Lecture 0 of 42

Artificial Intelligence: Course Organization and Survey

Monday, 21 August 2006

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-2006/CIS730
Instructor home page: http://www.cis.ksu.edu/~bhsu

Reading for Next Class:
Chapter 1, Russell and Norvig 2nd edition
Syllabus and Introductory Handouts

Course Outline

- 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
- Practicum (Short Software Implementation Project)

Course Administration

- Official Course Page (KSOL): http://snipurl.com/v9v3
- Class Web Page: http://www.kddresearch.org/Courses/Fall-2006/CIS730
- Instructional E-Mail Addresses
  - CIS730TA-L@listserv.ksu.edu (always use this to reach instructor)
  - CIS730-L@listserv.ksu.edu (this goes to everyone)
- Instructor: William Hsu, Nichols 213
  - Office phone: (785) 532-7905; home phone: (785) 539-7180
  - IM: AIM/YIM/MSN hsuwh & rizanabsith, ICQ 28651394 & 191317559
  - Office hours: after class Mon/Wed/Fri; other times by appointment
- Graduate Teaching Assistant: TBD
  - Office location: Nichols 124
  - Office hours: to be announced on class web board
- Grading Policy
  - Midterm: 20% (in-class, closed-book); final (open-book): 30%; quiz: 3%
  - Machine problems, problem sets (8 of 10): 16%; term project: 26%
  - Class participation: 5% (3% attendance, 1% questions, 1% answers)

How To Get an A in This Course

- A Story from Dr. Gerard G. L. Meyer, Johns Hopkins University
- Ask Questions
  - Ask for (more) examples, another explanation, etc. if needed ("don't be shy")
  - All students (especially remote students): post in class web board
    - Unclear points – bring to class as well
    - "When will *X* happen"?
  - Fastest way to reach instructor: instant messaging (ICQ, MSN Messenger)
  - Notify TA, KDD system administrators of any computer problems
- Be Aware of Resources
  - Check with instructor or GTA about
    - Handouts, lectures, grade postings
    - Resources online
  - Check with classmates about material from missed lecture
- Start Machine Problems (and Problem Sets) Early
  - How to start virtuous (as opposed to vicious) cycle
  - Don't cheat

Homework Assignments: Problem Sets and Machine Problems

- PS1 assigned Wed 23 Aug 2006, due Wed 06 Sep 2006
- MP2 assigned Wed 06 Sep 2006, due Fri 22 Sep 2006
  - Submit using K-State Online
  - HW page: http://www.kddresearch.org/Courses/Fall-2006/CIS730/Homework
- Model solutions: 2 class days after due date
- Graded assignements: ≤ 7 days after due date
- Machine Problem: Search
  - Problem specifications to be posted on homework page before 06 Sep 2006
  - Languages: C/C++ & Java
  - MP guidelines
    - Work individually
    - Generate standard output files and test against partial standard solution
    - No late submissions except with documented excusal (medical, etc.)
  - See also: state space, constraint satisfaction problems

Questions Addressed

- Problem Area
  - What *are* intelligent systems and agents?
  - Why are we interested in developing them?
- Methodologies
  - What kind of software is involved? What kind of math?
  - How do we develop it (software, repertoire of techniques)?
  - Who uses AI? (Who are practitioners in academia, industry, government?)
- Artificial Intelligence as A Science
  - What is AI?
  - What does it have to do with intelligence? Learning? Problem solving?
  - What are interesting problems to which intelligent systems can be applied?
  - Should I be interested in AI (and if so, why)?
- Today: Brief Tour of AI History
  - Study of intelligence (since classical age), AI systems (1940-present)
  - Viewpoints: philosophy, math, psychology, engineering, linguistics

What is AI? [1]

- Four Categories of Systemic Definitions
  - 1. Think like humans
  - 2. Act like humans
  - 3. Think *rationally*
  - 4. Act *rationally*
- Thinking Like Humans
  - *Machines with minds* (Haugeland, 1985)
  - Automation of "decision making, problem solving, learning…" (Bellman, 1978)
- Acting Like Humans
  - *Functions* that *require intelligence when performed by people* (Kurzweil, 1990)
  - Making computers do things *people currently do better* (Rich & Knight, 1991)
- Thinking Rationally
  - Computational models of mental faculties (Charniak & McDermott, 1985)
  - Computations that make it possible to *perceive*, *reason*, and *act* (Winston, 1992)
- Acting Rationally
  - Explaining, emulating intelligent behavior via computation (Schalkoff, 1990)
  - Branch of CS: automating intelligent behavior (Luger, 2005)

