



Lecture 1 of 42

Intelligent Agents Overview Discussion: Problem Set 1, Term Projects 1 of 3

Wednesday, 22 August 2007

William H. Hsu
Department of Computing and Information Sciences, KSU

KSOL course page: <http://snipurl.com/v9v3>
Course web site: <http://www.kddresearch.org/Courses/Fall-2007/CIS730>
Instructor home page: <http://www.cis.ksu.edu/~bhsu>

Reading for Next Class:

Sections 1.3 – 1.5, p. 16 – 29, Russell & Norvig 2nd edition
Sections 2.1 – 2.2, p. 32 – 38, Russell & Norvig 2nd edition
Syllabus and Introductory Handouts



Lecture Outline

- **Reading for Next Class: Sections 1.3 – 1.5 & 2.1 – 2.2, R&N 2^e**
- **Today and Friday: Intelligent Agent (IA) Design, Chapter 2 R&N**
 - * **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**
- **Next Week: Problem Solving and Search, Chapter 3**
 - * **State space search handout (Nilsson, *Principles of AI*)**
 - * **Search handout (Ginsberg)**





Problems and Methodologies (Review)

- **Problem Solving**
 - * Classical search and planning
 - * Game-theoretic models
- **Making Decisions under Uncertainty**
 - * Uncertain reasoning, decision support, decision-theoretic planning
 - * Probabilistic and logical knowledge representations
- **Pattern Classification and Analysis**
 - * Pattern recognition and machine vision
 - * Connectionist models: artificial neural networks (ANNs), other graphical models
- **Data Mining and Knowledge Discovery in Databases (KDD)**
 - * Framework for optimization and machine learning
 - * Soft computing: evolutionary algorithms, ANNs, probabilistic reasoning
- **Combining Symbolic and Numerical AI**
 - * Role of knowledge and automated deduction
 - * Ramifications for cognitive science and computational sciences




Intelligent Agents (Review)

- **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





Generic Intelligent Agent Model (Review)



Term Project Topics, Fall 2007 (review)

- **1. Game-playing Expert System**
 - * “Borg” for Angband computer role-playing game (CRPG)
 - * <http://www.thangorodrim.net/borg.html>
- **2. Classic Trading Agent Competition (TAC)**
 - * **Supply Chain Management (TAC-SCM) scenario**
 - * <http://www.sics.se/tac/>
- **3. Machine Learning in Bioinformatics**
 - * **Evidence ontology for genomics or proteomics**
 - * <http://bioinformatics.ai.sri.com/evidence-ontology/>



Homework 1: Problem Set

- Assigned: 23:00 CDT Wed ²² ~~23~~ Aug ²⁰⁰⁷ ~~2006~~
- Due: before midnight CDT Wed ⁰⁶ ~~05~~ Sep ²⁰⁰⁷ ~~2006~~
- Topics
 - * Intelligent agents concepts
 - * State space representations
 - * Informed search
- To Be Posted
 - * KSOL web site
 - * KDDresearch.org (URL mailed to class mailing list)
- Questions and Discussion
 - * General discussion on class mailing list: CIS730-L@listserv.ksu.edu
 - * Questions for instructor: CIS730TA-L@listserv.ksu.edu
- Outside References: On Reserve (Cite Sources!)



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?
 - * First: how, when to evaluate success?
 - * Later: representing / reasoning with uncertainty, 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 effectors, in context of agent knowledge
 - * **Specification (cf. software specification):** pre/post-conditions
- **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 Success**
 - * **Definition:** criteria that determine how successful agent is
 - * **Depends on**
 - ⇒ Agents
 - ⇒ Environments
 - * **Possible measures?**
 - ⇒ **Subjective** (agent may not have capability to give accurate answer!)
 - ⇒ **Objective:** outside observation
 - * **Example: web crawling agent**
 - ⇒ **Precision:** did you get only pages you wanted?
 - ⇒ **Recall:** did you get all pages you wanted?
 - ⇒ **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.)
 - * **Episodic? Milestones? Reinforcements?** (e.g., games)





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: agent may or may not have state, i.e., 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 def'n later
 - * **Basis for action**
 - ⇒ Evidence provided by percept sequence
 - ⇒ Built-in knowledge possessed by the agent
- **Ideal Mapping from Percepts to Actions (Figure 2.1 p. 33 R&N 2^e)**
 - * Mapping p : percept sequence → action
 - * Representing p as list of pairs: infinite (unless explicitly bounded)
 - * Using p : 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!





Knowledge and Bounded Rationality

- **Rationality versus Omniscience**
 - * **Nota Bene (NB):** not the same
 - * **Omniscience:** knowing *actual* outcome of all actions
 - * **Rationality:** knowing *plausible* outcome of all actions
 - * **Example:** is it too risky to go to the supermarket?
- **Key Question**
 - * What is a *plausible* outcome of an action?
 - * **Related questions**
 - ⇒ How can agents make rational decisions given beliefs about outcomes?
 - ⇒ What does it mean (algorithmically) to “choose the best”?
- **Bounded Rationality**
 - * What agent *can* perceive and do
 - * What is “likely” to be right – not what “turns out” to be right



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**)
 - * $Agent = architecture + program$
 - * This course (CIS730): primarily concerned with p
- **Applications**
 - * Chapter 22 (NLP/Speech), 24 (Vision), 25 (Robotics), R&N 2e
 - * Swarm intelligence, multi-agent systems, IAs in cybersecurity



