CPS331 Lecture: Intelligent Agents   last revised July 25, 2018

Objectives:

1. To introduce the basic notion of an agent
2. To discuss various types of agents

Materials:

1. Projectable of Russell and Norvig (2015) Figure 2.1
2. Projectable of Russell and Norvig (2015) Figure 2.4
3. Projectable of Russell and Norvig (2015) Figure 2.5
4. Projectable of Russell and Norvig (2015) Figure 2.6
5. Projectable of Russell and Norvig (2015) Figures 2.9, 2.10 together
6. Projectable of Russell and Norvig (2015) Figures 2.11, 2.12 together
7. Projectable of Russell and Norvig figure 2.13
8. Projectable of Russell and Norvig figure 2.14

I. **Introduction**

   A. The notion of an "agent" provides a framework for thinking about intelligent systems.

   B. Definition: An agent is a computational system that performs some task with a measure of autonomy. An agent has appropriate means of perceiving its environment (sensors) and acting upon its environment (effectors).

      1. Autonomy is a key characteristic - what distinguishes an "agent" from a more conventional program such as a word processor or web-browser.

         a) An agent is typically given a task to perform on behalf of a user - either a human user, or another agent that uses its services to accomplish its goals. It does so with minimal interaction with that user - though it may eventually report its discoveries back to its user.

            Examples: web search agents, GPS

b) As noted, in many cases agents may need to interact with other agents to perform their task.

2. An agent typically functions in a complex and changing environment, and is expected to adjust its behavior appropriately.

3. We can think of an agent as having <u>sensors</u> which allow it to perceive its environment and <u>actuators</u> which allow it to change its environment.

   PROJECT Russell and Norvig Figure 2.1

   a) We call the information it is receiving from it sensors at a given moment of time its <u>percepts</u> and what is doing through its actuators its <u>actions</u>.

   b) We speak of the information it receives from its sensors over a period time as its <u>percept sequence</u>.

4. An agent can be thought of as realizing a function that maps a percept sequence into the action to be taken at a given instant of time.

   Percept sequence ----> AGENT ----> Action

   It is the task of an AI program to implement a desired agent function.

5. Russell and Norvig suggest that an agent can be described using the acronym PEAS (Performance Measure, Environment, Actuators, Sensors). For example, the desired behavior of an autonomous taxi might be described as follows;

   PROJECT Russell and Norvig figure 2.4

6. A figure that appears in their book shows how the requirements for various kinds of AI systems can be modeled using PEAS

   PROJECT Russell and Norvig Figure 2.5

C. Two broad types of agents

1. Pure software agents - "softbots"

2. Hardware agents - "robots"

3. We will focus on the former for most of the course.

D. Some examples:

ASK

1. Softbots:

   a) spam filters and email distribution agents that route incoming email to appropriate boxes and/or notify the user of email from an "interesting" sender

   b) detection of viruses and other kinds of attack. (Actually, for that matter a virus or worm is also a sort of softbot!)

   c) algorithmic stock trading software

2. Robots:

   a) NASA has used robotic rovers on the moon and Mars. While humans are still involved in making decisions about what to explore, movement needs to be autonomous due to long signal propagation times.

   b) Police robots: Google search on "police robot"

E. Actually "agent" is an overarching term - an agent is a system that combines appropriate AI methods to fulfill its particular role.

   1. The notion of an agent has emerged in AI as an alternative to classic models of strong intelligence.

     One concept that underlies this is a notion first espoused in a 1986 book by Marvin Minsky called *The Society of Mind* - which he states this way: "we'll show you that you can build a mind from many little parts, each mindless in itself" (quoted on Lucci p. 199)

       a) In essence, an agent implements a form of applied or weak AI.

       b) The argument in Minsky's book was that strong AI is essentially the product of a lot of weak AI agents working together.

   2. Some AI books use the notion of an "agent" as a unifying theme

     Example: the best-selling AI text (Russell and Norvig: *Artificial Intelligence - a Modern Approach*) includes the following phrase in its description: "It uses an intelligent agent as the unifying theme throughout"

II. **Properties of Environments**

An agent is designed to function in a particular type of environment. These environments can be characterized in a various ways. (we'll illustrate using different types of game, but the notion is general)

A. Fully observable versus partially-observable - do the agent's sensors provide access to the complete state of the environment (that is everything that is relevant to the agent's decision-making)?

  1. Observable: most board games

  2. Partially-observable: poker

B. Single Agent vs Multi-agent: Is the state of the environment entirely determined by an agents actions, or are there other agents that affect the environment, too? In the latter case, are the other agents competitive or cooperative or some combination of the two.

