Representational Dimensions

CPSC 322 Lecture 2

Lecture Overview

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

Agents acting in an environment

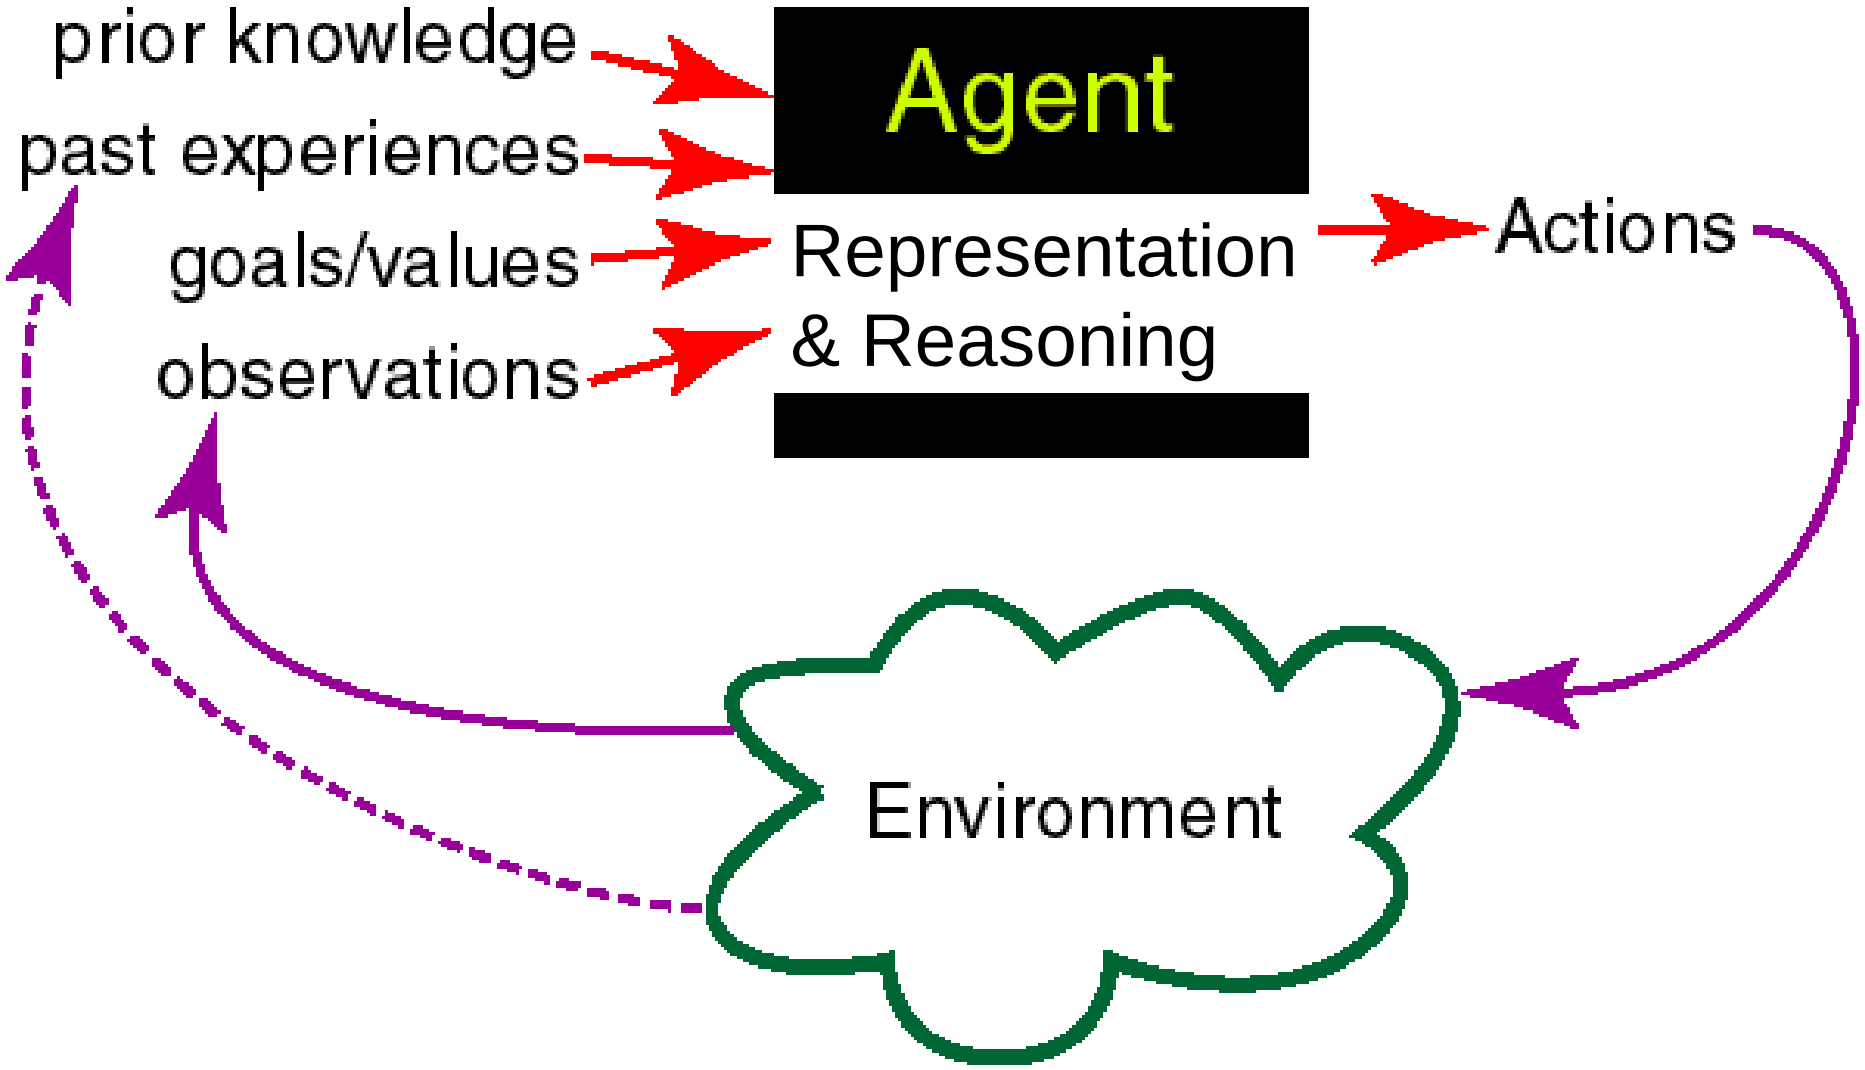

Lecture Overview

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

What do we need to represent ?

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:** *can only write one exam at once*
  - **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?*

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 suffer from chicken pox?*
- **Planning – Find sequence of actions to reach a goal state / maximize utility.** *E.g., Navigate through an 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)

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 direct effects of its actions are?

Factory Floor

Poker

Chess

Doctor Diagnosis/Treatment

Clicker Question: Chess and Poker

Stochastic if the answer to at least one of these is "No"

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

A. Poker and Chess are both stochastic
B. Chess is stochastic and Poker is deterministic
C. Poker and Chess are both deterministic
D. Chess is deterministic and Poker is stochastic
E. Quit trying to make me think about stuff

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

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

R&R systems we'll cover in this course

		Environment	
Problem		**Deterministic**	**Stochastic**
Static	Constraint Satisfaction	*Variables + Constraints* Search Arc Consistency Local Search	
