

The democratization of data science education

Sean Kross* Roger D. Peng* Brian S. Caffo* Ira Gooding** Jeffrey T. Leek*
kross@jhu.edu rdpeng@gmail.com bcaffo@gmail.com iragooding@jhu.edu jtleek@gmail.com

* Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health

** Center for Teaching and Learning, Johns Hopkins Bloomberg School of Public Health

Abstract

Over the last three decades data has become ubiquitous and cheap. This transition has accelerated over the last five years and training in statistics, machine learning, and data analysis have struggled to keep up. In April 2014 we launched a program of nine courses, the Johns Hopkins Data Science Specialization, which has now had more than 4 million enrollments over the past three years. Here the program is described and compared to both standard and more recently developed data science curricula. We show that novel pedagogical and administrative decisions introduced in our program are now standard in online data science programs. The impact of the Data Science Specialization on data science education in the US is also discussed. Finally we conclude with some thoughts about the future of data science education in a data democratized world.

What is the Johns Hopkins Data Science Specialization (timeline)

Data has become dramatically cheaper to collect and store over the last several decades. This phenomenon has touched fields ranging from the government, to social sciences, to molecular biology (1). As data have become commoditized and democratized, there has been a dramatically increased demand to analyze and extract information from these data. The demand for data analytic training was made clear when three Stanford computer science professors -- Daphne Koller, Andrew Ng and Sebastian Thrun -- opened their Stanford courses focusing on data analysis for free to the world and enrolled hundreds of thousands of students (2).

Recognizing the demand for these classes -- later labeled massive online open courses (MOOCs) - these three professors launched companies, Coursera (3) and Udacity(4), to offer similar courses in partnership with professors at other universities. These companies, and the nonprofit EdX(5) built on earlier efforts, such as OpenCourseware(6), for delivering top level content to a broad audience.

As biostatistics professors focusing on the development and in particular the application of data analytic methods to solve problems in biomedicine, we were also enthusiastic and equipped to teach this type of class. The effort that culminated in the Johns Hopkins Data Science Specialization occurred over a period of approximately 1.5 years. It was the work of three faculty at the Johns Hopkins University Department of Biostatistics with the assistance of key staff members who will be mentioned during the description of the program.

Our online educational efforts were spurred by a partnership that was formed between Johns Hopkins and Coursera (5a). Brian Caffo, Roger Peng, and Jeff Leek elected to teach three classes on Coursera: *Mathematical Biostatistics Boot Camp*, *R programming*, and *Data Analysis*. These three classes reflected course materials that were developed for students at Johns Hopkins in mathematical statistics, R programming, and the art of the analysis of data. In development, the new MOOCs were direct translations of the associated courses to this new format. *Mathematical Biostatistics Boot Camp* and *Computing for Data Analysis* launched in September 2012 and *Data Analysis* launched in January 2013.

The lectures were recorded and distributed as videos. All three courses included multiple choice assignments. Two of the courses included unique assessment features. The R programming class included programming exercises where code output was automatically evaluated. The Data Analysis class included a final project that was peer assessed. Each submitted project was randomly assigned to a small number of other students, who graded the project using a fixed rubric (7). Students then received the median score from these peer assessments on their project. These assessment tools, and related ones based on automation and crowdsourced grading, are common in MOOC platforms to address scalability of project based learning. They played a significant role in the development of the Johns Hopkins Data Science Specialization.

Propelled by the popularity of early MOOCs (8), the enrollment in these first three classes was extraordinary (Table 1) and was an early indication of the demand for data science focused courses. This demand has since been validated by the popularity of not only online courses in data science, but a host of boot camps (9), internships, and a dramatic increase in enrollment across statistics, machine learning, and data science undergraduate programs (10).

Course	Enrollment	Completions
Mathematical Biostatistics Boot Camp 1	109,789	4,150
Mathematical Biostatistics Boot Camp 2	23,842	944
Computing for Data Analysis	243,987	21,069
Data Analysis	193,126	6,500
Biostatistics Case Study	39,140	3,322

Table 1. Number of Students in Initial Online Courses

The size of the student enrollment and completion numbers from the initial courses we created for Coursera.org strongly encouraged our development of the Data Science Specialization.

