

Full syllabus (from 2022): <https://cseweb.ucsd.edu/classes/fa22/cse258-a/>

## Basic Info

### Intro:

CSE 158 and 258 are undergraduate and graduate courses devoted to current methods for recommender systems, data mining, and predictive analytics. No previous background in machine learning is required, but all participants should be comfortable with programming (all example code will be in Python), and with basic optimization and linear algebra.

### Lectures:

Lectures will be broadcast (in Twitch or similar live-streaming platform) and recorded on UCSD podcasts.

### Textbook:

The textbook for this course is [Personalized Machine Learning](#). A (draft) pdf of the textbook is available [for download](#).

Additional references are provided from [Pattern Recognition and Machine Learning \(Bishop\)](#), and from Charles Elkan's 2013 [course notes](#). Links are also provided to our [Coursera Specialization](#), which covers similar (though more introductory) material.

### Office hours:

Office hours (and instructions to access) for each class are posted to Piazza

### Assessment:

- **Homework 1:** due Monday week 3
- **Homework 2:** due Monday week 5
- **Midterm:** (roughly) November 1
- **Homework 3:** due Monday week 7
- **Assignment 1:** due Monday week 8
- **Homework 4:** due Monday week 9
- **Assignment 2:** due Tuesday week 10

See course webpage for last year's versions.

### Grading:

- Each **Homework** is worth 8%. Your lowest (of four) homework grades is dropped (or one homework can be skipped).
- The **(take-home) Midterm** is worth 26%.
- Each **Assignment** is worth 25%.
- **Assignment 2** is a **group assignment**. All other assessments must be completed individually.
- All assessments are due **before** the Tuesday lecture on the due date. Late submissions are not accepted.

## Wk 0/1

### Supervised Learning: Regression

- Least-squares regression
- Overfitting and regularization
- Feature Engineering
- Training, validation, and testing

#### References:

- [Personalized Machine Learning](#), Chapter 2
- [Chapter 2 Workbook](#)
- [Notebook from lecture](#)

## Wk 1/2

### Supervised Learning: Classification

- Logistic regression
- Other Classification Algorithms
- How to evaluate classifiers

#### References:

- [Personalized Machine Learning](#), Chapter 3
- [Chapter 3 Workbook](#)

- [Notebook from lecture](#)

**Wk  
3/4/5**

## **Recommender Systems**

- Collaborative Filtering
- Latent Factor Models
- Recommender Systems Evaluation
- Deep Learning for Recommendation

### **References:**

- [Personalized Machine Learning](#), Chapter 4
- [Personalized Machine Learning](#), Chapter 5
- [Chapter 4 Workbook](#)
- [Chapter 5 Workbook](#)

**Wk 6/7**

## **Text Mining**

- Sentiment analysis
- Bags-of-words
- TF-IDF
- Stopwords, stemming, and low-dimensional representations of text

### **References:**

- [Personalized Machine Learning](#), Chapter 8
- [Chapter 8 Workbook](#)

## Wk 7

# Content and Structure in Recommender Systems

- Factorization Machines
- Group- and Socially-Aware Recommendation
- Online Advertising
- (Maybe some tools/libraries/misc topics, depending on time)

### References:

- [Personalized Machine Learning](#), Chapter 6
- [Chapter 6 Workbook](#)

## Wk 8

# Modeling Temporal and Sequence Data

- Sliding windows and autoregression
- Temporal dynamics in recommender systems

### References:

- [Personalized Machine Learning](#), Chapter 7
- [Chapter 7 Workbook](#)

## Wk 9

# Visual Recommendation

- Complementary Item Recommendation
- Fashion and Outfit Recommendation
- Fit Prediction

**References:**

- [Personalized Machine Learning](#), Chapter 9
- [Chapter 9 Workbook](#)

**Wk  
9/10**

## **Ethics and Fairness**

- Filter Bubbles and Recommendation Diversity
- Calibration, Serendipity, and Other "Beyond Accuracy" Measures
- Algorithmic Fairness

**References:**

- [Personalized Machine Learning](#), Chapter 10
- [Chapter 10 Workbook](#)