

# Teaching Statement

Ting-Hao (Kenneth) Huang

## Teaching Philosophy

Communication lies at the core of my teaching and mentoring philosophy. I believe good communication can build a transparent and inclusive environment for students to learn and grow. Through my decades of experience as an amateur fiction writer and a contestant of impromptu speech competitions, I learned to deliver complex concepts to people from a variety of backgrounds. Telling a story is similar to teaching: good teachers, like good storytellers, must engage audiences emotionally, challenge audiences intellectually, and deliver the core message to audiences effectively. To accomplish these goals, I apply the lessons I have learned from narrative writing in my teaching.

- **Set Expectations Transparently.** Good stories balance expectations and surprises. Thus, I express my expectations to my students, and have my students express their expectations to me.
- **Learn to Argue Constructively.** Critique helps writers improve. Learning to express thoughts constructively challenges both me and my students to put ourselves in someone else's shoes, from which we can begin to consider alternative perspectives.
- **Demonstrate Support & Confidence Explicitly.** Emotion is essential for a good narrative or speech. Academic work comes with challenges and, inevitably, frustrations. As a teacher, I believe showing support to students is essential to helping them grow as learners.
- **Have Conversations Constantly.** As Stone, Heen, and Patton write in *Difficult Conversations* [1], things that are hard to talk about are often things that are important. I believe it is essential to encourage students to ask hard questions, to ask for help, or even to ask for forgiveness.

## Teaching Experience

At Carnegie Mellon University, I served as Teaching Assistant for Prof. Carolyn Rosé's Applied Machine Learning (11-663 Fall 2016), for which I hosted weekly office hours and led recitations to answer students' questions. This class was designed for students with limited experience in computer science, and thus I learned to break down the concepts of classic machine-learning algorithms into small intuitions. For example, I developed a lesson on Decision Tree where I asked students to imagine they are zookeepers who try to distinguish mammals from other animals in the zoo by using a set of features; I also served as Teaching Assistant for Dr. Ralf Brown and Prof. Florian Metze's Language Technologies Institute Colloquium (11-700 Fall & Spring 2014). My role in this seminar course was to coordinate between speakers, students, and supporting staff members. I further volunteered to recreate the website and built the first YouTube channel for this course. Currently, I am co-teaching Crowd Computing (05-499 Fall 2017) with Prof. Jeffrey Bigham. In this graduate-level course, I am responsible for lectures about real-time crowdsourcing and crowd-powered systems, which has prepared me to teach this topic in the future.

## Mentoring Experiences

I have been fortunate to mentor many talented students in multiple research projects. At National Taiwan University, I co-mentored a master's student, Ho-Cheng Yu, in the NLP Lab with Prof. Hsin-Hsi Chen, and led a project of opinion mining in news articles. In this experience, I learned how to set goals and develop a research idea iteratively with students who are passionate but have less research experience. I also witnessed how a student with a good work ethic can achieve great, difficult goals. This experience was so positive that we continued to collaborate remotely and held weekly virtual meetings after I came to the U.S. in 2011. Together, we published two conference papers (ROCLING 2012, COLING 2012) and one

journal paper (IJCLCLP). This valuable experience inspired me to pursue a career as a professor, since seeing a student learning and growing is so rewarding.

At Carnegie Mellon University, I have also mentored students across multiple research projects. During the early stages of deploying the crowd-powered conversational assistant Chorus, I led a weekly brainstorming sessions with undergraduate students to create prototypes of the system, which later inspired many aspects of the actual implementation. In 2017, I co-mentored two undergraduate students, Katia Villevald and Jennifer Lee, in CMU’s summer “Research Experience for Undergraduates” program (REU), with Prof. Jeffrey Bigham and visiting Prof. Saiph Savage. In this program, I guided them to create automatic text annotators from scratch, starting with developing the data annotation schema, building the interface for crowd workers to annotate text, and eventually using machine-learning tools to train supervised-learning annotators using the labelled data.

