What is Machine Learning?

- What is learning?

  "changes in [a] system that ... enable [it] to do the same task or tasks drawn from the same population more efficiently and more effectively the next time." (Simon 1983)

  "any computer program that improves its performance at some task through experience." (Mitchell 1997)

There are two ways that a system can improve:

1. By acquiring new knowledge
For example: acquiring new facts or skills

2. By adapting its behavior
For example: solving problems more accurately or efficiently

Adaptive User Interfaces

- **Definition**: An *adaptive user interface* is a software artifact that improves its ability to interact with a user by constructing a user model based on partial experience with that user.

- AUI are really at the intersection of Human Computer Interaction and Machine Learning (the latter is an area within Artificial Intelligence)

Input to an ML System

- Most ML algorithms use a **training set** of examples as the basis for learning.
- Each example is encoded using the instance representation chosen for the problem.
- Examples presented to a machine learning algorithm are typically represented as **attribute-value** pairs.
  - An **attribute** is a general property associated with an object. For example, we might describe animals with the attributes: *size, color, temp, has_tail, has_beak, covering*
  - A **value** is one possible instantiation of the attribute.
    A CANARY might be represented as:
    size=small, *color*=yellow, *temp*=warm_blooded, *has_tail*=true, *has_beak*=true, *covering*=feathers

User Interface Dimensions

- More than one form of presentation modality (speech, vision, sound)
- More than one form of interaction (press buttons, type, speak)
  - example: Organization on web page different for blind users: key information at top of page where it's read to user sooner;
  - example: Customized on-screen keyboard for disabled users
- Different content for different users
  - example: different levels of expertise at using interface and at the knowledge level

Dimensions of Adaptation

- Data processing level
- Information Filtering level
- Information Presentation level

Training Experience

- Direct or indirect?

  Direct experience consists of individual states/items and their individual classifications.

  Indirect experience consists of a sequence of states and a final classification only. *Credit assignment* is a major issue.

- Teacher or self-selected training?
  - A teacher may select informative training examples.
  - The learner may propose training examples that would be especially useful to learn from.
- Distribution of training data: Generally assume training data is representative of the examples to be judged on when tested for final performance.

AUI: an Application of Machine Learning

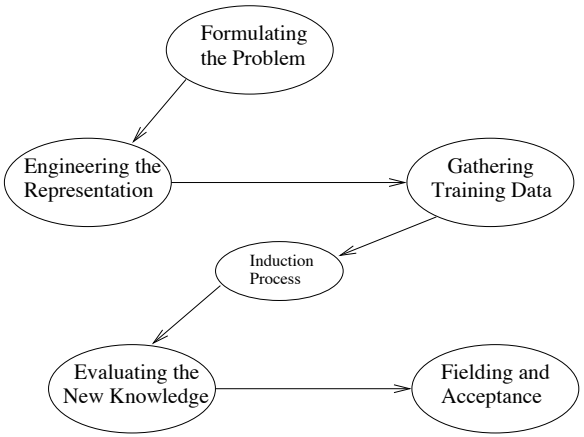

Adaptive Interfaces Discussion Points

- Dimensions along which interfaces could adapt?
- What challenges do you anticipate for Artificial Intelligence when applied to interface adaptation?

Syskill & Webert, cont'd

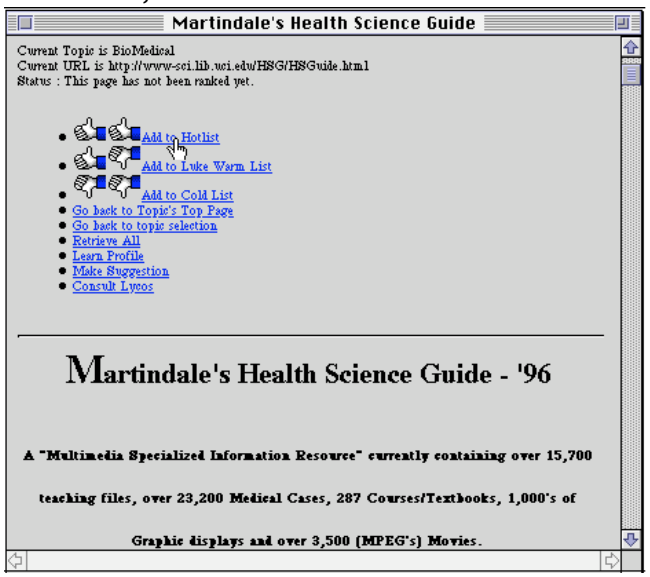

Examples of Adaptive User Interfaces

- Information filtering:
  - Syskill & Webert
  - NewsWeeder
- Generating new knowledge to satisfy user's goals:
  - scheduling
  - part layout
- Optimization:
  - route advisor
  - scheduling
- Information entry:
  - predict keystrokes or Unix commands
  - form filling

