Introduction to Artificial Intelligence (AI)

Computer Science cpsc502, Lecture 1

Sep, 8, 2011

People

Instructor

- **Giuseppe Carenini** (carenini@cs.ubc.ca; office CICSR 129)

Teaching Assistant

- **Shafiq Joty**   rjoty@cs.ubc.ca

This course is designed as a "breadth" introduction to AI.

It is suitable for those with

- no AI background ,
- or with only one undergraduate course in AI (or Machine Learning)

Today Sept 8

Overview of the field – Key definitions

Overview of course

- Background knowledge
- Topics
- Activities and Grading
- Administrative Stuff

Assignment-0

What is Intelligence?

problem solving

reasoning

decision

instance-of

classification

max utility

knowledge

learning

adaptive

language

planning

to achieve goals

Thinking and Acting Humanly

Model the cognitive functions of human beings

- Humans are our only example of intelligence: we should use that example!

Problems:

- But... humans often think/act in ways that we don't consider intelligent **(why?)**

emotions

cognitive limitations

memory

tired

incorrect

missing

knowledge

- And... detailed model of how people's minds operate not yet available

Thinking Rationally

Rationality: an **abstract "ideal" of intelligence**, rather than ``whatever humans think/do''

- Ancient Greeks invented *syllogisms*: argument structures that always yield correct conclusions given correct premises
  - This led to **logic**, and **probabilistic reasoning** which we'll discuss in this course
- But correct sound reasoning is not always enough "to survive" "to be useful"…

Acting (&thinking) Rationally

This course will emphasize a view of AI as building **agents**: artifacts that are able to think and act rationally in their environments

Rationality is **more cleanly defined** than human behavior, so it's a better design objective

(Eg: "intelligent" vacuum cleaner: maximize area cleaned, minimize noise and electricity consumption)

Agents that can answer queries, plan actions and solve complex problems

And when you have a rational agent you can always tweak it to make it irrational!

Why do we need intelligent agents?

learn

planning

Help people

decision making

work more effectively

Doctor

Robotics

driving car

space exploration

Autonomous

dangerous/boring tasks

(Adversarial) Search: Checkers

Game playing was one of the first tasks undertaken in AI

Arthur Samuel at IBM wrote programs to play checkers (1950s)

- initially, they played at a strong amateur level
- however, they used some (simple) machine learning techniques, and soon outperformed Samuel

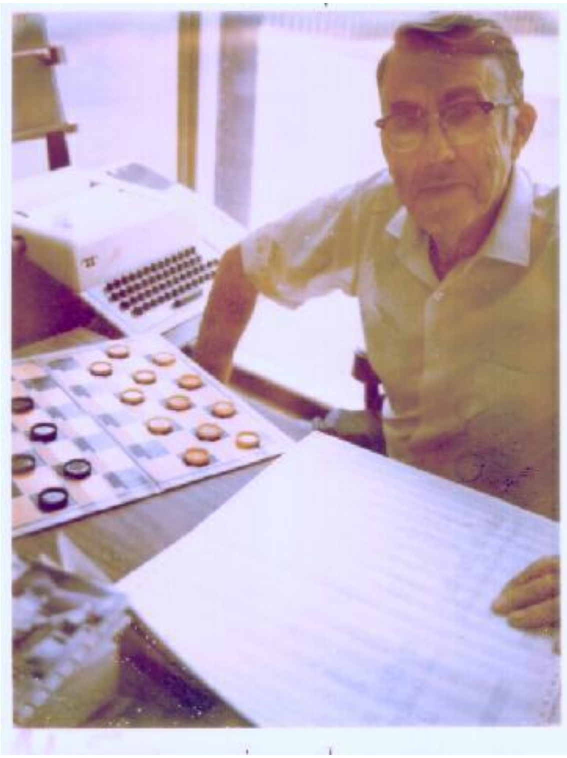

Source: *IBM Research*

Chinook's program was declared the Man-Machine World Champion in checkers in 1994!

…and **completely solved** by a program in 2007!

(Adversarial) Search: Chess

In 1996 and 1997, Gary Kasparov, the world chess grandmaster played two tournaments against Deep Blue, a program written by researchers at IBM

Source: *IBM Research*

(Adversarial) Search: Chess

Deep Blue’s Results in the first tournament:

- won 1 game, lost 3 and tied 1
  - ✓ first time a reigning world champion lost to a computer

Source: *CNN*

(Adversarial) Search: Chess

Deep Blue's Results in the second tournament:

- second tournament: won 3 games, lost 2, tied 1

- 30 CPUs + 480 chess processors
- Searched 126.000.000 nodes per sec
- Generated 30 billion positions per move reaching depth 14 routinely

CSPs: Crossword Puzzles

Daily Puzzles

370 puzzles from 7 sources.

Summary statistics:

- 95.3% words correct (miss three or four words per puzzle)
- 98.1% letters correct
- 46.2% puzzles completely correct

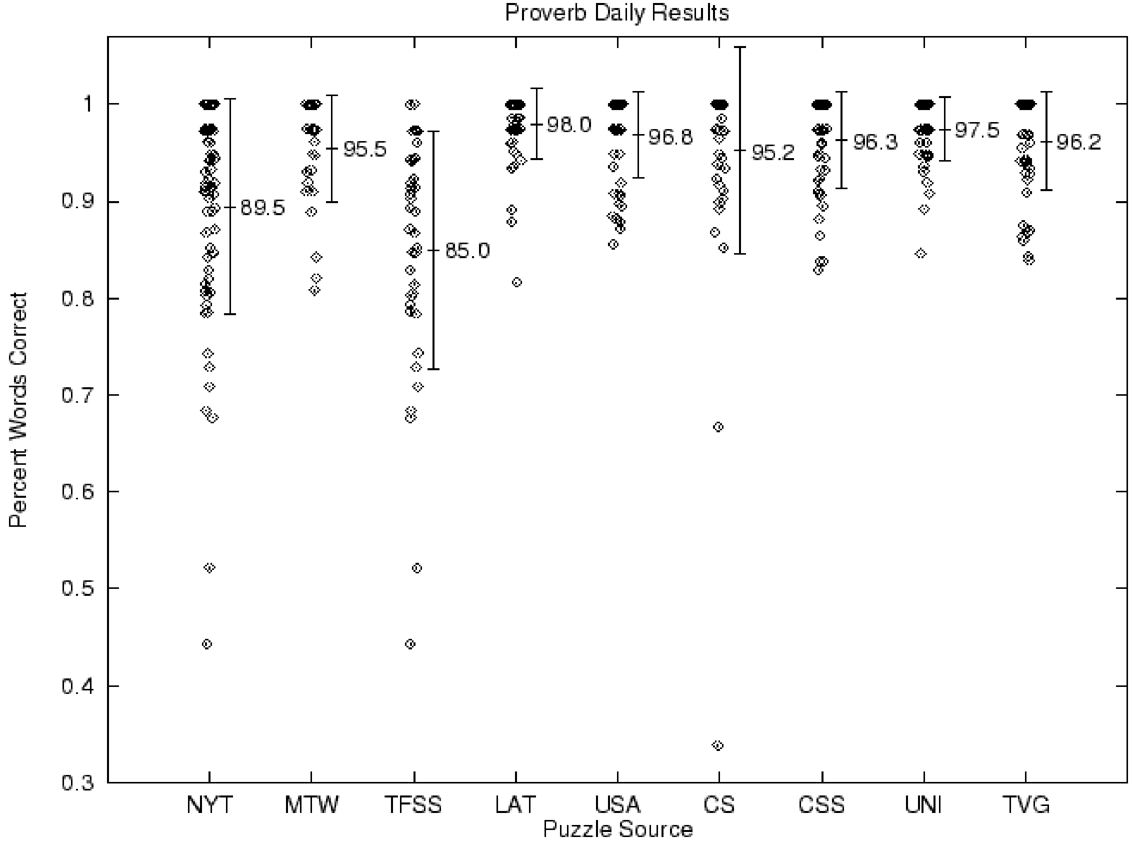

Source: *Michael Littman*

CSPs: Radio link frequency assignment

Assigning frequencies to a set of radio links defined between pairs of sites in order **to avoid interferences**.

Constraints on frequency depend on **position of the links** and on **physical environment** .

Source: *INRIA*

Sample Constraint network

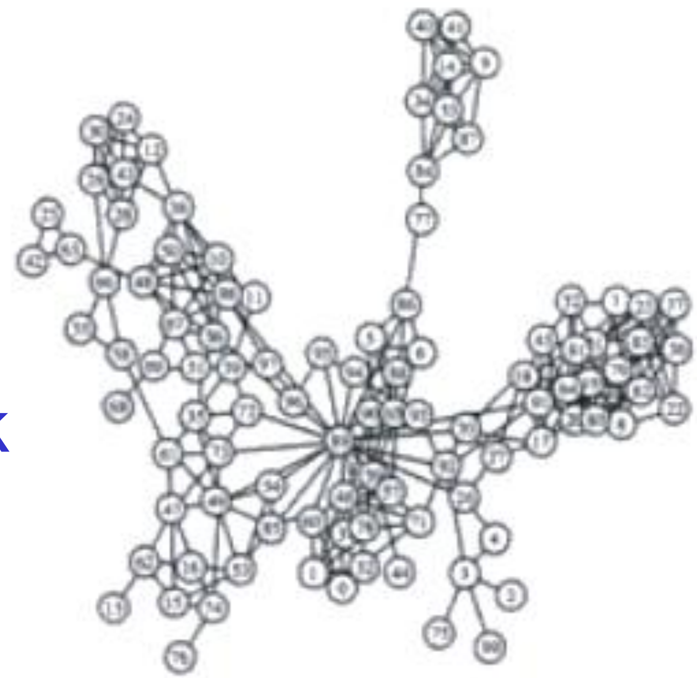

Example: SLS for RNA secondary structure design

RNA strand made up of four bases: cytosine (C), guanine (G), adenine (A), and uracil (U)

2D/3D structure RNA strand folds into is important for its function

Predicting structure for a strand is “easy”: $O(n^3)$

But what if we want a strand that folds into a certain structure?

