

# Teaching Tactics

[Eric Hehner](#)

University of Toronto

I like teaching. As an undergraduate I was a teaching assistant (TA) in a couple of math and physics courses. As a graduate student in computer science I TA'd more courses, and taught my first two courses. By then I was hooked on teaching. I also like research, so my choice to be a university professor was easy, and I was hired by the University of Toronto, Department of Computer Science.

My duties were the same as any other beginning professor in my department in those years (1974). I taught 4 half-courses, had some administrative duties, supervised a couple of graduate students' research, plus my own research. My three undergraduate courses were at a suburban campus, and I had no TAs. One of those courses was csc108, programming for non-computer-science students — 262 of them, the capacity of the largest lecture hall. I wanted everyone to understand everything. As I slogged through the interminable marking, it was overwhelmingly clear that many of them understood very little. With each wrong answer, I would agonize over how I had failed to make the matter clear, and how I could improve my explanations. Marking an assignment took much longer than the time between assignments.

In that first term, my feeling of failure grew; I would cry when marking tests; I thought I might have to resign. I saved myself by radically lowering my criterion of success. My new criterion became: almost everyone should understand something, and someone should understand almost everything.

A few years later, after I had tenure, I was assigned two sections of a course, and I decided that was the perfect opportunity to perform an experiment. In the first section, I was the expert. My explanations were authoritative. I solved each example problem smoothly, without error. Students took notes. But I think the material going in their eyes and ears and coming out their writing hand bypassed the part of the brain that thinks and remembers. In the second section, I acted unsure of myself. I would say that I think it's something like this, but you better read the textbook to get it right. When solving a problem, I would stop, scratch my head, and ask if anyone could help me out. Seeing the prof in trouble, students engaged and focussed, some to try to help, others to show up the stupid prof. Both sections did the same assignments, test, and exam. In the end, there were two points of comparison. One difference was the teaching evaluations: the first section said I was a great prof who knew his stuff; the second section said I was a terrible prof who didn't know what he was doing. The other difference was in the grades: the first section did not do as well as the second section. I don't claim that this one experiment has any statistical significance, and I won't make any big conclusions from it. But it seemed clear to me that I had to decide whether my main purpose was to be admired and to receive good teaching evaluations, or to provide the best learning environment for the students. I don't mean that the two goals are mutually exclusive, but they are not the same; sometimes, as my experiment demonstrates, there is a trade-off.

Teachers often like to make their lectures more enjoyable, and themselves more popular, by making their lectures entertaining. There's no harm in presenting course material in an entertaining way. But there is harm in adding entertainment to the course material. For example, PowerPoint (or Keynote) presentations are an excellent way to present material, but there's a strong temptation to make it pretty, with colors and border designs and cartoons. These distract students from the message. Students enjoy the lectures more, but what they remember is not the message, it's the added entertainment. If the lecture material is not interesting by itself, it does not become more interesting by adding interesting distractions.

I have never been comfortable with the policy on cheating. Some students like to work together and help each other, and that seems good to me. These days, when a problem is assigned, it seems likely that its solution can be found on the internet. Not all learning comes from the prof. Any rule that tries to distinguish between legitimate help and cheating has grey areas. And there will always be cheaters who don't get caught. My solution has been to separate education and evaluation.

Problem assignments are educational and I give out lots of assignments; the students can take their pick which ones they do. The solutions are all available, on the course website, from the start of term. If they cheat and look at the solutions before doing the problems, they only cheat themselves. They can work in groups if they want; I encourage it. Each week the tutorial hours are spent doing problems, with the TA there to help anyone who wants help. Their solutions are not graded. There are three tests and a final exam, held with the usual invigilation; they are not educational; they are for evaluation; they are graded. The test and exam questions are designed so that those who did the assignments will do well on the tests and exam, and those who did not do the assignments will have trouble on the tests and exams. I realize that my solution to the cheating problem cannot be used for all courses. I also realize that it disadvantages those people who do not do well under exam pressure, but do very well on assignments. But it ends the unfairness and unpleasantness of cheating, and allows group help, TA help, and internet help as part of learning.

