

COMP 640: Graduate Seminar in Machine Learning

- **Instructor:** Anshumali Shrivastava (anshumali AT rice.edu)
- **Class Timings:** Monday 3pm-4:30pm
- **Location:** Duncan Hall 3076
- **Office Hours:** Monday 4:30pm - 5:30pm, Duncan Hall 3118

Structure

This research seminar is intended to discuss recent advances and trends in machine learning. We will be presenting and discussing 1-2 recent technical papers each week. The focus will be on modern techniques for massive datasets. The aim is to understand the key ideas and concepts with an aim of generalizing them and stimulating research. Whenever necessary, some concepts will be introduced for clarification and to make connections.

Grading and Logistics

Class participation and one presentation for 1 credit. In addition students can undergo a semester long project for 3 credits.

Prerequisite

Familiarity with basics in linear algebra, probability, and machine learning is required.

Paper Assignment

[Paper Assignments](#)

Schedule

- 08/24 : Introduction, Logistics. [slides](#)
- 08/31 : Locality Sensitive Hashing, Sublinear Algorithms for Approximate Near Neighbor Search. [slides](#)
 - Book: Mining Massive Datasets (Chapter 3) [pdf](#)
 - **Optional** : Near-Optimal Hashing Algorithms for Approximate Nearest Neighbor in High Dimensions (CACM) [pdf](#) (Skip Section 4)
 - **Optional** : Similarity Search in High Dimensions via Hashing [pdf](#)
- 09/7 : **Labor Day**
- 09/14 : SVMs, Kernels and Random Projections
 - Support-Vector Networks [pdf slides](#)
 - Random Projection, Margins, Kernels, and Feature-Selection [pdf slides](#)
 - **Optional** : An Elementary Proof of a Theorem of Johnson and Lindenstrauss [pdf](#)
- 09/21 : Fast Random Projections
 - Very Sparse Random Projections [pdf](#)
 - Approximate Nearest Neighbors and the Fast Johnson-Lindenstrauss Transform [pdf](#) (Only Sections 1 and 2)
- 09/28 : Hashing for Large Scale Learning (**Timing Change 5:30pm - 7:00 pm**)

- Hash Kernels [pdf](#)
- Hashing Algorithms for Large Scale Learning [pdf](#)
- 10/5 : Fast Minwise Hashing
 - Densifying One Permutation Hashing via Rotation for Fast Near Neighbor Search [pdf](#)
 - Improved Densification of One Permutation Hashing [pdf](#)
- 10/12 : **Midterm Recess**
- 10/19 : **No Class**
- 10/26 : Applications
 - Randomized Algorithms and NLP: Using Locality Sensitive Hash Functions for High Speed Noun Clustering [pdf](#)
 - Fast, Accurate Detection of 100,000 Object Classes on a Single Machine [pdf](#)
- 10/26 : **Midterm Project Presentations**
- 11/2 : Deep Networks
 - On the difficulty of training Recurrent Neural Networks [pdf](#)
 - Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift [pdf](#)
 - We will assume knowledge of the vanilla Backpropagation Algorithm for Deep Networks [wiki](#)
- 11/09 : **Guest Speaker: Ankit Patel** "A Probabilistic Theory of Deep Learning" [pdf](#)
- 11/16 : Topic Models and Variational Inference
 - Latent Dirichlet Allocation [pdf](#)
 - Stochastic Variational Inference [pdf](#)
- 11/23 : Scalable MCMC and Sampling
 - An Architecture for Parallel Topic Models [pdf](#)
 - Reducing the Sampling Complexity of Topic Models [pdf](#)
- 11/24 : **Final Project Presentations (Timing Change 6:00pm - 7:30 pm)**

Students with Disability

If you have a documented disability that may affect academic performance, you should: 1) make sure this documentation is on file with Disability Support Services (Allen Center, Room 111 / adarice@rice.edu / x5841) to determine the accommodations you need; and 2) meet with me to discuss your accommodation needs.