

Applied Mathematics 120: Applied linear algebra and big data

Course web page: <https://canvas.harvard.edu/courses/4766> (*Spring 2016*)

Last updated: May 19, 2016.

Feel free to write call or visit us with any questions.

1 Administrative

Instructor: [Eli Tziperman](mailto:eli@seas.harvard.edu) (eli at seas.harvard.edu);

TFs: Please see course web page,

Day & time: Tue,Thu 10:00-11:30

Location: Jefferson 250

Sections: see course web page for times and places.

1st meeting: Tuesday, Jan 26, 2016

Office hours: Eli: Please see course web page for times; 24 Oxford, museum building, 4th floor, room 456. TFs: please see course web page.

Textbooks: Page or section numbers from the relevant textbook for any given lecture are given in the detailed syllabus below. The main textbooks to be used are:

Str Strang, G., *Linear Algebra and its Applications*, 4th ed., 2005.

MMD J. Leskovec, A. Rajaraman, J. D. Ullman, *Mining of Massive Datasets*,
[download](#), see also [on-line MOOC](#)

Other textbooks used:

Mitch Mitchell, Tom, *Machine learning*, McGraw-Hill Science/ Engineering/ Math,
1997, 432 pages.

GMW Practical optimization, Gill, Murray and Wright (1981)

Supplementary materials/ Sources directory: Additional materials from several additional textbooks and other sources, including Matlab programs used in class, may be found [in the Sources directory](#). Follow links below for the specific source material for each lecture. **In order to access these materials from outside the Harvard campus (possibly also from the Harvard wireless network), you'll need to use the VPN software which can be downloaded from the [FAS software download site](#).**

Please note: Course materials are the property of the instructional staff, Harvard University, or other copyright holders, and are provided for your personal use. You may not distribute them or post them on websites without permission of the course instructor.

Prerequisites: Applied Mathematics 21a and 21b, or equivalent; CS50 or equivalent.

Computer Skills: Programming knowledge (CS50 or equivalent experience) is expected for this class, Matlab experience would be particularly helpful. Course demonstrations, sections and homework assignment will be Matlab-based, some of the homework assignments will involve significant code writing effort. Students will gain additional experience with Matlab as part of the course.

Students are asked to download and install Matlab on their computers from the [FAS software download site](#). If you have not been exposed to Matlab or would like a refresher, we recommend the *Matlab boot camp* which involves 3-4 lectures and takes place during the beginning of the term, you need to register in advance at <https://canvas.harvard.edu/courses/1858> and the number of students may be limited.

Sections: Regular times for sections will be scheduled at the beginning of the semester. Each TF will hold a weekly section and have weekly office hours. During the sections, the TFs will discuss and expand on the lecture material and solve additional problems. Although these sections are not mandatory, you are strongly encouraged to attend. Occasionally the TFs will explain material covered in the reading material but not in lectures.

Homework: Homework will be assigned every Tuesday, and will be due the following Tuesday in class unless otherwise noted. The homework assignments are meant to help you better understand the lecture material and introduce you to come important extensions. It is essential that you actively engage in problem solving using the assigned HW and other problems from the course textbooks. Continuously practicing the lecture material on a weekly basis via such problem solving is the only way to become comfortable with the subjects covered in the course.

You are encouraged to post any questions, homework related or others, to the on-line forum on the course web page and contribute answers when relevant.

Collaboration policy: Discussion and the exchange of ideas are essential to doing academic work, and we strongly encourage you to discuss and work on homework problems with other students (and with the teaching staff, of course). However, after discussions with peers, you need to work through the problem yourself and ensure that any answers you submit for evaluation are the result of your own efforts, reflect your own understanding and are written in your own words. In the case of assignments requiring programming, you need to write and use your own code for solving HW problems, code sharing is not allowed. In addition, you must cite any

books, articles, websites, lectures, etc that have helped you with your work using appropriate citation practices.

Quizzes, final, grading: Homework: 40%; three quizzes, *tentatively* scheduled to

1. Tuesday March 1, 7-9pm, Jefferson 250
2. Tuesday March 29, 7-9pm, Jefferson 250
3. Thursday April 21, 7-9pm, William James Hall B-1, 33 Kirkland St.

(all in the evening): 30% together; final: 30%;

Readings: Reading material from the textbooks or other sources will be assigned occasionally. This material will complement the lectures, is therefore an important part of the course, and can be found under the course [downloads](#) web page.

Sections notes: please see course web page

This document:

<http://www.seas.harvard.edu/climate/eli/Courses/APM120/2016spring/detailed-syllabus-apm120.pdf>

Contents

1 Administrative	1
2 Outline	4
3 Syllabus	4
Introduction, overview	4
Linear equations	5
Eigenvalues, eigenvectors	5
Principal component analysis, Singular Value Decomposition	7
Optimization, linear programming	9
Data mining overview	11
Similar items and frequent patterns	11
Unsupervised learning: cluster analysis	12
Supervised learning: classification	15
Modeling and prediction	16
Review	16
4 Useful links	17

2 Outline

Topics in linear algebra which arise frequently in applications, including in the analysis of large data sets: linear equations, eigenvalue problems, principal component analysis, singular value decomposition, quadratic forms, linear inequalities, linear programming, optimization, linear differential equations, modeling and prediction, data mining methods including frequent pattern analysis, classification, clustering, outlier detection. Examples will be given from physical sciences, biology, climate, commerce, internet, image processing, economics and more.

