Welcome to CSCE 478/878!

- Please check off your name on the roster, or write your name if you're not listed
  - Indicate if you wish to register or sit in
- **Policy on sit-ins:** You may sit in on the course without registering, but not at the expense of resources needed by registered students
  - Don't expect to get homework, etc. graded
  - If there are no open seats, you may have to surrender yours to someone who is registered
- You should have two handouts:
  - Syllabus
  - Copies of slides
- In addition, check out **Homework 0** on the web

Override Policy

Option 1

Priority given to
- Undergraduate CSE majors graduating in December or May
- CSE graduate students who need it for research
- If you want an override, fill out the sheet with your name, ugrad/grad, major, and why this course is important to you

Option 2

Machine Learning Lecture 1: Introduction

What is Machine Learning?

- Building machines that automatically *learn* from experience
  – Sub-area of artificial intelligence
- (Very) small sampling of applications:
  – Detection of fraudulent credit card transactions
  – Filtering spam email
  – Autonomous vehicles driving on public highways
  – Self-customizing programs: Web browser that learns what you like and seeks it out
  – Applications we can't program by hand: E.g., speech recognition

What is Learning?

- Many different answers, depending on the field you're considering and whom you ask
  – Artificial intelligence vs. psychology vs. education vs. neurobiology vs. …

Does Memorization = Learning?

- Test #1: Thomas learns his mother's face

 Memorizes: 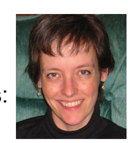  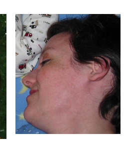

But will he recognize:  

Thus he can generalize beyond what he's seen!

Does Memorization = Learning? (cont'd)

- Test #2: Nicholas learns about trucks

Memorizes:

But will he recognize others?

- So learning involves **ability to generalize** from labeled examples
- In contrast, memorization is trivial, especially for a computer

What is Machine Learning? (cont'd)

- When do we use machine learning?
  - Human expertise does not exist (navigating on Mars)
  - Humans are unable to explain their expertise (speech recognition; face recognition; driving)
  - Solution changes in time (routing on a computer network; driving)
  - Solution needs to be adapted to particular cases (biometrics; speech recognition; spam filtering)
- In short, when one needs to generalize from experience in a non-obvious way

What is Machine Learning? (cont'd)

- When do we *not* use machine learning?
  - Calculating payroll
  - Sorting a list of words
  - Web server
  - Word processing
  - Monitoring CPU usage
  - Querying a database
- When we can definitively specify how all cases should be handled

More Formal Definition of (Supervised) Machine Learning

- Given several *labeled examples* of a *concept*
  - E.g., trucks vs. non-trucks (binary); height (real)
- Examples are described by *features*
  - E.g., number-of-wheels (int), relative-height (height divided by width), hauls-cargo (yes/no)
- A machine learning algorithm uses these examples to create a *hypothesis* that will *predict* the label of new (previously unseen) examples

Machine Learning Definition (cont'd)

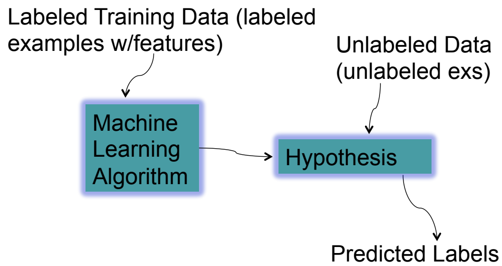

- Hypotheses can take on many forms

Hypothesis Type: Decision Tree

- Very easy to comprehend by humans
- Compactly represents if-then rules

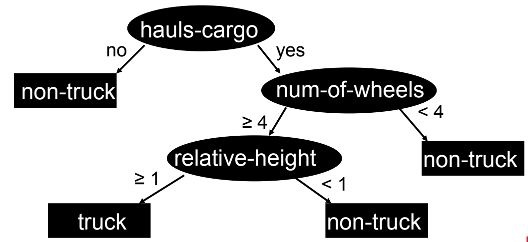

Hypothesis Type: Artificial Neural Network

- Designed to simulate brains
- "Neurons" (processing units) communicate via connections, each with a numeric weight
- Learning comes from adjusting the weights

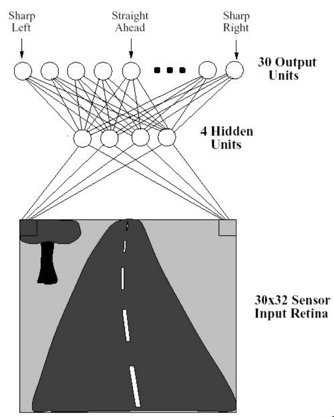

Hypothesis Type: *k*-Nearest Neighbor

- Compare new (unlabeled) example x_q with training examples
- Find k training examples most similar to x_q
- Predict label as majority vote

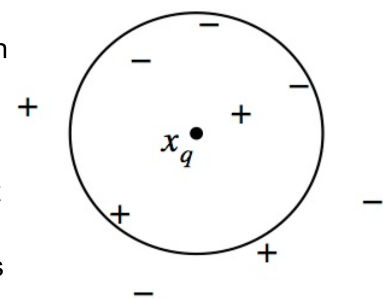

Other Hypothesis Types

- Support vector machines
  - A major variation on artificial neural networks
- Bagging and boosting
  - Performance enhancers for learning algorithms
- Bayesian methods
  - Build probabilistic models of the data
- Many more

Variations

- Regression: real-valued labels
- Probability estimation
  - Predict the probability of a label
- Unsupervised learning (clustering, density estimation)
  - No labels, simply analyze examples
- Semi-supervised learning
  - Some data labeled, others not (can buy labels?)
- Reinforcement learning
  - Used for e.g., controlling autonomous vehicles
- Missing attributes
  - Must some how estimate values or tolerate them
- Sequential data, e.g., genomic sequences, speech
  - Hidden Markov models
- Outlier detection, e.g., intrusion detection
- And more …

Issue: Model Complexity

- Possible to find a hypothesis that perfectly classifies all training data
  – But should we necessarily use it?

