Mathematical and Machine Learning approaches to Context

Good news: Machine Learning has some techniques that are relevant to dealing with the context problem

Bad news:

- Vocabularies are specialized and difficult to understand
- Tools have very specific strengths and weaknesses, so be careful when applying techniques

References for Machine Learning

Elements of Machine Learning

Pat Langely, MorganKauffman, 1996

Machine Learning

Tom Mitchell, McGraw-Hill, 1997

Machine Learning is often described as Classification

Classification Problem:

- Given a "training sequence" of examples, figure out if a new example belongs to the set or not
- Supervised vs. Unsupervised learning

Other problems can be "reduced to classification", but it's a pain

Leaning problems for context-aware applications

Action Description

Data Description

Sequence Prediction

Inferring Preferences

Feature Selection

Traditional mathematical inference techniques

First-order logical inference almost guaranteed to fail in problems where context is important

Deduction insists on sound inference

Contextually based inference formulated as a problem of induction, not deduction

Induction from specific to general not sound

Watch out for hidden assumptions in machine learning

- Everything appears at onvs. Incremental
- Batchvs. Interactive
- Pretending the user doesn't exist

Generalization Lattice

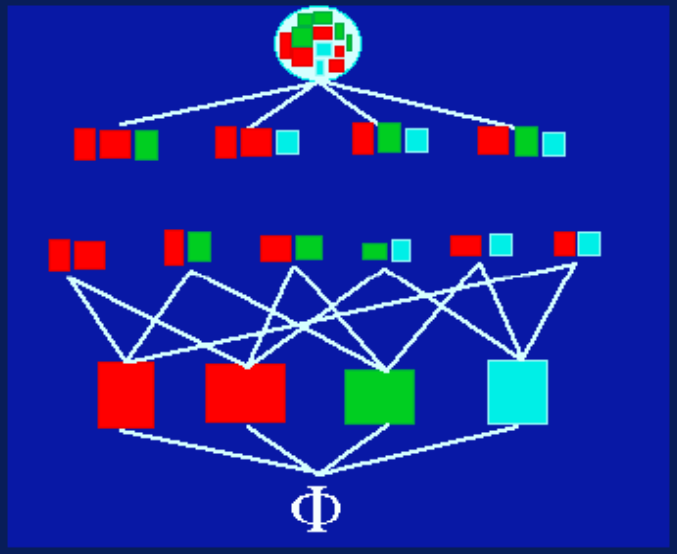

Version Spaces

Lattice of hypotheses ordered by generalization

Each step proceeds by

- Making the Most Specific hypothesis a little more general
- Making the Most General hypothesis a little more specific

Version Spaces

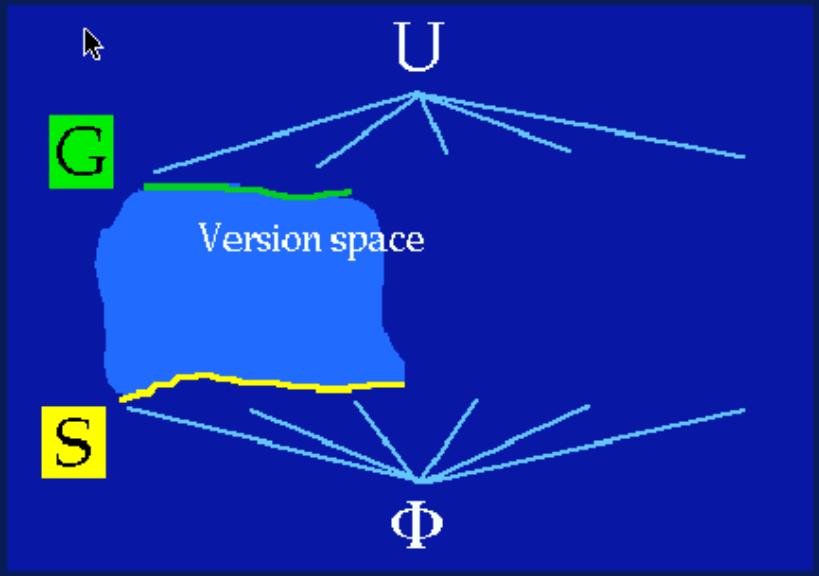

Case-Based Reasoning

Index specific examples and solutions in a case base

When you get a new example:

- Retrieve the closest case
- Modify the solution to fit the current case

Works by *delaying* generalization

Explanation-Based Generalization

Present an example, a generalization, and show the example satisfies the generalization

- New examples can make analogies to the old explanation

Works by recording dependencies, and propagating generalizations through dependencies

Learning by Analogy

Find successful previous experience

Map *roles* in past examples to new example

Gentner, Structure-Mapping

Sequence prediction

Given examples of a sequence, what's the next element?