Please see [here](#) for a presentation with a review of example applications.

3 Syllabus

Follow links to see the source material and Matlab demo programs used for each lecture under the appropriate section of the course [downloads](#) web page.

1. INTRODUCTION, OVERVIEW. [sources](#).
We'll discuss some logistics, the course requirements, textbooks, overview of the course, what to expect and what not to expect ([presentation](#)).

2. LINEAR EQUATIONS. [sources](#).

- (a) Notation
- (b) **Motivation**: matrices and linear equations arise in the analysis of electrical network, chemical reactions, large ones arise in network analysis, Leontief economic models, ranking of sports teams (**Str**§2.5 p 133-134; also **Str**§8.4), numerical finite difference solution of PDEs, and more.
- (c) Reminder: row and column geometric interpretations for linear equations $A\mathbf{x} = \mathbf{b}$, $a_{ij}x_j = b_i$ (**Str**§1.2, 2d example on pp 4-5; [Matlab demo](#)). Solving linear equations using Gaussian elimination and back substitution (**Str**§1.3 pp 13-14). Cost (number of operations, **Str**, pp 15-16).
- (d) Solution of large linear systems via direct vs iterative techniques
 - i. Direct method: LU factorization (**Str**§1.5 pp 36-43, [Matlab demo](#); detailed 4×4 example with partial pivoting (row only) on p 3-4 in [notes](#) by Lothar Reichel).
 - ii. Iterative methods: Jacobi, Gauss-Seidel, SOR (**Str**§7.4, pp 405-409; [Matlab example](#); [SOR derivation](#); convergence is further discussed in [notes](#) by RAPETTI-GABELLINI Francesca, and typically systems based on matrices that are either diagonally-dominant, or symmetric positive definite, or both, tend to converge best).
- (e) Does a solution exist and is it sensitive to noise/ round-off error? Two examples from (**Str** p 70) showing the effects of ill conditioned matrix and of using wrong algorithm even with a well conditioned matrix. (Matrix norm and condition number to be discussed later.)
- (f) Dealing with huge systems:
 - i. Special cases: sparse, banded and diagonal matrices ([wikipedia](#) and [Matlab example](#)) [HW: solving tridiagonal systems]. Bad news: LU factorization of a sparse matrix is not necessarily sparse ([Figure](#)), so might be best to use an iterative method to solve the corresponding linear system of eqns.
 - ii. Google's MapReduce (Hadoop) algorithm: general idea (**MMD**§2 intro, pp 21-22). Examples: calculating mean daily flight delay ([Matlab](#)) and corresponding output file; matrix-matrix multiplication using one MapReduce step (**MMD**§2.3.10 pp 39-40, video and text [links](#)); (Time permitting:) the more efficient two step approach (**MMD**§2.3.9).

3. EIGENVALUES, EIGENVECTORS. [sources](#).

- (a) **Motivation**: Google's PageRank; partitioning (clustering) of graphs/ networks; differential equations (**Str**§5.1 p 258-259) and explosive development of weather systems.

- (b) Reminder: eigenvalue problems $A\mathbf{x} = \lambda\mathbf{x}$, finding eigenvalues through $\det(A - \lambda I) = 0$, then finding eigenvectors by solving $(A - \lambda_i I)\vec{e}_i = 0$ (**Str**§5.1, pp 260-261). Similarity transformation and diagonalization of matrices with a full set of eigenvectors (**Str**§5.2, pp 271-273) and of symmetric matrices (**Str** 5S, p 328). Orthogonality and projection of vectors (projection of \vec{b} in the direction of \vec{a} is $(\vec{b} \cdot \vec{i}_a)\vec{i}_a = (|b| \cos \theta)\vec{i}_a$ using the unit vector $\vec{i}_a = \vec{a}/|a|$). Orthonormal base, Gram-Schmidt orthogonalization (**Str**§3.4, pp 195, 200-203).
- (c) Google's PageRank algorithm and finding the first eigenvector efficiently via the power method: first, Google vs BMW: [this](#) and [this](#). Modeling the Internet via a random walker and the PageRank algorithm from p 1-7 [here](#). See [Matlab demo](#). It turns out that PageRank is the eigenvector with the largest eigenvalue of the transition matrix. The theoretical background, proving that there is a PageRank and that it is unique is the Perron-Frobenius theorem stating that a stochastic matrix (each row sums to one) with all positive elements has a single largest eigenvalue equal to one. See Wikipedia for the [theorem](#) and for [stochastic matrices](#);
- (d) The power method: (i) calculating the largest eigenvalue/ vector; (ii) including the calculation of the largest p eigenvalues/ vectors using the block power method; (iii) the inverse power method for calculating the smallest eigenvalue/ eigenvector; (iv) the more efficient shifted inverse power method (**Str**§7.3 pp 396-397; [Matlab demo](#)); it seems that the block method should work only for normal matrices, whose eigenvectors are orthogonal, although Strang does not mention this);
- (e) More on networks and matrices: Transition matrix was covered already as part of the PageRank algorithm above (**MMD** example 5.1, p 166). Adjacency matrix (example 10.16, p 363), Degree matrix (example 10.17, p 364), Laplacian matrix (example 10.18, p 364). [consider also **Str**§8.4 and snowflakes as an example from [snowcrystals.com](#)]
- (f) Spectral clustering (partitioning) of networks via eigenvectors of corresponding Laplacian matrices ([Matlab demo](#) and [notes](#), expanding on **MMD**§10.4.4 and example 10.19, pp 364-367).
- (g) (Time permitting:) Solving large eigenvalue problems efficiently: QR (Gram-Schmidt) factorization and Householder transformations (**Str**§7.3)
- (h) Generalized eigenvalue problems, $A\mathbf{x} = \lambda B\mathbf{x}$, arise in both differential equations and in classification problems (see later in the course). If A, B are symmetric, it is not a good idea to multiply B^{-1} to obtain a standard eigenproblem because $B^{-1}A$ is not necessarily symmetric. Instead, transform to a regular eigenvalue problem using Cholesky decomposition ([Matlab example](#) and [notes](#)).
- (i) Linear ordinary differential equations and matrix exponentiation (**Str**§5.4, pp 294-295, remark on higher order linear eqns on p 296, heat PDE example on p

