BUILD YOUR OWN STATISTICS COURSE FOR STUDENTS IN A NON-QUANTITATIVE FIELD[1]

Alexis Lerner[2] and Andrew Gelman[3]

July 13, 2020

Abstract

Statistics courses are typically in mathematics or statistics departments or in social and natural sciences such as economics, political science, psychology, and biology. Here we discuss how to construct a statistics course for students in non-quantitative fields, with a goal of integrating the statistical material with students' substantive interests, using modern teaching methods to increase student involvement. We demonstrate with the example of an introductory applied statistics class at the University of Toronto's Centre for Jewish Studies.

Keywords

Quantitative Methods, Pedagogy, Undergraduates, Course Design

---

[1] We thank Yufan Yang and Tomoko Takahashi for their helpful feedback.
[2] Department of Political Science and Centre for Jewish Studies, University of Toronto, and Jordan Center for the Advanced Study of Russia, New York University, Alexis.Lerner@mail.utoronto.ca
[3] Department of Statistics, Columbia University, New York, gelman@stat.columbia.edu

Introduction

Non-mathematically-minded undergraduates express hesitancy about enrolling in statistics courses (Slootmaeckers et. al., 2013; Bradstreet, 1993; Rumsey, 2017). Some students have inadequate quantitative training, others lack numerate confidence, and still others fail to see the relevance of statistics to them. And yet, one does not need to be a statistician to depend on numbers in everyday life. Journalists, non-profit leaders, and teachers use quantitative tools and statistical methods to measure trends, identify patterns, and interpret results. And news outlets habitually use graphs and charts to illustrate information for their readers. So how can educators best help students to build a numerate life?

In this article, we present a Build Your Own Statistics Course template following Gelman (2019) which can be modified to appeal not only to students in the social sciences and humanities, but also to students affiliated with theological institutions, centers for ethnic or diaspora studies, and journalism schools. Using our modified template, educators can accomplish two separate and complementary goals:

1. Teach introductory statistics to non-statistics students
2. Teach substantive content through the perspective of statistics.

Template

In the following section, we outline the major components of this course. We suggest that the objective of such a course is some combination of a critical ability to interpret data and secondary analysis skills (e.g., basic regression, data illustration). Follow this template and you will have a course.

What materials do you need?

Students should have access to computers, loaded with basic computing and visualization software, such as Tableau and R. Particular exercises may require further materials; for example, when we teach students how to compute and understand error terms, we use 10-15 photographs of individuals for the age-guessing demonstration of Gelman and Nolan (2017). Instructors may also find it useful to have access to an overhead projector, in order to provide visuals to students. For collaborative assignments, we recommend a cloud-based platform that students can edit in real time (e.g., Google Drive, Dropbox, Microsoft One Drive).

Evaluation: How should you assess your students?

If the goal is getting students to learn, instructors need tools for evaluation. We suggest that course instructors employ a variety of summative (how much someone has learned) and formative (how someone is learning) assessment tools. They are described as follows:

Pretest and Posttest

Pretest/posttest followup is a useful pedagogical tool for student learning because it is both a summative and formative evaluation model. First, the pretest can provide early feedback to the instructor, as it shows where the students have substantial pre-existing knowledge and where they are limited. A pretest can help an instructor craft their semester, determining whether more or less introductory material is necessary. Further, a posttest can operate both as a final exam, when it helps to leverage final judgment on student learning and retention, as well as a formative tool for showing and illustrating student growth, and therefore for reinforcing future student learning in the realm of statistics.

For students, the pretest/posttest model is useful because it is low-stakes. For example, one can grade the pretest on a pass-fail basis, with the grade accounting for approximately five percent of their overall grade. This means that, on day one, each student has an A+ in the course, which alleviated some of their anxieties about taking a statistics class. This early confidence seems to help students to be more willing to take chances, make mistakes, and engage with course material. On the second-to-final day of class, the posttest should be administered. We suggest that this be graded and account for a larger percentage of students' overall grade. Every question on the posttest should also appear on the pretest and must be addressed over the course of the semester through course materials, lectures, and assignments. A posttest can be administered on the second-to-last day of class so that, on the last day of class, an instructor can redact student names and go over the results of the posttest collectively. In our experience, students enjoy this process because they are able to see where they did well and, if they got a question wrong, that others got it wrong, too.

