Theory Outline

- Week 1 – Jan 13
  - What is Machine Learning?
  - Single Layer NN
  - The Loss Function
- Week 2 – Jan 20
  - Gradient Descent Optimization
  - Mathematical calculation of gradient
  - Matrix interpretation of gradient computation
- Week 3 – Jan 27
  - Tensors
  - Gradient Descent for Single Layer NN
  - Local and Global Minima
- Week 4 – Feb 3
  - Optimization of Deep Functions
  - Gradient Descent on Acyclic Graphs
  - General Loss Functions
- Week 5 – Feb 10
  - Convolutional Neural Networks
  - Gradients for CNNs
- Week 6 – Feb 17
  - Probability and Estimation
  - Frequentist versus Bayesian Estimation
  - Bias and Variance
- Week 7 – Feb 24
  - Training and Generalization
  - Regularization and Dropout Methods
- Week 8 – March 3
  - Exam #1
- Week 9 – March 10
  - Stochastic Gradient Descent
  - Batches and Epochs
  - Learning Rate and Momentum
- Week 10
  - Spring break
- Week 11 – March 24
  - Widely Used DL Techniques
  - Vanishing Gradients; Skip Connections
  - Batch Normalization
  - Transfer Learning and Data Augmentation
- Week 12 – March 31
  - Recurrent Neural Networks
  - LSTM and GRU Networks
- Week 13 – April 7
  - Autoencoders
  - Unsupervised training
  - Concept of a generator
- Week 14 – April 14
  - Adversarial Learning
  - The generator and discriminator
  - Generative Adversarial Networks (GAN)
  - Nash equilibrium
- Week 15
  - Lecture - April 19
  - GAN convergence
  - Wasserstein and Conditional GANs
  - Exam #2 – April 21
- Week 16 – April 28
  - Reinforcement learning
  - Various approaches: Q learning, double Q learning, Actor Critic etc.
  - Model and curiosity based RL

Exams in Weeks 8 and 15