- Local search over strands
  - ✓ Search for one that folds into the right structure
- Evaluation function for a strand
  - ✓ Run $O(n^3)$ prediction algorithm
  - ✓ Evaluate how different the result is from our target structure
  - ✓ Only defined implicitly, but can be evaluated by running the prediction algorithm

RNA strand

GUCCCAUAGGAUGUCCCAUAGGA

Easy ↓ ↑ Hard

Secondary structure

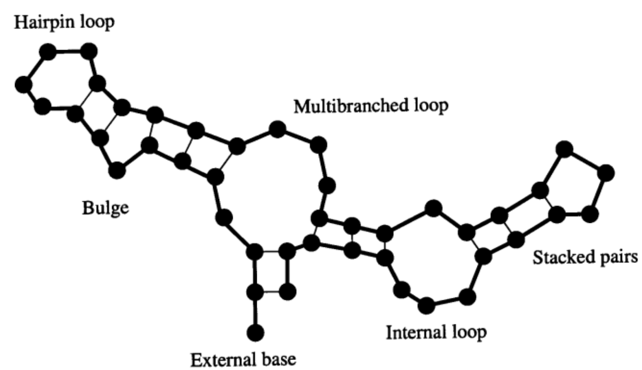

Best algorithm to date: Local search algorithm RNA-SSD developed at UBC [Andronescu, Fejes, Hutter, Condon, and Hoos, Journal of Molecular Biology, 2004]

Constraint optimization problems

Optimization under side constraints (similar to CSP)

E.g. mixed integer programming (software: IBM CPLEX)

- Linear program: max $c^T x$ such that $Ax \leq b$
- Mixed integer program: additional constraints, $x_i \in \mathbb{Z}$ (integers)
- NP-hard, widely used in operations research and in industry

Transportation/Logistics:
SNCF, United Airlines
UPS, United States
Postal Service, …

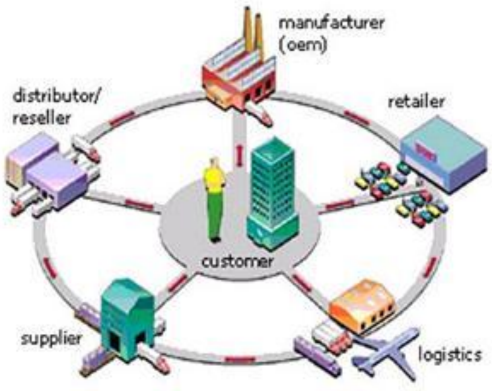

Supply chain
management
software:
Oracle,
SAP,…

Production planning
and optimization:
Airbus, Dell, Porsche,
Thyssen Krupp,
Toyota, Nissan, ...

Planning & Scheduling: Logistics

Dynamic Analysis and Replanning Tool (Cross & Walker)

- logistics planning and scheduling for military transport
- used in the 1991 Gulf War by the US
- problems had 50,000 entities (e.g., vehicles); different starting points and destinations

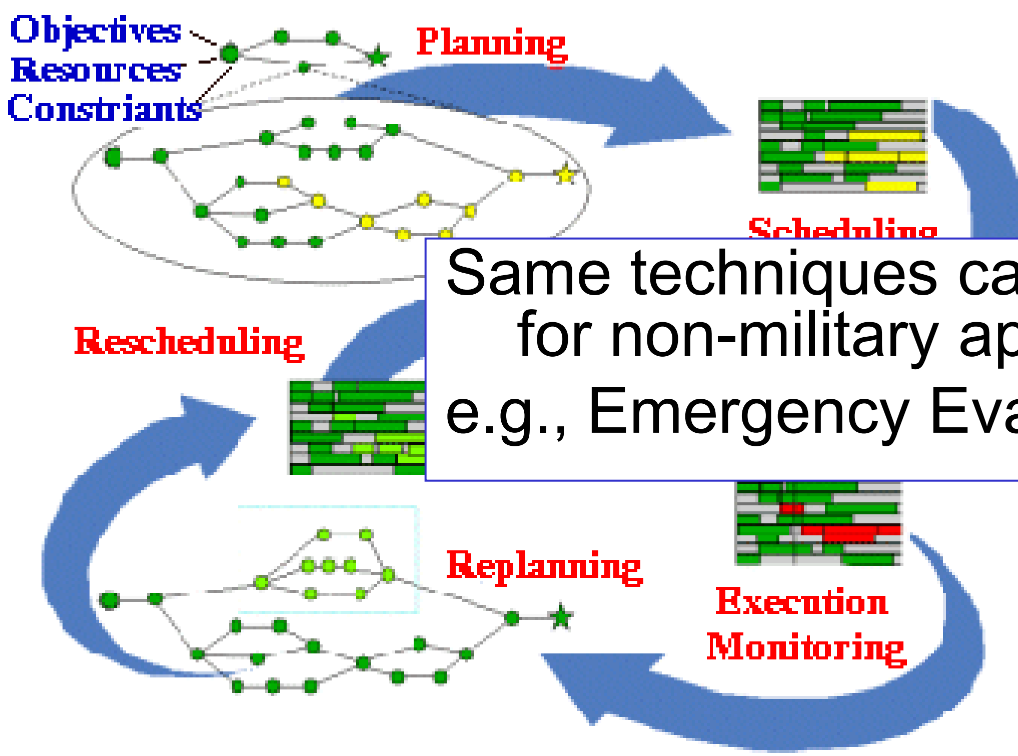

Same techniques can be used for non-military applications:
e.g., Emergency Evacuation

Source: *DARPA*

CSP/logic: formal verification

Hardware verification
(e.g., IBM)

Software verification
(small to medium programs)

Most progress in the last 10 years based on:
Encodings into propositional satisfiability (SAT)

Logic: CycSecure

"**scans a computer network** to **build a formal representation** of the network, based on Cyc's pre-existing ontology of networking, security, and computing concepts:

Excerpted from: *Shepard et al., 2005*

This formal representation also allows users to interact directly with the model of the network, allowing testing of proposed changes."

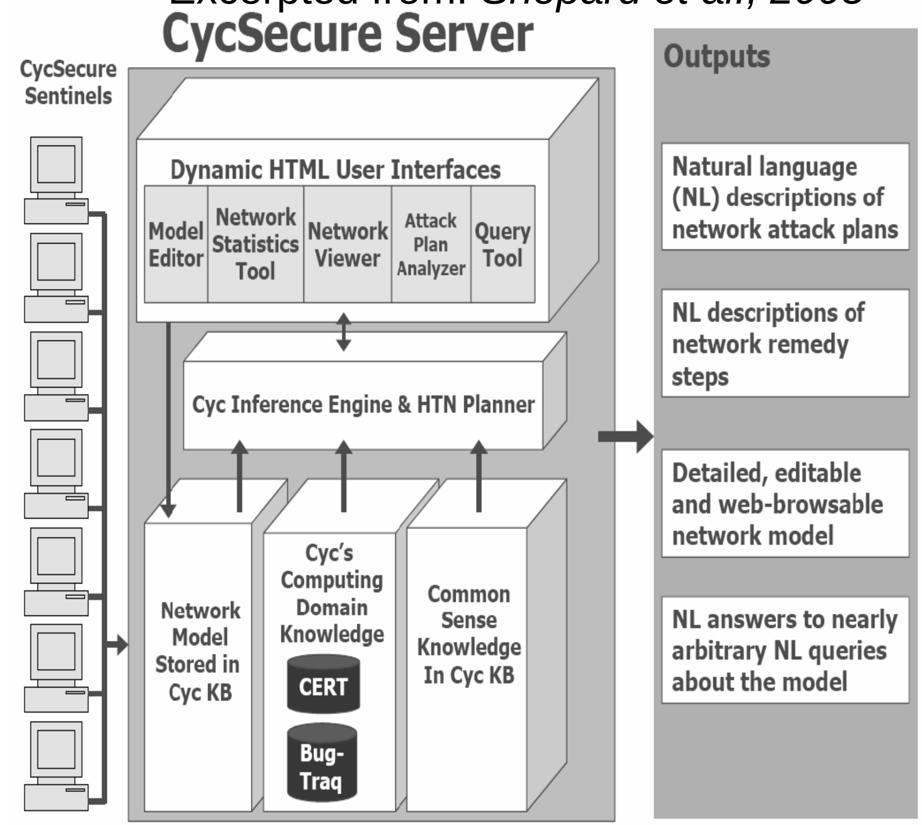

- Knowledge Representation
- Semantic Web !

Planning: Spacecraft Control

NASA: Deep Space One spacecraft
operated autonomously for two days in May, 1999:

- determined its precise position using stars and asteriods
  - ✓ despite a malfunctioning ultraviolet detector
- planned the necessary course adjustment
- fired the ion propulsion system to make this adjustment

