



Department of Engineering
Seminar Program
Wednesday, February 24, 2016
Shanahan Teaching and Learning Center
Lecture Hall 1430, 4:15pm

“Audio Hashprints: Theory & Application”

Timothy J. Tsai

Summary:

We rely heavily on search engines like Google to navigate millions of webpages, but a lot of content of interest is multimedia, not text data. One important class of multimedia data is audio. How can we search a database of audio data? One of the main challenges in audio search and retrieval is to determine a mapping from a continuous time-series signal to a sequence of discrete symbols that are suitable for reverse-indexing and efficient pairwise comparison. This talk introduces a method for learning this mapping in an unsupervised, highly adaptive way, resulting in a representation which we call audio hashprints. We will discuss the theoretical underpinnings that determine how useful a particular representation is in a retrieval context, and we show how hashprints are a suitable representation for tasks requiring high adaptivity. We investigate the performance of the proposed hashprints on two different audio search tasks: synchronizing consumer recordings of the same live event using audio correspondences, and identifying a song at a live concert.

Bio:

TJ Tsai completed his BS and MS in electrical engineering at Stanford University in 2006 and 2007. From 2008 to 2010, he worked at SoundHound, a startup that allows users to search for music by singing or humming a part of the song. He will be receiving his PhD in May 2016 from the electrical engineering and computer science department at the University of California Berkeley, in joint collaboration with the International Computer Science Institute.