Introduction to Artificial Intelligence (AI)

Computer Science cpsc322, Lecture 1

May, 8, 2012

Lecture Overview

- Course Essentials
- What is AI?
- Representation and Reasoning
- Course Overview
- AI applications……

People

Instructor

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

Teaching Assistants

- **Mahsa Imani** mimani@cs.ubc.ca

MSc

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

PhD

- **Nathan Tomer** ntomer@cs.ubc.ca

MSc

Course Essentials(1)

- **Course web-pages:**

  www.cs.ubc.ca/~carenini/TEACHING/CPSC322-12/index.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!

- **Lectures**:
  - Cover basic notions and concepts known to be hard
  - I will try to post the slides in advance (by 8AM).
  - After class, I will post the same slides inked with the notes I have added in class.
  - Each lecture will include a set of learning goals:

Student can….

Course Essentials(2)

- **Textbook**: *Artificial Intelligence*, 2nd Edition, by Poole, Mackworth.
  - It's free!
  - It's available electronically

    http://www.cs.ubc.ca/~poole/aibook/
  - We will cover at least Chapters: 1, 3, 4, 5, 6, 8, 9

Course Essentials(3)

- **WebCT:** discussion board
  - Use the discussion board for questions about assignments, material covered in lecture, etc. That way others can learn from your questions and comments!
  - Use email for private questions (e.g., grade inquiries or health problems).
- **AIspace** : online tools for learning Artificial Intelligence http://aispace.org/
  - Under development

  here at UBC!

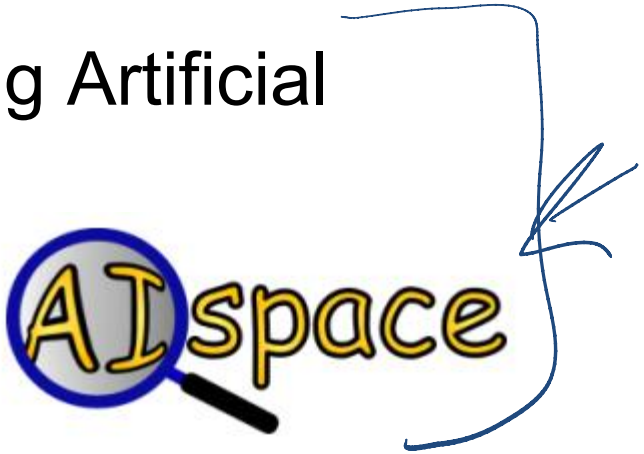

Course Elements

- **Practice Exercises:** 0% (we may do some in class)
- **Assignments:** 20%
- **Midterm:** 30%
- **Final:** 50%

If your final grade is >= 20% higher than your midterm grade:

- Assignments: 20%
- Midterm: 15% ↓
- Final: 65% ↑

Assignments

- **There will be four assignments in total**
  - They will not necessarily be weighted equally

- **Group work**
  - code questions:
    - ✓ you can work with a partner
    - ✓ always hand in **your own piece of code** (stating who your partner was)
  - written questions:
    - ✓ you may discuss questions with other students
    - ✓ you may not look at or copy each other's written work
    - ✓ you'll be asked to sign an honour code saying you've followed these rules

Assignments: Late Days

- **Hand in by 9am on due day** (in class or electronically)
- **You get three late days** ☺
  - to allow you the flexibility to manage unexpected issues
  - additional late days will not be granted except under truly exceptional circumstances
- **A day is defined as:** all or part of a 24-hour block of time beginning at 9 am on the day an assignment is due
- if you've used up all your late days, **you lose 20% per day**

Missing Assignments / Midterm / Final

Hopefully late days will cover almost all the reasons you'll be late in submitting assignments.

- However, something more serious like an extended illness may occur ☹

- **For all such cases:** you'll need to provide a note from your doctor, psychiatrist, academic advisor, etc.
- **If you miss:**
  - **an assignment,** your score will be reweighted to exclude that assignment
  - **the midterm,** those grades will be shifted to the final. (Thus, your total grade = 80% final, 20% assignments)
  - **the final,** you'll have to write a make-up final as soon as possible.

How to Get Help?

- Use the course **discussion board** on WebCT for questions on course material (so keep reading from it)

- Go to office hours (newsgroup is NOT a good substitute for this) –
  - **Shafiq**: Mon2pm (learning Center X150)
  - **Giuseppe:** Tue 2pm (CICSR #129)
  - **Nathan:** Wed 2pm (learning Center X150)
  - **Mahsa:** Thu 2pm (learning Center X150)

Can schedule by appointment if you can document a conflict with the official office hours

Getting Help from Other Students? From the Web? (Plagiarism)

- **It is OK to talk with your classmates about assignments; learning from each other is good**
- **But you must:**
  - Not copy from others (with or without the consent of the authors)
  - Write/present your work completely on your own (code questions exception)

- **If they use external source (e.g., Web)** in the assignments. Report this.

e.g., "bla bla bla….." [wikipedia]

Getting Help from Other Sources? (Plagiarism)

When you are in doubt whether the line is crossed:

- Talk to me or the TA's
- See **UBC official regulations** on what constitutes plagiarism (pointer in course Web-page)
- Ignorance of the rules will not be a sufficient excuse for breaking them

Any unjustified cases will be **severely dealt with** **by the Dean's Office** (that's the official procedure)

- My advice: better to skip an assignment than to have "academic misconduct" recorded on your transcript and additional penalties as serious as expulsion from the university!

To Summarize

- All the course logistics are described in the course Webpage

www.cs.ubc.ca/~carenini/TEACHING/CPSC322-12/index.html

WebSearch: Giuseppe Carenini

(And summarized in these slides)

- Make sure you carefully read and understand them!

What is Intelligence?

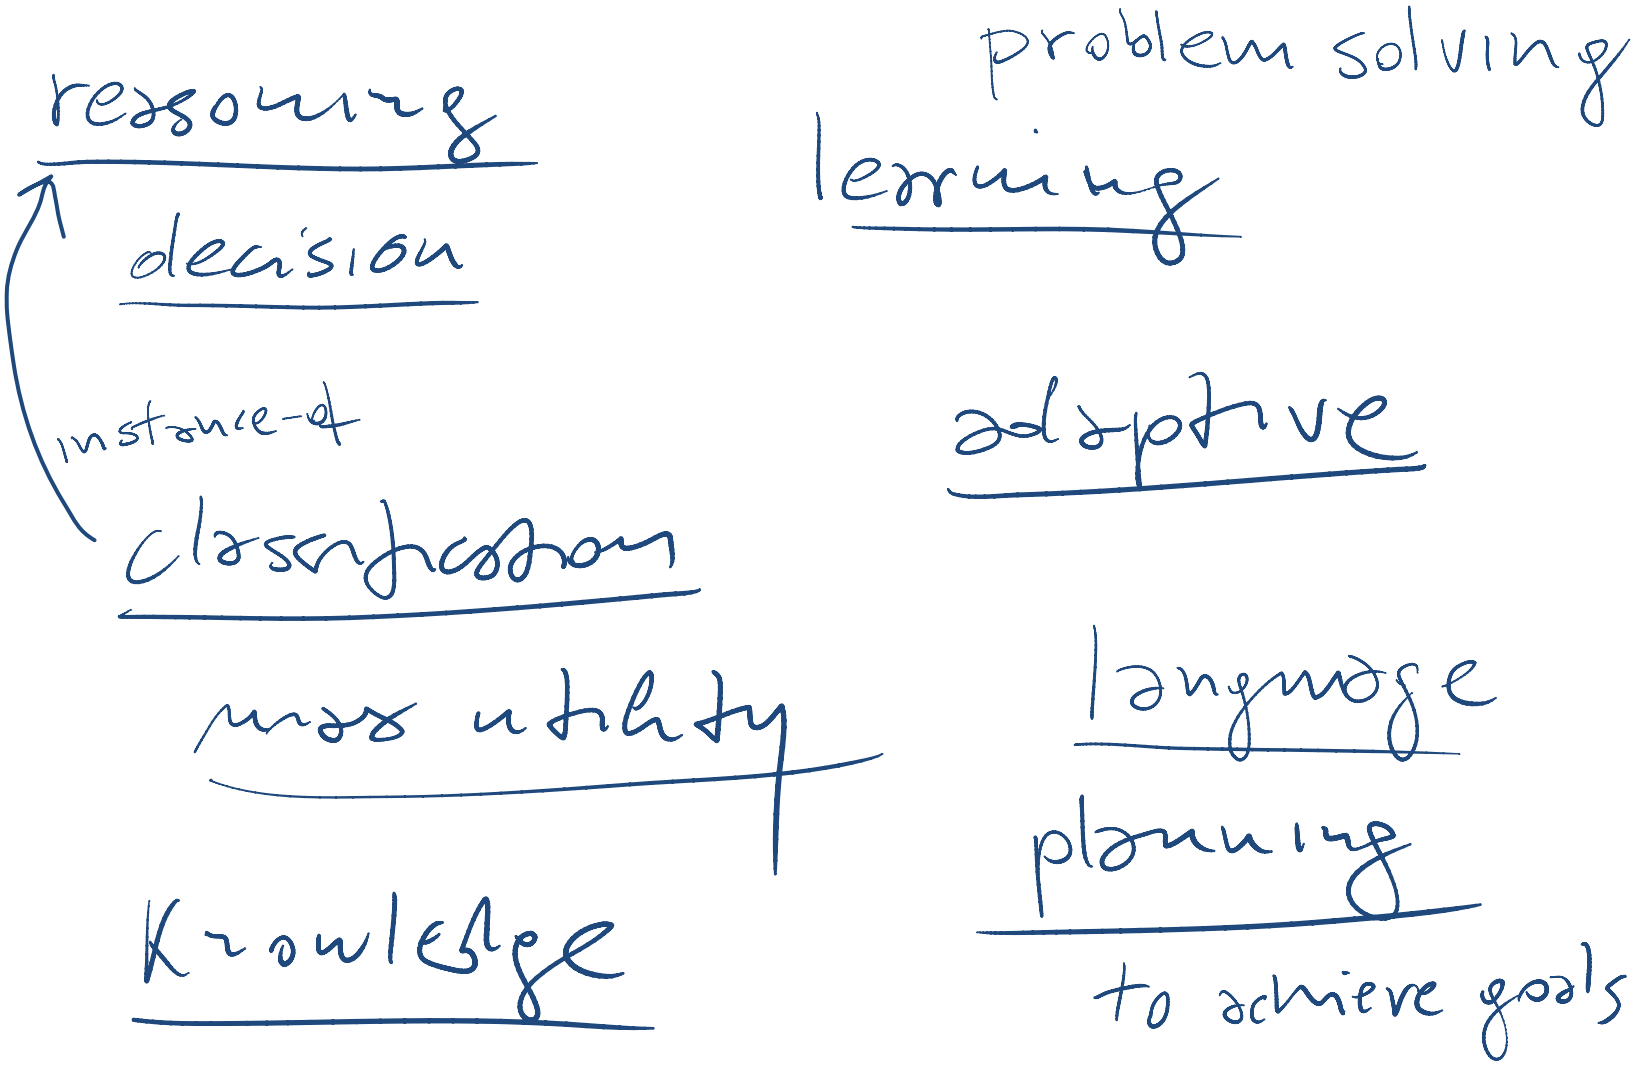

What is Artificial Intelligence?

Two definitions that have been proposed:

- Systems that **think** and **act** like **humans**

- Systems that **think** and **act rationally**

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?

