

Textbook:

There is no textbook for the course. We will make all readings available.

Grading Structure:

While there's no guarantee that the same pattern will hold, in the past grad courses that I have taught, the C, B, B+ and A cutoffs have roughly been 60, 70, 80 and 90 respectively. Grading is based on a curve with gaps in the histogram being natural cutoffs.

Prerequisites:

- Basic calculus, limits, integration, partial differentiation, minimization and maximization of functions.
- Basic probability theory including Bayes theorem, joint and conditional distributions, independence, probability density functions.
- Basic statistics including mean, variance, Gaussian (normal) distributions, expectation.
- Basic linear algebra including linear systems, matrix inverse, rank, null space, determinants, trace, eigenvectors and eigenvalues.

Syllabus:

1. **Linear Algebra and Vector Spaces:** *vector spaces, matrix algebra, singular value decomposition (SVD), norms, subspaces, linear least-squares.*
2. **Constrained Optimization:** *Convexity, Lagrange parameters, Karush-Kuhn-Tucker (KKT) conditions.*
3. **Probability Theory:** *Functions of a random variable, maximum likelihood, MAP estimation, basic inference.*
4. **Neural Networks:** *Entropy, divergence measures (Kullback-Leibler and Bregman), nonlinear optimization, multilayer perceptrons.*

Assignments:

Almost all assignments will be due Mondays at midnight.