IMPACT OF AUTO-GRADING ON AN INTRODUCTORY COMPUTING COURSE [*]

Mark Sherman, Sarita Bassil, Derrell Lipman, Nat Tuck, Fred Martin
Department of Computer Science
University of Massachusetts Lowell
Lowell, MA 01852
978-934-{1964, 3911, 1964, 1964, 1964}
{msherman, sbassil, dlipman, ntuck, fredm}@cs.uml.edu

ABSTRACT

This project presents and assesses the impact of a pedagogical tool, called Bottlenose, deployed in an introductory Computer Science course. Bottlenose is a web-based framework that accepts student submissions and presents immediate feedback. Students may submit any number of times before the assignment due date. We expect them to use the feedback from previous submissions to improve their subsequent submissions. We compared student behavior on assignments against previous semesters, which used the same assignments, but with no automated feedback system. We observed that students, when using the feedback system, make more submissions per assignment, indicating that students were leveraging feedback to improve their programs.

INTRODUCTION

We introduced a web-based framework, called Bottlenose, to provide students with immediate, automated feedback in an introductory computing course, and assessed the impact of that system on student learning.

During the course of the semester the students submitted approximately 50 C programming exercises. The course typically has 40 to 60 individual programming exercises, and 40 to 60 students in a course section. Unlike other web-based feedback systems like BOSS [5], CourseMaker [3], and Web-CAT [2], Bottlenose is extremely

---

lightweight and encodes very few pedagogical assumptions. Tests for a particular assignment can be specified in any scripting language, and the testing programs can invoke any tools available on the system on which Bottlenose is installed.

In the existing course, students must complete three or four small assignments per week. The high frequency with which students are required to submit assignments means that students may submit several assignments before receiving any feedback. Instructors and graders may spend time correcting and penalizing the same error on several assignments, when earlier feedback might have allowed a student to understand and correct the error on all or most of their assignments. This prevents the graders from giving feedback on more subtle errors, and tends to demoralize students and slow their progress.

From this perspective, Bottlenose was designed to achieve the following features:

- Accept student code.
- Run assignment-specific tests cases on student code.
- Allow any testing methods, frameworks, or tools the instructor desires in tests.
- Display test results to the student nearly immediately.
- Store the results of test cases for instructor (or human grader).
- Store student's code for manual review by instructor (or human grader).
- Provide an efficient interface for a human grader to review tests, read code, and provide written feedback.

This is a common situation for institutions teaching computer science [4]. The above features are aligned with the desirable features described by Ala-Mutka [1]. The system described in this paper represents a secure, lightweight, and easy-to-use solution. Specifically, this system is an online automatic assessment tool, which runs pre-made, custom tests on student submissions, but also allows for easy human assessment of the code and the test results.

Technical Description of System

Bottlenose was initially developed to support the teaching of a "flipped" course, where students watch video lectures online before class to prepare for classroom questions and discussion. It also included online submission and grading of programming assignments, which are useful functions for traditional courses, as is examined in this paper.

The system was built using the Ruby on Rails [6] web application development framework. It runs on a Linux server, providing two main benefits: the system utilizes many Linux security mechanisms, and, the tests for student code can access the variety of Linux-based applications and tools. Student submissions are stored on the server file system, and other data, including grades, are stored in a database.

A simple process for online submission of assignments is provided. Students are emailed unique authentication links that bring them to their list of assignments, and identify the students to the system. Assignments are submitted by uploading the programming code directly in the web browser. Bottlenose supports assignments requiring submission of a single source file, as well as those assignments requiring multiple files, submitted as a compressed archive file. The automated grading process

begins immediately when an assignment is submitted, giving students feedback within a few seconds. Students may attempt submissions multiple times.

