
ROCHESTER INSTITUTE OF TECHNOLOGY

COLLEGE OF SCIENCE
CENTER FOR IMAGING SCIENCE

COS-IMGS-789* Special Topics: Deep Learning for Vision

1.0 Course Information

a) Catalog Listing

Course title (100 characters)	<i>Special Topics: Deep Learning for Vision</i>
Transcript title (30 Characters)	<i>Deep Learning for Vision</i>
Credit hours	3
Prerequisite(s)**	Graduate standing in science or engineering, or permission of the instructor
Co-requisite(s)	

b) Terms(s) offered

X	Fall
	Spring
	Summer
	Other
	Offered biannually

If "Other" is checked, explain:

c) Instructional Modes

	Contact hours	Maximum students/section
Classroom	3	20
Lab		
Studio		
Other (specify, i.e. online, workshop seminar, etc.)		

2.0 Course Description (as it will appear in the bulletin)

Deep learning is an area of machine learning that has enabled enormous progress on long-standing problems in computer vision. This course will review neural networks and related theory in machine learning that is needed to understand how deep learning algorithms work. After gaining the prerequisite background knowledge, the course will review the latest algorithms that use deep learning to solve problems in computer vision and machine perception, and students will read recent papers on these systems. Beyond reviewing state-of-the-art systems, students will implement and evaluate one or

more of these systems and then apply them to problems that match their interests. Students are expected to be comfortable with linear algebra and calculus. No prior background in machine learning or pattern recognition is required. (Prerequisites: Graduate Standing or permission of instructor). Class 3, Credit 3 (F)

3.0 Goals of the Course

1. To understand how deep learning algorithms work and how to train them.
2. To review recent state-of-the-art applications of deep learning to problems in computer vision and machine perception.
3. To gain experience using deep learning to solve problems in computer vision and machine perception.

4.0 Intended course learning outcomes and associated assessment methods

Course Learning Outcome	Assessment Method
Outline prerequisite deep learning background material	Homework
Critique papers in deep learning	Presentations; Project
Apply deep learning to problems in computer vision	Homework; Project

5.0 Topics (should be in outline format)

1. Mathematical background
 - a. Linear Algebra
 - b. Probability
2. Machine Learning Basics
 - a. Supervised vs. Unsupervised Learning
 - b. Classifiers
 - c. Performance Evaluation
3. Feed-forward Neural Networks
4. Regularization and Sparsity
5. Optimization for Training Deep Models
6. Convolutional Networks
7. Recurrent Networks
8. Autoencoders
9. Generative Models
10. Transfer Learning
11. Current Limitations of Deep Learning
12. Applications
 - a. Object Recognition
 - b. Object Detection
 - c. Semantic Segmentation
 - d. Instance Segmentation
 - e. Object Tracking
 - f. Visual Question Answering
 - g. Perception for Robotics
 - h. Deep Reinforcement Learning

6.0 Possible Resources

Book: Deep Learning. Ian Goodfellow, Yoshua Bengio, Aaron Courville, MIT Press.
<http://www.deeplearningbook.org/>

Recent papers published in premier machine learning and computer vision conferences, e.g., NIPS, CVPR, ICCV, Nature, etc.

7.0 Program outcomes and/or goals supported by this course)

- To provide students with a depth and breadth of imaging science.
- Prepares graduate students in science and engineering for careers in imaging science by educating them in the construction of sophisticated mathematical models of realistic problems associated with imaging systems and the implementation of computations with practical programming tools for testing, evaluation, and refinement of designs.
To develop the student's skills in applying mathematical techniques and scientific reasoning to different laboratory situations.