Representational Dimensions

Computer Science cpsc322, Lecture 2

(Textbook Chpt1)

Sept, 7, 2012

Lecture Overview

- **Recap from last lecture**
- Representation and Reasoning
- An Overview of This Course
- Further Dimensions of Representational Complexity

Course Essentials

- **Course web-page :** CHECK IT OFTEN!
- **Textbook**: Available online!
  - We will cover at least Chapters: 1, 3, 4, 5, 6, 8, 9
- **Connect:** discussion board, grades
- **AIspace** : online tools for learning Artificial Intelligence
  http://aispace.org/
- **Lecture slides…**
- **Midterm exam, Mon, Oct 29**(1 hours, regular room)

Agents acting in an environment

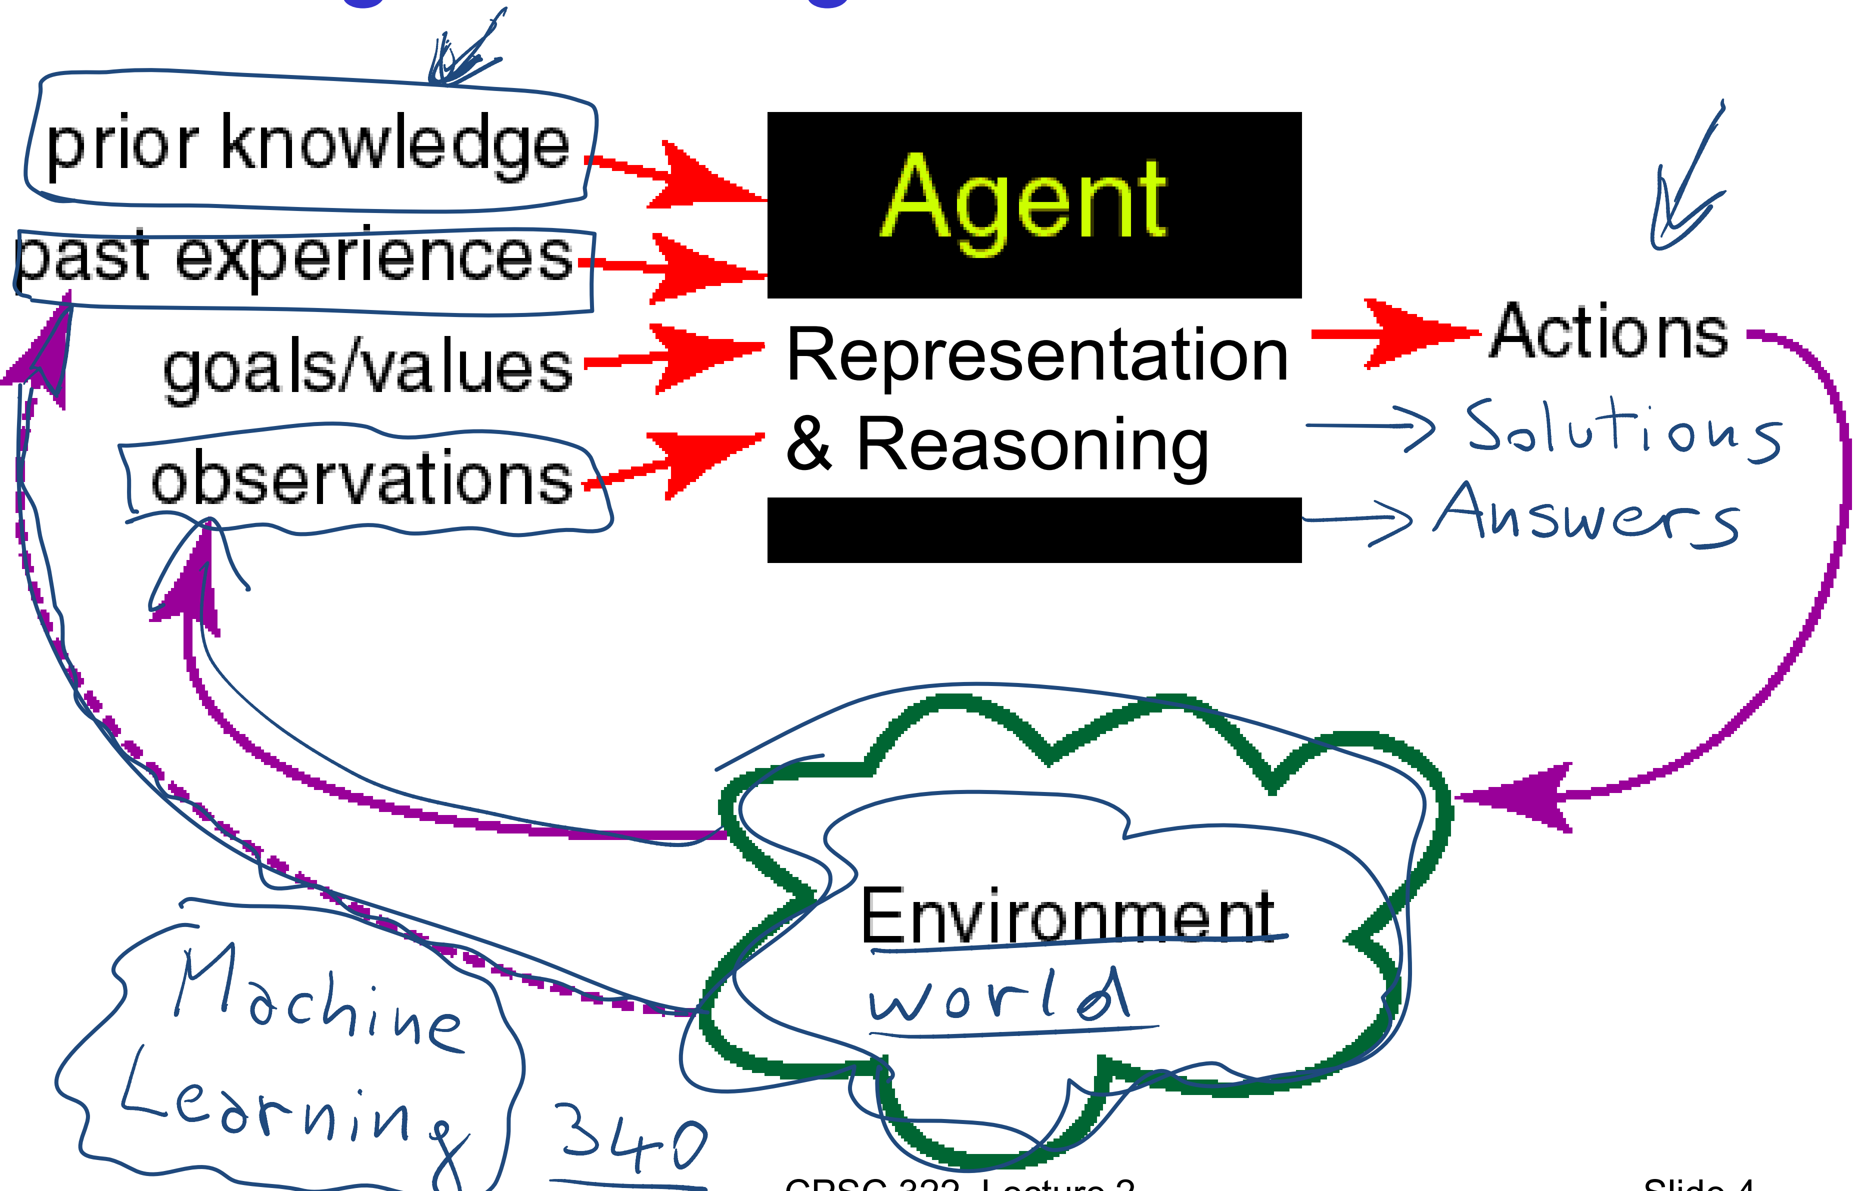

Lecture Overview

- Recap from last lecture
- **Representation and Reasoning**
- An Overview of This Course
- Further Dimensions of Representational Complexity

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*
  - **Causal:** *what are the causes and the effects of brain disorders?*
  - **Actions** preconditions and effects: *when can I press this button? What happens if I press it?*

electric circuit

medicine

Corresponding Reasoning Tasks / Problems

- **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. Collect gems and avoid monsters*

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?

