Lecture 1

The Intelligent Agent Framework

Friday 22 August 2003

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

Reading for Next Class:
Chapter 2, Russell and Norvig

Lecture Outline

- Today's Reading: Chapter 2, Russell and Norvig
- Intelligent Agent (IA) Design
  - Shared requirements, characteristics of IAs
  - Methodologies
    - Software agents
    - Reactivity vs. state
    - Knowledge, inference, and uncertainty
- Intelligent Agent Frameworks
  - Reactive
  - With state
  - Goal-based
  - Utility-based
- Thursday: Problem Solving and Search
  - State space search handout (Winston)
  - Search handout (Ginsberg)

Review: Course Topics

- Overview: Intelligent Systems and Applications
- Artificial Intelligence (AI) Software Development Topics
  - Knowledge representation
    - Logical
    - Probabilistic
  - Search
    - Problem solving by (heuristic) state space search
    - Game tree search
  - Planning: classical, universal
  - Machine learning
    - Models (decision trees, version spaces, ANNs, genetic programming)
    - Applications: pattern recognition, planning, data mining and decision support
  - Topics in applied AI
    - Computer vision fundamentals
    - Natural language processing (NLP) and language learning survey
- Implementation Practicum – 1-2 Students per Team

Intelligent Agents: Overview

- **Agent: Definition**
  - Any entity that perceives its environment through sensors and acts upon that environment through effectors
  - Examples (class discussion): human, robotic, software agents
- **Perception**
  - Signal from environment
  - May exceed sensory capacity
- **Sensors**
  - Acquires percepts
  - Possible limitations
- **Action**
  - Attempts to affect environment
  - Usually exceeds effector capacity
- **Effectors**
  - Transmits actions
  - Possible limitations