Seeing the demand for data science spiking, we were inspired to create an entire data science curriculum that would deviate from both the in person program at Johns Hopkins and other statistics and biostatistics programs nationally. This coincided with an effort by Coursera to create sequences of MOOCs or “Specializations”. The launch date for the original set of specializations was April 2014. Using our lightweight production process we created a series of nine courses and launched one of the first Specializations on the Coursera platform. The classes in our program were:

1. **The Data Scientist’s Toolbox** (11) - covering version control and the most common tools used by data scientists.
2. **R programming** (12) - covering the fundamentals of how to program in R
3. **Getting and cleaning data** (13) - covering how to obtain data from databases, the web, and other sources, then clean it up.
4. **Exploratory data analysis** (14) - covering plotting and other initial explorations of a data set.
5. **Reproducible research** (15) - covers the basics of R markdown, literate programming, and the principles of reproducible research
6. **Statistical inference** (16) - covers probability and statistical inference.
7. **Regression models** (17) - covers linear models and their application to real data.
8. **Practical machine learning** (18) - covers the basics of machine learning and predicting in R.

9. Developing data products (19) - covers technologies for building data products like R packages, swirl courses, and Shiny apps.

This program has now been running since April 2014 and has enrolled more than 4 million students and over 875,000 people completing a course. It is the largest MOOC program in history.

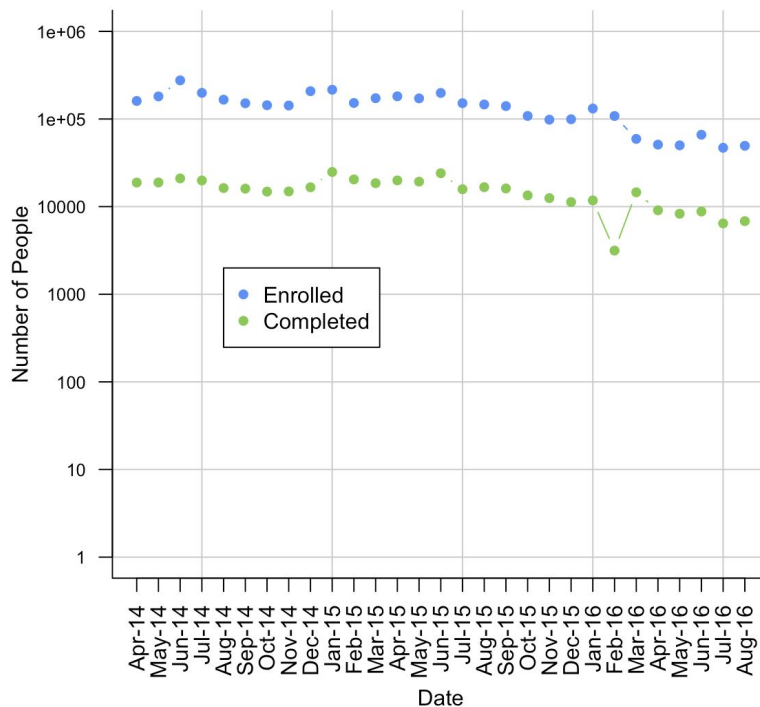


Figure 1. Number of students in the Johns Hopkins Data Science Programs

The number of students enrolling (blue) and completing (green) courses in MOOCs created by the Johns Hopkins Biostatistics Department over time. Coursera launched their new platform in February 2016.

Over time enrollments have appeared to reach a steady state (Figure1). A large number of students continue to enroll in the program, now one of the most widely recognized sequences of MOOCs. Despite the well-known criticism that MOOCs have low completion rates(20), the number of completers of the *JHU Data Science Specialization* - at the time of this writing (n = 6,235 completers) - is almost double the entire output of masters students from the ten largest programs in statistics and the ten largest programs in biostatistics combined (n = 2,056 degrees) over the years 2011-2013 (21).

