Lecture 1: Introduction

Victor R. Lesser

CMPSCI 683

Fall 2004

Acknowledgements

- Many Slides Courtesy of Shlomo Zilberstein
- Also Slides Courtesy of Milind Tambe (USC), Tuomas Sandholm (CMU) and many others!

Contact Information

- Instructor: Victor R. Lesser
  413-545-1322, fax 413-545-1249
  Email: lesser@cs.umass.edu
  Office hours:
  Tuesdays, 1:00 - 2:30pm, CS306
- T.A.: Mike O'Neill
  Email: mpo@cs.umass.edu
  Office hours:
  Wednesdays 1:00 - 3:00, TBA
- *Course Web Page: http://mas.cs.umass.edu/classes/cs683/*
- *Learning Support Services will contain DVD on the 10th floor of the DuBois Library*

Course Prerequisites

- Nominally you need an undergraduate course in AI
- Not necessary to be successful in course
- Will move quickly over elementary material

Some General Comments

- **Encouragement to ask questions during class**
  - *Without your feedback, it is impossible for me to know what you don't know*
  - *There is no reason not to ask questions during class*
  - *Of course, you could also send email, or meet in person*
- ***Encouragement to read course material prior to class***

Course Objectives

After this course you will be able to...

- **understand state-of-the-art AI techniques**
- **construct simple AI systems**
- **read the AI literature**
- **evaluate AI-related technology claims**
- **apply AI techniques in non-AI settings**
- **pursue specialized AI courses and research**

What is AI about?

Historical Definition of AI

"The study of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."

– Dartmouth Workshop, Summer of 1956

AI has Multiple Goals

- **Constructing intelligent machines for a wide range of applications**
  - Augmenting human problem solving
- **Formalizing knowledge and mechanizing intelligence**
- ***Using computational models to understand complex behavior***
- ***Making computers as easy to work with as people***

Why is AI Important?

Complexity of Work Place (Rappaport, 1991)

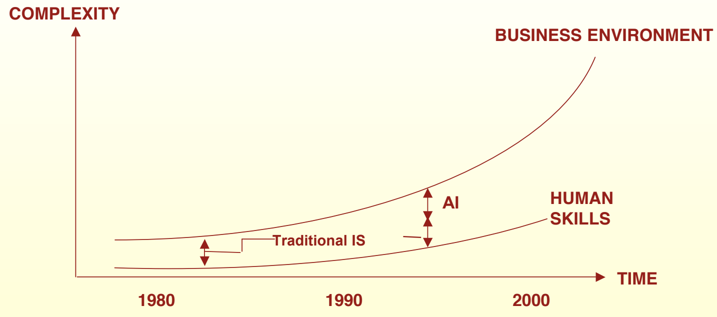

IS = Information Systems

Why is AI Practical?

Cheaper Sensors and Better Sensor Processing

- vision, speech and text

Orders of magnitude increase in computing power

- Commodity processors (>billion instructions per second)
- Special purpose processors for sensor processing
- Parallel computing (fine and large grain)
- Distributed processing and high-speed communication (NII)
- Large on-line knowledge bases

Tremendous Progress in Providing Formal and Practical Underpinning to the Discipline over the last 20 years

Why AI is Hard?

- **Ambiguity: "**I dropped the egg on the table and **it** broke."
- **Resource limitations: "**Is there a winning opening move in chess?"
  - Many applications are NP-complete
- **Conflicting goals and trade-offs:** "Computer, I need more free disk space!"
- **Uncertainty and missing information**
  - Missing information, probabilistic actions, open and dynamic environments.