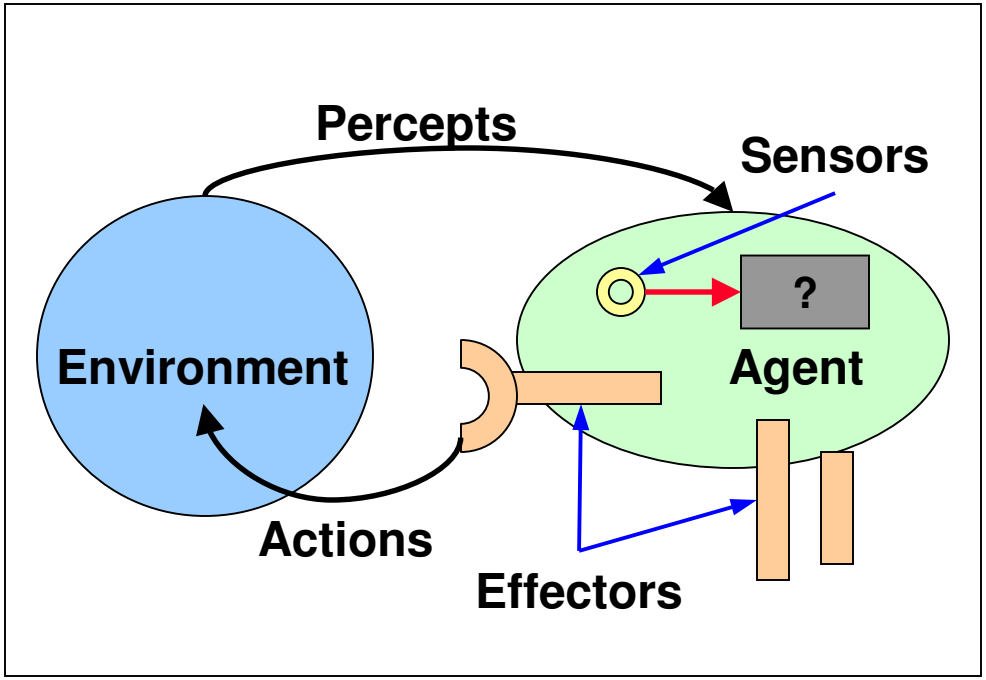

How Agents Should Act

- **Rational Agent: Definition**
  - Informal: "does the right thing, given what it believes from what it perceives"
  - What is "the right thing"?
    - First approximation: *action that maximizes success of agent*
    - Limitations to this definition?
  - Issues to be addressed now
    - How to evaluate *success*
    - When to evaluate *success*
  - Issues to be addressed later in this course
    - Uncertainty (in environment, in actions)
    - How to express beliefs, knowledge
- **Why Study Rationality?**
  - Recall: aspects of intelligent behavior (last lecture)
    - Engineering objectives: optimization, problem solving, decision support
    - Scientific objectives: modeling correct inference, learning, planning
  - Rational cognition: formulating *plausible* beliefs, conclusions
  - Rational action: "doing the right thing" given beliefs

Rational Agents

- "Doing the Right Thing"
  - Committing actions
    - Limited to set of effectors
    - In context of what agent knows
  - Specification (cf. software specification)
    - Preconditions, postconditions of operators
    - Caveat: not always perfectly known (for given environment)
    - Agent may also have limited knowledge of specification
- Agent Capabilities: Requirements
  - Choice: select actions (and carry them out)
  - Knowledge: represent knowledge about environment
  - Perception: capability to sense environment
  - Criterion: *performance measure to define degree of success*
- Possible Additional Capabilities
  - Memory (internal model of state of the world)
  - Knowledge about effectors, reasoning process (reflexive reasoning)

Measuring Performance

- **Performance Measure**: *How* to Determine Degree of Sucesss
  - Definition: *criteria that determine how successful agent is*
  - Clearly, varies over
    - Agents
    - Environments
  - Possible measures?
    - Subjective (agent may not have capability to give accurate answer!)
    - Objective: *outside observation*
  - Example: web crawling agent
    - Number of hits
    - Number of *relevant* hits
    - *Ratio* of relevant hits to pages explored, resources expended
    - Caveat: "you get what you ask for" (issues: redundancy, etc.)
- When to Evaluate Success
  - Depends on objectives (short-term efficiency, consistency, etc.)
  - Is task episodic? Are there milestones? Reinforcements? (e.g., games)

Knowledge in Agents

- **Rationality versus Omniscience**
  - Nota Bene (*NB*): not the same
  - Distinction
    - Omniscience: knowing *actual* outcome of all actions
    - Rationality: knowing *plausible* outcome of all actions
    - Example: is crossing the street to greet a friend too risky?
  - Key question in AI
    - *What is a plausible outcome?*
    - Especially important in knowledge-based expert systems
    - Of practical important in planning, machine learning
  - Related questions
    - *How can an agent make rational decisions given beliefs about outcomes of actions?*
    - *Specifically, what does it mean (algorithmically) to "choose the best"?*
- **Limitations of Rationality**
  - Based only on what agent *can* perceive and do
  - Based on what is "likely" to be right, not what "turns out" to be right

What Is Rational?

- **Criteria**
  - Determines what is rational *at any given time*
  - Varies with agent, environment, *situation*
- **Performance Measure**
  - Specified by outside observer or evaluator
  - Applied (consistently) to (one or more) IAs in given environment
- **Percept Sequence**
  - Definition: *entire history* of percepts gathered by agent
  - NB: may or may not be retained fully by agent (issue: state and memory)
- **Agent Knowledge**
  - Of environment – "required"
  - Of self (reflexive reasoning)
- **Feasible Action**
  - What can be performed
  - What agent believes it can attempt?

Ideal Rationality

- **Ideal Rational Agent**
  - Given: any possible percept sequence
  - Do: ideal rational behavior
    - Whatever action is expected to maximize performance measure
    - NB: expectation – informal sense (for now); mathematical foundation soon
  - Basis for action
    - Evidence provided by percept sequence
    - Built-in knowledge possessed by the agent
- **Ideal Mapping from Percepts to Actions**
  - Figure 2.2, R&N
  - Mapping *p*: *percept sequence* → *action*
    - Representing *p* as list of pairs: infinite (unless explicitly bounded)
    - Using *p*: specifies ideal mapping from percepts to actions (i.e., ideal agent)
    - Finding explicit *p*: in principle, could use trial and error
    - *Other (implicit) representations may be easier to acquire!*