In order to automatically grade student programs, submissions are compiled and run on the server. Allowing students to run arbitrary code on the server is clearly a potential security issue [4]. Bottlenose uses a sandbox mechanism to prevent student programs from causing trouble. Five major techniques are used to isolate student programs from the rest of the system:

- **Separate system user** - Each student program is run under a separate, limited system user.
- **Run in a sandbox** - Student programs can only access specific, white-listed parts of the file system.
- **Disposable working directory** - Each program is executed in a temporary file system, which ceases to exist when the grading process is finished.
- **Resource limits** - Limits on a variety of resources, including RAM, child processes, and created file size, are enforced on student programs.
- **Watchdog timer** - A grading process is terminated if it lasts more than a set time limit.

This sandbox mechanism may be vulnerable to a clever student intentionally trying to defeat it. Thus far we have not yet identified any exploit. It does perform adequately at preventing the grading server from being disrupted by common student mistakes, like infinite memory allocation loops, without the need for exotic isolation systems, virtual machines, or additional servers.

Writing Tests

In preparation for student submissions of assignments to the system, the instructor writes a set of test scripts, which automatically evaluate specific aspects of students' programs. A test may evaluate, for example, that given specific input, the program generates proper output; it may check whether certain required functions are provided by the student's program; it can determine whether the program exited successfully or crashed; it can execute unit tests; and more, as discussed below.

Tests can be written in any language. The only requirement for tests is that they conform to the Test Anything Protocol (TAP) [7]. With TAP, a test outputs an indication of the number of tests that will be run, e.g., 1..4 for four tests, followed by a line of output for each test, which indicates whether the test, e.g., test #2, succeeded (ok 2) or failed (not ok 2). The tests are otherwise free to provide input in any form to the program under test and to retrieve output via their standard output mechanism or via a file. Tests are also free to provide additional output, such as a description of the test to be run, or student's and expected output, when a program does not yield the correct results.

The feedback that students receive is entirely written by the instructor, and is generally displayed as a function of the success or failure of a particular test. To present feedback to the student, the instructor simply prints to the screen in the test script. The instructor may write feedback text based on any programmatic conditions.

Given the flexibility of this testing harness, it is possible to run some interesting tests on students' programs to let them know whether they are correctly accomplishing their program's goal or not. The most common test we have written parses the output of the student's program. Parsing could be a direct comparison, a regular expression, or a custom program to process complex output. We have modified the student's source code and executables, replacing system calls with reference implementations that provide feedback of their use. We have isolated and unit-tested functions in the student code. We have run some student programs inside memory-checker programs, and parsed the output of the checker and returned it to the student. The test framework flexibility of Bottlenose makes possible a plethora of other testing scenarios.

The simplicity of this platform allowed for any test to be written that can be realized through another program. This makes the testing mechanism more flexible and robust than those used in other C-language studies, such as Sterbini and Temperini [8], which largely depend on comparing student code and results against a reference implementation. With Bottlenose, the tests are written by the instructor to custom-fit the assignment, which can require a significant investment in time and effort by the instructor, but provides unsurpassed flexibility. Once written, tests can be re-used for subsequent course offerings, and often, new tests can be largely based on those written previously.

Now, we will discuss research on student behavior on assignments that we conducted after Bottlenose was introduced.

METHODOLOGY

Analysis was performed on historic course data, with IRB-approved protocols to de-identify data and protect participants. The data included grades the students received on their submissions to assignments, as well as the submissions themselves, which were C source code files. Previous course offerings utilized a rudimentary, no-feedback electronic submission system, which provided the basis against which the Bottlenose offerings were compared.

Two forms of analysis were conducted: (1) an aggregate analysis, tallying all submissions per assignment, per course offering, resulting in numbers of student submissions per assignment for each course section, and (2) a fine-grain submission rate analysis, where the distribution of submission rates among students was visible.

The analysis used data from seven offerings of the same introductory computing course, spanning four semesters and three instructors. The data included participating students, defined as those who submitted for at least one programming assignment. Assignments that were not programming assignments were also removed from the data, such as in-class paper exercises, quizzes, and tests. All sections of the course used a common, core set of assignments. Assignments that deviated greatly from the common set, such as those with novel specifications or dependent on new concepts, were removed from the data.

