CS324-Artificial Intelligence

Lecture 3: Intelligent Agents

Waheed Noor

Computer Science and Information Technology,
University of Balochistan,
Quetta, Pakistan

Outline

Outline

What is a good behavior?

We know!

Agent is in the environment, generating sequence of actions according to the percepts it receives. The action sequence is changing the *states* of the environment generating sequence of states.

Good behavior

The agent is said to have good behavior if the state sequence is desirable, and evaluation of desirable state sequence is done by a *performance measure*

Performance measure

It depends on the task and agents, therefore you should design it according to the circumstances. It is not easy!

General Rule

Recall the vacuum-cleaner agent example, and design a performance measure.

General Rule

It is better to design performance measures according to what one actually wants in the environment, rather than according to how one thinks the agent should behave.

General Rule

Recall the vacuum-cleaner agent example, and design a performance measure.

General Rule

It is better to design performance measures according to what one actually wants in the environment, rather than according to how one thinks the agent should behave.

How to Achieve Rationality

At any given time

1. Performance measure (Success criteria)
2. Agent's prior knowledge of the environment
3. Set of actions
4. Percept sequence (Up to given time)

Rational Agent

What we already know about rational agent

That acts to achieve best outcome or best expected outcome in the presence of uncertainty. Act sufficiently, if not optimally, in all situations

Rational Agent: In Technical Terms

A rational agent should select an action for each possible percept sequence that maximizes its performance measure given to date percept sequence, and whatever built-in knowledge agent has.

Is the vacuum-cleaner agent is rational?

Rational Agent

What we already know about rational agent

That acts to achieve best outcome or best expected outcome in the presence of uncertainty. Act sufficiently, if not optimally, in all situations

Rational Agent: In Technical Terms

A rational agent should select an action for each possible percept sequence that maximizes its performance measure given to date percept sequence, and whatever built-in knowledge agent has.

Is the vacuum-cleaner agent is rational?

Rational Agent

What we already know about rational agent

That acts to achieve best outcome or best expected outcome in the presence of uncertainty. Act sufficiently, if not optimally, in all situations

Rational Agent: In Technical Terms

A rational agent should select an action for each possible percept sequence that maximizes its performance measure given to date percept sequence, and whatever built-in knowledge agent has.

Is the vacuum-cleaner agent is rational?

Vacuum-Cleaner Agent

Example (Assume)

Performance measure: One point reward for each clean square at each time step (say lifetime = 1000 time step)

Knowledge of the environment: is given, but agent does not know where the dirt is, and clean square will remain clean.

Set of Actions: Left, right, and clean

Perception: Agent can correctly perceive its location, and identify dirt.

- Is this agent rational?
- Can you imagine different circumstances where this agent is not rational.

Vacuum-Cleaner Agent

Example (Assume)

Performance measure: One point reward for each clean square at each time step (say lifetime = 1000 time step)

Knowledge of the environment: is given, but agent does not know where the dirt is, and clean square will remain clean.

Set of Actions: Left, right, and clean

Perception: Agent can correctly perceive its location, and identify dirt.

- Is this agent rational?
- Can you imagine different circumstances where this agent is not rational.

Omniscience Vs. Rationality

- Omniscience agent knows the actual outcome of its action, i.e., perfection (Impossible in reality)
- Omniscience maximizes actual performance, while rationality maximizes expected performance.
- Rational agent action depends on the percept sequence to date, it also needs to perform **information gathering**.
- **Exploration** is also necessary when the environment is unknown to the rational agent.
- The rational agent may also be **learning** from its experience.
- The rational agent when getting experience, it should also become autonomous, i.e., compensating partial or incorrect knowledge.

Outline

Task Environment

To build an agent, we should first describe *task environment* that is actually the description of performance measure, environment, actuators, and sensors (PEAS).

