

Time and Attention: Students, Sessions, and Tasks

Andrew Arnold, Richard Scheines, Joseph E. Beck, Bill Jerome

Center for Automated Learning and Discovery and
Open Learning Initiative
Carnegie Mellon University
{aarnold@cs, scheines@andrew, joseph.beck@cs, wjj@andrew}.cmu.edu

Abstract

Students in two classes in the fall of 2004 making extensive use of online courseware were logged as they visited over 500 different “learning pages” which varied in length and in difficulty. We computed the time spent on each page by each student during each session they were logged in. We then modeled the time spent for a particular visit as a function of the page itself, the session, and the student. Surprisingly, the average time a student spent on learning pages (over their whole course experience) was of almost no value in predicting how long they would spend on a given page, even controlling for the session and page difficulty. The page itself was highly predictive, but so was the average time spent on learning pages in a given *session*. This indicates that local considerations, e.g., mood, deadline proximity, etc., play a much greater role in determining student pace and attention than do intrinsic student traits. We also consider the average time spent on learning pages as a function of the time of semester. Students spent less time on pages later in the semester, even for more demanding material.

1. Introduction

Students vary their attention and pace when they engage in any learning activity at a given time, be it reading material, online course material or an intelligent tutor [Brusilovsky]. Some of this variation is due to individual differences. That is, some students are disposed to attend longer than others or go through material more quickly than others. Another part of this variation may be due to differences within an individual student across different times. That is, in some learning sessions a given student might read slowly and intensely and in another skim quickly. Finally, some part of the attention given to an activity will be a function of the activity itself, or of the stakes put on work associated with the activity, or of the time in the semester in which the activity occurs, etc [Stern].

In this paper we discuss and analyze log data collected from over 40 students in two classes in the fall of 2004,

each of which covered, primarily online, at least six weeks of a course on Causal and Statistical Reasoning (CSR)¹.

The CSR courseware is delivered in modules, each of which is approximately one week's worth of material in a college course. Each module is separated into approximately ten to twenty learning pages. Each learning page has some combination of text, pictures, simulations, interactive virtual labs, and one or two voluntary multiple-choice comprehension checks. At the end of each module the online course randomly constructed a quiz from a pool of items. Students may take each module quiz up to three times, and they must exceed a given percentage (e.g., 70%) on at least one attempt by a certain date to get credit for the module.

The learning pages vary considerably in length, with some involving only two or three paragraphs of text and others involving up to five pages of text, simulations, virtual labs and comprehension checks. So clearly the demand on a student's attention and time will vary substantially over different learning pages.

Students' behavior on the pages varies considerably as well, with some students skimming quickly through the text, barely stopping to open a simulation or a voluntary comprehension check, while others do everything slowly and thoroughly. Some students eschew the online setting and opt to print out the modules instead of doing them online, but unless these students return to the online setting they miss all the interactive content. From previous research [Scheines], we know that more successful students behave quite differently than unsuccessful ones, where success is measured by post-module quiz performance and midterm exam scores.

Successful students do not necessarily take more time, but they do the comprehension checks and interactive exercises. In data collected from both the University of Pittsburgh and UC San Diego, the correlation between the proportion of voluntary comprehension checks attempted and the average module quiz score was above .8. Printing out modules was negatively associated with exam

¹ See Causal and Statistical Reasoning at: www.cmu.edu/oli

performance, even controlling for pre-test and attendance at discussion sections.

In this work we begin to explore what we can learn from analyzing the time students spend in the online setting. We believe that by attending to the duration of student engagement on a variety of activities we can learn a lot about their general learning patterns and local mood, and eventually react intelligently to real-time analysis based on this research.

2. Measures

Our data consists of the experience of 45 users in two different classes in the fall of 2004. One class, offered at the University of Pittsburgh as a way to satisfy the quantitative reasoning requirement, covered 13 of 16 possible modules over the course of an entire semester. Another, offered at Carnegie Mellon as an upper level seminar on Causation and Social Policy, covered 16 modules over the course of 6 weeks.

Each time a student logs onto the system, every action the student takes before logging off the system is considered to be an action in a single "session." Each time a student requests a new learning page, simulation, multiple-choice question, lab activity, or quiz from the course server, we create a time stamped log of the event. In this paper we focus on the time students spend on learning pages, without focusing on the time spent on the various activities that might be embedded in these pages, e.g. simulations and virtual lab exercises.

Each time a student visited a new learning page during a given session, we computed the time spent on that page by computing the difference in time between when they visited that page and when they next visited a different learning page. Since in some instances this led to recorded durations of several hundred minutes, we threw out any duration that exceeded 10 minutes. We discarded such data rather than Winsorize them since large durations were not necessarily evidence the student spent a large amount of time perusing a page; he could have simply loaded the page and walked away from the computer.

