Machine Learning
Lecture 1: Introduction

Stephen Scott

Adapted from *Future Problem Solvers* lecture, January 2004

What is Machine Learning?

- Building machines that automatically *learn* from experience
  - Important research goal of artificial intelligence
- (Very) small sampling of applications:
  - Data mining programs that learn to detect fraudulent credit card transactions
  - Programs that learn to filter spam email
  - Autonomous vehicles that learn to drive on public highways
  - Recognizing handwritten characters for mail and check sorting
  - Modeling users to aid web browsing, shopping, etc.

What is Learning?

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

Does Memorization = Learning?

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

Memorizes:

But will he recognize:

Welcome to 478/878!

- Please check off your name on the roster, or write your name if you're not listed
  - If you write your name, indicate if you plan to register
- You should have 2 handouts:
  1. Syllabus
  2. Copies of slides
- In addition, check these out on the web (mandatory!):
  1. Homeworks 0 and 0.5
  2. Prerequisite Test

CSCE 478/878 Lecture 0: Administrivia

Stephen D. Scott

August 25, 2008

Again, what is Machine Learning?

- Given several *labeled examples* of a *concept*
  - E.g. trucks vs. non-trucks
- Examples are described by *features*
  - E.g. *number-of-wheels* (integer), *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
- Similar to a very simplified form of human learning
- Hypotheses can take on many forms

Thus he can generalize beyond what he's seen!

Hypothesis Type: Decision Tree

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

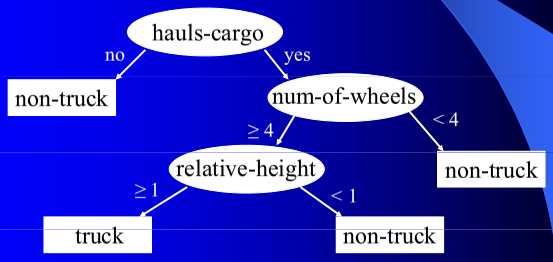

Does Memorization = Learning? (cont'd)

- Test #2: Nicholas learns about trucks & combines

Memorizes:

But will he recognize others?

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

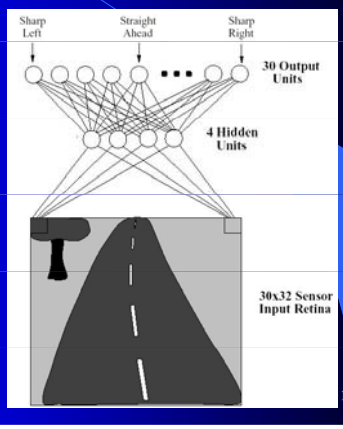

So learning involves ability to generalize from labeled examples (in contrast, memorization is trivial, especially for a computer)

Why Machine Learning? (cont'd)

- 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

Other Hypothesis Types

- Nearest neighbor
  - Compare new (unlabeled) examples to ones you've memorized
- Support vector machines
  - A new way of looking at artificial neural networks
- Bagging and boosting
  - Repeatedly apply your favorite learning algorithm and combine the results
- Bayesian approaches
  - Build probabilistic models of the concept
- Many more
  - See your local machine learning instructor for details

Machine Learning vs. Expert Systems

- 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"

Why Machine Learning?

- (Relatively) new kind of capability for computers
  - Data mining: extracting new information from medical records, maintenance records, etc.
  - 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, handwriting recognition, autonomous driving

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*

Why Machine Learning? (cont'd)

- Understanding human learning and teaching:
  - Mature mathematical models might lend insight
- The time is right:
  - Recent progress in algorithms and theory
  - Enormous amounts of data and applications
  - Substantial computational power
  - Growing industry (e.g. Google's and other companies' research groups)

Content-Based Image Retrieval (cont'd)

User's Query:

System's Response:

User Feedback: Yes Yes Yes NO!

Relevant Disciplines

- AI: 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

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

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"

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

Content-Based Image Retrieval (cont'd)

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

Variations on the Theme

- Classification vs. regression
  - Discrete- vs. real-valued labels
- Supervised vs. semi-supervised vs. unsupervised
  - How many of the training examples are labeled (all of it vs. some of it vs. none of it)
  - There's use to be made of unlabeled data!
- Noise in attributes and/or labels
- We'll emphasize binary classification, fully supervised, noise-free

Other Applications of ML

- The Google search engine uses numerous machine learning techniques
  - Spelling corrector: "spehl korector", "phonitick spewling", "Brytney Spears", "Brithney Spears", …
  - Grouping together top news stories from numerous sources (news.google.com)
  - Analyzing data from over 3 billion web pages to improve search results
  - Analyzing which search results are most often followed, i.e. which results are most relevant

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

Other Applications of ML (cont'd)

- ALVINN, developed at CMU, drives autonomously on highways at 70 mph
  - Sensor input only a single, forward-facing camera

Other Applications of ML (cont'd)

- SpamAssassin for filtering spam e-mail
- Data mining programs for:
  - Analyzing credit card transactions for anomalies
  - Analyzing medical records to automate diagnoses
- Intrusion detection for computer security
- Speech recognition, face recognition
- Biological sequence analysis
- Each application has its own representation for features, learning algorithm, hypothesis type, etc.