CONVENTIONAL SOFTWARE VS. AI

Conventional software can be contrasted with AI systems in a number of ways:

- **Language: arithmetic vs. (logical, probalistic)**
- **Data: numbers vs. symbols**
- **Coding: procedural (code) vs. declarative (sentences)**
- **Operations: calculations vs. reasoning (symbolic and decision theorectic)**
- **Knowledge: formulas vs. heuristics**
- **Program output: deterministic vs. nondeterministic (certain/well-defined vs. uncertain/search)**
- **Problem specification: precise vs. imprecise/relative**
- **Solution quality: exact/optimal vs. approximate/satisficing**

What can AI systems do today? - 1

- Scheduling and Planning:
  - DARPA's DART system used in Desert Storm and Desert Shield operations to plan logistics of people and supplies.
  - American Airlines rerouting contingency planner.
  - European space agency planning and scheduling of spacecraft assembly, integration and verification.
- Speech Recognition:
  - PEGASUS spoken language interface to American Airlines' EAASY SABRE reservation system.
- Computer Vision:
  - Face recognition programs in use by banks, government, etc.
  - The ALVINN system from CMU autonomously drove a van from Washington, D.C. to San Diego (all but 52 of 2,849 miles), averaging 63 mph day and night, and in all weather conditions.
  - Handwriting recognition, electronics and manufacturing inspection, baggage inspection, automatically construct 3D geometric models.

What can AI systems do today? - 2

- Diagnostic Systems:
  - Microsoft Office Assistant.
  - Pathfinder for diagnosis of lymph-node diseases
    - Outperforms experts that designed it; approved by AMA
  - Whirlpool customer assistance center.
- System Configuration:
  - DEC's XCON system for custom hardware configuration.
- Financial Decision Making:
  - Fraud detection and transaction approval by credit card companies, mortgage companies, banks, and the U.S. government.
  - Improving prediction of daily revenues and staffing requirements for a business.
  - Help desk support systems to find the right answer to any customer's question.

What can AI systems do today? - 3

- Classification Systems: **NASA's system for classifying very faint areas in astronomical images into either stars or galaxies with very high accuracy by learning from human experts.**
- Mathematical Theorem Proving: **Use inference methods to prove new theorems.**
- Game Playing: **Computer programs beat world's best players in chess, checkers, and backgammon.**
- Machine Translation:
  - AltaVista's translation of web pages.
  - Translation of Caterpillar Truck manuals into 20 languages.

Agent-centric view of AI

An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through effectors.

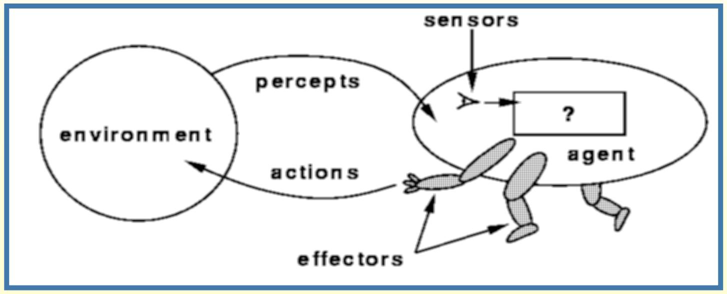

Basics: What is an Agent?

PROPERTY	MEANING
– *Situated*	Sense and act in dynamic/uncertain environments
– *Flexible*	Reactive (responds to changes in the environment) Pro-active (acting ahead of time)
– *Autonomous*	Exercises control over its own actions
– *Goal-oriented*	Purposeful
– *Learning*	Adaptive
– *Persistent*	Continuously running process
– *Social*	Interacts with other agents/people
– *Mobile*	Able to transport itself
– *Personality*	Character, Emotional state

Types of Agents

- Simple reflexive agents
- Agents that keep track of the world
- Deliberative agents
  - Explicit representation of domain knowledge
  - Run-time problem-solving and reasoning
  - Modeling the world
  - Self-modeling and meta-reasoning

Agents

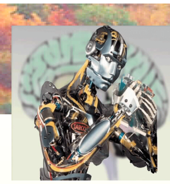

- Applications:
  - *Information gathering, integration*
  - *Distributed sensors*
  - *E-commerce*
  - *Distributed virtual organization*
  - *Virtual humans for training, entertainment*
- ***Rapidly growing area***