We suggest that students be required to include for each question a self-reported confidence level (on a range of 0-10, where 0 is low and 10 is high) for their answer. This confidence level allows students to signal to their instructor when—and to what degree—they were guessing, in order to account for lucky guesses and to assess overconfidence. This model of assessment also helps to measure student growth and skill acquisition without relying on post-course student evaluations, which have been proven to be biased in regard to female and minority educators.

Just-in-time-teaching (jitt) tasks

In *just-in-time-teaching* (jitt), the course is supplemented by a short online assignment to be done before each class period (Novak et al., 1999, Simkins and Maier, 2010, Watkins and Mazur, 2010). For our jitts, we typically include three short questions, one to check on the readings, one that is a short problem, and one feedback on the class. The jitt is intended to take about 15 minutes, and students are graded not on their correctness but just for seriously attempting it. We expect that students often will do their jitts right before class, and this will get them in the mood for the class period. Gelman (2013) provides some discussion and examples of how we implement jitts in our class.

These short ungraded online tasks are great pedagogical tools for getting students to think critically about course materials prior to class and to provide feedback to the instructor. Both authors of this paper use jitts in all of our classes by opening student answers on the screen and going through the answers together as a class. The answers appear to the students as anonymous, which allows students to see what a correct answer looks like without being publicly shamed for an incorrect answer. This benefits students by boosting their confidence, not only because they are able to see how they did in relation to their peers, but also because they are able to fail without severe penalty. Jitt answers are discussed anonymously in class, but names are attached to answers for the instructor. This allows the instructor to grade jitts for effort on a pass-fail basis. Populating the syllabus with frequent, small assignments—such as these jitts—can help motivated students to earn higher grades; in a statistics class like this one, these frequent, small assignments help students feel like their hard work is, to some degree, paying off. It is also possible to fully redact names and identifying information (and instructor feedback) from the jitts, and to share the answers with students as a study tool.

Papers and Projects

A combination of solo, partnered, and collaborative projects is a good way to promote student learning. After all, if a student can solve a homework assignment correctly on their own, they demonstrate proficiency. However, a student that can *explain their thought-process to a peer* demonstrates mastery while improving their communication skills. This is the principle behind *peer instruction* (Mazur, 1997, Crouch and Mazur, 2001). Collaborative assignments are particularly interesting because of their enmeshed free rider problem, where some group members tend to do more work than others, but the group earns a uniform grade, thereby allowing weak group members to free ride on the coattails of stronger students.

In a statistics course, free riding in collaborative projects operates in a somewhat different manner. For example, say you are teaching a lesson about regressions. In order to increase student interest, it is helpful to involve them in every step of the process. Therefore, instead of using a premade dataset, have students create a codebook together and decide on how different observable traits will be operationalized (Boger 2001). Students can collect this data quite easily using participant observation or survey questions. If you have a class of 15 students and each student surveys five strangers on campus, you have 75 observations. You can use the jitt to request that students approve you to use their data in subsequent iterations of the course, which makes those 75 observations into 150, and so on.

After building their collaborative dataset, students can practice uploading it into R (or whichever computing program you prefer), cleaning it, and using it to run models or to illustrate exploratory findings. You can even use this process to teach library science fundamentals regarding the organization and preservation of data for optimal readability across researchers. You can conclude this collaborative assignment with a short solo-authored paper on findings, takeaways, and illustrations. This kind of collaborative assignment solves the free rider problem, because a free rider would produce missing data, which negatively impacts the results for every single person in the class, including the free rider.

What textbooks and readings should you use?