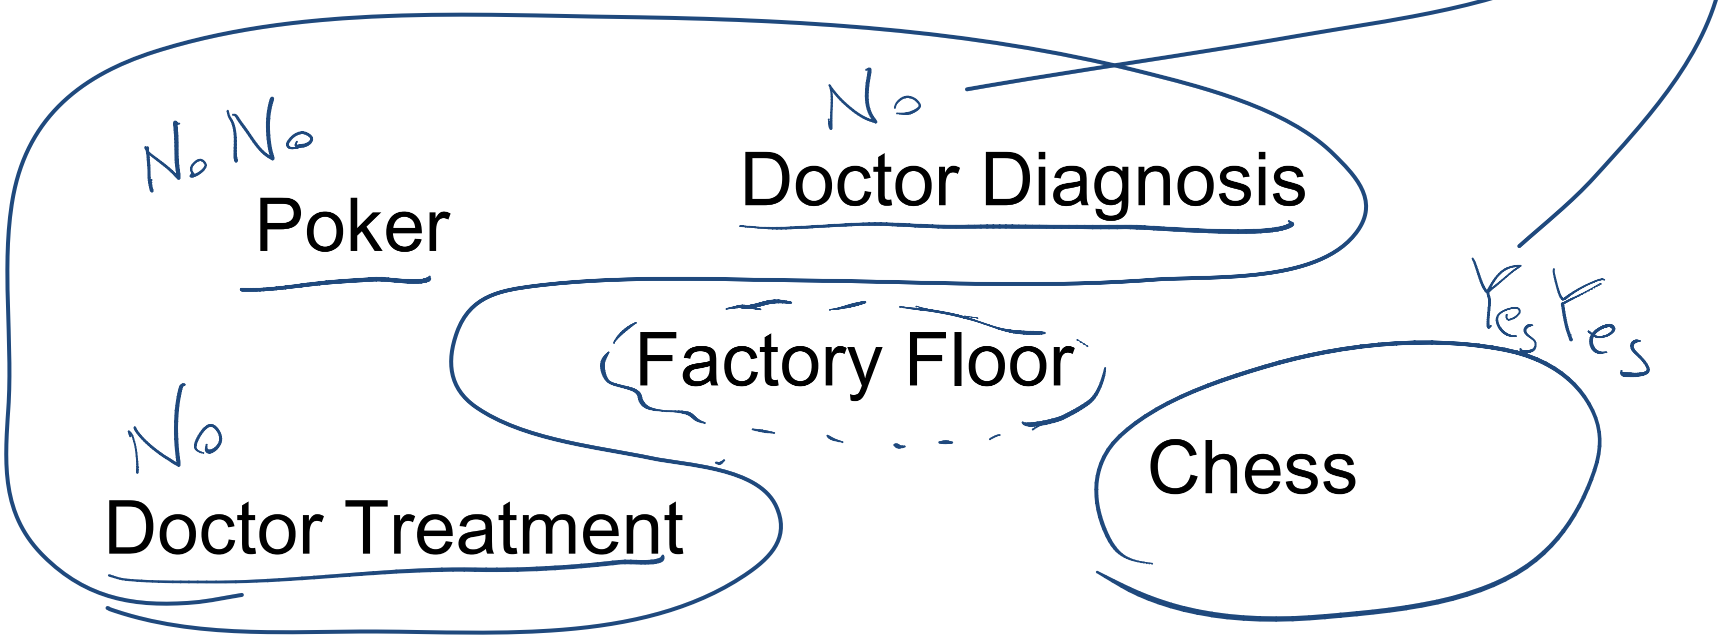

Deterministic vs. Stochastic Domains

Historically, AI has been divided into two camps: those who prefer 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

- Recap from last lecture
- Representation and Reasoning
- **An Overview of This Course**
- Further Dimensions of Representational Complexity

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

Problem		Environment: Deterministic	Environment: Stochastic
Static	Constraint Satisfaction	*Vars + Constraints* Arc Consistency Search	
Static	Query	*Logics* Search	*Belief Nets* Var. Elimination
Sequential	Planning	*STRIPS* Search	*Decision Nets* Var. Elimination *Markov Processes* Value Iteration

Representation

Reasoning Technique

Lecture Overview

- Recap from last lecture
- Representation
- An Overview of This Course
- **Further Dimensions of Representational Complexity**

Dimensions of Representational Complexity

We've already discussed:

- Problems /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

features

are binary features

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

LocX LocY

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) → $R(s_1, c_2) \in A_c$

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!

