CSCI5525: Advanced Machine Learning

Ju Sun
Computer Science & Engineering
University of Minnesota, Twin Cities

January 20, 2021

What 5525 is about?

advanced introduction to modern machine learning

What 5525 is about?

advanced introduction to modern machine learning

- what: state-of-the-art ML methods

What 5525 is about?

advanced introduction to modern machine learning

- what: state-of-the-art ML methods
- why: they work
  * key ideas/intuitions
  * basic theoretical insights

  $\Longrightarrow$ mathematical

Difference from 5521:

5521 focuses on classic topics and tell what and not much into why

1. Intro: What is Machine Learning (Chap 1)
   HW0 due Friday Sept 11 ##
2. Supervised Learning: Some basic concepts (Chap 2)
3. Bayes Decision Theory: Conditional Probability (Chap 3)
   . . . Discriminant Functions, Normal Dist. (Chap 3)
   HW1 due Friday September 25 ##
4. Estimating Unknown Probability Densities, (Chap 4)
   . . . Parametric Classification (Chap 4)
5. Multivariate Methods: estimation and classification (Chap 5)
   Take Home Exam 1, due Friday October 16. ##
6. . . . *continued*
7. Dimensionality Reduction: feature selection PCA (Chap 6)
8. Unsupervised Clustering: K-means EM (Chap 7)
   HW2, now due Friday October 30 ##
9. Support Vector Machines, (Linear and Kernel) (chap 13)
10. . . . *continued*
    HW3 due Friday November 13 ##
11. Linear Discriminant - the Perceptron (Chap 10)
    Take Home Exam 2, due November 17 ##
12. Multilayer Perceptrons (Chap 11)
    . . . *continued*
13. Decision trees; (Chap 9) random forests (bagging)
    HW4, due Wednesday December 2 ##
14. Review
    HW5 due Thursday December 10 ## NEW
15. *Take Home Exam 3, due Tuesday December 15*

Prerequisites

- highly recommended: CSCI5521 or 5523 or equivalents (may become official Prerequisites in the future)
- mathematics: maturity in **multivariate calculus**, **linear algebra/matrix analysis**, **probability**, and hopefully also **optimization**
- programming: Python 3 (Numpy, Scipy), possibly a bit of deep learning (PyTorch or Tensorflow)

Tentative topics

Overview
Linear prediction
Supervised learning and generalization
Support vector machines and kernel methods
Boosting
Decision trees
Neural networks
Dimension reduction: PCA and kernel PCA
Unlock relations: ICA and CCA
Clustering (spectral, hierarchical)
Generative models (GMM, GAN, variational autoencoder)
Online/Reinforcement learning

About deep learning

We'll be brief to minimize the overlap with the deep learning course

Logistics

- **When & Where:** Mon/Wed 1:00–2:15pm, Keller 3-210
- **Who:**

  Professor Ju Sun (Instructor) Email: jusun@umn.edu
  Le Peng (TA) Email: peng0347@umn.edu
  Tiancong Chen (TA) Email: chen6271@umn.edu

- **Office Hours:** Tue/Thur 5–6pm (Instructor) Mon/Wed 3–4pm (Le) Fri 3–4pm (Tiancong)
- **Course Website:** https://sunju.org/teach/ML-Spring-2021/ All course materials, including course schedule, lecture slides and notes, homework assignments, project description, will be posted on the course website. Enrolled students are encouraged to check the website on a regular basis.
- **Communication—minimizing emails: Canvas** is the preferred and most efficient way of communication. Please post all questions and discussions related to the course in Canvas instead of sending emails. If you have to use emails, please begin the subject line with "CSCI 5525", and expect delay in response.

Logistics

- **When & Where:** Mon/Wed 1:00–2:15pm, Keller 3-210
- **Who:**

Professor Ju Sun (Instructor) Email: jusun@umn.edu
Le Peng (TA) Email: peng0347@umn.edu
Tiancong Chen (TA) Email: chen6271@umn.edu

- **Office Hours:** Tue/Thur 5–6pm (Instructor) Mon/Wed 3–4pm (Le) Fri 3–4pm (Tiancong)
- **Course Website:** https://sunju.org/teach/ML-Spring-2021/ All course materials, including course schedule, lecture slides and notes, homework assignments, project description, will be posted on the course website. Enrolled students are encouraged to check the website on a regular basis.
- **Communication—minimizing emails:** **Canvas** is the preferred and most efficient way of communication. Please post all questions and discussions related to the course in Canvas instead of sending emails. If you have to use emails, please begin the subject line with "CSCI 5525", and expect delay in response.

Teaching format

- UNITE (https://cse.umn.edu/unite) will live stream the lecture and also maintain recorded lectures in both UNITE and Canvas

Teaching format

- UNITE (https://cse.umn.edu/unite) will live stream the lecture and also maintain recorded lectures in both UNITE and Canvas
- Caveat 1: there could be 30–60 secs latency; we may use Youtube later which is better

Teaching format

- UNITE (https://cse.umn.edu/unite) will live stream the lecture and also maintain recorded lectures in both UNITE and Canvas
- Caveat 1: there could be 30–60 secs latency; we may use Youtube later which is better
- Caveat 2: mostly whiteboard presentations; Zoom participation is not ideal

References

Advanced machine learning textbooks

Foundations of Machine Learning (2e) by Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. https://cs.nyu.edu/~mohri/mlbook/

Understanding Machine Learning: From Theory to Algorithms by Shai Shalev-Shwartz and Shai Ben-David. https://www.cs.huji.ac.il/~shais/UnderstandingMachineLearning/

The Elements of Statistical Learning: Data Mining, Inference, and Prediction by Trevor Hastie, Robert Tibshirani, and Jerome Friedman. https://web.stanford.edu/~hastie/ElemStatLearn/

Machine Learning : A Probabilistic Perspective by Kevin P. Murphy. click here (UMN library access; login required)

Mathematics foundations

Mathematics for Machine Learning by Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong. https://mml-book.github.io/

Linear Algebra and Optimization for Machine Learning by Charu C. Aggarwal. https://rd.springer.com/book/10.1007/978-3-030-40344-7 (UMN library access; login required)

Hand-on in Python

Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems by Aurélien Géron (2ed). O'Reilly Media, 2019. click here (UMN library access; login required)

Assessment

- Homework 60%: 6 Homeworks (12% each), the top 5 scores will count toward the final grade. LaTeXing not required but strongly encouraged

Assessment

- Homework 60%: 6 Homeworks (12% each), the top 5 scores will count toward the final grade. LaTeXing not required but strongly encouraged
- Takehome mid-term 15%

Assessment

- Homework 60%: 6 Homeworks (12% each), the top 5 scores will count toward the final grade. LaTeXing not required but strongly encouraged
- Takehome mid-term 15%
- Group course project 25% (2 students): proposal (5%) + final report (20%). The project can be survey of a chosen topic not covered in detail in the class, implementation and comparison of existing methods, or novel foundational or applied research