Agent Type	Performance Measure	Environment	Actuators	Sensors
Taxi driver	Safe, fast, legal, comfortable trip, maximize profits	Roads, other traffic, pedestrians, customers	Steering, accelerator, brake, signal, horn, display	Cameras, sonar, speedometer, GPS, odometer, accelerometer, engine sensors, keyboard

Figure : PEAS description for Autonomus Car Agent

Class Activity: What will be the task environment for a left luggage detection agent on the air port?

Outline

Structure of Agents

Agent structure is consist of:

- Its architecture, i.e., computing device, actuators and sensors
- Agent program that implements agent function

Softbot

You can also imagine your agent as an intelligent program, called software robot or softbot. There are plenty of examples such as a website with recommendation agent or scanning news on the Internet and displaying only the one user is interested in.

Explore: Verbot `http://verbots.sourceforge.net/`

Outline

Types of Intelligent Agents

Agents

- Simple Reflex Agents
- Model-Based Reflex Agents
- Goal-Based Agents
- Utility-Based Agents

Simple Reflex Agents

Definition

These agents select actions on the basis of the current percept, ignoring the rest of the percept history.

Example

Vacuum-Agent
Autonomus Car Agent

Advantage: Simple
Disadvantage: Needs fully observable environment.

Model-Based Reflex Agents

Definition

An agent that keeps track of part of the environment that it can not observe at a particular time using its best guess. This information is called the internal state and maintained by an internal model.

Example

In Autonomous car example, breaking if car is closure than a moment ago, if it can not observe the red lights.

Updating internal state information encoded in the agent program depends on:

- How the world evolves without the agent.
- Agent's action affecting the world.

Goal-Based Agents

Definition

With current state description, the agent needs some sort of goal information that describes situations that are desirable. They are less efficient but more flexible.

Example

Game playing agent needs to consider a large sequence of percept before actually applying an action to find a way to win the game. Autonomous car agent may consider the sequence of percept for possible routes to reach the destination.

Searching and *planning* provides some techniques to find the action sequence that achieve the goal of the agent.

Reflex Vs. Goal-Based Agents

Rules Vs. Decision Making

Simple reflex agents action is based on rules that maps percept to actions, where as goal based agents make decisions by reasoning " What will happen if I do this."

Example

The reflex agent brakes when it sees brake lights. A goal-based agent, in principle, could reason that if the car in front has its brake lights on, it will slow down. Given the way the world usually evolves, the only action that will achieve the goal of not hitting other cars is to brake.

Discussion: How both agent would behave if it starts raining in the above example.

Utility Based Agents

Example

Consider, a goal-based agent will get the taxi to its destination, achieving the goal even if some routes are block. But how about safer, more cheaper, and quicker...

Goal-based can make distinction between good vs. bad, but can not answer how much.

Definition

A rational utility-based agent chooses the action that maximizes the expected utility of the action outcomes that is, the utility the agent expects to derive, on average, given the probabilities and utilities of each outcome. The agent should include an explicit internal utility function and agent should choose action with agreement of the performance measure and utility function.

Outline

Assignment 1

Q1: Fully observable, partially observable, deterministic, stochastic, discrete and continuous are the characteristics of an environment, describe them.

Q2: Today AI is doing many things around us. Examples include game playing such as IBM's Deep Blue, Logistic Planning, Robotics, etc. Describe two such real world applications of AI in detail of your own choice.

Q3: Do you see the application of AI in following:

1. Supper market bar code scanner
2. Web search engines
3. Voice-activated telephone menus
4. Internet routing algorithms that respond dynamically to the state of the network.

Important

Important Note

- Due date: 15th-April-2016
- Explain everything in your own words, copy if proved will penalize your marks.
- You have to defend you work in a presentation therefore better to make presentation slides as well.

References I

Stuart Russell and Peter Norvig.
Artificial Intelligence: A Modern Approach.
Pearson Education, Inc, USA, second edition, 2003.

George F. Luger.
Artificial Intelligence: Structures and Strategies for Complex Problem Solving.
Addison-Wesley Publishing Company, USA, 6th edition, 2008.