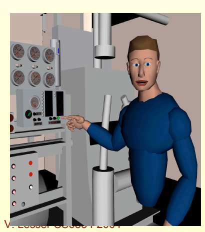

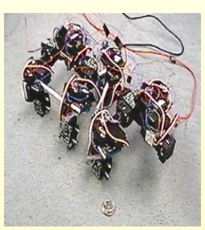

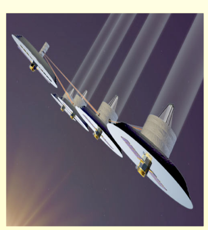

Traditional Systems vs. Intelligent Agents

- **How are agents different from the traditional view of system definition?**
  - When sensor = static input, effectors = fixed output, the goal is to produce the correct output, and environment is static/ irrelevant, we fall into the category of traditional systems.
- **BUT:**
  - Environment may be dynamic
  - Sensing may be an on-going situation assessment process
  - Effectors may require complex planning
  - Goal may be defined with respect to current state of environment
- **As a result**:
  - Deriving the input/output mapping (from the goal) is not obvious!

How to decide what to do?

An Ideal Agent

- Given:
  - The performance measure
  - The percept sequence
  - The agent's knowledge
  - The set of available actions
- An ideal agent will outperform any other agent in maximizing the performance measure.

A Lofty Goal But AI is a Very Long Way from Being Able to Address this Question for Complicated Applications

Example 1: Information Gathering Agent

- BIG (resource Bounded Information Gathering)
  - Takes role of human in support of decision process.
  - Integration of Planning, scheduling, text processing and interpretation style reasoning.
  - Helps pick software packages.

Motivation

- Rapid growth of WWW.
- Growth has outstripped technology.
- Information Retrieval technology a start.
  - Efficient, fast, general.
  - Access to enormous amount of data.
  - Browsing & processing documents manually non-trivial.

Sample Query

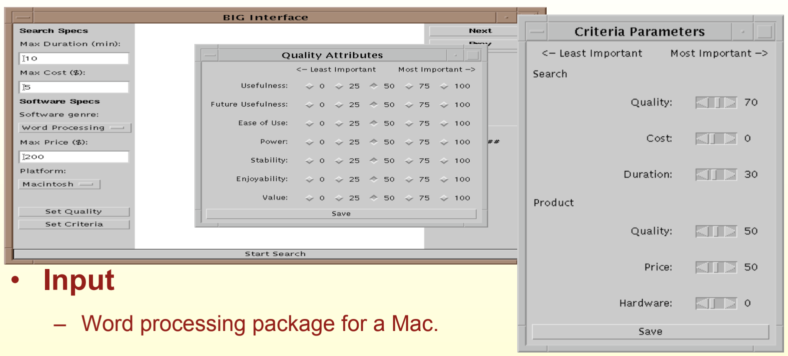

- Input
  - Word processing package for a Mac.
  - $200 price limit.
  - Search process should take 10 min. & cost less than $5.

Sample Trace, Cont.

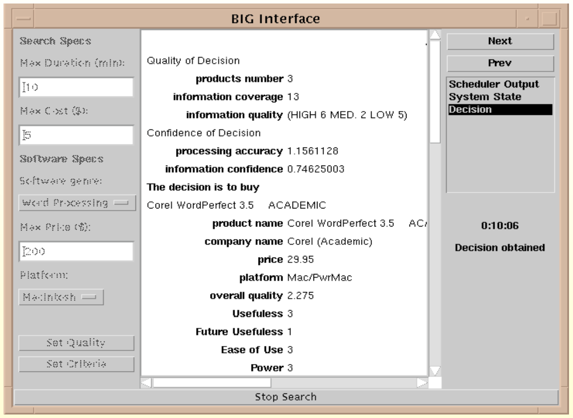

- BIG recommends Corel WP3.5

The BIG Agent: Salient Features

- Active search and discovery.
  - WWW pages on the internet
- Resource Bounded Reasoning.
  - Can only process a limited amount of information with in resource constraints (time and money).
- Goal-driven and Opportunistic control
  - Focused on specific goals but need to react to emerging information.
- Information extraction
  - Structured and unstructured information
- Information fusion.
  - Incomplete, unreliable, contradictory information

