

1  
2  
3  
4       **Abstract**  
5

6           We present empirical findings from using an interactive electronic textbook (eTextbook)  
7  
8       system named OpenDSA to teach sophomore- and junior-level Computer Science courses. The  
9  
10       web-based eTextbook infrastructure allows us to collect large amounts of data that can provide  
11  
12       detailed information about students' study behavior. In particular we were interested in seeing if  
13  
14       the students will attempt to manipulate the electronic resources so as to receive credit without  
15  
16       deeply going through the materials. We found that a majority of students do not read the text. On  
17  
18       the other hand, we found evidence that students voluntarily complete additional exercises (after  
19  
20       obtaining credit for completion) as a study aid prior to exams. We determined that visualization  
21  
22       use was fairly high (even when credit for their completion was not offered). Skipping to the end  
23  
24       of slideshows was more common when credit for their completion was offered, but also occurred  
25  
26       when it was not. We measured the level of use of mobile devices for learning by CS students.  
27  
28       Almost all students did not associate their mobile devices with studying. The only time they ac-  
29  
30       cessed OpenDSA from a mobile device was for a quick look up, and never for in depth study.  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40

41       **Keywords:** eTextbook, learning behavior, mobile learning, computing education  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

## 1. Introduction

Recent interest in MOOCs (Massive Open Online Courses) and the popularity of practice systems like Khan Academy and Code Academy is just the latest in a trend toward increased use of interactive online course materials at all levels of education. MOOCs especially drive a need for scalable methods of assessment that do not require direct involvement of limited teaching resources. While some MOOCs attempt to address the need for assessment through crowd sourcing (especially to students in the class), automated assessment of exercises provides an excellent way to scale up assessment when it is possible. For many traditional courses, a longstanding problem is lack of sufficient practice exercises with feedback to the student. Again, automated assessment provides a way to increase the number of exercises on which students can receive feedback.

Online tutorials and interactive exercise systems provide the opportunity to automatically log massive amounts of user interaction data at fine detail, to the level of individual mouse events. This wealth of information might be used in a number of ways to improve the pedagogical value of online materials. Developers can hope to learn from interactive log data simple things like what platforms or devices are most often used by students, which exercises are taking an unreasonable amount of time, or which ones appear too easy because students never get them wrong. But careful analysis of log data can hope to deduce more complex behavior. For example, by examining the times of various interactions, we should be able to determine whether students are reading the materials before attempting the associated exercises. We can hope to tell whether students are viewing all of the parts of a given visualization, or skipping through it. We should be able to tell whether students are going back to the materials to use them to study for a test by the fact that they look at them even after receiving credit for completing them.

1  
2  
3  
4 The OpenDSA project (Fouh et al., 2014; Shaffer, Karavirta, & Naps, 2011; Shaffer, Naps,  
5 & Fouh, 2011) provides a collection of online, open-source tutorials for Data Structures and Al-  
6 gorithms (DSA) courses. They combine textbook-quality text with algorithm visualizations (AV)  
7 and randomly generated instances of interactive examples and exercises to provide students with  
8 unlimited practice. OpenDSA collects log data for all user interactions occurring on an OpenD-  
9 SA web page. Between Spring 2013 and Spring 2014, we collected and analyzed usage data  
10 from eight courses totaling about 700 students. We present an analysis that demonstrates how log  
11 data might be used to understand how students use online tutorial systems.  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22

## 26 **2. Material and methods**

27  
28  
29  
30

31 The basic functional unit for OpenDSA materials is a module, which represents a single top-  
32 ic or part of a typical lecture, such as a single sorting algorithm. Each module is a complete unit  
33 of instruction and typically contains AVs, interactive assessment activities with automated feed-  
34 back, and textbook quality text. Modules can be grouped together into chapters, such as might be  
35 found in traditional paper books. OpenDSA content is built using HTML5 and JavaScript, mak-  
36 ing it device and browser independent. AVs are built using the JavaScript Algorithm Visualiza-  
37 tion (JSAV) library (Karavirta & Shaffer, 2013). Many OpenDSA exercises are “algorithm simu-  
38 lations”. These require that the student manipulate a data structure to show the changes that an  
39 algorithm would make on it, such as clicking to swap elements in an array or clicking on appro-  
40 priate nodes in a tree or graph. The AVs and algorithm simulation exercises are specifically de-  
41 signed to be manipulated either through mouse and pointer interactions or touch interactions  
42 when using touchscreen devices. We generally refer to algorithm simulation exercises, as “profi-  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

1  
2  
3  
4 proficiency exercises". This type of exercise were inspired and built in collaboration with the team  
5 that created the TRAKLA2 system (Malmi et al., 2004). We make use of the Khan Academy  
6 framework (<http://github.com/Khan/khan-exercises>) to provide support for multiple choice, T/F,  
7 and custom interactive exercises that we call ``mini-proficiency" exercises. We will refer to these  
8 collectively as "KA exercises". All exercises are automatically graded and provide feedback to  
9 the user. Students can repeat exercises as many times as they want until they get credit, or even  
10 work them again after receiving credit as a study aid.  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20

