Chapter 1

Active learning

This book supplies material for a course or self-study program in regression and causal inference, aimed at students who have already taken an introductory statistics course or have the need to use statistics and want to jump right into applied methods. But the principles of active learning can be applied at any level, from introductory to advanced courses.

1.1 Flipped classroom and collaborative learning

Here is a description of the "flipped classroom" idea:[1]

> "Binding credit for student participation to pre-class preparatory work enables class time to be spent in answering informed student questions and participating in in-depth discussions. Ensuring that students are prepared for class maximizes engagement during class time: the session becomes an interactive and dynamic experience where student pedagogical needs and reflections are explored."

That's how we think a course should go. Before each class, the students are supposed to read the book and do warmup assignments and homeworks. During class time, students are involved in activities and discussions and work together on problems—kind of like a high school math class, with the instructor as a leader and a coach, not a lecturer.

Much has been written on the benefits of active learning—classroom interactions that involve students doing things, talking with each other, and solving problems together.[2] Here are some key features of cooperative learning:[3]

- "Positive interdependence, that is, all members of a learning team are responsible for the learning of other members."
- "The teacher designs the learning activities and monitors the groups as they are engaged in team learning. Rather than functioning solely as an expert, dispensing knowledge to students, the teacher in collaborative learning serves as a facilitator."
- "Explicit attention to social skills. Students are required to cooperate with one another and are often given explicit rules and guidelines for appropriate social skills."
- "Face-to-face verbal problem solving, which holds advantages for both skilled and less skilled

[1]From Massachusetts Institute of Technology Office of Digital Learning (2022), https://openlearning.mit.edu/mit-faculty/residential-digital-innovations/student-pre-class-preparation-enhances-class-time/.

[2]Regarding teaching in general, see, for example, Robert Slavin (1980), Cooperative learning, *Review of Educational Research* 50, 315–342, Donald Bligh (1990), *What's the Point in Discussion*, and Catherine Crouch and Eric Mazur (2001), Peer Instruction: Ten years of experience and results, *American Journal of Physics* 69, 970–977. Regarding statistics education more specifically, see George Cobb (1992), Teaching statistics, in *Heeding the Call for Change: Suggestions for Curricular Action*, edited by Lynn Steen, 3–34, Deborah Nolan and Terence Speed (2000), *Stat Labs: Mathematical Statistics Through Applications*, and Allan Rossman and Beth Chance (2001), Teaching contemporary statistics through active learning, http://www.rossmanchance.com/pbs/pbs.html.

[3]From Jim Cooper and Randall Mueck (1990), Student involvement in learning: Cooperative learning and college instruction, *Journal of Excellence in College Teaching* 1, 68–76.

https://doi.org/10.1017/9781009436243.002 Published online by Cambridge University Press

students. Good students benefit from serving as tutors to the other members of the group; less proficient students receive diagnostic and remedial help from their teammates."

- "Students who are reluctant to participate in large class discussion are often quite comfortable contributing to small group interactions."

There is evidence that students learn more with active learning but feel like they learn less,[4] and this can be reflected in teaching evaluations.

So it can make sense for an instructor to get students prepared for the flipped classroom. This should start before the semester begins by making the structure clear in the course description, and it should continue on the first day of class with active student participation in the story and activity, as explained in Section 1.2 below. New teachers should also discuss the instructional plan with their supervisors or department chairs.

1.2 What happens during the semester?

First day of class

The instructor should *not* begin with, "Welcome to Statistics 200. I'm Professor Vehtari and in this course we will . . . ." It is better to just start the story for the first class (see Section 3.1 or, for first class of the second semester, Section 4.14), with lots of student involvement, stopping at various places to ask questions and have students work in pairs to think about them, and conclude with a path of how the story relates to the week's reading. Then continue with the class-participation activity for the first class. All this will set a norm of student participation throughout the semester.