Example 2: Sophisticated Auction Agent *[Courtesy Tuomas Sandholm]*

- Auctioning multiple *distinguishable* items when bidders have preferences over combinations of items
- Example applications
  - Allocation of transportation tasks
  - Allocation of bandwidth
    - Dynamically in computer networks
    - Statically e.g. by FCC
  - Manufacturing procurement
  - Electricity markets
  - Securities markets
  - Liquidation
  - Reinsurance markets
  - Retail ecommerce: collectibles, flights-hotels-event tickets
  - Resource & task allocation in operating systems & mobile agent platforms

Combinatorial reverse auction *[Courtesy: Tuomas Sandholm]*

- Example: procurement in supply chains
- Auctioneer wants to buy a set of items (has to get all)
- Sellers place bids on how cheaply they are willing to sell bundles of items

Combinatorial Auctions

Auctioneer's winner determination problem

- Set of items, $M = \{1, 2, \ldots, \#items\}$
- Set of bids, $\beta = \{B_1, B_2, \ldots, B_{\#bids}\}$
- $B_j = \langle S_j, p_j \rangle$, where $S_j \subseteq M$ is a set of items and p_j is a price.
- $S_j \neq S_k$ (if multiple bids concern the same set of items, all but the highest bid can be discarded by a preprocessor)
- Problem: Label the bids as winning ($x_j = 1$) or losing ($x_j = 0$) so as to maximize auctioneer's *revenue* such that each item is allocated to, at most, one bid:

$$\max \sum_{j=1}^{\#bids} p_j x_j \quad \text{s.t.} \quad \sum_{j \mid i \in S_j} x_j \leq 1 \qquad i = 1, 2, \ldots, \#items$$

$$x_j \in \{0, 1\}$$

Courtesy of Tuomas Sandholm

Space of Allocations

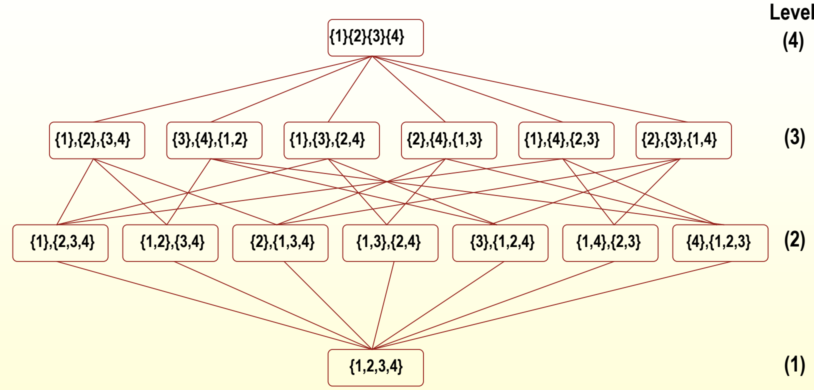

NP-complete

Question For the Next Class

- **What are the ideas that you can use to develop a search algorithm for winner determination for 1000's of bids and 1000's of items in a reasonable time frame (seconds) ?**
  - Optimal, anytime, satisficing

THEME OF COURSE:

A Computational Perspective on the Design of Intelligent Agents

- What is the form of the computational structures that are required?
- How does this differ from, or relate to, other computational problems?

A Perspective on the Design of Intelligent Agents

- **There is no universal approach to the design of an agent.**
- **We will be exploring the design space.**
  - Components and architectures
- **Different approaches…**
  - For different classes of problems
  - For different environments
  - For different criteria for success

Another Way of Looking at the Course

- Dealing with the Ubiquity of Uncertainty
  - **Environment, Sensing, Action and Knowledge**
- Dealing with Limited Resources
  - **Computation, memory, communication bandwidth, etc.**

Course Structure

- Text Book: **"Artificial Intelligence: A Modern Approach," (2nd Edition) Stuart Russell and Peter Norvig, Prentice Hall, 2003.**
  - Book's website <http://aima.cs.berkeley.edu>
  - *Will Augment with Material Not Covered in Book*
- Additional Readings (Available on the course web site):
  - Required readings
  - Suggested readings
  - Will add to suggested readings to ensure you have plenty of reference material
  - Not necessary to read through suggested readings