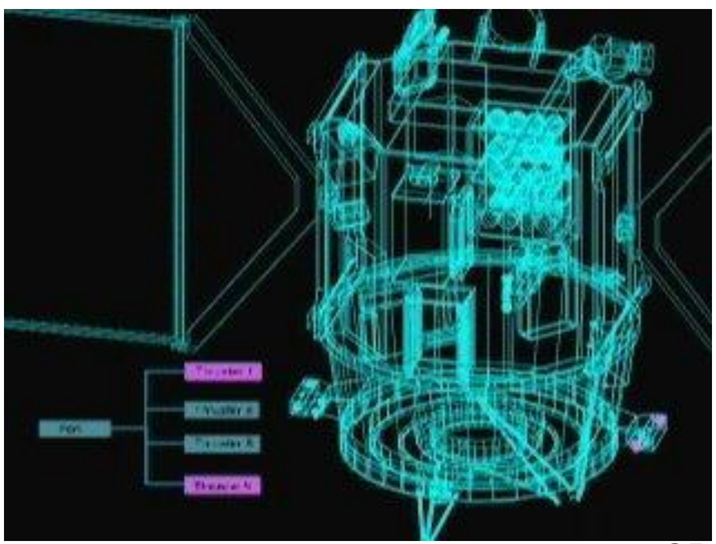

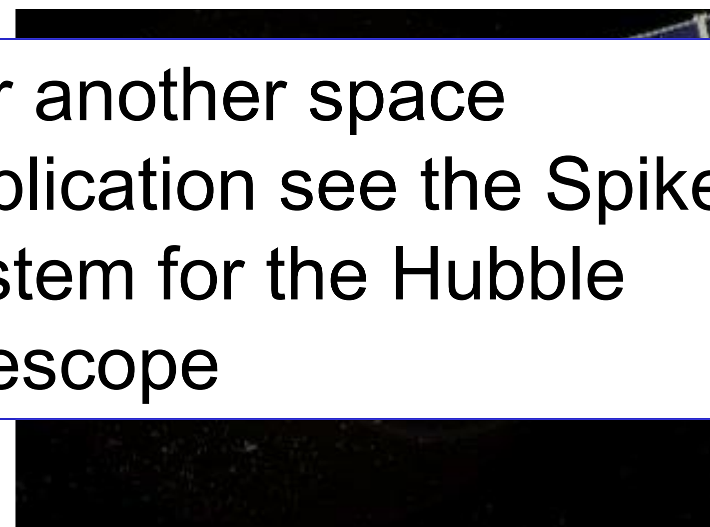

For another space application see the Spike system for the Hubble telescope

Source: *NASA*

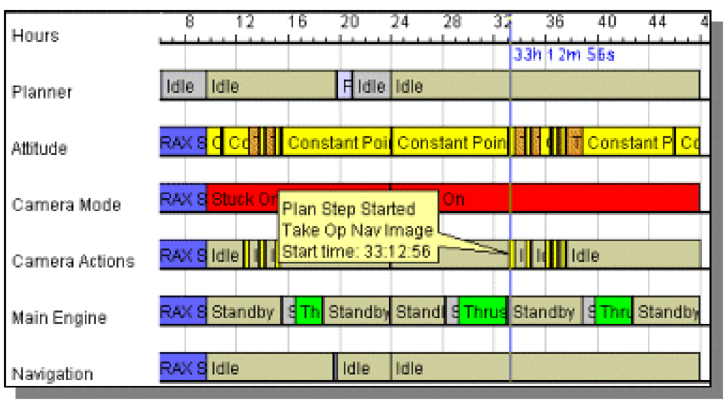

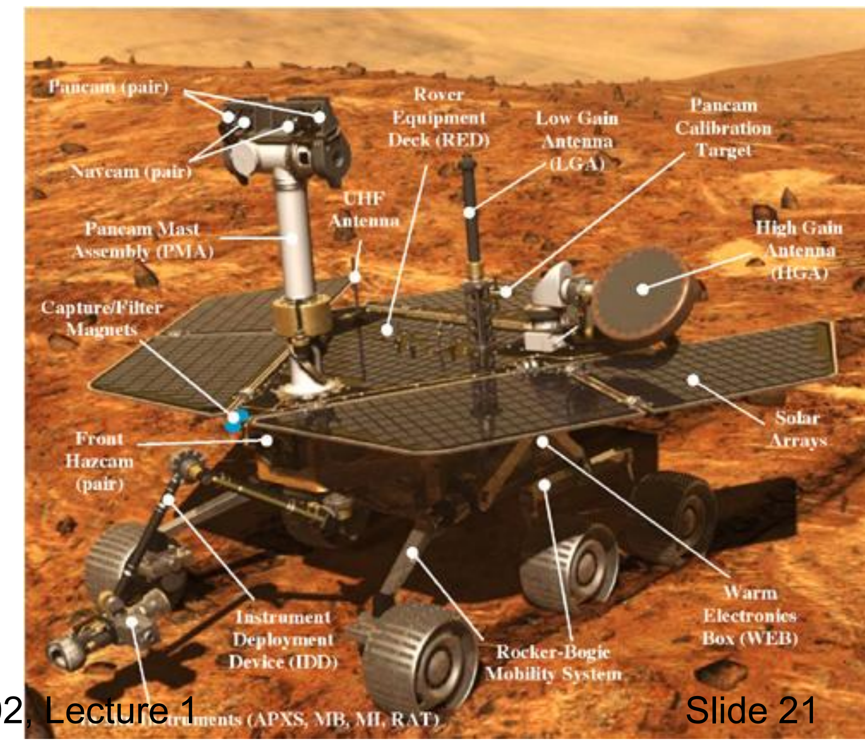

Source:
cs221 stanford

Reasoning under Uncertainty: Diagnosis

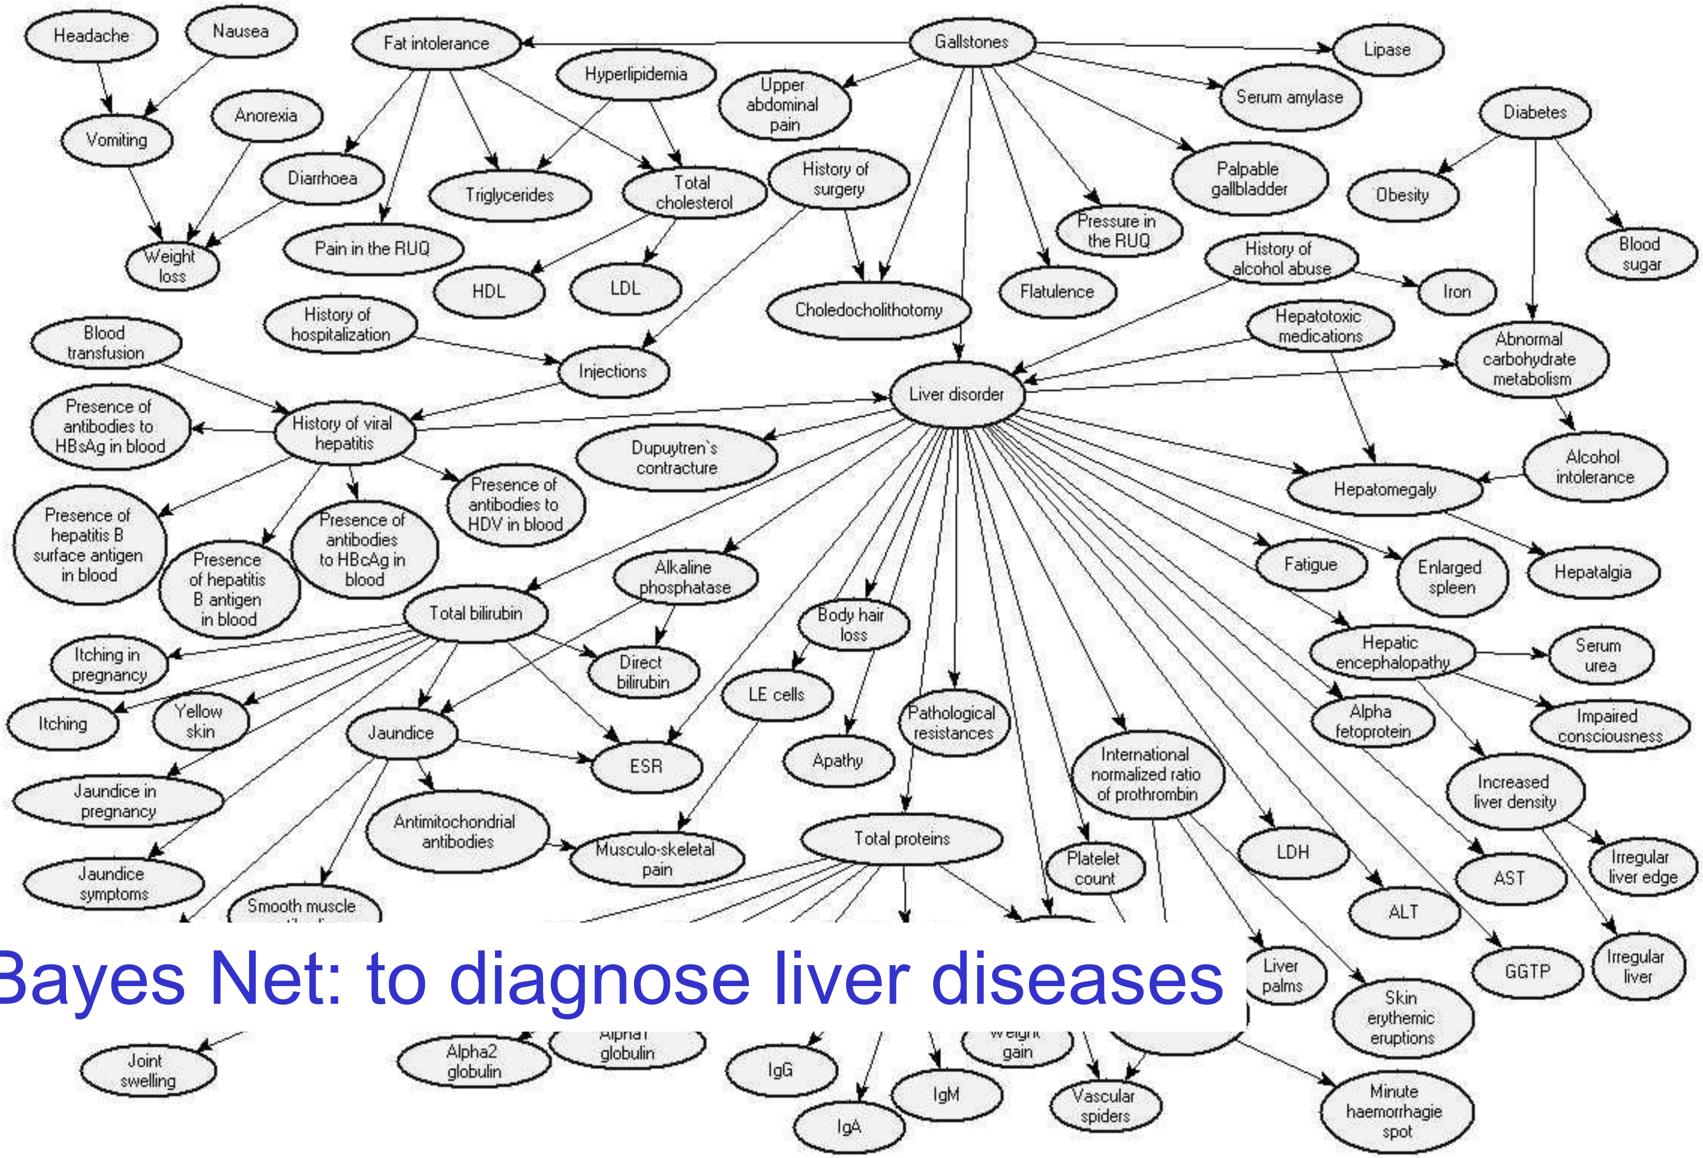

Bayes Net: to diagnose liver diseases

Reasoning Under Uncertainty

Texture classification using Support Vector Machines

- foliage, building, sky, water

Source: *Mike Cora, UBC*

Reasoning Under Uncertainty

E.g. motion tracking: track a hand and estimate activity:

