Music 421N: Deep Learning for Music and Audio

Center for Computer Research in Music and Acoustics (CCRMA)
Department of Music, Stanford University
Stanford, California 94305

Fall Quarter, 2017-2018

Overview

This research seminar will discuss advances in deep learning applied to music and audio, and related fields such as speech/image processing. The coverage parallels that of other Stanford courses pertaining to Vision, NLP, and Genomics. The presentations will start with preliminaries about neural nets, signal processing, machine learning, supervised and unsupervised learning, and then we will look at a selection of cutting edge research over the past $2 - -3$ years.

Topics will include teaching computers how to compose music, recognizing emotions from sound, content-based music recommendation (Spotify/Pandora), speech recognition, speech/instrument synthesis, and speech enhancement. Papers studied will be drawn from conferences such as ICASSP, NIPS, ICML, ICLR, etc. We have invited a few guest lectures. This seminar course will introduce students to a number of cutting-edge research topics in this field.

The ideal audience is students having interest in AI/deep-learning, and/or its applications to music/audio/speech signals. The concepts studied in this course are applicable to fields outside of music and audio. The material will be approximately self contained, overlapping with other courses such as CS 231N, CS 224S, EE 264, and Music 421.

Prerequisites: Elementary linear algebra and probability. CS229 and CS 231n will help but are not required.

- Units: 1
- Grading: S/NC
- Time: Friday afternoons (probably 5pm: see the Time Schedule for the latest word)
- Venue: Knoll 217 ("CCRMA Classroom")
- Course website: `https://ccrma.stanford.edu/courses/mus421n/`
- Piazza page: `https://piazza.com/class/j86ay8rctn743k`