

10716: Advanced Machine Learning: Theory and Methods

Instructor:

Pradeep Ravikumar, pradeepr@cs.cmu.edu

Teaching Assistants:

Leqi Liu <leqil@cs.cmu.edu>

Boxiang (Shawn) Lyu <blyu@andrew.cmu.edu>

Ritika Mulagalapalli <rmulagal@andrew.cmu.edu>

Karthika Nair <knair@andrew.cmu.edu>

Biswajit Paria <bparia@cs.cmu.edu>

Yao-Hung Tsai <yaohungt@cs.cmu.edu>

Lectures: POS 160, Tuesdays and Thursdays, 1:30 – 2:50 PM

Office Hours:

Pradeep Ravikumar: GHC 8111, Thursdays 3:00 – 4:00 PM

The office hours of the TAs will be posted on Piazza.

Course Description:

Advanced Machine Learning is a graduate level course introducing the theoretical foundations of modern machine learning, as well as advanced methods and frameworks used in modern machine learning. The course assumes that students have taken graduate level introductory courses in machine learning (Introduction to Machine Learning, 10-701 or 10-715), as well as Statistics (Intermediate Statistics, 36-700 or 36-705). The course treats both the art of designing good learning algorithms, as well as the science of analyzing an algorithm's computational and statistical properties and performance guarantees. Theorems are presented together with practical aspects of methodology and intuition to help students develop tools for selecting appropriate methods and approaches to problems in their own research. We will cover theoretical foundation topics such as computational and statistical convergence rates, minimax estimation, and concentration of measure. We will also cover advanced machine learning methods such as nonparametric density estimation, nonparametric regression, and Bayesian estimation, as well as advanced frameworks such as privacy, causality, and stochastic learning algorithms.

Pre-requisites:

Basics of probability and statistical inference, and basics of machine learning (such as regression, classification, clustering).

This course is for students who have already taken introductory courses in machine learning and statistics, and who are interested in deeper theoretical foundations of machine learning, as well as advanced methods and frameworks used in modern machine learning.

Course Goals

1. Understand statistical and computational considerations in machine learning methods.
2. Develop the skill of devising computationally efficient and yet statistically rigorous algorithms for solving machine learning problems.
3. Understand the science of modern statistical analysis.
4. Develop the skill of quantifying the statistical performance of any new machine learning method.

Key Topics:

Statistical Decision Theory, Non-asymptotic Statistical Analysis, Computational Complexity of Estimation, Computational vs Statistical Efficiency Tradeoffs, Optimization and Estimation, Online Learning, High-dimensional Prediction, Nonparametric Prediction, Nonparametric Bayesian Inference, Deep Density Estimation, Nonparametric Density Estimation

Class Website:

<http://www.cs.cmu.edu/~pradeepr/716>

The class schedule, and lecture materials will be posted there.

Discussion, Announcements:

We will use **Piazza** for announcements, providing resource materials, as well as the discussion board for the class.

Textbooks:

Lectures are intended to be self-contained. The following references might be useful:

- JB: Statistical Decision Theory and Bayesian Analysis, by James O. Berger
- MW: High-Dimensional Statistics: A Non-Asymptotic Viewpoint, by Martin J. Wainwright
- BL: Prediction, Learning, and Games, by Nicolo Cesa-Bianchi, Gabor Lugosi
- W: All of Nonparametric Statistics, by Larry Wasserman
- AB: Computational Complexity: A Modern Approach, by Sanjeev Arora, Boaz Barak
- N: Introductory Lectures on Convex Optimization, by Yurii Nesterov

Homeworks:

There will be 5 homework assignments, approximately evenly spaced throughout the semester. The assignments will be posted on the course website, and on Piazza. We will use Gradescope for submitting, and grading assignments. You will get a late day quota of 10 days, which you can distribute among the five homeworks as you wish, subject to a maximum of 3 days per homework. Homeworks submitted after your late day quota will lose all points. The homework schedule is posted right at the beginning of the semester, so please plan in advance. We expect you to use the late day quota for conference deadlines and events of the like, so we cannot provide an additional extension for such cases. In the case of an emergency (sudden sickness, family problems, etc.), we can give you a reasonable extension. But we emphasize that this is reserved for true emergencies.

Collaboration Policy:

The homeworks are structured to give you experience in written mathematical exercises. While it is completely acceptable for you to collaborate with other students in order to solve the problems, we assume that you will be taking full responsibility in terms of writing up your own solution: failing to do this will count as an academic integrity violation and will be reported to CMU authorities. You must indicate on each homework the students with whom you collaborated.

Exams:

There will be two exams, one scheduled halfway through, and the other at the last scheduled class for the semester. The precise dates are on the course website. The exam will consist of multiple choice and true/false questions, as well as short-answer questions and will be open book and open notes.

Class project:

There will be a class project. You can form groups of up to 2 students. Further details on the project will be provided on the website.

Grading:

Exam 1	15%
Exam 2	15%
Homeworks	50%
Project	20%

Accommodations for Students with Disabilities:

If you have a disability and are registered with the Office of Disability Resources, I encourage you to use their online system to notify me of your accommodations and discuss your needs with me as early in the semester as possible. I will work with you to ensure that accommodations are provided as appropriate. If you suspect that you may have a disability and would benefit from accommodations but are not yet registered with the Office of Disability Resources, I encourage you to contact them at access@andrew.cmu.edu.

Take care of yourself:

Take care of yourself. Do your best to maintain a healthy lifestyle this semester by eating well, exercising, avoiding drugs and alcohol, getting enough sleep and taking some time to relax. This will help you achieve your goals and cope with stress.

All of us benefit from support during times of struggle. There are many helpful resources available on campus and an important part of the college experience is learning how to ask for help. Asking for support sooner rather than later is almost always helpful.

If you or anyone you know experiences any academic stress, difficult life events, or feelings like anxiety or depression, we strongly encourage you to seek support. Counseling and Psychological Services (CaPS) is here to help: call [412-268-2922](tel:412-268-2922) and visit their website at <http://www.cmu.edu/counseling/>. Consider reaching out to a friend, faculty or family member you trust for help getting connected to the support that can help.

If you or someone you know is feeling suicidal or in danger of self-harm, call someone immediately, day or night:

- CaPS: [412-268-2922](tel:412-268-2922)
- Re:solve Crisis Network: [888-796-8226](tel:888-796-8226)

If the situation is life threatening, call the police

- On campus: CMU Police: [412-268-2323](tel:412-268-2323)
- Off campus: 911

If you have questions about this or your coursework, please let me know. Thank you, and have a great semester.