



# Lecture 1: Introduction

Victor R. Lesser

CMPSCI 683

Fall 2010

# 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
  - Email: [lesser@cs.umass.edu](mailto:lesser@cs.umass.edu), 413-545-1322
  - Office hours: Tuesdays, 1:15 - 2:45pm, CS306
- ◆ T.A.: Hala Mostafa
  - Email: [hmostafa@cs.umass.edu](mailto:hmostafa@cs.umass.edu)
  - Office hours: TBA
- ◆ Course Web Page:  
<http://mas.cs.umass.edu/classes/cs683/>
  - username: cs683      passwd: cs683protected

# Course Prerequisites

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

# Course Structure

- ◆ Text Book: “Artificial Intelligence: A Modern Approach,” (2nd or 3<sup>rd</sup> Edition) Stuart Russell and Peter Norvig, Prentice Hall, 2003.
  - Book’s website <http://aima.cs.berkeley.edu>
- ◆ 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.

# 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.”

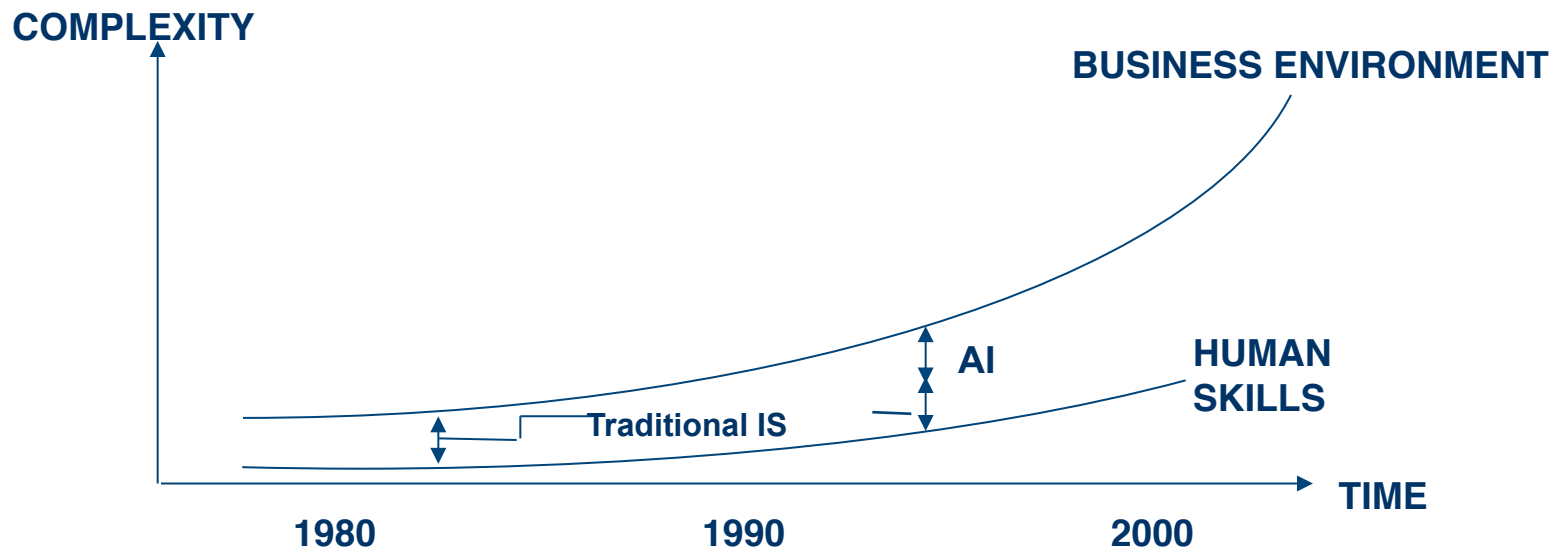
- Dartmouth Workshop, Summer of 1956

# 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 human behavior*
- ◆ *Making computers as easy to work with as people*

# Why is AI Important?

Complexity of Work Place (Rappaport, 1991)



