ZHIYAO DUAN TEACHING STATEMENT

The main reason that I joined the faculty of a major research university after completing my PhD is that it allows me to be a researcher and an educator at the same time. It provides me the opportunity to introduce the latest research directly into the classroom, to observe how research concepts shape our students' minds, and to get inspiration from them. Through the integration of research and teaching, both are enriched.

My teaching philosophy can be summarized as three points:

- Foster passion: Interest is the best teacher. Fostering passion, in my view, is the most important goal of all teaching activities.
- Promote active learning: I deeply agree with the saying of "the purpose of teaching is not to teach". Students' active learning is the ultimate goal of education.
- Tailor to students' aptitude: Each student is unique. They come with different backgrounds and have different strengths and weaknesses. A good teacher should interact with students, pay attention to their feedback, and adjust teaching/learning agenda to their aptitude.

Since coming to UofR in 2013, I have been mainly teaching two courses: AME/ECE 277/477 - Computer Audition (Fall, graduate-level) and AME/ECE 272/472 - Audio Signal Processing (Spring, AME core course). I was/am the academic advisor of AME Class of 2018 and Class of 2022. I have mentored more than 20 undergraduate students on research, many of whom were supported by NSF REU supplements, REU Site, Xerox Engineering Fellows, McNair Fellows, and UofR's Discover Grant.

In addition, I have given guest lectures in Computer Vision and Human Computer Interaction. I organized two reading groups beyond my own lab: Deep Learning (Summer 2016) and Reinforcement Learning (Summer 2019). I also taught a mini course multiple times to high school students to attract them to STEM education. The remainder of the document highlights some of these achievements.

AME/ECE 277/477 – Computer Audition

(Course website: http://www.ece.rochester.edu/~zduan/teaching/ece477/index.html)

I created this graduate course upon my arrival at UofR in 2013 with the aim of training students to do research in this field. In the first half of the semester, this course covers fundamental concepts including short time Fourier transform (STFT), auditory models, audio features (pitch, rhythm, timbre) and audio modeling techniques (non-negative matrix factorization, hidden Markov models, deep neural networks). Students work on five programming assignments to practice these concepts. In the second half of the semester, the course covers state-of-the-art research topics in the field, including multi-pitch analysis, source separation, source localization, instrument identification, speaker recognition, etc. Students work on a final project that allows them to **experience the full cycle of research**, from project proposal, to algorithm implementation, poster presentation, paper writing and peer reviewing.

Critical thinking is one of the most important qualities of researchers. I designed **the following assignments to train students to think critically**:

- Reviewing papers: Each student is required to write a 1-page review for each of seven research papers that they choose from the course reading list, focusing on providing constructive feedback on the weaknesses of the paper.
- Reviewing peer's work: Each student needs to write a 1-page review of three other students' final projects, focusing on providing constructive feedback for project authors to improve.
- Presenting literature review: Students need to present a literature review of a research topic that they choose, focusing on expressing their own opinions.

This course has proven successful. The enrollment has been continuously growing from 7 in 2013 to 36 in 2019. Students reported that they learned a lot about this field as well as conducting research from this course. Students' final projects covered a wide range of topics, from source localization to music transcription. Several projects developed into independent research projects after the course, which eventually resulted in peer-reviewed conference papers:

- Bochen Li and Zhiyao Duan, "Score following for piano performances with sustain-pedal effects," in *Proc. International Society for Music Information Retrieval Conference (ISMIR)*, 2015, pp. 469-475.
- Yichi Zhang and Zhiyao Duan, "Retrieving sounds by vocal imitation recognition," in *Proc. IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, 2015.
- Yujia Yan, Ethan Lustig, Joseph Vaderstel, and Zhiyao Duan, "Part-invariant model for music generation and harmonization," in *Proc. International Society for Music Information Retrieval Conference (ISMIR)*, 2018.