After going through the story and activity, the instructor can take a breath and introduce the course, explaining while writing on the board the components of the course (readings, homeworks, feedback sheet, classes, and final exam) and the structure of each class period (story, activity, discuss reading and homework, computer demo, drill, discussion problem). Emphasize that the class period is all about motivation, exploration, problem solving, and answering questions. There will be no "lectures" and no slides, beyond certain images and instructions displayed as an aid to discussion. Go over students' responsibilities: they are expected before class to do the reading and homework assignment and to contribute to the shared document, and to show up and participate during every class. Discuss the goals of the course and what students will be expected to be able to do once the semester is over. Discuss the roles of mathematics, computing, and applications in the course. Pause to give students a chance to ask questions about the content or structure of the course.

Then is the time to go back and do the computer demo, the drill, and the discussion problem for the first class. At the end of the class, each student can be given a hard copy of the syllabus and reminded of the course website. Throughout the course, the most important things should still be written down for students. Don't assume that everything said is said in class is "heard," and also establish early for students that they will be responsible for being aware of online material.

Later classes

The other classes during the semester should have a similar structure, except that, in the middle of class, instead of giving an overview of the course, the instructor can talk about some aspect of the week's reading. This mini-lecture can be prepared ahead or in response to students' questions on the shared document, or the instructor can simply answer some questions in an unstructured fashion.

Most of each class period can be spent on the prepared material: the story, activity, computer demonstration, drill, and discussion problem. Instructors can use what is in this book or develop their own materials that more directly capture their experiences and perspectives, as well as the interests

[4] Louis Deslauriers, Logan McCarty, Kelly Miller, and Greg Kestin (2019), Measuring actual learning versus feeling of learning in response to being actively engaged in the classroom, *Proceedings of the National Academy of Sciences* 116, 19251–19257.

https://doi.org/10.1017/9781009436243.002 Published online by Cambridge University Press

and goals of the students. For example, many of our activities have a political science focus but could be changed to work with other subjects. In any case, the instructor should repeatedly explain the connections to the week's reading and homework. Students will not necessarily see these connections without being told, and it's important for these not to just be fun stories, activities, and so forth, but also to advance understanding of the core material.

Projector and slides

We don't like slides with bullet points, but we use the projector for stories, computer demonstrations, and class discussions. For computer demonstrations we set up an RStudio window and then go through code line by line, typing code into the text editor window and then copy-pasting it into the console to execute the R code.

For each semester of our course we have prepared an accompanying pdf slide deck, which includes all the images we display in our stories, class-participation activities, drills, and discussion problems. By design, the slides are minimal: they present material for the instructor to point to, but that is all.[5] More detailed slides can be helpful if this allows the instructor to better structure the course. We do not, however, recommend that slides be so detailed that they would be used by students as a replacement or shortcut to reading the textbook. The purpose of the slides should be to help the class period go more smoothly, not to present the material. When adapting the course, for example by introducing new examples, instructors can separately prepare what needs to be said in class and keep the slides minimal.

Blackboard or whiteboard

We find a blackboard or whiteboard to be helpful in discussions. Here are some tips:

- Start at the upper-left corner of the board and use the space efficiently; don't just start writing from the middle outward. Whatever's on the board should be organized: students aren't always paying attention at all times, so when they look up, what's on the board should make sense.
- When speaking, write key words on the board. For example, when outlining the structure of the course on the first day, as discussed in Section 1.2, write a list on the board: Story, Activity, Discuss reading, etc. When asking students to do a task together, write that task on the board so they don't have to memorize what they need to do.
- Conversely, when typing code into the computer or writing on the board—for example, sketching a graph or working through formulas—talk it through at the same time. Don't write silently. Provide that second channel of communication.
- If the setup allows projecting images onto the board, the course can include some fun things such as projecting a scatterplot and then having students go to the board to draw the regression line.
- If the classroom doesn't have a proper board, or switching between the board and the projector is inconvenient, you can simulate the board by typing with a big font.

What should students be doing in class?

Students should be working in pairs, talking with each other, and participating in class discussions. The instructor should keep an eye on the students to make sure they are focused. If necessary, ask them to put away their phones or computers (although computers can also be used in many in-class activities for making notes or computations). After asking students to work in pairs, make sure they all pair up. If any students are sitting alone or staring into space, connect them with existing pairs and ask them to move over if needed. When students are working together, walk around the room

[5] These are posted at http://www.stat.columbia.edu/~gelman/regression/ along with the rest of the material for the course.

https://doi.org/10.1017/9781009436243.002 Published online by Cambridge University Press

taking a look at what they are doing. Make sure they have pens and paper out so they can write out their ideas while talking. Ask them if they understand what they're being asked to do, and feel free to intervene in their conversations to point them in useful directions.

