CENG 466
Artificial Intelligence

Lecture 7

Knowledge and Reasoning

Topics

- Knowledge-based Agents
- Knowledge Base
- Knowledge Representation
- Knowledge-based Agent Description Levels
- Inference
- Example

Intelligent Agents

- An agent is something that perceives and acts in an environment
- An ideal agent always takes actions that maximizes its performance
- An agent adopts a goal and searches the best path to reach that goal

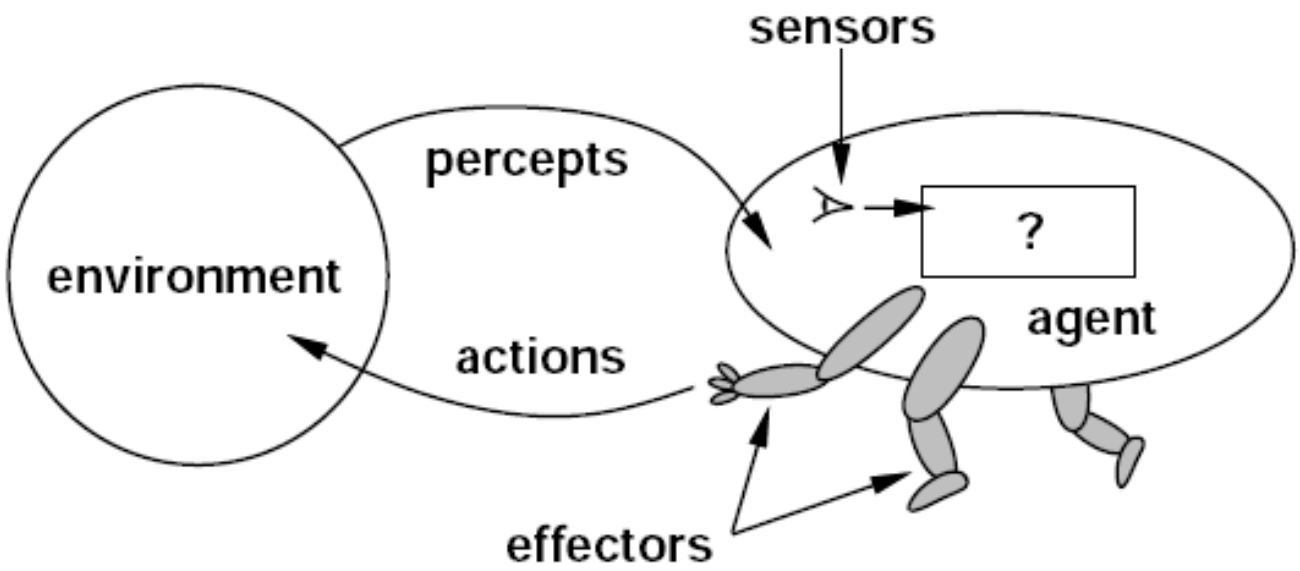

Intelligent Agents

- An agent that has goals and searches for solutions to the goals can do better than one that just reacts to its environment.
- Therefore, an agent should have some knowledge about its goal, its environment, and its actions.
- In addition, the agent should be capable of general logical reasoning.
- It uses logical reasoning to:
  - describe the world as new percepts arrive,
  - decide the action to achieve its goals.

Knowledge-based Agents

- knowledge-based agents are agents which
  - can be seen as *knowing* about their world,
  - and *reasoning* about their possible action.

Knowledge Base (I)

- The central component of a knowledge-based agent is its **knowledge base**, or KB.
- Informally, a knowledge base is a set of facts about the world.
- Each fact in a KB is called a **sentence**.

Knowledge Base (II)

- The sentences are expressed in a language called a **knowledge representation language**.
- There must be a way to add new sentences to the knowledge base, and a way to query what is known.
- The standard names for these tasks are TELL and ASK, respectively.

Knowledge Base (III)

- The fundamental requirement on TELL and ASK is that when one ASKS a question of the KB, the answer should follow from what has been TOLD to the KB previously.

- Determining what follows from what the KB has been TOLD is the job of the **inference** mechanism, the other main component of a knowledge-based agent.

How does the Knowledge-based Agent Act?

- The agent maintains a *KB*, which contains some **background knowledge.**
- **The agent** TELLS the knowledge base what it perceives.
- Then, it ASKS the KB what action it should perform.
- To answer this query, logical reasoning is used to prove which action is better than all others, given:
  - what the agent knows
  - what its goals are.
- The agent then performs the chosen action.

