Machine Learning Foundations and Applications (AI42001)

Announcements

Theory: Course Logistics and Discussion

Course Logistics (8-9 AM, Tue, 2-1-2024)

Teachers: Sudheshna Sarkar, Somdyuti Paul, Mahesh Mohan M R
References
• Introduction to Machine Learning, Ethem Alpaydin, MIT Press (available online)
• Machine Learning, Tom M Mitchell, Mc Graw Hill Education (available online)
Theory
• Monday – 11:00-11:55 AM, Tuesday – 8:00-9:55 AM
• Evaluation: Class Test-1, Midsem, Class Test-2, and Endsem (5%+20%+5%+30%)
Laboratory
• Thursday – 02:00 - 05:00 PM
• Demo + Assignment Discussion + Problem solving
• Evaluation:
1. 12-15 program assignments (20%)
2. Midsem Lab Test + Endsem Lab Test (10% +10%)

Plagiarism: No tolerance policy. Binary marking (both parties).
Attendance: Compulsory to avoid deregistration.

Class Discussion: Applications of ML in Different Fields (9:05-9:55 AM, Tue, 2-1-2024)

Robotic arm based manufacturing
Drone for package delivery and agricultural pesticide spraying
Banking -- Loan allotment rules based on Decision tree
Medical Field -- Smart phone based tele-medicine and cancer detection from CT
Captcha
Recommendation systems
Production line to detect faulty products

Course Logistics

Lab: Introduction to Python Programming

Introduction to Python (2-5 PM, Thu, 4-1-2024)

Python: Fun to Use
Intro to Google Collab
Intro to Numpy: Matrices/Vectors and Mathematical Manipulations
Intro to Pandas: Dataset structure and easy analysis
Intro to Matplotlib: Plotting Data

Slides: Intro to Python

Class Tutorial: Intro to Python

Lab Assignment 1: Intro to Python

Theory: Introduction to Machine Learning

Intro to ML: Part 1 (11-11:55 AM, Mon, 8-1-2024)

Motivations of ML. Why ML is trending now?
Classification of ML: Supervised, Unsupervised and Reinforcement Learning
Algorithm and its ineffectiveness for problems intuitive to humans (e.g., digit classification)
Basics of Supervised Learning

Intro to ML: Part 2 (8-9:15 AM and 9:20 AM to 9:55 AM, Tue, 9-1-2024)

Basics of Supervised Learning (contd): Labelled Data, Model, Loss, and Parameter Optimization
Basics of Unsupervised Learning: Clustering and Association
Basics of Reinforcement Learning: Optimal sequence of steps to an objective based on states

Reading Exercise : A very detailed and informative account on the Introduction to ML can be found in Secs 1.1 and 1.2 of the Ethem Alpaydin's text book (i.e., our first reference text)!

Intro to ML: Part 1

Intro to ML: Part 2

Practice Problems: Intro to ML

Theory + Lab: K Nearest Neighbor

K Nearest Neighbor (Theory: 2:15 PM to 3:55 PM)

KNN basics
Weighted KNN
Different Distance Metrics: Eucledian, Manhattan, Chebyshev, Cosine and Hamming.
Normalization: Min-max and Z-score normalization
Drawbacks: Curse of Dimensionality, Expensive and Storage Need

Lab Tutorial: 4 - 4:55 PM

Intro to Scikit Learn
KNN Tutorial using Scikit Learn

Reading Exercise: Sections 8.1 and 8.2 of the Tom Mitchell Text Book (second reference). More important, it contains the equations that we presented in the class, but not included in the slide.

K-NN Theory

K-NN Lab Tutorial

Lab Assignment 2: K-Nearest Neighbor

Linear Algebra for Machine Learning

Basic Vector/Matrix operations

Addition and scaling
Product: Hadamard product, Inner product, Outer Product

Vector-Vector Multiplication --- Inner product and Outer product interpretation
Matrix-Vector Multiplication --- ""
Matrix-Matrix Multiplication --- ""
Interpreting Neural Network using Linear Algebra: What do network weights convey?
Single-Output Linear Network via Vector-Vector Inner product
Multi-Output Linear Network via Matrix-Vector Inner product
Special Matrices and Eigen values/vectors

Diagonal and Identity Matrices
Inverse and Pseudo-Inverse Matrices
Matrix transformations and their use in Image dataset augmentation (translation, rotation, scaling, sheering, etc)
Determinant & Systems of algebraic equations
Eigen values and vectors. Their calculation.
Reading exercise: https://minireference.com/static/tutorials/linear_algebra_in_4_pages.pdf

Linear Algebra Slides

Theory + Lab: Linear Models for Classification and Regression

Linear Models (Theory: 2:10 PM to 3:25 PM)

Linear Classifier

Linear Mapping
Signum function for binary class
Threshold for multiple classes

Linear Regressor or Perceptron Algorithm

Linear Mapping (no thresholds)
Importance of Bias

Optimizing parameters of (Incrementally) Linear Mapping

Closed form solution
Gradient Descent

Linear Models Tutorial (Theory: 3:35 PM to 4:55 PM)

Linear Classifier using Scikit Learn
Linear Regressor using Scikit Learn
Reading Exercise : Chapter 10.1-10.6 of the Ethem Alpaydin's text book!