297-298). Eigenvalues and stability (p 298; phase space plots from Strogatz, Romeo and Juliet).

- (j) Dramatic surprises on the path to tranquility: Non-normal matrices, transient amplification and optimal initial conditions ([notes](#), [Matlab](#)).
- (k) Jordan form and generalized eigenvectors: when straightforward diagonalization using standard eigenvectors doesn't work because they are not independent.
 - i. Start with simple example of the issue using [Matlab demo](#).
 - ii. Definition and statement of the ability to always transform to Jordan normal form (**Str**, 5U p 329-330).
 - iii. Extreme sensitivity to round-off error: HW.
 - iv. Connection to ODEs: exponential of Jordan form, let $\hat{I}_{k \times k}$ be a matrix with 1 above the diagonal, so that one Jordan block is $J = \lambda I + \hat{I}$; now $e^{Jt} = (Ie^{\lambda t})e^{\hat{I}t}$; Noting that $\hat{I}^k = 0$ one finds that $e^{\hat{I}t}$ is given by a finite power series leading to a form as in eqn 10 on p 331.
 - v. Example of second order ODE equivalent to $d\mathbf{x}/dt = J\mathbf{x}$, with $J = [-2, 1; 0, -2]$: $\dot{x} = -2x + y$, $\dot{y} = -2y$; take the derivative of the first $\ddot{x} = -2\dot{x} + \dot{y}$, add twice the first equation to find $\ddot{x} + 2\dot{x} = (-2\dot{x} + \dot{y}) + 2(-2x + y)$ use the second equation to find $\ddot{x} + 4\dot{x} + 4x = 0$; note that this is equivalent to $0 = (d_t - (-2))^2 x = (d_t + 2)(\dot{x} + 2x) = (\ddot{x} + 2\dot{x}) + 2(\dot{x} + 2x) = \ddot{x} + 4\dot{x} + 4x$; Solution is $x = c_1 e^{-2t} + c_2 t e^{-2t}$; To find this, substitute $x = e^{at}$ in eqn to find $a^2 + 4a + 4 = (a + 2)^2 = 0$, so that $a_{1,2} = -2$ and therefore must add the secular term $t e^{-2t}$ as a second solution.
 - vi. How to find the Jordan form using the matrix of generalized eigenvalues (**Str** App B, pp 463-468; detailed [example](#) of a simple case, and [notes](#) on the more general case by Enrico Arbarello).
 - vii. (Time permitting:) Proof by recursion that a Jordan form can always be found is also in **Str** Appendix B.

4. PRINCIPAL COMPONENT ANALYSIS, SINGULAR VALUE DECOMPOSITION. [sources](#).

- (a) **Motivation**: dimension reduction, e.g., image compression, face recognition, El Niño; comparing structure of folded proteins; more unknowns than equations
- (b) Principal Component Analysis (PCA; also known as Factor Analysis or Empirical Orthogonal Functions), calculation from correlation matrix ([notes](#) and [Matlab demos](#), based on based on appendix B of Peixoto and Oort (1992)).
- (c) Example code for PCA analysis in Matlab [using covariance matrix](#) (Example code using SVD also given below).
- (d) Singular Value Decomposition (SVD): statement and examples of SVD decomposition, $X = U\Sigma V^T$ (**Str**§6.3 pp 364-365 including remarks 1,2,4,5 and