I never use true/false questions on tests and exams, but those who do use them know that you have to subtract the number of wrong answers from the number of right answers. Otherwise random guessing would get 50%, instead of 0%. I never use multiple choice questions, but those who do use them do not all know that the same principle applies. If a question has 5 choices, then a right answer gets 1 mark, and a wrong answer gets  $-1/4$  mark. (You can scale this as you wish; for example, a right answer gets 4 marks, and a wrong answer gets  $-1$  mark.) All my test and exam questions require written answers, and the same principle applies! Here's the standard paragraph at the top of my tests and exams.

The value of each question is indicated in square brackets.

A blank answer is worth about one-third of the marks;  
to that, marks will be added for readable and relevant and correct information,  
and marks will be subtracted for unreadable or irrelevant or incorrect information.

When a student does not know an answer, I don't want the student to write random nonsense in the hope that the marker will find something of value in it. A student who knows that they do not know the answer knows more than a student who does not know that they don't know the answer, and they should be rewarded for saying "I don't know". More generally, the student is rewarded for all they do know, and penalized for writing what they don't know. As a nice side-benefit, the marker's time is saved when the student writes less nonsense.

A few years ago, it occurred to me that I tell students about the wonders of modern technology using stone age technology: chalk on slate (stone on stone). For many centuries, teachers have stood at the front of a class, copying from their notes onto the chalkboard, while the students copy from the chalkboard into their notebooks. In effect, the teacher and students are a copying machine, making a copy of the teacher's notes for each student. This was sensible before the invention of photocopiers, but not now.

In 2006 I was asked by the U.S. National Technological University (now part of Walden University, Minneapolis) to develop an online course in formal methods of software engineering. I was sent to Orlando, Florida for two weeks to record the course with the help of Disney Studios. They had the best equipment and technical expertise; they had teleprompters and special lighting and the right camera angles, and I had to wear the right color of clothes. They

had put a lot of thought into how to present and run an online course. I was fortunate to be an early online course creator, and I learned a lot. Back at the University of Toronto I decided to put one of my own courses [online](#). I didn't need fancy equipment; I just needed my laptop computer. For the parts of the lectures that show me, I put my notes on my computer screen (I see them; the students don't), and the built-in camera and microphone record me. For the parts that don't show me, the PowerPoint (or Keynote) presentation on my computer screen is recorded, together with my voice, together with the cursor pointing at anything I want to point at. It's easy. I recorded it all at home. The recording quality isn't as good as Disney Studios quality, but it's good enough.

Rather than trying to fit the material to an allocated 50 minute period, I make each lecture the right length for the material; it turned out to be 34 lecture segments with an average segment length of 18 minutes, which matches a student's attention span better. The students can watch the lectures wherever and whenever they want, even on their smartphones. No more tired students sleeping through lectures. No more missed lectures due to illness and religious holidays. If they miss a point, they can replay it as many times as necessary. They can replay a lecture when studying for the exam. At any point when they want to think about what was just said, they can pause. This solves an old lecturing problem: when you ask a question, how long do you wait for the students to think before you continue? If you wait only a short time, the slower students don't have enough time, and they don't even try. If you wait long enough for the slow students, the quick thinkers get bored. But now, each student can pause the lecture for the right length of time for them.

Some professors think that the way to make an online course is to lecture, as usual, to a classroom of students, and to have it recorded. That misses many of the benefits. An online lecture isn't speaking to a roomful of students; it's speaking to just one student. And that changes the way you talk and what you say. It's more like a personal, one-on-one talk, and that makes it easier to listen to. If you make a mistake, or think of a better way to say it, you just record over the bit that needs to be improved. (The recording program has an editor, just like your word processor, with cut and paste, so the new bit doesn't have to be the same length as the bit it replaces.)

Putting lectures online is better for both the lecturer and the students, but it does not cancel the need for personal contact. In my online course, I use all three of my allocated classroom hours per week, but not for lecturing: they are all problem-solving sessions. People work in groups, or individually, while I and the TAs wander around asking people how they are doing, giving hints, and answering questions. When I have to be out of town for a conference, or if I am ill, there is no scramble to cover my absence; the class and the TAs can carry on perfectly well without me. Outside class hours, students can email me to ask questions or request an appointment, just the same as in a traditional offline course. There is also the usual online chat room (aka bulletin board).

I don't claim to have made any great innovations, or to express insights that have not been expressed before. I have just related some of my experience, in case it is of interest to you.

[other essays](#)