The average course rating over the past offerings is 4.65 out of 5.0. Some selected student comments are:

- "Must take for those interested in audio"
- "I have take over 60 courses during my academic years, and this is the course i feel more proud of finishing it. Very time demanding but very rewarding if you finish it."
- "Very current research is being taught and the work on reading is really important for grad students."
- "This instructor explains very clearly. He seems very nice and approachable and he is interested in his students."
- "Professor Duan is very intelligent and has good lectures."

AME/ECE 272/472 – Audio Signal Processing

(Course website: http://www.ece.rochester.edu/~zduan/teaching/ece472/index.html)

This is a core required course in the AME program and the ECE master's program in the concentration of musical acoustics and signal processing. Prior to my arrival, this course was taught by Prof. Mark Bocko. After I took it over, I have made significant revisions. This course is a survey of audio digital signal processing fundamentals and applications. Topics include sampling and quantization, analog to digital converters, time and frequency domains, spectral analysis, vocoder, digital filters, audio effects, music audio analysis and synthesis, and other advanced topics in audio signal processing.

The main idea behind the course design is to **build intimate connections between theory and practice**; this is based on my belief that practicing is critical in mastering engineering concepts. These connections are realized through 8 programming assignments (5 using MATLAB and 3 using C++ on DSP boards) practicing theories covered in lectures. Finally, students also complete a group project on a topic that they choose. They need to write a project proposal, update with me their project status, write a final paper and make a poster presentation. The projects covered a wide range of topics in audio signal processing, from analysis to synthesis, from offline to real time.

The enrollment of the course also increased continuously from 15 in 2014 to 35 in 2019. The average course rating is 4.28 out of 5.0. Some selected student comments are:

- "Professor Duan is brilliant and you can tell he enjoys what he teaches."
- "Amazing professor! DSP board knowledge is a great job skill."
- "This course really focuses on experiments. I think it is good for the students, since it is more practical to really do it rather than just studying the concepts."

- "Must-take course for audio concentration. Enjoy your pain and come back stronger."
- "Prof. Duan was always prepared and his lectures were extremely well organized."
- "Duan is extremely smart and talented. He has sparked my interested in digital signal processing and I plan to pursue that in the coming years. Also the jokes made the year!"
- "He is really an expert in this field. I was so happy that I met him and took his course. It helped me a lot to improve myself as an engineer."

Mentoring Student Research

One advantage that students at a research university have is the opportunity to get involved in research. Over the past years, I have mentored more than 20 undergraduate students on independent research projects. Through the **Xerox Engineering Fellows** program, I mentored ECE undergraduates, Haowen Pan (summer 2014) on analyzing how pop music evolved over the past 50 years, Karan Vombatkere on developing an automatic lyrics display system for choral concerts, and Yiting Zhang on developing a sound search engine using vocal imitation as queries. In particular, the automatic lyrics display project was further extended by my PhD students and the system was successfully deployed in annual concerts of the Chinese Choral Society of Rochester in 2017-19. Through **NSF REU supplements**, I mentored our ECE undergraduates Ryan Bhular and Ayumi Yuasa on improving query-by-example sound search. Through **NSF REU site** on Computational Methods for Music, Media and Minds, I mentored our ECE undergraduate Arlen Fan and U Central Florida undergraduate Andrew Smith on building a web-based interactive music transcription system in the summer 2017. Through the same program, I mentored U Kentucky undergraduate Daniel Dopp on music harmonization and Marlboro College undergraduate Nick Creel on score following for augmented concerts in summer 2019. Through the **UofR Discover Grant** for Undergraduate Research, I mentored ECE undergraduate Steven Belitzky and CS undergraduate David Fink and Jovan-Gianni Lee. Through **research credits**, I mentored ECE undergraduates Yukun Chen and Meixiao Han. I have also mentored a number of self-funded visiting students from China and Kazakhstan, including Almas Abdibayev, Zhihan Zhou, Junyi Fan, Peizhe Gao, Hangyu Li, Yinghao Ma, and Mingrui Yuan. Many of the abovementioned students went on to pursue MS and PhD degrees in our department and other research universities including Indiana U, Georgia Tech, UCSD, Northwestern, and Dartmouth.