- drawing, erasing/shading, other

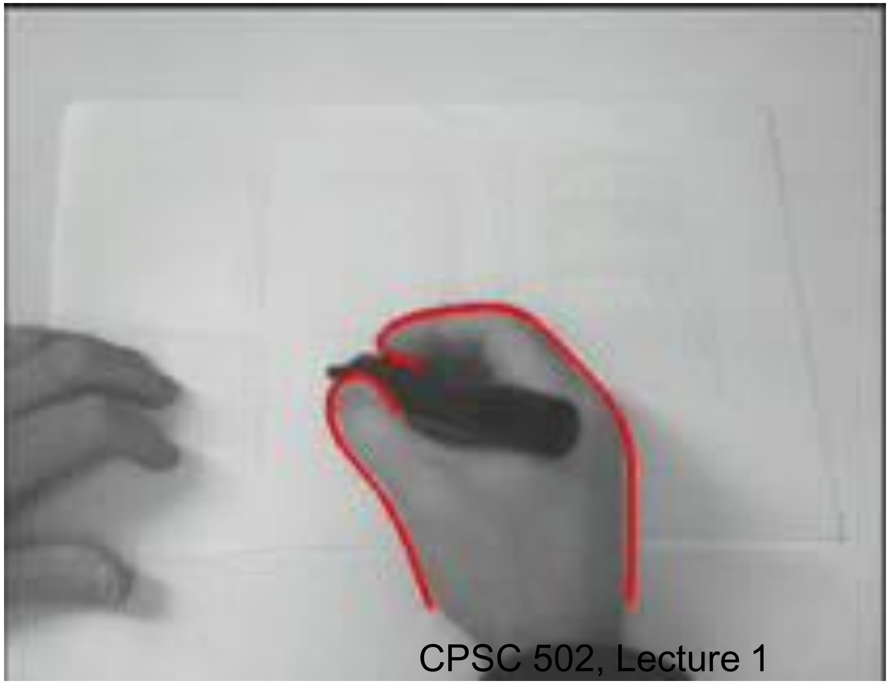

Source: *Kevin Murphy, UBC*

Computer Vision (not just for robots!)

Jing, Baluja, Rowley, Google: Finding Canonical Images

Source:
cs221 stanford

Compare low level features

(a) A v.s. B

(b) A v.s. C

(c) A v.s. D

(d) B v.s. C

Source:
cs221 stanford

Induced Graph

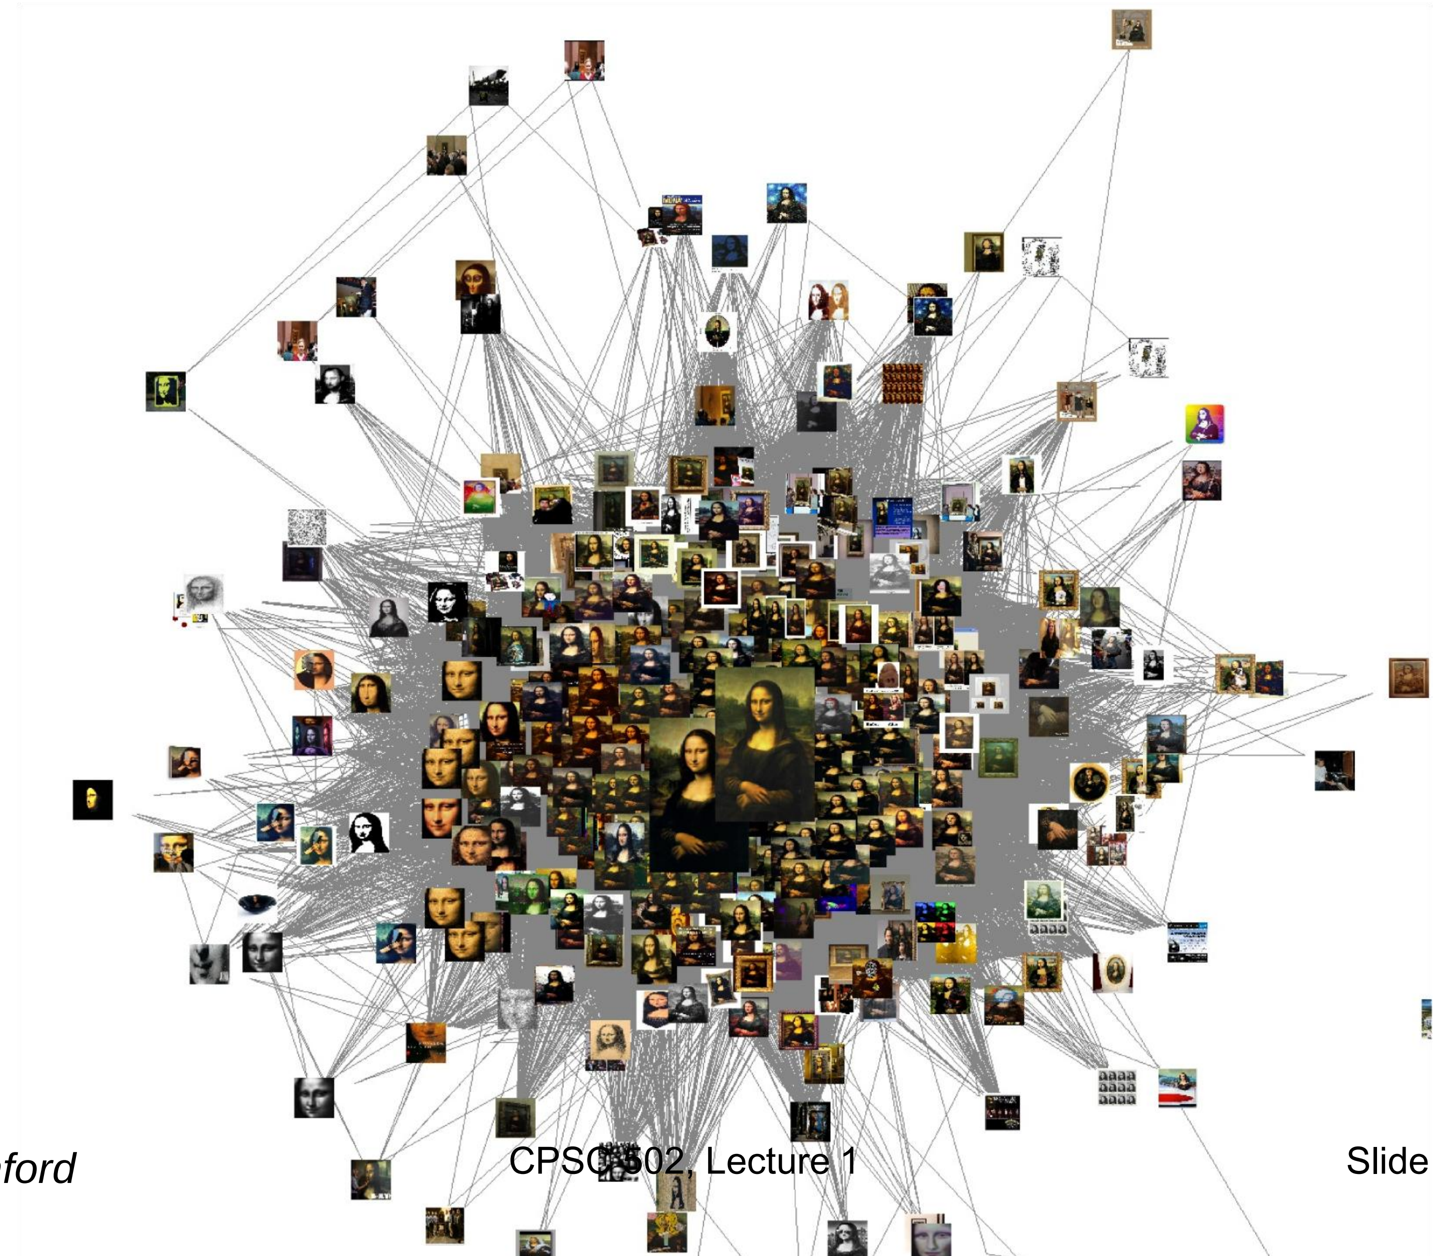

Source:
cs221 stanford

AI - Machine Learning @google

- Spam/Porn Detection
- Which ad to place given a query
- Train Speech to search on mobile
- Machine Translation
- …..

- Highly Parallelizable EM + Map Reduce (simple code to write)
- Stochastic Gradient Descent

JEOPARDY! The IBM Challenge

Watson : analyzes natural language questions and content well enough and fast enough to compete and win against champion players at Jeopardy!

"This Drug has been shown to relieve the symptoms of ADD with relatively few side effects."

