CS5270: NUMERICAL LINEAR ALGEBRA FOR DATA ANALYSIS

Jan-Apr 2015

INSTRUCTOR: Vineeth N Balasubramanian

Objective:
An in-depth understanding of many important linear algebra techniques and their applications in data mining, machine learning, pattern recognition, and information retrieval.

Description:
Much of machine learning and data analysis is based on Linear Algebra. Often, the prediction is a function of a dot product between the parameter vector and the feature vector. This essentially assumes some kind of independence between the features. In contrast, matrix parameters can be used to learn interrelations between features: The (i,j)th entry of the parameter matrix represents how feature i is related to feature j. This richer modeling has become very popular. In some applications, like PCA and collaborative filtering, the explicit goal is inference of a matrix parameter. Yet in others, like direction learning and topic modeling, the matrix parameter instead pops up in the algorithms as the natural tool to represent uncertainty. The emergence of large matrices in many applications has brought with it a slew of new algorithms and tools. Over the past few years, matrix analysis and numerical linear algebra on large matrices has become a thriving field. The course also introduces iterative methods for solving linear equations systems and eigenvalue problems of large dimensions. For more moderate size problems matrix factorization methods and their implementation are introduced.

Sample Topics:

- Review of relevant properties of matrices as transformations; Orthogonal transformations; Matrix and Vector Norms
- Singular Value Decomposition; Principal Component Analysis, Linear Discriminant Analysis
- QR Decomposition; Least Squares Problem and Linear Regression; Reduced Least Squares Problem; Gaussian elimination and other related direct methods for linear systems, iterative methods for linear systems, numerical solution of linear least-squares problems
- Eigenvalues, -vectors and eigenvalue problems in dimensionality reduction: the power method, the bisection method, and the QR method
- From Gauss to LU factorization, Positive definite matrices and Cholesky factorization
- Tensor decomposition; Spectral relaxation of K-Means clustering, and matrix approximations via clustering; Non-negative matrix factorization and Applications in Clustering and Recommender Systems
- Linear Bandits and Matrix Completion Methods
- Sparse linear algebra: Sparse matrices and sparse solutions
- Mining the web: PageRank

Pre-requisites:

- Prior coursework in Calculus, Linear Algebra
- Programming experience (for course projects)

Relevant Resources/References:

- Numerical Linear Algebra by Lloyd N. Trefethen and David Bau, III, SIAM, Philadelphia, ISBN 0-89871-361-7.
- Applied Numerical Linear Algebra by J. Demmel, published by SIAM, 1997.

- Matrix Methods in Data Mining and Pattern Recognition, by Lars Elden, SIAM, 2007.
- Gene H. Golub and Charles F. Van Loan, Matrix Computations, The Johns Hopkins University Press, any edition
- An Introduction to Numerical Linear Algebra, C.G. Cullen, PWS, Boston
- The Matrix Cookbook, by Kaare B. Petersen and Michael S. Pedersen. Available on-line at http://www2.imm.dtu.dk/pubdb/views/publication_details.php?id=3274.
- Pattern Recognition and Machine Learning, by Christopher M. Bishop, 2006
- The Elements of Statistical Learning: Data Mining, Inference, and Prediction, by T. Hastie, R. Tibshirani, and J. Friedman, 2001