Syskill & Webert, cont'd

Profile can be used to:

- suggest which links a user might be interested in
- construct a Lycos query to find interesting (to user) pages

Learns a profile for each topic and user

Functionality is added to pages to collect user feedback:

- hot (2 thumbs up),
- lukewarm (1 up, 1 down), or
- cold (2 thumbs down);

But learner only uses 2 classes by combining lukewarm & cold

Syskill & Webert in Depth

(Pazzani et.al., 1996)
Technique:

- Recommends web pages on a user-specified topic
- Accepts user feedback about pages the user selects
- Represents each web page as a "bag of words"
- Uses naive Bayes to adapt to user's page preferences

Evaluation: 80% accurate on predicting user's movie opinions after training on only 35 pages

Issue: Where to get training data

- past user actions
- domain/application-specific information
- context
- what other users have done

Syskill & Webert Discussion

What are the advantages of this method?

What are the disadvantages?

How might we improve the interface? The learning?

Issue: Nature of User Feedback

- Least specific: Binary feedback: WebWatcher, Syskill & Webert
- Intermediate: ordering on choices: Adaptive Route Advisor, INCA
- Most specific: scoring: Firefly
- Passive versus Active

Form Filling in Depth

(Schlimmer & Hermens, 1993)

- Collects traces of a secretary filling out vacation forms
- Learns rules that predict some fields based on others
- Suggests default entries for fields that the user can override

Evaluation: Reduced keystrokes by 87%

Evaluation Discussion

What aspects of the system would you want to improve as a result of its learning about you?

How easy are these to measure?

Is there a difference between real (quantitative) and perceived (qualitative) improvements? Which would be more important to you?

Email Alerting in Depth

(Horvitz, et.al., 1999)
Send an alert when important email arrives

- Compares expected cost of interruption to expected cost of delaying notification
- Bayes net for assessing user's focus of attention
- Learns to assess criticality of a message from email with user-labeled criticality values

Related Application: email filtering whether to send to user's remote device

Dimensions of Evaluation

- Dependent measures
  - solution speed
  - solution quality
  - amount of effort reduced
  - user satisfaction
  - predictive accuracy
- Independent variables
  - number of user interactions
  - characteristics of system, user, and task

General Design Issues

- What to predict? (how can we best help users)
- Demand on user (obtrusive versus unobtrusive)
- Should user be able to examine and change their model?
- Should user be able to override the system's changing itself?
- Evaluation
- Design of interface more tied to design of adaptation process?

Dialog Systems

- Bring in techniques from natural language understanding, speech recognition, human factors, and problem solving and inference
- Most existing systems are specific to a *task*
- Developing the knowledge base needed to control a system is time and resource intensive
- We focus on decreasing that time by improving a simple system, online, after it is built
- Our first effort improves along the dimension of asking better questions when helping the user find information

Ethical Issues

- Privacy
- Etiquette
- Is it legal?
- Could we do anonymous user models?
- Where is the model stored and who owns it?
- Is the interface trying to influence the user?
  - Adaptive tutor vs.
  - Selling you a product you would buy anyway vs.
  - (Trying to) Change your buying habits or preferences

Demo

Adaptive Place Advisor discussion and demonstration

User modeling and Intelligent Interfaces

- Stereotypes
- Dialog systems
- Just do it? Or ask first?
- Personification

Personalization on the Web

- Firefly
- WiseWire
- News Dude
- Many others

What is the Future?

- Aides, not Agents
- Intelligent Environments
- Wearable computers
- Conversational interfaces
- Commercialization

Issues for Web Personalization

- Of 1400 random web sites, 92% collected “great amounts of personal data” (US Federal Trade Commission, 1998)
- Many users do not wish to register with web sites, or if they do they provide fake information
- Some countries don’t allow usage logs to be kept from session to session
- Security of your personal information!

Extra Slides Follow

Music Recommendation in Depth

(Shardanand & Maes, 1994)

- Users provide ratings of music items
- System provides recommendations of other items the user may also enjoy
- Compare ratings to other user's ratings and suggest music accordingly