21       OpenDSA has been used to teach the second semester fundamental data structures and algo-  
22 rithms course (CS2), and also a more advanced data structures, algorithms, and analysis course  
23 (CS3) at a total of five higher education institutions in the US, Egypt, and Finland. Different in-  
24 structors have used OpenDSA in different ways, but typically they used OpenDSA exercises for  
25 graded homework, and/or used AVs from OpenDSA as a lecture aid. During Spring 2013,  
26 OpenDSA was used in three course sections: two offerings of Virginia Tech's version of CS3  
27 (we will refer to these as C1 and C2), and one offering of a CS3 course at Alexandria University  
28 in Egypt (which we will refer to as C3). OpenDSA was used as the primary resource to teach  
29 topics related to Hashing in C3, followed by a midterm that included questions on hashing.  
30  
31 Course C1 used OpenDSA as the primary source of material for sorting and hashing (about three  
32 weeks of material). For this section, AVs (slideshows) were not given credit. There was a single  
33 assignment to complete the exercises for the two chapters, so all exercises were due at the same  
34 time. In contrast, course C2 used another textbook for the initial presentation of material on sort-  
35 ing. However, OpenDSA exercises were then assigned after the lecture period on sorting was  
36 complete. C2 used OpenDSA as the primary source for Hashing. C2 students received a small  
37 amount of credit for completing slideshows. We deliberately made this distinction in slideshow  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

1  
2  
3  
4 credit between the two sections in order to study how that affected student learning behavior, as  
5  
6 described below.  
7

8  
9 We relied on the logged data to infer students' behavior. The OpenDSA system records  
10  
11 about 200 different types of events. The events can be grouped into the following categories:  
12

- 13  
14 • Registration and login interactions (all actions such as student registration, log-  
15  
16 ging into and out of the system);  
17
- 18  
19 • Static Content interactions (when a student loads a module page, follows a hyper-  
20  
21 link, or navigates to another page using the navigation menu or the table of contents);  
22
- 23  
24 • Interactive Activities interactions (when the student clicks to advance a slide-  
25  
26 show, clicks within an AV, etc.);  
27
- 28  
29 • Assessment Activities interactions (all interactions involving loading an exercise,  
30  
31 submitting an answer, completing a step of a proficiency exercise, etc.); and  
32
- 33  
34 • Gradebook interactions (when students load the gradebook page to check their  
35  
36 score).  
37

38  
39 All events are recorded with a timestamp. To assess the relationship between the use of  
40  
41 OpenDSA and students performance, we used both log data and performance data (stu-  
42  
43 dents' scores in written tests) from the Fall 2013 offering of a CS3 course at VT (which we  
44  
45 refer to as C4). OpenDSA was used as the main course material, and the instructor regularly  
46  
47 used the visualization in the classroom as lecture aide. The students performed regular  
48  
49 mandatory OpenDSA homework and took two midterms and one final examination.  
50  
51  
52 OpenDSA-based homework accounted for 20% of the course final grade.  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

### 3. Student Learning Behavior

We sought to measure the prevalence of several behaviors that we collectively term as “credit-seeking”. We believe this is different from (and possibly in conflict with) “learning” behavior. Examples of “credit-seeking” behavior include jumping straight to exercises without reading the associated text, clicking through slideshows (such as shown in Figure 1) as quickly as possible, skipping to the end of slideshows to get credit, and using the results from AVs to complete exercises.

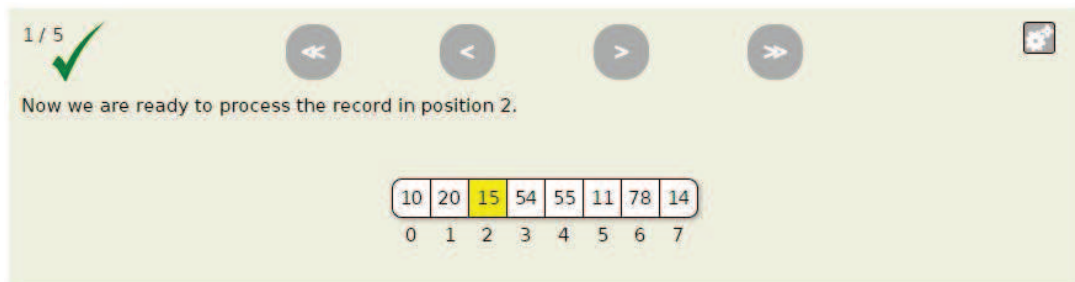


Figure 1: Example of an OpenDSA slideshow. Standard controls allow the user to advance the slideshow by one slide, back up one slide, jump back to the beginning, or jump to the end.

#### 3.1.1 Not Reading

We observed two general patterns of behavior related to reading the text. Ideally, students would read most or all of the text, completing exercises as they encounter them within the module. But students often skip directly to the exercises, only reading as required to get exercise credit. Note that proficiency and KA exercises are placed “behind a button” by default, in that they require that the student clicks on an interface button to reveal the exercise.

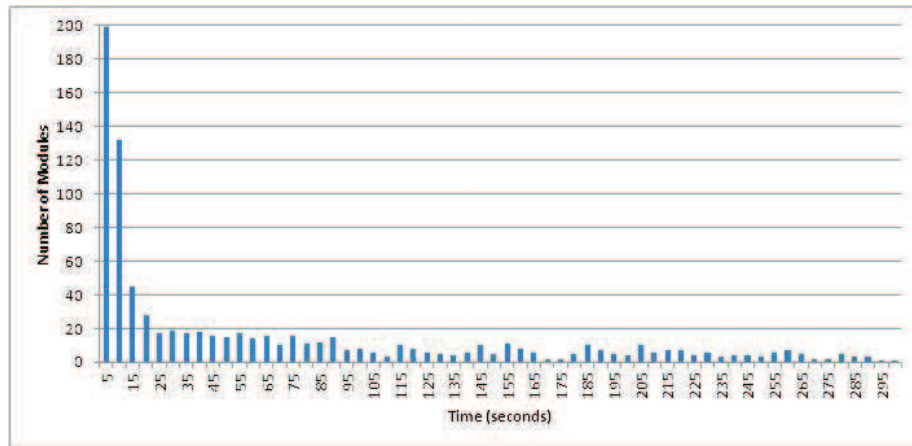


Figure 2: Time between initial module page load and the first exercise event for C1, using 5-second bins.

