What is Learning?

Learning denotes changes in the system that are adaptive in the sense that they enable the system to do the same task or tasks drawn from the same population more effectively the next time. -- Simon, 1983

Learning is making useful changes in our minds. -- Minsky, 1985

Learning is constructing or modifying representations of what is being experienced. -- McCarthy, 1968

Learning is improving automatically with experience. -- Mitchell, 1997

Why Machine Learning?

- Data, Data, DATA!!!
  - Examples
    - World wide web
    - Human genome project
    - Business data (WalMart sales "baskets")
  - Idea: sift heap of data for nuggets of knowledge
- Some tasks beyond programming
  - Example: driving
  - Idea: learn by doing/watching/practicing (like humans)
- Customizing software
  - Example: web browsing for news information
  - Idea: observe user tendencies and incorporate

Analysis/Prediction Problems

- What kind of direct mail customers buy?
- What products will/won't customers buy?
- What changes will cause a customer to leave a bank?
- What are the characteristics of a gene?
- Does a picture contain an object (does a picture of space contain a metereorite -- especially one heading towards us)?
- … Lots more

Tasks too Hard to Program

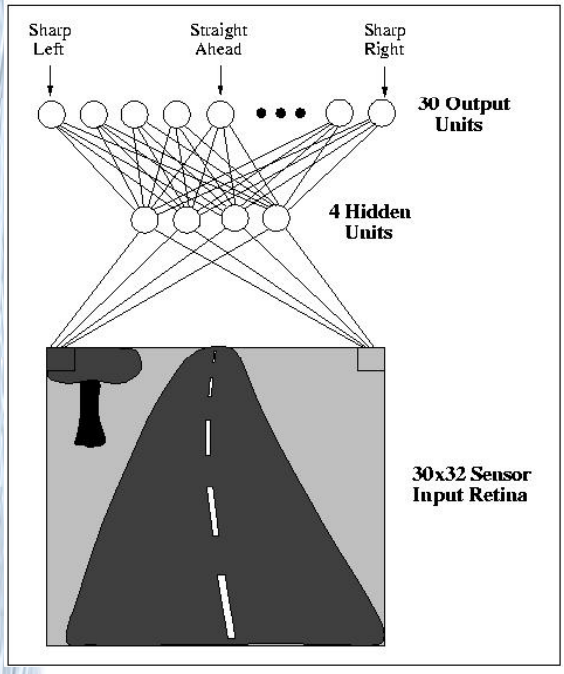

ALVINN [Pomerleau] drives 70 MPH on highways

STANLEY: Stanford Racing

- http://www.stanfordracing.org
- Sebastian Thrun's Stanley Racing program
- Winner of the DARPA grand challenge
- Incorporated learning/learned components with planning and vision components

Software that Customizes to User

Some Areas of Machine Learning

- Inductive Learning: inferring new knowledge from observations (not guaranteed correct)
  - Concept/Classification Learning - identify characteristics of class members (e.g., what makes a CS class fun, what makes a customer buy, etc.)
  - Unsupervised Learning - examine data to infer new characteristics (e.g., break chemicals into similar groups, infer new mathematical rule, etc.)
  - Reinforcement Learning - learn appropriate moves to achieve delayed goal (e.g., win a game of Checkers, perform a robot task, etc.)
- Deductive Learning: recombine existing knowledge to more effectively solve problems

Classification/Concept Learning

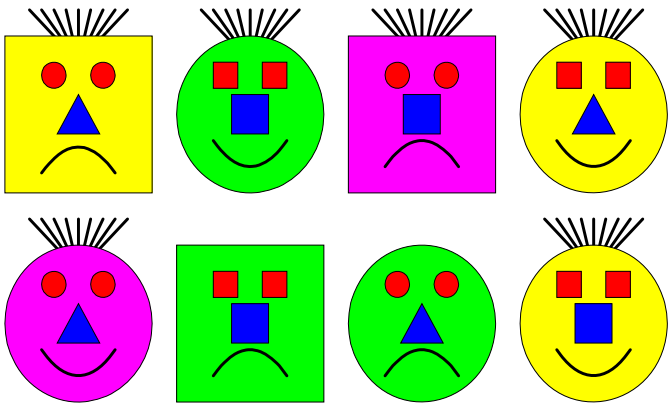

- What characteristic(s) predict a smile?
  - Variation on Sesame Street game: *why are these things a lot like the others (or not)?*
- ML Approach: infer model (characteristics that indicate) of why a face is/is not smiling