examples 1,2; note that $A\mathbf{u}_i = \sigma_i\mathbf{v}_i$ and $A^T\mathbf{v}_i = \sigma_i\mathbf{u}_i$ these are therefore “right and left eigenvectors”). A useful proof for SVD derivation: showing that eigenvectors of a symmetric matrix $A = A^T$ are orthogonal: consider two eigenvectors $A\mathbf{x} = \lambda\mathbf{y}$, $A\mathbf{y} = \mu\mathbf{x}$; then, consider the difference between the following two identical scalars,
 $0 = \mathbf{y}^T A\mathbf{x} - (\mathbf{y}^T A\mathbf{x})^T = \mathbf{y}^T A\mathbf{x} - \mathbf{x}^T A\mathbf{y} = \lambda\mathbf{y}^T \mathbf{x} - \mu\mathbf{x}^T \mathbf{y} = \mathbf{y}^T \mathbf{x}(\lambda - \mu)$, from which we conclude that if the eigenvalues are different, the eigenvectors must be orthogonal.

- (e) Practical hints on calculating SVD: Choose the smallest of $A^T A$ or AA^T ; calculate its eigenvalues and eigenvectors to find the singular values and the smaller set of singular vectors; use $AV = \Sigma U$, or $A^T U = \Sigma V$ to find the first part of the larger set of singular vectors; complete the rest of the larger set by starting with random vectors and using Gram-Schmidt orthogonalization. [Matlab demo](#).
- (f) Geometric interpretation of SVD for the special case of a real square matrix with a positive determinant (see [animation](#) and caption from Wikipedia by Kieff, with some more details [here](#); an additional interpretation is given below in terms of PCA);
- (g) (Time permitting:) [proof of existence](#), not really needed after the above remarks in **Str**.
- (h) SVD applications:
 - i. Image compression, low-rank approximation, (**Str** p 366, [Matlab demo](#)); variance explained (let $X_{n \times m} = f(x, t)$, $x = x_1, \dots, x_n$, $t = t_1, \dots, t_m$; $X^T X = (U\Lambda V^T)^T(U\Lambda V^T) = V\Lambda^2 V^T$; variance is sum of diagonal elements of $C = X^T X/N$, e.g., $C_{ii} = \sum_j X_{ij} X_{ij}/N = \sum_j V_{ij} V_{ji} \Lambda_{ii}^2/N = \Lambda_{ii}^2/N$; total variance is sum of singular values squared, explained variance by first k modes is $\sum_{i=1}^k \Lambda_{ii}^2 / \sum_{i=1}^n \Lambda_{ii}^2$).
 - ii. Effective rank of a matrix (**Str** p 366, [Matlab demo](#)). Matrix condition number and norm (**Str**§7.2, p 388-392).
 - iii. Polar decomposition (**Str** p 366-367). Applications exist in continuum mechanics, robotics, and, our focus here: bioinformatics, the SVD-based Kabsch Algorithm for comparing protein structures using the root-mean-square deviation method ([notes](#) by Lydia E. Kaviraki, p 1-5; [Matlab example](#)).
 - iv. When number of unknowns is different from number of equations:
 - A. Medical tomography as an example application which may lead to either under or over determined systems ([notes](#), section 1).
 - B. Overdetermined systems: more equations than unknown and least squares. (i) Brief reminder ([notes](#), section 2). (ii) Using QR decomposition (cover it first if it was not covered in the eigenvalue/

vector section) for an efficient solution of least-square problems (**Str**§7.3).

C. Under-determined systems, more unknowns than equations: Pseudo inverse solution using SVD and a short proof that it is indeed the smallest-norm solution (section 3 of [notes](#), including Matlab example there; or see **Str** p 369-370).

v. PCA using SVD: [notes](#) based on [Hartmann](#), and example [Matlab code](#) for PCA using SVD).

vi. Multivariate Principal Component Analysis and Maximum Covariance Analysis (MCA): analysis of two co-varying data sets. E.g., M stocks from NY and L stocks from Tokyo, both given for N times: Y_{mn}, T_{ln} .

A. First using “multivariate” PCA of a combined $(M + L) \times N$ data set $F = [Y; T]$ (after removing mean and normalizing each time series by std). Second, MCA, by calculating SVD of the covariance matrix $C = YT^T$ whose elements are $c_{ml} = \sum_{n=1}^N Y_{mn}T_{ln}$. ([notes](#) from Dennis Hartman, p 99-103;

B. Then a simple [Matlab demo](#) of both multivariate PCA and MCA, and a more elaborate [Matlab demo](#)).

C. Finally, the failure of multivariate PCA to analyze correlations between two data sets and the advantages of MCA are demonstrated by the following [Matlab demo](#).

vii. The Netflix challenge part I: latent factor models and SVD (first highlighted text and Figs. 1 and 2 on pp 43-44 of Koren et al. (2009); then highlighted parts of section 6.1 of Vozalis and Margaritis (2006); both available [here](#); Instead of eqn (4) in Vozalis, let the predicted rating of movie a by user j be $pr_{aj} = \sum_{i=1}^n sim_{ji}(rr_{ai} + \bar{r}_a) / (\sum_{i=1}^n |sim_{ji}|)$, where sim_{ji} is the similarity between the ratings of movies i, j by all users, and the sum is over movies)

5. OPTIMIZATION, LINEAR PROGRAMMING. [sources](#). (Time permitting)

Many applications can be formulated as a minimization of a quadratic expression (optimization), or as a set of linear inequalities plus a minimization problem (linear programming).

