

CS 499 DL: Deep Learning
Spring 2022
Department of Computer Science, George Mason University

Time and location:

- Friday, 1:30 pm - 4:10 pm
- East 122

Instructor:

- Dr. Shuochao Yao
- Email: shuochao@gmu.edu
- Office hours: Tuesday 3:00 pm – 4:00 pm, or by appointment
- Zoom link:
<https://gmu.zoom.us/j/99089635051?pwd=UzUvU0FyR3VhanByU3UyVldydmQydz09>

Description:

- This class provides a practical introduction to deep learning, including theoretical motivations and how to implement it in practice. As part of the course we will cover multilayer perceptrons, backpropagation, automatic differentiation, and stochastic gradient descent. Moreover, we introduce convolutional networks for image processing, starting from the simple LeNet to more recent architectures such as ResNet for highly accurate models. Secondly, we discuss sequence models and recurrent networks, such as LSTMs, GRU, and the attention mechanism. The goal of the course is to provide both a good understanding and good ability to leverage deep learning technique. The course loosely follows Dive into Deep Learning in terms of notebooks, slides

Books:

- Aston Zhang, Zack C. Lipton, Mu Li and Alex J. Smola, "Dive into Deep Learning" (online available: <https://d2l.ai/>)

Graduate teaching assistant (TA):

- Ting Li
- Email: tli21@gmu.edu
- Office hours: Thursday 9:00 am – 10:00 am
- Zoom link: <https://gmu.zoom.us/j/4556011224>

Topics:

- Basic Linear Algebra, Probability and Statistics
- Gradients, Basic Optimization, and Information Theory
- Multilayer Perceptron, Weight Decay, Dropout, and Numerical Stability

- Basic and Advanced Convolution Neural Network, Image Augmentation, and Fine Turning
- Object Detection
- Sequence models
- Basic and Advanced Recurrent Neural Network
- Embedding (Word2vec, FastText, GloVe, Sentiment Analysis)
- Encoder-Decoder, Seq2seq, Machine Translation
- Attention, Transformer, BERT

Grading:

- Your grade will be calculated using the following percentages:
 - Project: 50% (Proposal 10%, Final Report 40%)
 - Can be done with a team of 2-3 students
 - Mid-term (15%) & Final Project Presentation (15%)
 - With the project team
 - Paper Summary (15%)
 - Participation (5%)

Course Schedule:

Note that the schedule is tentative and is subject to change.

Date	Topic
Jan-28	Introduction, Basic Linear Algebra, Probability and Statistics
Feb-04	Gradients, Basic Optimization, and Information Theory
Feb-11	Multilayer Perceptron, Weight Decay, Dropout, and Numerical Stability
Feb-18	Convolution Neural Network
Feb-25	Advanced Convolution Neural Network, Image Augmentation, and Fine Turning
Mar-04	Mid-term Paper Presentation
Mar-11	Mid-term Paper Presentation
Mar-18	Spring Recess—No Class
Mar-25	Object Detection; Sequence models
Apr-01	Recurrent Neural Network
Apr-08	Advanced Recurrent Neural Network; Embedding
Apr-15	Encoder-Decoder, Attention, BERT
Apr-22	Model Compression and Machine Learning System
Apr-29	Final Project Presentation
May-06	Final Project Presentation

Honor Code:

- Please see the Office for Academic Integrity (<https://oai.gmu.edu/>) for a full description of the code and the honor committee process, and the Honor Code Policies of the

Department of Computer Science (<https://cs.gmu.edu/resources/honor-code/>) regarding the course project. GMU is an Honor Code university. The principle of academic integrity is taken seriously and violations are treated gravely. If you rely on someone else's work in an aspect of the course project, you should give full credit in the proper, accepted form. Another aspect of academic integrity is the free play of ideas. Vigorous discussion and debate are encouraged in this course, with the firm expectation that all aspects of the class will be conducted with civility and respect for differing ideas, perspectives, and traditions. When in doubt (of any kind) please ask for guidance and clarification.