

Assessment of Programming Skills of First Year CS Students: Do they know how to program?

A Working Group at ITiCSE 2001

Working group goals

As we stated in the proposal, the goals of the working group are:

1. Develop a shared understanding of the expected programming skills of first year students.
2. Determine if students are operating at the level of expected skills.
3. Generate a set of assessments that can be used in computer science curricula to collect more extensive data and to generate proposed interventions to improve programming skills.

We haven't deluded ourselves into thinking we can completely answer the questions proposed in the goals. What we hope to accomplish is the beginning of a dialogue to develop a more comprehensive set of assessments that can be used by computer science educators. To that end, we have proposed a multi-step process that begins with participants collecting a set of data from their own classes and sharing that data with other working group participants at the conference. We hope to produce a report that offers others a set of assessments that they can begin to use in their classes, and an analyzed set of data that benchmarks what we know about student's skills.

A historical note and the basic premise for the working group's activities

When I wrote this proposal my own parochial interests and biases were included. My research has convinced me that a large part of programming expertise starts with basic skill development. There are definitely knowledge elements, but at the introductory level, skill development is critical. My observations and informal experiments on our first year students convinced me that the level of skill we hoped they would attain doesn't exist. I should note that we teach our first semester class to every student at Georgia Tech, and our second semester class to primarily computer science and computer engineering students. My evaluations have been done on the second semester group. My data is incomplete, my observations are primarily anecdotal, and my concerns are to a large degree intuitive. But, with that set of caveats my conclusion is they can't program. That is what prompted me to write the proposal for the working group. I spoke with friends in the computer science education community, and they had similar observations and experiences, so I could at least feel somewhat comfortable that this wasn't a Georgia Tech unique issue.

With that entrée I proposed the working group to the ITiCSE program committee and it was accepted. I wanted the working group to begin to get a grasp on how extensive this problem is, what might characterize it, and what if anything we might do to attempt to mitigate it. The point of the working group's activities is to get enough data so that we can say, "we've got a problem here", or "Mike's crazy and our kids can program even if he thinks his can't". Either way, I think your participation will be valuable to you and looks to me like fun.

Activities prior to the conference

We are starting this activity with one simple programming problem for you to administer to your class. We think it gives us some insight into programming skills of second semester students, but you decide where to administer it. I should note I made that last comment with some trepidation. I want to re-emphasize that we developed the problem with our students in mind. But, a set of data you gave to juniors won't be of much use. That's why we developed three versions. The three versions of the

problem exist so that you may select the "level" of the problem appropriate to where you think your students ought to be in the class you administer it in. We obviously can make no claim that we have found the quintessential problem. You may argue that it doesn't assess, x or y or z. I accept those criticisms, and I hope that we will use those criticisms at the conference to develop a better set of problems. The problem needs to be given under the following constraints mainly for consistency across the data set.

1. You need to characterize the students who you collect data from. You or others may use some of the data I am asking you to collect for other research if you wish to share it. Demographics should include:
 - Number of students in sample and total number of students in the potential sample. (For example, you gave the problem to one of the three sections of your second semester programming class).
 - Position in the program (year and semester)
 - Major (This would be characterized as. Sample is 300, 100 are CS majors, 200 are history majors).
 - Male/Female ratio
 - Student characteristics that may skew the data. I have to leave this to you to decide what to do, but some examples of what might skew the data, or at least we should be aware of could include: Your population is an adult population. Your population is predominantly part-time students. Your population is non-CS majors.
 - School. When you submit your data to me, I need to identify where it came from. Once I put it in the database, your identification (you and school) will be removed.
 - If there are other data you wish to provide, please do, but make sure you don't violate any privacy issues.
2. You need to tell me which problem you had the students do. It would be of great interest if you have multiple sections and you decide to give a version to different sections. Make sure if you do that, you ensure the population is relatively uniformly distributed. Also, you will need to give me demographics on each group separately.
3. Make sure you get the appropriate approvals for human subject research. At Georgia Tech I applied for essentially a non-review review. What that means is I am collecting data from "normal" classroom activities and no subjects will be revealed. I needed to inform our administration of this, but that was all. If you furnish data and haven't run it through the proper channels, it could spoil the fun for everyone. The application for my approval took me about 15 minutes and they got back to me with my approval in about a week.
4. You need to score the problems with the scoring criteria we give you. Again that helps insure consistency.
5. You need to get me your data no later than May 15 for inclusion in the data I will bring to the conference. If you don't get your part done before that date, we can try to include it at the conference, but that will mean we will spend time trying to fit your data in, and have less time to talk about it.
6. Based upon your findings, generate questions and ideas for discussion at the conference. I'll try to summarize those if you send them to me prior to May 15.
7. Please, to make this of any use we need to be careful. Selecting the easiest problem to make your students look better won't help. Having TA's help the students won't help. Not confining the work to the allotted time won't help. Not ensuring that the work is individual won't help. All of these things and many others will just invalidate the data.

That covers what I want you to do, now what am I going to do prior to the conference?

1. When I receive your data as well as my data, I will enter it into a database/spreadsheet that will allow me to perform some rudimentary analysis on it. I'll do some simple comparative looks and summary statistics. Primarily I plan on bringing the raw data with me for us to look at and discuss.
2. If I get the data in time, I'll send you the summaries prior to the conference.
3. I'll also try to summarize any criticisms/complaints for use at the conference.

What we will do at the conference

Meet to share analysis results and develop benchmarks from the data.

Critique the first instrument and develop more refined instruments if necessary.

Develop a report outlining the results, the concerns, proposed instruments and interventions and the plans for future work.