IS = Information Systems

## Complexity of the Operating Environment

- Open vs Close Systems

# Why is AI Practical?

- ◆ Cheaper Sensors and Better Sensor Processing
  - Vision, speech, sound and text
- ◆ Orders of magnitude increase in computing power
  - Commodity processors (>billion instructions per second)
  - Special purpose processors for sensor processing
  - Parallel/Cloud/Distributed computing
- ◆ Availability of Tremendous Amount of Information
  - Large on-line knowledge bases
  - World Wide Web – ability to process unstructured information
  - Crowd Sourcing

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

# Why AI is Hard?

- uncertainty in the information that is sensed from the environment
- the situation specificity of knowledge
  - there is rarely common sense knowledge that can be applied uniformly over a broad spectrum of real-world situations
- the lack of complete theories that fully explain naturally occurring phenomena
- the bounded rationality of computation (NP hardness) that makes it impossible to fully assess all options.

*Ubiquity of Uncertainty*

# AI at 50 Years

'Artificial intelligence has accomplished more than people realize ... It permeates our economic infrastructure. Every time you place a cell phone call, send an e-mail, AI programs are directing information.' ... AI technology is used by banks to police transactions for fraud, by cell phone companies for voice recognition, and by search engines to scour the web and organize data. Beyond business, programs like Artificial Intelligence in Medicine help doctors diagnose and treat patients, while vision-recognition programs scan beaches and pools and alert lifeguards to signs of drowning. ... “ (futurist Ray Kurzweil. )

# 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.
  - Artificially intelligent robots drive autonomously over rugged terrain and long distances in the desert
  - 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.



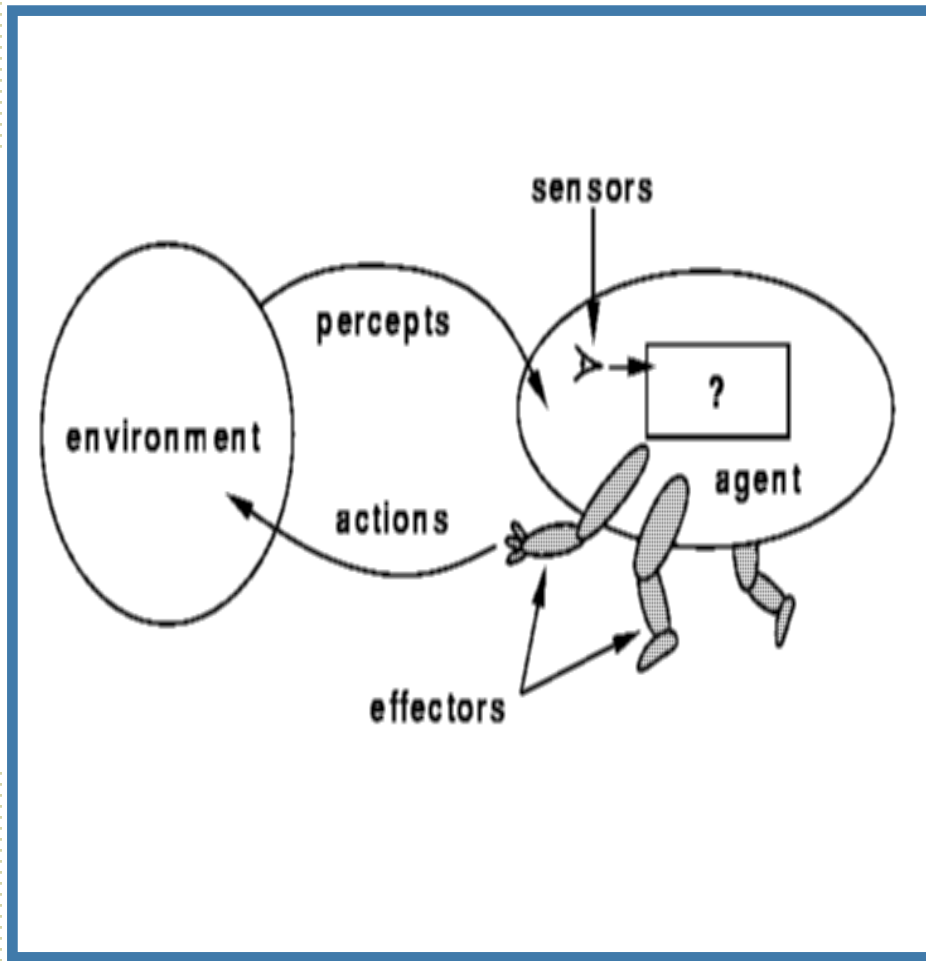
# What can AI systems do today? - 3

- ◆ **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.
- ◆ **Playing "Jeopardy"** – question answering