Slides: Theory

Linear Regression Lab Tutorial

Lab Assignment 3: Linear Regression

Linear Models (Pre- and Post- Inner product)

Probability for ML and Naive Bayes Model

Basics of Probability
Probability and Random Variables
Probability Distribution (Joint, Marginal, and Conditional)
Bayes' Theorem
Independence
Mean, Variance, and Covariance

Naive Bayes' Model
Main assumption: Independence. Why?
Naive Bayes' training
Numerical Stability
Merits and Demerits
Reading Exercise : Appendix and Secs 3.1 and 3.4 of the Ethem Alpaydin's text book! Also for Naive Bayes, refer to Sec. 6.9 and Sec 6.10 of Tom M Mitchell's text book.

Probability_naivebayes
Naive_bayes_tutorial
Naive_Bayes Assignment

Bias-Variance Tradeoff

Bias Variance Tradeoff
Derivation of MSE in terms of Variance and Bias^2
Implications: Underfitting and Overfitting
Lasso and Ridge Regression
Ridge: Optimization via Closed form and Gradient Descent
Lasso: Optimization via Gradient Descent
Intro to Boosting and Bagging
Reading Exercise : Secs 4.6 and 4.9 of the Ethem Alpaydin's text book!

Bias_variance_tradeoff_ons
Bagging and Boosting
Bias-variance tradeoff detailed proof

Decision Tree

Decision Tree for Classification
Motivation for Decision Tree: Interpretability
Elements of Decision Tree: Root, Nodes, and Leaves
When to Split a Node?
Concept of Impurity -- Which feature to consider in a given node
Entropy Measure
Dealing with categorical as well as numerical features
Rule Extraction from Decision Trees
Other Impurity Measures
Reading Exercise: Chapter 9, Ethem Alpaydin.

Decision Tree - Part 1
Decision Tree - Part 2
Decision tree Tutorial

MLE and MAP

MLE_MAP

Midsem Lab Test

https://www.kaggle.com/datasets/uciml/pima-indians-diabetes-database?resource=download

Midsem Lab Test

Support Vector Machines

Intuition for maximal margin classifiers
Functional and geometric margins
Deriving maximal margin classifier
Primal forms
Lagrange duality and KKT conditions
Dual form
Solution in the form of support vectors
Non-linear SVMs
Kernel trick
Commonly used kernels
Kernel SVM
Soft-margin SVMs

Lecture slides: SVM

Jupyter notebook on support vector machines

Lab Assignment: Support Vector Machines

Feedforward Neural Networks

Artificial neurons as building blocks of neural networks
The perceptron algorithm
Limitation of single layer perceptron and multi-layer perceptrons'
Representation power of multi-layer perceptron and need for deeper network architectures
Activation functions
Feedforward neural networks as universal function approximators
Learning network parameters through backpropagation

Lecture slides: Feedforward Neural Networks

Jupyter notebook for feedforward neural networks

Convolutional and Recurrent Neural Networks

Early Motivations for CNNs
Precursors to modern CNNs
Feature design vs. feature learning
Convolutional filters and convolutional layers
Pooling layers
Hierarchical abstractions learned by CNNs
Applications of CNNs
Modeling sequential data with RNNs
Vanilla RNNs
Backpropagation through time (BPTT)
Vanishing and exploding gradient issues in BPTT
Long-short term memory (LSTM)
Bidirectional RNNs

Lecture slides: Convolutional and Recurrent Neural Networks

Jupyter notebook on CNN implementation

Lab Assignment - Neural Networks

Ensemble Learning

Introduction
Error analysis with ensembling
Bagging
Bias-variance trade-off with bagging
Random Forests
Out-of-bag error and feature importance
Boosting
Adaptive Boosting (Adaboost)
Face detection with Adaboost

Lecture slides: Ensemble learning

Jupyter notebook on ensemble learning

Lab Assignment: Ensemble Learning

Clustering

Introduction to Clustering
Kmeans algorithm
Evaluation of Clustering
Agglomerative Hierarchical Clustering
Gaussian Mixture Model and EM Algorithm

Slide: Introduction to Clustering and kmeans

Slide: Evaluation of Clsutering

Slide: Agglomerative Hierarchical Clustering

Slide: Gaussian Mixture Model and EM Algorithm

PCA and LDA

Principal Component Analysis

PCA and LDA

Final Lab Test

Final Lab Test

Hidden Markov Model

Definiton, Components

Evaluation problem and details

Decoding problem and Viterbi algorithm

Bief overview of Learning Task

Hidden Markov Model

Hidden from students

Hidden Markov Model Updated Slides

Large Language Model Introduction

Not in endsem syllabus

LLM and Chatgpt

Class Test 2 Scores

Class Test 2 Scores

Moodle Docs for this page

You are logged in as Dr. Mahesh Mohan MR (Log out)
Home

Data retention summary