List of Topics

- Some learning theory: What is unsupervised (supervised) learning? Give an example of each.

- Some Basic Stats

  Know how to compute the mean, variance, and correlation by hand and using Matlab. Be able to double-center a matrix (one way is fine, you don't need to know all the alternatives), and mean-subtract data stored as a matrix (using Matlab). Know how to define a discrete probability density function.

  Know the definition of the covariance matrix. Be able to prove things about the variance and correlation (like in the exercises).

- Linear Regression: Be able to find the line of best fit using the normal equations. Be able to set up linear regression problems for other linear models (for example, if the model equation was $y = c_1 \cos(x) + c_2 \sin(2x)$ instead of $y = mx + b$).

  Matlab questions: Be able to plot distinct points in the plane. Be able to plot a given function (you have to construct a nice domain and range!).

- $n-$armed bandit: Be able to discuss the general problem (including supervised versus unsupervised learning), and the details for how we "solved" the problem. In particular, be able to discuss the greedy algorithm, the $\epsilon-$greedy algorithm, the softmax strategy and the pursuit (or win-stay, lose-shift) method. Understand the technique we used to "select action i with probability P_i" (meaning be able to describe it in words, not using Matlab).

  Be able to analyze what happens to our parameters as we iterate the algorithm (as in the exercises).

  In Matlab: Understand the difference between a *script* and a *function*. Be able to write a script for doing the homework (using double percents to separate sections).

- Basic linear algebra: Be able to define *coordinates* with respect to a particular basis. Given a basis, be able to compute the coordinates for a given vector. Be able to talk about the difference between the "high dimensional representation" of a given vector versus its "low dimensional representation".

  Define an "orthogonal" matrix. Be able to compute the projection of a vector onto another vector, and onto a subspace.

  Be able to compute a projection matrix: For a projection to a vector, For a projection to a subspace spanned by orthonormal columns in a matrix. To project to the column space of a general matrix A.

  Know the four fundamental subspaces associated to a matrix A (be able to draw a diagram like we did in class). Given the rank of A, give the dimensions of all four subspaces. Given a matrix and its rref, be able to construct a basis for the column space, null space and row space.

Show that the null space is orthogonal to the row space. Prove that if Q has orthonormal columns, then $\|Q\mathbf{x}\| = \|\mathbf{x}\|$, and that the dot product between Qx and Qy is the same as the dot product between x and y.

- Eigenvalues/eigenvectors: Recall the three main equations used in eigenvalue/eigenvector computations and proofs. Be able to compute eigenvalues and eigenvectors for various matrices. Know the definition of an eigenspace (E_λ), and understand why it is a subspace (therefore, "solving for an eigenvector" is not well defined). Be able to prove some basic facts about eigenvalues. Is there a relationship between eigenvalues and invertibility?

- The Spectral Theorem: We don't need all of the details here, mainly the following:

  "If and $n \times n$ matrix is symmetric, then it has n real eigenvalues and the set of eigenvectors can be formed so that they form an orthonormal basis for $\mathbb{R}^n$." Additionally, understand the relationship between the rank and the number of non-zero eigenvalues.

- The SVD: Be able to state the Singular Value Decomposition, and define "singular values". Be able to compute the SVD by hand for "simple" matrices. Be able to use Matlab to compute the SVD, and from that, be able to find a basis for all four fundamental subspaces. From that, be able to project vectors to any of the four subspaces.

  In obtaining the reduced SVD, we need to compute the rank of the matrix, which is sometimes not clearly numerically determined. In that case, we looked at the normalized eigenvalues as a method of determining the dimension in terms of the percent variance captured.

- The pseudoinverse: Given the reduced SVD, be able to compute the pseudoinverse. Describe how the pseudoinverse works in terms of the four fundamental subspaces.