# Emerging Innovative Applications of AI

- *Intelligent Web Services*: Auto-categorization of Web-based pictures
- *Antiterrorism / Emergency Response*: Simulator for teaching police force allocation in urban environments; and a training tool to coordinate various disaster/emergency response entities (e.g., fire engines, police cars, etc.)
- *Health Care*: Autonomous long-term patient health monitoring using a network of lightweight sensors; and recognize user activity (e.g., falls, shock, etc.) from sensors
- *Computer Science*: Network security — risk assessment and monitoring; Autonomic self-reconfiguring computer systems;
- *Electric Power Industry*: Multiobjective Optimization of Power Distribution System Operations
- *Aerospace*: Deep Space Network Scheduling

# Agent-centric view of AI: AI operating in an Open Environment



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

# 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?

# CASA: Collaborative Adaptive Sensing


**CASA:** dense network of low power, *adaptive* radars: sense lower 3 km of earth's atmosphere

*collaborating radars:*

- improved sensing
- improved detection, prediction

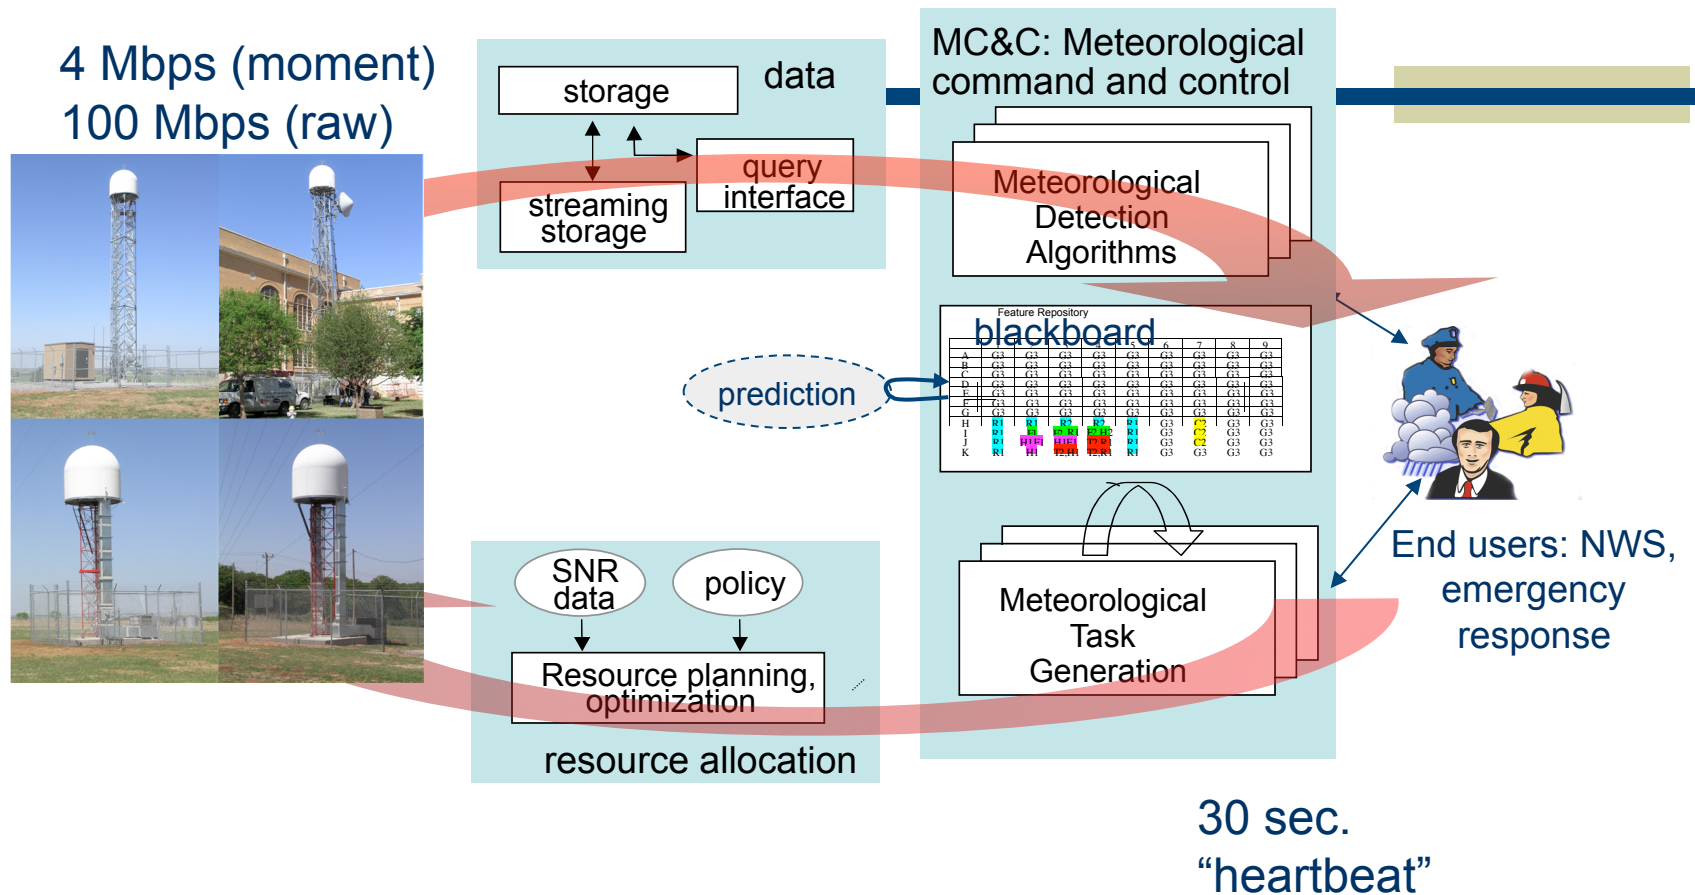
finer spatial resolution

responsive to multiple *end-user needs*



*“Sample atmosphere when and where end-user needs are greatest”*


