
10-2-2017

Greater Data Science at Baccalaureate Institutions

Amelia McNamara
Smith College

Nicholas J. Horton
Amherst College, nhorton@smith.edu

Benjamin S. Baumer
Smith College, bbaumer@smith.edu

Follow this and additional works at: https://scholarworks.smith.edu/sds_facpubs

Recommended Citation

McNamara, Amelia; Horton, Nicholas J.; and Baumer, Benjamin S., "Greater Data Science at Baccalaureate Institutions" (2017). Statistical and Data Sciences: Faculty Publications, Smith College, Northampton, MA. https://scholarworks.smith.edu/sds_facpubs/35

This Article has been accepted for inclusion in Statistical and Data Sciences: Faculty Publications by an authorized administrator of Smith ScholarWorks. For more information, please contact scholarworks@smith.edu

Greater data science at baccalaureate institutions
Amelia McNamara, Nicholas J Horton, Benjamin S Baumer
October 25, 2017

Donoho’s paper is a spirited call to action for statisticians, who he points out are losing ground in the field of data science by refusing to accept that data science is its own domain. (Or, at least, a domain that is becoming distinctly defined.) He calls on writings by John Tukey, Bill Cleveland, and Leo Breiman, among others, to remind us that statisticians have been dealing with data science for years, and encourages acceptance of the direction of the field while also ensuring that statistics is tightly integrated.

As faculty at baccalaureate institutions (where the growth of undergraduate statistics programs has been dramatic [2]), we are keen to ensure statistics has a place in data science and data science education. In his paper, Donoho is primarily focused on graduate education. At our undergraduate institutions, we are considering many of the same questions.

We enthusiastically concur with Donoho’s description of a “Greater Data Science” comprised of

1. Data Gathering, Preparation, and Exploration
2. Data Representation and Transformation
3. Computing with Data
4. Data Modeling
5. Data Visualization and Presentation
6. Science about Data Science

and aim to have our students develop all these key capacities in our courses and major programs.

In considering our curriculum development, we have been guided by the 2014 American Statistical Association (ASA)’s *Curriculum Guidelines for Undergraduate Programs in Statistical Science* [1] and the 2016 *Guidelines for Assessment and Instruction in Statistics Education (GAISE) College Report* [13]. Both documents highlight the need for students to work with real problems, messy data, and complex models.

Even more recently, a working group (including Baumer) developed the *Curriculum Guidelines for Undergraduate Programs in Data Science*, which have now been endorsed by the ASA [15]. This forward-thinking document addresses one of Donoho’s primary concerns with data science education—that it may end up being a piecemeal collection of extant courses, with little “long-term direction.” While [15] does provide guidance to institutions working with existing courses, it also outlines a model curriculum with a number of new and reformulated courses.

Data science developments at our institutions

Both the Smith College major in statistical and data sciences and the Amherst College major in statistics have been explicitly structured to introduce, extend, and integrate work in all six of the areas of Greater Data Science. Real problems have been interwoven into our courses at multiple

levels. This has required extensive revision of existing courses along with the creation of a number of new and courses with complementary learning outcomes.

At both Smith College and Amherst College, the introductory course touches on all six GDS elements, with an increased emphasis on visualization and modeling [6, 21]. In subsequent courses like *Multiple Regression* or *Intermediate Statistics*, students explore, prepare, clean, transform, and visualize data. In the *Communicating with Data*, *Visual Analytics*, and *Multivariate Data Analysis* courses, students learn principles of data visualization and presentation of data. Modeling is reinforced in *Multiple Regression* and *Machine Learning*. Capstone courses help to integrate prior course work with project-based learning while further refining computing and communication skills.

Existing Amherst College theory courses such as *Probability and Theoretical Statistics* have been restructured to integrate computing as an explicit learning outcome (e.g., how to write a function, how to perform simulations, how to undertake empirical problem solving to complement analytic results, and how to collaborate in groups using GitHub).

At Smith College, *Introduction to Data Science*, *Communicating with Data*, *Visual Analytics*, and *Machine Learning* are all new offerings guided by our understanding of data science as its own discipline.

We would like to draw particular attention to *Introduction to Data Science*, a successor to the course described in [4] that is offered at both institutions. Donoho makes reference to this course, which teaches data visualization, data wrangling, ethics, SQL, and communication, using a new textbook [7]. The course is tied together by *liberal arts modules*, where professor from other disciplines outline a question relevant to their discipline, and the students seek to address it using their new-found data skills.