We suggest organizing your syllabus by weekly statistical objectives. These should begin with basic theoretical concepts that answer the general questions of, 'what is data?' and 'who is involved in constructing a dataset?' To answer questions such as these, you should include not only texts that explain the concept from a statistical perspective (whether something written for the non-STEM reader such as Wheelan (2012) or something more traditional like Gelman, Hill, and Vehtari (2020) but also texts that consider data from a post-positivist perspective (e.g., including Fujii (2010) on the topic of meta-data). Supplementary readings can cover specific topics on research design (e.g., case selection, avoiding bias, and/or regression versus classification problems) and statistical concepts (such as means comparisons with t-tests or *how* to perform regression and/or classification analyses).

Most importantly, a conversation about what makes data *data* must include, well, data. Therefore, we recommend that you include a weekly or biweekly dataset in your syllabus that complements concept-based readings; in the case above, examining an archive works well because of its specificity. In more advanced discussions about, for example, multilevel models or sample sizes, you can introduce the responses from most any survey-based dataset.

With each new dataset that you introduce on your syllabus, you ought to include some primary materials from its researchers (e.g., an Executive Summary or other reports, details regarding the scope and methodology of the study, and/or a list of questions asked in the case of a survey). We also recommend including on your syllabus some secondary materials (e.g., public lectures, opinion pieces in the press, blog posts from knowledgeable sources) about how the results of the study were received. Assuming that students learn in a number of different ways, we suggest that you are more likely to engage your students—and, in turn, increase their confidence—by diversifying the texts included on your syllabus. By joining traditional materials with niche datasets, students can actively apply the statistical material that they learn to real-world examples, which strengthens and deepens their understanding of both.

What happens in class?

One primary objective of this approach is for non-STEM students to learn introductory statistics. In order for these students to take the intellectual risks necessary to succeed in a statistics course, they need to increase their confidence and buy-in. One way to increase both of these is to provide lots of low-cost opportunities for participation through discussion, collaboration, and in-class activities. We recommend that statistical lessons are couched in experiential education whenever possible. These exercises can be adapted to match the course's theme; for instance, in the exercise that teaches about error terms, discussed at the beginning of this paper, one might estimate the ages of individuals known with a discipline (e.g., guessing the age of Barbara Streisand for students in Jewish Studies or the age of Steve Jobs for business school students).

A secondary goal of a data literacy course for non-STEM students is substantive within a particular discipline, whether Italian Studies or History. As such, some of each class period will inevitably need to be lecture-based. This is necessary, for example, when you provide background information about a particular dataset or its context. We also recommend that students are incentivized to do a few close readings each term, by earning credit for introducing a text to the class and posing a few questions for further discussion.

From a practical perspective, each class meeting begins with a review of anonymized jitt answers. This process can take up to about 15 minutes, as you may want to use correct and incorrect jitt answers as opportunities to teach. Students are then invited to go over the exercise problems that they received as homework. Students are incentivized to work together on problems and to submit work collaboratively, as it increases not only the likelihood that they will *complete* the homework, but also that they are successful in doing so. In both of our classrooms, we often invite guest lecturers. For example, Andrew invited Amanda Cox, a statistician and expert data illustrator from the New York Times, to talk with his Communication in Statistics course at Columbia University about illustrating uncertainty and writing about statistics for a non-STEM audience. Alexis brought in Betsy Anthony, a researcher and administrator at the US Holocaust Memorial Museum, to speak with her Applied Statistics and Data Science course at the University of Toronto about her data-related work with the International Tracing Service

Archive and the International Committee of the Red Cross. Visitors need not be face-to-face, especially given the prevalence of video-meeting technology. Former students and colleagues are also welcome to workshop relevant papers and ideas with students. By introducing new voices into the classroom, students were encouraged to explore career paths related to statistics and data science, discuss niche questions with area experts, and apply their knowledge in a topical manner.

