

MIDTERM EXAM GUIDE

The first half the class concentrated on the following three modules:

1. Basics of Artificial Neural Networks
2. Convolutional Neural Networks
3. Recurrent Neural Networks

Hence, the midterm exam will cover all materials contained in the lectures. Topics covered in the lectures are listed in detail below:

Basics of Artificial Neural Networks

- Lecture 1: Introduction to Deep Learning
compositionality, end-to-end learning, distributed representations
- Lecture 2: Machine Learning Overview
linear models, loss functions, linear regression, gradient descent, overfitting and generalization, regularization, cross-validation, bias-variance tradeoff, maximum likelihood estimation
- Lecture 3: Multi-layer Perceptrons
feed-forward neural networks, activation functions, chain rule, backpropagation, computational graph, automatic differentiation, distributed word representations
- Lecture 4: Training Deep Neural Networks
data preprocessing, weight initialization, normalization, regularization, model ensembles, dropout, optimization methods

Convolutional Neural Networks

- Lecture 5: Convolutional Neural Networks
convolution layer, pooling layer, evolution of depth, design guidelines, residual connections, semantic segmentation networks, addressing other tasks
- Lecture 6: Understanding and Visualizing Convolutional Neural Networks
transfer learning, interpretability, visualizing neuron activations, visualizing class activations, pre-images, adversarial examples, adversarial training

Recurrent Neural Networks

- Lecture 7: Recurrent Neural Networks
sequence modeling, recurrent neural networks (RNNs), RNN applications, vanilla RNN, training RNNs, long short-term memory (LSTM), LSTM variants, gated recurrent unit (GRU)

Things to bring:

- 1 sheet of hand-written notes on A4 paper (you can use both sides of this sheet).
- Scientific calculator