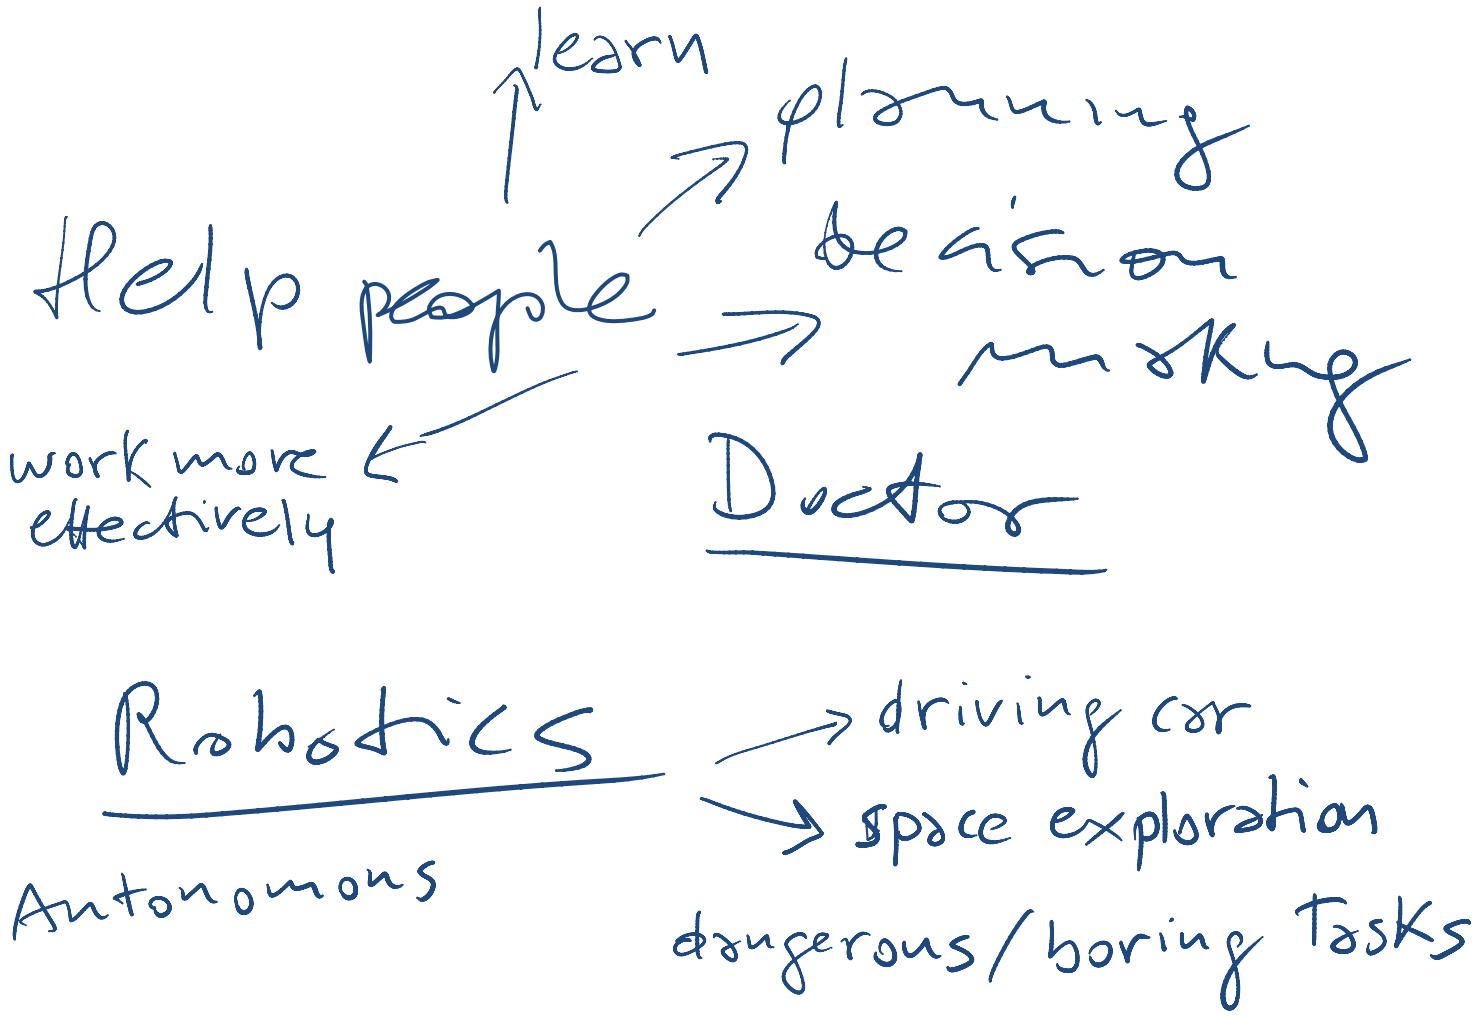

Agents acting in an environment

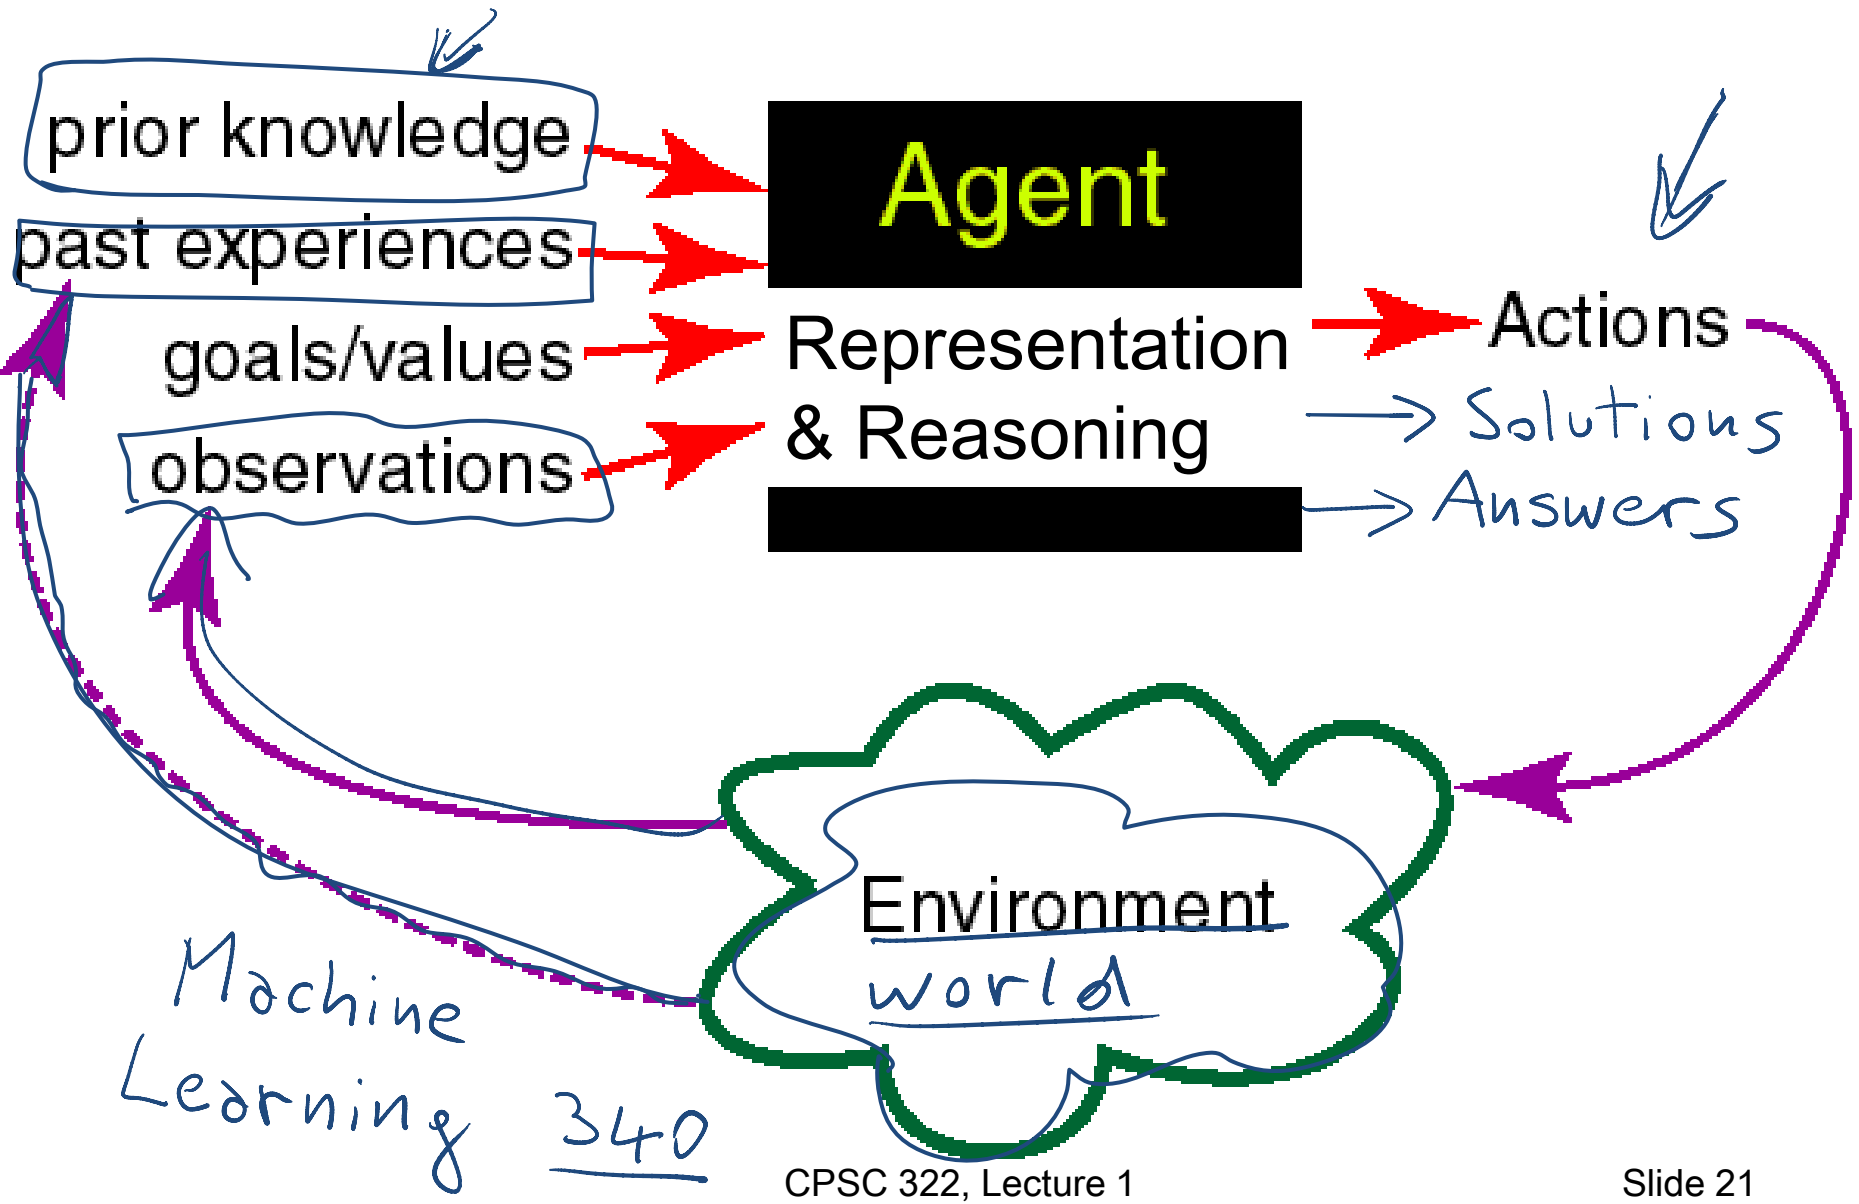

What is an agent?

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, send an email)*
- 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)

Possible Group activity

Search on the Web for an example of an AI system and try to answer some of these questions

- What does the application actually do**?** Was it evaluated? Is it a **fielded system**?
- Why is it **intelligent?** Is it **learning?**
- What is its **environment?**
- What are its **observations ?**
- What are its **actions?**
- Does it model **goals or preferences?**
- What **AI technology** does it use?

Lecture Overview

- Course Essentials
- What is AI?
- **Representation and Reasoning**
- Course Overview
- AI applications……

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

sequential

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 know for sure what the effects of its actions are?

