Theory Outline

Fall 2026 Syllabus

- Week 1 – Jan 15
  - What is Machine Learning?
  - Single Layer NN
  - The Loss Function
- Week 2 – Jan 22
  - Gradient Descent Optimization
  - Mathematical calculation of gradient
  - Matrix interpretation of gradient computation
- Week 3 – Jan 29
  - Tensors
  - Gradient Descent for Single Layer NN
  - Local and Global Minima
- Week 4 – Feb 5
  - Optimization of Deep Functions
  - Gradient Descent on Acyclic Graphs
  - General Loss Functions
- Week 5 – Feb 12
  - Convolutional Neural Networks
  - Gradients for CNNs
- Week 6 – Feb 19
  - Probability and Estimation
  - Frequentist versus Bayesian Estimation
  - Bias and Variance
- Week 7 – Feb 26
  - Training and Generalization
  - Regularization and Dropout Methods
- Week 8 – March 5
  - Exam #1
- Week 9 – March 12
  - Stochastic Gradient Descent
  - Batches and Epochs
  - Learning Rate and Momentum
- Week 10
  - Spring break
- Week 11 – March 26
  - Widely Used DL Techniques
  - Vanishing Gradients; Skip Connections
  - Batch Normalization
  - Transfer Learning and Data Augmentation
- Week 12 – April 2
  - Multi-Head Attention Layers
  - Transformer Layers
  - The Visual Transformer
- Week 13 – April 9
  - Recurrent Neural Networks
  - LSTM and GRU Networks
  - Unsupervised Learning and Autoencoders
- Week 14 – April 16
  - Adversarial learning
  - Generative Adversarial Networks (GAN)
  - Nash equilibrium
  - GAN convergence, theory and practice
  - Gibbs Distributions, and Score
- Week 15 – April 23
  - Exam #2
- Week 16 – April 30
  - Denoising Score Matching, and Langevin dynamics
  - Generative Diffusion Models
  - DALL-E Generative Diffusion Model

Exams in Weeks 8 and 15