To avoid poor estimates of mean duration, we set a minimum of five observations to include an item in our data set: we excluded sessions that contained fewer than five learning pages, learning pages that were seen fewer than five times during the year, and students who saw fewer than five pages during the year. Overall, this gave us 356 separate sessions, 308 distinct pages visited, and 6,431 different page visits for 28 unique students.

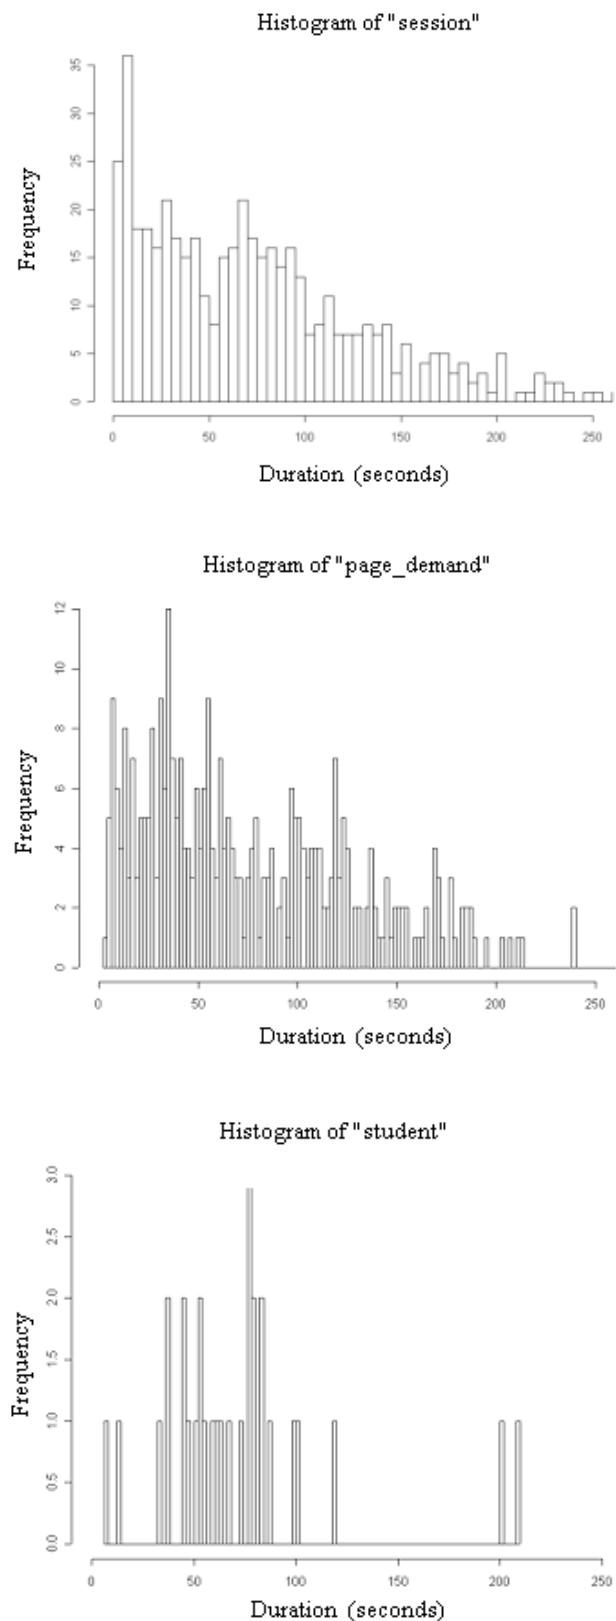


Figure 1. Histograms of frequency vs. duration in seconds

We computed the following measures from these data:

- $page_demand(i)$: mean time spent over all visits to learning page i
- $session(j)$: mean time spent over all visits to learning pages during session j
- $student(k)$: mean time spent over all visits to learning pages by student k

$Page_demand$ varied considerably across different pages, as did $session$ across different sessions and $student$ across different students. In Figure 1 we show histograms of each of these variables.

3. Results

We examine what impacts the amount of time students spend on a page, and investigate both static properties of pages and students, as well as temporal effects for time of semester.

3.1 Predicting Time Spent on a Learning Page

To predict the duration of a particular visit to a particular page, we estimated the following function, where \mathcal{E} is a noise term capturing factors not measured and random noise.

$$Visit_Length = f(page_demand(i), session(j), student(k), \mathcal{E})$$

Using multiple linear regression, we get the following results:

	Beta	Std. Err	t value	P
$page_demand$	0.83656	0.02075	40.32	$< 2e-16$
$session$	0.83826	0.02181	38.43	$< 2e-16$
$user$	0.14119	0.04681	3.016	0.00256

The p-values from the regression model must be taken with a grain of salt, since the observations are correlated with each other (e.g. a single student has multiple rows in the data set used to estimate the regression parameters). Therefore, the p-values are under estimated. However, the qualitative picture of each variable's contribution to the regression model is well estimated. Although the mean time spent by the user ($user$) is statistically significant, at this sample size (~9,000), the predictive impact of $user$ is extremely small, and is completely dwarfed by the effect of $session$ and $page$. The adjusted R-square for the same model with only $page_demand$ and $session$ as predictors is .314; adding $user$ raises it by .001 to .315. Similarly, if we use locally-weighted regression with $\alpha=.5$, with just $page_demand$ and $session$, we get a general cross validation mse of 2918, which is relatively unchanged (2921) after including $user$. Clearly, the page current

page and the student's current purpose or mood have a much bigger effect on how long they spend on a given page than does the identity (and thus general disposition) of the student.