We measured the total time that a module was open before a button was clicked to reveal an exercise. We assumed that once a student clicked the button to reveal an exercise, they attempted it immediately. Figure 2 represents the distribution of all the modules with exercises attempted by C1 students, grouped by the time between the moment the page is loaded and when the students uncover the exercises. All classes showed a peak at the beginning of the histogram, indicating that a large number of students opened modules and immediately started exercises. We quantified how much the class as a whole read by identifying the number of modules in quartiles ranked by time between page load and first exercise attempt for each class. Table 1 shows that the first quartile ranges from 10 to 20 seconds and the second quartile from 70 to 120 seconds. This means that students begin exercises on 25% of modules within 20 seconds of loading the page and on 50% of modules within two minutes. The third quartile is roughly five to seven minutes. Note that at least some measures from the upper 25% quartile are likely to include instances where the user left the browser page open without being active. So the amount of skipping directly to exercises could be greater than the quartiles indicate. While there is naturally a

1  
2  
3  
4 high variability between the speeds at which different people read, and we lack the ability to  
5  
6 measure passive interaction (that is, time when a student is actually reading without manipulating  
7  
8 the webpage), we can say with reasonable confidence that students read only about half of the  
9  
10 modules, at least on their first view. In other words, typically the first time that a student works  
11  
12 on a module, it is for credit purposes (to get credit for completing homework) rather than for the  
13  
14 purpose of understanding the material.  
15  
16  
17

18 Table 1: Quartiles for time spent between loading a module and starting an exercise.  
19  
20

Class	Quartile 1	Quartile 2	Quartile 3
C3	0-15 sec	16-90 sec	91-440 sec
C1	0-10 sec	11-70 sec	71-290 sec
C2	0-20 sec	20-120 sec	121-420 sec

21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35

### 36 3.1.2 Clicking through Slideshows 37 38 39

40  
41 We also observe two patterns of behavior related to viewing slideshows. We term “learning”  
42  
43 behavior to be when students read slide descriptions and navigate backward and forward as ne-  
44  
45 cessary to reexamine the content. We term “rushing” behavior to be when students click through  
46  
47 the slides as quickly as possible in order to obtain credit, without paying attention to the material.  
48  
49 Note that there are two forms of “receiving credit”. One is that a student actually receives points  
50  
51 (typically one tenth of a point per slideshow, with a module worth typically 2 to 5 total points)  
52  
53 when some criterion is reached to define “completion”. The other is that a green checkmark ap-  
54  
55 pears on the slideshow once the criterion is reached to define “completion”. In addition, once all  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65



1  
2  
3  
4 activities on a module page (including slideshows and exercises) have completion credit, the  
5  
6 module shows a “module complete” message on the page, and the gradebook and table of con-  
7  
8 tents also mark the module as “complete”. For some students, “checking off” the module seems  
9  
10 to be intrinsically important (separate from perhaps learning the material).  
11  
12

13  
14 To determine whether students rushed through slideshows, we examined the mean time per  
15  
16 slide on the occasions where a student first completed a slideshow. In this study, we did not ana-  
17  
18 lyze whether students who rush through slideshows return later and exhibit learning behavior  
19  
20 after obtaining completion credit. We will analyze that aspect of student behavior in subsequent  
21  
22 studies.  
23  
24

25  
26 Each slide within a slideshow contains text, most often a single line. Based on the time spent  
27  
28 we infer reading of the text or not as follows:  
29  
30

31 - If time spent on a slide is less than 8 seconds, we conclude that students did not  
32  
33 read the text on the slide (no reading).  
34

35  
36 - If time spent on a slide is between 8 and 15 seconds, we infer that the student read  
37  
38 the slide thoroughly (reading).  
39

40  
41 - If time spent is greater than 15 seconds, we conclude that the student stepped  
42  
43 away from the computer while the slideshow was up (stepped away).  
44

45  
46 Even though we choose our threshold arbitrarily, we think it provides a good estimate for  
47  
48 the occurrence of reading.  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

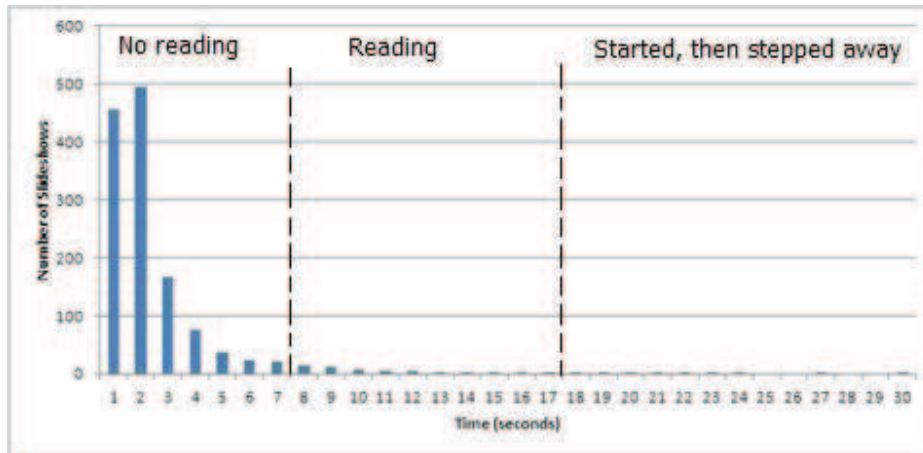


Figure 3: Mean time per slide, C1

Figure 3 shows the distribution of times for the C1 section. We see that more than 90% of the slide have a view time of less than 7 seconds. We also examined the class behavior through the use of quartiles of slideshows ranked by view time. Table 2 shows the different quartiles for C1, C2 and C3 courses.

