

A Fast and Accurate Open-Source Solo Musical Instrument Classifier

Yip, Hanna (School: The Spence School)

The purpose of my research was to build an accurate solo musical instrument classifier that assigns a probability distribution of instrument class labels to a given audio signal. Last year, I compared the bias-variance tradeoff of musical instrument classifiers that used four different learning algorithms -- Random Forest, Adaboost, bagging on Adaboost, and Support Vector Machine -- on a small dataset of 215 audio files from 10 instruments. Random Forest achieved the highest accuracy of 84%. This year, I worked to improve the performance of the Random Forest classifier and quantify the relative contribution of each aspect of the new model to improvement in classification accuracy. The first alteration was extracting Mel Frequency Cepstral Coefficients (MFCC) first and second order differences in addition to MFCCs as input features. The second was training on a larger, more diverse dataset of audio files which includes around 7000 audio clips of 18 instruments from the MedleyDB and Philharmonia library. The Random Forest classifier achieved a testing accuracy of 96%. Training on the larger dataset that includes both short sound samples and full-length songs improved the classifier's ability to perform well on real-world data. The first 40 MFCCs contributed most to accurate classification, while MFCC first and second order differences only incrementally improved the accuracy. The classifier had reasonable confusions between instruments with similar timbre. To test the classifier on real-world data, I implemented the classifier into an interactive web application that identifies the instrument present in live recordings or stems of new multitracks.