What is AI? [2]
Thinking and Acting Like Humans

- Concerns: Human Performance (Figure 1.1 R&N, Left-Hand Side)
  - Top: thought processes and reasoning (learning and inference)
  - Bottom: behavior (interacting with environment)
- Machines With Minds
  - Cognitive modelling
    - Early historical examples: problem solvers (see R&N Section 1.1)
    - Application (and one driving force) of cognitive science
  - Deeper questions
    - What is intelligence?
    - What is consciousness?
- Acting Humanly: The Turing Test Approach
  - Capabilities required
    - Natural language processing
    - Knowledge representation
    - Automated reasoning
    - Machine learning
  - Turing Test: can a machine appear indistinguishable from a human to an experimenter?

What is AI? [3]
Viewpoints on Defining Intelligence

- Genuine versus Illusory Intelligence
  - Can we tell?
    - ⇨ If so, how?
    - ⇨ If not, what limitations do we postulate?
  - The *argument from disability* ("a machine can never do X")
- Turing Test Specification
  - Objective: develop intelligent system "indistiguishable from human"
    - ⇨ Blind interrogation scenario (no direct physical interaction – "teletype")
    - ⇨ 1 AI system, 1 human subject, 1 interrogator
    - ⇨ Variant: total Turing Test (perceptual interaction: video, tactile interface)
  - *Is this a reasonable test of intelligence?*
  - Details: Section 26.3, R&N
  - See also: Loebner Prize page
- Searle's Chinese Room
  - Philosophical issue: is (human) intelligence a pure artifact of symbolic manipulation?
  - Details: Section 26.4, R&N
  - See also: consciousness in AI resources

What is AI? [3]
Thinking and Acting Rationally

- Concerns: Human Performance (Figure 1.1 R&N, Right-Hand Side)
  - Top: thought processes and reasoning (learning and inference)
  - Bottom: behavior (interacting with environment)
- Computational Cognitive Modelling
  - Rational ideal
    - ⇨ In this course: rational agents
    - ⇨ Advanced topics: learning, utility theory, decision theory
  - Basic mathematical, computational models
    - ⇨ Decisions: automata (Chomsky hierarchy – FSA, PDA, LBA, Turing machine)
    - ⇨ Search
    - ⇨ Concept learning
- Acting Rationally: The Rational Agent Approach
  - Rational action: *acting to achieve one's goals, given one's beliefs*
  - Agent: entity that perceives and acts
  - Focus of next lecture
    - ⇨ "Laws of thought" approach to AI: correct inferences (reasoning)
    - ⇨ Rationality not *limited to* correct inference

What is AI? [4]
A Brief History of The Field

- Philosophy Foundations (400 B.C. – present)
  - Mind: dualism (Descartes), materialism (Leibniz), empiricism (Bacon, Locke)
  - Thought: syllogism (Aristotle), induction (Hume), logical positivism (Russell)
  - Rational agentry (Mill)
- Mathematical Foundations (c. 800 – present)
  - Early: algorithms (al-Khowarazmi, 9th century mathematician), Boolean logic
  - Computability (20th century – present)
    - ⇨ Cantor diagonalization, Gödel's incompleteness theorem
    - ⇨ Formal computuational models: Hilbert's Entscheidungsproblem, Turing
    - ⇨ Intractability and NP-completeness
- Computer Engineering (1940 – present)
- Linguistics (1957 – present)
- Stages of AI
  - Gestation (1943 – c. 1956), infancy (c. 1952 – 1969)
  - Disillusioned early (c. 1966 – 1974), later childhood (1969 – 1979)
  - "Early" (1980 – 1988), "middle" adolescence (c. 1985 – present)

Why Study Artificial Intelligence?