Autonomy

- **Built-In Knowledge**
  - ***What if agent ignores percepts?***
  - **Possibility**
    - **All actions based on agent's own knowledge**
    - **Agent said to *lack* autonomy**
  - **Examples**
    - **"Preprogrammed" or "hardwired" industrial robots**
    - **Clocks**
    - **Other sensorless automata**
    - **NB: to be distinguished from closed versus open loop systems**
- **Justificiation for Autonomous Agents**
  - **Sound engineering practice: "Intelligence demands robustness, adaptivity"**
    - **Example: dung beetle (Egyptian scarab)**
    - **Ethological and evolutionary bases of knowledge**
  - **This course: mathematical and CS basis of autonomy in IAs**

Structure of Intelligent Agents

- **Agent Behavior**
  - Given: sequence of percepts
  - Return: IA's actions
    - Simulator: description of results of actions
    - Real-world system: committed action
- **Agent Programs**
  - Functions that implement *p*
  - Assumed to run in computing environment (architecture)
    - Hardware architecture: computer organization
    - Software architecture: programming languages, operating systems
  - *Agent = architecture + program*
    - This course (CIS730): primarily concerned with *p*
    - CIS540, 740, 748: concerned with architecture
    - See also: Chapter 24 (Vision), 25 (Robotics), R&N
- **Discussion: "Real" versus "Artificial" Environments**

Agent Programs

- **Software Agents**
  - Also known as (*aka*) software robots, softbots
  - Typically exist in very detailed, unlimited domains
  - Example
    - (Real-time) critiquing, automation of avionics, shipboard damage control
    - Indexing (spider), information retrieval (IR; e.g., web crawlers) agents
    - Plan recognition systems (computer security, fraud detection monitors)
  - See: Bradshaw (*Software Agents*)
- **Focus of This Course: Building IAs**
  - Generic skeleton agent: Figure 2.4, R&N
  - function *SkeletonAgent* (*percept*) returns action
    - static: *memory*, agent's memory of the world
    - *memory* ← *Update-Memory* (*memory, percept*)
    - *action* ← *Choose-Best-Action* (*memory*)
    - *memory* ← *Update-Memory* (*memory, action*)
    - return *action*

Example: Automated Taxi Driver

- Agent Type: Taxi Driver
- Percepts
  - Visual: cameras
  - Profilometer: speedometer, tachometer, odometer
  - Other: GPS, sonar, interactive (microphone)
- Actions
  - Steer, accelerate, brake
  - Talk to passenger
- Goals
  - Safe, legal, fast, comfortable
  - Maximize profits
- Environment
  - Roads
  - Other traffic, pedestrians
  - Customers
- Discussion: Performance Requirements for Open Ended Task

Terminology

- Artificial Intelligence (AI)
  - Operational definition: study / development of systems capable of "thought processes" (reasoning, learning, problem solving)
  - Constructive definition: expressed in artifacts (design and implementation)
- Intelligent Agents
- Topics and Methodologies
  - Knowledge representation
    - Logical
    - Uncertain (probabilistic)
    - Other (rule-based, fuzzy, neural, genetic)
  - Search
  - Machine learning
  - Planning
- Applications
  - Problem solving, optimization, scheduling, design
  - Decision support, data mining
  - Natural language processing, conversational and information retrieval agents
  - Pattern recognition and robot vision

Summary Points

- Artificial Intelligence: Conceptual Definitions and Dichotomies
  - Human cognitive modelling vs. rational inference
  - Cognition (thought processes) versus behavior (performance)
  - Some viewpoints on defining intelligence
- Roles of Knowledge Representation, Search, Learning, Inference in AI
  - Necessity of KR, problem solving capabilities in intelligent agents
  - Ability to reason, learn
- Applications and Automation Case Studies
  - Search: game-playing systems, problem solvers
  - Planning, design, scheduling systems
  - Control and optimization systems
  - Machine learning: pattern recognition, data mining (business decision support)
- More Resources Online
  - Home page for AIMA (R&N) textbook
  - CMU AI repository
  - KSU KDD Lab (Hsu): http://www.kddresearch.org
  - Comp.ai newsgroup (now moderated)