Example: Knowledge Base

- Wumpus was an early computer game, based on an agent who explores a cave consisting of rooms connected by passageways.

Example: The Wumpus

- Somewhere in the cave is a beast (the wumpus) that eats anyone who enters its room.
- Some rooms have bottomless pits that will trap anyone who enters into these rooms (except for the wumpus, who is too big to fall in).
- There are heaps of gold at some of the rooms.

Wumpus Environment

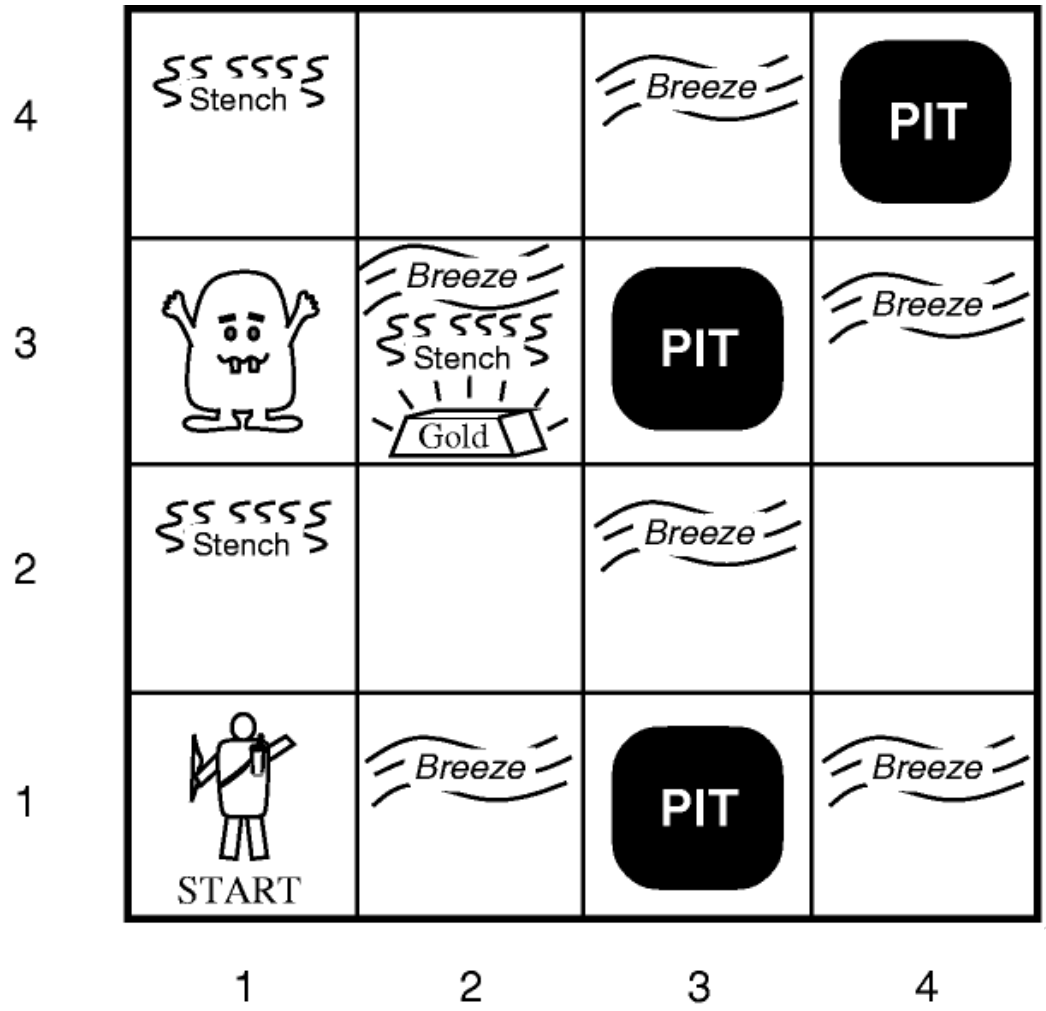

Wumpus' Knowledge Base (I)

- In the room with the wumpus and in the directly (not diagonally) adjacent rooms the agent will perceive a stench.
- In the rooms directly adjacent to a pit, the agent will perceive a breeze.
- In the room where the gold is, the agent will perceive a glitter.
- When an agent walks into a wall, it will perceive a bump.
- When the wumpus is killed, it gives out a scary scream that can be perceived anywhere in the cave.