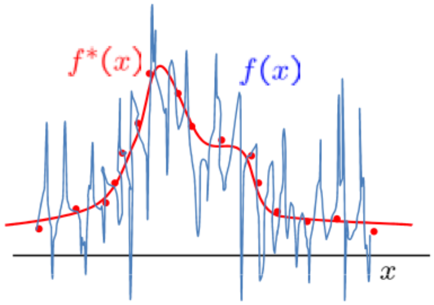

Model Complexity (cont'd)

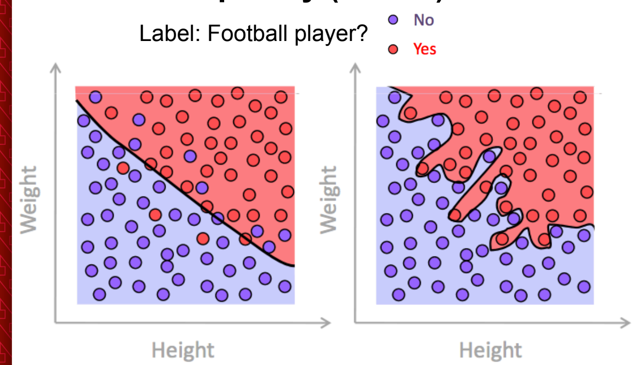

➔ To generalize well, need to balance ***accuracy*** with ***simplicity***

Issue: What If We Have Little Labeled Training Data?

- E.g., billions of web pages out there, but tedious to label

Conventional ML approach:

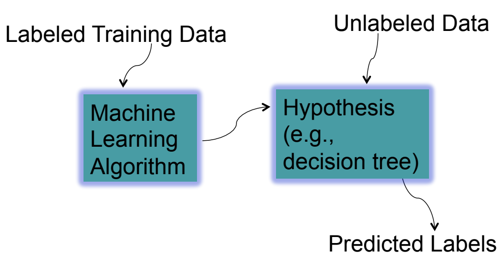

What If We Have Little Labeled Training Data? (cont'd)

Active Learning approach:

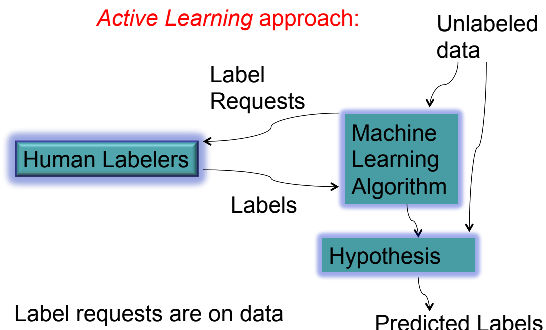

- Label requests are on data that ML algorithm is unsure of

Machine Learning vs Expert Systems

- Many old real-world applications of AI were *expert systems*
  – Essentially a set of if-then rules to emulate a human expert
  – E.g. "If medical test A is positive and test B is negative and if patient is chronically thirsty, then diagnosis = diabetes with confidence 0.85"
  – Rules were extracted via interviews of human experts

Machine Learning vs Expert Systems (cont'd)

- ES: Expertise extraction tedious;
  ML: Automatic
- ES: Rules might not incorporate intuition, which might mask true reasons for answer
- E.g. in medicine, the reasons given for diagnosis x might not be the objectively correct ones, and the expert might be unconsciously picking up on other info
- ML: More "objective"

Machine Learning vs Expert Systems (cont'd)

- ES: Expertise might not be comprehensive, e.g. physician might not have seen some types of cases
- ML: Automatic, objective, and data-driven
  - *Though it is only as good as the available data*

Relevant Disciplines

- Artificial intelligence: Learning as a search problem, using prior knowledge to guide learning
- Probability theory: computing probabilities of hypotheses
- Computational complexity theory: Bounds on inherent complexity of learning
- Control theory: Learning to control processes to optimize performance measures
- Philosophy: Occam's razor (everything else being equal, simplest explanation is best)
- Psychology and neurobiology: Practice improves performance, biological justification for artificial neural networks
- Statistics: Estimating generalization performance

More Detailed Example: Content-Based Image Retrieval

- Given database of hundreds of thousands of images
- How can users easily find what they want?
- One idea: Users query database by image *content*
  - E.g., "give me images with a waterfall"

Content-Based Image Retrieval (cont'd)

- One approach: Someone annotates each image with text on its content
  - Tedious, terminology ambiguous, may be subjective
- Another approach: *Query by example*
  - Users give examples of images they want
  - Program determines what's common among them and finds more like them

Content-Based Image Retrieval (cont'd)

User's Query


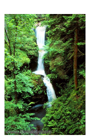


System's Response




User feedback: Yes | Yes | Yes | **NO!**

Content-Based Image Retrieval (cont'd)

- User's feedback then labels the new images, which are used as more training examples, yielding a new hypothesis, and more images are retrieved

How Does The System Work?

- For each pixel in the image, extract its color + the colors of its neighbors

- These colors (and their relative positions in the image) are the features the learner uses (replacing, e.g., number-of-wheels)
- A learning algorithm takes examples of what the user wants, produces a hypothesis of what's common among them, and uses it to label new images

Conclusions

- ML started as a field that was mainly for research purposes, with a few niche applications
- Now applications are very widespread
- ML is able to automatically find patterns in data that humans cannot
- However, still very far from emulating human intelligence!
- Each artificial learner is task-specific