## Lecture 1

### Analytical Learning and Data Engineering: Overview

Wednesday, January 19, 2000

William H. Hsu
Department of Computing and Information Sciences, KSU
http://www.cis.ksu.edu/~bhsu

Readings:
Chapter 21, Russell and Norvig
Flann and Dietterich



## Lecture Outline

- Quick Review
  - Output of learning algorithms
    - What does it mean to *learn a function*?
    - What does it mean to *acquire a model* through (inductive) learning?
  - Learning methodologies
    - Supervised (inductive) learning
    - Unsupervised, reinforcement learning
- Inductive Learning
  - What does an inductive learning problem specification look like?
  - What does the "type signature" of an inductive learning algorithm mean?
  - How do inductive learning and inductive bias work?
- Analytical Learning
  - How does analytical learning work and what does it produce?
  - What are some relationships between analytical and inductive learning?
- Integrating Inductive and Analytical Learning for KDD



## Introductions

- Student Information (Confidential)
  - Instructional demographics: background, department, academic interests
  - Requests for special topics
    - Lecture
    - Project
- On Information Form, Please Write
  - Your name
  - What you wish to learn from this course
  - What experience (if any) you have with
    - Artificial intelligence
    - Probability and statistics
  - What experience (if any) you have in using KDD (learning, inference; ANN, GA, probabilistic modeling) packages
  - What programming languages you know *well*
  - Any specific applications or topics you would like to see covered



## In-Class Exercise

- Turn to A Partner
  - 2-minute exercise
  - Briefly introduce yourselves (2 minutes)
  - 3-minute discussion
  - 10-minute go-round
  - 3-minute follow-up
- Questions
  - 2 applications of KDD systems to problem in your area
  - *Common* advantage and obstacle
- Project LEA/RN™ Exercise, Iowa State [Johnson and Johnson, 1998]
  - Formulate an answer *individually*
  - Share your answer with your partner
  - Listen carefully to your partner's answer
  - Create a new answer through discussion
  - Account for your discussion by being prepared to be called upon



## About Paper Reviews

- 20 Papers
  - Must write at least 15 reviews
  - Drop lowest 5
- Objectives
  - To help prepare for *presentations and discussions* (questions and opinions)
  - To introduce students to current research topics, problems, solutions, applications
- Guidelines
  - Original work, 1-2 pages
    - *Do not just summarize*
    - Cite external sources properly
  - Critique
    - Intended audience?
    - Key points: *significance* to a particular problem?
    - Flaws or ways you think the paper could be improved?



## About Presentations

- 20 Presentations
  - Every registered student must give at least 1
  - If more than 20 registered, will assign duplicates (still should be original work)
  - First-come, first-served (sooner is better)
- Papers for Presentations
  - Will be available at 14 Seaton Hall by Monday (first paper: online)
  - May present research project in addition / instead (contact instructor)
- Guidelines
  - Original work, ~30 minutes
    - *Do not just summarize*
    - Cite external sources properly
  - Presentations
    - Critique
    - Don't just read a paper review: *help the audience understand significance*
    - Be prepared for 20+ minutes of questions, discussion

## Quick Review: Output of Learning Algorithms

- **Classification Functions**
  - Learning hidden functions: estimating ("fitting") parameters
  - Concept learning (e.g., chair, face, game)
  - Diagnosis, prognosis: medical, risk assessment, fraud, mechanical systems
- **Models**
  - Map (for navigation)
  - Distribution (query answering, *aka* QA)
  - Language model (e.g., automaton/grammar)
- **Skills**
  - Playing games
  - Planning
  - Reasoning (acquiring representation to use in reasoning)
- **Cluster Definitions for Pattern Recognition**
  - Shapes of objects
  - Functional or taxonomic definition
- ***Many Problems Can Be Reduced to Classification***