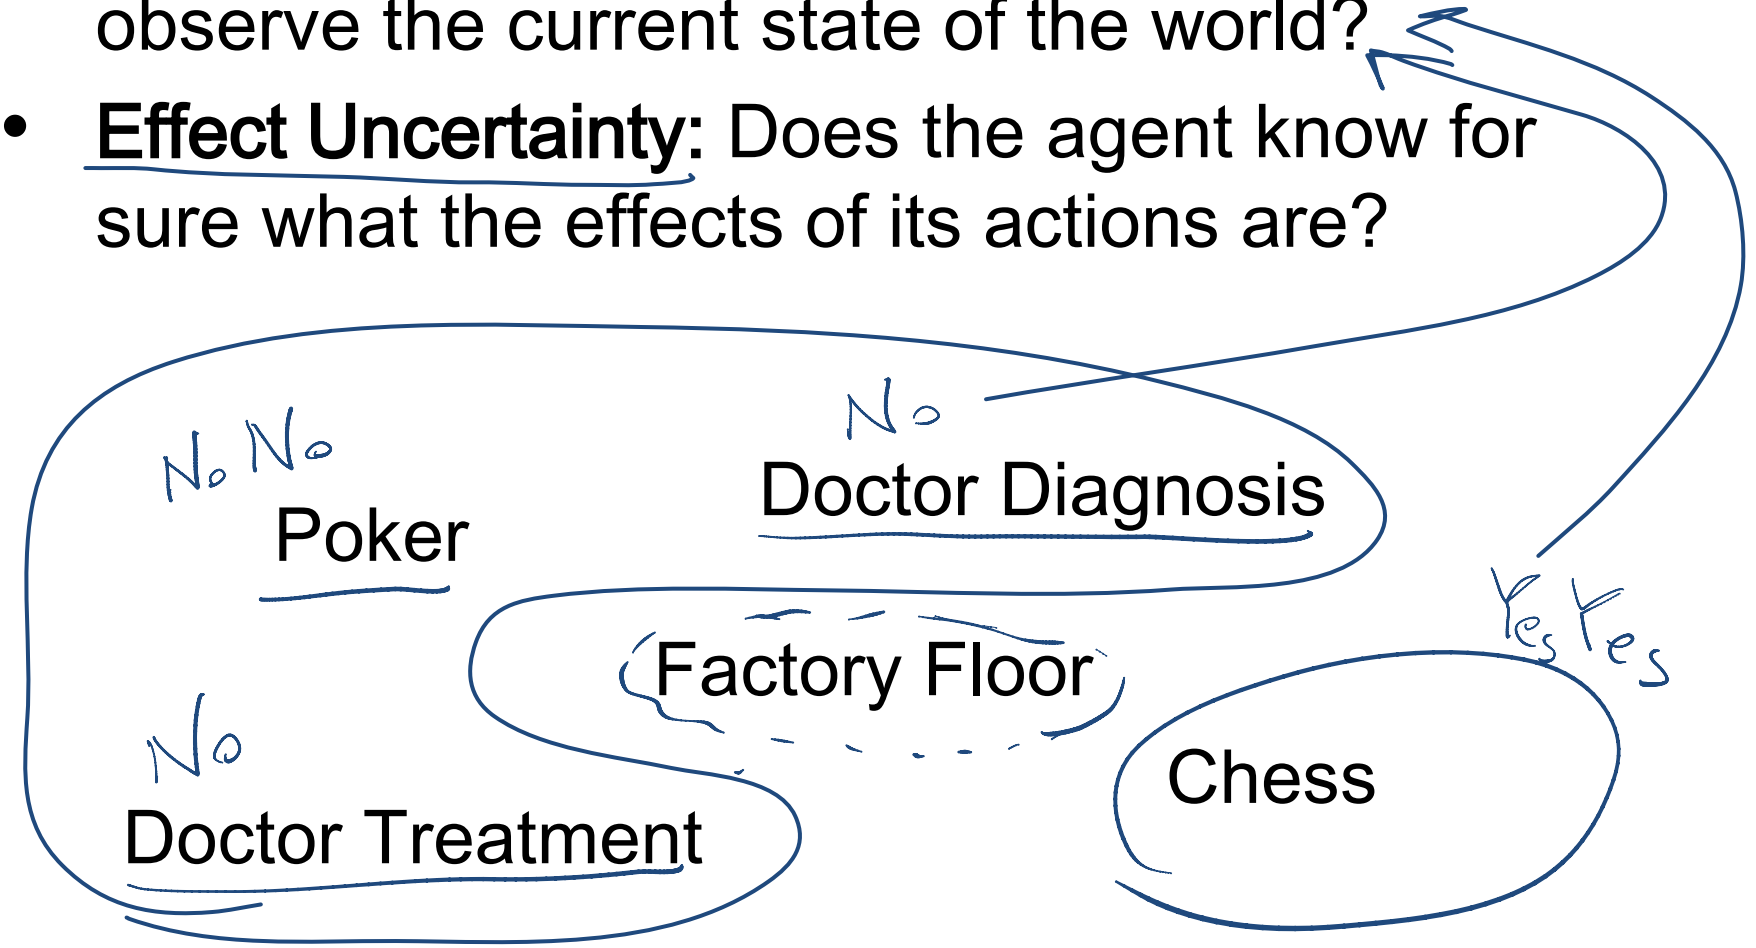

Deterministic vs. Stochastic Domains

Historically, AI has been divided into two camps: those who focus on representations based on **logic** and those who prefer **probability**.

A few years ago, CPSC 322 covered logic, while CPSC 422 introduced probability:

- now we introduce both representational families in 322, and 422 goes into more depth
- this should give you a better idea of what's included in AI

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

Lecture Overview

- Course Essentials
- What is AI?
- Representation and Reasoning
- **Course Overview**
- AI applications……