# IP1 Oklahoma Testbed



- Soft real-time control loop
- Radar Scan Optimization Based on Evolving Multi-Time Step Goals
- Detection/Analysis/Prediction of Emerging Weather Phenomenon

# GENIE

- A project sponsored by NSF to investigate next-generation Internet infrastructure
- Exposes network-accessible APIs for consumers to lease virtualized resources
- Resources donated by universities and research labs
  - machines, block devices, sensors, mobile devices, network links



# Combinatorial Auctions

## Auctioneer's winner determination problem

- ◆ Set of items,  $M = \{1, 2, \dots, \#items\}$
- ◆ Set of bids,  $\beta = \{B_1, B_2, \dots, B_{\#bids}\}$
- ◆  $B_j = \langle S_j, p_j \rangle$ , 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:

$$\begin{aligned} \max \quad & \sum_{j=1}^{\#bids} p_j x_j \quad \text{s.t.} \quad \sum_{j|i \in S_j} x_j \leq 1 \quad i = 1, 2, \dots, \#items \\ & x_j \in \{0, 1\} \end{aligned}$$



# Aggregation of Information to Resolve Uncertainty

- ◆ Search – uncertainty about which action to take
  - Connect together a sequence of actions
- ◆ Uncertainty Reasoning – uncertainty about the correct interpretation of a piece of information
  - Connect together a set of interrelated uncertain information
- ◆ Learning – uncertainty about the implication of an observed event
  - Connect together observed events that are similar

