

# Learning Goals 576

Theja Tulabandhula

# Overall Goals

Intuition for different deep learning models

Confidence while considering use of deep learning methods

Cognizance of choices being made in an deep learning workflow

Successfully use packages in Python with large datasets


# A Note on Practice

The goal of the assignments is to get you up to speed with Python and practice with a few datasets

The goal of the lectures is to ask 'why' for every step we take (or choice we make) in a deep learning workflow

These two complement each other so the learner can get a better understanding of the subject.





Make use of the resources, reach out to the  
teaching staff and have fun learning!



# 576: Goals for Lecture 1

1

Understand the key components to set up a classification task

3

Understand how chain rule works

2

Relate business problems to machine learning methods

4

Understand why multiclass logistic regression may not work well even for 2D data

# 576: Goals for Lecture 2

1

Get acquainted with the basics of Python

3

Convolution layer as collection of filters applied to input tensors

2

Understand the notion of hidden layers and nonlinearities

4

Why pooling helps in reducing parameters downstream

# 576: Goals for Lecture 3

1

Understand how to transfer parameters previously learned for a new task

3

Be aware of the different engineering tricks such as dropout, batch normalization

2

Know the different ways to debug a deep network

4

Learn why image datasets can be enhanced using data augmentation

# 576: Goals for Lecture 4

1

Understand how natural language elements (such as words) are processed in an analytics workflow

2

Understand the shortcomings of methods such as Naive Bayes, Latent Dirichlet Allocation

3

Realize that a CNN can also be used for a NLP task (sentence classification/sentiment analysis)

4

What is word2vec and how does it help in NLP tasks?



# 576: Goals for Lecture 5

1

What prediction tasks have sequential dependencies?

3

A high level idea about LSTMs

2

The basic RNN architecture

4

Applications of 'sequential to sequential' models

# 576: Goals for Lecture 6

1

Meaning of generative modeling

3

What are variational autoencoders (VAEs) and where can they be used?

2

The intuition behind generative adversarial networks (GANs)

4

Differences between GANs and VAEs

# 576: Goals for Lecture 7

1

What is online learning? How is it different from supervised learning?

3

The multi armed bandit problem and solutions

2

Relation between forecasting and decision making

4

Contextual bandits

# 576: Goals for Lecture 8

1

What is reinforcement learning?

3

Policies, Values and how to think about these two objects

2

Basics of a Markov Decision Process

4

Bellman Expectation Equation and Bellman Optimality Equation

# 576: Goals for Lecture 9

1

Function approximation in Q-learning

3

AlphaZero for Go, Chess and Shogi

2

DQN for Atari games

4

Course roundup!

# 576: Goals for Lecture 10

1

What are graphical models?

3

The mechanics of directed graphical models

2

Relating a distribution to a graph and reading the implied conditional independence assumptions

4

The mechanics of undirected graphical models

# 576: Goals for Lecture 11

1

What is meant by inference? What are some example inference tasks?

3

Using belief propagation for inference

2

The notion of factor graphs

4

The basics of sampling (MCMC, Gibbs) for inference

# 576: Goals for Lecture 12

1

How to learn the parameters of a graphical model?

3

The EM algorithm for tough MLE problems

2

The Maximum likelihood estimation problem

4

Learning parameters in applications