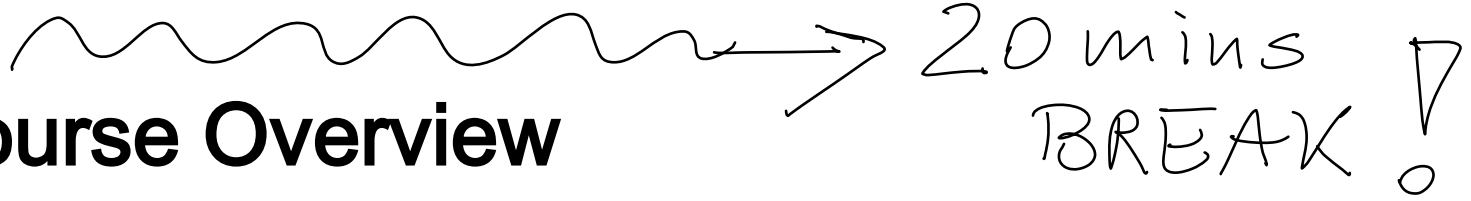

Modules we'll cover in this course: R&Rsys

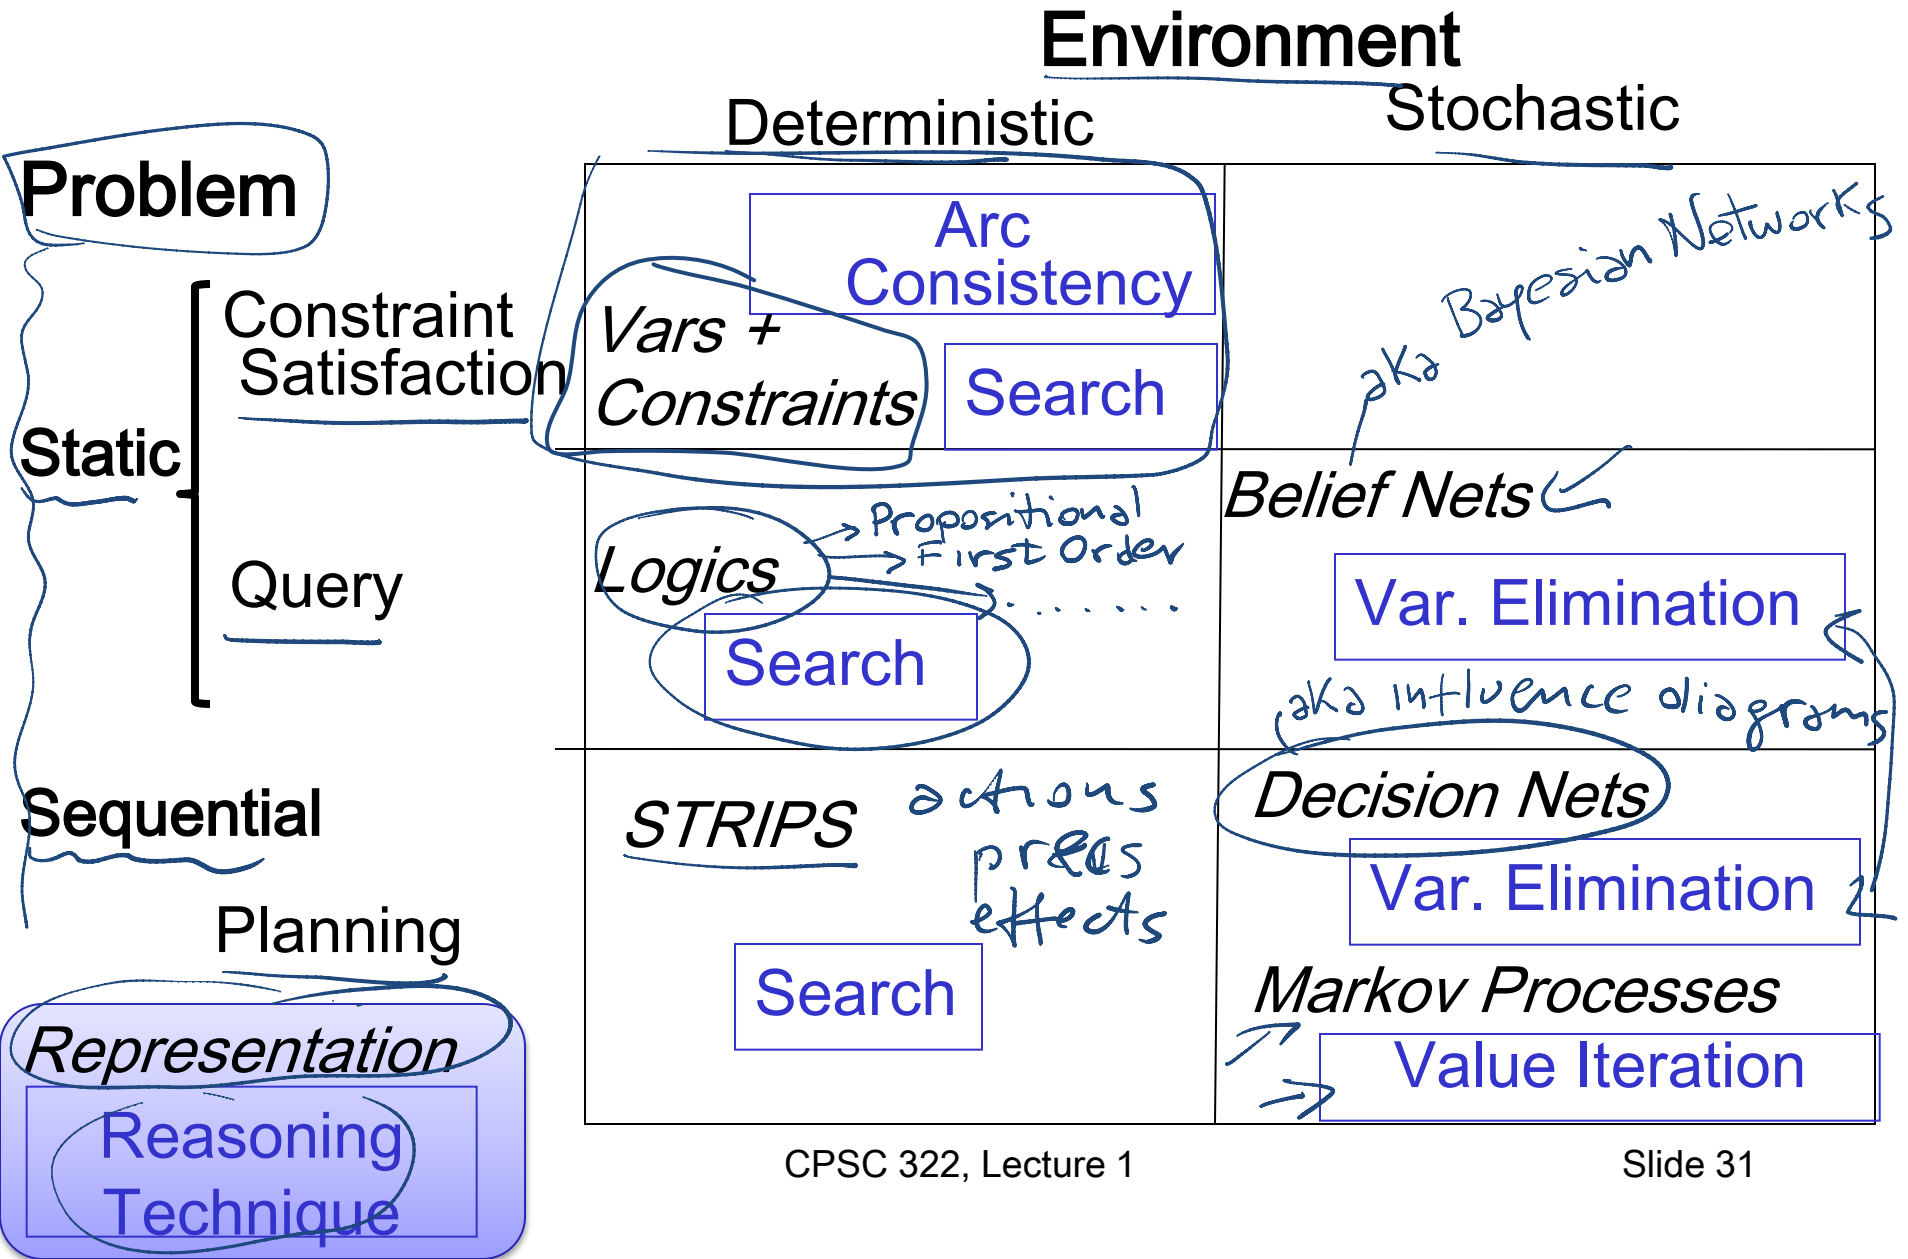

Dimensions of Representational Complexity

We've already discussed:

- Reasoning tasks (Static vs. Sequential)
- Deterministic versus stochastic domains

Some other important dimensions of complexity:

- Explicit state or propositions or relations
- Flat or hierarchical
- Knowledge given versus knowledge learned from experience
- 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* {S, C}
- *Temperature* {−40..+40}
- *LocX* 0-359 *LocY* 0 179

sunny cloudy

one possible state {S, +35, 30°, 110°}

2 * 81 * 360 * 180

number of possible states

mutually exclusive

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

1 relationship

Registred(S,C) → $R(S_1, C_2)$ etc

Students (S) = { s_1 s_2 s_3 s_4 }

Courses (C) = { c_1 c_2 c_3 }

individuals/objects

12 propositions

2^{12} states

- 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*

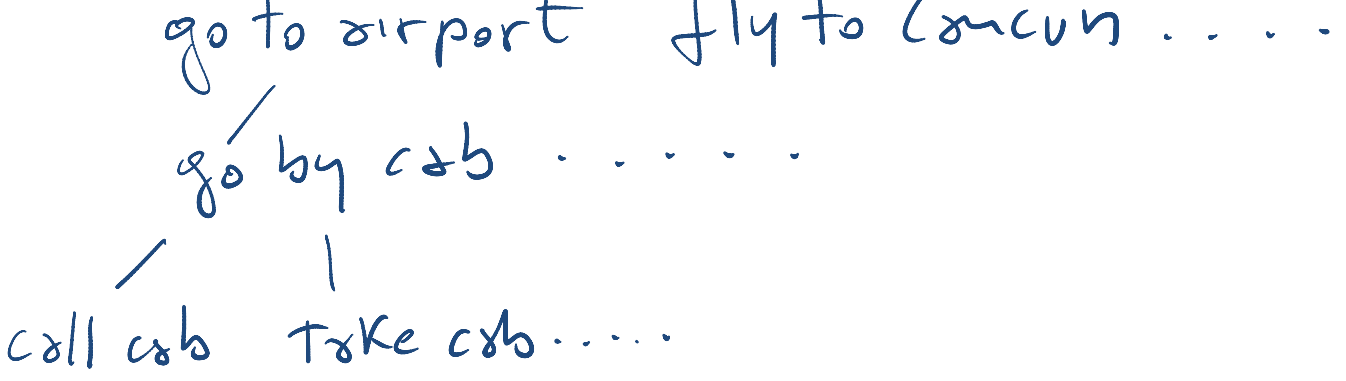

Knowledge given vs. knowledge learned from experience

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

not in this course

- 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**

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**

- 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*

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 in CPSC322

not in this course

- Reasoning tasks (Constraint Satisfaction / Logic&Probabilistic Inference / Planning)
- Deterministic versus stochastic domains

Some other important dimensions of complexity:

- Explicit state or features or relations
- Flat or ~~hierarchical~~
- Knowledge given versus knowledge ~~learned from experience~~
- Goals vs. (~~complex~~ simple) preferences
- Single-agent ~~vs. multi-agent~~ grad course

Lecture Overview

- Course Essentials
- What is AI?
- Representation and Reasoning
- Course Overview
- **AI applications……**

(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

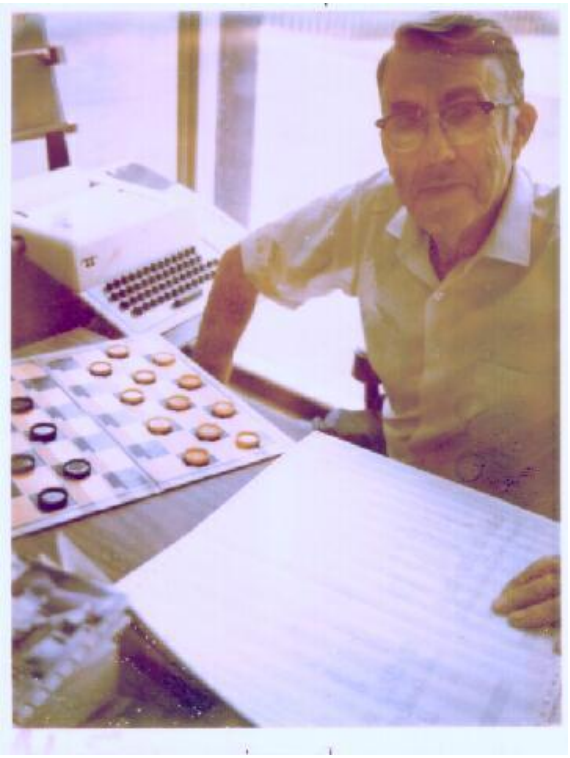

Source: *IBM Research*

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

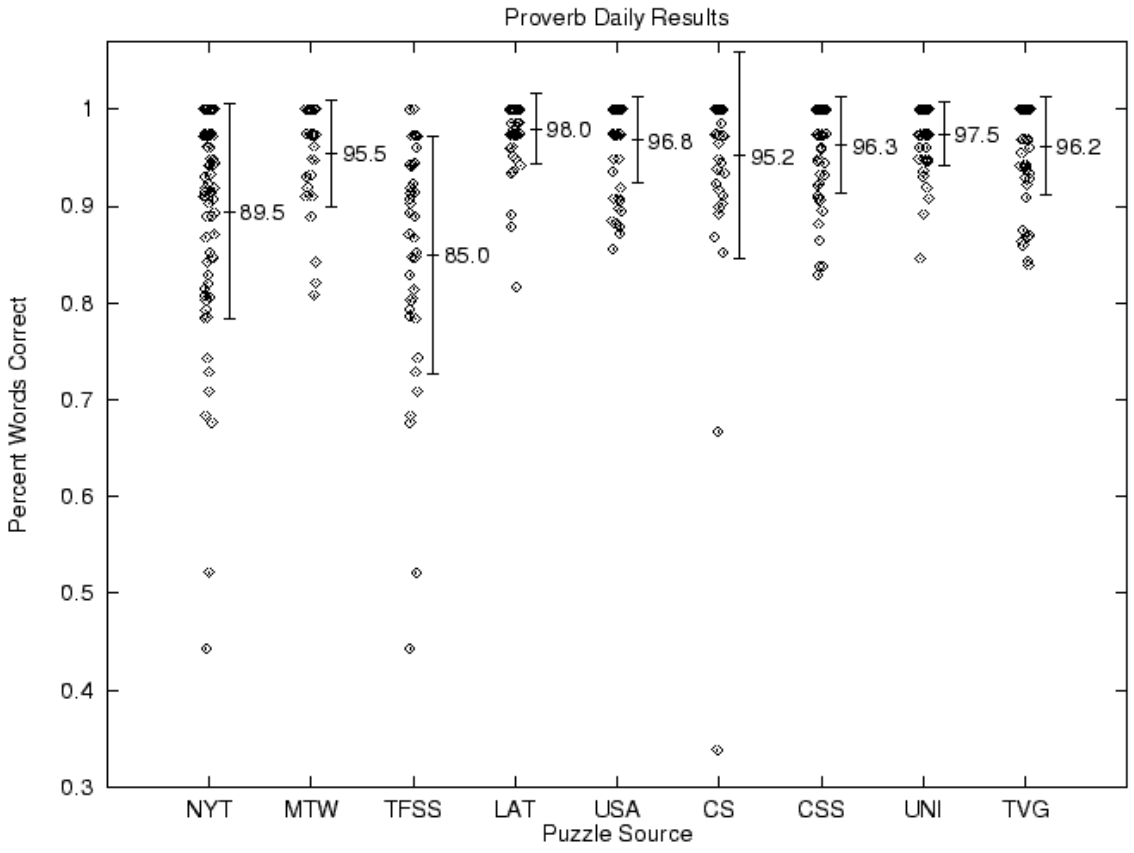

P	O	L	O	N	E	■	P	A	L	O	M	I	N	O
A	S	I	M	O	V	■	I	S	O	L	A	T	E	D
S	L	E	E	V	E	■	T	H	W	A	R	T	E	D
T	I	G	G	E	R	■	C	O	R	N	Y	■	■	■
A	N	E	A	L	E	■	A	R	I	D	■	J	A	M
■	■	■	■	E	S	P	I	E	S	■	L	O	G	O
S	E	A	O	T	T	E	R	■	E	E	E	N	O	N
A	B	B	O	T	■	A	N	A	■	U	S	A	G	E
B	O	O	Z	E	S	■	S	N	A	P	S	H	O	T
E	N	V	Y	■	P	L	I	N	T	H	■	■	■	■
R	Y	E	■	H	I	E	S	■	T	E	A	S	E	T
■	■	■	K	A	R	E	L	■	I	M	P	A	L	E
M	A	R	I	N	A	R	A	■	M	I	A	S	M	A
A	B	E	R	D	E	E	N	■	E	S	C	H	E	R
B	H	N	K	Y	A	R	D	■	S	M	E	A	R	S