Application: a course in Jewish Studies

In the spring of 2019, and in the spring and fall of 2020, we used this model in an introductory applied statistics and data science course at the University of Toronto (Lerner, 2020, Pitic, 2020, Jankovic, 2019, Csillag, 2019). Each iteration of the 13-week, seminar-style course hosted approximately 15 students, who came from disciplines across the university, including English Literature, Russian Literature, Peace and Conflict Studies, Political Science, Economics, Geography, and Business. The first year of this course's offering included several students of geography, and so included a special week on how GIS is used in Holocaust and Genocide Studies (e.g., recent research by Anne Knowles and Paul Jaskot). Students were mostly sophomore, junior, and seniors and were able to use the course to satisfy a faculty-wide mathematical breadth requirement for undergraduates. As such, this course also posed a benefit to the department—in this case, the Anne Tanenbaum Centre for Jewish Studies; through the course, the centre could offer students a foray into the social sciences and an in-unit way to complete their breadth requirement.

As this was a course geared toward students of Jewish studies, the datasets used to teach general concepts were chosen with the goal of engaging Jewish studies as a discipline. For example, students used the Arolsen Archives to learn about 'where data comes from,' the Anti-Defamation League's Global 100: Index of Anti-Semitism to discuss problems related to operationalization, or the process of turning a nuanced concept (in this case, anti-Semitism) into a measureable factor, and the 2018 Survey of Jews in Canada data to practice simple and multiple linear regression. By the end of the course, students gained not only new information about demographics, politics, and culture within the scope of Jewish studies, but they also gained applicable tools for making and evaluating empirical claims within its purview.

The course also sought to teach students technical skills, such as how to turn archival and observational materials into a usable dataset, how to illustrate the correlations, patterns, and stories in that dataset, and how to discuss candidly the limitations of the dataset and/or method of analysis. In particular, students use text mining and content analysis skills learned in class to code one English-language Holocaust survivor testimony from a pre-selected subset of the USC Shoah Foundation's Visual History Archive (VHA), a collection of 55,000 videotaped testimonies given by survivors of the Holocaust and eight other genocides (e.g., interviews with survivors of the Rwandan Genocide, the Nanjing Massacre, and so on). Using VHA testimonies, students built and analyzed a collaborative dataset in four stages:

Stage One: Students completed a short paper about the archive and the meta-data of the testimony that they chose to focus upon. They answered questions, such as: who built the archive, who conducted the interview and in which language, what are the contents of the archive, and whether there are any potential ethical concerns related to the archive, the interview method, or the use of these testimonies. Students also identified possible testable hypotheses, considered the operationalization of their selected variables, and explained why the VHA would be an appropriate 'dataset' for the proposed question.

Stage Two: After designing a codebook together in class (which demanded that students put into practice their understanding of data collection and management), students inputted their code into a shared, cloud-based spreadsheet. As such, students were able to download and submit a .csv file of not only their results, but also the results of their classmates and course alumni. The instructor secured written approval from each student individually, so that their code might be used in future years by their classmates.

Stage Three: After learning about data visualization best practices and software skills in class, students submitted an original illustration, including a title, a key (if relevant), and a caption that explains the main features of the graph and highlights what the reader should take away from the illustration. Students were invited to workshop their illustration with their peers, who used this opportunity to provide knowledgeable and empathetic feedback on the effectiveness of an illustration. Students could use any data analysis and illustration platform for this assignment, such as ggplot2 in R, Tableau, or D3.

Stage Four: Students produced a short (800-word) essay that incorporated their tested hypothesis, their workshopped illustration, and an answer to the following question:

> *Do quantitative methods, organizational systems, and data visualization tools help scholars to learn new things about archival or observational materials?*

Overall, students reported that this assignment helped them to build confidence in data literacy, substantive knowledge in Jewish studies, and hard skills in quantitative methods. One student reported,

> "I am not one to usually write course evaluations. But in terms of this course, I felt compelled to describe how much this course has shifted the path of my academic

> endeavours. At first, I was not even signed up for this course, but after some research I signed up for it on a whim. It was the best decision I ever made. Before enrolling in this course, I thought I had a strong knowledge on the Holocaust but this completely proved me wrong in the best way possible. I was able to learn about different perspectives from groups I was not even aware were impacted by this genocide. I also gained hard skills and learned about opportunities where I could utilize quantitative methods in the realm of humanities…"

Other Possible Applications

This course model could be applied in many other settings, which could be used to attract potential students. Here we list a few that fall outside the usual domains of mathematics, natural and social sciences, and engineering.

