



Center for
**LifeLong
Learning
& Design**

University of Colorado at Boulder

**Wisdom is not the product of schooling
but the lifelong attempt to acquire it.
- Albert Einstein**

Learning

—

Part 6 in Russell / Norvig Book

Gerhard Fischer

AI Course, Fall 1996, Lecture October 14

Overview

- in the book: learning is **machine learning** = subfield of AI concerned with programs that learn from experience
- we (as a research center) have been interested in computational media and environments in support of **human learning**
- chapter 18: learning from observations
 - improve behavior by analysis of own behavior
- chapter 19: learning in neural networks — Guest Lecture by **Michael Mozer**
 - train complex networks of simple computing elements
- chapter 20: reinforcement learning — Guest Lecture by Satinder Singh
 - learn from success and failure, reward and punishment
- chapter 21: knowledge in learning
 - take advantage of prior knowledge

What Does Learning Mean?

- acquisition of new declarative knowledge
- development of motor and cognitive skills through instruction and practice
- organization of new knowledge into general effective representation
- discovery of new facts and theories through observation and experimentation
- example: playing Othello

Objectives of Machine Learning

- **applied learning systems — a practical necessity?**
 - to overcome the tedious work of programming
 - the ultimate form of knowledge acquisition in knowledge-based systems
 - example: “do not put appliances with the door against the wall!”
- **machine learning as a science**
 - understand human learning well enough to reproduce aspects of that learning behavior in computer systems
 - computer enforces a commitment to fine-structure process-level detail
 - insights into the principles underlying human learning abilities has the potential to lead to more effective educational techniques —example: student models in intelligent tutoring systems
 - exploration of alternative learning mechanism complementing human learning methods
- **knowledge acquisition versus skill refinement**
 - knowledge acquisition (example: learning physics) — learning new symbolic information coupled with the ability to apply that information in an effective manner
 - skill refinement (example: riding a bicycle, playing the piano) — occurs at a subconscious level by virtue of repeated practice

The History of Machine Learning

- **neural modeling**
 - general purpose learning systems
 - little or no initial structure or task-oriented knowledge
 - example: self-organizing systems
- **decision-theoretic techniques**
 - acquisition of linear, polynomial, or other forms of discriminant functions from a given set of training examples
 - example: Samuel's Checkers program (1959)
- **symbolic concept-oriented learning**
 - logic or graph structure representations
 - strong structural assumptions about the concepts to be acquired
 - examples: Feigenbaum's EPAM (= "Elementary Perceiver and Memorizer"; 1961) and Winston's learning description of complex objects (1970)
- **knowledge intensive approaches:**
 - to acquire new knowledge a system must already possess a great deal of initial knowledge
 - example: Lenat's AM (1977)

A Taxonomy of Machine Learning Research

- classification based upon the underlying learning strategy
- classification according to the type of knowledge acquired
- classification by domain of application

Classification Based upon the Underlying Learning Strategy

- **classification by the amount of inference:**
 - no inference (programmed, all effort on the side of the teacher)
 - substantial amount (burden on the teacher decreases)
- **taxonomy:**
 - rote learning or direct implanting of new knowledge
 - learning from instruction
 - learning by analogy
 - learning from examples
 - learning from observation and discovery
- **example: learning from examples**
 - source of the example — the teacher, learner, or external environment
 - examples available — positive, positive and negative
- **example: learning from observation and discovery**
 - passive observation
 - active experimentation

Classification According to the Type of Knowledge Acquired

- parameters in algebraic expressions
- decision trees
- formal grammars
- production rules
- formal logic-based expressions and related formalisms
- graphs and networks
- frames and schemas
- computer programs and other procedural encodings
- taxonomies
- multiple representations

Chapter 18: Learning from Observations

- **inductive learning** = constructing a description of a function from a set of input/output examples
- **example** = pair $(x, f(x))$
 - x: input
 - $f(x)$: output
- induction: given a collection of examples of f , return a function h that approximates f — h is called a **hypothesis**
- **bias** = a preference for one hypothesis over another (beyond mere consistency with the examples)

Learning Decision Trees

- decision tree
 - input: object or situation described by a set of properties
 - output: yes/no decision (decision trees are Boolean functions)
- inducing decision trees from examples
 - positive examples
 - negative examples
- assessing the performance of the learning algorithm
 - training sets
 - test sets

Human Learning: Current Theories

- learning is a process of **knowledge construction**, not of knowledge recording or absorption
- learning is **knowledge-dependent**; people use their existing knowledge to construct new knowledge
- learning is highly **tuned to the situation** in which it takes place
- learning needs to account for **distributed cognition** requiring to combine knowledge in the head with knowledge in the world
- learning is affected as much by **motivational issues** as by cognitive issues

Issues in Human Learning

- learning
 - is not only done in schools
 - learning can take place without being taught
- reflection in action and learning on demand
Norman — *“Real Learning: The way we learn is trying something, doing it and getting stuck. In order to learn, we really have to be stuck, and when we’re stuck we are ready for the critical piece of information. The same piece of information that made no impact at a lecture makes a dramatic impact when we’re ready for it.”*
- lifelong learning
- collaborative learning
- organizational learning
- distributed cognition (learning on demand versus using on demand)