## Quick Review: Learning Methodologies

- **Supervised**
  - What is learned? Classification function; other models
  - Inputs and outputs? Learning: examples $\langle x, f(x) \rangle \rightarrow$ approximation $\hat{f}(x)$
  - How is it learned? Presentation of examples to learner (by teacher)
- **Unsupervised**
  - Cluster definition, or *vector quantization* function (*codebook*)
  - Learning: observations $x \times$ distance metric $d(x_1, x_2) \rightarrow$ discrete codebook $f(x)$
  - Formation, segmentation, labeling of clusters based on observations, metric
- **Reinforcement**
  - Control policy (function from states of the world to actions)
  - Learning: state/reward sequence $\{\langle s_i, r_i \rangle : 1 \le i \le n\} \rightarrow$ policy $p : s \rightarrow a$
  - (Delayed) feedback of reward values to agent based on actions selected; model updated based on reward, (partially) observable state



## Example: Inductive Learning Problem

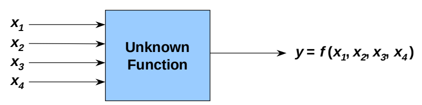


| Example | $x_1$ | $x_2$ | $x_3$ | $x_4$ | $y$ |
|---|---|---|---|---|---|
| 0 | 0 | 1 | 1 | 0 | 0 |
| 1 | 0 | 0 | 0 | 0 | 0 |
| 2 | 0 | 0 | 1 | 1 | 1 |
| 3 | 1 | 0 | 0 | 1 | 1 |
| 4 | 0 | 1 | 1 | 0 | 0 |
| 5 | 1 | 1 | 0 | 0 | 0 |
| 6 | 0 | 1 | 0 | 1 | 0 |

- $x_i$: $t_i$, $y$: t, $f$: $(t_1 \times t_2 \times t_3 \times t_4) \rightarrow t$
- Our learning function: Vector $(t_1 \times t_2 \times t_3 \times t_4 \times t) \rightarrow (t_1 \times t_2 \times t_3 \times t_4) \rightarrow t$



## Quick Review: Inductive Generalization Problem

- **Given**
  - Instances $X$: possible days, each described by attributes *Sky, AirTemp, Humidity, Wind, Water, Forecast*
  - Target function $c \equiv$ *EnjoySport*: $X \rightarrow H \equiv$ {{Rainy, Sunny} $\times$ {Warm, Cold} $\times$ {Normal, High} $\times$ {None, Mild, Strong} $\times$ {Cool, Warm} $\times$ {Same, Change}} $\rightarrow$ {0, 1}
  - Hypotheses $H$: e.g., conjunctions of literals (e.g., *<?, Cold, High, ?, ?, ?>*)
  - Training examples $D$: positive and negative examples of the target function

$$\langle x_1, c(x_1) \rangle, \ldots, \langle x_m, c(x_m) \rangle$$

- **Determine**
  - Hypothesis $h \in H$ such that $h(x) = c(x)$ for all $x \in D$
  - Such $h$ are consistent with the training data
- **Training Examples**
  - Assumption: no missing $X$ values
  - Noise in values of $c$ (contradictory labels)?



## Inductive Bias

- **Fundamental Assumption: Inductive Learning Hypothesis**
  - Any hypothesis found to *approximate the target function well* over a *sufficiently large* set of training examples will also approximate the target function well over other *unobserved* examples
  - Definitions deferred
    - *Sufficiently large, approximate well, unobserved*
    - Statistical, probabilistic, computational interpretations and formalisms
- **How to *Find* This Hypothesis?**
  - Inductive concept learning as *search* through hypothesis space $H$
  - Each point in $H \equiv$ subset of points in $X$ (those labeled "+", or positive)
- **Role of Inductive Bias**
  - Informal idea: preference for (i.e., restriction to) certain hypotheses by structural (syntactic) means
  - Prior assumptions regarding target concept
  - Basis for inductive generalization



