Teaching Statement

Daniel Khashabi

December 2021

Teaching Philosophy

Throughout my journey, both as a student and as an instructor, I have experienced many teaching styles and have developed a taste for what I believe is an effective teaching experience. The three principles guiding my teaching style are as follows:

(1) Interaction is key to keeping students engaged. Incorporating an interactive exercise into each lecture could help students' understanding and give them a break to digest everything learned so far. One way to do this is to pose a thoughtful question and let the class discuss their reflections in small groups or with the whole class. Instructor-led activities have the potential to expose students' reasoning flaws and address their misconceptions. For example, to motivate recurrence in neural networks, I once asked, *"how would you feed variable-length sentences to a neural network?"* I was delighted to see many students were eager to discuss issues beyond what I foresaw. In some cases, students pointed out others' mistakes and helped them identify better solutions. I let the conversation go on for 5-10 minutes, before laying out my proposal. Finally, when explaining the answer, I tried to cover how it addressed the issues of the earlier ideas by the students.

(2) Abstraction of ideas to their high-level picture is important in every step of teaching. We all forget about details, but the "big picture" stays with us. Rigor matters. However, exclusive focus on it can drown us in the swamp of details and equations. Take the Backpropagation algorithm, for example, which is a key algorithmic tool in deep learning. Jumping into the details of this algorithm, without a proper build-up, could result in a dramatic failure as it contains many notations and seemingly complicated equations. A successful approach should first focus on building a solid intuition. My usual approach is to solidify the following three insights first: (a) chain rule from calculus, (b) gradient descent, and (c) dynamic programming. With these tools at hand, describing the Backpropagation algorithm *in plain English* is pretty simple. Once a student understands the algorithm at this level they know everything, even if they fail to follow the details or forget them.

(3) Repetition is a key learning mechanism to internalize ideas. Repetition should happen at every level of the curriculum. Each lecture should open and close with a reflection on the course trajectory and a brief look into the next sections so that they do not lose sight of the journey and put what they have learned into context. It is often helpful to reference ideas from the prior lectures and their connections with the newly presented material. For example, when introducing *decision tree* learning, one can relate and contrast it with prior ideas such as *gradient descent* learning of differentiable functions. The homework assignments should also serve as a way to remind students about what was discussed and help them apply the learned ideas to slightly novel settings. Rehearsing difficult new ideas over time gradually makes them easy and familiar.

Teaching Experience

Over the course of my recent career I have taken part in a variety of teaching experiences:

Classroom teaching — I have extensive experience as a teaching assistant (TA'ed for 7 classes) for a variety of subjects: Programming, Statistics, Signal Processing, Algorithms, and Machine Learning. In these roles, I helped design and grade homework assignments and exams. Classes like "Fundamental Algorithms" or "Machine Learning" also involved facilitating review sessions that went over the key learning items from the main class and discussed relevant problem sets. On several occasions, I **guest-lectured large classes** (>100 enrollment) at UIUC and UPenn.

Invited lectures — I have given technical lectures in various occasions (over 10 venues) as **invited** (e.g.,

Georgetown, Yale NLP seminars) and **conference talks** (e.g., IJCAI'16, AAAI'18).

Future Teaching Plans — I look forward to more teaching opportunities, particularly on topics such as natural language processing (NLP), artificial intelligence (AI), and machine learning (ML), that are foundational to my research. I would be interested in teaching these topics in either *elementary* or *advanced* level courses. Additionally, since AI now has a diverse audience one can form classes to fit different mindsets and aspirations. For example, a curriculum for those interested in the practical application of the existing technology, a curriculum for those who are research-minded and want to learn about the latest developments that might not have any immediate application yet, and so on.

Mentoring & Advising

During my **advising experience**, I had the opportunity to mentor over 10 undergraduate or junior graduate students with a variety of backgrounds. Collaborations with these students have led to publications at prestigious conferences such as EMNLP (2018, 2019, 2020), ACL (2020, 2021), and NAACL (2019), with several still under review.

In my current role as a postdoc at the Allen Institute for AI, I have continued my mentorship efforts. Over the course of this past year, I have mentored several interns: Tao Li, Jieyu Zhao, Swaroop Mishra on various projects at the intersection of NLP and machine ethics. Each of these projects typically involved starting from high-level research questions, brainstorming ways to turn them into pilot experiments, and forming a story tying together the key findings of the experimental observations.

Moving forward, I look forward to building a longer-term mentoring relationship with my students as they evolve to become more seasoned researchers. The experience of mentoring several students has taught me the necessity of adapting my advising style to meet the needs of each person. I believe that the best way to nurture the development of my students as researchers is by forming groups to work on problems as a team, since science is a social endeavor.

Outreach & Broad Participation

Teaching is our way of projecting our ideas to the outside world. I have fulfilled this duty in several ways:

Volunteering — I have dedicated time to teach basic AI or computer science to underprivileged middle-school or high school students to help **reduce the accessibility barrier** to science.

Community organizing — In several conferences, I have volunteered as a mentor (NAACL'19 & WiNLP'20) to help junior students from all over the world. This year, I am the volunteer chair at NAACL'22 where diversity is one of our key criteria in distributing the awards. In 2019, I co-chaired the ACL student workshop which provided a forum for student researchers interested in computational linguistic an excellent opportunity to present their work and receive valuable feedback from the international research community — especially those who did not have easy access to top-notch mentors, typically found in top schools.

Broad visibility — I have been fortunate to have my work covered by various news publications, including VentureBeat, GeekWire, and Vox. I believe in the necessity of projecting our work to the world.

Looking ahead, I hope to continue communicating computer science and AI to a broad non-technical audience. This serves not only as a way of showing the world what we have done but also helping alleviate the representation problem by making computer science more accessible.