

CSCI567 Machine Learning (Fall 2024)

Prof. Dani Yogatama

University of Southern California

Aug 30, 2024

Introduction

About

- Modern machine learning methods used in real-world AI applications.
- Focus on conceptual understanding of these methods.

Objectives

- Develop skills to grasp abstract ML concepts and think critically.
- Practice with hands on programming tasks.
- Preparation for studying advanced machine learning topics.

Prerequisites

- Undergraduate level training in probability and statistics, linear algebra, multivariate calculus.
Important: attend today's discussion session to see if you have the required background.
- Programming in Python
Not an intro-level CS course, no training of basic programming skills.

Logistics

- Lectures: Fridays 2-4.20pm (SGM 123)
- Discussions: Fridays 4.30-5.20pm
- Course website: <https://usc-tamagotchi.github.io/csci-567/24f/>
- TAs: Yavuz Faruk Bakman, Robby Costales, Xiao Fu, Isabelle Lee, Sajjad Shahabi, Xinyan Yu
- No required textbook, lecture slides will be posted before class.
Mathematics for Machine Learning <https://mml-book.github.io/>
Deep Learning <https://www.deeplearningbook.org/>

Grade

- Quiz 1 (10/4): 40%
- Quiz 2 (12/6): 40%
- Project report: 20%

Initial cut-offs

$B^- = [70, 75)$, $B = [75, 80)$, $B^+ = [80, 86)$

$A^- = [86, 92)$, $A = [92, 100)$

Final cut-offs will not be released

Course Project

An implementation focused project to assess real world machine learning skills done in a group of four.

- Choose a popular benchmark that is approved by at least two TAs (5 points).
- Implement a top performing baseline method *from scratch* that achieves comparable performance to the original implementation (10 points).
- Improve on the baseline (5 points).
- Write an intelligible report to present your results (5 points).

What is machine learning?

It is the fuel that powers state-of-the-art AI models.

AI is a civilization-altering technology that is going to transform the way we live.

Consumer products. speech-to-speech chatbot, search engine, stock price prediction, wearable devices.

Scientific applications. protein structure prediction, nuclear fusion, social network analysis.

The frontier of AI

- Pseudo artificial general intelligence.
- Chatbots that passed the Turing test.
- Realistic image generation.
- Short video generation.

Why study ML today?

Acquire the knowledge to apply state-of-the-art methods to solve real-world problems.

Analyze how existing models work to understand how to control them better.

Make progress towards superintelligence.

Foundations of Modern Machine Learning

- Data
- Model
- Loss function (objective function)
- Evaluation metric

Outline

- 1 Overview of machine learning


What is machine learning?

One possible definition (cf. Murphy's book)

a set of methods that can automatically *detect patterns* in data, and then use the uncovered patterns to *predict future data*, or to perform other kinds of *decision making under uncertainty*


Example: detect patterns

How the temperature has been changing?



Example: detect patterns

How the temperature has been changing?




Patterns

- Seems going up
- Repeated periods of going up and down.

How do we describe the pattern?


Build a model: fit the data with a polynomial function



- The model is not accurate for individual years
- But collectively, the model captures the major trend

Predicting future

What is temperature of 2030?



- Again, the model is probably inaccurate for that specific year
- But then, it might be close enough

What we have learned from this example?

Key ingredients in machine learning

- Data
collected from past observation (we often call them *training data*)

What we have learned from this example?

Key ingredients in machine learning

- Data
collected from past observation (we often call them *training data*)
- Modeling
devised to capture the patterns in the data
 - The model does not have to be true — “All models are wrong, but some are useful” by George Box.

What we have learned from this example?

Key ingredients in machine learning

- Data
 - collected from past observation (we often call them *training data*)
- Modeling
 - devised to capture the patterns in the data
 - The model does not have to be true — “All models are wrong, but some are useful” by George Box.
- Prediction
 - apply the model to forecast what is going to happen in future

A rich history of applying statistical learning methods

Recognizing flowers (by R. Fisher, 1936)

Types of Iris: setosa, versicolor, and virginica



Huge success 30 years ago

Recognizing handwritten zipcodes (AT&T Labs, late 1990s)

true class = 7



true class = 2



true class = 1



true class = 0



true class = 4



true class = 1



true class = 4



true class = 9




true class = 5



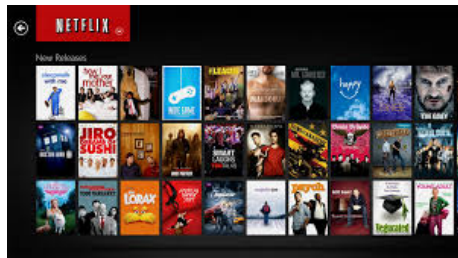
More modern ones, in your social life

Recognizing your friends on Facebook



It might know more about you than yourself

Recommending what you might like



A screenshot of the Netflix recommendation engine interface. It features two main sections: 'More Works to Consider' and 'Related to Items You've Viewed'. Each section displays a grid of movie and TV show covers with associated titles and prices.

More Works to Consider

- Frank & Jesse: What to Cook #10** (TV-14) - \$17.99
- What to Cook and How to Cook #9** (TV-14) - \$12.49
- Whiskin' (Unrated)** (TV-14) - \$12.99
- It, Jay's Parents, Bucked with Will...** (TV-14) - \$13.99
- SOLO: The Ice Age Adventure to the Edge of the World** (TV-14) - \$11.99
- From: Peppino** (TV-14) - \$25.99

Related to Items You've Viewed

- What the Saw Fell On** (TV-14) - \$9.99
- Wings** (TV-14) - \$9.99
- Providence** (TV-14) - \$9.99
- Love Is a Four Letter Word** (TV-14) - \$10.99
- Management** (TV-14) - \$9.99
- Love Is a Four Letter Word** (TV-14) - \$9.99

What is in machine learning?

Different flavors of learning problems

- Supervised learning
Aim to predict (as in previous examples)

What is in machine learning?

Different flavors of learning problems

- Supervised learning
Aim to predict (as in previous examples)
- Unsupervised learning
Aim to discover hidden and latent patterns and explore data

What is in machine learning?

Different flavors of learning problems

- Supervised learning
Aim to predict (as in previous examples)
- Unsupervised learning
Aim to discover hidden and latent patterns and explore data
- Decision making (e.g. reinforcement learning)
Aim to act optimally under uncertainty

What is in machine learning?

Different flavors of learning problems

- Supervised learning
Aim to predict (as in previous examples)
- Unsupervised learning
Aim to discover hidden and latent patterns and explore data
- Decision making (e.g. reinforcement learning)
Aim to act optimally under uncertainty
- Many other paradigms