Part of the reason for this success was simply the serendipity of launching the program at the right time. But some of the innovations produced by the program made it unique and have now led other groups to build on those ideas. In the following sections we summarize the differences

we think gave us an advantage and propose some ideas for the future of data science education.

Philosophical principles behind our innovations

One of the key reasons for the success of the *Johns Hopkins Data Science Specialization* is our focus on a few key principles:

1. focusing on developing large quantities of quality open source content,
2. targeting broad audiences as opposed to niche experts,
3. thinking about the users of online education.

From the conception of the program, content was focused on our hypothesized motivations of online data science students. This led to several strategic decisions - for example to reduce production time to increase the speed of generating content. While the resulting reduction in production value could be seen as a negative, it was turned positive by developing open source, modular, and public content that could be updated quickly. A second strategic decision was to orient programs for introductory and general audiences, rather than focusing on more technical material only accessible to experts with many prerequisites. This is part of a key general strategic decision to focus on the typical needs of online education users. For example, the modal MOOC students take courses in their spare time and arrive at the material with little background (22) (23). Such learners would not be interested in courses and programs that have typical university requirements that a full time student would be able to meet.

Based on these motivating principles we adapted our platform, content, structure and community to maximize our reach. We now detail several choices we made in each of these categories.

Innovations in Platform

Leanpub

The high cost of textbooks with frequent edition updates is one of the biggest justifiable complaints from students (24). From the beginning we had released all course material under an open source license on Github. However, we needed a mechanism for releasing affordable books and course notes. We used the Leanpub (25) platform, which offered a variable pricing model with no minimum payment required. Textbooks for several of our courses including *Statistical Inference for Data Science*, *The Elements of Data Analytic Style* (26), and *R Programming for Data Science* (27) have sold collectively hundreds of thousands of free and paid copies on the platform. The self-publishing model and markdown based writing format made it easy to translate notes written for courses into textbooks and rapidly release both at the same time. Since then, Brian Caffo, Roger Peng, and Jeff Leek have published a collective 12 books on Leanpub, all available for free on a variable pricing model.

Capstones

After developing 9 courses for the data science specialization, all that was left was to build a capstone project in which learners could apply all of the skills that they had learned through the series in a larger project. The first capstone project was a collaboration between Johns Hopkins University and SwiftKey, a company that made a predictive keyboard app for Android phones. Partnership with industry brought on several political and legal challenges but provided the opportunity for students to work on a project that was guided by experts in the field.

The basic idea behind the capstone project was that students would have to take a large corpus of text data and build a predictive model for text, essentially a streamlined version of a task SwiftKey would have to do at production scale. Students had to build a web application using the Shiny (28) framework for R. Shiny allows students to rapidly prototype and share web applications which are written purely in R code. We instructed students to build a web application where a user could enter a string of text into a textbox, and then the user would see the next predicted word. In order to build the app students had to parse and clean a corpus of text data. In order to make an algorithm students created n-gram models using any method of their choosing, though they were encouraged to consider the runtime and the amount of memory that their algorithm would require. Ultimately, they had to build a Shiny app that demonstrated the algorithm to others. Each Shiny app was hosted on RStudio's shinyapps.io platform(29), and RStudio was generous to provide limited-time free hosting for all learners in the capstone. Students graded each other's assignments via Coursera's peer grading system. The rubric that students used to perform the peer assessment includes criteria for their app responding quickly and being able to process input from new articles and social media. SwiftKey's engineers helped to develop the project and were gracious with their time by offering to a live Google Hangout "office hours" where they answered questions from the learners in the capstone course.

Since our original capstone project, we've developed a second capstone with Yelp, using a similar process. Now, the capstone projects rotate between the two, giving learners an opportunity to choose what type of project they wish to complete. Throughout the development of capstone projects we focused on making the problems real, as ambiguous as possible so that students would have to struggle with the same problems as practicing data scientists, and interesting so students would be motivated to complete them.