Disciplinary Applications

This course can be modified to be used in a non-STEM department (e.g., literature, anthropology, or religious studies), as well as in a non-numerical institute (e.g., a theological seminary or a business school). The first step is to consider what level of substantive knowledge students are expected to have, and then to select datasets accordingly. For example, an introductory course might include more basic demographics or feeling thermometer-type datasets, whereas a more advanced group of students might be able to focus on a nuanced theme (e.g., data on a particular ethno-national conflict or on a specific sub-group).

Language Study

This course can also be modified to teach introductory statistics within the context of an advanced language course. For example, in the four-stage deliverable outlined earlier in this article, students learning Russian at an advanced-level could focus on a Russian-language survivor testimony. This would allow them to practice close listening in a foreign language. Instructors could also require the illustration and/or papers to be submitted in the language of instruction. This could help students expand their language learning beyond literature and current events-based study.

Level of Quantitative Expertise

The third avenue of modification is in the level of quantitative expertise. For example, in the four-stage deliverable model, instructors could substitute more advanced analytics in stage-three (e.g., requiring that students clean the dataset or build a multilevel model). Of course, this would likely require a larger dataset in order to go beyond a classroom exercise and to reach compelling conclusions. Therefore, this approach could even be adapted for introductory computer science exercises—e.g., web scraping or the design of natural language processing algorithms. Lab sections in STATA, R, or Stan can also be added to supplement the course.

Remote Instruction

During the spring and fall of 2020, this course was adapted for online instruction when the University of Toronto moved online during COVID-19. Remote instruction can be complicated because it assumes that all students have computer and internet access, as well as a physical environment conducive to learning. However, if students do have these conditions, the course's cloud-based collaborative tools, focus on digitally-available databases, and library access make this course easily adaptable for remote instruction.

Conclusion

In this paper, we presented a template for teaching statistics and data science to non-mathematically-minded undergraduates. We suggest that data literacy is a vital component of any higher education program across the hard sciences, the social sciences, and the humanities. Graduates of history and engineering should be equally capable of reading and assessing numerical information, such as in the charts and graphs used to supplement news articles. We suggest that the best way to establish data literacy is to build confidence and skills in quantitative methods, and provide a template—from what materials are needed, to evaluation methods, and lesson plans—for achieving this with a non-numerate audience. We include a case application of how this template has been used in Jewish studies, before concluding with a number of possible applications for student variation in discipline, language skill, and quantitative expertise. We also outline how this course can be modified for remote instruction.

References

"ADL's Global 100™ Anti-Semitism Index." n.d. ADL Global 100. Accessed July 9, 2020. https://global100.adl.org/map.

"Arolsen Archives - International Center on Nazi Persecution." n.d. Arolsen Archives. Accessed July 9, 2020. https://arolsen-archives.org/en/.

Boger, Pam. 2001. "The Benefit of Student-Generated Data in an Introductory Statistics Class." *Journal of Education for Business* 77 (1): 5–8. https://doi.org/10.1080/08832320109599663.

Bradstreet, Thomas E. 1996. “Teaching Introductory Statistics Courses So That Nonstatisticians Experience Statistical Reasoning.” *The American Statistician* 50 (1): 69–78. https://doi.org/10.1080/00031305.1996.10473545.

Brym, Robert, and Keith Neuman. n.d. “2018 Survey of Jews in Canada.” Berman Jewish Databank. Accessed July 9, 2020. https://www.jewishdatabank.org/databank/search-results/study/1043.

Crouch, Catherine H., and Mazur, Eric. 2001. "Peer Instruction: Ten Years of Experience and Results." *American Journal of Physics* 69 (9), 970-977.

