Machine Learning

COS 598 - Spring 2020

Tue/Thu 9:30-10:45 am

Instructor: Salimeh Yasaei Sekeh

salimeh.yasaei@umaine.edu

The course addresses the foundations of machine learning with applications in real-world problems. The ML techniques recognize hidden patterns that often provide dramatic competitive advantages with significant computational improvements in learning tasks never imagined before.
By 2020, IT departments will be monitoring 50 times more data than they are today. This tidal wave of data has driven unprecedented demand for those with the skills required to manage and leverage these very large and diverse data sets. Advances in ML are helping to learn the required skills and solve a very broad variety of problems in many fields including: engineering and AI, computer vision, advertising, health care, social media, robotic, economics, security, agriculture, and cool new industries like self-driving drones/cars and highly efficient automated homes.

Our goal in this course is to help you to:

- Understand fundamentals of ML
- Learn technical details of ML algorithms
- Learn how to implement some important algorithms
- Use ML algorithms for your research and applications

Enrollment: This is a graduate level course but qualified undergraduate students are allowed to enroll under instructor's permission.

Note that a good background in the following topics are required:

- Matrix operations, rank, eigenvalue/vector, nullity, linear independence, inner products, orthogonality, positive (semi-) definite matrices, eigenvalue decomposition.
- Jointly distributed random variables, multivariate densities and mass functions, expectation, independence, conditional distributions, Bayes rule, the multivariate normal distribution.
- Partial derivatives, gradients, chain rule

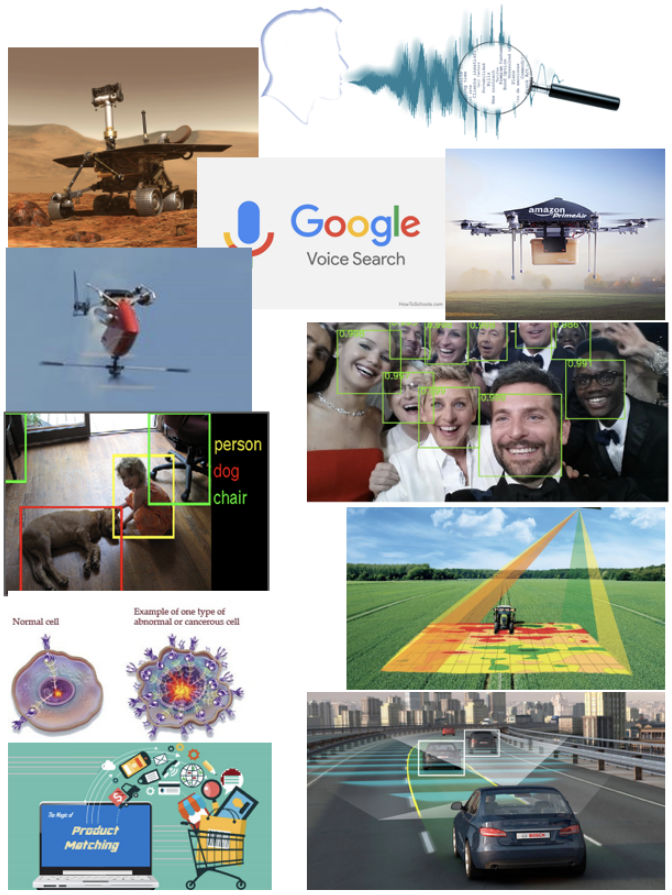

Course organization: The course is split into two parts: 1) ML algorithms and their theoretical analysis, and 2) Applications of ML techniques in real-world problems. The class has a final project, which will provide you with the opportunity to apply the material to an advance topic of your interest or even your own research.

We'll cover:

- **Supervised Learning**
  Classification, Regression, Empirical Risk Minimization, Regularizations, Sparsity (Lasso), Feature Selection, Kernel Methods, Support Vector Machine, etc.
- **Unsupervised Learning**
  Principle Components Analysis (PCA), Clustering, K-Means, Gaussian Mixture Models, The Expectation Maximization Algorithm, Latent variable models, etc.
- **Neural Networks**
  Deep Neural Network, Backpropagation, Convolutional Neural Network, Autoencoder, etc.
- **Decision Trees, Random Forest**
- **Advance Topics** Online Learning, Markov Decision Process, Active Learning, Reinforcement Learning, etc.