Teaching and Advising Statement

Bo Waggoner
Warren Center for Network and Data Sciences, University of Pennsylvania
Fall 2017

Teaching

Teaching is an exciting opportunity to help others gain new skills and knowledge. It is also a continual challenge to improve one's own understanding and communication ability. I would identify the key components of my approach as (a) adapting style to context; (b) preference for "narrative" learning processes rather than teach-then-quiz; and (c) skepticism and openness to change.

(a) An hour's instruction can be used in many ways – dry-but-dense lectures, Socractic Q&A, student presentations, quizzes, small-group problem-solving – and I prefer to vary my style to fit the context and goals of the class. In general, students who are more motivated and prepared to learn this particular material are better served by an efficient lecture format, while others (or the same students, at the start of the semester) benefit from more interaction and background/motivation.

In particular, I see a distinction between the philosophy of undergraduate versus graduate courses. Graduate courses are usually focused on effectively disseminating knowledge to students who are motivated and prepared to learn about that topic. Undergraduate courses fall under an institution's larger educational mission and also serve students with varying career paths (industry, medicine, law, ...). The instructor is generally more responsible for motivation as well as, in my opinion, connections to the ethics and consequences of the topic to society at large.

(b) People learn best when: (1) they are confronted by a problem that (2) they are motivated to solve, which requires that (3) they obtain new knowledge. I try to structure my teaching in this way whenever possible, including lectures or even problem sets (example from Harvard's CS236r): rather than imparting knowledge, then checking it with questions, we set a goal or challenge for the students, then help guide them in acquiring the knowledge needed to achieve that goal.

Such narratives can come from history of the subject or simply posing a challenge. Picturing the human "computers" of the 1930s using pen and paper, we can ask: what *is* a mathematical definition of computation? By guiding student discussion, the definition of a Turing Machine can unfold relatively naturally. Similarly, after motivating linear programming, we can draw a polytope and ask: How would you navigate from an arbitrary corner to the optimal one? The idea of the simplex algorithm can arise in discussion, after which students can be challenged to formalize it.

(c) In evaluating my own approach, I recall the philosophy of a physician I once met: to be "aggressive with diagnosis, but conservative with treatment". It is important to acknowledge challenges in an era of video lectures and laptop distractions; and opportunities of methods such as active learning and uses of technology in (and out of) the classroom. However, this desire to improve does not necessarily mandate drastic changes in style. By keeping my eyes open and occasionally experimenting, I hope to gradually expand and hone my repertoire of teaching methods.

Teaching and service. As a professor and researcher, I view part of my role to be distilling knowledge from courses into slides, lecture notes, or less-formal materials such as blog posts that can be helpful to others beyond my own classroom. I am also interested in other ways to help expand access to knowledge, either via online courses or learning opportunities aimed at my local community.

Topics - core and introductory. I am comfortable teaching undergraduate and graduate courses on algorithms; theory of computation (experience: Salil Vadhan's CS121 at Harvard); machine learning (particularly theoretical); optimization (e.g. linear programming); artificial intelligence; and game theory / mechanism design.

Topics - seminars and special topics. Here there are many possibilities and the choice can depend on the needs of the department. Some examples:

- General algorithmic game theory or "EconCS" course; can be more broad and interdisciplinary (e.g. David Parkes' CS136 at Harvard or Tim Roughgarden's CS269I at Stanford); or more advanced and theoretical (e.g Aaron Roth and Rakesh Vohra's NETS412 at Penn, which I will be teaching in spring 2018).
- More specialized AGT topics course, e.g. information elicitation, crowdsourcing, and ML (experience: Yiling Chen's CS236r at Harvard); or online learning, prediction, and game theory (e.g. Jacob Abernethy's EECS 598 at Michigan).
- Online learning course covering basic models of adversarial and distributional data arrival, experts problems, online convex optimization, bandits, etc.
- Differential privacy (e.g. Salil Vadhan's CS229r, Aaron Roth's CIS800, etc.).
- Models and algorithms for networks, undergraduate (e.g. Michael Kearns' NETS 112 at Penn) or graduate (experience: Yaron Singer's CS286r at Harvard).
- Randomized and online algorithms; or similarly, a seminar for first-year theory students surveying "tools and concepts", particularly probabilistic tools such as tail bounds.

Topics - interdisciplinary. I envision an algorithmic game theory/mechanism design course that is taught jointly with the economics department, both in order to expose computer scientists to the econ perspective and also to attract economics students to the algorithmic one.

Perhaps a bit later in my career, I also have hopes for "softer" interdisciplinary courses focusing on privacy and ethics of computation in every day life, somewhat similar to Owen Astrachan's CS342s at Duke.

Experience. I have acted as a teaching assistant (designing and grading homeworks, leading recitations, guest lectures) for undergraduate theory of CS (Salil Vadhan) and graduate seminars on networks (Yaron Singer) and information elicitation (Yiling Chen); I have also been invited to guest lecture in a number of courses including optimization and algorithmic game theory. In spring 2018 I will teach NETS412 (Algorithmic Game Theory) at Penn.

Advising

An advisor of graduate students has significant influence on their success in graduate school, the trajectory of their future career, and often their mental health and well-being. I take these responsibilitites extremely seriously, while recognizing that students must ultimately take them into their own hands.

Students enter graduate school from a huge variety of backgrounds and, in my experience, are generally expected to gradually "patch" their knowledge and skill sets as holes or obstacles to myopic research goals are discovered. I think this is, to some extent, a mistake, and some effort should be devoted to covering certain bases from early on in the PhD: writing and structuring papers, speaking and giving presentations, mathematics and programming skills as appropriate, and so on.

A significant and perhaps under-appreciated challenge in advising is selecting good problems for students to work on. Good problem selection for one's students can be quite challenging (balancing difficulty, interest level, and publishability) and can also have a significant impact on the students' trajectories. My philosophy is to keep my ears open for potential problems; keep students on at least one low-risk problem in a main area of my expertise; to otherwise diversify; and to be willing to pause or let go unfruitful projects.

There is a saying, "it takes a village to raise a child", and I think this remains true at higher levels of education as well. I expect a large part of my advisees' graduate educations to come from their peers and from my colleagues: via courses, informal interactions, seminars and research groups, collaboration, and internships. Similarly, I view myself as partially responsible to all graduate students in my department whether in the classes I teach, for general advice and guidance, or as a potential collaborator.

Experience. Toward the end of my PhD and throughout my postdoc, I have been fortunate to work with students in more of a mentorship role, including undergraduates, masters students, and PhD students. In graduate school, I was fortunate to work with undergraduates along with my advisor, Yiling Chen. These include Adam Su on an undergraduate thesis, Shuran Zheng on a paper currently in submission, and Jerry Anunrojwong in ongoing work. My time at Penn has given me the opportunity to work with many graduate students with various years of experience and exposure to additional advising styles of the faculty.

I do not expect to have all the answers and plan to reach out to more experienced colleagues, as well as my current network of mentors, for advice and guidance as I gain experience in advising students.