

Teaching Statement – Dongyeop Kang

Teaching experience. During my graduate studies at CMU, I worked as a teaching assistant (TA) for the course *Machine Translation and Sequence to Sequence Models* taught by Prof. Neubig. Because the course was taught by the professor for the first time, the two TAs myself and one other, helped assisted him in effectively organizing the class. We gave detailed feedback on all class materials [5] before the class began, assembled quiz questions of each lecture, implemented reference codes for assignment algorithms, and addressed questions from students in both the online platform and during office hours.

As a senior Ph.D. student, I worked as one of the two main TAs for the course *Computational Semantics for NLP* taught by Prof. Hovy and Prof. Mitamura. In addition to grading and office hours, I took on more responsibilities, such as giving a lecture on “Distributional Compositionality and Logics”, which included my latest research [7,9]. Over the semester I regularly consulted with eight groups of students to help them brainstorm research ideas for course projects and gave feedback on their intermediate and final reports. Half of the groups successfully published their works at top NLP venues. I was also invited to the same class next year, giving an hour lecture on “Relational Semantics on Word Embedding”.

As well as the class activities, I actively participated giving talks about my research at conference presentations such as ACL, EMNLP, and NAACL, lab meetings, internship companies, and other small seminars at CMU. After I graduated from CMU, I was invited to give talks at the CMU LTI seminar, Berkeley AI Research (BAIR) workshop, and colloquiums at KAIST, GIST, and POSTECH in Korea.

Teaching interests. Computer science, particularly natural language processing, is one of fastest-growing fields. Half of the technical knowledge that I learned at CMU was developed within the last few decades. Therefore, teaching for NLP classes is not just for conveying knowledge in textbooks but teaching students how the field of NLP has evolved and motivating them to be a part of it. Another challenge of teaching NLP would be providing the right direction to students. The field of NLP has been often hyped by public attention and recently trapped into chasing the leaderboard competitions. The leaderboards or shared task competitions are great practices for junior students to prototype what they have learned from their classes. However, they should not be the long-term research goals for students. Instead, I like to teach students to deeply understand the nature of the problems, rather than to chase the scores only.

I am interested in building active classes. In particular, I have two proposals designed to make my future classes more interactive: First, I like to build a class where a group of students extensively interact with each other about a specific topic and learn about other topics from the other groups. Based on my experience, I do believe teaching is always the best way to learn, and peers are the best teachers. Secondly, I like to make my lectures and course project outputs from students publicly available, if the department permits. This publicity will motivate me to maintain the quality of my lectures and interact with researchers and students outside of the department. I believe this openness will be helpful to students outside of the department who are looking for resources, particularly in this pandemic situation.

I plan to build from general to focused courses in NLP. The course “Introduction to NLP” would be the general one that can be taken by both undergraduate and graduate students, and students outside of the CS department. For this type of introductory class, I will cover as many broad topics as I can, including algorithms, tasks, and their applications in NLP. NLP as a secondary field of linguistics, I would like to develop more computational linguistics (CL) focused classes such as “Computational Semantics for NLP”. This class will cover more classical but seminal linguistic theories and build assignments to replicate those theoretical concepts with modern machine learning techniques. Lastly, I plan to make a very specialized class focusing on my main research field: “Natural Language Generation”. This class will cover the broader aspects of language generation, from theories (e.g., rhetorical structure theory, script theory), to classical and modern techniques, to real-world applications where NLG models are applied.

Advising. I enjoy mentoring junior students and helping them find the right problem to solve and define the exact scope of it. During my graduate and postdoc studies, I have been formally and informally

advising several graduate students, such as Taehee Jung at the University of Pittsburgh (UPitt), Shirley Hayati at the University of Pennsylvania (UPenn), Vivek Aithal at the UC Berkeley, Steven Feng at CMU, Hwiyeol Jo at Seoul National University, and other students who recently contacted me. They all have different interests and backgrounds in NLP, machine learning, social science, or statistics. Most of the students that I have advised has published or submitted to peer-reviewed NLP conference papers [1–4, 8].

The first student I advised is Taehee Jung, a Statistics Ph.D. student at UPitt. Despite her interest in NLP, she didn't have any CS background. I taught her the modern NLP techniques weekly and took the introductory NLP class at CMU together. The first project I gave her was to develop a text summarization system, by maximizing the volume size of convex hull summary sentences in continuous space. Over the two months of experiments, however, it turned out that our volume maximization hypothesis didn't have a strong correlation with the summarization performance, even though the diversity of summaries does. She was quite depressed by the negative result, but I taught her how joyful this scientific investigation process is, and that the negative findings are useful not only for other researchers not to follow the same path, but also for ourselves guiding to a positive direction. Based on this experience, we further explored other aspects i.e., diversity, importance, position, and empirically showed how current summarization corpora and systems are biased to each aspect. Surprisingly, we could complete and submit this new research motivated by the negative result in *two weeks*. This work was presented at EMNLP 2019 [8]. This advising experience made me think that the research goals given to students should be the progressive exploration of long-term scientific questions, not only paper acceptance. As the intersection of statistics and NLP, I helped her find novel ways of optimizing MLE training with posterior calibration objectives in NLP tasks. This work was presented at ACL 2020 [4], and she is currently extending this idea in her dissertation research where I am still a part of it.

