Representational Dimensions

CPSC 322 Lecture 2

January 6, 2006

Lecture Overview

Recap from Last Lecture

Representation

An Overview of This Course

Further Dimensions of Representational Complexity

Essentials

- Course web site: http://cs.ubc.ca/~kevinlb/teaching/cs322
  - This is where most information about the course will be posted, most handouts (e.g., slides) will be distributed, etc.
- Textbook: *Computational Intelligence, 2nd Edition*, by Poole, Mackworth and Goebel. Still under development; electronic version only.
- WebCT: used for textbook, discussion board
  - Use the discussion board for questions about assignments, material covered in lecture, etc, rather than email
  - If you do use email, you'll just be asked to repost to the message board, and your answer will take longer!
- Office hours:
  - Kevin Leyton-Brown: Thursdays 2:15 — 3:45, CICSR 185
  - Wei Li: Monday 10:00 – 11:00, Learning Center
  - Mark Crowley: Tuesday 1:00 – 2:00, Learning Center
  - David Thompson: TBA, Learning Center

Agents acting in an environment

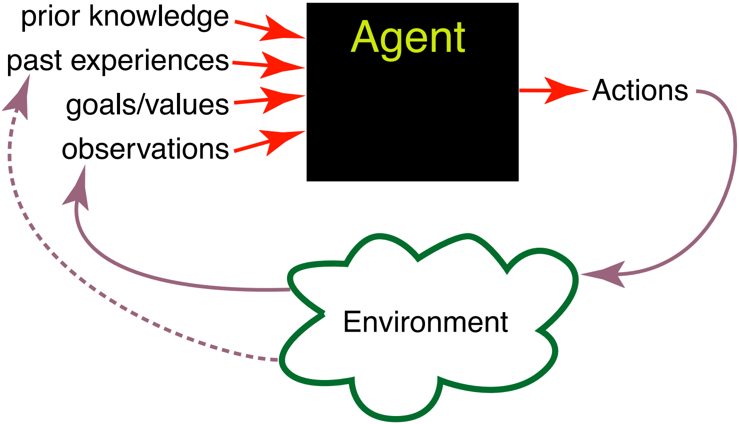

Representation

It turns out that when we want to think clearly and precisely about action, representation is critical:

- What different configurations can the world be in, and how do we denote them in a computer?
- What sorts of beliefs can we have about what configuration the world is in, and are these beliefs certain?
- How would the world be changed if we were to take some given action: what are the system dynamics?

What do we want from a representation?

We want a representation to be:

- rich enough to express the knowledge needed to solve the problem.
- as close to the problem as possible: compact, natural and maintainable.
- amenable to efficient computation; able to express features of the problem we can exploit for computational gain.
- learnable from data and past experiences.
- able to trade off accuracy and computation time.

Representation and Reasoning System

Problem $\Longrightarrow$ representation $\Longrightarrow$ computation

A representation and reasoning system (RRS) consists of

- Language to communicate with the computer.
- A way to assign meaning to the symbols.
- Procedures to compute answers or solve problems.

Example RRSs:

- Programming languages: Fortran, C++,...
- Natural Language

We want something between these extremes.

Overview of this course

This course will emphasize two main themes:

Reasoning

How should an agent act given the current state of the world and its goals?

Representation

How should the world be represented in order to help an agent to reason effectively?

Representations considered this course

Furthermore, the course will consider three main representational dimensions:

1. Static vs. sequential domains
2. Deterministic vs. stochastic domains
3. Single-agent vs. multiagent domains

1. Static vs. Sequential Domains

How many decisions does the agent need to make?

- The agent needs to take a single action
  - solve a Sudoku
  - diagnose a patient with a disease
- The agent needs to take a sequence of actions
  - navigate through an environment to reach a goal state
  - bid in online auctions to purchase a desired good

2. Deterministic vs. Stochastic Domains

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

- Is the environment deterministic or stochastic?
- Is the agent's knowledge certain or uncertain?

In past years, CPSC 322 has covered logic, while CPSC 422 has introduced probability

- this year I'll introduce both representational families in the same course
- 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.

3. Single-agent vs. Multiagent domains

One final representational question is whether the environment includes 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
- Agents can have their own goals: cooperative, competitive, or a bit of both

We'll end the course by looking at multiagent systems

Modules we'll cover in this course

1. Making single decisions in deterministic environments
   - Search, CSPs
2. Making sequential decisions in deterministic environments:
   - Planning
3. Richer representations in deterministic environments:
   - Logic
4. Making single decisions in stochastic environments:
   - Bayes Nets, Influence Diagrams
5. Making sequential decisions in stochastic environments:
   - MDPs
6. Multiagent systems
   - Zero-sum games; Nash equilibria

Dimensions of Representational Complexity

We've already discussed:

- Static versus sequential domains
- Deterministic versus stochastic domains
- Single-agent versus multiagent 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
- Perfect rationality versus bounded rationality

Explicit State or propositions or relations

How do we model the world?

- 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 features.
  - 30 binary features can represent $2^{30} = 1,073,741,824$ states.
- Features can be described in terms of objects and relationships.
  - There is a feature for each relationship on each tuple of individuals.
  - 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 system at one level of abstraction: flat
- You can model the system at multiple levels of abstraction: hierarchical
- Example: Planning a trip from here to a resort in Cancun, Mexico

Knowledge given versus knowledge learned from experience

How much do we know about the world in advance?

- The agent is provided with a model of the world before it starts to act
- The agent must 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

If an agent doesn't want to achieve anything, it has no reason to act. How do we represent an agent's desire(s)?

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

Perfect rationality versus bounded rationality

We've defined rationality as an abstract ideal. Is the agent able to live up to this ideal?

- Perfect rationality: the agent can derive what the best course of action is.
- Bounded rationality: the agent must make good decisions based on its perceptual, computational and memory limitations.

Modules we'll cover in this course

1. Making single decisions in deterministic environments
   - Search, CSPs
2. Making sequential decisions in deterministic environments:
   - Planning
3. Richer representations in deterministic environments:
   - Logic
4. Making single decisions in stochastic environments:
   - Bayes Nets, Influence Diagrams
5. Making sequential decisions in stochastic environments:
   - MDPs
6. Multiagent systems
   - Zero-sum games; Nash equilibria