Representational Dimensions

CPSC 322 – Intro 2

Textbook §1

Lecture Overview

1. Recap from Last Lecture
2. What is an Agent?
3. Representation
4. An Overview of This Course
5. Further Dimensions of Representational Complexity

Essentials

- Course web site: `http://www.ugrad.cs.ubc.ca/~cs322`
  - This is where most handouts and slides will be distributed
- Textbook: *Computational Intelligence, 2nd Edition*, by Poole, Mackworth and Goebel. Still under development; electronic version only.
- WebCT: used for textbook, discussion board, solutions to assignments
  - Use the discussion board for questions about assignments, material covered in lecture, etc, rather than email
  - Use email for private questions (e.g., health problems)
- You get 4 late days for assignments
- Assignment 0 is due on Monday, no late days
  - submit it using WebCT.

Lecture Overview

1. Recap from Last Lecture
2. What is an Agent?
3. Representation
4. An Overview of This Course
5. Further Dimensions of Representational Complexity

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
- It can make observations
  - perhaps imperfectly
- It is able to act
  - perhaps within constraints
- It has goals or preferences
- It may have prior knowledge or beliefs, and some way of updating beliefs based on new experiences

Agents acting in an environment

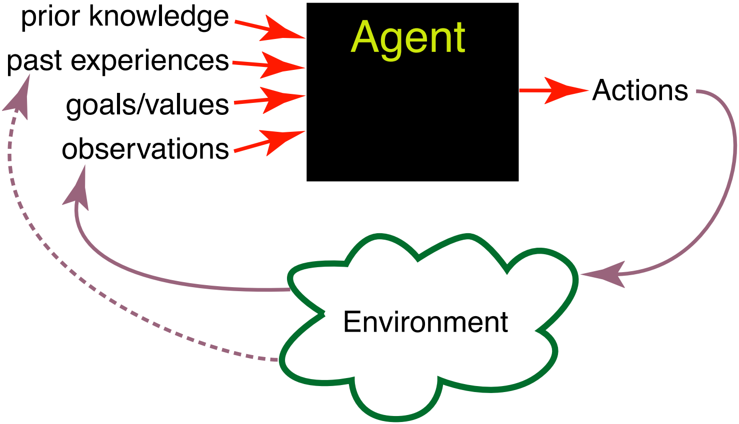

Examples

Which of these things is an agent, and why or why not?

- A soccer-playing robot?

Examples

Which of these things is an agent, and why or why not?

- A soccer-playing robot?
- A rock?

Examples

Which of these things is an agent, and why or why not?

- A soccer-playing robot?
- A rock?
- A Google web crawler?

Examples

Which of these things is an agent, and why or why not?

- A soccer-playing robot?
- A rock?
- A Google web crawler?
- A thermostat?

Examples

Which of these things is an agent, and why or why not?

- A soccer-playing robot?
- A rock?
- A Google web crawler?
- A thermostat?
- A dog?

Examples

Which of these things is an agent, and why or why not?

- A soccer-playing robot?
- A rock?
- A Google web crawler?
- A thermostat?
- A dog?
- A car?

Assignment 0

- Your first assignment asks you to find two examples of fielded AI agents, and to explain some high-level details about how they work.
  - you get bonus marks if you're the only one in the class to describe a given application.
- The assignment is available from the course web page
  - Remember: `http://www.ugrad.cs.ubc.ca/~cs322`
- It's due on Monday, and you can't use late days
  - I'll show some pictures and videos of cool applications in that class, and will give you an opportunity to discuss the applications that you discovered

Lecture Overview

1. Recap from Last Lecture
2. What is an Agent?
3. Representation
4. An Overview of This Course
5. Further Dimensions of Representational Complexity

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?

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?

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.

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.

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.

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.

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

- A language in which a model of the world can be described symbolically ("syntax")

Representation and Reasoning System

Problem $\Longrightarrow$ representation $\Longrightarrow$ computation

A representation and reasoning system (RRS) consists of

- A language in which a model of the world can be described symbolically ("syntax")
- A way to assign meaning to the symbols ("semantics")

Representation and Reasoning System

Problem $\Longrightarrow$ representation $\Longrightarrow$ computation

A representation and reasoning system (RRS) consists of

- A language in which a model of the world can be described symbolically (“syntax”)
- A way to assign meaning to the symbols (“semantics”)
- Computational procedures to compute answers or solve problems

Representation and Reasoning System

Problem $\Longrightarrow$ representation $\Longrightarrow$ computation

A representation and reasoning system (RRS) consists of

- A language in which a model of the world can be described symbolically ("syntax")
- A way to assign meaning to the symbols ("semantics")
- Computational procedures to compute answers or solve problems

Example RRSs:

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

Representation and Reasoning System

Problem $\Longrightarrow$ representation $\Longrightarrow$ computation

A representation and reasoning system (RRS) consists of

- A language in which a model of the world can be described symbolically ("syntax")
- A way to assign meaning to the symbols ("semantics")
- Computational procedures to compute answers or solve problems

Example RRSs:

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

This course will explore alternatives that fall between these extremes.

Lecture Overview

1. Recap from Last Lecture
2. What is an Agent?
3. Representation
4. An Overview of This Course
5. Further Dimensions of Representational Complexity

Overview of this course

This course will emphasize two main themes:

Reasoning

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

Representation

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

Representations considered this course

Furthermore, the course will consider two main representational dimensions:

1. Deterministic vs. stochastic domains
2. Static vs. sequential domains

1. 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?

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.

2. Static vs. Sequential Domains

How many actions does the agent need to select?

- The agent needs to take a single action
  - solve a Sudoku
  - diagnose a patient with a disease

2. Static vs. Sequential Domains

How many actions does the agent need to select?

- 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. Static vs. Sequential Domains

How many actions does the agent need to select?

- 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
- Important caveat:
  - in deterministic domains, the distinction between static and sequential settings may seem somewhat artificial
    - we can redefine actions (e.g., fill in individual numbers in the Sudoku vs. solving the whole thing)
    - indeed, some of the same techniques work in both settings
  - the same cannot be said about stochastic domains

Modules we'll cover in this course

1. Making single and sequential decisions in deterministic environments
   - Search, CSPs, Planning
2. Richer representations in deterministic environments:
   - Logic
3. Making single decisions in stochastic environments:
   - Bayes Nets, Influence Diagrams
4. Making sequential decisions in stochastic environments:
   - Influence Diagrams, MDPs

Lecture Overview

1. Recap from Last Lecture
2. What is an Agent?
3. Representation
4. An Overview of This Course
5. Further Dimensions of Representational Complexity

Dimensions of Representational Complexity

We've already discussed:

- Deterministic versus stochastic domains
- Static versus sequential domains

Dimensions of Representational Complexity

We've already discussed:

- Deterministic versus stochastic domains
- Static versus sequential domains

Some other important dimensions of complexity:

- 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, described as assignments of values to one or more variables (features).
  - 30 binary features can represent $2^{30} = 1,073,741,824$ states.

Explicit State or propositions or relations

How do we model the world?

- You can enumerate the states of the world, described as assignments of values to one or more variables (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

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 and sequential decisions in deterministic environments
   - Search, CSPs, Planning
2. Richer representations in deterministic environments:
   - Logic
3. Making single decisions in stochastic environments:
   - Bayes Nets, Influence Diagrams
4. Making sequential decisions in stochastic environments:
   - Influence Diagrams, MDPs