1.3 Active learning in class

Stories

Stories are fun to tell and can be fun to listen to. We recommend beginning each class with a story. Whether instructors use the stories presented here or choose their own, they should explain the source of every story; consider its difficulty level; include student involvement during the storytelling process; and conclude by connecting to the week's reading, the themes of the course, and the assumptions implicit in the storyline.

To get a sense of these general principles, stop and read the story for our first class period, the Wikipedia example in Section 3.1.

The *source* of that particular story was a consulting experience of ours, a serendipitous event that illustrates the principle that just about every problem is of statistical interest when looked at carefully. Reality is fractal. There is a tendency in textbooks to smooth over real-life difficulties; here we go into some of the hairy details.[6]

The *difficulty level* of the story is mixed. There are some quantitative steps involved in computing standard errors, and this uses formulas and concepts that aren't covered until several weeks later in the course, but the main points of the story should be clear even to students who hadn't seen those topics before.

When *telling the story*, the point is to give a sense of the use of experimentation and statistical analysis to attack an applied problem (in this case, increasing donations to Wikipedia) and then to see how we reacted when things went wrong. *Student involvement* comes when you stop the story at various places and ask questions to the students. After showing students Figures 2 and 3 and explaining that the estimated treatment effect is implausibly large, we ask them to look at the screenshot carefully and discuss in pairs to figure out what went wrong. Then we display Figure 4 and ask how this explains the large treatment effect. After going over this error, we ask the students to discuss in pairs how the problem could be fixed.

We conclude each story by *connecting to the week's reading and to the themes of the course*. For the story on the first day of class, the relevant topics come in Chapter 1 of *Regression and Other Stories*, and after finishing the story and class discussion, we ask the students to discuss in pairs the specific concerns in this example regarding generalizing from sample to population, from control to treatment, and from observed data to underlying constructs of interest. We then mention the larger themes of the connection between data collection and analysis, the importance of looking at data and checking assumptions and the relevance of mathematical calculations (in this case, the standard error of estimated proportions).

We also recommend the book, *Telling Stories with Data* by Rohan Alexander, which covers a wide range of topics on statistical communication, programming, and modeling to supplement any statistics course or self-learning program.[7]

Class-participation activities

For each class period we have constructed an activity that involves active student participation. Some of these activities involve collecting and analyzing data from the students in the class; others require students to estimate numbers or assess uncertainty. As with the stories, each of these activities has a

[6]For some examples of confusion and misrepresentation of data even in well-regarded texts, see Andrew Gelman (2011), Going beyond the book: Towards critical reading in statistics teaching, *Teaching Statistics* 34 (3), 82–86.

[7]Rohan Alexander (2022), *Telling Stories with Data*, CRC Press, https://tellingstorieswithdata.com.

direct point, but it is important also to draw the connection to the week's readings and homeworks. We often use the trick of throwing questions back at the students, asking *them* why they think we did a particular activity and how it relates to the larger themes of the course. By doing active learning every class period, we try to establish the norm of student participation. On the occasion when we throw a question to students and no one responds, we call on individual students, giving them the option to go along to the line to others if they are not ready to participate.

When teaching the course, we have not required students to hand in any responses to stories and class discussion. The instructor should do make sure that everyone is involved and participating while these activities are happening, talking to individual students as necessary to keep them active.

Hands-on computer demonstrations

A short demonstration is a great way to demonstrate the feel of live data analysis while providing an opportunity to field questions about programming, statistical analysis, and graphics using R (or whatever software is being used in the course). There is no need for the examples to be elaborate: even a few lines of code can be enough, and remember that this is just a single component of a 75-minute class. Everything is projected on the screen or board, and the instructor should speak the code aloud while typing it, stopping after each paragraph (as indicated by a blank line in the code) to explore and answer questions. For example, after reading in or simulating a vector called `x`, we might type `print(x)` into the console to take a look at its values. We post the code on the course website so that students can see it ahead of time or review it after class if they would like.