Table 2: Quartiles for mean time per slide.

Class	Quartile 1	Quartile 2	Quartile 3
C3	1 sec	4 sec	8 sec
C1	1 sec	2 sec	3 sec
C2	1 sec	2 sec	3 sec

For all three classes, 25% of slideshows were completed with a mean-time-per-slide of one second or less. The second and third quartiles for both Virginia Tech classes indicate that 50% were completed with a mean less than two seconds and 75% with a mean less than three seconds. Both the second and third quartiles for the C3 class were larger than the third quartile of the Vir-

1  
2  
3  
4 ginia Tech classes, indicating that students in this class tended to spend longer on slideshows.

5  
6  
7 We hypothesize that this occurs because these students are non-native English speakers.

8  
9 Overall, it appears that a sizeable number of students in all classes quickly click through the  
10 slideshows. Calculating mean time per slide assumes constant time is spent on each slide, which  
11 is unlikely. Given the small means, this technique is unable to differentiate between students who  
12 click through all the slides quickly and those who initially exhibit learning behavior and then  
13 quickly click through the rest of the slides once they grasp the necessary concepts. Additionally,  
14 distractions could cause a single slide time to be significantly longer than the others, distorting  
15 the overall mean time. Future work using more reliable event data should examine individual  
16 slide times to account for these factors.  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27

28 An important strategy that we have begun incorporating into the design of our modules is to  
29 use many short slideshows that each illustrates one particular aspect of the topic, rather than a  
30 single, longer slideshow. We hypothesize that this will lead to more attention given to each slide-  
31 show.  
32  
33  
34  
35  
36  
37  
38  
39  
40

### 41 **3.1.3 Skipping to the End of Slideshows**

42  
43  
44

45 We knew prior to our experiment that we had a bug that allows students to obtain credit for  
46 slideshows without viewing all the slides by just jumping to end of the slideshow and clicking on  
47 the “forward” button once to get credit. We intentionally left this bug in place during our expe-  
48 riment in order to quantify how much students discovered and took advantage of it to artificially  
49 gain credit through this “skipping” behavior. (We have since changed completion criteria to re-  
50 quire that the student actually click through every slide in the slideshow, which of course will  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

1  
2  
3  
4 change behavior from what we observed in this study.) There was a small amount of slideshow  
5 skipping in C3 and C1, and a moderate amount in C2. The key difference is that in C2, credit  
6 was given for completing a slideshow. In C3, five students (24%) skipped slideshows. While one  
7 student skipped 8 of 16 slideshows, three of the five only skipped a single slideshow. In C1, 12  
8 of 51 (24%) students skipped slideshows. While 75% of the group skipped 5 or fewer slide-  
9 shows, three students skipped 8, 15 and 28 of the 44 slideshows. The median was 3.5 slideshows  
10 skipped. In C2, 23 of 65 students (35%) skipped at least one slideshow. 31% of the C2 group  
11 skipped 5 or fewer slideshows; the class median was 8 slideshows skipped. Inspection of the  
12 event data revealed two types of behavior. Some students viewed a portion of the slideshow,  
13 possibly moving backward and forward again to examine a specific operation, and then skipped  
14 to the end. Other students skipped straight to the end without viewing any slides.  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35

### 36 **3.1.4 Using AVs to Complete Exercises**

37  
38  
39

40 Our proficiency exercises require students to perform actions that closely resemble those  
41 demonstrated by some algorithm that is present using a related AV. We know from anecdotal  
42 data that some students ran the AVs using the exercise input, and then mimicked the AV output  
43 in order to complete the exercises. We sought to quantify this behavior. We determined that eight  
44 C1 students (16%) and seven students in C2 (11%) used AVs to assist in exercise completion.  
45 Table 3 provides details. We considered how the group of students who used AV assistance  
46 compared with their non-assisted classmates in terms of total time required for proficiency. Us-  
47 ing each student's median time for proficiency exercises, we found the assisted group in C1 had a  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

1  
2  
3  
4 median of 127.5 seconds, and the non-assisted group had a median of 84 seconds. The assisted  
5  
6 group in C2 had a median of 119 seconds, and the non-assisted group had a median of 77.5  
7  
8 seconds. In other words, copying the results from the AV to the exercise slowed the students  
9  
10 down (though we cannot determine from this how long the other students spent in gaining the  
11  
12 necessary knowledge to successfully complete the exercise).  
13  
14  
15  
16  
17  
18


19 Table 3: The number of exercises attempted, completed, completed with a score above the  
20  
21 proficiency threshold, and the number of exercises where students used an AV for assistance  
22

<b>Cl ass</b>	<b>Exercise</b>	<b>At- tempts</b>	<b>Com- pletions</b>	<b>Profi- ciency</b>	<b>As- sisted</b>
<b>C1</b>	Shellsort	671	67	57	1
	Mergesort	263	181	181	0
	Quicksort	1105	71	52	14
<b>C2</b>	Shellsort	561	90	68	2
	Mergesort	127	68	67	0
	Quicksort	1065	76	60	19


### 3.1.5 When do students study

44  
45  
46  
47  
48  
49  
50  
51 We analyzed data from the CS2 course offering at Virginia Tech during spring 2014  
52  
53 to find out when students actually study. We will refer to this course as C5. In C5, all the  
54  
55 assignments were due Sunday at 11:59PM, while the lecture sessions happened on Mon-  
56  
57 days and Wednesdays. The students usually had a week to complete a given assignment  
58  
59  
60  
61  
62  
63  
64  
65

