CS 536: Machine Learning

Fall 2003: Michael L. Littman
TA: Yihua Wu

Welcome!

My first class here and first time teaching ML.

- Good news: I'm still energetic
- Bad news: You're my guinea pigs.

Textbook

Machine Learning by Tom M. Mitchell

http://www-2.cs.cmu.edu/afs/cs.cmu.edu/user/mitchell/ftp/mlbook.html

- Chapters 1, 3-10, 13
- Supplementary material on SVMs, clustering
- Feedback from you.

Course Goals

- Introduce the background and "lore" of machine learning.
- Prepare you to be able to use ML tools, contribute to the field.
- Introduce you to the research process.

Not: Find out details of the state of the art…

Course Mechanics

- Problem sets
- Term project, short paper, peer review!
- Midterms
- Final

Repeat with me: "*I am not an undergrad…*"

Mini-Conference (icml-2003)

Choose a research project
Do the research and some background reading
Prepare a 4-page paper (ICML format)
Submit the paper after midterms
Each paper will have 3 anonymous reviewers
Reviewers submit their reviews, discuss
Best ~25% of papers presented to class.
Additional 25% presented as posters.
All authors revise papers based on reviews.

Example: Snowballs

The snowball game: You are having a snowball battle with a neighbor over your backyard fence. You have a catapult with force and angle controls. You need to hit your neighbor's fort 10 times to win.

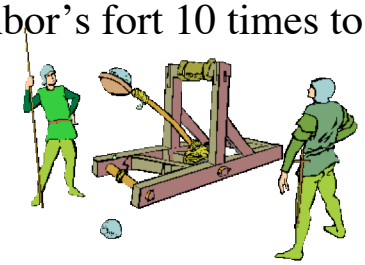

Snowballs: Details

You catapult's controls (force, angle) are set randomly. You decide:

- don't fire (costs 1 unit)
- fire (costs 2 units)
- get distance travelled (costs 5 units)
- check hit or miss (costs 100 units)

How minimize cost to achieve the objective?

Some Types of Learning

Supervised, classification: vector → Boolean

Supervised, regression: vector → real

Supervised, general: vector → vector

Unsupervised, discrete: vector → cluster

Reinforcement, associative: vector → action, via real

Reinforcement, temporal: vector → action, via delayed real

Project Ideas (pg. 1)

Given a web page that (probably) contains glossary entries and definitions, extract the fields.

- http://www.twjc.co.uk/glossary.html
- http://www.cs.unc.edu/~helser/juggler-0.81/glossary.html

Given multiple database with addresses, create a unified database of places.

Create a more accurate battery power indicator.

Extract titles, authors, references from pdf files.

Self organization of a peer-to-peer network.

Project Ideas (pg. 2)

Predict server response time for nodes in a wireless network.

In RL, there are several algorithms that trade off exploration and exploitation in a theoretically motivated way. Evaluate them empirically.

Compare existing RL techniques for "mountain car" or Tetris.

Figure out how to beat a fixed set of TAC agents.

Compare techniques for merging probability distributions theoretically.

Project Ideas (pg. 3)

To solve multiple choice synonym questions, we've shown that multiple experts is a smart way to do this. Training is done using supervised data. Can the multiple modules be used to train each other? ("Labeling via collectives of sufficiently accurate modules").

How about modules for RL? Is there an advantage for doing policy search, table, neural net all together?

Project Ideas (pg. 4)

Applications: Natural language dialog, robotics, financial trading, network diagnosis, object recognition, combining speech and commands and images, problem solving (Sokoban), video game

What data do *you* have?

Meet the Class (pg. 1)

Signed up:
Vasilios Daskalopoulos (CS)
Jixin Li (CS)
Mark Rogaski (CS)
Georgios Sakkis (CS)
Mark Sharp (SCILS)
Alexander Strehl (CS)
Zhi Wei (CS)
Jiankuan Ye (CS)

Permission:
Bing Bai (CS/SCILS)
Chan su Lee (CBIM)
Igor Gierymski (Bus.) X
Lu Liu (SCILS)
Paul Batchis (CS)
Christopher Peery (CS)
Rong Xu (CS)
Rui Huang (CS/CBIM)
…

Meet the Class (pg. 2)

Sabrina R. Li (Stat)
Dave LeRoux (CS)
Xiaoxia Ren (CS)
Zhiguo Li (CS/CBIM)

Sign up Sheet

Next Time

Me:

- syllabus
- web page
- snowball simulator

TA

- available ML resources

You:

- Ch. 1