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

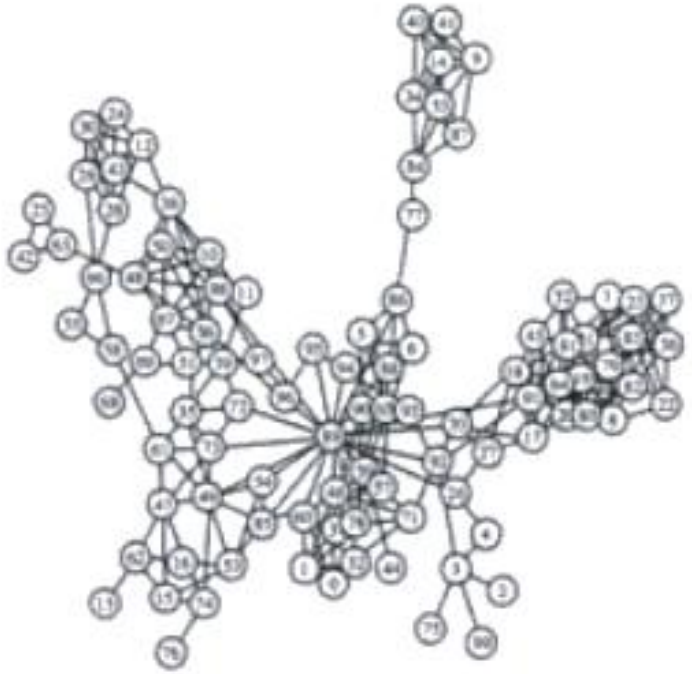

Example: SLS for RNA secondary structure design

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

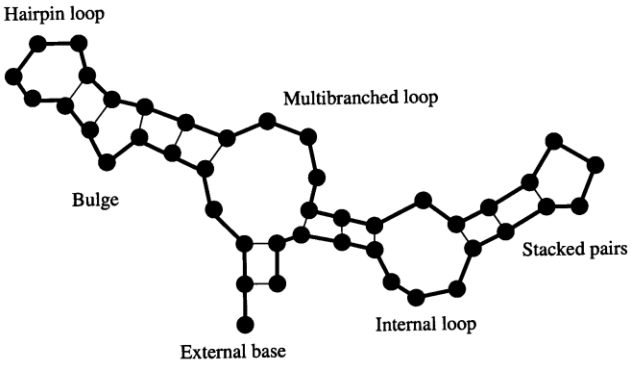

Best algorithm to date: Local search algorithm RNA-SSD developed at UBC [Andronescu, Fejes, Hutter, Condon, and Hoos, Journal of Molecular Biology, 2004]

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, …

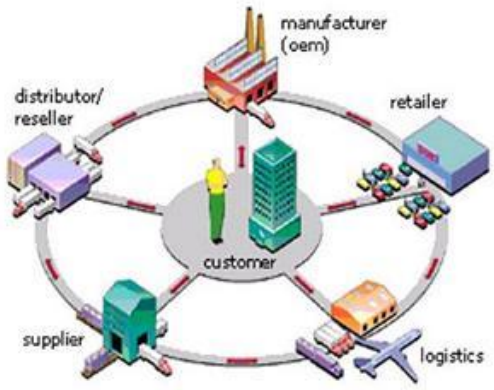

Supply chain management software:
Oracle,
SAP,…

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

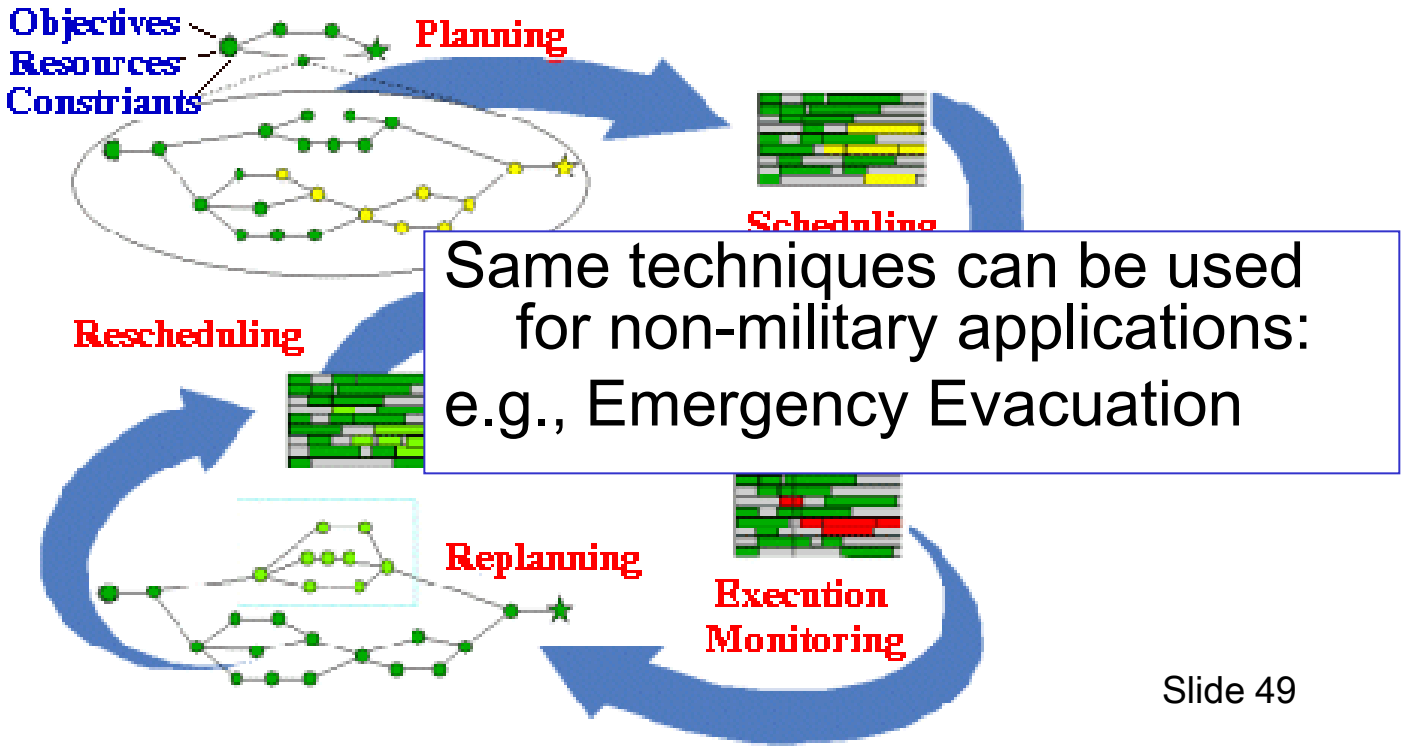

Source: *DARPA*

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*

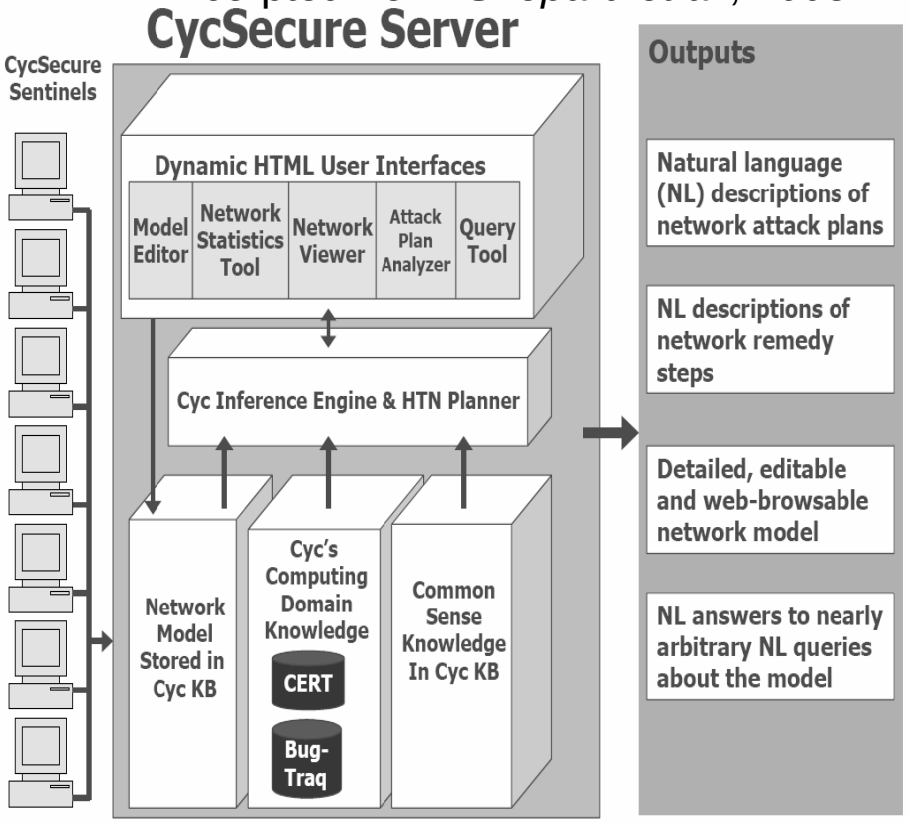

This formal representation also allows users to interact directly with the model of the network, allowing testing of proposed changes."

