Aligned Hierarchies for Sequential Data

Katherine M. Kinnaird

Abstract

We present aligned hierarchies, a novel solution to the dimension reduction problem, representing high-dimensional and noisy sequential data as a low-dimensional object that encodes relevant information. In this thesis, we motivate our presentation and discussion of aligned hierarchies through the lens of Music Information Retrieval (MIR), constructing aligned hierarchies by finding, encoding, and synthesizing all repeated structure present in a song. In the method for building these aligned hierarchies, we present a novel geometric-based approach to the parameter used for thresholding self-similarity and self-dissimilarity matrix representations for songs. Given a particular MIR task, such as locating the chorus of a given musical song or finding all copies of a particular recording of a song, we compare songs based on their aligned hierarchies. For the fingerprint task and the cover song task, we present comparisons for two music data sets, one based on digitized score and the other one performances of scores. Results from these comparisons on clean data had very high precision-recall values and provide a proof of concept for the aligned hierarchies. Results on noisy data were not as successful, but demonstrate the effect of different pre-processing techniques on the comparisons.