

I am committed to fostering diversity and inclusion, which is of primary importance in educational and research environments. To do so, we must invite and allow everyone—regardless of background or identity—to be themselves, and feel supported in doing so. Without that, diversity is token, a message to the outside world but not an internalized truth. Being the first person in my family born in the United States, I am well aware of issues of equity—from my father getting frequently chosen for “randomized” searches at airports to people making fun of my mother’s accent and asking where she is “really from”. These issues are minor compared to the enormous obstacles many of my colleagues face, and those who might have been my colleagues but were confronted with unfair standards and legacies.

Higher education can be a powerful vehicle to reduce this inequality, and Computer Science is uniquely positioned to act as a ladder for social mobility. My doctoral institution—the University of Washington—did a great job in creating a helpful on-ramp and supportive environment. From the pre-application mentorship service that pro-actively helps students from all over the world apply to Ph.D. programs, to the Q++, the computer science school’s active and engaged LGBTQ student organization. The result is a positive feedback loop, in which the school becomes known as a welcoming community, and in turn attracts more diverse students and faculty, enriching our community. This shows up in hard statistics, for example in the latest computer science cohort 34% of undergraduates and 30% of graduate students identify as female, significantly higher than the national average. I would like to help emulate these efforts and bridge the gap at my next institution by participating in these kinds of organized efforts, especially ones that reach further up the pipeline to fix systemic issues.

I believe mentorship is the key leverage point for retaining diversity: establishing trusting relationships is the basis for real inclusion and the most effective barometer of potential issues. During my Ph.D, I mentored a diverse group of undergraduate students, Master’s students, and Ph.D. students—including women and first-generation immigrants who eventually enrolled in Ph.D. programs for Natural Language Processing at Johns Hopkins University and the Georgia Institute of Technology. I found that the most important gaps to bridge among incoming students are often not purely institutional, but cultural: although students gained formal access to the institution they aimed for, many still found themselves cut-off from informal channels that are critical to traversing a community and a career. Mentors generally shy away from mentioning tips about things as simple as how to politely set boundaries and expectations in a collaboration, even though these are critical skills. I offered informal mentoring, and many of these mentorship relationships have spanned my entire grad school experience. The natural mentorship position of a professor will give me an opportunity to use these experiences, and I will seek out opportunities to be a resource to undergraduate and graduate students from different backgrounds, underprivileged circumstances, and whose identities or perceived identities cause bias to stand in their way.

Relevant Research My research has investigated some of these issues, for instance my first paper measured bias in how people of different genders are portrayed in movie scripts, an experience that helped me realize how inundated we are with biases that are hard to even see let alone remedy. In recent work, I examined how we can train NLP models so that bias removal is easier than it is with current, monolithic neural models: we factor different data distributions into parallel networks in the model, letting us remove specific distributions that we purposefully train to capture undesirable parts of the data (e.g. hate speech). More generally, in text generation there is a commonly assumed but fictional “neutral point of view” that generated text is expected to stylistically mimic, similar to how the stereotypical Los Angeles accent is often perceived as “neutral” due to its place in media. I intend to pursue future research in dismantling such fictions and instead ask: Who are text generators mimicking? And how can we capture more diverse expression rather than collapsing onto a perceived-majority voice?