As part of the demonstration it is important that the instructor types the code live, or, if you are doing self-study, that you type the code rather than just copying it from an online source. For class the instructor should get the code working ahead of time so that the main part of the demonstration goes smoothly. Then there should be some time for improvisation, at which point some errors can arise. Some of the mistakes that we and our colleagues have made in class include forgetting how basic functions work (even some of those functions that we wrote!), not finding a data file where we expect, and messing up brackets or variable assignments and spending five minutes searching for the error. When these and other mistakes happen, watching the instructor resolve the problem is one of the most important lessons of the activity. Students need to learn not just how to do things right but also how to identify and fix mistakes.

Discussion of questions related to readings and homeworks

There is much more in any textbook than can realistically be covered in class, especially given that we are not following a lecture format. *Regression and Other Stories* includes discussion of general principles and methods, stories, worked examples, and code. During each class period we remind students of the week's reading and take some time to clarify points that confuse them, while recognizing that we can select only a subset of topics for discussion.

Before each class, students are required to contribute to a shared online file such as a Google doc, either asking a question on the readings or homework assignment due for that class, asking a more general question about the topics of the course, or responding to a question that another student has already asked. Before the class we quickly scan the day's document and choose a few questions to discuss during class. This is yet another way to promote student involvement as well as an opportunity to resolve confusion. And, when in doubt, use the tactic of asking rather than telling. When a student asks a question, we can reply, What do you think? Eventually we have to answer the question (or admit we don't know), and we don't want to drag this out too much, but occasionally bouncing a question back can help clarify what the student's question really is.

When responding to questions or getting into discussions, we avoid digressions that don't advance the main thread of the course. It's fine to skip a question or to put it off until later, and, again, when answering questions and leading discussions it is a good idea to repeatedly touch back to the week's readings and homeworks. Connections might seem obvious to the instructor but not so clear to

https://doi.org/10.1017/9781009436243.002 Published online by Cambridge University Press

students. If the class has a teaching assistant, these questions can also serve as the basis for further clarification in section meetings.

Mini-lecture

Responses to student questions will often include a mini-lecture during which the instructor works out a problem on the blackboard. For example, an important but challenging skill in logistic regression is to go back and forth between the coefficients (a, b) and the curve, $\text{logit}^{-1}(a + bx)$. When we reach logistic regression in the course, we demonstrate this activity in both directions: given the coefficients, drawing the curve, and given the curve, extracting the coefficients; see for example Figure 13.1 in *Regression and Other Stories*. We do this on the board, drawing the curves freehand and performing the computations approximately, emphasizing that these are skills we expect the students to be able to perform themselves for the exams and more generally when interpreting fitted models in the literature.

Drills

In past years we have often found students to have difficulties with what we had considered to be basic concepts routine computations. In response, we try to spend some time every class on a drill: a set of short easy problems that we write on the board and ask students to work on in their notebooks and then discuss in pairs. This book has a set of drill questions for each class period, with a solution given for the first problem in each group.

Sometimes a drill is just a series of related questions; other times the questions build up in complexity to get at different angles of a problem. In presenting a drill, we solve the first problem on the board, and students can use this as a template to solve the remaining problems on the list. After the drill is done, we can discuss its relevance to the week's readings and the course more generally.

Discussion problems

Finally, we like to conclude each class with a prepared discussion problem on which students can work together in pairs or small groups. These are similar to "concept tests" in peer instruction.[8] A discussion problem should be challenging, so that students can work for a few minutes together to think about it. Some discussion problems are open ended; others have precise answers. Often the challenge is not in solving the problem but more in setting it up. And, again, end the discussion by tying the problem back to the day's readings and also situated it within the larger plan of the course.

Substitutions

This book presents material for 52 class periods, with warmup assignments, homework assignments, a story, a class-participation activity, a computer demonstration, drills, and a discussion problems for each. There is not always a sharp division between all of these. The stories should be told with many pauses for student participation, the activities have aspects of storytelling, the demonstrations often connect to the stories and activities, and the discussion problems often also include stories, as well as being opportunities for students to work together, often with coding. The warmup and homework assignments and drills represent different ways to practice the course material.

