

CPSC 540: Machine Learning

Mark Schmidt

Machine Learning

- Automatically detecting patterns in data to:
 - make predictions.
 - make decisions.

Machine Learning

- **Automatically detecting patterns in data to:**
 - make **predictions**.
 - make **decisions**.
- **One of the fastest-growing areas:**
 - we are collecting unprecedented amounts of data.
 - ML is a key tool in making sense of this data.

Machine Learning

- **Automatically detecting patterns in data to:**
 - make **predictions**.
 - make **decisions**.
- **One of the fastest-growing areas:**
 - we are collecting unprecedented amounts of data.
 - ML is a key tool in making sense of this data.
- **Key technology behind:**
 - spam filtering, credit card fraud detection, Amazon product recommendation, Microsoft Kinect, speech recognition, object detection and self-driving cars (coming soon).
 - your graduate research???

CPSC 540

- Big class:
 - 72 students registered (more coming).
 - students from many departments (math, stats, eng., sci.)
 - nice classroom: AERL 120

CPSC 540

- Big class:
 - 72 students registered (more coming).
 - students from many departments (math, stats, eng., sci.)
 - nice classroom: AERL 120
- Course Objective:
 - you will be able to not only use existing machine learning methods, but be able to develop your own machine learning models.
 - deeper than undergrad course: do the math and implementation.

Course Outline

- Classic machine learning material:
 - regression, classification, model selection, regularization, kernels and Gaussian processes, bootstrapping/boosting and random forests, mixtures and latent variable models, missing data and semi-supervised learning.

Course Outline

- Classic machine learning material:
 - regression, classification, model selection, regularization, kernels and Gaussian processes, bootstrapping/boosting and random forests, mixtures and latent variable models, missing data and semi-supervised learning.
- Advanced material:
 - convex and stochastic optimization, graphical models, Bayesian inference, deep learning.

Course Outline

- Classic machine learning material:
 - regression, classification, model selection, regularization, kernels and Gaussian processes, bootstrapping/boosting and random forests, mixtures and latent variable models, missing data and semi-supervised learning.
- Advanced material:
 - convex and stochastic optimization, graphical models, Bayesian inference, deep learning.
- Not covered:
 - effects of actions (causality, active learning, reinforcement learning).
 - limits of learning (VC dimension).

Administrivia

- Prerequisites:
 - multivariate calculus, linear algebra, probability, CS, statistics.
 - weekly **tutorials** (Thursdays, 3-4) to help with the background.
 - possibility to **audit** (or take undergrad course).

Administrivia

- Prerequisites:
 - multivariate calculus, linear algebra, probability, CS, statistics.
 - weekly **tutorials** (Thursdays, 3-4) to help with the background.
 - possibility to **audit** (or take undergrad course).
- Course work:
 - **assignments**: written and Matlab programming
 - **midterm**: on assignment questions.
 - **final project**: small research project.
 - **coding project**: we write a new machine learning package.

Administrivia

- Prerequisites:
 - multivariate calculus, linear algebra, probability, CS, statistics.
 - weekly **tutorials** (Thursdays, 3-4) to help with the background.
 - possibility to **audit** (or take undergrad course).
- Course work:
 - **assignments**: written and Matlab programming
 - **midterm**: on assignment questions.
 - **final project**: small research project.
 - **coding project**: we write a new machine learning package.
- I hope to see you there!