Improving Current and Future Offerings of a Data Science Course through Large-Scale Observation of Students

Tabitha Belshee[1], Adam Chang[1], Nebil Ibrahim[1], Mikako Inaba[1], Nikoo Karbassi[1], Angelo Kayser-Browne[1], Hye Jee Kim[1], Rachel Kim[1], Seungjae Ryan Lee[1], Natalia Orlovsky[1], Michael Guerzhoy[1,2,3]

Princeton University[1], University of Toronto[2], Li Ka Shing Knowledge Institute[3]

Abstract

We delivered a large Introduction to Data Science course with a team of undergraduate Teaching Assistant-Researchers (TARs) who both helped students in the lab and collected qualitative observations about student learning. TARs concurrently participated in a senior-level Pedagogy of Data Science seminar. We present a strategy for collecting and systematizing a large body of qualitative observations. We then apply this methodology to our collected data and present actionable conclusions about students' learning trajectories and outcomes, which can be used to improve future offerings of the course. Finally, we present evidence that participating in the study raised student performance on an end-of-semester test by 0.4σ (CI: [0.1σ, 1.8σ], $p = 0.02$), where σ is the class standard deviation.

Introduction to Data Science

Our *Introduction to Data Science* course is aimed at a broad audience. The course covers programming for data science, predictive modelling, and the foundational level of statistical inference. The course overlaps substantially with traditional CS1, and does not require programming experience as a prerequisite. In total, about 130 students were enrolled in the course during spring 2020. Of this group, 30 students opted into our observational study.

We take a functional approach to teaching programming in the R programming language. Since we focus on processing data stored in data frames (tables), repeated computation tasks involve processing every row of a data frame, or every element of a column. This means that it is natural to use higher-order functions such as **dplyr**'s **summarize** and **mutate** and R's **sapply** (the *map* function). In contrast to the usual approaches in CS1, although our students practice using repeated computation every week, they only see for-loops in passing, and they do not see recursion at all. This approach is consistent with the usual programming style using the **tidyverse** libraries

Pedagogy of Data Science

A team of 10 Undergraduate Teaching Assistant-Researchers (TARs) participated in the delivery of the course. In parallel to working as teaching assistants, they were enrolled in a senior-level seminar, Pedagogy of Data Science. In the seminar, we discussed approaches to teaching data science and programming for data science, pedagogical research techniques, organizing and conducting this study, as well as topics in data science and programming for data science from an advanced point of view.

Background: Qualitative Studies in CS Ed

Several styles:

- A moderate number of interviews that are coded to systematize the findings (e.g. Petersen et al.)
- Rich observations of as few as two study subjects (e.g. Lewis)

Large-Scale Observation of Students

- Detailed observations made by 10 TARs of about 30 students
- Each TAR proposed *constructs* related to their observations
- TARs engaged with each others' constructs
- A synthesis of the TARs' observations was produced

Large-Scale Observation of Students

- Detailed observations made by 10 TARs of about 30 students
- Each TAR proposed *constructs* related to their observations
- TARs engaged with each others' constructs
- A synthesis of the TARs' observations was produced
- Large-scale observations supported by the fact that TARs are enrolled in a seminar that deals with research methods

Some Constructs and Observations

- **Programming ``from the ground up" vs. modifying sample code.**
  - Some students over-rely on sample code
  - Other students don't reach for sample code well enough
    - Particularly in the context of repeated computation with **sapply** and simulation-based inference
  - **Conclusion:** explicitly discuss in class when and how to use examples
- **Working with partners**
  - Observed working relationships evolving over time, trust developing and resulting in a freer interaction
  - Observed insight developing through explaing code to a partner
- **Confidence and self-efficacy**
  - Observed students gaining confidence, and stopping to ask for validation from their partner
  - Observed confidence boost when students completed all the problems in the lab
    - **Conclusion:** consider trimming down labs
- **Virtual labs**
  - Observed students collaborating more closely because they could share screens
    - **Conclusion:** consider replicating this aspect in person

Using the Observations

- Obtained much more systematic and large-scale observations than would be available to a course instructor
- Interesting to note which observations were *not* made
  - E.g., TARs did not observe that mathematical background affected student performance

Effect of Participating in the Study

- Want to estimate the effect of participating in the study
- Did not directly measure prior experience or enthusiasm for learning
- Predict performance on Test 2
  - Participation in the study
  - Test 1 score
    - Idea: Test 2 score is only influenced by prior experience through the Test 1 score
  - Extra credit assignment score
    - Proxy for enthusiasm for learning
  - Participation in the study associated with an increase in Test 2 scores by .4σ (CI: [0.1σ, 1.8σ], $p = 0.02$), where σ is the class standard deviation (controlling for Test 1 and bonus score)
  - Pre-registered design
  -

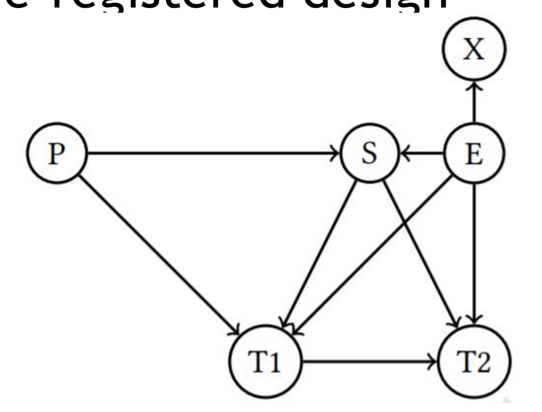

A possible causal model that treats enthusiasm for learning as an unobserved confounder. E represents enthusiasm for learning in the course, and X represents a challenging extra-credit assignment.

Conclusions

- We present a way to collect large-scale qualitative observations by running a pedagogy seminar
- Evidence that students in the introductory class benefited from participating in the study
- Actionable observations were collected

Colleen M. Lewis. 2019. A Case Study of Qualitative Methods. In *The Cambridge Handbook of Computing Education Research*, Sally A. Fincher and Anthony V. Robins (Eds.). Cambridge University Press, 875âĂŞ894.

Andrew Petersen, Michelle Craig, Jennifer Campbell, and Anya Tafliovich. 2016. Revisiting why students drop CS1. In Proceedings of the 16th Koli Calling International Conference on Computing Education Research. 71-80