- (a) **Motivation**: the stable marriage problem, production planning, portfolio selection (**Str**, p 380); the transportation problem, minimization of production costs, minimization of environmental damage and maximization of profits ([here](#))
- (b) Test for the minimum of a general smooth nonlinear multivariate function based on second derivative matrix (Hessian), positive definite matrices, bilinear/quadratic forms (**Str**§6 p 343 after remark 6; p 345-346 from “Higher dimensions: linear algebra” until eq 6, including examples 3 and 5; p 349 beginning of §6.2; p 350 theorem 6B and proofs only until and including sentence after eq 2)

- (c) How to find the minimum of a cost function of many variables: large-scale optimization problems.
- i. A general decent algorithm, requires search direction and line search algorithms (**GMW**§4.3.1)
 - ii. Search direction first try: Steepest descent and its poor convergence ([figure](#), [slide](#), or **GMW**§4.3.2.2).
 - iii. Improved search direction: conjugate gradient algorithm ([notes](#) based on **GMW**§4.8.3 including Fig 4o).
 - iv. Graphical interpretation of clustering of Hessian eigenvalues: cost surfaces of $J = x^2 + y^2$ vs $J = x^2 + y^2/1000$ and corresponding Hessian matrices $\text{diag}(1, 1)$ vs $\text{diag}(1, 1/1000)$; similarly $J = \sum_{i=1}^{100} x_i^2$ which will converge in one step, vs $J = \sum_{i=1}^{50} x_i^2 + \sum_{i=51}^{100} x_i^2/1000$ has two clusters of eigenvalues and would require two iterations to converge. Transformation of coordinates that leads to perfect conditioning: if $J = \mathbf{x}^T H \mathbf{x}$, let $H = R^T R$ and define $y = R x$, then $J = \mathbf{x}^T H \mathbf{x} = \mathbf{x}^T R^T R \mathbf{x} = \mathbf{y}^T I \mathbf{y}$.
 - v. Mention briefly: c-g with restarts (§4.8.3.4); preconditioned c-g using iteratively improving approximation to the Hessian to increase eigenvalue clustering (**GMW**§4.8.5)
 - vi. Line search and step length δ : e.g., assume cost is quadratic along search direction $J(\delta) = J(\delta = 0) + a\delta + b\delta^2$ and use initial cost and gradient plus the cost at some trial step δ_1 to find the expected minimum location. Or use value and gradient at two positions and a cubic approximation.
- (d) Linear programming (**Str**§8, Sultan (1993))
- i. Definition: Minimize $c^T x$ subject to $Ax \geq b$ and $x \geq 0$, where x is n dimensional and A is m rows and n columns. (**Str**p 419)
 - ii. Example applications: Diet problem, portfolio problem and production problem. (**Str** p 417, 418)
 - iii. Preliminaries:
 - A. Inequality constraints, convex region, linear cost;
 - B. Corners as a meeting point of n hyper-planes ($n - 1$ dimensional) at a point, e.g., two lines on the plane or three planes in 3d;
 - C. Feasible solutions (in the allowed subspace) and basic solutions (corners).
 - D. Theorem that minimum and maximum occur at a corner, geometric interpretation
 - E. Geometric method in 2d: find cost at all corners, look for smallest value (first part of the [Matlab demo](#))
 - F. Bringing to standard canonical form, slack variables
 - iv. Simplex algorithm: using the two-phase method

- A. Definition and geometric interpretation of phase I and phase II, moving along edges.
 - B. Phase II using the tableau ([notes](#), section 2; [Matlab](#) program using `linprog` for solving the phase II problem with slack variables, and another [Matlab](#) example code for manipulating the tableau step by step)
 - C. Phase I using slack variables and another cost definition ([notes](#), section 3)
 - v. (Time permitting:) dual problem (**Str**§8.3)
 - vi. Brief mention: Interior point methods
 - vii. Application to network models (**Str**§8.4, p 441-442)
 - viii. Application to the marriage problem (**Str**§8.4, p 443-444)
 - (e) (Time permitting:) Least-square minimization with equality and inequality constraints using SVD (Tziperman and Hecht, 1988).
6. DATA MINING OVERVIEW. [sources](#), [wikipedia](#).
Brief overview of subjects that will be covered in more detail below. [slides](#).
- (a) Similar items and Frequent patterns/ itemsets (association rule learning)
 - (b) Unsupervised learning: classification
 - (c) Supervised learning: clustering
 - (d) (Time permitting:) Outlier/ anomaly detection, Summarization
7. SIMILAR ITEMS AND FREQUENT PATTERNS. [sources](#).
- (a) **Motivation for similar items**: face recognition, fingerprint recognition, comparing texts to find plagiarism, Netflix movie ratings. (**MMD**§3)
 - (b) Similar items:
 - i. Jaccard Similarity index (**MMD**§3.1.1 p 74; [Matlab demo](#) for logicals, numbers, text files).
 - ii. Converting text data to numbers: Shingles, k -shingles, hashing, sets of hashes (**MMD**§3.2 p 77-80; section 1 of [notes](#) and corresponding [Matlab demo](#) of an oversimplified hash function; another [demo](#) for the Cyclic Redundancy Check (crc32) hash function)
 - iii. MinHash algorithm for comparing sets (**MMD**§3.3 p 80-86, and section 2 of [notes](#) with summary of MinHash steps)
 - A. Matrix representation of sets (**MMD**§3.3.1 p 81)
 - B. Minhashing: creating a similarity-conserving signature matrix that is much smaller than the original data matrix, and that allows for an efficient comparison of sets. Signature matrix is based on a set of random permutations of the rows of the data matrix (**MMD**§3.3.2,§3.3.4 p 81-83)