Swirl

One of the challenges of providing personalized activities for large numbers of students is the lack of interaction with an instructor. One approach we took to solving this problem was to develop the *swirl* platform (30). The *swirl* R package walks students through a linear interactive programming tutorial in the R console. Students progress through *swirl* lessons by typing R command into the console, answering multiple choice questions, or by writing their own R

scripts. If a student's answer is correct according a set of answer tests provided by a *swirl* lesson author, then the student is allowed to move on to the next exercise. If the student's answer is incorrect then *swirl* offers them a hint. Lessons for *swirl* can be written by anyone by combining a yaml structure for each lesson that specifies the order in which questions are asked, the question type, and the correct answer for a question.

swirl was built so student could learn programming, R and statistics in R ("learn R in R" is the catchphrase for *swirl*). It was developed by Nick Carchedi, Brian Caffo, Lauren Williams, Ethan Schwartz, Gina Grdina, Sean Kross and Bill Bauer.

A key step was to incorporate *swirl* into Coursera via open education APIs. *swirl* has since become a main component of the program and one of the most popular R packages with over 600,000 downloads and it has garnered a community that has translated the package into eight languages.

Innovations in Content

There are several elements in our data science program that are not formally addressed as a part of a traditional statistics curriculum. The first course in our program, called The Data Scientist's Toolbox, aims to set up the computational environment for students who are new to scientific computing. In this course we address material that is usually left to students to figure out on their own or in office hours including how to configure R and RStudio, the command line, Git, and GitHub, all across different operating systems. Consider the rate at which new data science technologies are being released, we believe that teaching students to set up their own programming environment is an essential skill to have as a data scientist.

Institution	Program	Inference	Modeling	Programming	Data Products	Data Cleaning	Reproducibility	Exploratory Analysis
Stanford (31)	MS Statistics	Introduction to Statistical Inference	Regression Models and Analysis of Variance	Programming Methodology	NA	NA	NA	NA
CMU (32)	Master's of Statistical Practice	Advanced Methods for Data Analysis	Applied Linear Models	Statistical Computing	Statistical Practice	NA	NA	NA
NYU (33)	MS in Applied Statistics	Applied Statistical Modeling and Inference	Applied Statistical Modeling and Inference	Statistical Computing	NA	NA	NA	NA
Columbia (34)	MA in Statistics	Multivariate Statistical Inference	Regression and Multi-Level Models	Statistical Computing and Intro to Data Science	NA	NA	NA	Topics in Modern Statistics – Statistical

								Graphics
Harvard (35)	AM in Statistics	Statistical Inference	Linear and Generalized Linear Models	Statistical Computing	NA	NA	NA	NA
Illinois (36)	MS in Statistics	Statistical Analysis	Applied Regression and Design	Statistical Computing	NA	NA	NA	NA
Georgia Tech (37)	MS in Statistics	Math Statistics I	Regression Analysis	Computational Statistics	NA	NA	NA	NA
Indiana (38)	MS in Applied Statistics	Introduction to Statistical Theory	Applied Linear Models	Statistical Computing	NA	NA	Managing Statistical Research	Exploratory Data Analysis
Johns Hopkins	Data Science Specialization	Statistical Inference	Linear Models	R Programming	Developing Data Products	Getting and Cleaning Data	Reproducible Research	Exploratory Data Analysis
UBC (39)	Master of Data Science	Statistical Inference and Computational I	Regression I	Programming for Data Science	Capstone Project	Data Wrangling	Data Science Workflows	Data Visualization I

Table 2. Comparison of Traditional Statistics Programs to the Data Science Specialization

The role of statistical practices is integral to data science, however the scope of data science programs extends beyond many of the courses offerings by traditional statistics programs. We include the Data Science Specialization for comparison in addition to the Master of Data Science program from the University of British Columbia.

