Announcements

- Last day of classes, you've almost made it!
- Final exam is **6-8pm on December 22**
- Today is the last day of office hours
- We will still respond on Ed Discussion

Lecture 27: Review Part 2

CIS 4190/5190

Fall 2022

Final Exam Tentative Format

- Similar length/format to the practice exam
  - ≈ 15 questions in increasing difficulty
  - Fewer questions but more parts per question
  - Will require **less memorization** than the practice exam
  - But no cheat sheet

- Make sure you know the written homework well!
  - Also questions from practice final exam that we cover today

Final Exam Format

- We will provide any **complicated** equations necessary
  - **You do not need to know:** Entropy, sigmoid function, logistic NLL, neural network model families
  - **You should know:** Linear regression model family, decision tree model family
  - **You should know:** How to compute a derivative, probability identities, etc.
  - **You should know:** K-means clustering algorithm structure, gradient descent algorithm structure, AdaBoost structure (but not the detailed formulas)

- You should also know how different design choices/hyperparameters affect performance of each algorithm
  - E.g., k in kNN, λ in linear/logistic regression, feature dimension d, number of examples n, AdaBoost iterations T, random forest base models k, etc.

Incomplete List of Potential Topics

- **Models/algorithms**
  - What is the model family? How does its decision boundary look?
  - What is the loss function? How does it compare to the “true” loss (e.g., NLL vs. accuracy for logistic regression)?
  - What is the optimizer? Is it guaranteed to find the global optimizer?

Incomplete List of Potential Topics

- **Models/algorithms**
  - Linear/logistic regression
  - KNNs
  - Decision trees
  - Random forests, gradient boosted decision trees
  - Feedforward neural networks, convolutional neural networks
  - K-means clustering
  - PCA
  - Bayesian networks
  - Q iteration, Q learning, epsilon-greedy exploration
  - Collaborative filtering

Incomplete List of Potential Topics

- **Concepts**
  - Supervised vs. unsupervised vs. reinforcement learning
  - Loss minimization framework
  - Maximum likelihood framework
  - Bias-variance tradeoff
  - Regularization
  - Exploration in reinforcement learning
  - ML ethics

Good Luck!!!