Grading

- **Factors:**
  - **Homework (40%)**
    - **Will include programming assignments**
  - **Midterm exam (30%)**
  - **Final exam (30%)**
- **Late Policy: Usually each assignment has two weeks time to finish; Assignments will not be accepted later without the express permission of the instructor or the teaching assistant.**

Syllabus Outline

- **Introduction -1**
- **Problem solving using sophisticated search - 7**
- **Reasoning under uncertainty - 8**
- **Learning - 7**
- **Intelligent Systems - 3**
- **Summary -1**

- **Won't be Covered**
  - Adversarial Game Playing
  - Elementary Logical Formalism for Knowledge Representation
- **May add in lectures on Planning, Knowledge Representation**

Detailed Syllabus -1

- **Introduction**
  - Lecture 1 - Introduction and Course Information
- **Search**
  - Lecture 2 - Overview of Issues in Heuristic Search
  - Lecture 3 - Heuristic Search
  - Lecture 4 - Search Complexity and Applications
  - Lecture 5 - Time and Space Variations of A*
  - Lecture 6 - Abstraction, Approximation, and Real-Time Search
  - Lecture 7 - Iterative Improvement Search/GSAT
  - Lecture 8 - Constraint Satisfaction and Genetic Algorithms

Detailed Syllabus -2

- **Reasoning Under Uncertainty**
  - Lectures 9 and 10 - Blackboard Systems as an Architecture for Interpretation
  - Lecture 11 - Representing and Reasoning with Uncertain Information
  - Lectures 12 and 13 - Probabilistic Reasoning with Belief Networks
  - Lecture 14 - Alternative Models of Uncertainty
  - Lecture 15 - Decision Theory
  - Lecture 16 - Decision Networks

May add lecture Hidden Markov Models (HMMs)

Detailed Syllabus -3

- **Learning**
  - Lecture 17 - Learning from Observations
  - Lecture 18 - Learning Techniques
  - Lecture 19 - Neural Networks
  - Lectures 20 and 21 - Markov Decision Processes and Reinforcement Learning
  - Lecture 22 - Data Mining
  - Lecture 23 - Analytical Learning and Planning

May add lecture on Kernel Machines

Detailed Syllabus - 4

- **Intelligent Systems**
  - Lecture 24 - Resource-Bounded Reasoning Systems
  - Lectures 25 and 26 - Intelligent Agent Architectures
  - Lecture 27 - Multi-Agent Systems and Summary

A Model for Computation in the 21st Century

Network of Cooperating, Intelligent Agents (people/machines)

- **Constructionist perspective**
  - build out of heterogeneous systems
  - high-level artificial language for cooperation
  - problem solving for effective cooperation will be as or more sophisticated than the actual domain problem solving
    - reasoning about goals, plans, intentions, and knowledge of other agents

A Model for Computation in the 21st Century

- **Operate in a "satisficing" mode**
  - Do the best they can within available resource constraints
  - Deal with uncertainty as an integral part of network problem solving
  - Complex organizational relationships among agents
- **Highly adaptive/highly reliable**
  - Learning will be an important part of their structure (short-term/long-term)
  - Able to adapt their problem-solving structure to respond to changing task/environmental conditions

Profound implications Computer Science!

How Does This Course Relate to This Model?

- **"Satisficing" Computation/Bounded Rationality**
  - Computational framework that allows you to trade off the quality of the answer derived with the amount of resources used to derive it
- **Uncertainty/inconsistency as integral part of problem solving**
  - Computational framework that allows you to live with it — rather than eliminate it
- **Intelligent Control**
  - Computational framework that allows you to effectively manage your resources to satisfy the given goals
- **Agency/Semi-Autonomous Agent**
  - Computational framework that allows agents to interact autonomously with the world in terms of sensing, perceiving, planning, effecting and communicating

Next Lecture

- **Why is search the key problem-solving technique in AI?**
- **Formulating and solving search problems.**
- **Reading: Sections 3.1-3.7.**

Recent Focus of Lab's Activity: Vehicle Tracking Distributed Sensor Network

- Small 2D Doppler radar units (30's)
  - Scan one of three 120° sectors at a time
- Commodity Processor associated with each radar
- Communicate short messages using one of 8 radio channels
- Triangulate radars to do tracking