1, 4, 9, 16, 25, ...

4, 14, 34, 42, 59, 72, 96, 125,...

Loop detection

Grammar induction

- Sequitur - Neville-Manning&Maulsby, Witten

Approaches that use large numbers of examples

- Genetic Algorithms
  - Analogy to Evolutionary Biology
- Neural Networks
- Bayesianandprobablistidearning
- Fuzzy sets

Large-example-set approaches appropriate for context problems?

Problem is, user interactions don't usually generate large numbers of examples

Large-scale techniques often have trouble with explaining "what happened", control issues

Best for unsupervised learning, data mining

Users don't have patience for slow learning

Good for "noisy data"

Statistical approaches: Clustering

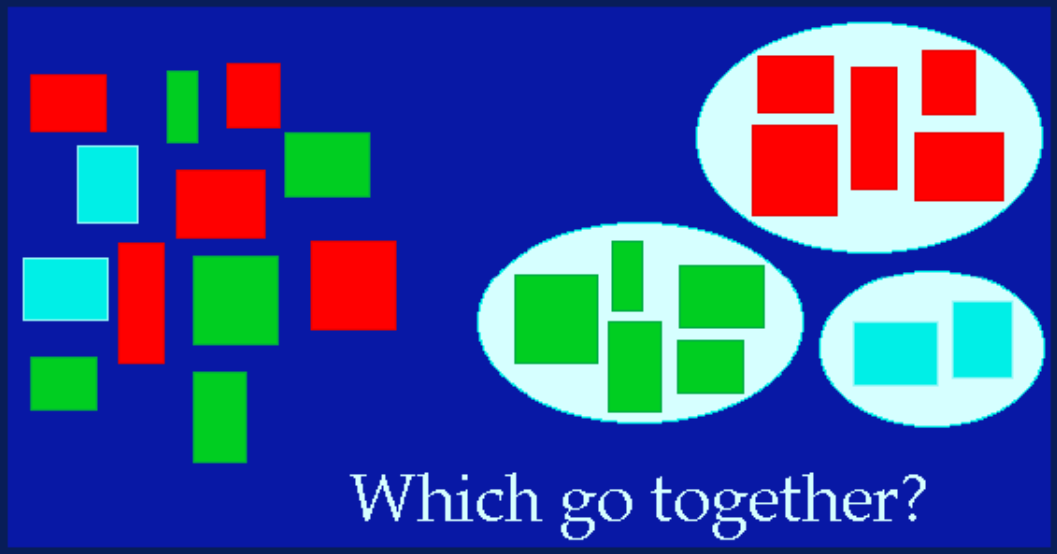

Discrimination Net

Attributes:

Domain? Animal, Vegetable, Mineral

Coverage? Feathers, Hair, Bald

Can fly? Yes, No

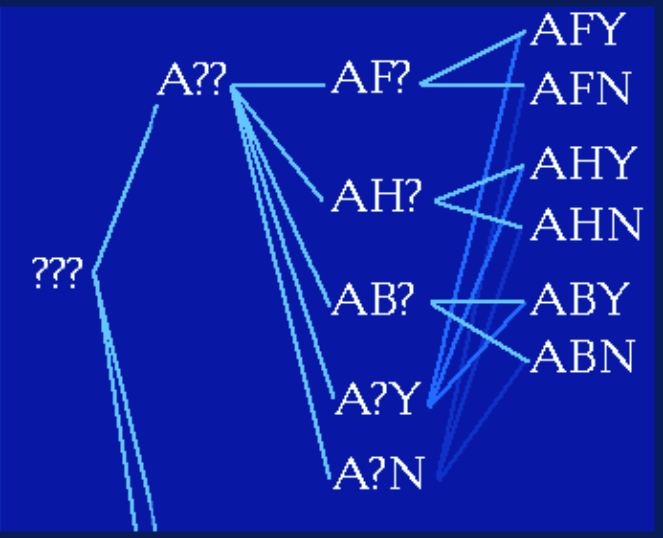

Feature selection problem

Out of the all the possible features that could affect a decision, which actually do?

Pick relevant aspects of context

Cimafeature selection Maulsby

Classify, Find, Generate, Modify

The Frame Problem in AI

When you change the context of an assertion, what changes?

Difficult to say all the possible things that might affect an assertion

Circumscription

Birds fly.

Tweetyis a bird.

Therefore,Tweetyflies.

... unless Tweetyis a penguin or ostrich,Tweety broke his wing,Tweetyis underwater,Tweetyis a stuffed bird, ...

Inferences are circumscribed by context

Everything takes [implicit] context argument

Default reasoningnonmonotoicreasoning

Inference rules take "extra argument" of context