- 3 secs

- 1000s of algorithms and KBs,

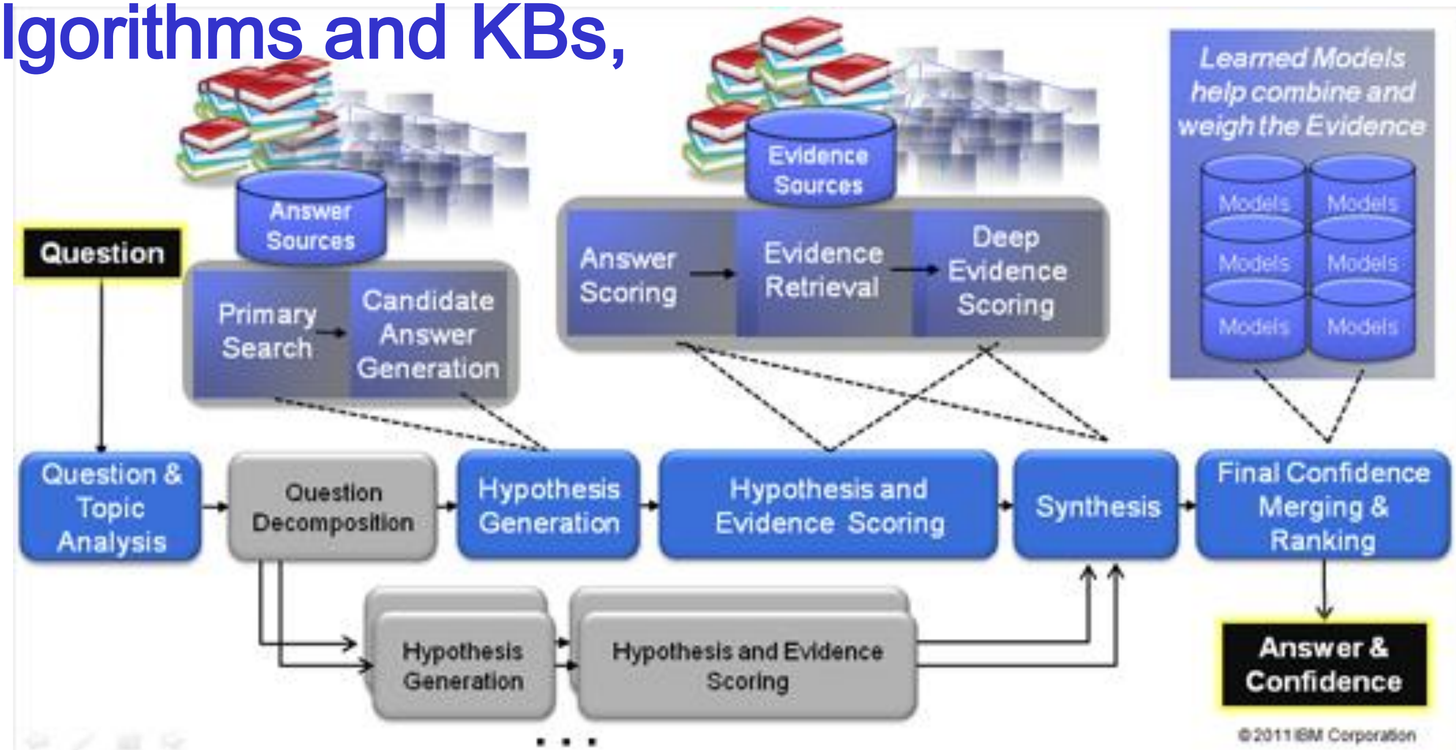

Source: *IBM*

Planning Under Uncertainty

Learning and Using POMDP models of Patient-Caregiver Interactions During Activities of Daily Living

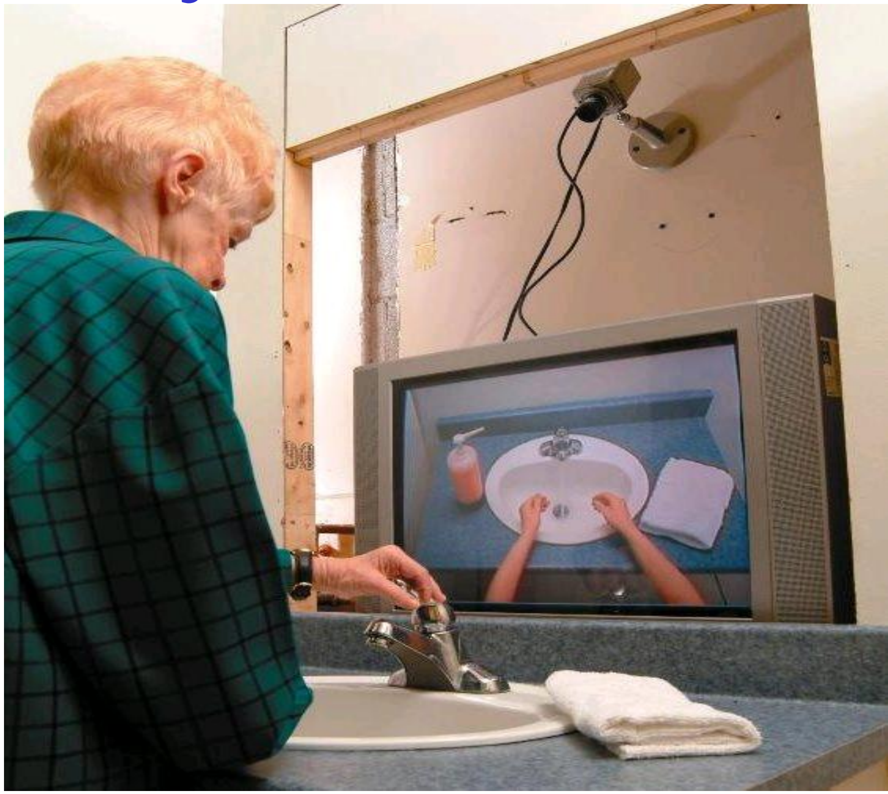

Goal: Help Older adults living with cognitive disabilities (such as Alzheimer's) when they:

- **forget the proper sequence of tasks** that need to be completed
- **they lose track of the steps** that they have already completed.

Source: *Jesse Hoey UofT 2007*

Military applications: ethical issues

- Robot soldiers
  - Existing: robot dog carrying heavy
    materials for soldiers in the fie
  - The technology is there
- Unmanned airplanes
- Missile tracking
- Surveillance
- …

Planning Under Uncertainty

Helicopter control: MDP, reinforcement learning

States: all possible positions, orientations, velocities and angular velocities

Source: *Andrew Ng 2004*

Final solution involves

Deterministic **search**!

Decision Theory: Decision Support Systems

E.g., Computational Sustainability

New interdisciplinary field, AI is a key component

- Models and methods for decision making concerning the management and allocation of resources
- to solve most challenging problems related to sustainability

Often constraint optimization problems. E.g.

- Energy: when are where to produce green energy most economically?
- Which parcels of land to purchase to protect endangered species?
- Urban planning: how to use budget for best development in 30 years?

Source: http://www.computational-sustainability.org/

Multiagent Systems: Poker

Search Space: 1.2 quintillion nodes

"In full 10-player games Poki is better than a typical low-limit casino player and wins consistently; however, not as good as most experts

New programs being developed for the 2-player game are quite a bit better, and we believe they will very soon surpass all human players"

Source: *The University of Alberta GAMES Group*

Multiagent Systems: Robot Soccer

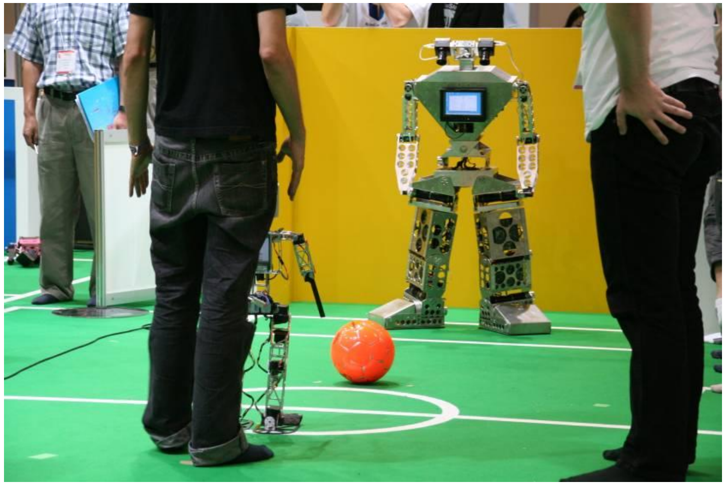

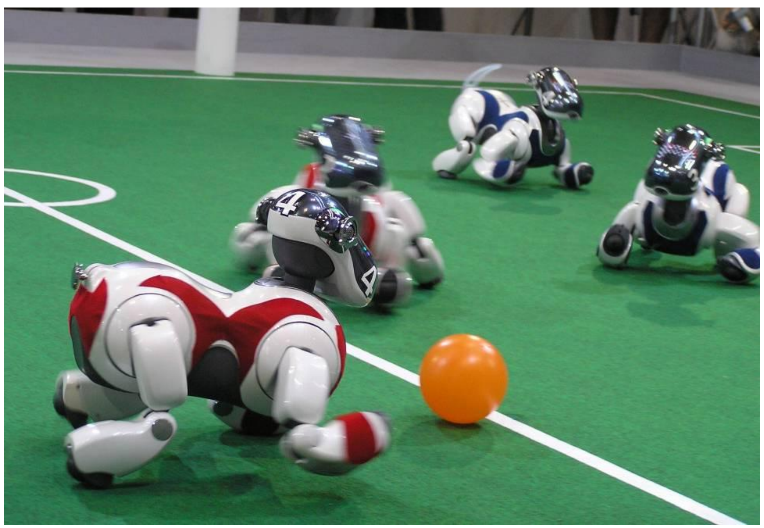

Extremely complex

- Stochastic
- Sequence of actions
- Multiagent

robotic soccer competition was proposed by LCI (UBC) in 1992 (which became *Robocup* in 1997).

Source: *RoboCup web site*

Statistical Machine Translation

SEHR GEEHRTER GAST!
KUNST, KULTUR UND
KOMFORT IM HERZEN
BERLIN.

↔

DEAR GUESTS,
ART, CULTURE AND
LUXURY IN THE HEART
OF BERLIN.

DIE ÖRTLICHE
NETZSPANNUNG
BETRÄGT 220/240 VOLT
BEI 50 HERTZ.

↔

THE LOCAL VOLTAGE
IS 220/240 VOLTS 50 HZ.

