מכון ויצמן למדע

THE WEIZMANN INSTITUTE OF SCIENCE

FACULTY OF MATHEMATICS AND COMPUTER SCIENCE

Vision and Robotics Seminar

Room 141 ,Ziskind Building
on Monday, Jan 12, 2015
at 14:00

NOTE UNUSUAL ROOM, DAY, TIME

Karen Livescu
Toyota Technological Institute, Chicago

Multi-view representation learning: A tutorial introduction and applications to speech and language

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 includes a variety of techniques that use multiple views of data to learn improved models for each of the views. The views can be multiple measurement modalities (like the examples above) but also can be different types of information extracted from the same source (words + context, document text + links) or any division of the data dimensions into subsets satisfying certain learning assumptions. 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 (what is noise in one view is not in the other). In this talk, I will focus on multi-view learning of representations (features), especially using canonical correlation analysis (CCA) and related techniques. I will give a tutorial overview of CCA and its relationship with other techniques such as partial least squares (PLS) and linear discriminant analysis (LDA). I will also present extensions developed by ourselves and others, such as kernel, deep, and generalized ("many-view") CCA. Finally, I will give recent results on speech and language tasks, and demonstrate our publicly available code.

Based on joint work with Raman Arora, Weiran Wang, Jeff Bilmes, Galen Andrew, and others.