RESULTS

Aggregate analysis showed a substantial increase in the number of submissions students made when using Bottlenose, compared to the no-feedback electronic submission system. Shown in Table 1, course sections M1-M4 were electronically submitted but provided no immediate results, and sections A1-A3 used Bottlenose (M = manual/non-automatic feedback; A = automatic feedback). The numbers of total submissions made, participating students, and assignments were divided to create a "Submissions per Student per Assignment" descriptor, which indicated the general rates of submissions per assignment in each section. A value of 1.0 would indicate that, on average, every student made one submission for every assignment. For any given assignment, there was a subset of students who did not submit an attempt for it, which lowered the average submission rates.

Course Section	# Submissions	# Students	Submits/Student/Assignment	Average	Difference
M1	2711	49	0.9706	0.9831	1.829x
M2	2672	47	0.9974		
M3	2425	50	0.9898		
M4	386	44	0.9747		
A1	3578	45	2.0387	1.7982	
A2	3119	47	1.7016		
A3	2561	36	1.6544		

Table 1: Submission rates among course sections

We also used a finer analysis technique, where the average submission rate per assignment was computed for each student. With this, we could see the distribution of submission rates within the course sections. In Figure 1, we can see the changes in that distribution between the sections that used Bottlenose (automatic) and those that did not (manual). The sections that used Bottlenose show a slightly wider, right-tailed distribution, with the peak at a higher submission rate (1.77) than the manually graded sections (1.26).

DISCUSSION

The standard deviation of the Submission/Student/Assignment rates of the manually graded sections, as seen in Table 1, was 0.013, indicating that, despite spanning multiple semesters and multiple instructors, submission rates before the introduction of Bottlenose were consistent. The course sections that used Bottlenose had higher submission rates than those that did not. This increase was significant ($p<.001$), showing that students, on average, made nearly twice as many attempts on an assignment when using the feedback-providing automatic assessment system.

As shown in Figure 1, in the course sections that did not provide automatic feedback, a large portion of the students submitted, on average, just over once per assignment for every assignment. In those sections, almost none of the students had averages above two submissions per assignment. In the sections that used Bottlenose, a greater number of students submitted more than once, and many more submitted more

than 1.5 times per assignment. Some students submitting more than twice per assignment, or more, with Bottlenose, which occurred rarely in sections without the feedback system. The lack of significantly higher average submission rates (six or more) indicated that students were not using Bottlenose as their only compiler, or abusing the re-submission mechanism, which were common concerns in the literature.