EN

DE

Source:
cs221 stanford

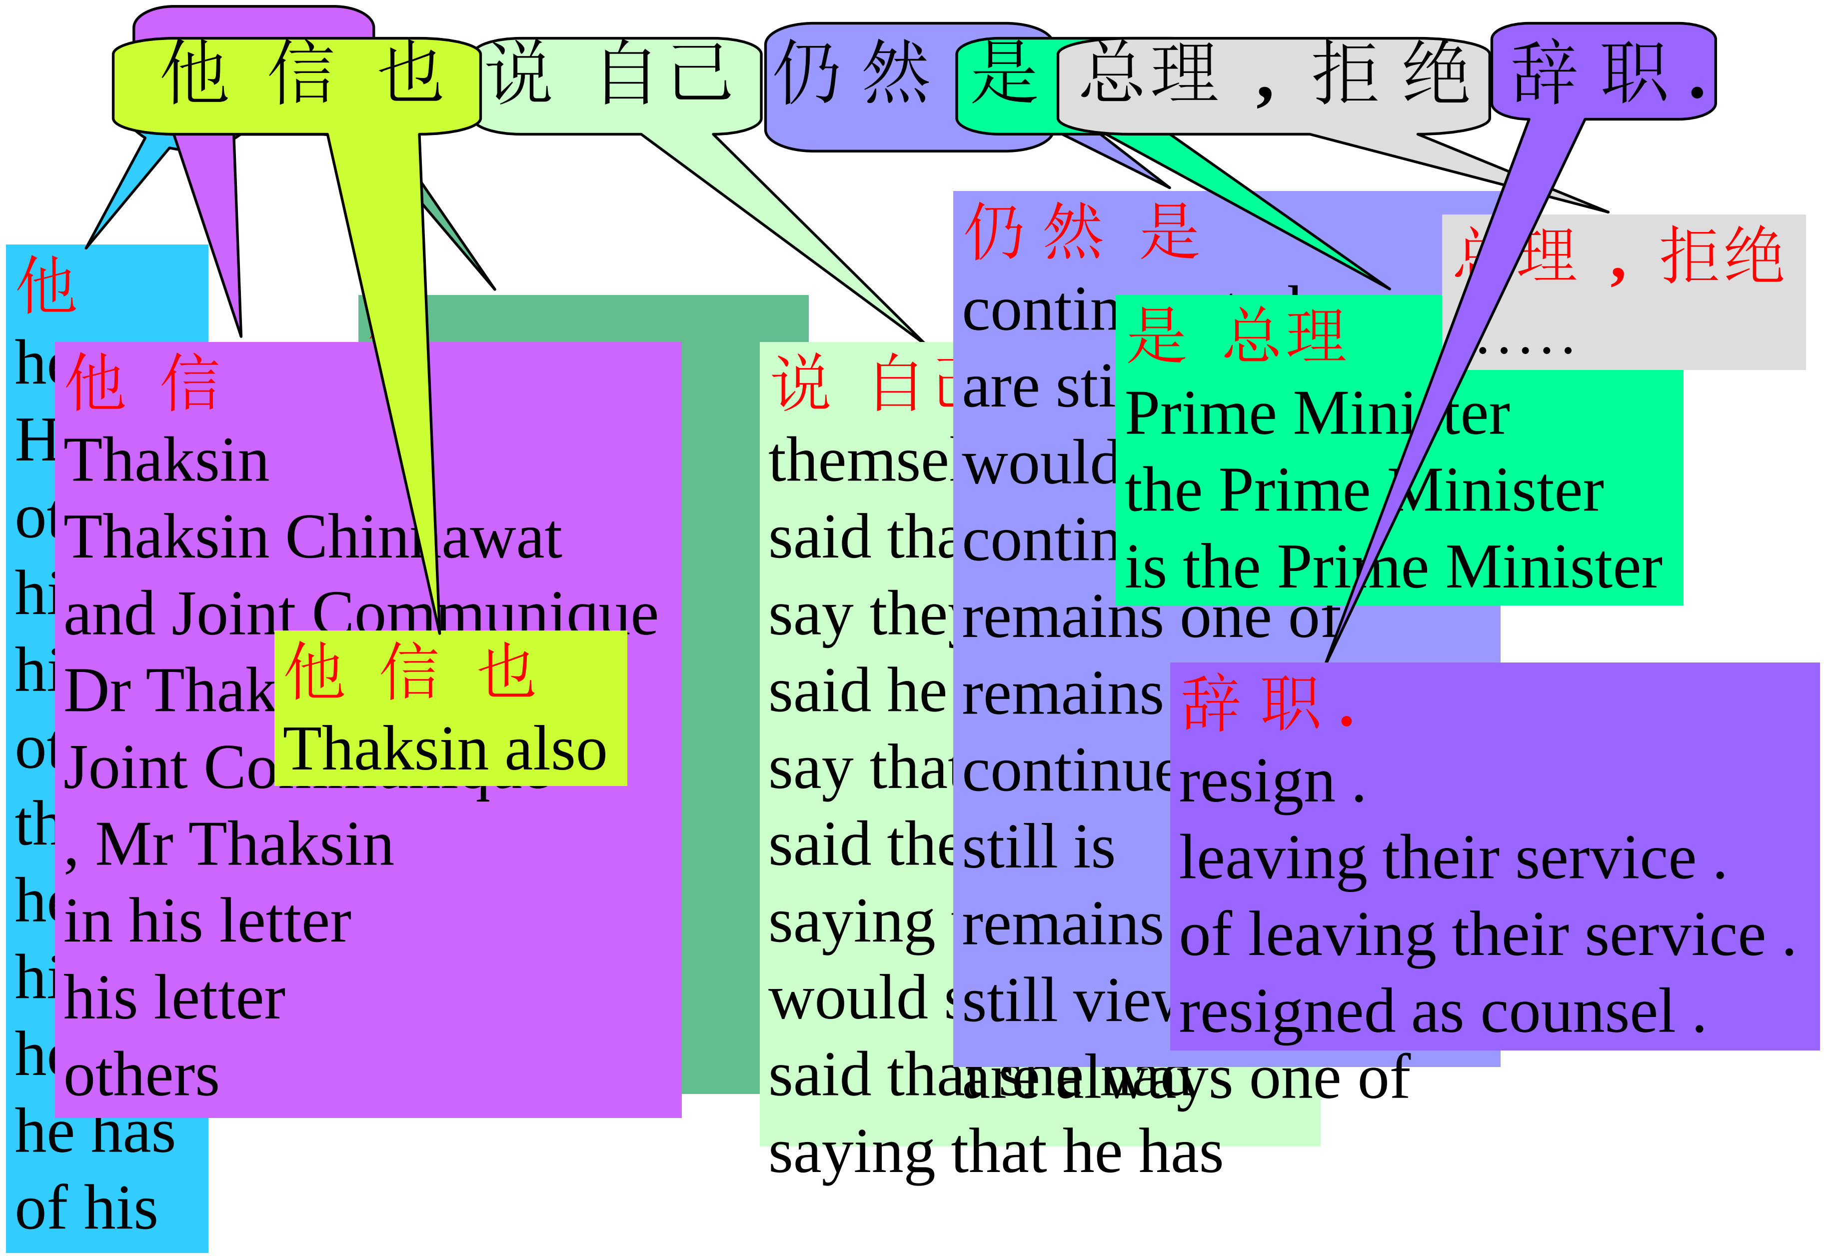

Source:
cs221 stanford

Zite: a personalized magazine

… that gets smarter as you use it

AI magazine
YAHOO! RESEARCH
AI Music Composition & Performance
AI Paintings
Robot Soccer
Humanoid Robot
How Can AI Systems Learn From Experience?
Social Simulation Game
How Can AI Systems Have Common Sense?
Autonomous Vehicle
Robot Tour Guide
Characters for Virtual Worlds
Smart Environmental Controls
Intelligent Tutoring System
Vehicle Navigation System
Smart Desk with Gesture Recognition
Autonomous Space Exploration
Personalized TV Guide
How Can AI Systems Solve Problems Creatively?
Robots for Education
Robot Vacuum Cleaner
Recommender System
How Can AI Systems Cooperate Effectively With Humans?
Robotic Surgery
Disease Diagnosis
Aristotle
Drug Design
AI LAB
Fraud Detection
Leibniz
Descartes
Whitehead
Lovelace
Russell
Turing
Smart Wheelchair
Web Search
BOOKS & COFFEE
BOOKS
Handwriting & Sketch Recognition
Spam Filtering
Machine Translation
1956
See the AI timeline and more at www.aaai.org/AILandscape
See the AI timeline and more at
www.aaai.org/AILandscape
The AI Landscape
David Leake, Indiana University, Poster Development Committee Chair
Poster Design: Giacomo Marchesi www.GiacomoMarchesi.com