- C. “Proof” that the probability of having similar MinHash signatures of two sets is equal to the Jaccard similarity of the two sets (**MMD**§3.3.3 p 82-83)
 - D. MinHash signature estimated using a set of random hash functions acting on the data matrix (**MMD**§3.3.5 p 83-86)
 - E. Additional resources: [Matlab example](#) of calculating signature matrix and using it to estimate Jaccard similarity; A more elaborate example [python code](#) by Chris McCormick, run using `~/bin/ipython_qtconsole`)
- iv. (Time permitting:) Locality-Sensitive Hashing (LSH, **MMD**§3.4-3.8)
- (c) **Motivation for frequent patterns:** market basket analysis: hot dogs and mustard, diapers and beer; frequent combined Internet searches: Brad and Angelina; medical diagnosis: biomarkers in blood samples and diseases; detecting plagiarism. (**MMD**§6)
- (d) Frequent patterns and association rules.
- i. Mining frequent patterns (and association rule learning): support for set I (number of baskets for which I is a subset); I is frequent if its support is larger than some threshold support s ; (**MMD**§6.1 p 201-206)
 - ii. Association rules $I \rightarrow j$ between a set I and an item j ; confidence (fraction of baskets with I that also contain j) and interest (difference between confidence in $I \rightarrow j$ and fraction of baskets that contain j); (**MMD**§6.1.3-6.1.4)
 - iii. Apriori algorithm: (**MMD**§6.2, highlighted parts on p 209-217)
 - A. Baskets as sets of numbers (**MMD**§6.2.1 p 209)
 - B. Monotonicity of itemsets (**MMD**§6.2.3 p 212)
 - C. A-priori first pass; renumbering of relevant itemsets between passes; and second pass to identify frequent pairs (**MMD**§6.2.5; a simple [Matlab example](#))
 - D. Beyond frequent pairs: larger frequent itemsets (**MMD**§6.2.6)
 - E. Example of finding association rules via A-priori algorithm, [Matlab code](#) by Narine Manukyan, run using `demoAssociationAnalysis`;

8. UNSUPERVISED LEARNING: CLUSTER ANALYSIS. [sources](#).

- (a) **Motivation:** *Archaeology/Anthropology*: group pottery samples in multiple sites according to original culture; *Genetics*: clustering gene expression data groups together genes of similar function, grouping known genes with novel ones reveals function of the novel genes; *TV marketing*: group TV shows into groups likely to be watched by people with similar purchasing habits; *Criminology*: clustering Italian provinces shows that crime is not necessarily linked to geographic

- location (north vs south Italy) as is sometimes believed; *Medical imaging*: measuring volumes of cerebrospinal fluid, white matter, and gray matter from magnetic resonance images (MRI) of the brain using clustering method for texture identification; *Internet/ social networks*: identify communities; *Internet search results*: show relevant related results to a given search beyond using keywords and link analysis; *Weather and climate*: identify consistently re-occurring weather regimes to increase predictability.
- (b) Overview: Two main approaches to clustering: hierarchical (each point is an initial cluster, then clusters are being merged to form larger ones) and point-assignment (starting with points that are cluster representatives, centroids, and then adding other points one by one). Other considerations: Euclidean vs non, and large vs small memory requirements (MMD§7.1.2, p 243).
 - (c) Distances/ metrics:
 - i. Requirements from a distance: MMD§3.5.1, p92-93.
 - ii. Examples of distance functions (MMD§3.5, p 93-97): Euclidean (L_2 distance), L_r distance, Manhattan (sum of abs values, L_1 norm), maximum (L_∞), Hamming distance between two strings of equal length or between vectors of Booleans or other vectors, cosine (difference between angles), Jaccard distance (one minus Jaccard similarity), edit. Noting that “average” distance does not necessarily exist in non Euclidean spaces (p97).
 - (d) Curse of dimensionality: problems with Euclidean distance measures in high dimensions, where random vectors tend to be far from each other and perpendicular to each other, making clustering difficult (MMD§7.1.3 p 244-245, [Matlab demo](#))
 - (e) Hierarchical clustering: intro and example (MMD§7.2.1, Figs 7.2, 7.3, 7.4, 7.5, and the resulting dendrogram in Fig 7.6), efficiency (MMD§7.2.2, p 248-249), merging and stopping criteria (MMD§7.2.3), in non Euclidean spaces using clustroids (MMD§7.2.4, p 252-253). (Use this [Matlab script](#) to run three relevant demos: First detailed hand calculation, than the example script run first with argument (2) and then with (20). The “simpler” version there is a bare bone version that can be useful in HW)
 - (f) K-means algorithms: these are point-assignment/ centroid-based clustering methods. Basics (MMD§7.3.1), initialization (MMD§7.3.2; e.g., initialize centroids on k farthest neighbors), choosing k (MMD§7.3.3). (run two demos using [Matlab script](#), first a detailed hand calculations and then the a more detailed example)
 - (g) Self-organizing maps (a type of an artificial neural network!):
 - i. First the idea and an example application of identifying weather patterns in Figs 1 and 2 (including the one paragraph discussion of the Sammon mapping scheme on p 6357) of Johnson et al. (2008);

