CS 445/545
Machine Learning
Winter, 2010

- Course overview:
  - Instructor
    - Melanie Mitchell
    - Office hours: 4-5pm M,W, or by appointment
  - Readings
    - Textbook: *Machine Learning: An Algorithmic Approach* by Stephen Marsland
    - Other readings will be downloadable from the Web

Homework

- One assignment for each topic
- Do on your own; no teamwork, but you can discuss general ideas with classmates

Quizzes

- Weekly short in-class quizzes on reading assignments, lectures

Midterm exam

- Take-home open-book exam

Final exam

- Take-home open-book exam

Presentation

- You'll be assigned a paper to present in class (presentation about 15 minutes)

- Grading
  - Homeworks: 40%
  - Quizzes: 15%
  - Midterm: 15%
  - Final: 20%
  - Presentation: 10%
- Late homework policy:
  - Students must request and be granted an extension for any assigment *before* that assignment is due. Otherwise 5% of the assignment grade will be subtracted for each day the assignment is late.
- Mailing list: MachineLearning2010@cs.pdx.edu
  (please give me your e-mail address)

- Academic Integrity
  - You can discuss concepts with other students, but all work you hand in must be your own.
  - You may not copy any text from other students, papers, the internet, or any other source, without clearly referencing it.
  - See syllabus for PSU policy.

- Use of laptops / PDAs in class

  Please don't, unless you are using it to take notes or view class slides.

What is machine learning?

- Textbook definitions of "machine learning":
  - Detecting patterns and regularities with a good and useful approximation ("model" or "hypothesis")
  - Execution of a computer program to optimize the parameters of the model using the training data or past experience.

- **Next term course:** CS 446/546 Advanced Topics in Machine Learning

A B C D E

A B C D E

A B C D E

A B C D E

A B C D E

- Example:
  - Task: Letter recognition
  - Performance measure: Fraction of letters recognized correctly
  - Training experience: Different examples of letters in different typefaces

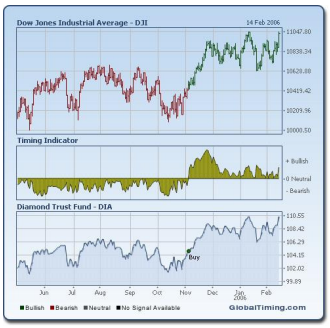

• Example:

– Task: Stock price prediction

– Performance measure: Amount of profit earned

– Training experience: Past time series of stock prices

Types of machine learning tasks

- Classification
  - Output is one of a number of classes (e.g., 'A')
- Regression
  - Output is a real value (e.g., '$35/share")

Difference between “artificial intelligence” and “machine learning”?

Key Ingredients for Any Machine Learning Method

- Underlying representation for “hypothesis”, “model”, or “target function”:
  - mathematical expression, bit string, neural network, decision tree, logical description, if-then rules
- Features (or “attributes”)
  - Elements of underlying representation that describe which aspects of problem instances (or training examples) should be used in learning.
- Space to search:
  - coefficient values, bits, weights, topologies of networks, topologies of trees, possible logic sentences
- Search method:
  - gradient descent, genetic algorithm, greedy algorithm, etc.

Key Ingredients for Any Machine Learning Method

- Data: Divide into three parts.
  - Training data
    - Used to train the model
  - Validation data
    - Used to select model complexity, to determine when to stop training, or to alter training method
  - Test data
    - Used to evaluate trained model

Types of Machine Learning Methods

- Supervised
  - provide explicit training examples with correct answers
    - e.g. neural networks with back-propagation
- Unsupervised
  - no feedback information is provided
    - e.g., unsupervised clustering based on similarity

Types of Machine Learning Methods, continued

- "Semi-supervised"
  - feedback information is provided, but is not detailed
    - e.g.,
      - genetic algorithm: calculates single-valued "fitness" of individual in population
      - reinforcement learning: reinforcement single is single-valued assessment of current state

Assumption of all ML methods:

Inductive learning hypothesis:

Any hypothesis that approximates target concept well over sufficiently large set of training examples will also approximate the concept well over other examples outside of the training set.

Difference between "induction" and "deduction"?

What questions does ML ask?

- Which is the best method for a given learning problem?
- What can be proved about convergence, performance, computational complexity?
- How much training data or feedback is needed for good generalization performance?

- How to select/design training examples or feedback?
- How can prior knowledge be used to guide learning?
- How can meta-learning be used? (I.e., learner automatically modifies its representation, selection of training examples, etc.)

Syllabus

Feature Extraction

- Example: Text classification
- Brainstorm about features