1  
2  
3  
4 associated with the prior week's lectures. We analyzed our log data to see when the stu-  
5  
6  
7 dents usually work on OpenDSA exercises.  
8  
9  
10  
11



12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27 Figure 4: Daily distribution of exercises performed  
28



29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49 Figure 5: Daily distribution of all interactions  
50

51  
52 Figures 4 and 5 show that most exercises were done on the due date. The graph in Figure 6  
53 confirms that most interactions with OpenDSA happened on Sundays. The relative surge of ac-  
54 tivity on Tuesday, May 13 corresponds to the course final examination. This activity represents  
55 use of the materials as a study aid, rather than for completing graded assignments.  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

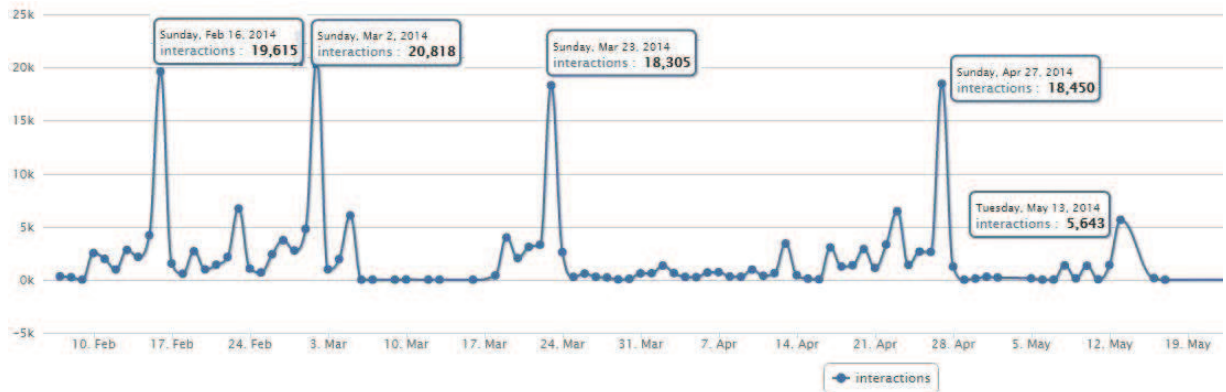


Figure 6: C5 semester-long OpenDSA activity distribution graph.

On Sundays, most of the work is done between 2pm and midnight as shown in Figure 7. The above results confirmed that most students waited until the last day to do their homework.

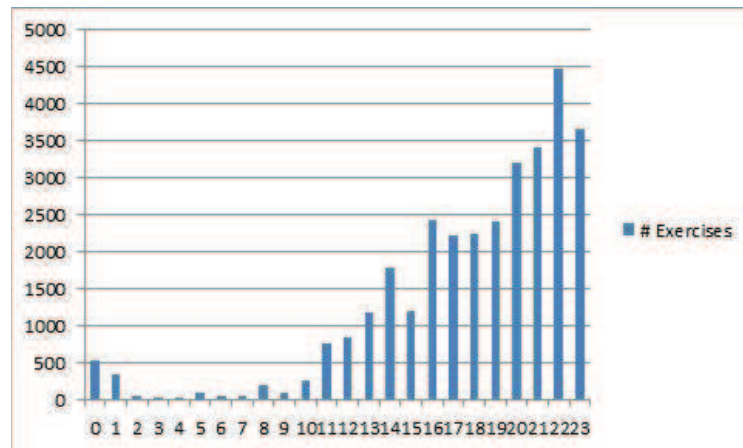


Figure 7: Hourly distribution of exercises performed on Sundays

#### 4. Student performance

In our opinion, the single most important pedagogical feature of interactive eText-books and OpenDSA is the interactive exercises. They provide the students with a mastery-based experience on a topic, since they allow the students to try exercises as many

1  
2  
3  
4 times as they want. We investigated the relationship between OpenDSA exercises and  
5  
6 student performance on tests. We only focused on students' declarative knowledge, hence  
7  
8 only students' performance on written tests are included in our analysis.  
9

10  
11 OpenDSA includes three categories of exercises: **Simple Questions**, **Proficiency**  
12 **exercises**, and **Programming exercises**. Simple questions are similar to traditional T/F,  
13  
14 multiple choice, or fill-in-the-blank questions. Proficiency exercises are visual algorithm  
15  
16 simulation activities. We collected the following data for each student:  
17  
18  
19

- 20  
21 - Total number exercises attempted (Total.KA for simple questions, Total.KAV for  
22  
23 KA-based proficiency exercises, and Total.PE for JSAV-based proficiency exercises).  
24  
25
- 26 - Total number of completed exercises (Correct.KA for simple questions, Cor-  
27  
28 rect.KAV for KA min proficiency exercises, and Correct.PE for JSAV-based proficiency  
29  
30 exercises)  
31  
32  
33

34  
35 We collected and analyzed the data from C4. The course's OpenDSA instance had a total of  
36  
37 95 required exercises, of which 36 were KA simple questions, 30 were KA mini proficiency ex-  
38  
39 ercises, 26 were JSAV-based proficiency exercises, and 3 were other interactive activities.  
40  
41

42  
43 We divided the data into two groups: M1 for all activity until the first midterm examination  
44  
45 and M2 for all activity between the first and second midterms. The final exam was on sections  
46  
47 with very little or no interactive elements in OpenDSA, so it was not included in our analysis.  
48  
49

50  
51 Before Midterm 1, students were required to solve a total of 15 KA simple question sets, 12  
52  
53 KA proficiency exercises, and 6 JSAV proficiency exercises. Between Midterm 1 and Midterm  
54  
55 2, they had to complete 19 KA simple question sets, 18 KA mini proficiency exercises and 20  
56  
57 JSAV full proficiency exercises. Figures 8 and 9 show the performance distribution. The x-axis  
58  
59  
60  
61  
62  
63  
64  
65



of the Correct.KA, Correct.KAV, and Correct.PE histograms represent the number of exercises correctly completed. In the case of KA simple questions and KA proficiency exercises it is the number of times a student reached the proficiency score. For JSAV proficiency exercises, it is the number of times a student reached the proficiency threshold (correctly simulated 90% of an algorithms step).

