Mathematical Algorithms for Artificial Intelligence and Big Data Analysis (Spring 2017)

Course: MAT 180
CRN:
Title: Mathematical Algorithms for Artificial Intelligence and Big Data Analysis
Class: MWF 1:10pm-2:00pm, location: MSB 1147

Instructor: Thomas Strohmer
Office: 3144 MSB
Email:"my last name" at math.ucdavis.edu
Office Hours: TBA

Course Objective:

> Experiments, observations, and numerical simulations in many areas of science nowadays generate massive amounts of data. This rapid growth heralds an era of "data-centric science," which requires new paradigms addressing how data are acquired, processed, distributed, and analyzed. This course will cover mathematical concepts and algorithms that can deal with some of the challenges posed by Artificial Intelligence and Big Data. The course may be of interest to all those students dreaming of becoming a billionaire by starting the next Google--or if you simply want to get a job at Google, Facebook, Twitter, one of the hundreds of startups in AI, ...

Prerequisite:

> Linear algebra as well as basic experience in programming (preferably Matlab) will be required. Some basic knowledge in probability and optimization is helpful but not required. This class targets seniors or advanced juniors with knowledge how to write proofs, say at the level of MAT 125A.

List of topics: (subject to minor changes)

- Brief overview of the aims of Artificial Intelligence and Machine Learning
- Principal Component Analysis, Singular Value Decomposition.
- Curse and Blessings of dimensionality. Strange phenomena in high dimensions
- Data clustering, k-means
- Dimension reduction. Johnson-Lindenstrauss, sketching, random projections.
- Graph Laplacian, spectral clustering, nonlinear dimension reduction.
- Diffusion maps, manifold learning, intrinsic geometry of massive data sets.
- Basic concepts of Deep Learning

Textbooks:

There is no required textbook. The following books contains some material on these topics. But there is definitely no need to buy these books. I will list later some material that is available for free online.

- C. Bishop. Pattern Recognition and Machine Learning.
- F. Cucker, D. X. Zho. Learning Theory: an approximation theory viewpoint.
- T. Hastie, R. Tibshirani, and J. Friedman, The Elements of Statistical Learning: Data Mining, Inference and Prediction.
- Yoshua Bengio. Learning Deep Architectures for AI.

Grading Scheme:

- 50% Homework
- 50% Final Project

Homework:

I will assign homework about every other week. A subset of these problems will be graded. The homework will be announced at the homework page. Late homework will not be accepted.

Final Project:

For the Final Project you need to write a report on one of the following topics:

- Find a practical application yourself (not copying from papers/books) using the methods you learned in this course; describe how to use them; describe the importance of that application; what impact would you expect if you are successful?
- A report describing a thorough numerical comparison of existing algorithms related to one of the topics of this couse for a specific application or problem.