Raw data is often not easily parsed or interpreted by computer programs, and preparing raw data for analysis, modeling, and visualization is often the job of a data scientist. We created a course called Getting and Cleaning Data in order to emphasize the importance of these data preparation skills. Data preparation involves two general steps: first the data must be imported into a programming environment where it can be manipulated, and then the data is transformed into a structure that eases its later use. The importance of this data janitorial work often neglected by traditional programs despite the fact that it often occupies a significant amount of the time spent working with data (40).

The importance of ensuring that one's research is reproducible has entered the scientific mainstream in recent years (41), however the topic is often presented as yet another infrastructure burden that is expected of a data scientist. We designed our Reproducible Research course through a different philosophy by emphasizing how tools like Markdown and knitr can aid a data scientist while developing an analysis. This topic has been increasing in popularity and is now fundamental to many data analytic training programs (42).

The results of a data analysis deserve to be communicated as widely and coherently as possible, and this belief motivated the creation of our Developing Data Products course. A large proportion of this course is dedicated to teaching Shiny, a web application framework for R. Shiny allows students to create interactive web applications simply by writing R code, which can then be shared online. This allows our students to share models and data visualizations that app users can use without any knowledge of R. We also teach students how to create their own stand-alone websites with R Markdown, featuring interactive graphics and maps powered by the plotly and leaflet libraries for R. We believe that the creation of data products is integral to the skill set of any data scientist.

While many of these skills - programming, getting and cleaning data, reproducible research, and data products - were covered very well in individual classes at different universities or in one off programs (43), they were not integrated into core curricula in most statistics, biostatistics, and machine learning programs. One advantage we had in distributing our content via Coursera was the ability to avoid vetting our program through a curriculum committee, which allowed us to quickly insert content we believed was important but was not considered “core” to any specific discipline.

Innovations in Structure

Applying the lessons learned from their first three courses proved prescient. Several choices in the structure of the program have contributed to the popularity of the Johns Hopkins Data Science Specialization.

- **Timing** - the courses were created on an expedited time scale so the JHU DSS was one of the only programs available at the time of its release.
- **Cost** - we insisted the course certificate be priced at \$50 for a total program cost of \$500 including the capstone course.
- **Frequency** - we decided to run every course in their program every month. This was in contrast to previous MOOCs that ran once or at irregular intervals modeled after university semesters.
- **Length** - we designed their courses to be one month long with a 2 week project involved in each course. This fit into the schedule of career changers and people doing the courses part time.
- **Integration** - the sequence was not cobbled together from existing courses, but designed to follow the data science pipeline from beginning to end.
- **Portfolios** - each class had a project that was pushed to Github and peer graded, resulting in a portfolio of work for each student completer.

Several of these innovations have now been adopted as standard practice at Coursera and other MOOC platforms (44). For example, all classes are now being offered “on demand” rather than as one off or irregularly scheduled courses. At the time, this approach was strongly

discouraged because of the labor required to migrate content and set up nine new courses sessions on a monthly basis. Coursera's new platform and course format makes high-frequency sessions the default across its offerings. All classes on Coursera are also organized in one month units based on the design principles of the Johns Hopkins Data Science Specialization.

Running courses frequently was greatly appreciated by students and resulted in no drops in quality of student experience and enrollment. Another lesson was that making content more atomic across the board was preferable in this format. In other words, whether for readings, lectures, modules or classes, splitting was always better than lumping. Atomizing content created flexibility in terms of organizing the course, so that if we felt that the order of course modules should be changed partway through developing a course, we could easily rearrange the content.

Innovations in Community

One of the most inspiring parts of this work has been the opportunity to witness the development of the community that has grown up around the specialization. What began with the free-for-all discussion forums on the course sites quickly spilled over into other online venues like blogs and social media. Thanks to Coursera's efforts to harness this dynamic community's eagerness to both give back and deepen their own knowledge by helping others, we've had the pleasure of recruiting and relying on some truly incredible Community Mentors (originally called Community TAs) who volunteer their time to moderate the forums and guide new learners. We also learned that many of these mentors and other learners were starting to produce help documents, lecture notes, study guides, and other materials to supplement the specialization and assist new learners. Some people worried that this somehow constituted cheating, but we saw it as a great way for the community to help guide one another and fill in some of the gaps that we missed when developing the courses. In order to encourage these activities, we created the Data Science Specialization Community Site on GitHub [github](#) (21a) and asked learners to contribute their work so that others could easily find and use the materials.