- New Computational Capabilities
  - Advances in uncertain reasoning, knowledge representations
  - Learning to act: robot planning, control optimization, decision support
  - Database mining: converting (technical) records into knowledge
  - Self-customizing programs: learning news filters, adaptive monitors
  - Applications that are hard to program: driving, speech recognition
- Better Understanding of Human Cognition
  - Cognitive science: theory of knowledge acquisition (e.g., through practice)
  - Performance elements: reasoning (inference) and *recommender* systems
- Time is Right
  - Recent progress in algorithms and theory
  - Rapidly growing volume of online data from various sources
  - Available computational power
  - Growth of AI-based industries (e.g., data mining, robotics, web search)

Artificial Intelligence: Some Problems and Methodologies

- 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: 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

A Generic Intelligent Agent Model

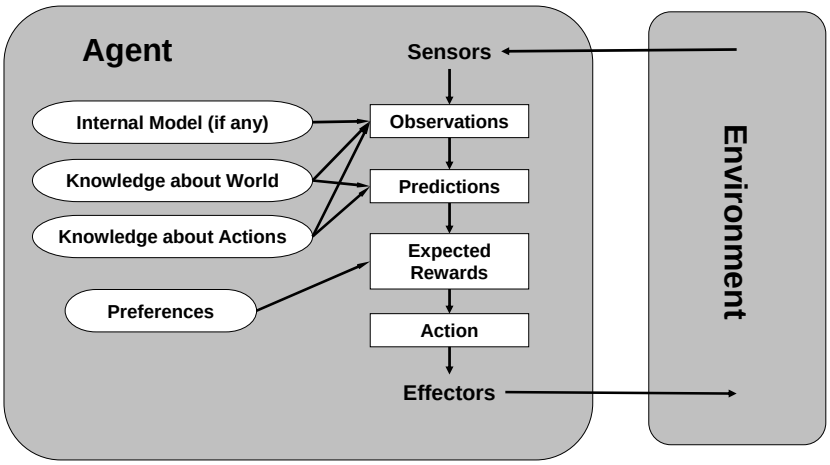

Term Project Topics Fall 2006

- 1. Game-playing Expert System
  - "Borg" for Angband computer role-playing game (CRPG)
  - http://www.thangorodrim.net/borg.html
- 2. Trading Agent Competition (TAC)
  - Supply Chain Management (TAC-SCM) scenario
  - http://www.sics.se/tac/page.php?id=13
- 3. Knowledge Base for Bioinformatics
  - Evidence ontology for genomics or proteomics
  - http://bioinformatics.ai.sri.com/evidence-ontology/

Term Project Guidelines

- Due: 17 Nov 2006
  - Project milestones: initial (plan), interrim (interview), final (presentation)
  - Presentations, peer review outside class
- Individual Projects
  - Topic selection due Fri 01 Sep 2006
  - First draft of project plan due Fri 08 Sep 2006
- Grading: 260 points (out of 1000)
  - Proposal: 20 points
  - Interview: 20 points
  - Presentation: 20 points
  - Project content: 160 points
    - Originality: 40 points
    - Functionality: 40 points
    - Development effort: 40 points
    - Completeness: 40 points
  - Writeup: 40 points

Related Online Resources

- Research
  - KSU Laboratory for Knowledge Discovery in Databases http://www.kddresearch.org (see especially Group Info, Web Resources)
  - KD Nuggets: http://www.kdnuggets.com
- Courses and Tutorials Online
  - At KSU
    - CIS732 *Machine Learning and Pattern Recognition* http://www.kddresearch.org/Courses/Fall-2005/CIS732
    - CIS830 *Advanced Topics in Artificial Intelligence* http://www.kddresearch.org/Courses/Spring-2005/CIS830
    - CIS690 *Implementation of High-Performance Data Mining Systems* http://ringil.cis.ksu.edu/Courses/Summer-2005/CIS690
  - Other courses: see KD Nuggets, www.aaai.org, www.auai.org
- Discussion Forums
  - Newsgroups: comp.ai.*
  - Recommended mailing lists: *Data Mining*, *Uncertainty in AI*
  - KDD Group Mailing List (KDD-L@listserv.ksu.edu)

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, information retrieval and extraction (IR/IE)
  - 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 (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)