

To me, teaching and advising is equally enjoyable as doing cutting-edge research. One of the most important reasons for applying for faculty is that I'm confident that I'll be a good teacher and mentor. My confidence and passion is gradually built during my career from experience, acknowledgements from my advisors or mentees, and self-discoveries. In the meantime, I have formed philosophy that I find to benefit my students the most. I detail them below.

Advising and Mentoring During my graduate career as a master's student at Shanghai Jiao Tong University (SJTU), a PhD student at Massachusetts Institute of Technology (MIT), and a postdoc at the University of Washington (UW), I am fortunate to have had the privilege to closely mentor 6 undergrad students, 3 master's students and 1 PhD student. Most of our collaborations have resulted in at least one paper published, and in three of the published papers [1–3] I am recognized as the last author. I highlight some experiences below:

Jingyu (Jack) Zhang, undergrad: Jack was an undergrad in Johns Hopkins University (JHU). I have been mentoring Jack for around 2 years. Our first project resulted in a paper [2] on controllable generation, which got the best paper award in the 2022 NeurIPS-ENLSP workshop. We then collaborated (co-first-author) on a second project on language generation evaluation metrics [4], which got accepted by ACL2023. We both think that project is one of our favorite projects. Thanks to Jack's strong recommendation, I got contacted by another JHU undergrad, Abe Hou. The three of us worked on a semantic watermark project [5], which is currently under submission. When Jack was applying for PhD, he told me: "*Before I met you, I was not sure whether I wanted to do a PhD. Working with you, I felt the genuine joy and excitement from doing cutting-edge research. And now I'm determined.*" Jack is now a PhD student in JHU.

Xiao (Sophia) Pu, undergrad: Sophia is a female undergrad in Peking University who I have been mentoring for 6 months, and will continue to work with in the near future. Under my guidance, her first first-author paper on generalization of machine-generated text detectors has been accepted by EMNLP-finding [3]. In light of her improved research skills during that collaboration, I encouraged her to propose her own project. We spent the next two months brainstorming on new research ideas. Sophia did a lot of literature search and presented multiple new ideas. I gave feedback by analyzing its feasibility and significance. In the end, we reached consensus on an on-going project about emergent communication. It was a very rewarding experience to me as I witnessed my mentee gradually become an independent researcher. Sophia is currently applying for PhD.

Advising Philosophy: I want to empower my students and help them realize their research and career goals. Below I expand on three aspects: research skills, connections, and mental health.

- **Research Skills:** It is important for a junior researcher to gradually form good research taste and technical skill-sets. I exchange outstanding papers with my mentees, discuss the reasons behind their success, and encourage them to keep an eye on what's going on in other fields. During meetings, I help them check the soundness of project motivations and logical flows. Being a technical researcher myself, I conduct code/derivation reviews with my mentees and encourage them to design various tests. When they have acquired enough experience, I urge them to form their own directions and gradually become a fully independent researcher.
- **Connections:** Depending on my mentees' skill level and research topic, I connect them with professors or experts I know so that they can get a wider range of advice or acknowledgements. For example, Sophia once proposed an idea that is out of my expertise, and I connected her with a HCI expert in UW. I also encourage them to form inside or outside collaborations. For example, Jack helped significantly in Abe's and Sophia's projects [5, 3].

- **Mental Health:** Doing cutting-edge research can sometimes be very frustrating. I shared a lot of my own failure stories to my mentees to help them cope with theirs. Every time I observe my mentee's improvement, I praise them with huge excitement. When they make any mistake, I gently point it out and encourage them to improve.

Teaching I have been teaching since I was an undergrad. So far, I have done 6 times of teaching assistant (TA), 7 guest lectures, and numerous recitation/tutorial/project sessions. Among the 6 times of TA, I was a “head TA” twice. In that position (which is a special tradition from the SJTU ACM Class), I led a group of TAs to compose and teach a full semester-long programming course consisting of tutorials, exercises, and final projects. Below I highlight several experiences which also demonstrate my teaching philosophy.