As Donoho reminds us, some academic statisticians have long been guilty of eschewing data analysis. But even some programs in data science focus more on tools and skills rather than developing the capacity to solve real problems. We believe our positions at liberal arts colleges give us a particular ability to reach across disciplines, connecting to data in the sciences, social sciences, and the humanities. The integration of liberal arts modules in *Introduction to Data Science* can be used as a model for similar courses.

Another learning outcome in all of our courses is to produce students who learn how to learn. As with many disciplines, data science is evolving quickly. The tools we teach our students today may not be relevant in five years. In fact, several of the R packages referenced by Donoho (`reshape2` and `plyr`) have now been supplanted by others (`tidyr` and `dplyr`) [26, 27]. As instructors, we do our best to stay on top of the current computational trends to provide our students with the most contemporary methods, which requires us to continually modify our curriculum. However, the focus is on generalized problem-solving that can be applied using different tools in different settings.

Ethical precepts are an important part of any data science program. Donoho alludes to this with his detailed coverage of the University of California–Berkeley Master’s program, which includes a course now titled “Behind the Data: Humans and Values” (formerly “Legal, Policy, and Ethical Considerations for Data Scientists”) [23]. At Amherst ethics is now included as a learning outcome in the *Intermediate Statistics* course with subsequent extension and reinforcement in elective and capstone courses. Ethics is also a component of the *Introduction to Data Science* courses. Students consider questions like those posed in [8]: what are the ethical implications of data science products? Who has access to data science, and who does not? What are our ethical obligations to our clients,

ourselves, and our subjects? These higher-level questions make up a key part of the capstone courses.

At all levels, our courses emphasize best practices of statistical computing and reproducible research. These efforts build upon scholarly work that goes back at least to Don Knuth’s literate programming [17] and Donoho’s previous work on reproducibility [12]. Baumer and McNamara are former faculty fellows on Project TIER: Teaching Integrity in Empirical Research [3], which aims to spread good computing and data practices to the social sciences. We are now seeing evidence of adoptions at our institutions, and others, where faculty members in economics, psychology, and environmental science and policy integrate reproducible research into their coursework, further strengthening our pool of data-capable students.

Data science scholarship

Beyond our interest in the pedagogy of data science, we are also researchers. However, this is an area that is also undergoing development. Since it is an emerging field, institutions must determine how to judge new types of scholarly production. Like many problems of data science, this is something that applied statisticians have been wrestling with for decades. However, not all data science work is precisely applied statistics (thus, the new degree programs and scholarship).

Much like Donoho’s notion of Science about Data Science, Jeff Leek has been proposing the idea of Data Science as a Science [18]. While Donoho’s examples focus on meta-analysis, Leek’s conception includes hands-on research. Calling on examples like Cleveland’s study of graphical perception [14], Leek advocates for data scientists experimenting to learn how software syntax impacts learning, and how practitioners are actually working (like [22]).

As a case study of scholarly production in data science, consider Hadley Wickham’s many contributions. Wickham’s work often centers on a profoundly useful R package. However, each piece of software fits into a higher-level framework of intellectually-weighty ideas. The ideas behind `ggplot2` were articulated in a book on implementing Wilkinson’s Grammar of Graphics [25, 28]. In addition to `tidyr`, Wickham wrote a article in the *Journal of Statistical Software* on the concept of tidy data, which transcends the language it is implemented in [26]. Although these works are highly-cited, they do not fit cleanly into the traditional fields of statistics (having nothing to do with modeling, estimation, or inference) nor computer science (software engineering?). We submit that these are early, influential works of scholarship in data science.

Another set of exemplary papers can be found in a recently-published collection of articles—curated by Jenny Bryan and Hadley Wickham—entitled *Practical Data Science for Stats* (to which the authors all contributed) [11]. These articles discuss meta-data science topics like how to package reproducible analytical work [19], how to organize data in a spreadsheet [9], how to share data for collaboration [16], and how to implement a version control system [10]. Our contributions discussed surviving as an isolated data scientist [5], and wrangling categorical data [20].

The collection also contains an article on evaluating scholarly work in data science, focusing particularly on data science faculty in traditional statistics and biostatistics departments [24]. Can these exemplary scholarly contributions in data science be neatly categorized into statistics or computer science research? If not, this further strengthens the notion that data science exists as a field of research unto itself.

Situating greater data science

This brings us to our final question. If Donoho’s vision of ‘Greater Data Science’ takes hold, one wonders whether the current academic departmental alignments will (or should) continue. Of the authors, one is situated within a Department of Mathematics and Statistics (Horton), while the other two are appointed in a Program in Statistical and Data Sciences. Which approach is most fruitful?