- ii. A first [Matlab demo](#) with an explicit demonstration of the algorithm;
- iii. Refinements: (1) allow for more representative points (hence more clusters), still in a 1d configuration; (2) allow a 2d arrangements of the representative points; (3) make learning rate decrease in time; (3) Make neighborhood kernel more restrictive in later iterations of the algorithm, i.e., make adjustment to representative points other than nearest be smaller; (4) make the learning rate at later stages be smaller for points that are further away from the representative point (i.e., make η be a function of the distance from the data point and the representative point).
- iv. Then the appendix (p 6367-6369 and Fig A1) for the actual method and a simple example;
 - v. A [Matlab example](#) using Matlab's SOM routines, demonstrating the method again and showing sensitivity of clustering results to assumed 2d topology.
 - vi. Sammon map of the results to find out which clusters are similar.
- (h) Mahalanobis distance: first for stretched data, diagonal covariance matrix, then non-diagonal, stretched and rotated ([notes](#)).
- (i) Spectral clustering into two sub-clusters. First a reminder of network clustering [notes](#). Then for clustering of other data: Form a distance matrix $s_{ij} = |\mathbf{x}_i - \mathbf{x}_j|$, defined here as the distance between points i and j in the set; then a "similarity" matrix (equivalent to adjacency matrix in network clustering) using, e.g., $w_{ij} = \exp(-s_{ij}^2/\sigma^2)$, then a diagonal degree matrix $d_i = \sum_j w_{ij}$, and finally the Laplacian matrix $L = D - W$ (highlighted parts in p 1-4 of Von-Luxburg (2007); [Matlab demo](#), try first with argument 2 and then 20). Comments: (1) [Wikipedia](#) adds that "The algorithm can be used for hierarchical clustering by repeatedly partitioning the subsets in this fashion". (2) Such clustering using eigenvector 2 was already covered for networks, using the Laplacian matrix of the network, in the eigenvalues/ eigenvectors section. (3) (Time permitting) Can also cluster into $k > 2$ parts using the first k eigenvectors, see Von-Luxburg (2007).
- (j) k -means on large data sets that cannot be fully contained in memory: BFR algorithm and Summarization (**MMD**§7.3.4 and 7.3.5, p 257 to middle of 261)
- (k) CURE (Clustering Using Representatives) algorithm, for clusters that have complex shapes, like concentric rings. This is a point-assignment clustering algorithm, like k -means, not relying on centroids but on a set of representative points that span the cluster, allowing the cluster to take more general shapes (**MMD**§7.4, p 262-265; and a [Matlab demo](#) of the first step via Hierarchical clustering using the appropriate distance measure)
- (l) (Time permitting:) Outlier/ anomaly detection: a brief overview only. Motivation: unusual credit activity as indication of credit card theft. Detection using statistical methods e.g., assuming Gaussian distribution;

- (m) (Time permitting) GRGPF algorithm combining hierarchical and point-assignment approaches, for large data sets (MMD§7.5). Clustering for Streams (MMD§7.6). Simrank; Density-based (DBSCAN).

9. SUPERVISED LEARNING: CLASSIFICATION. [sources](#).

(We stick to Euclidean distances for now, other options were discussed under cluster analysis).