Unsupervised Learning

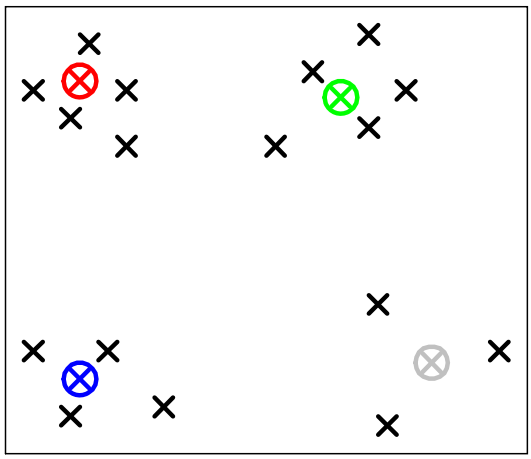

- Clustering - group points into “classes”
- Other ideas:
  - look for mathematical relationships between features
  - look for anomalies in data bases (data that does not fit)

Reinforcement Learning

Problem

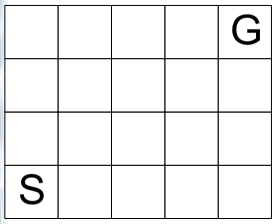

S - start
G - goal
Possible actions:
up left
down right

Policy

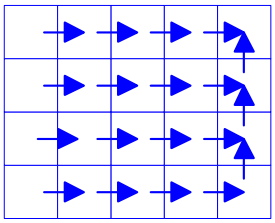

- Problem: feedback (reinforcements) are delayed - how to value intermediate (no goal states)
- Idea: online dynamic programming to produce policy function
- Policy: action taken leads to highest future reinforcement (if policy followed)

Analytical Learning

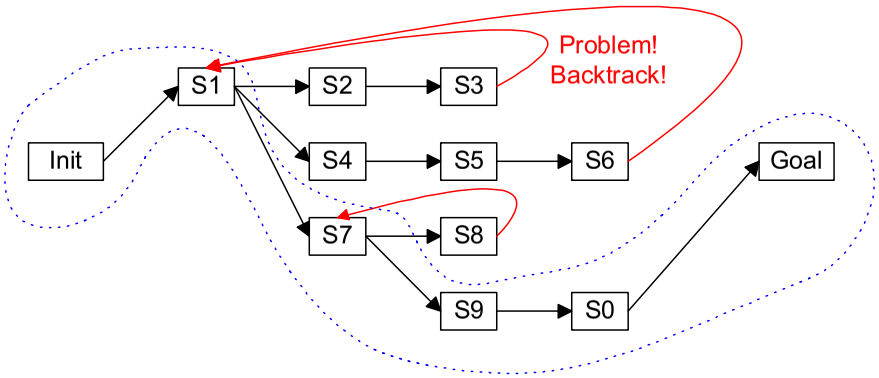

- During search processes (planning, etc.) remember work involved in solving tough problems
- Reuse the acquired knowledge when presented with similar problems in the future (avoid bad decisions)

The Present in Machine Learning

The tip of the iceberg:

- First-generation algorithms: neural nets, **decision trees**, regression, **support vector machines**, …
- Composite algorithms - ensembles
- Significant work on assessing effectiveness, limits
- Applied to simple data bases
- Budding industry (especially in data mining)

The Future of Machine Learning

Lots of areas of impact:

- Learn across multiple data bases, as well as web and news feeds
- Learn across multi-media data
- Cumulative, lifelong learning
- Agents with learning embedded
- Programming languages with learning embedded?
- Learning by active experimentation

What is Knowledge Discovery in Databases (i.e., Data Mining)?

- Depends on who you ask
- General idea: the analysis of large amounts of data (and therefore efficiency is an issue)
- Interfaces several areas, notably machine learning and database systems
- Lots of perspectives:
  - ML: learning where efficiency matters
  - DBMS: extended techniques for analysis of raw data, automatic production of knowledge
- What is all the hubbub?
  - Companies make lots of money with it (e.g., WalMart)

Related Disciplines

- Artificial Intelligence
- Statistics
- Psychology and neurobiology
- Bioinformatics and Medical Informatics
- Philosophy
- Computational complexity theory
- Control theory
- Information theory
- Database Systems
- ...

Issues in Machine Learning

- What algorithms can approximate functions well (and when)?
- How does number of training examples influence accuracy?
- How does complexity of hypothesis representation impact it?
- How does noisy data influence accuracy?
- What are the theoretical limits of learnability?
- How can prior knowledge of learner help?
- What clues can we get from biological learning systems?