We see from the histograms that most students correctly completed most exercises only one time regardless of the type of question. The exception was the JSAV proficiency assignments before the second midterm, where the majority of students correctly completed exercises more than once. We can explain the exception by the fact that Midterm 2 covered the Sorting and Hashing chapters. Both included concepts that were fairly new to students. Midterm 1 covered topics where the students had had previous exposure (linear data structures, algorithm analysis and binary trees).

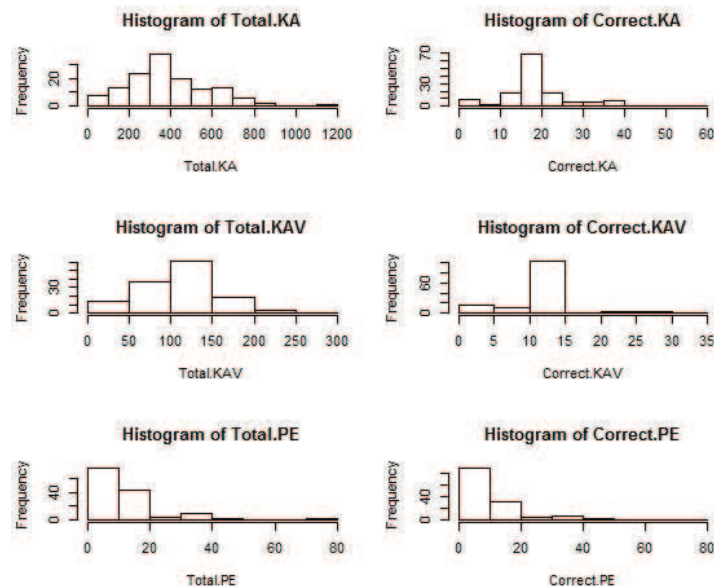


Figure 8: C4 - Midterm 1 distribution

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

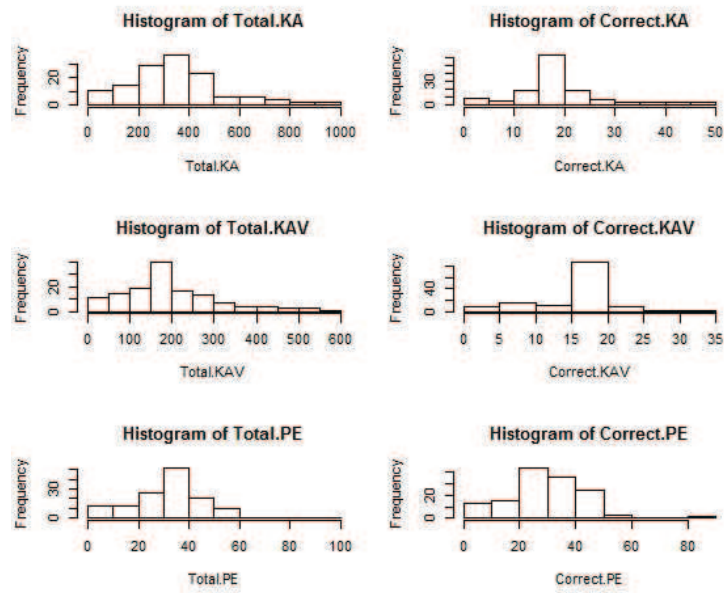


Figure 9: C4 - Midterm 2 distribution

Since there was no penalty associated with attempting an exercise several time, we expected to see a difference between the students who correctly completed the exercises “post-proficiency”.

We found a low correlation between the amount of correct exercises (regardless of the type) completed by the students and exam scores as shown in Table 4. The negative correlation between the total number of KA simple question attempted suggests that some students were just trying out all the provided answers until they find the right one, without deeply thinking about the exercise's answer, or reading the module text.

Table 4: C4 - Correlation table

Variables	Mid 1	Mid 2
Total.KA	-0.11	0
Correct.KA	0.32	0.29
Total.KAV	0.09	0.01

Correct.KAV	0.29	0.25
Total.PE	0.15	0.20
Correct.PE	0.19	0.20

We wanted to know if the analysis of student scores will show different patterns of OpenD-SA exercise use. We computed the average number of correct exercises performed by students grouped by midterm score quartile. As shown in Figure 10, students with higher scores correctly completed more exercises than those with lower grade. The difference in the average number of correct exercise between quartiles was significant at the  $p < 0.05$  level for KA simple questions and KA mini proficiency exercises. However, the difference was not statistically significant for JSAV proficiency exercises.