- **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

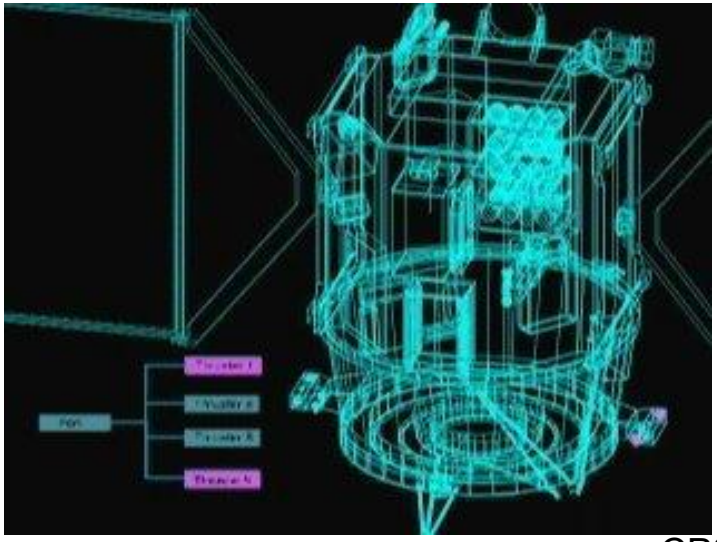

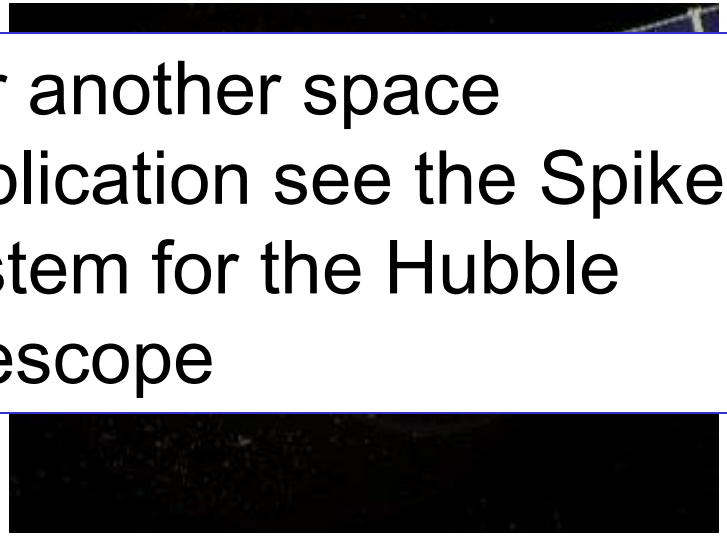

For another space application see the Spike system for the Hubble telescope

Source: *NASA*

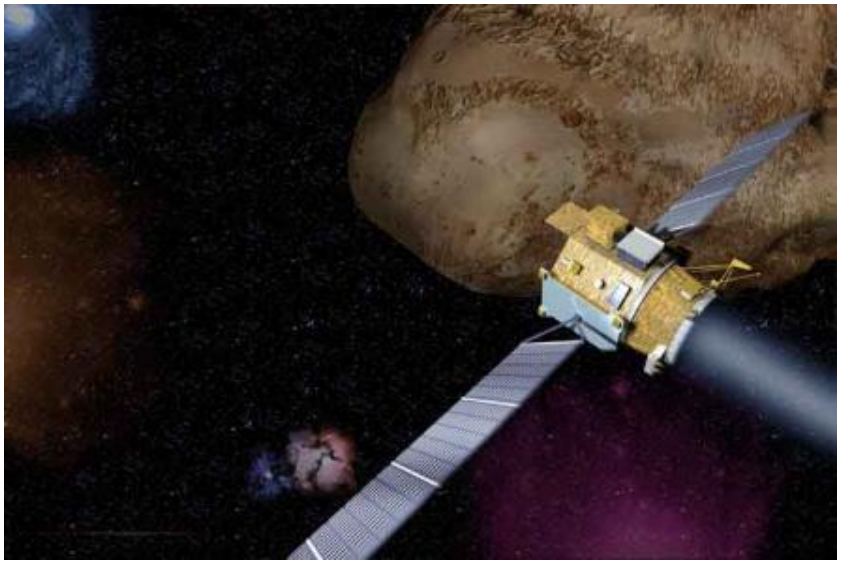

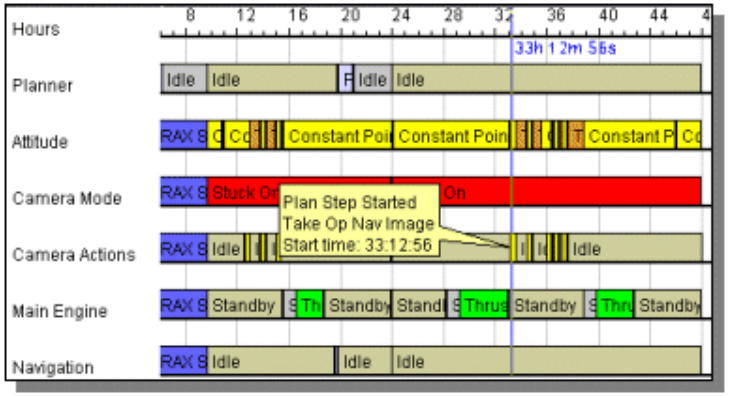

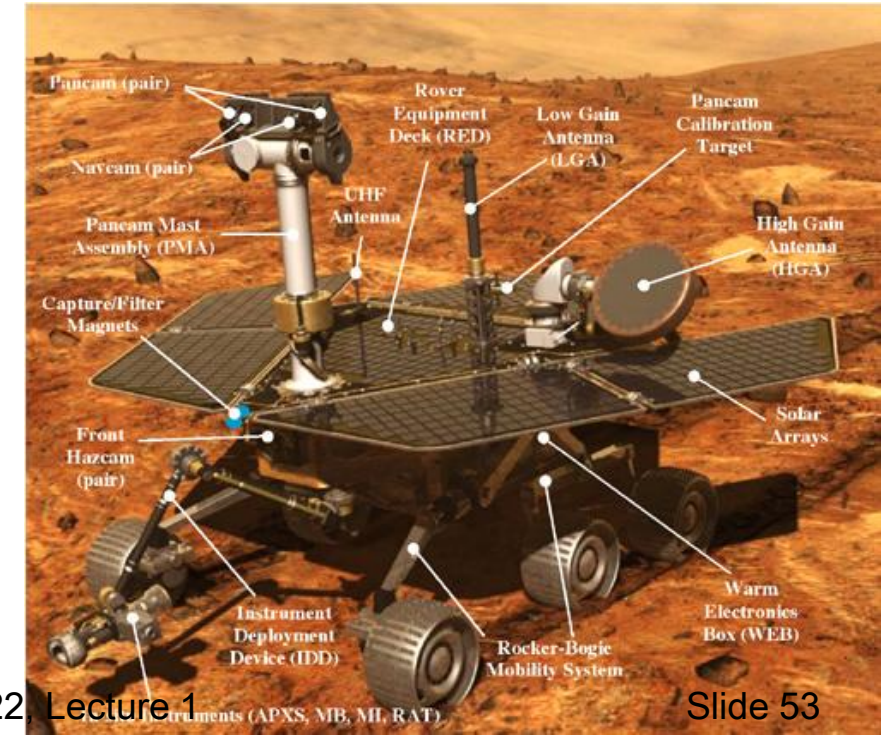

Source:
cs221 stanford

Reasoning under Uncertainty: Diagnosis

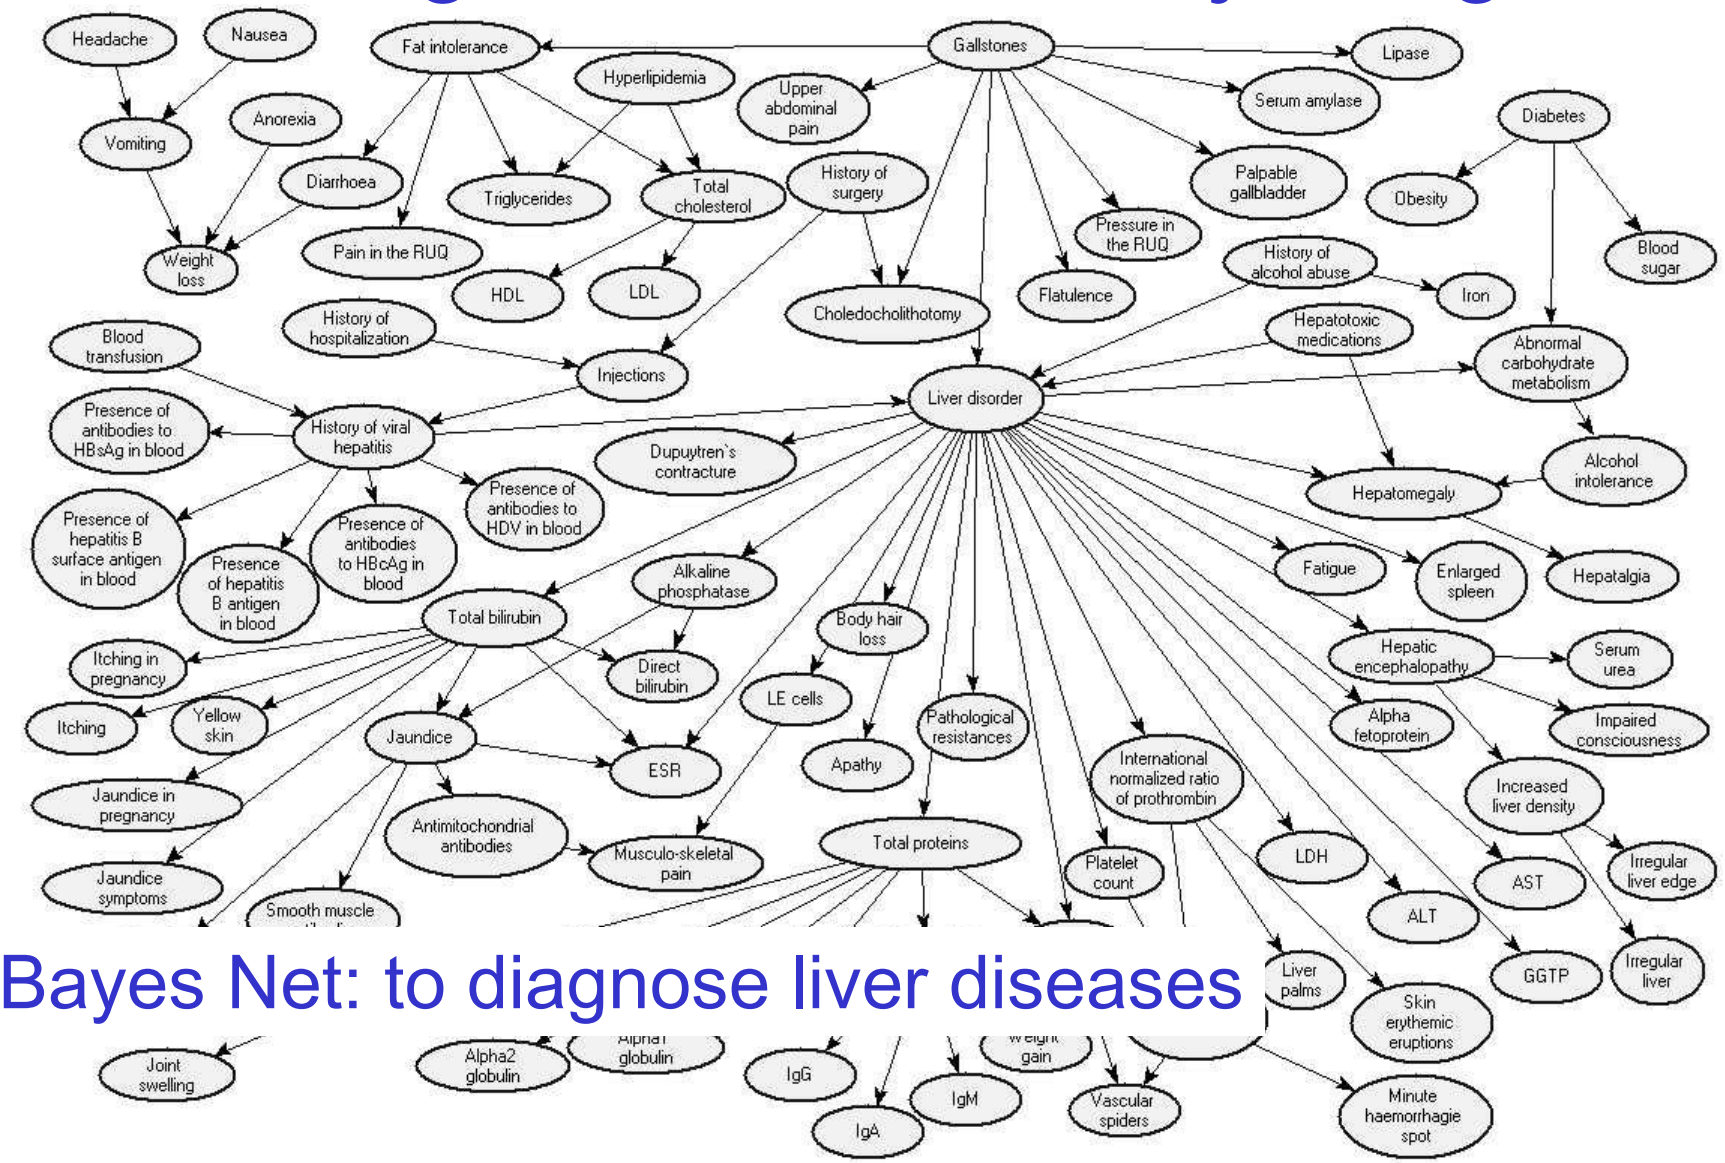

Bayes Net: to diagnose liver diseases

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