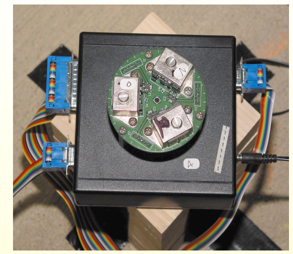

New Application: Real-Time Tornado Tracking (CASA)

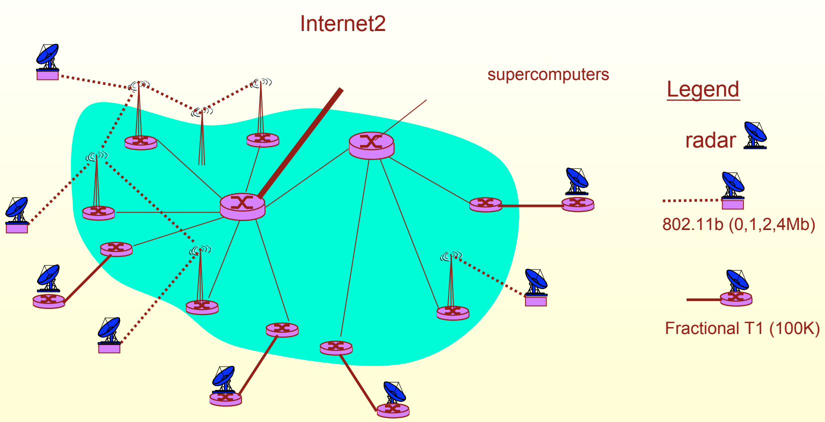

SRTA: Soft Real-Time Control Architecture

Multiple, Time and Resource Sensitive Goals, Planning, Scheduling and Execution

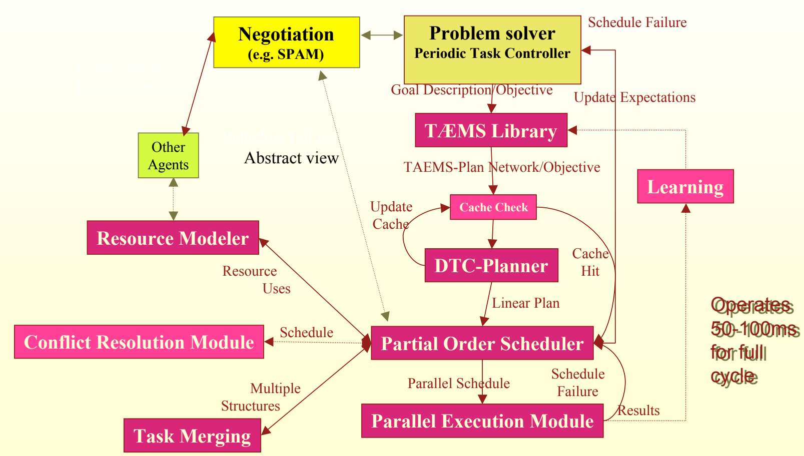

A Multi-agent Approach For Peer-to-Peer-based Information Retrieval Systems

An agent-view reorganization protocol to adapt the agent society topology to form loose content clusters in a distributed manner

A set of cooperative searching mechanisms, that take full advantage of the language model and the underlying topology, to fulfill information retrieval tasks

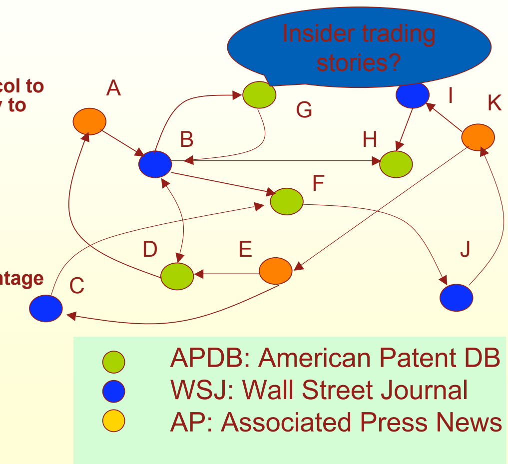

Prescriptive, knowledge-based design process