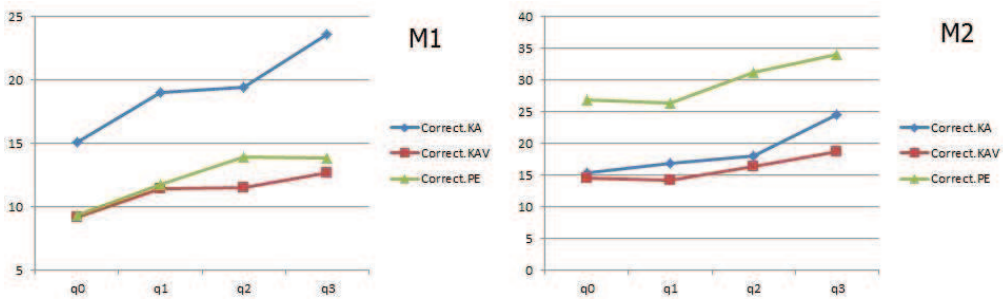


Figure 10: C4 - M1 and M2 variable means per quartile

We can say that when used as the main course material, a high number of correct exercises (above the minimum level required to receive homework credit) is associated with higher grades on written tests. That is, using the exercises for additional study appeared to help students on the midterm. The above findings can be used to detect struggling students, since they are less likely

1  
2  
3  
4 to redo an exercise. In addition, students with the lowest grades have the lowest “number of cor-  
5  
6  
7 rect exercises to number of exercises attempted” ratio.  
8  
9

## 10 11 **5. Students use of mobile devices** 12 13 14

15  
16 As handheld wireless mobile computer become more ubiquitous, their potential use and  
17  
18 benefit in education is being examined. Learning technology developers and users are excited  
19  
20 about the opportunities provided by mobile devices (Roschelle, 2003). In their meta-analysis of  
21  
22 mobile learning study, (Wu et al., 2012) concluded that most research in mobile learning studies  
23  
24 resulted in positive outcomes. Also, they found that mobile learning was most frequently used by  
25  
26 students in the professional and applied sciences field. However, Computer Science was the  
27  
28 second most investigated discipline (after languages and linguistic) in regard to mobile learning.  
29  
30 Other studies show that college students want to use mobile devices more for learning and hope  
31  
32 that their instructors and institution will create an environment that promotes the use of mobile  
33  
34 devices (Dahlstrom, Walker, & Dziuban, 2013).  
35  
36  
37  
38  
39

40  
41 OpenDSA gives us an opportunity to study the use of mobile device for learning in CS  
42  
43 courses. We looked at logged data to get a sense of the proportion of interactions originating  
44  
45 from mobile devices (smartphone, small size tablet, full size tablet). Since more students own  
46  
47 tablets (handheld devices with screen size above 7”), we are interested in measuring the use of  
48  
49 mobile device to access online tutorials. During Spring 2014, in addition to analyzing the logged  
50  
51 data, we surveyed students in C5 to uncover the factors that encouraged or prevented the students  
52  
53 from using their mobile device. In C5, students had to complete weekly OpenDSA assignments,  
54  
55 and OpenDSA homework accounted for 2% of the total course grade.  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

1  
2  
3  
4 Most students enrolled in courses using OpenDSA own a mobile device in addition to their  
5 Laptop/PC. We tracked the origin of user interactions between Fall 2013 and Spring 2014, and  
6  
7 found that less than 1.5% of the traffic originated from a mobile device.  
8  
9

10  
11 We surveyed students enrolled in C5 to gather their experiences and opinion about mobile  
12 devices and OpenDSA. We got a total of 51 responses. They revealed that:  
13  
14

- 15 - 68% of respondents own a laptop computer with no touch screen
  - 16 - 49% own a Touch screen laptop
  - 17 - 72% own a smartphone
  - 18 - 17% own a small tablet (approx. 7")
  - 19 - 15% own a Full-size tablet (approx. 10")
- 20  
21  
22  
23  
24  
25  
26  
27

28 The mobile devices ownership distribution in C5 is consistent with the general US popula-  
29 tion statistics<sup>1</sup>, and can be generalized at least to the other VT courses that we studied.  
30

31 From the log data, we found that out of 176 students enrolled in C5, only 5 accessed  
32 OpenDSA from a mobile device at least one time. The ratio of mobile users was equally low in  
33 other VT courses using OpenDSA. We did not expect smartphone owners to access the materials  
34 via their devices, because several factors make it difficult to effectively use them to study (screen  
35 size, battery life, and usability concerns). But even small- and full-sized tablet owners rarely  
36 tried to access the tutorials from their devices. This situation seemed to contradict previous study  
37 showing students positive attitude towards using full size tablets to study (Rossing, Miller, Cecil,  
38 & Stamper, 2012).  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51

52 We tried to determine if the low use of mobile devices was due to some serious usability bug  
53 within OpenDSA. When asked to ranked course resources by positive learning impact, students  
54  
55  
56  
57

---

58  
59 <sup>1</sup> <http://www.pewinternet.org/factsheets/mobile-technology-fact-sheet/> (accessed 09/13/2014)  
60  
61  
62  
63  
64  
65

1  
2  
3  
4 ranked OpenDSA as the second most important (after programming labs), thus proving that the  
5  
6 students did not have a negative opinion of the OpenDSA system. The low number of mobile  
7  
8 users seems to indicate that the reason not to use mobile devices was made independently of  
9  
10 OpenDSA, since the vast majority (97%) of students never accessed OpenDSA from a mobile  
11  
12 device.  
13  
14

