

# ADVANCED MACHINE LEARNING, HOMEWORK # 2

Due: March 25, 2015 before class

## Structured Support Vector Machines

In this homework, you will use structured prediction to perform chord transcription from music data as in the following paper:

Structured Prediction Models for Chord Transcription of Music Audio  
Adrian Weller, Daniel Ellis, Tony Jebara  
International Conference on Machine Learning and Applications, 2009  
<http://www.cs.columbia.edu/~jebara/papers/icmla09adrian.pdf>

You will need the  $SVM^{struct}$  or  $SVM^{hmm}$  code available here:

[http://svmlight.joachims.org/svm\\_struct.html](http://svmlight.joachims.org/svm_struct.html)

The data comes from the MIREX challenge. Your goal is to label music by the Beatles where each song is represented as a sequence of Chroma frames over time. Chroma frames are just 12-dimensional continuous features between [0,1] that compactly represent sound in a short time-window. There are 180 Beatles songs. Each was automatically converted from mp3 into Chroma frames. The songs vary in length from 77 to 1806 Chroma frames total. Each Chroma frame has been manually annotated with a discrete chord label. This label is between 0...24 since there are 25 types of chords: 12 major chords (like 'A major'), 12 minor chords (like 'D minor') and one 'no chord' label. You can download the labeled Chroma data here:

<http://www.cs.columbia.edu/~jebara/4772/CHORDS.zip>

This zip file contains each Beatles song as a .mat file. Each file contains an  $F$  matrix which is the Chroma frames across the length of each song and an  $L$  vector which is a list of the discrete labels across the length of each song.

Your goal is to build a structured prediction system that can learn from the annotated training data to then label remaining Beatles songs with the appropriate chord transcriptions. In particular, you will recreate the results in Table 1. You only need to recreate the first column with 30% training data. Explore using linear Chroma features, quadratic Chroma features, and using windows of Chroma features (forward in time and backward in time by up to +3 and -3). Explore different settings of  $C$  by cross-validation. Your goal is try to get similar results as in the left-most column of Table 1 in the paper. Submit your homework via Courseworks as a mini-paper that discusses your data-set, your experiments, your findings and plots of the accuracies you obtained, in particular, your own results for Table 1. Also, show your annotation on any two example Beatles songs in your testing data. Plot each song's Chroma features over time and annotate them with your predicted chord values for each frame as well as the true labels in the  $L$  vector.

BONUS: Explore variations and features with the structured prediction method to get a higher accuracy than 59.2%.

NOTE: It turns out that the data in the paper had some slight temporal misalignment errors which were fixed in the data-set CHORDS.zip. This may cause some slight discrepancy (in your favor) between your results and those in the 2009 paper.