Clearly, there are many other academic areas that use data and data science methods. As we’ve discussed, our colleagues across the disciplines are embracing it. However, if data science is its own discipline, it cannot be solely situated within data-generating departments. Its unique teaching and scholarship indicate it may need to become a separate entity.

References

- [1] American Statistical Association (2014). *2014 Curriculum Guidelines for Undergraduate Programs in Statistical Science*. <http://www.amstat.org/asa/education/Curriculum-Guidelines-for-Undergraduate-Programs-in-Stat>
- [2] American Statistical Association (2015). A peek into the largest, fastest-growing undergraduate statistics departments. http://magazine.amstat.org/blog/2015/02/01/undergraduatedepts_feb2015.
- [3] Ball, R. and Medeiros, N. (2012). Teaching integrity in empirical research: A protocol for documenting data management and analysis. *The Journal of Economic Education*, 43(2):182–189.
- [4] Baumer, B. (2015). A data science course for undergraduates: Thinking with data. *The American Statistician*, 69(4):334–342.
- [5] Baumer, B. (2017). Lessons from between the white lines for isolated data scientists. Technical report, PeerJ Preprints.
- [6] Baumer, B., Çetinkaya Rundel, M., Bray, A., Loi, L., and Horton, N. J. (2014). R Markdown: Integrating a reproducible analysis tool into introductory statistics. *Technology Innovations in Statistics Education*, 8(1).
- [7] Baumer, B., Kaplan, D. T., and Horton, N. J. (2017). *Modern data science with R*. Chapman & Hall/CRC.
- [8] boyd, d. and Crawford, K. (2012). Critical questions for big data. *Information, Communication & Society*, 15(5):662–679.
- [9] Broman, K. W. and Woo, K. H. (2017). Data organization in spreadsheets. *PeerJ Preprints*, 5:e3183v1.
- [10] Bryan, J. (2017). Excuse me, do you have a moment to talk about version control? *PeerJ Preprints*, 5:e3159v2.
- [11] Bryan, J. and Wickham, H., editors (2017). *Practical Data Science for Stats*. PeerJ.

- [12] Buckheit, J. B. and Donoho, D. L. (1995). Wavelab and reproducible research. Technical Report 474, Stanford University. <http://statistics.stanford.edu/~ckirby/techreports/NSF/EFS%20NSF%20474.pdf>.
- [13] Carver, R. et al. (2016). *Guidelines for Assessment and Instruction in Statistics Education: College Report 2016*. American Statistical Association.
- [14] Cleveland, W. S., McGill, R., et al. (1985). Graphical perception and graphical methods for analyzing scientific data.
- [15] De Veaux, R. D. et al. (2017). Curriculum guidelines for undergraduate programs in data science. *Annual Review of Statistics and Its Application*, 4(1):1–16.
- [16] Ellis, S. E. and Leek, J. T. (2017). How to share data for collaboration. *PeerJ Preprints*, 5:e3139v5.
- [17] Knuth, D. (1992). Literate programming. *CSLI Lecture Notes, Stanford University*, 27.
- [18] Leek, J. (2016). Data science as a science. In *Joint Statistical Meetings*.
- [19] Marwick, B., Boettiger, C., and Mullen, L. (2017). Packaging data analytical work reproducibly using R (and friends). *PeerJ Preprints*, 5:e3192v1.
- [20] McNamara, A. and Horton, N. J. (2017). Wrangling categorical data in R. *PeerJ Preprints*, 5:e3163v1.
- [21] Pruum, R., Kaplan, D. T., and Horton, N. J. (2017). The mosaic Package: Helping Students to ‘Think with Data’ Using R. *The R Journal*, 9(1):77–102.
- [22] Silberzahn, R. et al. (2017). Many analysts, one dataset: Making transparent how variations in analytical choices affect results.
- [23] UC Berkeley School of Information (2017). Master of information and data science: Curriculum. <https://datascience.berkeley.edu/academics/curriculum/>.
- [24] Waller, L. A. (2017). Documenting and evaluating data science contributions in academic promotion in departments of statistics and biostatistics. *PeerJ Preprints*, 5:e3204v1.
- [25] Wickham, H. (2009). *ggplot2: Elegant Graphics for Data Analysis*. Springer Verlag: New York, NY.
- [26] Wickham, H. (2014). Tidy data. *The Journal of Statistical Software*, 59(10). <http://vita.had.co.nz/papers/tidy-data.html>.
- [27] Wickham, H. and Francois, R. (2016). *dplyr: a grammar of data manipulation*. R package version 0.5.0.9000.
- [28] Wilkinson, L. (2005). *The Grammar of Graphics*. Statistics and computing. Springer Science + Business Media.