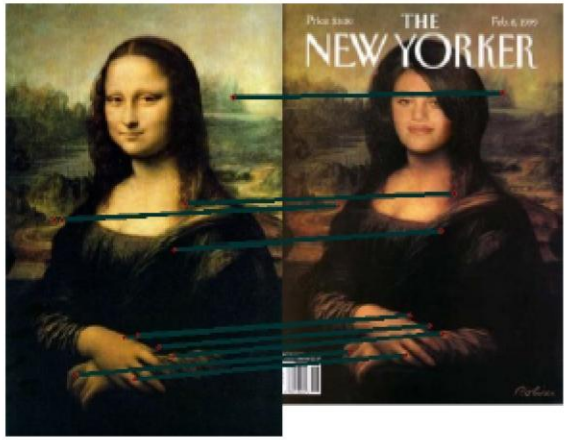

(a) A v.s. B

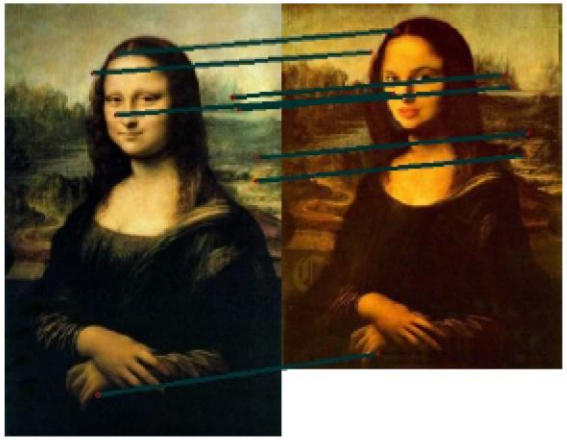

(b) A v.s. C

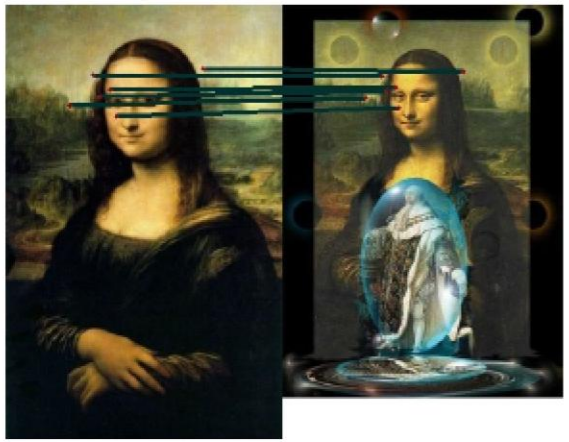

(c) A v.s. D

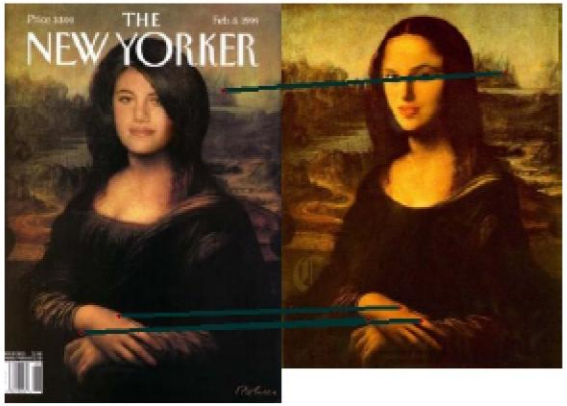

(d) B v.s. C

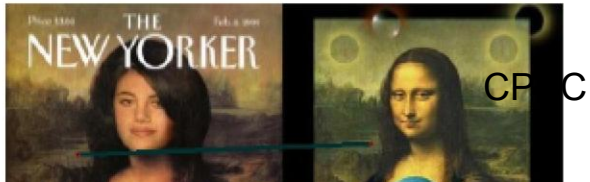

Source:
cs221 stanford

Induced Graph

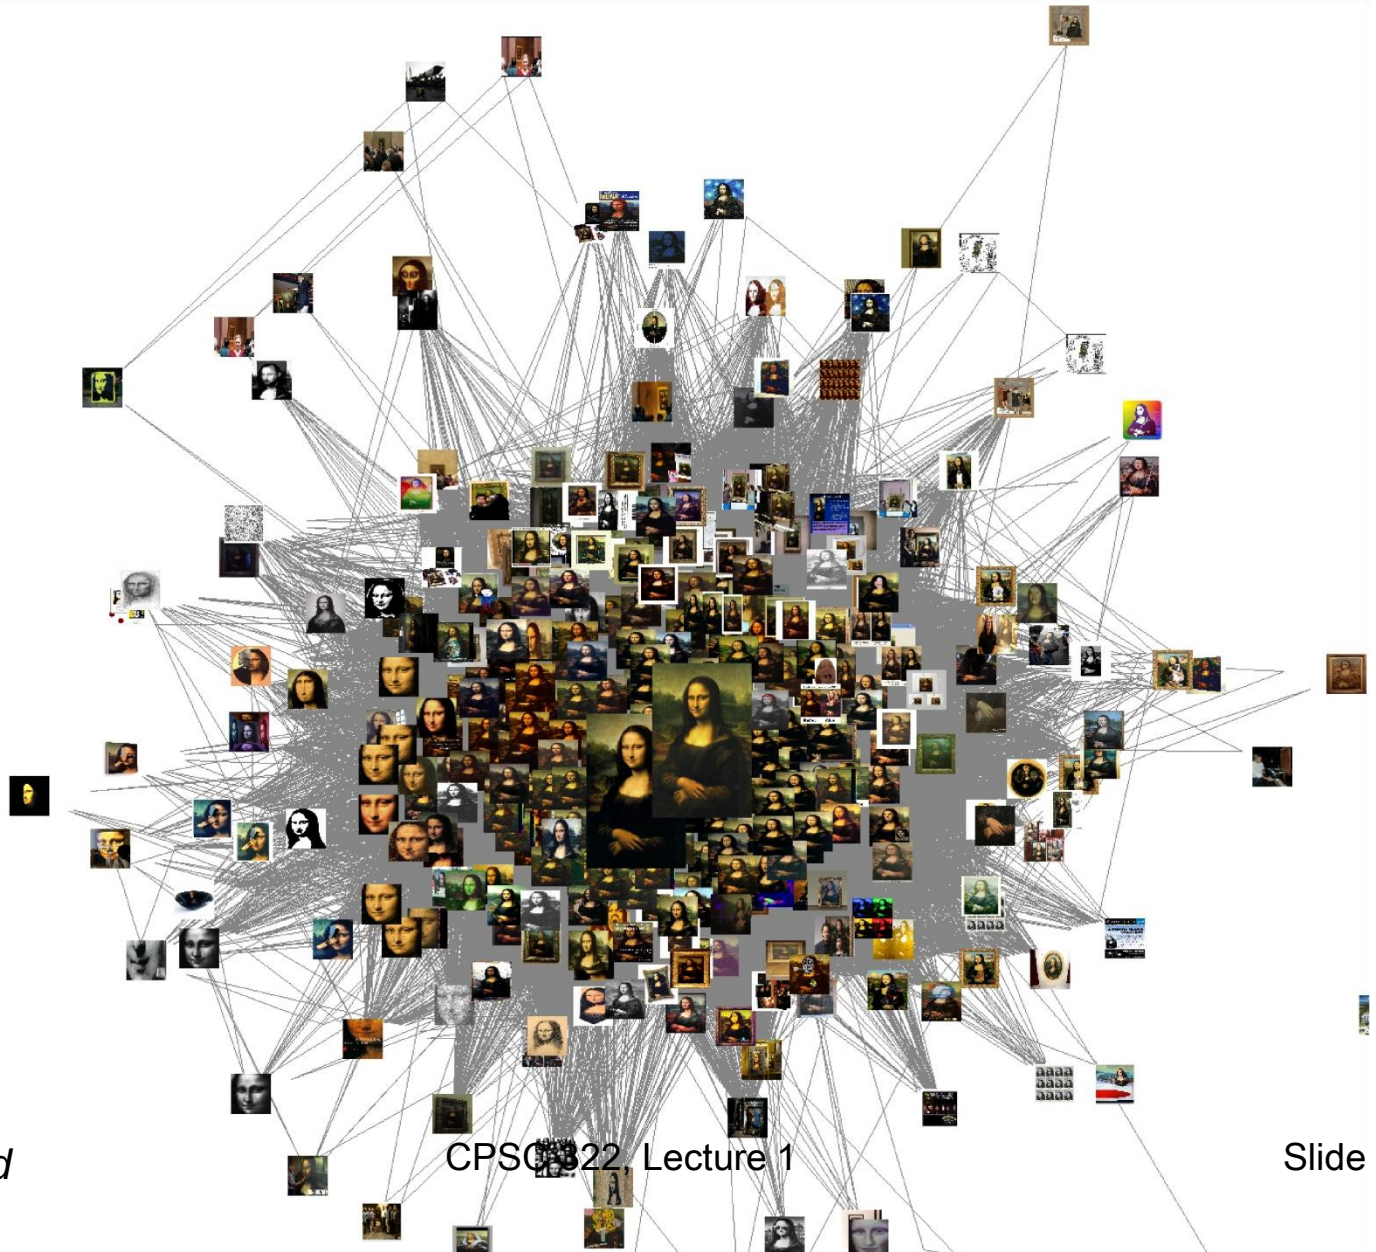

Source:
cs221 stanford

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,**

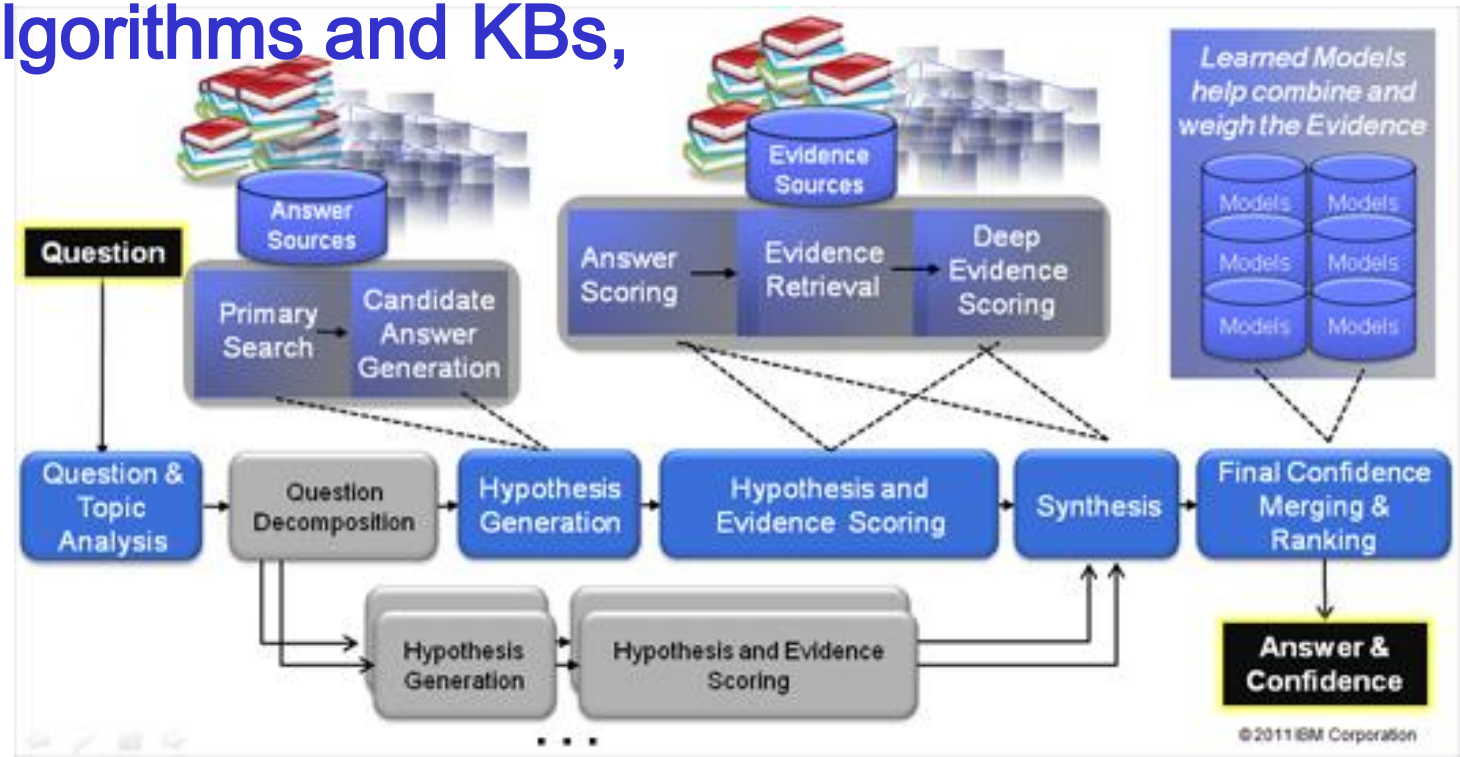

Source:
IBM

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

Final solution involves

Deterministic **search**!

