Building on the Shoulders of Bears: Next Steps in Data Science Education

Dustin Tingley^{1,2,3}

¹Department of Government, Harvard University, Cambridge, Massachusetts, United States of America,

²Harvard Initiative on Learning and Teaching, Vice Provost for Advances in Learning Data Science and Technology Group, Harvard University, Cambridge, Massachusetts, United States of America,

³ABLConnect, Program on Experienced Based Learning in the Social Sciences, Harvard University, Cambridge, Massachusetts, United States of America

Published on: Jun 07, 2021

DOI: <https://doi.org/10.1162/99608f92.46f66c22>

License: [Creative Commons Attribution 4.0 International License \(CC-BY 4.0\)](https://creativecommons.org/licenses/by/4.0/)

Adhikari et al. of UC-Berkeley (Bears) provide a thrilling tour of one of the most ambitious curricular experiments in modern history (“[Interleaving Computational and Inferential Thinking in an Undergraduate Data Science Curriculum](#),” this issue). They provide not just a conceptual tour de force taken from a myriad number of academic fields, but also an articulation of why they designed the series of courses, along with their connector courses, in the way that they did. The efforts described in their article are not to be underestimated. Shaping an interdisciplinary program from the ground up at any large academic institution is extremely challenging. And omitted from the article is a characterization of the many challenges they were able to overcome in doing so. While they document the impact that they have had at Berkeley, and through the delivery of Data 8 via the edX platform, omitted from the article is the profound impact they have had on other institutions and people. This has come in the form of other institutions directly drawing on their curriculum and design. But it has also come in the form of catalyzing numerous conversations at other institutions that do not directly draw on their work but are certainly inspired by them.

In light of these remarkable accomplishments, I would like to briefly offer some thoughts that are more in the realm of ‘what is next?’ My thoughts here are heavily influenced by both an administrative role to help advance teaching and learning, as well as a practicing data scientist operating within the social sciences.

First, there are numerous opportunities to reach learners that extend beyond those who are enrolled in B.A., M.A., and Ph.D. programs. As documented in this journal (Chen, 2020), there are substantial opportunities for younger students to begin exploring data science. Many years ago I taught at the high school level and had the opportunity to teach a hybrid course on statistics and game theory. Needless to say, the statistics components were dry and disconnected rehashings of basic statistical concepts. In contrast, the more modern data science turn can expose young learners to concepts around algorithms that are now part of their everyday life. And we can expose young learners to the fact that not all ‘data’ sits nicely in clean spreadsheets. Music is data, text is data, and so on. Exposing students to these ideas does not require detailed math or even programming. Of course, those components can be built into such curricula, but they need not be barriers. On the flip side are opportunities for reaching individuals who will never take a technical data science course but instead want to be literate in data science thinking. These individuals might even work with or manage data scientists. As such, this requires different types of content and pedagogy than what team Berkeley has so impressively built up.

This leads to my next observation—well probably more of an opinion. We need to make sure that we are leading with questions or problems, rather than data and algorithms. Indeed, the authors describe impressive ways in which they link their content to making ‘decisions.’ And their connector courses pull a lot of weight in this respect. But problem definition—What is the problem we are trying to solve?—is too often neglected in industry and academia when it comes to data science. This creates a trap wherein impressive resources are deployed, but any insights, inferences, and so on that come out are unable to be used in practice or even in principle.

Connecting to the first observation about broadening out who can be reached with a data science education, part of this problem is that managers, decision makers, are not posing problems in coherent ways that data scientists can then act upon. This plays out in academia as well. Too often I hear graduate students marveling over how long it took them to put together a massive new data set. And yet when you ask, ‘What questions will you be able to answer?’ the crickets chirp too often. Indeed, I sometimes catch myself in this trap. And the same concern even holds for those developing new algorithms.

Enforcing problem-centered approaches to educating aspiring data scientists themselves is also important because it will help them communicate more effectively—and with greater confidence—with managers and decision makers in this organization. And it might well be that encouraging data scientists to take purely non-data courses will help in this. I see great benefit in such students taking courses in microeconomics so that they better understand concepts like compliments and substitutes, thinking on the margin, and so on, game theory so that they understand the primacy of strategic interaction in many contexts, and ethics courses to understand different notions of freedom and liberalism. As such, establishing better ‘two-way’ communication between data scientists and those who work alongside them or those who draw on their talents is crucial. This focus forms part of the pedagogical bedrock of [a series of online courses](#) that Harvard is producing.

Finally, I think there are several content areas that data science education might do well to beef up. First, more discussion of data quality and the attendant data wrangling or cleaning is needed. These steps are often needed to transform data into a place where it can reasonably be used to answer a question or solve a problem. It is covered in the great Berkeley curriculum. But when in the trenches, it becomes a *massive* part of day-to-day life (indeed, a sort of inside joke among data scientists inspiring many social media memes). And this goes beyond the mechanical aspects of data wrangling (merges, filtering, reshaping, etc.). It includes thinking about data quality from the perspective of, what process generated this data. Are there selection effects such that causal claims might be difficult to support? What is this data representative of? Second, causal thinking is at times incorporated, but often is somewhat marginalized in favor of computational or inferential statistics topics. If you are making causal claims based on data with severe endogeneity problems, I do not really care all that much about the computational tools or the size of a test-statistic or out-of-sample performance. Research design and critical thinking becomes crucial (e.g., see Bueno de Mesquita & Fowler, 2021). Finally, how can we learn more by reflecting on the data we do not have? This “Dark Data” (Hand, 2020) is out there. Much can be done to reflect on what this lack of data implies about the problems we are trying to solve.

Disclosure Statement

Dustin Tingley has no financial or non-financial disclosures to share for this article.

References

Bueno de Mesquita, E., & Fowler, A. (2021). *Thinking clearly with data: A guide to quantitative reasoning and analysis*. Princeton University Press.

Chen, A. (2020). High school data science review: Why data science education should be reformed. *Harvard Data Science Review*, 2(4). <https://doi.org/10.1162/99608f92.1e28ce9e>

Hand, D. J. (2020). *Dark data: Why what you don't know matters*. Princeton University Press.

©2021 Dustin Tingley. This article is licensed under a Creative Commons Attribution (CC BY 4.0) [International license](#), except where otherwise indicated with respect to particular material included in the article.