Teaching philosophy

In my one-on-one teaching interactions, whether during TA office hours, informally mentoring younger students, or helping users on the statistics Stack Exchange site, I strongly believe that it is usually far more effective to guide learners through figuring things out for themselves. The job of a teacher there is to help figure out appropriate sub-goals when the problem seems too daunting, to give pointers in the right direction when needed, and to ask probing questions to get the student thinking about the right problems. Not only does this lead to learning "sticking" more, as evidenced by a long line of research in educational studies from John Dewey onwards, but building up the facility to be creative about solving problems with various tools is far more important than just memorizing what those tools are.

At the small liberal arts school I attended for undergrad, this philosophy is explicit in the course structure: students do most of the traditional "learning" on their own, then gather to discuss and present problem solutions or their own further investigations into related topics. The professor is there as a guide, helping when a student gets lost on something in particular but more importantly setting up the waypoints in a field that the students, without the context the professor has, might otherwise get lost within before even starting.

With a traditional classroom structure, and especially with larger class sizes, it's harder to adopt this approach wholesale. The best classes can still have the feel of the whole class – including the instructor – working together to solve a problem. Sometimes that is impractical, but the best courses can still certainly incorporate some of these lessons. Lectures can be used to overview the basics and to situate topics within the broader field and relate them to things that will not be covered in depth. The students then build their understanding through carefully-constructed problem sets and larger-scale projects. One of my favorite jobs as a TA was developing homework questions that can teach the whole class something interesting, even for classes with students from quite varied backgrounds. Working on these problems in groups can help students learn to effectively collaborate, and cement their own understanding by explaining ideas to one another. One technique I learned from graduate school is to encourage group work, but also request students to include a simple brief statement of who they worked with and to what extent, fostering habits of honesty and accountability in learning.

This theme of guiding students through thinking about things themselves shapes much of my work. I try to make my teaching, my research talks, and even my papers tell a story: motivate a problem, and then describe a sequence of various approaches, where each one falls short, and how the next approach addresses that problem. This not only often makes the presentation more engaging, but can also help avoid the all-too-common phenomenon in both introductory lectures and research talks where solutions are presented as if they sprang fully-formed from the mind of a genius. Students with that view of the world, when they don't immediately produce their own brilliant work in a flash of insight, can easily convince themselves that they aren't cut out for this subject, or that they're just stuck and need help from someone who "gets it" before moving on.

I've also learned through my informal mentoring experiences as a postdoc that even in advising research students, the careful balance to strike between simply handing the student problems to solve and leaving them entirely to their own devices has much in common with other forms of teaching.

Outreach and diversity

During my student years (and since), like many academics I identified strongly with my academic and intellectual pursuits. Unlike most academics, though, I was also struggling to accept another aspect of myself: an LGBT identity. Although most of the academic community was by that time very accepting, there were – and still are – very few role models who shared that identity, making it harder for me to fundamentally accept that the two were compatible.

For me, the academic identity won out, which made accepting myself for who I am a much more difficult process than it might have otherwise been. For others, their "competing" identities might be more settled than mine was, or for whatever reason more difficult to ignore. It's very easy for me to understand how this lack of role models and visible examples of success could discourage such students from entering a field in which they might otherwise excel – even if those in the field are entirely well-meaning and accepting.

I plan to be much more visible in my identity in the future, hopefully helping students with similar identities both to accept themselves and to see that careers in machine learning or STEM more broadly are accessible to them. As a first step towards that end, I've recently begun volunteering with the new Queer in AI group.

I will also strive to support those of other marginalized identities in whatever ways I can, from support for peer groups and outreach programs to attempting to be particularly mindful of potential implicit biases and assumptions that might not apply in dealing with students from varying backgrounds. Both my undergraduate and graduate institutions, Swarthmore and Carnegie Mellon, have made great efforts towards making undergraduate computer science more accessible to those outside traditional backgrounds for the field. I am very fortunate to have seen some of these strategies in practice, learning techniques I can help carry forward wherever I end up.

Teaching interests

My primary interest is of course machine learning, and I look forward to teaching it from the introductory undergraduate level (perhaps as part of a broader artificial intelligence course) to advanced graduate courses in particular topics such as statistical learning theory, nonparametric statistics, and deep learning. I would also be interested to teach courses on the many subjects used within and adjacent to machine learning, including core probability and statistics, optimization, and standbys like algorithms, data structures, and introductory courses in programming and practical data analysis. I've also found the journal club at Gatsby extremely useful for expanding my knowledge of the field, and look forward to running similar reading groups in the future.

Compared to many academics, I also have strong practical skills and an interest in good software practices, as evidenced e.g. by my involvement as a core developer of the conda-forge software packaging ecosystem. As a "side effect" of course projects and assignments, I hope to help teach students good software habits, which will benefit them whether they end up as researchers, software engineers, or anyone else who regularly writes code.