I advised another graduate student, Shirley Hayati as well. I helped her extend our recommendation dialogue system [6] to be more sociable, particularly focusing on designing the annotation guideline based on social science theories. This work was accepted to EMNLP 2020 [2] with very high review scores. While I worked with Shirley, she had the difficulty of being rejected from her paper submissions and Ph.D. admissions. In our regular meetings, I tried to encourage her to overcome it and told her my own failure stories of my academic years as a reference. Besides the mental support, I supported her in practical ways: planning together what she can do for next Ph.D. admissions, and writing recommendation letters for her job applications and Ph.D. applications. She has currently joined UPenn as a research assistant and has applied for Ph.D. programs for next year. Throughout the advising experience, I realize that being a good advisor is not only writing strong papers together but caring for students' mental and emotional status and supporting their future careers.

Generally, I help students set a feasible research goal, form hypothetical questions they enjoy exploring, validate the hypotheses either theoretically or empirically, convince others through communications, and iterate these steps until they come up with a complete research paper.

An important challenge of advising is to establish effective communications with students. Everyone has their own preferred working styles. For instance, some students prefer not having too specific feedback nor being micro-managed, others do. My advising style would be flexible, varying according to students' preferences and personalities. I often have regular meetings with students, particularly in the early stage of the project. However, as the project becomes stable and students are fully engaged, I would suggest that students take the full responsibility of making all technical decisions and taking the lead. My ideal goal of advising is to help students become independent scientists and leaders in the field.

When I advise students, I do my best to be an empathetic listener. As a former student, I do know how difficult it is to survive the graduate programs, to deal with rejections and pressures from peers, to make a balance between life and work, and to keep myself mentally healthy. I like to be a mentor to whom students are willing to share their concerns for the first time.

Everyone has different talents, and their potentials appear at different times. When students didn't make any research progress, I often have a casual chat or postpone the meeting until they are well prepared. I am patient, waiting for students to develop their maximum research capacities at their own pace. I saw many students and colleagues who are much more knowledgeable than me and learning from them

was the biggest joy of my research. I would not force students to go where I expect them to go, but let themselves decide where to go and encourage them to do so.

Academic Services. I enjoy reviewing papers and giving my feedback to anyone for their research. Since I have received helpful comments on the papers I submitted from several anonymous reviewers, I set a personal goal of contributing back to the community by reviewing papers more than the two times of the papers I submitted. From the beginning of my graduate study, every year I have reviewed papers at NLP venues such as ACL, EMNLP, and NAACL, machine learning venues such as ICML, NeurIPS, and ICLR, and many other workshops and journals. I was fortunate to be recognized as one of the top reviewers from both ACL 2018 and NeurIPS 2018.

References

- [1] Steven Y. Feng, Varun Gangal, **Dongyeop Kang**, Teruko Mitamura, and Eduard Hovy. Genaug: Data augmentation for finetuning text generators. In *Deep Learning Inside Out (DeeLIO) Workshop at EMNLP 2020*, Online, November 2020.
- [2] Shirley A. Hayati, **Dongyeop Kang**, Qingxiaoyang Zhu, Weiyan Shi, and Zhou Yu. Inspired: Toward social recommendation dialog systems. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Online, November 2020.
- [3] Hwiyeol Jo, **Dongyeop Kang**, Byoung-Tak Zhang, and Marti Hearst. Does bert know this word’s definition? a comparative study of contextual embedding representations and formal dictionary definitions. In *Submission*, 2020.
- [4] Taehee Jung, **Dongyeop Kang**, Hua Cheng, Lucas Mentch, and Thomas Schaaf. Posterior calibrated training on sentence classification tasks. In *2020 Annual Conference of the Association for Computational Linguistics (ACL)*, 2020.
- [5] Graham Neubig. Neural machine translation and sequence-to-sequence models: A tutorial. *CoRR*, abs/1703.01619, 2017.
- [6] **Dongyeop Kang**, Anusha Balakrishnan, Pararth Shah, Paul Crook, Y-Lan Boureau, and Jason Weston. Recommendation as a communication game: Self-supervised bot-play for goal-oriented dialogue. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Hong Kong, November 2019.
- [7] **Dongyeop Kang**^{*}, Taehee Jung^{*}, and Eduard Hovy. Don’t treat every relation the same! constraining word vectors to geometric regularities of inter-word relations. 2020.
- [8] **Dongyeop Kang**^{*}, Taehee Jung^{*}, Lucas Mentch, and Eduard Hovy. Earlier isn’t always better: Sub-aspect analysis on corpus and system biases in summarization. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Hong Kong, November 2019.
- [9] **Dongyeop Kang**, Tushar Khot, Ashish Sabharwal, and Eduard Hovy. Adventure: Adversarial training for textual entailment with knowledge-guided examples. In *The 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, Melbourne, Australia, July 2018.