Wumpus' Knowledge Base (II)

- The percepts will be given to the agent in the form of a list of five symbols;
- for example:

  [stench=Yes, breeze=Yes, glitter=Yes, bump=No, scream=No]

- The agent *cannot* perceive its own location.

Wumpus' Knowledge Base (III)

- Actions :
  - go forward, turn right by 90°, turn left by 90°, *Grab* (to pick up an object), *Climb* (to leave the cave; it is used only when the agent is in the start room), *Shoot* (to fire an arrow in a straight line in the direction the agent is facing.
  - The arrow continues until it either hits and kills the wumpus or hits the wall.
- The agent only has one arrow, so only the first *Shoot* action has any effect.
- The agent dies if it enters a room containing a pit or a live wumpus.
- It is safe (but smelly) to enter a room with a dead wumpus.

Wumpus in Action (I)

- The agent's goal is to find the gold and bring it back to the start as quickly as possible, without getting killed.
- An agent can do better by memorizing the sequence of actions that it takes.
- Assume a 4 x 4 grid surrounded by walls.
- The agent always starts in the room labeled (1,1), facing toward the right.

Wumpus in Action (II)

- An example scenario for the agent is as follows:
- From the fact that there was no stench or breeze in [1,1], the agent can infer that [1,2] and [2,1] are free of dangers.
- The agent marks with an *OK* to indicate this.
- From the fact that the agent is still alive, it can infer that [1,1] is also *OK*.
- Agent will only move into a room that it knows is *OK*.
- The agent decides to move forward to [2,1],
- The agent perceives a breeze in [2,1], so there must be a pit in a neighboring room, either [2,2] or [3,1].

Knowledge-based Agent Levels (I)

- A knowledge-based agent can be described at three levels:
- **The knowledge level**: we can describe the agent by saying what it knows.
  - For example, an automated taxi might be said to know that a bridge links street A and street B.
  - If TELL and ASK work correctly, then most of the time we can work at the knowledge level and not worry about lower levels.
- **The logical level**: the knowledge is encoded into sentences.
  - For example, the taxi might be described as having the logical sentence *Links(street A, street B)* in its knowledge base.

Knowledge-based Agent Levels (II)

- **The implementation level:** It is the level at which there are physical representations of the sentences at the logical level.
- For example, a sentence such as *Links(Street A, Street B)* could be represented in the KB by a "1" entry in a three-dimensional table indexed by road links and location pairs.
- The choice of implementation is very important to the efficient performance of the agent, but it is irrelevant to the logical level and the knowledge level.

Designing Knowledge Bases

- It is possible to understand the operation of a KB agent in terms of what it knows.
- It is possible to construct a KB agent by TELLing it what it needs to know. The agent's initial program, before it starts to receive percepts, is built by adding one by one the sentences.
- These sentences are the designer's knowledge of the environment.
- This is called the declarative approach to system building.
- Also, we can design learning mechanisms that output general knowledge about the environment given a series of percepts.
- By connecting a learning mechanism to a knowledge-based agent, we can make the agent fully autonomous.

Knowledge Representation

- The goal of **knowledge representation** is to express knowledge in computer-usable form.
- This format is used to help agents perform well.
- A knowledge representation language is defined by two aspects:
  - Syntax
  - Semantics

Knowledge Representation Syntax

- **The syntax** of a language describes the possible configurations that make sentences.

- A syntax is a set of rules that define the patterns of the sentences in a language.

Knowledge Representation Semantics

- The **semantics** determines the facts in the world to which the sentences refer.
- Without semantics, a sentence is just an arrangement of marks on a page.
- With semantics, each sentence makes a claim about the world.
- Semantics is the meaning or the concept that we want to store in a KB.

Inference (I)

- The terms "reasoning" and "inference" are generally used to cover any process by which conclusions are reached.
- **Logical inference or deduction** is a process that implements the entailment relation between sentences.
- P |= Q means P entails Q

Inference (II)

- If a sentence A entails B it means:
  - If A is true then B is true
  - E.g.
  - A: The temperature is 100 degrees
  - B: Water is boiling

Inference (III)

- Given a knowledge base *KB* a new sentence *a*, can be generated,

  OR

- Given a knowledge base *KB* and another sentence *a*, the inference system can report whether or not *a* is entailed by *KB*.

- An inference procedure that generates only entailed sentences is called **sound** or **truth-preserving**.

Knowledge Representation Methods

- Propositional Logic
- Predicate Logic

Questions?