Csillag, Ron. 2019. “New Course at U of T the First of Its Kind.” *The Canadian Jewish News* (blog). March 7. https://www.cjnews.com/news/canada/u-of-t-offers-new-jewish-stats-course.

Fujii, Lee Ann. 2010. “Shades of Truth and Lies: Interpreting Testimonies of War and Violence.” *Journal of Peace Research* 47 (2): 231–241. https://doi.org/10.1177/0022343309353097.

Gelman, Andrew. 2013. Setting up jitts online. Statistical Modeling, Causal Inference, and Social Science. September 27. https://statmodeling.stat.columbia.edu/2013/09/27/setting-up-jitts-online/.

Gelman, Andrew. 2019. “To Do: Construct a Build-Your-Own-Relevant-Statistics-Class Kit.” Statistical Modeling, Causal Inference, and Social Science. November 12. https://statmodeling.stat.columbia.edu/2019/11/12/to-do-construct-a-build-your-own-relevant-statistics-class-kit/.

Gelman, Andrew, Jennifer Hill, and Aki Vehtari. 2020. *Regression and Other Stories*. Cambridge: Cambridge University Press.

Gelman, Andrew, and Deborah Nolan. 2017. *Teaching Statistics: A Bag of Tricks*. Second edition. Oxford: Oxford University Press.

Jankovic, Jovana. 2019. “New ‘Jews: By the Numbers’ Course Introduces Students to Methods in Statistics and Digital Humanities.” University of Toronto Faculty of Arts & Science News. April 23. https://www.artsci.utoronto.ca/news/new-jews-numbers-course-introduces-students-methods-statistics-and-digital-humanities.

Lerner, Alexis. 2020. “Applied Statistics and Data Science in Jewish Studies.” Syllabus, University of Toronto. https://static1.squarespace.com/static/58a1878ff7e0abd069426c5d/t/5f0c7ea3553d1c4d55379 efb/1594654371396/Applied+Statistics+and+Data+Science+in+Jewish+Studies+Syllabus+Alexis+Lerner.pdf

Mazur, Eric. (1997). *Peer Instruction: A User's Manual.* Upper Saddle River, N.J.: Prentice-Hall.

Mazur, Eric, and Watkins, Jessica. 2010. "Just-in-Time Teaching and Peer Instruction." In *Just-in-time Teaching: Across the Disciplines, Across the Academy*, ed. Scott Simkins and Mark Maier, 39-62. Sterling, Virginia: Stylus.

Novak, G. M., Patterson, E. T., Gavrin, A. D., and Christian, W. 1999. *Just-in-time Teaching: Blending Active Learning with Web Technology.* Upper Saddle River, N.J.: Prentice-Hall.

Pitic, Badema. 2020. "Jews by the Numbers: An Introduction to Applied Statistics and Data Science for Students of Jewish Studies." *USC Shoah Foundation Center for Advanced Genocide Studies*, 2020, sec. Course Development Case Study. http://sfi.usc.edu/cagr/course_development/casestudy/jews-numbers-introduction-applied-statistics-and-data-science.

Rumsey, Deborah J. 2002. "Statistical Literacy as a Goal for Introductory Statistics Courses." *Journal of Statistics Education* 10 (3): 1-12. https://doi.org/10.1080/10691898.2002.11910678.

Simkins, Scott, and Maier, Mark. 2010. *Just-in-time Teaching: Across the Disciplines, Across the Academy*. Sterling, Va.: Stylus.

Slootmaeckers, Koen, Bart Kerremans, and Johan Adriaensen. 2014. "Too Afraid to Learn: Attitudes towards Statistics as a Barrier to Learning Statistics and to Acquiring Quantitative Skills." *Politics* 34 (2): 191–200. https://doi.org/10.1111/1467-9256.12042.

"USC Shoah Foundation Visual History Archive." n.d. USC Shoah Foundation. Accessed July 9, 2020. https://sfi.usc.edu/.

Wheelan, Charles. 2013. *Naked Statistics: Stripping Down The Dread From The Data*. New York: W.W. Norton.