Guest Lectures in UW 447+547 (NLP): My postdoc advisor (Prof. Yulia Tsvetkov) invited me to give two lectures on basics of neural network language models. I spent an entire week working on the slides. In these two lectures, I went from logistic regression all the way to transformer language models, with math, intuitions, connections and exercises. Knowing that the content would be challenging because this was the first time for students to encounter neural networks, I pause from time to time for review and tell students what are the most important messages to remember. Although the lectures were dense, I was extremely happy to see a lot of questions and students correctly answering the quiz I prepared. One student told me she really liked the organization and logic flow in the lectures, which helped her understand how this field had evolved. ¹

Guest Lectures for SJTU ACM Class: For 4 consecutive years, I was invited by Prof. Yong Yu to give lectures in a course named “Advanced Research Forum”. Around 12 professors from top universities (such as Stanford, CMU, UCSD, etc.) are invited to this lecture series, and I was the only one invited who was not (yet) faculty (confirmed by Prof. Yu). In the lectures, I talked about my research projects and stories (including failures), and gave advice on how to begin to conduct cutting-edge research. I revised my slides significantly every year to integrate my new research insights or lessons from my experience.

TA for MIT 6.864 (Advanced Natural Language Processing) 2020: As TA for MIT-6.864, I contributed with huge passion to almost all aspects of the course building and teaching: I gave one guest lecture on advanced language modelling, composed one homework set, did three one-hour interactive recitation sessions with well-prepared slides and exercises, contributed problems to the mid-term exam, etc. I went to every lecture and pondered on what additional material I could give in the recitations to aid the students' understanding. I responded timely to more than half of questions posted by students on piazza. One day there was an office hour and accidentally no other TA was available. Despite the fact that I just gave two recitation sessions the day before, I still volunteered and showed up in the office hour. I received sincere acknowledgements from the professors for my contributions.

Teaching Plan Due to my many TA experiences, I am experienced and willing to teach classical undergrad-level computer science courses such as programming, data structures, and algorithms. Given my research background, I am interested in teaching modern learning-related courses, including Machine Learning, Deep Learning, NLP, at either undergraduate or graduate level. Aligned with my specific research interests, I also look forward to running or leading advanced seminar classes or meet-ups on topics such as responsible/trustworthy Generative AI, interactive Agents, or efficiency in NLP in collaboration with other faculty.

¹Slides are available at https://cloudygoose.github.io/slides/uw_nlp_lecture.zip.

References

- [1] Jiabao Ji, Yoon Kim, James Glass, and **Tianxing He**. Controlling the focus of pretrained language generation models. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio, editors, *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3291–3306, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.findings-acl.260. URL <https://aclanthology.org/2022.findings-acl.260>.
- [2] Jingyu Zhang, James Glass, and **Tianxing He**. PCFG-based natural language interface improves generalization for controlled text generation. In *Proceedings of the 12th Joint Conference on Lexical and Computational Semantics (*SEM 2023)*, pages 295–313, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.starsem-1.27. URL <https://aclanthology.org/2023.starsem-1.27>.
- [3] Xiao Pu, Jingyu Zhang, Xiaochuang Han, Yulia Tsvetkov, and **Tianxing He**. On the zero-shot generalization of machine-generated text detectors. In *Findings of The 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP-Finding 2023)*, 2023.
- [4] **Tianxing He***, Jingyu Zhang*, Tianle Wang, Sachin Kumar, Kyunghyun Cho, James Glass, and Yulia Tsvetkov. On the blind spots of model-based evaluation metrics for text generation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12067–12097, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.674. URL <https://aclanthology.org/2023.acl-long.674>.
- [5] **Tianxing He***, Abe Bohan Hou*, Jingyu Zhang*, Yichen Wang, Yung-Sung Chuang, Hongwei Wang, Lingfeng Shen, Benjamin Van Durme, Daniel Khashabi, and Yulia Tsvetkov. Semstamp: A semantic watermark with paraphrastic robustness for text generation, 2023.