15  
16 We then sought to see if the nature of the course made it harder for students to use mobile  
17  
18 devices. We asked students if they ever used their mobile devices to study in other courses. We  
19  
20 got 15 positive answers out of 51 responses. But there were only 2 positive answers when asked  
21  
22 if they used mobile devices to access OpenDSA content. Further investigation will be necessary  
23  
24 to understand the difference the difference. The most common reason mentioned by students to  
25  
26 justify their preference for non-mobile devices was the “convenience”. Most students said that it  
27  
28 is more convenient to use a laptop or a desktop computer when working on OpenDSA assign-  
29  
30 ments. They mentioned the bigger screen and the full keyboard of laptops/desktops as key ele-  
31  
32 ments, confirming the findings in Rossing et al., 2012. As we mentioned earlier, many OpenDSA  
33  
34 exercises are proficiency exercises, suitable for touch interaction, and do not require typing text.  
35  
36 Even the KA exercises do not require the user to type in significant text.  
37  
38  
39  
40  
41  
42

43 We saw earlier that most students are “procrastinators”, they wait until late to do their  
44  
45 homework. Several studies have shown that procrastinators experienced greater stress especially  
46  
47 toward the end of the assignment period when they are rushing to complete the assignment on  
48  
49 time (Tice & Baumeister, 1997). For procrastinator it makes no sense to include a mobile device  
50  
51 into the study workflow. They will go for the device that will give them more capabilities (bigger  
52  
53 screen, possibility to open multiple tabs, full keyboard, etc.) to complete the homework in a  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

1  
2  
3  
4 small amount of time. Under those circumstances, mobile devices cannot compete with lap-  
5  
6 tops/desktops.  
7  
8  
9

## 10 11 **6. Conclusion** 12 13

14  
15  
16 We analyzed the behavior of CS students when using the OpenDSA eTextbook sys-  
17 tem. We found that a majority of students skip directly to the exercises without reading  
18 the text, which indicates that the existing text is not engaging enough or is too over-  
19 whelming for students. It is also possible that since most students waited until the due  
20 date to do the homework, reading the text was thought to “slow them down” so that they  
21 might not be able to submit the homework on time. However, perhaps seemingly in con-  
22 tradiction to this finding, some students voluntarily completed additional exercises as a  
23 study aid when preparing for the exams.  
24  
25

26  
27 When students were not given credit to complete slideshows, they completed fewer  
28 than the students who were given credit. However, all of the students used slideshows to  
29 some extent, and overall use was fairly high. Skipping to the end of slideshows occurs  
30 more often when completing slideshows is part of the scoring; however, the behavior was  
31 still observed even when students did not receive credit.  
32  
33

34  
35 The low use of mobile device might be explained by a combination of structural and  
36 behavioral factors. Structural factors include device characteristics and some usability is-  
37 sues, especially for small screen size devices. There was a lot of procrastination, thus  
38 leading to limited amount of time to complete homework. It is possible that this has an  
39 influence on the device selected.  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65

1  
2  
3  
4 **References**  
5  
6

7 Dahlstrom, E., Walker, J. D., & Dziuban. (2013). ECAR Study of Undergraduate Students and  
8 Information Technology, 2013: EDUCAUSE Center for Analysis and Research.

9 Fouh, E., Karavirta, V., Breakiron, D. A., Hamouda, S., Hall, S., Naps, T. L., & Shaffer, C. A.  
10 (2014). Design and architecture of an interactive eTextbook - The OpenDSA system. *Science of*  
11 *Computer Programming*, 88(0), 18.

12 Karavirta, V., & Shaffer, C. A. (2013). *JSAV: The JavaScript Algorithm Visualization Library*.  
13 Paper presented at the Conference on Innovation and Technology in Computer Science  
14 Education (ITCSE), Canterbury, UK.

15 Malmi, L., Karavirta, V., Korhonen, A., Nikander, J., Seppälä, O., & Silvasti, P. (2004). Visual  
16 Algorithm Simulation Exercise System with Automatic Assessment: TRAKLA2. *Informatics in*  
17 *Education*, 3(2), 267-288.

18 Roschelle, J. (2003). Keynote paper: Unlocking the learning value of wireless mobile  
19 devices. *Journal of Computer Assisted Learning*, 19(3), 260-272. doi: 10.1046/j.0266-  
20 4909.2003.00028.x

21 Rossing, J. P., Miller, W. M., Cecil, A. K., & Stamper, S. E. (2012). iLearning: The Future of  
22 Higher Education? Student Perceptions on Learning with Mobile Tablets. *Journal of the*  
23 *Scholarship of Teaching and Learning*, 12(2), 1-26.

24 Shaffer, C. A., Karavirta, V., & Naps, T. L. (2011). *OpenDSA: Beginning a Community*  
25 *Hypertextbook Project*. Paper presented at the Koli Calling International Conference on  
26 Computing Education Research, "Koli National Park, Finland.

27 Shaffer, C. A., Naps, T. L., & Fouh, E. (2011). *Truly Interactive Textbooks for Computer*  
28 *Science Education*. Paper presented at the Sixth Program Visualization Workshop, Darmstadt,  
29 Germany.

30 Tice, D. M., & Baumeister, R. F. (1997). Longitudinal Study of Procrastination, Performance,  
31 Stress, and Health: The Costs and Benefits of Dawdling. *Psychological Science*, 8(6), 454--458.  
32 doi: 10.2307/40063233

33 Wu, W.-H., Jim Wu, Y.-C., Chen, C.-Y., Kao, H.-Y., Lin, C.-H., & Huang, S.-H. (2012). Review of  
34 trends from mobile learning studies: A meta-analysis. *Computers & Education*, 59(2), 817-827.  
35 doi: <http://dx.doi.org/10.1016/j.compedu.2012.03.016>  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54  
55  
56  
57  
58  
59  
60  
61  
62  
63  
64  
65



## Acknowledgements

### Acknowledgements

This work is supported by the National Science Foundation, under grants DUE-1139861 and IIS-1258471.