1.4 Scheduling

Think of the course as presented here as a template: instructors can teach it as is or swap in alternative readings, warmup and homework assignments, stories, activities, computer demonstrations, drills,

[8]See, for example, Eric Mazur and Jessica Watkins (2009), Just-in-time teaching and peer instruction, *Just in Time Teaching: Across the Disciplines, Across the Academy*, edited by Scott Simkins and Mark Maier, 39–62, Sterling, Va.: Stylus.

https://doi.org/10.1017/9781009436243.002 Published online by Cambridge University Press

and discussion problems. The key is for all these pieces to fit together, providing a framework for teaching and learning.

Workload

The course as designed here is intended to average 10 hours per week, roughly divided as:

- 1 hour reading the textbook
- 0.5 hours on pre-class warmup assignments
- 2.5 hours in the classroom
- 6 hours on homework.

The semester concludes with a 3-hour final exam for which students are expected to study for about a week, that is, another 10 hours.

The point of giving these estimates is not to dictate how other courses should go but rather to provide a baseline: if students are expected to devote more or less than 10 hours each week, the workload should be adjusted accordingly. In addition, in describing in-class plans in this book, we have erred in the direction of providing more details, and realistically it will not be possible to do all the stories, activities, demonstrations, drills, and discussion problems given here for each class. Depending on how the class is going and how much time is needed to go over student questions, it will be necessary to skip some of the listed material. This is fine: the key points are to keep students actively involved and working together during the entire class period and to connect each week's material with real-world examples.

It's also important to control the workload of the instructor and any teaching assistants for the course. Indeed, one of our motivations in preparing this book is to make active learning accessible to busy instructors: we have put together these stories, activities, demonstrations, and problems so that instructors can be freed to spend their time helping students learn. In a large class, a big challenge is grading, especially if there are no teaching assistants to help. In that setting, it can make sense to set up a peer grading system with some subset of assignments checked directly by the instructor. The other challenge in a large class is meeting with students; here the scalable solution is scheduled help sessions where the instructor can help students figure things out together, rather than frustrating and time-consuming one-on-one meetings.

Fitting active learning into a busy class period

We have arranged the activities in this book to fit into a series of 75-minute classes, following roughly this plan:

1. (10 minutes) Story
2. (15 minutes) Class-participation activity
3. (15 minutes) Class discussion of questions related to readings and homeworks
4. (15 minutes) Computer demonstration
5. (10 minutes) Drill
6. (10 minutes) Small-group discussion problems.

These timings are only approximate, as the stories, activities, and demonstrations in this book vary in how much class time they require, and the time allocated for discussion of readings and homeworks is flexible. Sometimes a story can take 20 minutes of class time, especially if it is used as a springboard for class discussion. Realistically, we rarely include all the above segments in every class, as usually one piece or another takes longer than expected. Instructors can adapt based on the progress and struggles the students show in class each week.

Many of the stories, activities, and demonstrations in this book could take a lot more class time if you let them. These are fun examples—we only included things in this course that were interesting to

https://doi.org/10.1017/9781009436243.002 Published online by Cambridge University Press

us—and so, indeed, any single piece contains depths that there would not be time to explore without elbowing out other topics for the week.

With that in mind, you should consider the material here as resources more than as scripts. If a story would take too long to go through in detail, or if an activity has so many steps that it would feel like a distraction, then feel free to use shorter versions that will fit your needs. It's good to switch up activities every 10 to 20 minutes rather than getting stuck too long doing the same thing, even active learning. Similarly with the computer demonstrations: you can start one, entering and altering code until you feel that you've learned something, and then stop, with no requirement to go to the end. In this book we would rather include more detail than less, with the understanding that instructors and students will take what they need, and then the rest is always available for later study.

Pre-class warmup assignments

We have prepared a few quick questions due before each class to keep students on track regarding the week's reading. Instructors should feel free to adapt these to the needs of their classes. The warmup assignments are intended to be straightforward for students who have read the relevant textbook material (in other words, when the students look at the questions they are guided to read the relevant material).