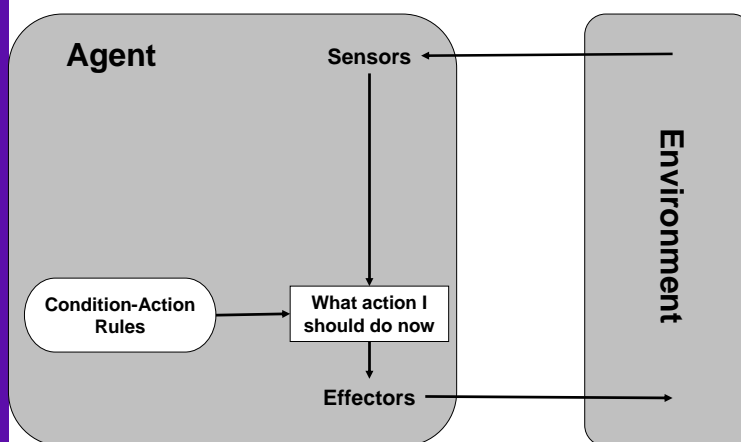


Example: Game-Playing Agent [2] Problem Specification

- **Angband**
 - * **Roguelike game** – descended from Rogue, Moria
See: <http://en.wikipedia.org/wiki/Roguelike>
 - * **v2.8.3**
Source code: <http://www.thangorodrim.net>
- **Automated Roguelike Game-Playing Agents**
 - * **Rog-O-Matic (1984)**
<http://en.wikipedia.org/wiki/Rog-O-Matic>
 - * **Angband Borks (1998-2001)**
<http://www.thangorodrim.net/borg.html>
- **Problem Specification**
 - * **Study Borks by Harrison, White**
 - * **Develop a scheduling, planning, or classification learning system**
 - * **Use White's APWBorg interface to develop a new Borg**
 - * **Compare it to the classic Borks**




Agent Framework: Simple Reflex Agents [1]

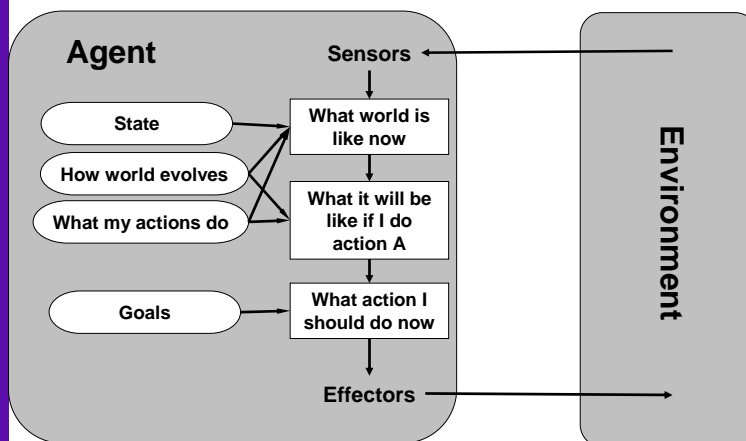




Agent Frameworks: (Reflex) Agents with State

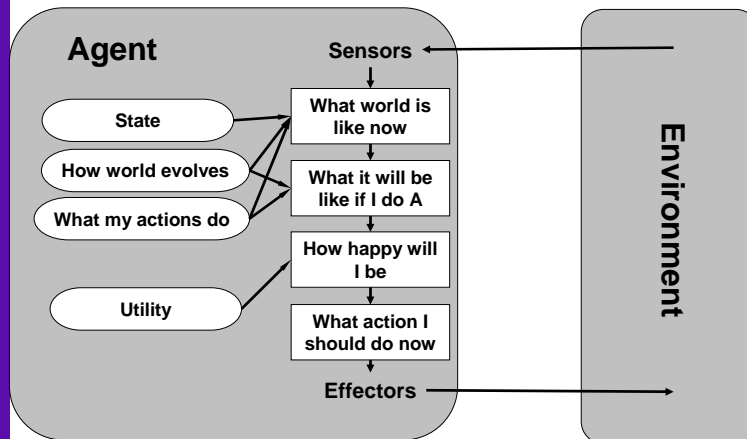


Agent Frameworks: Goal-Based Agents





Agent Frameworks: Utility-Based Agents



Course Topics Fall, 2007

- **Overview: Intelligent Systems and Applications**
- **Artificial Intelligence (AI) Software Development Topics**
 - * Knowledge representation
 - * Search
 - * Expert systems and knowledge bases
 - * Planning: classical, universal
 - * Probabilistic reasoning
 - * Machine learning, artificial neural networks, evolutionary computing
 - * Applied AI: agents focus
 - * Some special topics (NLP focus)
- **Implementation Practicum (≈ 40 hours)**



Terminology

- **Rationality**
 - * Informal definition
 - * Examples: how to make decisions
 - * Ideal vs. bounded
- **Automated Reasoning and Behavior**
 - * Regression-based problem solving (see p. 7)
 - * Goals
 - * Deliberation
- **Intelligent Agent Frameworks**
 - * Reactivity vs. state
 - * From goals to preferences (utilities)



Summary Points

- **Intelligent Agent Framework**
- **Rationality and Decision Making**
- **Design Choices for Agents (Introduced)**
- **Choice of Project Topics**
 - * 1. Game-playing expert system: Angband
 - * 2. Trading agent competition, supply chain management (TAC-SCM)
 - * 3. Knowledge base for bioinformatics: proteomics ontology
- **Things to Check Out Online**
 - * Resources page
<http://www.kddresearch.org/Courses/Fall-2007/CIS730/Resources>
 - * Course mailing list archives (class discussions)
<http://listserv.ksu.edu/archives/cis730-l.html>