Agents acting in an environment

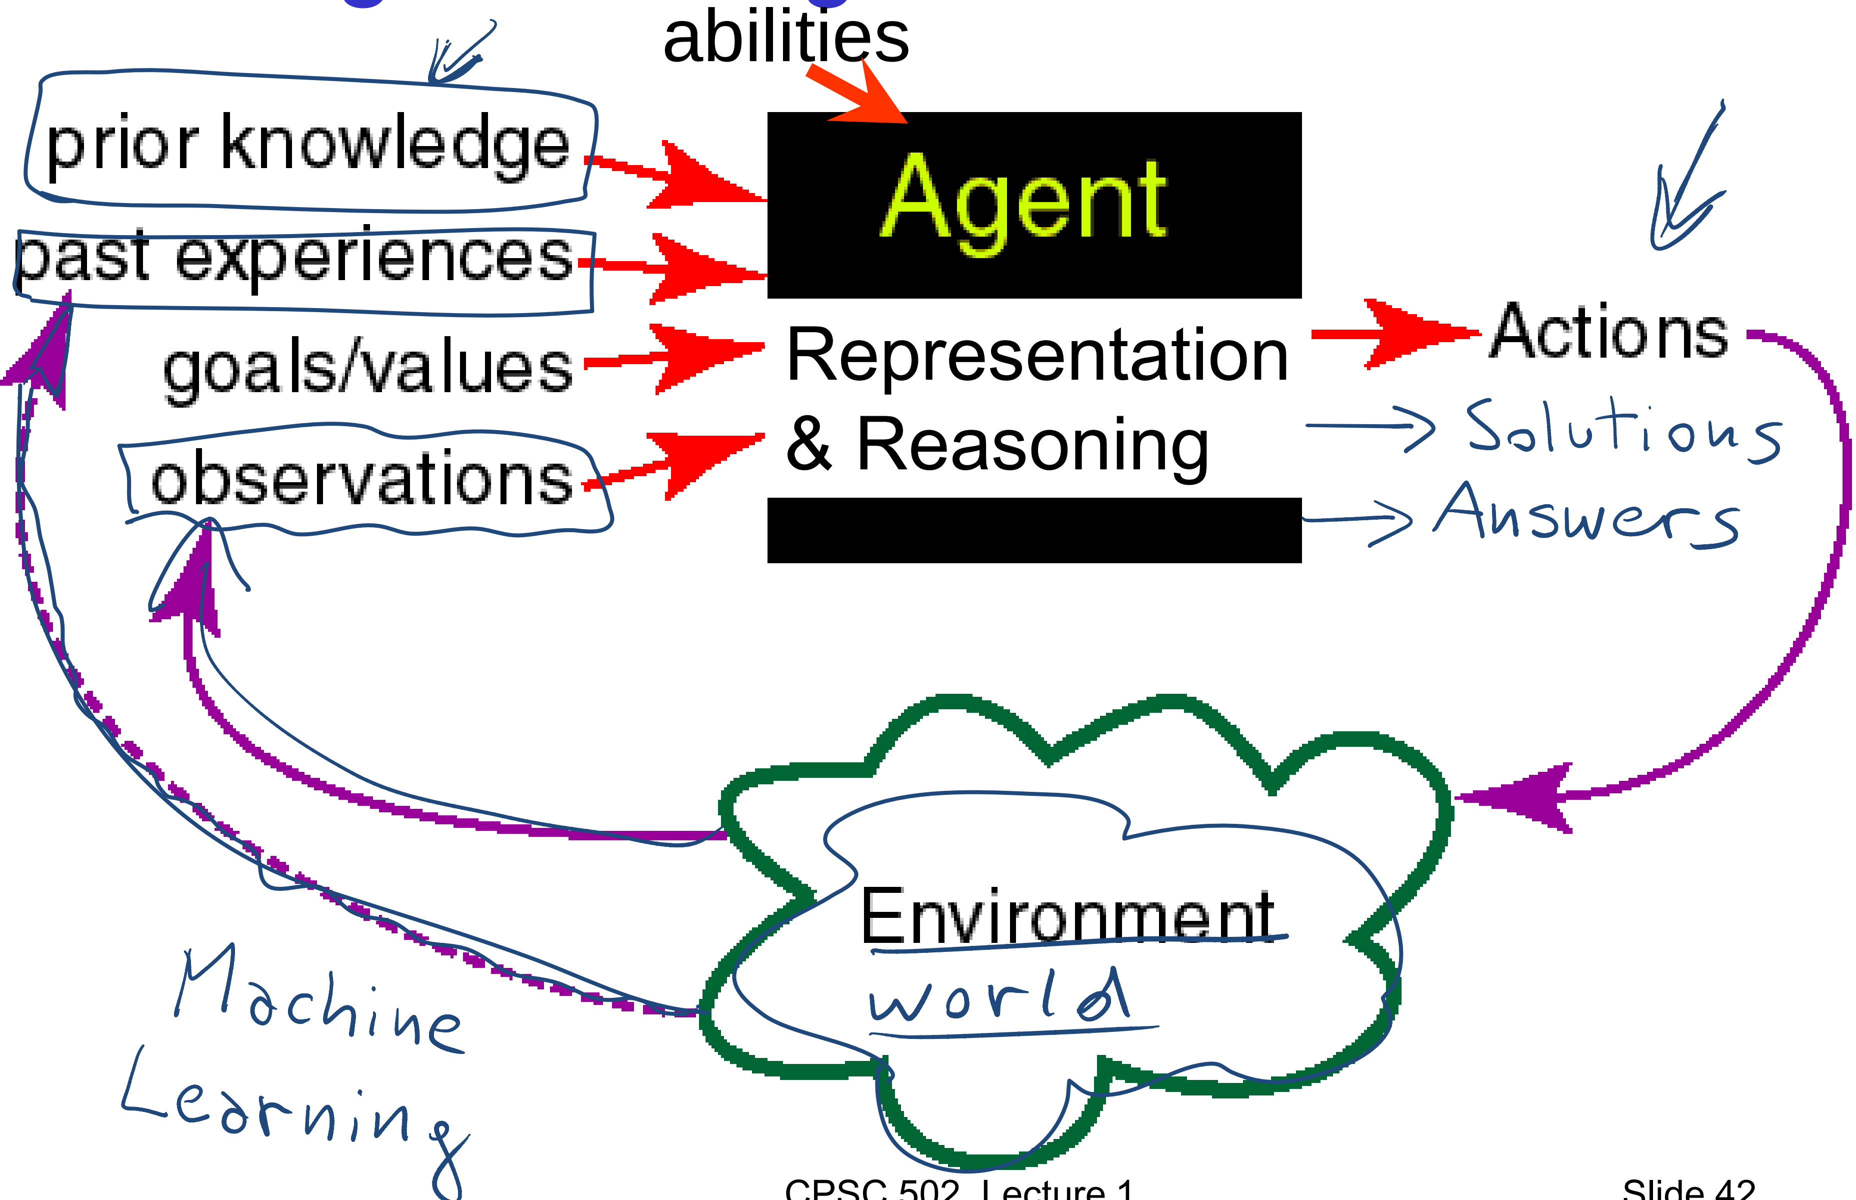

What is an agent?

It has the following characteristics:

- It is situated in some **environment**
  - does not have to be the real world---can be an abstracted electronic environment

Medical test / Eyetracking

- It can make **observations** *(perhaps imperfectly)*
- It is able to **act** *(provide an answer, buy a ticket)*
- It has **goals or preferences** *(possibly of its user)*

realestate advisor

- It may have **prior knowledge or beliefs**, and some way of **updating beliefs** based on new experiences (to reason, to make inferences)

Intelligent Agents in the World

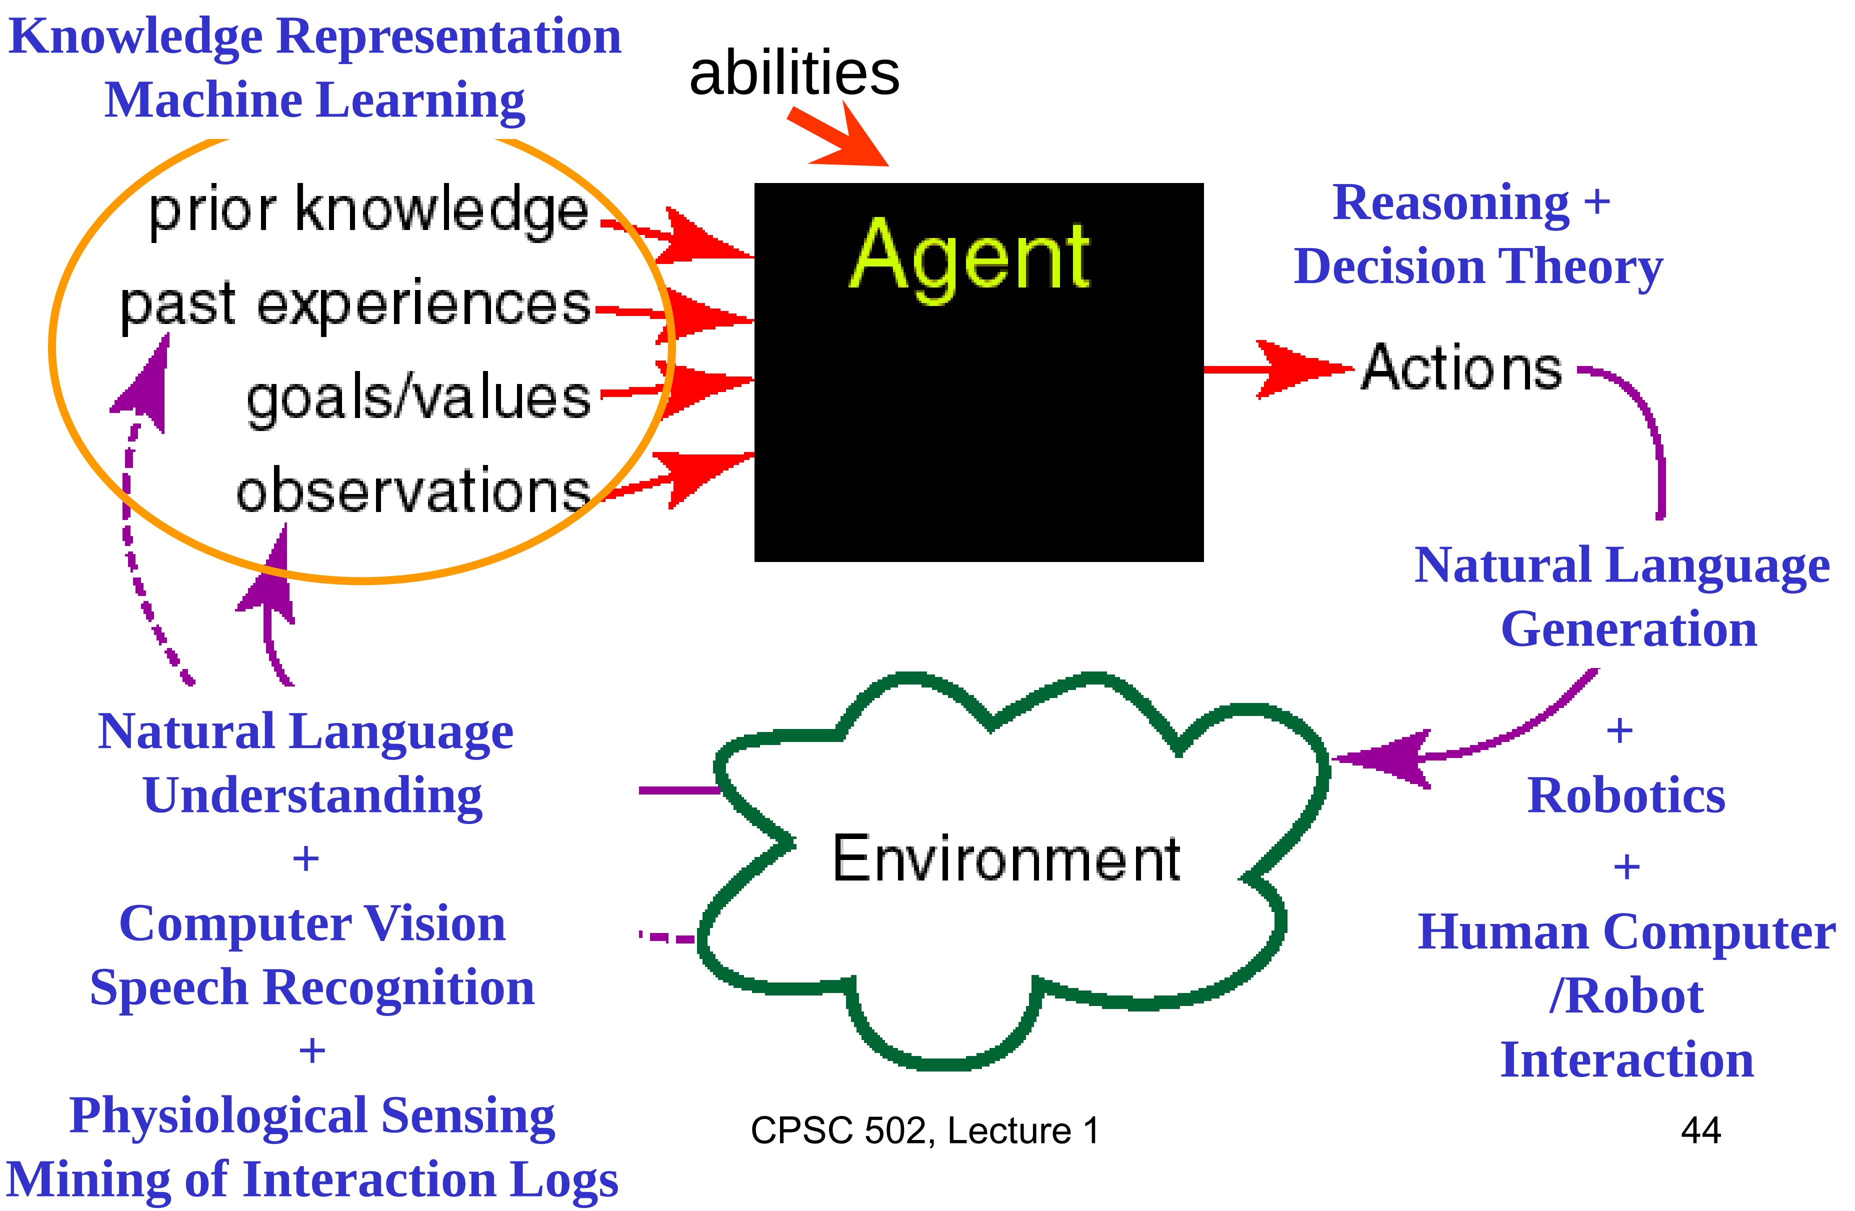

Today Sept 9

Overview of the field – Key definitions

Overview of course

- Background knowledge
- **Topics**
- Activities and Grading
- Administrative Stuff

Assignment 0

What do we need to represent ?

