
CS 536: Machine Learning

Fall 2004: Michael L. Littman
TA: Yihua Wu

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, Boosting, clustering, data mining

Welcome!

My second time teaching ML here.

This time:

- <http://www.cs.rutgers.edu/~mlittman/courses/ml04/>

Last time:

- <http://www.cs.rutgers.edu/~mlittman/courses/ml03/>

I encourage you to print lectures in advance.

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!
- Midterm
- Final

Repeat with me: “*I am not an undergrad...*”

Snowballs: Details


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

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

How minimize cost to achieve the objective?

ML 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.



Some Types of Learning

Supervised, classification: vector \rightarrow Boolean

Supervised, regression: vector \rightarrow real

Supervised, general: vector \rightarrow vector

Unsupervised, discrete: vector \rightarrow cluster id

Reinforcement, associative: vector \rightarrow action,
real-valued feedback

Reinforcement, temporal: vector \rightarrow action,
delayed real-valued feedback

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/~helsel/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. 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. 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. 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:

Steven Carroll (CS)
Carlos Diuk (CS)
Weijun He (CS)
Jason Keller (CS)
Nikita Lytkin (CS)
Joshua Morman (332?)
Zhen Qian (125?)
Ryan Wagner (CS)
Michael Wood (CS)

Contacted via email:

James Wu (chem-e)
Ahmed Chowdhury (CS)
Allan Kugel (CogSci)
Jacek Rawicki (CS)
Michael Cole (SCILS)
Qi Wei (CS)
Shu Chen (CS)

Email Sign Up Sheet

For announcements.

Next Time

Read Ch. 1.