



THE UNIVERSITY OF
CHICAGO

Department of Statistics

STATISTICS COLLOQUIUM

KAREN LIVESCU

TTIC & Department of Computer Science
University of Chicago

Multi-View Representation Learning for Speech and Language

MONDAY, February 1, 2016, at 4:00 PM

Eckhart 133, 5734 S. University Avenue

Refreshments following the seminar in Eckhart 110.

ABSTRACT

Many types of multi-dimensional data have a natural division into two "views", such as audio and video or images and text. Multi-view learning refers to techniques that use multiple views of data to learn improved models for each of the views. Theoretical and empirical results show that multi-view techniques can improve over single-view ones in certain settings. In many cases multiple views help by reducing noise in some sense. In this talk, I will focus on multi-view learning of representations (features), especially using canonical correlation analysis (CCA) and related techniques. I will present nonlinear extensions including deep CCA, where the learned representations are the outputs of deep neural networks, and other variants. Finally, I will give recent results on speech and language tasks.