how to scale the creation and distribution of the supplied packets. Currently the instructors involved in the experimental sections have limited the paper packet size to 60 slides which are printed on 15 two-sided 2-up pages. Even with this limitation the four sections involved in the experiment used

~2700 pages of paper each week. This equates to ~35,000 pages printed over the semester. Increasing this to cover all 14 sections seems untenable as it would require printing ~9500 pages each week, using ~122,500 pages a semester. Between the material and time costs, not to mention environmental impact, the current approach to class handouts is not recommended if the pilot is adopted for all sections.

Alternatives being considered are reduced packet sizes using condensed material, using electronic versions, or abandoning the fill-in-the-blank portion of the active learning. All options come with possible decreases in student outcomes which will need to be evaluated before a final decision can be made.

The other major concern is in preparation time. While not tracked directly, the pilot instructors expended significant amounts of time, measured in hundreds of hours, over the course of the semester preparing the materials. Currently each instructor for the course is responsible for preparing their own material for that week's lecture. It may be difficult to achieve full participation from all faculty with this amount of increased workload.

One approach that is being considered is to create starter presentations which include a sample lecture/activity set which could be distributed to faculty to aid in their preparation. This could limit the impact to individual instructors when switching to the new methods. This still leaves a few faculty members burdened with a greatly increased workload to maintain and update the sample presentation materials. The hope is that over time this additional workload would as the process becomes more streamlined, but it always expected to exist at some increased level.

Bibliography

[1] Scott Freeman, Sarah L. Eddy, Miles McDonough, Michelle K. Smith, Nnadozie Okoroafor, Hannah Jordt, and Mary Pat Wenderoth, *Active learning increases student performance in science, engineering, and mathematics*. PNAS vol. 111, no. 23, 2014

[2] Prince, Michael, *Does Active Learning Work? A Review of the Research.* Journal of Engineering Education, Jul 2004

[3] Xihui Zhang, Chi Zhang, Thomas F. Stafford, Ping Zhang, *Teaching Introductory Programming to IS Students: The Impact of Teaching Approaches on Learning Performance.* Journal of Information Systems Education, Vol. 24(2) Summer 2013

[4] Isabel C. Moura, Natascha van Hattum-Janssen, *Teaching a CS introductory course: An active approach.* Computers & Education 56 (2011)

[5] Sarah L. Eddy, Mercedes Converse, Mary Pat Wenderoth *PORTAAL: A Classroom Observation Tool Assessing Evidence-Based Teaching Practices for Active Learning in Large Science, Technology, Engineering, and Mathematics Classes.* CBE – Life Sciences Education Vol. 14, 1-16, Summer 2015

[6] Andala Rama Putra Barusman, *THE ROLE OF STUDENT SATISFACTION AT PRIVATE UNIVERSITIES*. Buisness and Entrepreneurial Review vol. 14, no. 1, 2014