	Query	*Logics* Search	*Bayesian (Belief) Networks* Variable Elimination
Sequential	Planning	*STRIPS* Search	*Decision Networks* Variable Elimination

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:

- Static (constraints, query) vs. Sequential (planning)
- Deterministic versus stochastic domains

Some other important dimensions of complexity:

- Explicit states, features/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 features/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}

How many states?

Relations

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

University Example

Registered(S,C) = { T, F } ← relationship

Students = { s1, s2, s3, s4 } ← objects

Courses = { c1, c2, c3 } ← objects

Number of propositions:

Number of possible states:

Relations (cont.)

One binary relation, ***likes(x,y)***, and 9 individuals.
How many states?

A. 81^2
B. 2^9
C. 2^{81}
D. 10^9
E. 42

Flat or hierarchical

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

- One level of abstraction: **flat**
- Multiple levels of abstraction: **hierarchical**
- *Example: Planning a trip to a resort in Cancun, Mexico*

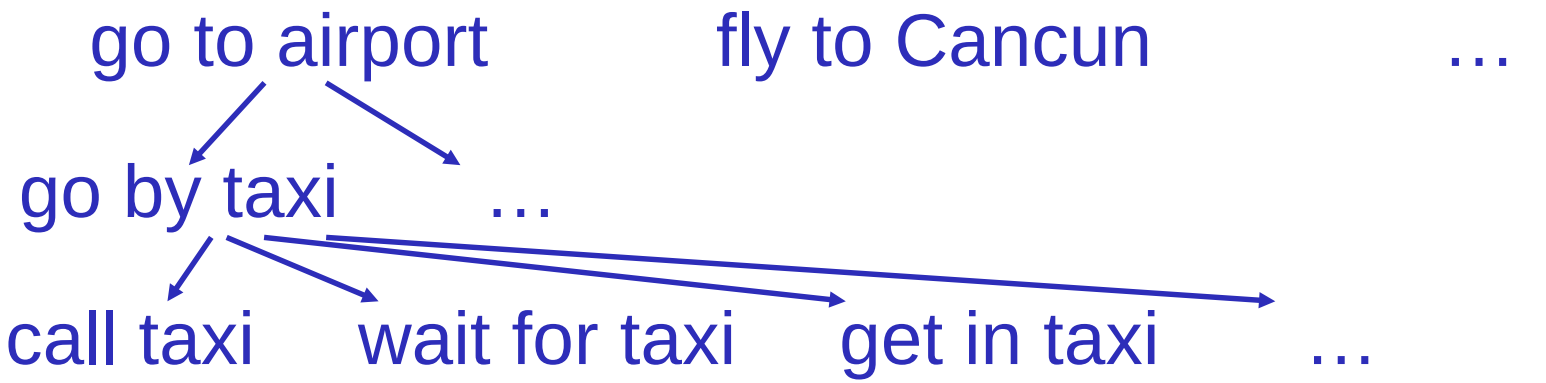

Knowledge given vs. knowledge learned from experience

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

OR

- 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*
- ***But Espresso is faster to make***

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

- **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) **preferences**
- **Single-agent** vs. multi-agent

not in this course