3.2 Time Spent as a Function of Point in Semester

Another way of examining how students use the website is to look at how their behavior changes over time and as a function of their class context. In the CSR course, content was divided into 16 modules which were presented chronologically in both the University of Pittsburgh (Pitt; Instructor=Smith) and the Carnegie Mellon (CMU; Instructor = Scheines) classes.

In the Pitt class, 13 modules were covered over an entire semester, with the class focusing entirely on the module content. In the CMU class, all 16 modules were covered in 6 weeks, but because of the fast pace students had a much lower threshold for passing the end-of-module quizzes. We had two predictions about student performance on the course material: first, we predicted that the CMU students should spend much less time on the material; second, that all students would spend more time on the material as the semester progressed, as the material is cumulative and learning it for the final exam loomed closer and closer.

The first prediction turned out to be accurate, but not the second. Some students decreased the amount of time they spent on a learning page and on a whole module as the course progressed.

This decrease in usage shows up not just in the number of times a module's pages were visited, but also in terms of the average amount of time spent by students looking at module pages. In Figure 2 we plot, for each module covered in the two courses, the average time per page visit, per student, for each module. Blue diamonds, with a solid linear trend line, mark Pitt students, while red squares, with a dashed trend line, mark CMU students.

First, CMU students had to follow such a tight schedule that they almost had to be efficient with their time from the outset, and so were at nearly a floor for time spent per page. Pitt students perhaps became more and more acclimated to the system over time and thus became progressively more efficient as time went on. This would also coincide with an increase in academic demand from other courses.

It should be noted that the amount of time spent by the Pitt students on module number 14 could be an artifact, since these students covered only 13 modules, skipping the surrounding modules number 13, 15 and 16. If this data point at module 14 were removed, the linear regression line for Pitt student time spent per module would be flatter, but still have a negative slope.

Module Usage vs. Module Number per Section




Figure 2. Time spent per module, per page, per student, for Pitt and CMU sections, over the semester

4. Conclusions, Limitations and Future Work

Clearly students change their behavior in response to different learning pages and behave quite differently from one online session to the next. Interestingly, engagement on a learning page seems to be predicted by the page itself and the local session they are in, and (once conditioned on page and session) practically independent of any intrinsic property of the student, as represented by the mean time spent per page per student.

Although the session was predictive, it is undoubtedly a proxy for other features of the student's life that make a big difference in how he attends to learning. For example, proximity to a deadline or test might matter, proximity to a social engagement might matter, or energy level in general might matter. In future analyses we will include more direct measures of these features whenever possible, and hopefully tease apart the factors that make one session different from another.

Over a semester, students should become more efficient both at processing the information in an online course and in determining which information is likely to be tested or

graded in their future. Thus it makes good sense that time spent per page decreases over the course of the semester.

Likewise, the total time spent per page can be further broken down into the time spent on various activities within that page, such as reading, doing lab exercises, watching videos, etc. By separating and comparing the time spent on these activities, along with the cumulative time spent, we hope to see a more nuanced picture of what specific features of the on-line course are most predictive of a student's learning.

In our next round of studies we will also include a performance measure in our analyses, such as pre and post tests. We left such analyses out of this paper because the connection between time spent and performance is not likely to be simple. Students who spend very little might be either lazy or smart enough to "get it" very quickly. Students who spend a lot of time on a page might be legitimately curious, dedicated to getting a good grade, or simply so confused that they want more information before proceeding.

Another complicating factor is our end of module quizzes. Our students are allowed to take these quizzes up to three times. Some students adopt a strategy in which

they skim through a module quickly, take the quiz as an early warning device, and then go back and spend real time only on those pages that offer material they didn't know on the quiz. Others will read everything carefully the first time through and then try to pass the quiz on the first attempt.

5. References

Brusilovsky, Peter: Adaptive and Intelligent Technologies for Web-based Education. *KI* 13(4): 19-25 (1999)

Scheines, R., Leinhardt, G., Smith, J., and Cho, K. 2005. Replacing Lecture with Web-Based Course Materials, *Journal of Educational Computing Research*, 32, 1.

Stern, M. K., and Woolf, B. P. (2000) Adaptive Content in an Online Lecture System. *In The Proceedings of Adaptive Hypermedia and Adaptive Web-Based Systems* (Trento, Italy, August 2000).