CS 520
Graduate Artificial Intelligence
Spring 2000

Matthew Stone
Department of Computer Science and
Center for Cognitive Science
Rutgers University

Artificial Intelligence

- **Engineering approach to**
- **constructing computational artifacts to**
- **act in the real world.**

Cool people call these artifacts agents.

Artificial Intelligence
THE REAL WORLD

- **In typical CS, an engineer freely designs**
  - the data and representations that the program uses
  - the actions that make up program execution

(Think of structuring classes and methods in object-oriented design)

Artificial Intelligence
THE REAL WORLD

- **Real-world computing is different**
  - data and actions are not constructed by machine or governed by uniform design
  - data and actions exist (and must make sense) independent of system

Artificial Intelligence
THE REAL WORLD

- **Sample real-world tasks**
  - control a physical robot moving around a populated office
  - carry one end of an information-seeking dialogue, in natural language, with a human partner
  - cull useful information from web pages that people have designed for one another

Artificial Intelligence
THE REAL WORLD

- **Real world tasks give AI a focus on modeling**
  - describing the real world mathematically
  - as a programmer, to inform a design or
  - concretely, to realize an implementation

Artificial Intelligence
COMPUTATION

- **A computational artifact:**
  - maintains symbolic representations that correspond to the real world (according to arbitrary conventions)
  - manipulates them according to form
- **This ideal distinguishes AI**
  - from bridge-building, and from closer neighbors like control theory and EE

Artificial Intelligence
ENGINEERING

- **Engineering dictates AI methodology:**
  - modeling the world mathematically
  - describing computations theoretically
  - constructing implementations
  - evaluating how well they work
    - validity of models (science)
    - performance of algorithms (computation)
    - usefulness for some overall task (application)

Artificial Intelligence
ENGINEERING

- **Does not mean human intelligence is irrelevant to your system**
  - on the contrary, interaction with people is (and will be) a focus of AI applications [dialogue, smart spaces, perceptual user interfaces, web technology,…]
  - for this work, you have to model what people want, think, do
- **Just means you care how well it works**

The goals of an AI course

- **Teaching useful techniques for designing and implementing models of the world; &**
- **since models encode assumptions**
  - explicitly, e.g., in the meaning of a representation maintained by an agent; or
  - implicitly, e.g., as requirements for the correctness of inference algorithms
- **instilling awareness of these assumptions and understanding of their overall impact**

AI Course A
LOGIC

- **Agent's representations take the form of a set of logical formulas (a knowledge base)**
- **Each formula corresponds to a proposition that will either be true or false in any possible situation**
- **The knowledge base (KB) embodies a claim about the world that each of these propositions is true.**

AI Course A
LOGIC

- **Techniques work by manipulating arguments that one formula follows logically from others to solve problems**
  - prediction: fact follows from KB
  - perception: sense data follows from KB plus assumptions of what agent senses
  - planning: desired state follows from KB plus assumptions of what agent could do

AI Course A
LOGIC

- **Assumptions derive from the agent's background theory of its environment**
  - a set of statements in KB that are constant and unquestioned
  - that play a key role in agent's reasoning

AI Course B
PROBABILITY

- **Representations describe the agent's uncertainty about its environment**
  - summarize the partial and conflicting evidence that's available to the agent
  - describe a set of situations that the agent regards as possible
  - weight each according to how likely the agent's evidence makes it

AI Course B
PROBABILITY

- **Key techniques allow these representations to be**
  - specified,
  - accessed to guide the agent's activity in its real world task, and
  - updated in response to new information

AI Course B
PROBABILITY

- **Assumptions take the form of**
  - statements of independence, so that two pieces of information give no evidence one for the other, or vice versa
  - models of processes, that set the form of functions assigning likelihood to situations
  - parameters for prior probabilities, in which a designer communicates background expectations about the world to the agent

A Syllabus on PROBABILITY in AI

- Simple pattern classification
  - Bayes decision theory and parameter estimation
- Structured discrete patterns
  - Hidden Markov models and probabilistic context-free grammars
- Structured continuous patterns
  - Kalman filters and particle filters
- Belief nets (Bayes nets or graphical models)
- Decision trees and Markov decision processes