

CS 6890: Deep Learning

Spring 2020

Class Meetings: Tue, Thu 1:30 – 2:50 pm, ARC 121

Instructor: Razvan Bunescu

Office: Stocker 341

Office Hours: Tue, Fri 3:00 – 4:00 pm, or by email appointment

Email: bunescu @ ohio edu

Course Website: <http://ace.cs.ohio.edu/~razvan/courses/dl6890>

Prerequisites:

Previous exposure to basic concepts in machine learning, such as: supervised vs. unsupervised learning, classification vs. regression, linear regression, logistic and softmax regression, cost functions, overfitting and regularization, gradient-based optimization. Experience with programming and familiarity with basic concepts in linear algebra and statistics.

Textbook:

There is no required textbook for this class. Slides and supplementary materials will be made available on the course website.

Machine learning introductions:

Listed on the course website.

Supplemental deep learning resources:

Listed on the course website.

Course Description:

This course will introduce the **multi-layer neural networks**, a common deep learning architecture, and gradient-based training through the **backpropagation algorithm**. Fully connected neural networks will be followed by more specialized neural network architectures such as **convolutional neural networks**, **recurrent neural networks with attention**, and **deep generative models**. The later part of the course will explore more advanced topics, such as **adversarial examples**, **deep reinforcement learning**, and **interpretability**. The lectures will cover theoretical aspects of deep learning models, whereas homework assignments will give students the opportunity to build and experiment with shallow and deep learning models, for which skeleton code will be provided.

Proposed Topics:

Logistic and Softmax Regression, Feed-Forward Neural Networks, Backpropagation, Vectorization, PCA and Whitening, Deep Networks, Convolution and Pooling, Recurrent Neural Networks, Long Short-Term Memory, Gated Recurrent Units, Neural Attention Models, Sequence-to-Sequence Models, Distributional Representations, Variational Auto-Encoders, Generative Adversarial Networks, Deep Reinforcement Learning.

Implementation Platforms:

Python, NumPy/SciPy, PyTorch or TensorFlow.

Grading:

30%: 5 Homework Assignments

10%: 1 Class Presentation.

30%: 2 Exams (Mar 5 and Apr 23, in class)

30%: Final Project

Important Dates:

Tuesday, Mar. 10: Spring break, no class.

Thursday, Mar. 12: Spring break, no class.

Thursday, Apr. 23: Last day of this class.

Course and Attendance policies:

Assignments: All homework assignments are due before the class. No late submissions will be accepted without prior approval.

Attendance: It is in your best interest to attend all the lectures. Some of the material will not be found in the reading materials or on the slides. Extra credit will be awarded for class activity. Also, be sure to check your OU email for important announcements on a regular basis.

Academic Dishonesty Policy:

All work must be the student's own. All external references used in reports must be properly cited. No credit will be given for duplicate or plagiarized work. Additional measures may be imposed by the Office of Community Standards and Student Responsibility, when conditions warrant. Students may appeal academic sanctions through the grade appeal process. The OU Student Code of Conduct Policy is available online at:

<http://www.ohio.edu/communitystandards/academic/students.cfm>

Disability-based Accommodation:

Any student who suspects s/he may need an accommodation based on the impact of a disability should contact the class instructor privately to discuss the student's specific needs and provide written documentation from the Office of Student Accessibility Services. If the student is not yet registered as a student with a disability, s/he should contact the Office of Student Accessibility Services.

Other Policies:

Be sure to notify the professor of any exam conflicts or other extenuating circumstances well in advance. No missed exams will be made up without prior approval. Medical excuse forms need to explicitly mention that the student could not have attended the exam at the specified time due to health concerns.