

My interest in teaching and advising was sparked during my undergraduate experience at Harvey Mudd, a liberal arts college. I was fascinated by how professors skillfully taught complex material: leading through examples to make concepts feel concrete and personal, and then following up with equations and theory. Through interdisciplinary research, my advisors instilled on me a different class of skills: how to formulate research questions, tackle them, and present the results.

My goal as an educator is to guide every student towards this kind of **creative** and **critical** thinking. During my PhD at the University of Washington's Paul G. Allen School of Computer Science & Engineering, I have served in a variety of teaching and mentorship roles, and I am excited to continue these efforts as a professor.

Classroom experience. I performed a variety of teaching roles during my PhD:

- *Natural Language Processing (NLP) for all.* In 2018, I served as a teaching assistant for NLP for Continuing Education students. I graded homework and advised students through regular office hours. Many students hadn't been in a classroom in years; I helped guide them towards technical understanding and self-confidence.
- *Leveraging students' strengths.* In 2019, I served as a teaching assistant for graduate-level NLP. I graded assignments, tutored students, and advised them on their research projects – reaching them *at their level*. For instance, I guided a student in programming languages towards completing a project she was passionate about – a prediction task over bug reporting forums.
- *Outreach towards first-generation students.* During my PhD, I formally tutored several first-generation students who were new to CS, and struggling with introductory classes (like Data Structures). I enjoyed giving them support and mentorship early on, which I saw pay off throughout their studies.
- *Guest lectures and external outreach.* I have given a variety of guest lectures in UW's NLP and deep learning-related classes over the years – covering topics like new NLP models (Transformers) as well as on my own research (such as grounded NLP, language and vision, and on applications like detecting neural fake news). Beyond the UW classroom, I have given a number of invited talks, including two at Microsoft Research (about benchmarks for grounding, and about learning grounding through language and vision). I gave a number of invited talks about both the *technical* and *policy aspects* of our defense against neural fake news, including at DEFCON, at Stanford's AI for Social Good class, and at several DARPA summits.

Teaching philosophy. I want to make grounded NLP, and AI more broadly, accessible to everyone – thinking back to the educators who sparked my initial interest in the field. Towards this goal, I think a lot about presentation when creating teaching slides: I try to communicate the high-level picture first, building intuition through examples, before filling in the gaps with lower-level details. I plan to take an active learning approach – both inside the classroom through discussions, and outside the classroom through student-chosen projects [1].

I strive to create an inclusive environment, where all students can feel comfortable learning technical concepts. I will keep up-to-date with best practices for inclusive teaching, like encouraging everyone to participate without letting certain individuals dominate the conversation [2]. In AI courses, I will include modules on ethics and society, encouraging students to think broadly.

Teaching interests. I am interested in teaching a variety of CS classes, including on NLP, AI, and Machine Learning. Beyond these, I would be excited to teach new and interdisciplinary courses. My research interests would align to a new class about connecting natural language to vision, audio, and the world. I would pitch this class to AI students broadly, promoting cross-pollination between fields. Among others, I would also be interested in teaching or co-teaching classes on the societal implications of AI, approaching issues from a security or public policy lens.

Advising. I have had the good fortune of mentoring a variety of junior students in research

while at UW, including junior graduate students, undergraduates, and one high-schooler. A few students, who published their work, include:

- *Jeff Da (BS student)*. I helped Jeff define, and carry out, research fitting his interests in NLP and computer vision. The idea was to teach machines to understand *intent* behind manipulated images – for instance, a harmful ‘deepfake’ differs from an enhanced vacation photo. I gave feedback at all stages, including collecting data, building a model, and writing the paper. Jeff’s paper was accepted to ACL [3], and he is planning on continuing research as a PhD student.
- *Ximing Lu (BS student)*. I advised Ximing through several projects, around her diverse interests in commonsense natural language and computer vision. With my help, Ximing introduced a new approach where symbolic constraints guide deep natural language generators; this paper was accepted to NAACL [6]. With my guidance, Ximing has grown significantly as a creative and independent researcher. In a recent collaboration, her work enabled video-based commonsense models to transfer to single-image tasks [7]. She is currently applying to PhD programs.
- *Gabriel Ilharco (PhD Student)*. I advised Gabriel on one of his first graduate school research projects: probing the amount of visual knowledge learned by language models. I helped Gabriel crystallize this abstract research question into something concrete and testable, and gave feedback on experiments, analysis, and presentation. The work was accepted to NAACL [5].

It has been enjoyable watching my students grow. As a professor, I am excited to advise students on a longer timeline, helping to nurture the next generation of independent researchers.

Advising philosophy. Through my advising, my goal is to enable students to grow as creative researchers who can direct and execute their own research agenda. Since this goal is quite long-term, I try to break it up into manageable pieces to prevent it from overwhelming students. At the start of research, I like to work with students to co-design a research project that excites them. I find it helps to start small and concrete at first – along the way, the students’ insights and observations often form the foundation of a top-tier conference paper.

I also acknowledge that impactful research is risky, which can be difficult for students. I promote the idea of embracing and growing from failure, rather than seeing it as the end of the world [4]. I do this both by leading through example, as well as trying to reach students where they are, technically and mentally. As a professor, I am excited to carry out this philosophy while also keeping it up-to-date, with best practices learned from my colleagues and students themselves.

References

- [1] Scott Freeman et al. “Active learning increases student performance in science, engineering, and mathematics”. *Proceedings of the National Academy of Sciences* (2014).
- [2] Lecia J Barker, Melissa O’Neill, and Nida Kazim. “Framing classroom climate for student learning and retention in computer science”. *SIGCSE*. 2014.
- [3] Jeff Da, Maxwell Forbes, Rowan Zellers, Anthony Zheng, Jena D Hwang, Antoine Bosselut, and Yejin Choi. “Edited Media Understanding Frames: Reasoning About the Intent and Implications of Visual Misinformation”. *ACL*. 2021.
- [4] Carol S Dweck. *Mindset: The new psychology of success*. 2008.
- [5] Gabriel Ilharco, Rowan Zellers, Ali Farhadi, and Hannaneh Hajishirzi. “Probing Text Models for Common Ground with Visual Representations”. *NAACL*. 2020.
- [6] Ximing Lu, Peter West, Rowan Zellers, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. “NeuroLogic Decoding: Unsupervised Neural Text Generation with Predicate Logic Constraints”. *NAACL*. 2021.
- [7] Rowan Zellers, Ximing Lu, Jack Hessel, Youngjae Yu, Jae Sung Park, Jize Cao, Ali Farhadi, and Yejin Choi. “MERLOT: Multimodal Neural Script Knowledge Models”. *NeurIPS*. 2021.