10715:Advanced Introduction to Machine Learning

Prerequisites *September 4, 2014*

1 Prerequisites

This course has no official prerequisites. However, we will assume a fair amount of expertise in several topics on Linear Algebra, Multivariate Calculus, Probability and Statistics. This course is intended for PhD students pursuing research in Machine Learning. If you do not have the required background or if Machine Learning is not your main focus you might consider the general Graduate Level Machine Learning course (10-701).

We have outlined some of the prerequisites needed for this course. You should be quite familiar with most of these topics. We have also listed a couple of resources in case you are not familiar with them.

1. Probability
   - Density Functions, Mass Functions
   - Variable Transformations
   - Basic Probability Distributions: Gaussian, Multivariate Gaussian, Exponential, Bernoulli, Binomial, Multinomial, Dirichlet etc.
2. Statistics
   - Point Estimation (e.g. Unbiased Estimators)
   - Linear Regression
3. Multivariate Calculus
   - The gradient and Hessian of multivariate functions.
   - Optimizing simple functions. For instance, you should know $\text{argmin}_\beta \|Y - X\beta\|_2 = (X^\top X)^{-1} X^\top Y$.
4. Linear Algebra
   - Vectors, norms, Inner Products
   - Euclidean Space, Subspaces, Basis
   - Range, Nullspace, Rank, Inverses of a matrix
   - Solving Linear Systems, Least Squares/ Least Norm Solutions
5. Other Topics
   - Taylor Series Expansion
   - Gradient Descent, Newton's Method

2 Resources

Here are some good resources for the above topics.

- Probability and Statistics: Some of Prof. Larry Wasserman's old Statistics notes are a good starting point: http://www.stat.cmu.edu/∼larry/=stat705/

- Linera Algebra: Prof. Zico Kolter has some video lectures to get you started on the basics: http://www.cs.cmu.edu/~zkolter/course/linalg/.
  Also, ML PhD Student Aaditya Ramdas has videos on some advanced topics such as SVD: http://www.cs.cmu.edu/~aramdas/videos.html.

3 Recitations

The TAs will be covering some background material that will not be covered in class. Below is a tentative list of topics.

1. Convex Optimization, Duality
2. Reproducing Kernel Hilbert Spaces
3. Information Theory
4. Bayesian/Markov Networks
5. Spectral Graph Theory