I have also mentored a large number of MS students on various projects. In particular, I supervised three MS theses: Andrew Trahan (2014) on drum kit transcription, Jonathan Downing (2016) on joint source separation and dereverberation, and Xinzhao Liu (2016) on audio-visual music performance analysis. Many of the MS students who worked in my lab joined high-tech companies after graduation, including Amazon, Bose, Apple, Harman, Cisco and Knowles. Some of them went on to pursue a PhD degree in research universities including Ohio State U, TU Delft, and RPI.

My main effort on student mentoring has been spent on PhD students. I have graduated two PhD students: Andrea Cogliati (Dec. 2017) and Emre Eskimez (Aug. 2019, co-advised with Prof. Wendi Heinzelman). Two more students, Bochen Li and Yichi Zhang, will graduate in Dec. 2019. Including these two soon-to-graduate students, there are eight PhD students in my lab at this moment (two 6th year, one 3rd year, two 2nd year, and three 1st year). I also hosted one visiting PhD student, Rui Lu, in 2017-18 and will host two visiting PhD students, Fei Jiang and Nan Jiang, in 2019-20.

Managing a large lab requires not only effort but also skills. I found that a "semi-supervised" way is the most effective. For PhD students, I set clear expectations at different stages of their PhD but also give them the maximal freedom to explore within the bounds set by funding sources. For example, I expect students to tackle a small-scale well-defined problem and submit a conference paper in the first year, dig deeper and define the thesis topic in the second and third years, and graduate with contributions summarized in three journal papers in the fifth/sixth year. I schedule 1-hour long weekly individual meetings with each student to track their progress and provide advice, but I never check whether or when they come to the lab each day. To promote the research atmosphere of the lab, I schedule a weekly lab

meeting for students to take turns to present their work or interesting papers. For mentoring MS and undergraduate students, I often use my PhD students as a proxy. This protects my own time, gains PhD students mentoring experiences, and ensures MS/undergraduate students to receive timely guidance. I think my lab management strategy and mentoring skills are effective and my lab is making steady progress each week.

Educational Outreach

I recognize the importance of exposing K-12 students STEM concepts to the society, and I enjoy spending some time on educational outreach. Through the Kearns center's UpwardBound program, I have taught a 4-lecture mini-course on "Music & Math" to high school students three times, in the summers of 2016, 2017 and 2019. This course covers topics on 1) Fundamentals of sound and vibration: basics of sinusoids; 2) Colors of musical instrumental sounds: the composition of harmonically related sinusoids; and 3) Beauty of musical scales: intricate relations between notes. Half of the class time was to have students practice these concepts by using a graphical programming language called PureData to create musical notes and chords. I am satisfied by the course outcome; many students showed excitement in the course, especially when they heard sounds generated from their programs. Three of my PhD students have assisted me in this course, and they all felt enriched through these unique experiences.

Conclusions and Future Plan

My teaching philosophy – fostering passion, promoting active learning, and tailoring to students' aptitude – have been manifested through my diverse teaching, mentoring and outreach activities. I feel satisfied when I see my students grow. I feel proud when I see my students succeed and surpass me.

In the future, I wish to expand my course repertoire by teaching other existing courses and creating new courses. These new courses include a seminar course on Interactive Music Systems to survey advances on building interactive systems that allow humans to play music with machines. Another new course I plan to teach is Audio-Visual Scene Understanding to introduce this emerging research area to students through lectures and hands-on projects on the analysis, generation and alternation of audio-visual scenes. On mentoring, I will continue improving my mentoring skills. In five years, I hope all of my current PhD students will have graduated or be close to graduation, with one or two students landing at faculty positions, practicing their own teaching philosophy.