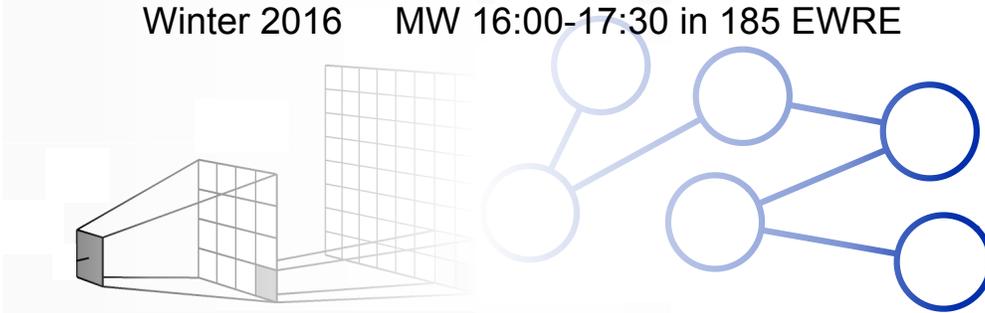


EECS 542

Advanced Topics in Computer Vision

Winter 2016 MW 16:00-17:30 in 185 EWRE



Course Overview: The course will focus on learning structured representations and embeddings for numerous problems in computer vision. Approaches for learning from unimodal and multimodal data will be covered and include topics from sparse coding, convolutional neural networks and others. Multimodal problems will include topics at the boundary of vision and language, vision and speech, and vision and robotics. The course will be a mix of lecture, student presentation, and discussion; students will complete some programming assignments, paper reviews and presentations, and a project. Prior experience from EECS 442 or EECS 5XX (Foundations of Computer Vision) expected.

Student Workload: The course will revolve around reading, reviewing, critiquing, presenting and discussing seminal and recent papers for learning structured representations and embeddings in vision and vision-and-X problems. To facilitate the advanced material, the professor will lecture on necessary background material. Students will read three papers and write two paper reviews each week; make at least two paper presentations during the semester (depending on the size of the course); implement certain methods via programming assignments in the first half of the term; and progress toward a course project in the second half of the term.

Target Audience: The course has been designed to further the progression of graduate students toward new work in computer vision. Prior experience in computer vision is required through EECS 442 or EECS 5XX Foundations of Computer Vision; otherwise instructor approval is required.

Prerequisites: Graduate standing; EECS 442 or 5XX (Foundations of Computer Vision); or Instructor Approval.

Instructor: Prof. Jason Corso, EECS

Website: <http://web.eecs.umich.edu/~jjcorso/t/542W16/>