- The environment /world : What different configurations (**states / possible worlds**) can the world be in, and how do we denote them?

Chessboard, Info about a patient, Robot Location

- How the world works ***(we will focus on)***
  - **Constraints:** *sum of current into a node = 0* (electric circuit)
  - **Causal:** *what are the causes and the effects of brain disorders?* (medicine)
  - **Actions** preconditions and effects: *when can I press this button? What happens if I press it?*

Corresponding Reasoning Tasks / Problems

STATIC

- Constraint Satisfaction – **Find state that satisfies set of constraints.** *E.g., What is a feasible schedule for final exams?*
- Answering Query – **Is a given proposition true/likely given what is known?** *E.g., Does this patient suffers from viral hepatitis?*

SEQUENTIAL

- Planning – **Find sequence of actions to reach a goal state / maximize utility.** *E.g., Navigate through and environment to reach a particular location*

Representation and Reasoning System

- A (**representation**) **language** in which the environment and how it works can be described

- Computational (**reasoning**) **procedures** to compute a solution to a problem in that environment (an answer, a sequence of actions)

But the choice of an appropriate R&R system depends on a key property of the environment and of the agent's knowledge

Deterministic vs. Stochastic (Uncertain) Domains

- **Sensing Uncertainty**: Can the agent fully observe the current state of the world?
- **Effect Uncertainty:** Does the agent knows for sure what the effects of its actions are?

No No

Poker

No

Doctor Diagnosis

Factory Floor

Yes Yes

Chess

No

Doctor Treatment

Deterministic vs. Stochastic Domains

Historically, AI has been divided into two camps:

- those who prefer representations based on **logic**
- those who prefer **probability**.

Note: Some of the most exciting current research in AI is actually building bridges between these camps.

R&Rsys we'll cover in this course

Environment

Problem		Deterministic	Stochastic
Static	Constraint Satisfaction	*Vars + Constraints*: Arc Consistency, Search	
	Query	*Logics* (→ Propositional, → First Order, → ......): Search	*Belief Nets* (aka Bayesian Networks): Var. Elimination, Approx. Inference, Temporal. Inference
Sequential	Planning	*STRIPS* (actions, preconds, effects): Search	*Decision Nets* (aka influence diagrams): Var. Elimination; *Markov Processes*: Value Iteration

Representation

Reasoning Technique

Knowledge given vs. knowledge learned from experience

The agent is provided with a model of the world once and for all

- The agent can learn how the world works based on experience
  - in this case, the agent often still does start out with some **prior knowledge**

- **Supervised Machine Learning**
- **Unsupervised Machine Learning**
- **Reinforcement Learning**

Dimensions of Representational Complexity

We've already discussed:

- Reasoning tasks (Static vs. Sequential)
- Deterministic versus stochastic domains
- Knowledge given versus knowledge learned from experience (Machine Learning)

Some other important dimensions of complexity:

- Explicit state or propositions or relations
- Flat or hierarchical
- Goals versus complex preferences
- Single-agent vs. multi-agent

Explicit State or propositions

How do we model the environment?

- You can enumerate the **states** of the world.
- A state can be described in terms of **features**
  - Often it is more natural to describe states in terms of assignments of values to features (variables).
  - 30 binary features (also called propositions) can represent 2^{30}= 1,073,741,824 states.

Mars Explorer Example

Weather

Temperature

Long *Lati*

Explicit State or propositions or relations

- States can be described in terms of **objects** and **relationships**.
- There is a proposition for each relationship on each "possible" tuple of individuals.

University Example

Registred(S,C)

Students (S) = { }

Courses (C) = { }

- Textbook example: One binary relation and 10 individuals can represents 10^2=100 propositions and 2^{100} states!

Flat or hierarchical

Is it useful to model the whole world at the same level of abstraction?

- You can model the world at one level of abstraction: **flat**
- You can model the world at multiple levels of abstraction: **hierarchical**
- *Example: Planning a trip from here to a resort in Cancun, Mexico*

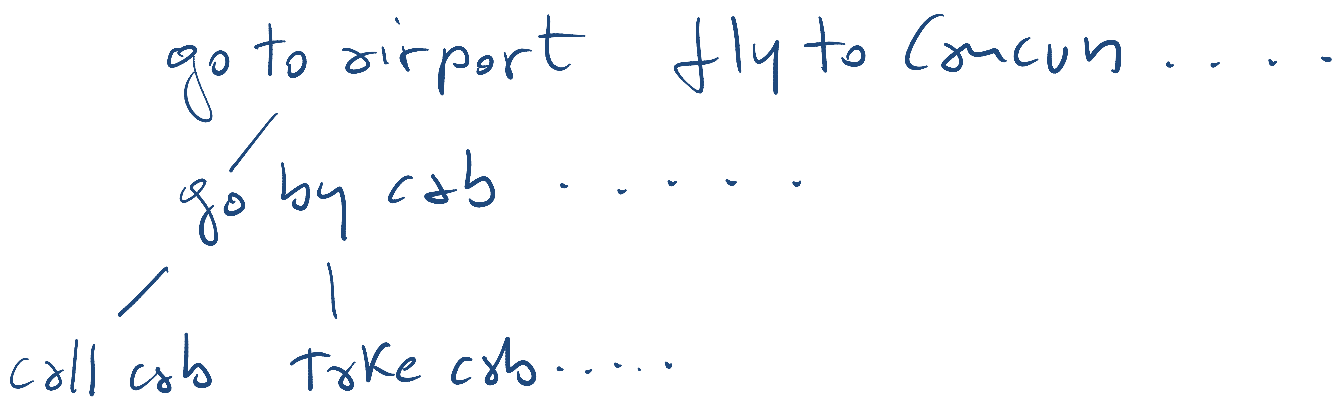

Goals versus (complex) preferences

An agent may have a **goal** that it wants to achieve

- e.g., there is some **state or set of states** of the world that the agent wants to be in
- e.g., there is some **proposition or set of propositions** that the agent wants to make true

An agent may have **preferences** [0, 1] (worst 0, best 1)

- e.g., there is some **preference/utility function** that describes how happy the agent is in each state of the world; the agent's task is to reach a state which makes it as happy as possible

Preferences can be **complex**…

What beverage to order?

- *The sooner I get one the better*
- *Cappuccino better than Espresso*

but Cappuccino takes 2 mins
Espresso takes 1 min
Agent must consider a TRADE-OFF

Single-agent vs. Multiagent domains

Does the environment include other agents?

Everything we've said so far presumes that there is only one agent in the environment.

- If there are other agents whose actions affect us, it can be useful to **explicitly model their goals and beliefs** rather than considering them to be part of the environment

- Other Agents can be: **cooperative**, **competitive**, or **a bit of both**

Dimensions of Representational Complexity

- Reasoning tasks (Constraint Satisfaction / Logic&Probabilistic Inference / Planning)
- Deterministic versus stochastic domains
- Knowledge given versus knowledge learned from experience

Some other important dimensions of complexity:

- Explicit state or features or relations
- Flat or ~~hierarchical~~
- Goals vs. (~~complex~~ simple) preferences
- Single-agent v~~s. multi-agent~~

grad course

Today Sept 8

Overview of the field – Key definitions

Overview of course

- Background knowledge
- Topics
- **Activities and Grading**
- **Administrative Stuff**

Assignment 0

Activities

- For the first 9 or 10 weeks we will cover an overview of the material in **lectures** and **assignments**
- There will be a **midterm** on the material covered in class
- The last few weeks will cover **current research topics presented by students**
- Students will **write a review paper** based on the lecture they gave (+ additional relevant papers), and **peer review** the papers of other students

Tentative Grading Scheme

- 5% Class Participation
- 25% Assignments
- 30% Midterm exam
- 10% Your Presentation
- 30% Your Essay

Textbook - Readings

Required

- Selected Chapters of **Artificial Intelligence: foundations of computational agents** by D. Poole and A. Mackworth, Cambridge University Press, 2010
- Between 15 and 20 research papers from recent conferences/journals in Artificial Intelligence

Reference

- **Artificial Intelligence : A Modern Approach**, by Russell and Norvig, 3rd Edition (Prentice-Hall, 2010)
- **Synthesis Lectures in Artificial Intelligence** webpage

Class Forum: Piazza

Join the class **asap** via the signup link below.

http://www.piazza.com/ubc.ca/fall2011/cpsc502

You need a **ubc.ca or cs.ubc.ca** email address to sign up. If you do not have one, please send an email to rjoty@cs.ubc.ca

TODO for next Tue

Read Chp 1 and 3 of textbook

Assignment 0: available from the course web page

- Join piazza (the class discussion forum)
- Read Course web-pages:

www.cs.ubc.ca/~carenini/TEACHING/CPSC502-11/502-11.html

WebSearch: Giuseppe Carenini

This is where most information about the course will be posted, most handouts (e.g., slides) will be distributed, etc.

CHECK IT OFTEN!