- (a) **Motivation:** Optical character recognition, handwriting recognition, speech recognition, spam filtering, language identification, sentiment analysis of tweets (e.g., angry/ sad/ happy), amazon book recommendation, Netflix challenge, on-line advertising and ad blocking on Internet sites, credit scores, predicting loan defaulting, and Mastering the game of Go!
- (b) Machine learning Introduction (MMD§12.1, p 439-443)
- (c) Perceptrons: Intro (MMD§12.2 p 447); zero threshold (MMD§12.2 first two paragraphs, 12.2.1, p 448-450); allowing threshold to vary (MMD§12.2.4, p 453); problems (MMD§12.2.7, simply show Figs. 12.11,12.12,12.13 on p 457-459). Use [Matlab demo](#), see comments at top of code for useful cases to show; for adjustable step I made step size (η) proportional to deviation of current data point that's not classified correctly ($\eta = |\mathbf{x} \cdot \mathbf{w} - \theta|$), but bounded on both sides, say $0.01 < \eta < 1$.
- (d) Support vector machines:
 - i. Introduction, formulation for separated data sets, formulation for overlapping data sets, solution via gradient method (MMD§12.3, p 461-467);
 - ii. Misc Matlab demos, run all relevant ones using [Matlab script](#).
 - iii. Note: in a case of data composed of two clusters that can be separated perfectly well, it is useful to choose a large $C = 1000$ or so to find an appropriate solution.
- (e) Multi-Layer Artificial Neural Networks (a brief introduction): these are based on a very powerful extension of the perceptron idea, and allow computers to perform image/ voice/ handwriting/ face recognition, as well as [Mastering the game of Go](#). See highlighted parts of the two introductory notes in the relevant [Sources sub-directory](#), and then [Matlab example](#) that demonstrates character recognition using Matlab's neural network toolbox.
- (f) k -nearest neighbors (k -NN):
 - i. Classification: finding a label of input data based on majority of k nearest training data neighbor(s) when label is discrete such as type of dog or sick vs healthy. Start with a single neighbor, including Voronoi diagram (MMD§12.4, p 472-474 including Fig. 12.21; then **Mitch**, Fig. 8.1, p 233

which shows how the results of nearest neighbor can be different from $k = 5$ nearest ones)

- ii. Locally-weighted kernel regression: e.g., estimating house prices as function of age and living area from similar houses (Section 1 of [notes](#) based on [Mitch§8.3.1](#) p 237-238; [Matlab demo](#))
 - iii. (Time permitting:) Using PCA for dimensionality reduction to avoid curse of dimensionality when looking for nearest neighbors in a high-dimensional space. (Section 2 of [notes](#))
 - iv. k -NN application: the Netflix challenge part II (presentation by Atul S. Kulkarni, [remote](#) and [local](#) links).
- (g) Decision trees:
- i. First, definition of entropy in information theory (from Wikipedia, [local](#)).
 - ii. ID3 algorithm: motivation and outline ([Mitch§3.1-3.3](#)); entropy measure ([Mitch§3.4.1.1](#)); information gain ([Mitch§3.4.1.2](#)); ID3 algorithm ([Mitch](#), Table 3.1, p 56); example ([Mitch§3.4.2](#), p 59-61, [here](#)).
 - iii. (Time permitting:) If the potential labeling variable is a continuous number, need to try all possible values to find the one that leads to the maximum entropy gain, as demonstrated for classifying houses into two neighborhoods based on house price and house area in the following [Matlab example](#).
- (h) (Time permitting:) Additional issues:
- i. Avoiding over-fitting, pruning and dealing with continuous variables and thresholds ([Mitch§3.7](#)).
 - ii. C4.5 algorithm for decision trees ([Mitch§3.7](#))
 - iii. Fisher's Linear discriminant analysis (LDA) leading to a generalized eigenvalue problem ([notes](#))
 - iv. From binary classification (two classes) to multiple classes: one vs rest and one vs one strategies ([here](#))
 - v. From linear to nonlinear classification, the kernel trick.
 - vi. Nearest neighbors using k-d trees.

10. MODELING AND PREDICTION. [sources](#).

- (a) **Motivation:** 'predictive analytics'; crowd-sourcing/ crowd wisdom for models: the Netflix challenge part III, global warming, El Niño; Kalman filter: used car pricing, guidance and navigation systems of cruise missiles, weather prediction;
- (b) Combining different methods/ models into a single, more effective, prediction. One example: sequential learning algorithms

11. REVIEW. [sources](#).

4 Useful links

1. http://psi.cse.tamu.edu/teaching/lecture_notes/

References

- Gill, P. E., Murray, W., and Wright, M. H. (1981). *Practical Optimization*. Springer-Verlag, Heidelberg,.
- Johnson, N. C., Feldstein, S. B., and Tremblay, B. (2008). The continuum of northern hemisphere teleconnection patterns and a description of the nao shift with the use of self-organizing maps. *Journal of Climate*, 21(23):6354–6371.
- Koren, Y., Bell, R., and Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37.
- Leskovec, J., Rajaraman, A., and Ullman, J. D. (2014). *Mining of massive datasets*. Cambridge University Press.
- Mitchell, T. (1997). *Machine learning*. McGraw-Hill Science/ Engineering/ Math.
- Peixoto, J. P. and Oort, A. H. (1992). *Physics of Climate*. American Institute of Physics.
- Strang, G. (2006). *Linear algebra and its applications*. 4th ed. Belmont, CA: Thomson, Brooks/Cole.
- Sultan, A. (1993). *Linear programming: an introduction with applications*. Elsevier.
- Tziperman, E. and Hecht, A. (1988). Circulation in the Eastern Levantine Basin determined by inverse methods. *J. Phys. Oceanogr.*, 18(3):506–518.
- Von-Luxburg, U. (2007). A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416.
- Vozalis, M. G. and Margaritis, K. G. (2006). Applying svd on generalized item-based filtering. *IJCSA*, 3(3):27–51.