Homework assignments

Everything that students are expected to learn should be in homework assignments and drills. There is no point in doing a derivation on the board, for example, if students are not required to reproduce it in some form by themselves. But with this realization comes a challenge, as is is difficult to include homework problems on all the important topics covered in the course while still respecting the limited time that students have to spend on any course. We suspect that the best solution here is to limit what is covered in class—better to cover one key concept and set of skills and make sure that students can do them well, than to hit several topics shallowly. For now, though, we'll just say that the instructor should affirmatively decide on what skills the students are expected to learn, give practice in these skills, and not to expect learning on topics that are mentioned in class or the textbook but not included in homework assignments and exams. It can be helpful to provide guidance on readings to help students decide where to focus their effort. Also it's important for the instructor to monitor progress during the semester to see which assignments are giving students trouble, to give a chance to reinforce these lessons.

We assign homework assignments with a range of difficulty levels. It's valuable to nail down key skills that can be used as building blocks, to have more involved problems that challenge students and promote independent thinking, and to give students practice collaborating on some of the computing challenges involved in working with real applications.

1.5 Assessment and feedback

Tracking students during the semester

Warmup assignments and homeworks are due for every class, and students should be actively involved during the class period, so there are many opportunities to observe and evaluate them and to talk with them when they run into trouble. If the class is large and the instructor does not have the time or resources to grade the warmup and homework assignments, they can be graded using peer grading. It is important that students try these problems so they can follow the discussions in class and learn from their mistakes.

https://doi.org/10.1017/9781009436243.002 Published online by Cambridge University Press

Final exams and practice final exams

We have prepared several multiple-choice exam questions for each chapter covered in the course. Having created this test bank, we sample one question at random from each chapter and put these together to form a practice exam to give to the class a few weeks ahead of time. We then create the exam itself by taking a new sample of two questions per chapter. The exam questions we prepared for the two semesters corresponding to Chapters 1–12 and 13–22 of *Regression and Other Stories* are in Appendices B.1 and B.2 of this book. Instructors can use these questions directly, create new questions by making small changes to the ones here, or write entirely new questions. The exam question bank can be open to the students, which guides the students to learn what is also tested. When creating new exam questions, it is good to think are what the learning objectives and how student learning is guided by these new questions.

Grading

Grading should be transparent and efficient. We use a weighted average of scores on pre-class assignments (graded based on completion rather than correct answers), homeworks, class participation, and the final exam.

At each stage, we recommend that grading be clear and simple. We grade each homework problem as 1 (correct), 0.5 (attempted but wrong), or 0 (not attempted), with no partial credit beyond that. Class participation for each period is a 1 (contributed meaningfully to the shared document) or 0 (did not contribute). The final exam is multiple choice, so each problem gets a grade of 0 or 1, with no partial credit. From a statistical point of view, this works because the grade is averaged over multiple homework problems for each class period, many class periods, and many final exam questions, so it is not necessary to make fine-grained decisions for each problem. Simple grading also makes it more feasible to have students correct and resubmit their homework assignments.

1.6 Some general issues in teaching and communication

We recommend *How to Talk So Kids Will Listen and Listen So Kids Will Talk* by Adele Faber and Elaine Mazlish.[9] That book, which originally appeared in 1980, is not explicitly about teaching, but we find its themes and techniques to be directly relevant to working with students and communication more generally. A lot of the ideas involve being direct about your goals and obstacles.

When communicating with students about the course itself, we have to be careful: we want students' feedback, but we don't want to turn them into theater critics either. We try to elicit their reactions and suggestions while making it clear that this is all within the bounds of the course as we have designed it.

There's also general advice for communication that works in class too, such as using people's names, answering questions with questions such as "What do you think?" that motivate students to stay focused in the conversation, and never describing a task as easy. Instead of saying, "We'll start with a simple example," or "Here's an easy problem for you," say, "Here's a problem that's challenging but doable, if you go about it the right way."

Pair discussion gives students a space to consider ideas they might not be ready to share with the whole class, and working in pairs gives them a chance to explain things to each other. Teaching can be an effective way to learn.

[9] Adele Faber and Elaine Mazlish (1980), *How to Talk So Kids Will Listen and Listen So Kids Will Talk*, New York: Simon and Schuster.