# CONVENTIONAL SOFTWARE VS. AI

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

- Language: arithmetic vs. (logical, probabilistic)
- Data: numbers vs. symbols
- Coding: procedural (code) vs. declarative (sentences)
- Operations: calculations vs. reasoning (symbolic and decision theoretic)
- Knowledge: formulas vs. heuristics
- Program control: deterministic vs. nondeterministic (certain/well-defined vs. uncertain/search)
- Solution quality: exact/optimal vs. approximate/satisficing

# Dealing with uncertainty in both data and control

- Delayed resolution of uncertainty through non-deterministic and assumption-based computation
- Asynchronous and opportunistic application of knowledge and its associated constraints
- Exploiting approximate knowledge and heuristics to focus problem-solving activities
- Combining diverse and extensive sources of knowledge to resolve uncertainty
- Self-aware computation and meta-level reasoning to provide more context for decision making under uncertainty.


# Why AI is not Considered the Best Thing Since Slice Bread!!

"There's a joke in the AI community that as soon as AI works, it is no longer called AI,"


'Once a technology leaves the research labs and gets proven, it becomes ubiquitous to the point where it is almost invisible,'

# Agents

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



TROW



# 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

# 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?



# 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

# Syllabus Outline

- ◆ Introduction -1
  - ◆ Problem solving using sophisticated search – 9
  - ◆ Markov Decision Processes - 4
  - ◆ Reasoning under uncertainty - 6
  - ◆ Learning – 4
- } *Decision Making With  
Probabilistic Information*
- ◆ Won't be Covered
    - Adversarial Game Playing
    - Elementary Logical Formalism for Knowledge Representation
    - Architecture of Intelligent Systems
    - Knowledge Representation .....

# Detailed Syllabus -1

- ◆ Search – *Non-determinism of Knowledge*
  - Lecture 2 - Introduction to Search Strategies
  - Lecture 3 - Heuristic Search as represented by A\*
  - Lecture 4 – Space Variations in A\*
  - Lecture 5 - Time and Hierarchical Variations of A\*
  - Lecture 6 – Local Search -- Iterative Improvement & GSAT
  - Lecture 7 – Genetic Search & Solving CSPs using Heuristic Search
  - Lecture 8 - Speeding up CSP Algorithms
  - Lecture 9 - Blackboard Systems – Multi-level Search
  - Lecture 10 - Planning as Search

# Detailed Syllabus -2

## ◆ Markov Processes

- Lecture 11: Markov Decision Processes – *uncertain outcome of actions*
- Lecture 12: Partial Observable Markov Decision Processes -- *uncertain outcome of actions and state of system*
- Lecture 13: Decision Making As An Optimization Problem
- Lecture 14: Hidden Markov Models -- *uncertain about interpretation of a time sequence of observations*

# Detailed Syllabus -3

- ◆ Reasoning Under Uncertainty
  - Lecture 15: Uncertainty in Intelligent Systems
  - Lecture 16: Introduction to Probabilistic Reasoning with Belief Networks – *uncertain about how to interpret a set of observations*
  - Lecture 17: Approximate inference for BNs
  - Lecture 18: Decision Theory
  - Lecture 19: Simple Decision Networks (Trees)
  - Lecture 20: Decision Networks

# Detailed Syllabus -4

## ◆ Learning

- Lecture 21: Introduction to Learning
- Lecture 22: Decision Tree Learning
- Lecture 23: Neural Networks
- Lecture 24: Reinforcement Learning

May add lecture on Kernel Machines

# 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 frameworks 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 frameworks that allows you to live with it — rather than eliminate it
- ◆ Intelligent Control
  - Computational frameworks that allows you to effectively manage your resources to satisfy given goals in the face uncertainty
- ◆ Agency/Semi-Autonomous Agent
  - Computational framework that allows agents to interact autonomously with the world in terms of sensing, perceiving, planning, effecting and communicating, and adapt (learn) to long-term changes in the environment

# Review of Lecture 1

- ◆ Dealing directly with Uncertainty in Computation
  - one of the aspects of AI that differentiates it from other sub-disciplines of CS
  - Focus of Course – Mechanisms for Dealing with different types of uncertainty
    - Search – Non-Determinism of Knowledge
    - Decision Making with probabilistic information
    - Learning

# Next Lecture

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