

Artificial Intelligence:
Representation and Problem
Solving

15-381

April 10, 2007

Introduction to Learning &
Decision Trees

Learning and Decision Trees to learning

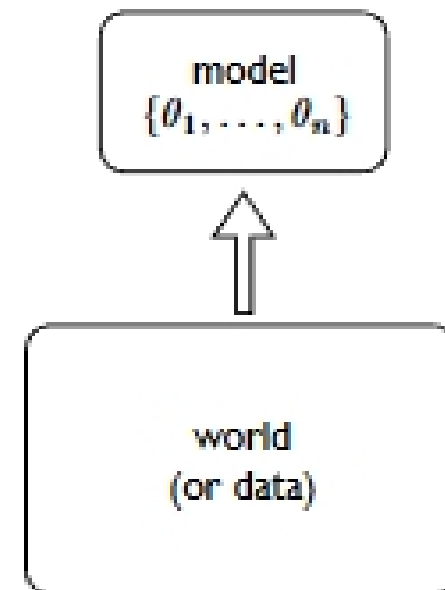
- What is learning?
 - more than just memorizing facts
 - learning the underlying *structure* of the problem or data
- A fundamental aspect of learning is *generalization*:
 - given a few examples, can you *generalize* to others?
- Learning is ubiquitous:
 - *medical diagnosis*: identify new disorders from observations
 - *loan applications*: predict risk of default
 - *prediction*: (climate, stocks, etc.) predict future from current and past data
 - *speech/object recognition*: from examples, generalize to others

aka:

- regression
- pattern recognition
- machine learning
- data mining

Representation

- How do we model or represent the world?
- All learning requires some form of representation.
- Learning:
adjust model parameters to match data




The complexity of learning

- Fundamental trade-off in learning:
complexity of model
vs
amount of data required to learn parameters
- The more complex the model, the more it can describe,
but the more data it requires to constrain the parameters.
- Consider a hypothesis space of N models:
 - How many bits would it take to identify which of the N models is 'correct'?
 - $\log_2(N)$ in the worst case
- Want simple models to explain examples and generalize to others
 - Ockham's (some say Occam) razor

Complex learning example: curve fitting

$$t = \sin(2\pi x) + \text{noise}$$



How do we model the data?

Polynomial curve fitting

$$y(x, \mathbf{w}) = w_0 + w_1x + w_2x^2 + \dots + w_Mx^M = \sum_{j=0}^M w_jx^j$$

$$E(\mathbf{w}) = \frac{1}{2} \sum_{n=1}^N [y(x_n, \mathbf{w}) - t_n]^2$$

