Teaching philosophy

True understanding includes not just knowledge of how to *apply* a new technique but also *why* it works, what its *limits* are, and how it *relates* to other techniques. I believe that students benefit from additional class time devoted to these topics, even at the inevitable cost of covering less material. A teacher's goal should be to maximize the amount of material *understood*, not merely the amount *covered*. Two examples of extra material of this type that I would include in an upper-level undergraduate course on probability in computer science (such as one using Mitzenmacher and Upfal's book) follow.

1. A few weeks into the course, present a lecture on the subtleties of modeling with probability. Describe the Monty Hall problem, the obvious incorrect solution, and then work with the class to create a correct analysis. After that present and justify a simple rule of decision making using probability: Condition on everything you know, including the fact that you know what you know. Justify this rule by proving that it yields decisions that are the best possible. Reinforce these concepts using homework problems on additional examples of tricky modeling, for example the inadvisability of testing on training data and the risk of publication bias in scientific experimentation.

2. After introducing Chernoff bounds and some examples of their uses, spend some time showing how close to tight they are. State the central limit theorem and then make it plausible by sketching a short proof that the binomial distribution is roughly normally distributed near its peak.[1] Next, show how Chernoff bounds and the normal approximation are identical up to constants in the exponent. Finally, show that the normal approximation is sometimes very wrong, using a trivial example: a single Poisson trial with small expectation. This explains why the standard normal-like Chernoff bound is only valid for deviations from the mean of magnitude less than the mean.

Teaching experience

As an undergraduate I was a "peer leader/tutor" for several years. In this role I led homework/study sessions for introductory physics and calculus courses, breaking people into groups as needed and answering questions.

My teaching experience broadened as a PhD student at Brown. I took a year-long (non-credit) course on teaching hosted by the Sheridan Center for Teaching and Learning. My teaching assistant duties at Brown included office hours, grading, homework question design, course note preparation, and an occasional lecture when the professor was out of town. I also did a number of lectures as part of my course work, including a couple that the professor asked me to do again the next time she taught the course.

As of mid January I will have presented six distinct works of mine at conferences, plus one work that was not mine (the authors could not attend the conference). I have also presented these works in a longer form at Brown Theory Lunches and various invited talks.

Teaching interests

The following are some of the courses I am especially interested in teaching (and am qualified to teach):

- **Early undergraduate:** Basic data structures and basic discrete math.
- **Late undergraduate / early graduate:** Algorithms, artificial intelligence, combinatorial optimization, and probability in computing.
- **Late graduate:** Approximation algorithms and algorithms related to economics/game theory.

I would be happy to teach other courses as needed.

[1] In short, use $\binom{n}{k+1} = \frac{n-k-1}{k+1}\binom{n}{k}$ repeatedly and proceed similarly to the analysis of the birthday paradox.