Future of data science education

Since the development of our program there has been an explosion of data science programs around the country (45). The ubiquity of data and the training gap in students and employees who can analyze this data means that data science education is likely to grow over the short term. But given the similarity between these programs we believe that data science education is becoming commoditized. The real value is no longer in generation of content - which is now abundant and available for free on the internet. In the future value will be driven by interactive education platforms such as *swirl* and *Datacamp*. It is also likely to be driven by the associated value of being around experts in the field who can help to improve the data literacy of students when they run into problems produced by rapidly changing open source software.

Most importantly the large majority of traditional data science programs target one major audience - technically minded undergraduate or early graduate students. But there is a major opportunity to develop data literacy beyond these groups - both extending earlier in the curriculum and reaching across different disciplines. We imagine that much of the growth in data science education will focus on these populations in the coming years. Regardless of the platforms, content, or outcomes, the democratization of data demands a democratization of data science education that started with the Johns Hopkins Data Science Specialization.

References

1. G. O'Connor, Moore's Law Gives Way to Bezos's Law. *GigaOm*, April. **19** (2014).
2. T. Lewin, MOOCs, Large Courses Open to All, Topple Campus Walls. *The New York Times* (2012), (available at <http://www.nytimes.com/2012/03/05/education/moocs-large-courses-open-to-all-topple-campus-walls.html>).
3. Coursera - Free Online Courses From Top Universities. *Coursera*, (available at <https://www.coursera.org/>).
4. Udacity - Free Online Classes & Nanodegrees, (available at <https://www.udacity.com/>).
5. edX. *edX*, (available at <https://www.edx.org/>).
6. M. OpenCourseWare, MIT OpenCourseWare | Free Online Course Materials, (available at <http://ocw.mit.edu/index.htm>).
7. jtleeek, jtleeek/dataanalysis. *GitHub*, (available at <https://github.com/jtleeek/dataanalysis>).
8. K. Jordan, Initial trends in enrolment and completion of massive open online courses. *The International Review of Research in Open and Distributed Learning*. **15** (2014), doi:10.19173/irrodl.v15i1.1651.
9. 2017 Coding Bootcamp Market Size Study, (available at <https://www.coursereport.com/reports/2017-coding-bootcamp-market-size-research>).
10. M. Guzdial, "Generation CS" drives growth in enrollments. *Commun. ACM*. **60**, 10–11 (2017).
11. Coursera - Free Online Courses From Top Universities. *Coursera*, (available at <https://www.coursera.org/learn/data-scientists-tools>).
12. Coursera - Free Online Courses From Top Universities. *Coursera*, (available at

