Is It Working? evaluating a particular approach to CS1

Kate Sanders, Ann Moskol, Erik Holder
Rhode Island College

Brenda Jochums-Slez
UMass/Lowell (emerita)

February 6, 2005

In order to understand the strengths and weaknesses of our approach to CS1, we are just beginning a four-semester project. During the next four semesters, we will be gathering extensive data in one section of CS1 per semester, including background information, attendance, number of times each student speaks in class, attendance at office hours, assignments given, and selected student work. At the end of the semester, we will compile a table summarizing the skills the students have demonstrated, similar to the following:

Concepts	Percent of Students Who		
	Define	Recognize	Use
Class			
Instance			
Instance variable			
Constructor			
Method			
Formal parameter			
...			

Each section has approximately 24 students. The students will (generally) be different each semester. There will be two instructors who alternate teaching the class, one in the spring and one in the fall. The topic order will be the same each semester, but details of assignments and time spent on particular topics may vary.

Research questions:

Based on these data, there are a lot of questions we could ask. Here are a few of the possibilities.

Given a particular approach to teaching introductory Java,

1. Which topics do the students demonstrate the greatest understanding of? the least?

2. Do students show greater understanding of object-oriented concepts (e.g., classes vs. instances, polymorphism) or of traditional CS1 topics (e.g., parameters and looping through arrays)?

3. Is there a correlation between success and gender?

4. Is there a correlation between success and student confidence level with computers at the start of the semester?

5. Is there a correlation between success and student major?

6. Is there a positive correlation between success and class attendance?

7. Is there a positive correlation between success and other indicators of engagement in the subject such as speaking out in class, attending office hours, etc.?

8. Is there a positive correlation between the amount of class time spent on and number of opportunities to demonstrate a particular concept and the overall percent of students who succeed in mastering it?

9. (A related question) Is there a correlation between the point in the semester when a concept is introduced and the overall percent of students who succeed in mastering it?

What evidence would convince me?

That depends on the question. Most of these questions compare the achievement levels of the students with fairly concrete information (how often they come to class, gender, major, etc.), so as long as the methods for assessing student achievement are reasonable (and reproducible), we'll be satisfied.

One question that's a little more problematic is the question of topic order. We can confirm whether (as expected) topics we present early are mastered more thoroughly than those we present later. What we can't say, without a control group using a different topic order, is whether topics presented early are *generally* mastered more thoroughly.

Questions for the reader

Is there anyone who would like to use the same tools to analyze a different CS1 course, in order to obtain more comparative data? Are there other data we should collect? Other questions we should ask? (Undoubtedly, there are many more things we should read.)

Acknowledgments

Thanks to the students in CS201 for participating in and to the NSF for funding this project.

References

[1] A. Robins, J. Rountree, and N. Rountree. Learning and teaching programming: a review and discussion. *Computer Science Education*, 13(2):137–172, 2003.