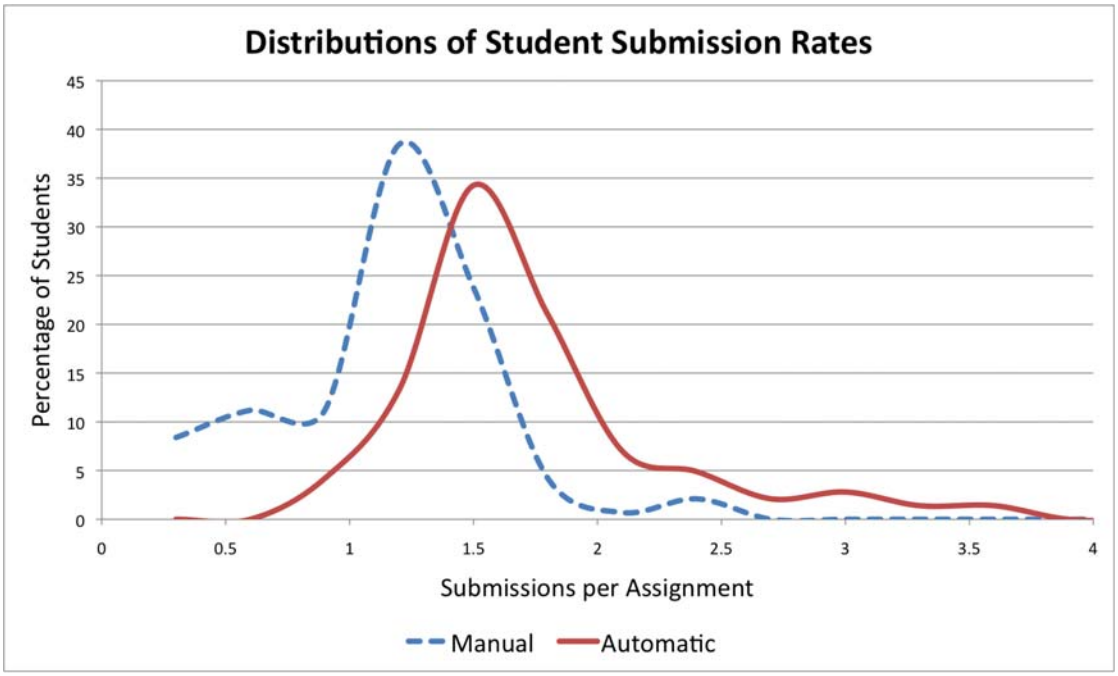

Figure 1: Distributions of how many students had similar submission rates, with and without the automatic Bottlenose feedback and submission system.

CONCLUSIONS AND FUTURE WORK

The results from the use of Bottlenose indicate a healthy change in mindset towards submitting, where students are using the system to help them make one or two additional iterations on their program.

There is a wealth of data from this study remaining to be analyzed, including the quality of student iterations and products. In using this system, many tests were written by the instructors, which include the feedback that the students see. There may be relationships between the types of tests used, the quality and kind of feedback provided to the student, and student performance. This avenue of study is promising. There are also numerous case studies to be extracted from the data, which show individual students using the feedback from Bottlenose to inform their iterations of work. These will be presented in a future publication.

ACKNOWLEDGEMENTS

We offer special thanks to Prof. Anne Mulhern for reviewing the paper and providing us with valuable input. Thanks also go to James DeFilippo for his work on the auto-assessment prototype project. This material is based upon work supported by the National Science Foundation under Grant No. DUE-1225719.

REFERENCES

[1] Ala-Mutka, K., A survey of automated assessment approaches for programming assignments, *Computer Science Education*, 15, (2), 83-102, 2005.

[2] Edwards, S., Perez-Quinones, M., Web-CAT: automatically grading programming assignments, In *Proceedings of the 13th Annual Conference on Innovation and Technology in Computer Science Education (ITiCSE '08)*, 2008.

[3] Higgins, C., Gray, G., Symeonidis, P., Tsintsifas, A., Automated assessment and experiences of teaching programming, *ACM Journal on Educational Resources in Computing*, 5, (3), 5, 2005.

[4] Ihantola, P., Ahoniemi, T., Karavirta, V., Seppälä, O., Review of recent systems for automatic assessment of programming assignments, In *Proceedings of the 10th Koli Calling International Conference on Computing Education Research (Koli Calling '10)*, 86-93, 2010.

[5] Joy, M., Griffiths, N., Boyatt, R., The BOSS online submission and assessment system, *ACM Journal on Educational Resources in Computing*, 5, (3), 2, 2005.

[6] Ruby on Rails, http://rubyonrails.org/ , retrieved November 15, 2012.

[7] Schwern, M., Lester, A., Documentation for the TAP format, 2003, http://search.cpan.org/~petdance/Test-Harness-2.64/lib/Test/Harness/TAP.pod , retrieved November 15, 2012.

[8] Sterbini, A., Temperini, M., Automatic correction of C programming exercises through unit-testing and aspect-programming, In *Proceedings of the 2nd International Conference on Educational Information Systems, Technologies, and Applications (EISTA '04)*, 6, 2004