APS360. Midterm Study Topics

This is a list of lecture/lab topics compiled by your instructor to help you study for the midterm. Note that the list may not be exhaustive.

In particular, the list of topics doesn't cover the skills that we develop throughout the course through the labs. For example, interpreting the training curve is a skill that we practiced throughout the course. As another example, you should be able to predict the output shape of PyTorch layers we discussed in class, like `nn.Conv2d`.

Week 1

Machine Learning Overview

- artificial intelligence, machine learning, deep learning
- type of problems where machine learning is successful
- supervised learning (regression, classification), unsupervised learning, reinforcement learning
- issues with deep learning: interpretability, fairness

Biological motivations

- biological neuron: dendrite, axon, synapse, cell body, the direction that information travel
- grandmother cell
- distributed encoding

Artificial Neuron

- artificial neuron: activation, weight, bias
- input layer, hidden layer(s), output layer
- fully-connected layer
- feed-forward neural network
- multi-layer perceptron (MLP)
- steps to training a neural network
- loss function (binary cross-entropy)
- optimizer (stochastic gradient descent)
- overfitting, training set, test set, ground-truth, prediction
- ethics and interpretability

Week 2

Artificial Neural Networks

- ReLU, sigmoid, tanh (know the behaviour)
- LeakyRelu (from Lab1)
- neural network architecture
- counting number of layers in a MLP
- parameters, weights, biases
- counting number of parameters / weights / biases
- credit assignment problem
- standard training/test set splits

Optimization and Hyperparameter

- batch size (effect of large / small batch size)
- iteration, epoch
- the use of PyTorch's `DataLoader` to batch data for us
- learning rate (effect of large / small learning rate)
- training curve (what's on each axis)
- training loss / accuracy / error
- interpreting training curves
- validation set
- validation loss / accuracy / error
- hyperparameter search
- examples of hyperparameters beyond learning rate and batch size
- model selection (based on validation)
- Checkpointing (from Lab2)

Week 3

Multi-Class Classification

- the use of PyTorch fully-connected `nn.Linear` layer
  - the meaning of `in_features` and `out_features`
  - changing the number of layers in a neural network
- the use of PyTorch activation functions `F.relu` and `F.softmax`
- the difference between PyTorch `CrossEntropyLoss` and `BCEWithLogitsLoss`
- the general training loop code:

```
out = model(imgs)             # forward pass
loss = criterion(out, labels) # compute the total loss
loss.backward()               # backward pass (compute parameter updates)
optimizer.step()              # make the updates for each parameter
optimizer.zero_grad()         # a clean up step for PyTorch
```

- interpreting training curves to identify overfitting
- being able to understand relevant PyTorch layers we discussed when the API documentation is given
- forward pass, backwards pass
- predict the output shape of a PyTorch layer

Troubleshooting

- the idea of overfitting on a small data set

Week 4

Convolutions

- reasons why fully-connected layers can be problematic
- locally-connected layers, weight-sharing
- convolutional kernels / filters and their sizes
- input and output channels / feature maps and their sizes
- convolutional arithmetics (computing forward pass of a convolution layer given the input and kenrel)
- PyTorch "NCHW" format
- PyTorch convolution layers `nn.Conv2D`, its parameters, and computing the output shape of a convolution
- zero-padding (the reason for their use, selecting padding size)
- max-pooling, and the `nn.MaxPool2D` layer

- average-pooling
- strided convolutions
- examples of the types of features computable using convolutions
- forward pass arithmetics of pooling and padded/strided convolutions

Architectures

- What is new in each architectures:
  - LeNet
  - AlexNet
  - GoogLeNet: repeated modules, multiple convolution kernels
  - ResNet: residual connections
- Global average pooling
- Fully-convolutional architectures

Week 5

Regularization

- why we don't talk about underfitting as much
- how to detect overfitting
- weight sharing
- data normalization
- data augmentation
- weight decay
- early stopping
- dropout
- marking whether a PyTorch model is in training mode or evaluation mode

Convolution Transpose

- upsampling
- convolution transpose arithmetics
  - effect of padding
  - effect of strides
  - effect of output_padding
- we did not cover counting the number of parameters of a transpose convolution layer

Autoencoders

- the idea of an embedding space, and different features/representation of the input
- mean-squared error loss (MSELoss)
- how to train an autoencoder
- how to use the decoder to generate new images
- the effect of interpolation in the embedding space vs interpolating in the pixel space
- possible uses of an autoencoder
- effect of size of latent space
- the general idea of an "embedding space"

Word Embeddings

- one-hot encodings
- word2vec idea: meaning of a word depends on its context

- general architecture of a word2vec model (input/output of the encoder/decoder)
- distances measures in embedding spaces (euclidean distance, cosine similarity)
  - intuition about how "close" points are in these measures
- effect of word embedding size
- how to use word embeddings as input
- limitations to using word embeddings (e.g. bias)

Bias

- learning the bias in the data set
- the type of bias in GloVe embeddings
- the type of bias can exist in any classifier