- <https://www.coursera.org/learn/r-programming>).
13. Coursera - Free Online Courses From Top Universities. *Coursera*, (available at <https://www.coursera.org/learn/data-cleaning>).
 14. Coursera - Free Online Courses From Top Universities. *Coursera*, (available at <https://www.coursera.org/learn/exploratory-data-analysis>).
 15. Coursera - Free Online Courses From Top Universities. *Coursera*, (available at <https://www.coursera.org/learn/reproducible-research>).
 16. Coursera - Free Online Courses From Top Universities. *Coursera*, (available at <https://www.coursera.org/learn/statistical-inference>).
 17. Coursera - Free Online Courses From Top Universities. *Coursera*, (available at <https://www.coursera.org/learn/regression-models>).
 18. Coursera - Free Online Courses From Top Universities. *Coursera*, (available at <https://www.coursera.org/learn/practical-machine-learning>).
 19. Coursera - Free Online Courses From Top Universities. *Coursera*, (available at <https://www.coursera.org/learn/data-products>).
 20. D. Clow, in *Proceedings of the Third International Conference on Learning Analytics and Knowledge* (ACM, New York, NY, USA, 2013), *LAK '13*, pp. 185–189.
 21. L. Higher Logic, Largest U.S. Master's Programs in Statistics and Biostatistics - amstat, (available at <http://community.amstat.org/blogs/steve-pierson/2014/02/09/largest-graduate-programs-in-statistics>).
 22. S. Zheng, M. B. Rosson, P. C. Shih, J. M. Carroll, in *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (ACM, New York, NY, USA, 2015), *CSCW '15*, pp. 1882–1895.
 23. T. Dillahunt, B. Chen, S. Teasley, in *Proceedings of the First ACM Conference on Learning @ Scale Conference* (ACM, New York, NY, USA, 2014), *L@S '14*, pp. 145–146.
 24. N. Nicholls, *The Investigation into the Rising Cost of Textbooks* (2009).
 25. Leanpub: Publish Early, Publish Often, (available at <https://leanpub.com/>).
 26. J. Leek, *The Elements of Data Analytic Style*. *Leanpub*, (available at <https://leanpub.com/datastyle>).
 27. R. D. Peng, *R Programming for Data Science*. *Leanpub*, (available at <https://leanpub.com/rprogramming>).
 28. W. Chang *et al.*, shiny: Web Application Framework for R, 2015. URL <http://CRAN.R-project.org/package=shiny>. *R package version 0.11* (2015).

29. shinyapps.io, (available at <http://www.shinyapps.io/>).
30. N. Carchedi, B. Bauer, G. Grdina, S. Kross, Swirl: Learn R, in R (2014).
31. Stanford Statistics M.S. Required Courses, (available at <https://statistics.stanford.edu/academics/ms-statistics-required-courses>).
32. CMU MS Statistics Curriculum, (available at <http://www.stat.cmu.edu/msp/curriculum>).
33. NYU MS in Applied Statistics, (available at http://steinhardt.nyu.edu/humsocsci/applied_statistics/sampleschedules).
34. Columbia University Masters of Arts in Statistics, (available at <http://stat.columbia.edu/ma-programs/approved-electives/>).
35. Harvard Graduate Statistics General Info, (available at <https://statistics.fas.harvard.edu/pages/graduate-statistics-general-information>).
36. Illinois Statistics Classes and Course Catalog, (available at <http://www.stat.illinois.edu/students/courses.shtml>).
37. Georgia Tech MS Statistics Curriculum, (available at <https://www.isye.gatech.edu/academics/masters/ms-statistics/curriculum>).
38. Indiana University Graduate Statistics Courses, (available at <https://stat.indiana.edu/graduates/courses.html>).
39. University of British Columbia Master of Data Science Courses, (available at <https://ubc-mds.github.io/descriptions/>).
40. R. M. John King, 2016 Data Science Salary Survey, (available at <http://www.oreilly.com/data/free/files/2016-data-science-salary-survey.pdf>).
41. V. Stodden, F. Leisch, R. D. Peng, *Implementing Reproducible Research* (CRC Press, 2014).
42. T. K. Teal *et al.*, Data Carpentry: Workshops to Increase Data Literacy for Researchers. *International Journal of Digital Curation*. **10**, 135–143 (2015).
43. J. Bryan, Stat 545 (2015), (available at <http://stat545.com/>).
44. Coursera, Coursera Update: Striking a Balance with Start Dates and Deadlines. *Coursera Blog* (2015), (available at <https://blog.coursera.org/coursera-update-striking-a-balance-with-start/>).
45. E. Tate, Data Analytics Programs Take Off. *Inside Higher Ed* (2017), (available at <https://www.insidehighered.com/digital-learning/article/2017/03/15/data-analytics-programs-taking-colleges>).
- 5a. I. Gooding, et al. Massive open online courses in public health. *Frontiers in Public Health Education and Promotion Front. Public Health*, 25 November 2013

<http://dx.doi.org/10.3389/fpubh.2013.00059>

21a. Data Science Specialization Community Site (available at <https://datasciencespecialization.github.io/>)