Source: *Andrew Ng 2004*

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

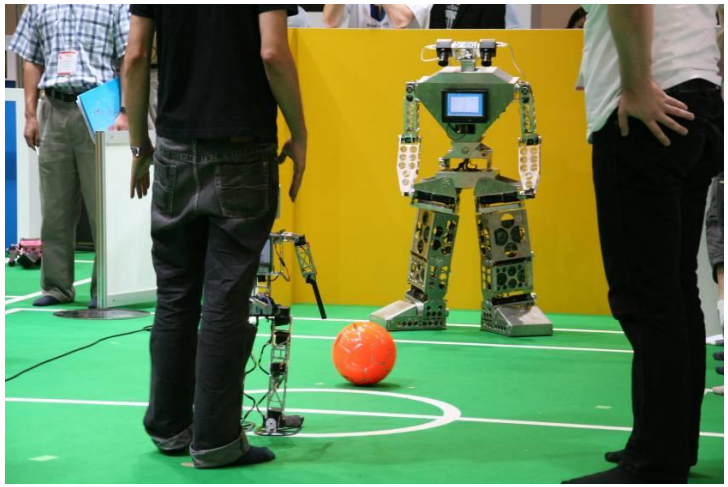

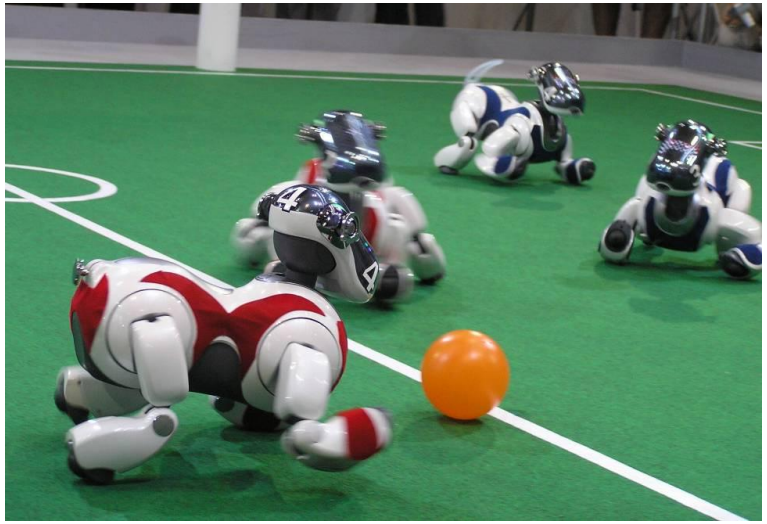

Extremely complex

- 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.

DE

EN

Source:
cs221 Stanford

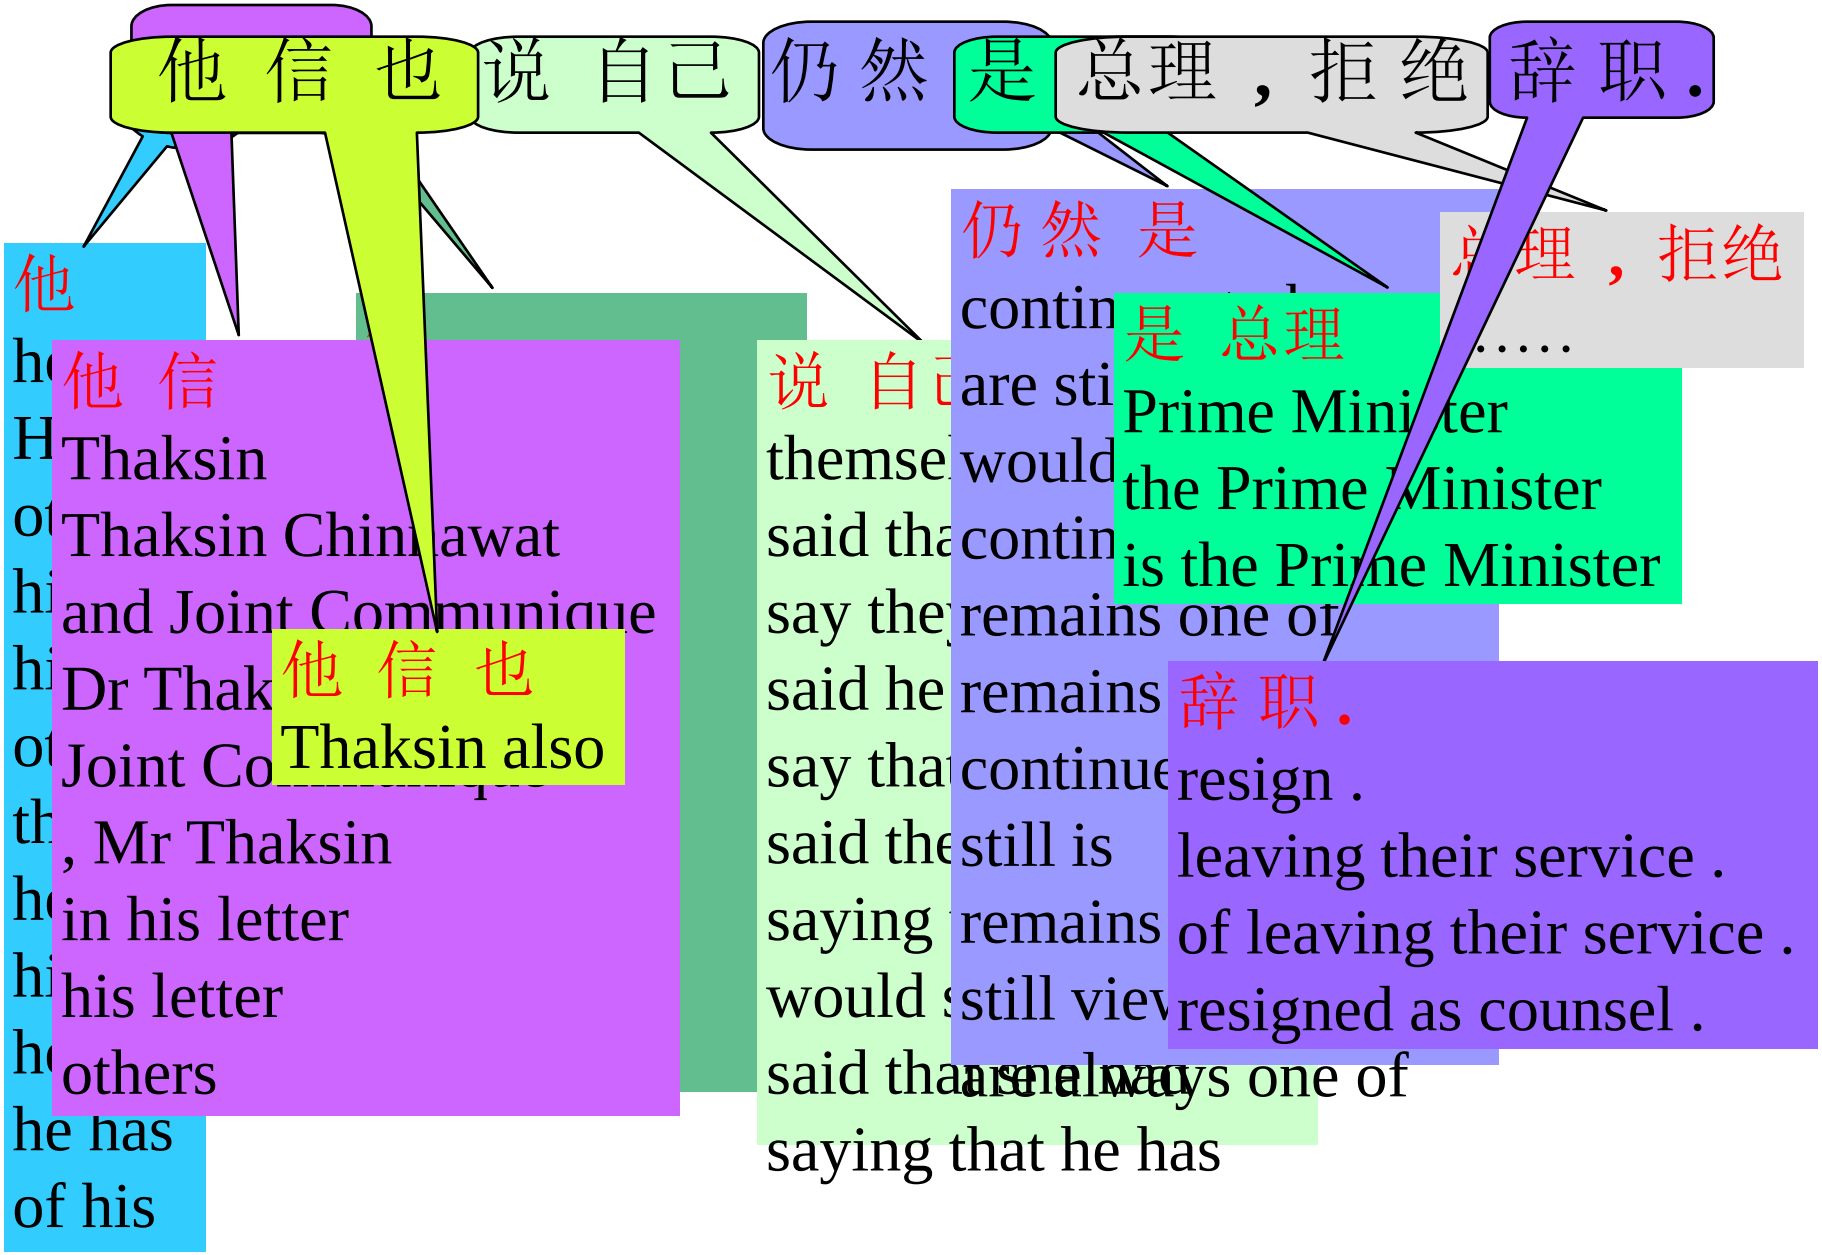
他 信 也
说 自己
仍 然 是
总理，拒绝
辞 职 .
他
he
H
of
hi
hi
of
th
he
hi
he
he has
of his
他 信
Thaksin
Thaksin Chinnawat
and Joint Communique
Dr Thak
Joint Co
, Mr Thaksin
in his letter
his letter
others
他 信 也
Thaksin also
说 自
themse
said tha
say the
said he
say tha
said the
saying
would s
said tha
saying that he has
仍 然 是
contin
are sti
would
contin
remains one of
remains
continue
still is
remains
still viev
are always one of
是 总理
Prime Mini ter
the Prime Minister
is the Prime Minister
总理，拒绝
.....
辞 职 .
resign .
leaving their service .
of leaving their service .
resigned as counsel .

Zite: a personalized magazine

… that gets smarter as you use it

David Poole, Indiana University, Poster Development Committee Chair
Poster Design: Giacomo Marchesi www.GiacomoMarchesi.com

TODO for Thurs

Read Chp 1

- Read carefully Section 1.6 : “Example Applications”
  - The Tutoring Systen
  - The trading agent
  - The autonomous delivery robot
  - The diagnostic assistant

- If your student Number is:
- 13950076 62462080 26750125 32404105
- Come and talk to me

Examples

Which of these things is an **agent**,

and why or why not?

- A soccer-playing robot?
- A rock?
- Machine Translator?
- A thermostat?
- A dog?
- A car?

Which of these things is an **intelligent agent**,

and why or why not?

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

- they act appropriately given goals and circumstances
- they are **flexible** to changing environments and goals
- they **learn** from experience
- they make appropriate choices given perceptual and computational limitations (sometimes they act without thinking!)
- They **gather information** (if cost less than expected gain)

Acting Humanly

The **Turing Test**

- Don't try to come up with a list of characteristics that computers must satisfy to be considered intelligent
- Instead, use an operational definition: consider it **intelligent when people can't tell a computer apart from other people**

The original test involved typing back and forth; the `**Total Turing Test** includes a video signal to test perception too

- But... is acting just like a person what we really want?
- For example, again, don't people often do things that we **don't** consider intelligent?