## Inductive Systems and Equivalent Deductive Systems

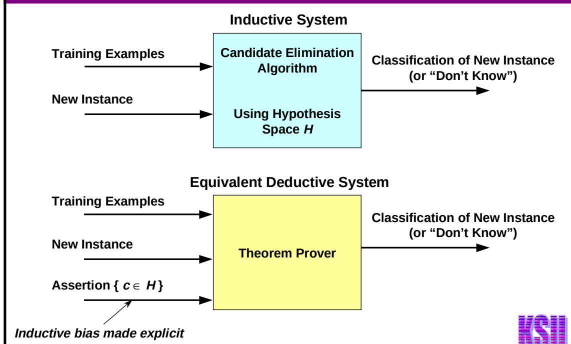

## Analytical Generalization Problem

- Given
  - Instances $X$
  - Target function (concept) $c: X \rightarrow H$
  - Hypotheses (i.e., hypothesis language *aka* hypothesis space) $H$
  - Training examples $D$: positive and negative examples of the target function $c$
  - Domain theory $T$ for explaining examples
- Domain Theories
  - Expressed in formal language
    - Propositional calculus
    - First-order predicate calculus (FOPC)
  - Set of assertions (e.g., well-formed formulae) for reasoning about domain
    - Expresses constraints over relations (predicates) within model
    - Example: *Ancestor* $(x, y) \leftarrow$ *Parent* $(x, z) \wedge$ *Ancestor* $(z, y)$.
- Determine
  - Hypothesis $h \in H$ such that $h(x) = c(x)$ for all $x \in D$
  - Such $h$ are consistent with the training data *and the domain theory T*



## Analytical Learning: Algorithm

- Learning with Perfect Domain Theories
  - Explanation-based generalization: Prolog-EBG
  - Given
    - Target concept $c: X \rightarrow$ boolean
    - Data set $D$ containing $\{x, c(x) \in$ boolean$\}$
    - Domain theory $T$ expressed in rules (assume FOPC here)
- Algorithm *Prolog-EBG* ($c$, $D$, $T$)
  - *Learned-Rules* $\leftarrow \varnothing$
  - FOR each positive example $x$ not covered by *Learned-Rules* DO
    - Explain: generate an explanation or proof $E$ in terms of $T$ that $x$ satisfies $c(x)$
    - Analyze: *Sufficient-Conditions* $\leftarrow$ most general set of features of $x$ sufficient to satistfy $c(x)$ according to $E$
    - Refine: *Learned-Rules* $\leftarrow$ *Learned-Rules* + *New-Horn-Clause*, where *New-Horn-Clause* $\equiv$ [$c(x) \leftarrow$ *Sufficient-Conditions*.]
  - RETURN *Learned-Rules*



## Terminology

- Supervised Learning
  - Concept – function: observations to categories; so far, boolean-valued (+/-)
  - Target (function) – true function $f$
  - Hypothesis – proposed function $h$ believed to be similar to $f$
  - Hypothesis space – space of all hypotheses that can be generated by the learning system
  - Example – tuples of the form $<x, f(x)>$
  - Instance space (*aka* example space) – space of all possible examples
  - Classifier – discrete-valued function whose range is a set of class labels
- Inductive Learning
  - Inductive generalization – process of generating hypotheses $h \in H$ that describe cases not yet observed
  - The inductive learning hypothesis – basis for inductive generalization
- Analytical Learning
  - Domain theory $T$ – set of assertions to *explain* examples
  - Analytical generalization - process of generating $h$ consistent with $D$ and $T$
  - Explanation – proof in terms of $T$ that $x$ satisfies $c(x)$



## Summary Points

- Concept Learning as Search through $H$
  - Hypothesis space $H$ as a state space
  - Learning: finding the correct hypothesis
- Inductive Leaps Possible Only if Learner Is Biased
  - Futility of learning without bias
  - Strength of inductive bias: proportional to restrictions on hypotheses
- Modeling Inductive Learners
  - Equivalent inductive learning, deductive inference (theorem proving) problems
  - Hypothesis language: syntactic restrictions (*aka* representation bias)
- Views of Learning and Strategies
  - Removing uncertainty ("data compression")
  - Role of knowledge
- Integrated Inductive and Analytical Learning
  - Using inductive learning to acquire domain theories for analytical learning
  - Roles of integrated learning in KDD
- Next Time: Presentation on Analytical and Inductive Learning (Hsu)