Moreover, in early 2016, I began collaborating with Dr. Lun-Wei Ku and her NLP/SA Lab at Academia Sinica, Taiwan. During weekly meeting with students in the NLP/SA lab, I introduced the perspective of HCI to Dr. Ku’s group, whose focus was originally only on natural language processing. The students learned to use their technical skills to explore HCI problems. I guided them to design the protocol for user study, to recruit appropriate users, and eventually to conduct the first user study in their lab. To date, we have published four papers at EMNLP Demo, AAAI UX-ML symposium, COLING Demo, and LREC.

Chorus Paper Prototype  
Jason Chen  
(2015)

Pollyanna Phenomena  
(COLING Demo 2012)  
Ho-Cheng Yu

**Finding the Ideal End of a Conversation For a Crowd-Powered Conversational Agent**  
Jennifer Lee, George Melnik, Kenneth Heun, Saiph Savage

**Abstract**  
We developed a machine learning model for a crowd-powered chatbot, Chorus, that can detect when the user has nothing else to ask of the crowd workers, for model uses (C) to identify the 'ideal' end of the conversation (EOC).

**Method**  
1. Defined 'ideal' end of a conversation (EOC)  
2. Created database: collect 643 EOC labels from MTurk workers  
3. Experiment: Conditional Random Fields (CRF) with word anchors  
4. Using OffRate, a CRF model is trained on our collected database that are critical to finding the ideal EOC. Ideal workers are a critical attribute for identifying EOC.

**Results**  
- Mean F1 score: 0.82  
- Mean AUC: 0.92  
- Mean Precision: 0.85  
- Mean Recall: 0.79  
- Mean Accuracy: 0.81

**Conclusion**  
- Mean F1 score: 0.82  
- Mean AUC: 0.92  
- Mean Precision: 0.85  
- Mean Recall: 0.79  
- Mean Accuracy: 0.81

Detecting the end of a conversation Jennifer Lee (REU 2017)

**Identifying requests in conversations**  
Katia Villevald

**Abstract**  
We developed a machine learning model for a crowd-powered chatbot, Chorus, that can detect when the user has nothing else to ask of the crowd workers, for model uses (C) to identify the 'ideal' end of the conversation (EOC).

**Method**  
1. Defined 'ideal' end of a conversation (EOC)  
2. Created database: collect 643 EOC labels from MTurk workers  
3. Experiment: Conditional Random Fields (CRF) with word anchors  
4. Using OffRate, a CRF model is trained on our collected database that are critical to finding the ideal EOC. Ideal workers are a critical attribute for identifying EOC.

**Results**  
- Mean F1 score: 0.82  
- Mean AUC: 0.92  
- Mean Precision: 0.85  
- Mean Recall: 0.79  
- Mean Accuracy: 0.81

**Conclusion**  
- Mean F1 score: 0.82  
- Mean AUC: 0.92  
- Mean Precision: 0.85  
- Mean Recall: 0.79  
- Mean Accuracy: 0.81

Identifying requests in conversations Katia Villevald (REU 2017)

## Example Courses

- **Crowd Computing:** An introductory course on crowdsourcing, crowd workflows, and crowd-powered systems that focuses on classic techniques, mechanisms, and systems introduced in human computation and their broader impacts.
- **Crowdsourcing & Natural Language Processing:** A project-based course that is designed to introduce the common practice of crowdsourcing to students with an NLP or ML background. The goal is to teach graduate students to use crowd technologies effectively and ethically.
- **Crowd-AI Systems:** An advanced class that focuses on the challenges and common practices, such as the multi-arm bandit problem and reinforcement learning, of creating fast-paced Crowd-AI systems.

## References

[1] Douglas Stone, Sheila Heen, and Bruce Patton. 2010. *Difficult conversations: How to discuss what matters most*. Penguin.