Complete Example

R = Registred $R(s,c)$

Courses = $\{c_1, c_2\}$ $|Courses| = 2$

Students = $\{s_1, s_2, s_3\}$ $|Students| = 3$

All possible instantiations of R } each of these can be T = true OR F = false

$3*2=6$

$R(s_1, c_1)$	T	T	T	T		F
$R(s_1, c_2)$	T	F	T	T		F
$R(s, c_1)$	T	T	F	T	......	F
$R(s_2, c_2)$	T	T	T	F		F
$R(s_3, c_1)$	T	T	T	T		F
$R(s_3, c_2)$	T	T	T	T		F

$2^6 = 64$ possible worlds

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*

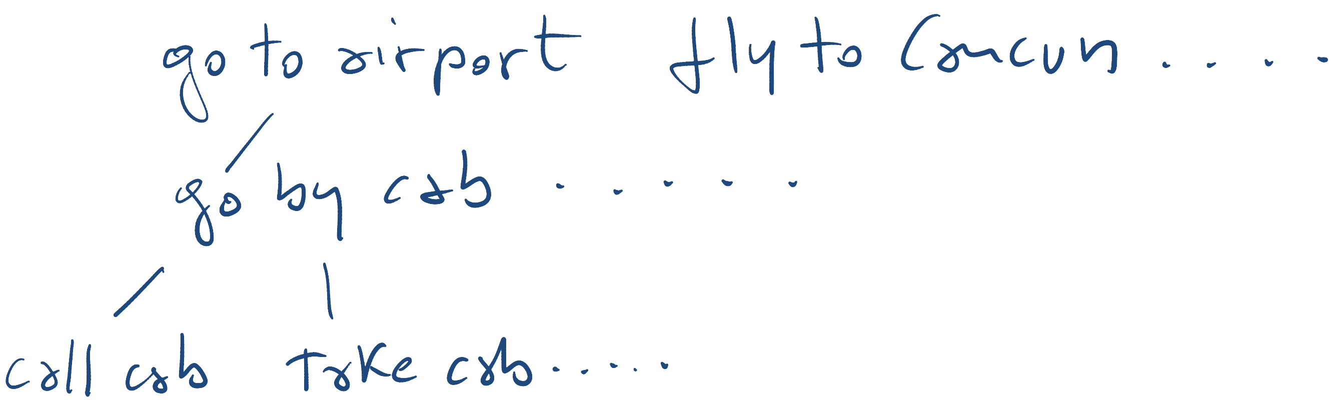

Knowledge given vs. knowledge learned from experience

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** [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 2mins
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 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

Next class (SEE AI LANDSCAPE ON NEXT SLIDE)

- Assignment 0 due: **submit electronically** and **you can't use late days**
- **Hint:** **AAAI is the main AI association**
- Come to class ready to discuss the **two examples of fielded AI agents** you found
  (or experimental)
- I'll show some pictures of cool applications in that class

- Read carefully Section 1.6 on textbook: "Example Applications"
  - The Tutoring System
  - The trading agent
  - The autonomous delivery robot
  - The diagnostic assistant

AI magazine
YAHOO! RESEARCH
NSF
AI Music Composition & Performance
AI Paintings
Humanoid Robot
How Can AI Systems Learn From Experience?
Robot Tour Guide
Characters for Virtual Worlds
Autonomous Vehicle
Robot Soccer
Social Simulation Game
Smart Environmental Controls
Intelligent Tutoring System
Vehicle Navigation System
Personalized TV Guide
Autonomous Space Exploration
Robot Vacuum Cleaner
Smart Desk with Gesture Recognition
How Can AI Systems Solve Problems Creatively?
Robots for Education
AI LAB
Recommender System
Robotic Surgery
Disease Diagnosis
Drug Design
Fraud Detection
Web Search
Smart Wheelchair
Handwriting & Sketch Recognition
BOOKS & COFFEE
BOOKS
Spam Filtering
Machine Translation
Aristotle
Leibniz
Descartes
Whitehead
Lovelace
Russell
Turing
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