  1. Single agent: solitaire

  2. Multiple agents - cooperative: LaVida ropes activities

  3. Multiple agents - competitive: soccer, football, baseball etc.

  4. Multiple agents - mixed - driving. (Other drivers cooperate to avoid accidents, but may compete for a parking place.)

C. Deterministic vs stochastic: is everything that happens in the environment a result of the actions of the agents (note plural)?

Note: a partially-observable environment may appear stochastic even though it isn't!

  1. Deterministic: most board games (assuming we regard all the players as agents)

  2. Stochastc:

     a) backgammon (use of dice)

     b) poker (inaccessible therefore appears nondeterministic - and also depends on the randomness of the shuffling of the deck.)

D. Episodic vs sequential: does time divide into episodes involving the agent perceiving, deciding and acting, with the agent's decision based only on what it perceives in that episode, or do episodes depend on agents' actions in previous episodes - or is time is continuous? (Note: if a task is non-episodic then the agent doesn't have to "plan ahead")

1. Episodic: games are usually not episodic since they involve planning strategy. Tasks like part-picking on an assembly line are episodic.

2. Sequential: board games with any strategy aspect; ping-pong - continuous

E. Static vs dynamic: Can the environment change while the agent is deciding? (If the environment is static, the agent can sense the environment and then decide what to do, but if it is dynamic the agent needs to keep sensing while deciding

1. Static: most board game

2. Dynamic: capture the flag

F. Discrete vs continuous: is there a discrete set of possible percepts and actions?

1. Discrete: most board games

2. Continuous: most forms of athletics and games of skill

G. Known versus unknown: The agents state of knowledge about the "physics" of the environment - i.e. the effect of its actions. Note that this is not the same as whether or not the environment is observable - it has to do with the effects of actions, not the state of the environment.

1. Known: most board games

2. Unknown: when playing a new video game for the first time one may not know exactly how the system will respond to certain actions

H. Russell and Norvig illustrate these for a wide variety of types of problem

PROJECT Russell and Norvig Figure 2.6

III. Basic Types of Agents

A. Table-driven agents

1. We could imagine a table-driven agent being created to function in a static, episodic environment. Each decision the agent makes would be made by looking up its current percepts in a table of specifying the correct action to take for that percept.

2. Example: We might envision an agent the plays the game of tic-tac-toe being given a table showing the correct move to make for any possible board configuration. (There are only 5565 possible configurations, and the number can be reduced significantly by taking advantage of symmetries. Thus, this is is feasible)

3. Of course, such agents are rarely practical since they would require a table that covers every possible percept sequence. (E.g. you could imagine creating such an agent to play tic-tac-toe, where the agent senses the complete current state of the board - but a table-driven chess-playing agent would likely be unfeasible and a table-driven ping-pong playing agent would be impossible.

   In any case, we would not consider a program realized this way to be an example of artificial intelligence, though it may be a manifestation of the intelligence of the author who created the table!

B. Reflex agents

1. The agent bases its decision completely on its current percepts, but chooses an action by using some sort of computation

   PROJECT: Russell & Norvig figures 2.9 + 2.10 (together)

2. Example: A system that can classify images based on their what they contain - e.g. what you get when you search of images of something on Google.

C. Model-based agents

1. The agent maintains some internal state and bases its decisions on both its state and its current percepts. Thus, the agents are shaped by both its history and its model of the current state of the world.

   PROJECT: Russell & Norvig figures 2.11 + 2.12 (together)

2. Example: an email filter on a personal computer makes use of a model of what kind of mail does and does not interest its user

D. Agents with goals

1. In addition to having state, the agent has one or more goals that it is seeking to achieve, and bases its decisions on its goals, its state, and its current percepts.   (Thus, this kind of agent also uses some sort of model of the world)

   PROJECT: Russell & Norvig figure 2.13

2. Example: the planning component of a GPS system

E. Utility-based agents

1. The agent has multiple, possibly conflicting goals, and has some means of comparing the relative "goodness" of different states in order to handle tradeoffs between goals.   (Thus it also needs goals and an internal model.)

   (The technical term for "goodness" is utility.)

   PROJECT: Russell & Norvig figure 2.14

2. Example: a chess-playing program might actually work on several different goals:

a) Checkmate the opposing king

b) Keep its own king out of check

c) Capture an opponent's piece where possible

d) Keep its own pieces from being captured.

There is a hierarchy the goals - e.g. (a) trumps all others, and generally (but not always) the remaining goals are preferred in the order listed.

3. Example: the ordering of pages presented in response to a query by a search engine like Google

a) Relevance to the query

b) Google uses a "page rank" algorithm to give preference to pages that are deemed more reliable or useful on the basis of references to them from other pages.

c) The first results returned by a query often return references to a few pages from paying advertisers.

F. Additional Types of Agents (not mutually-exclusive with the above)

1. Learning Agents

a) For much of the early history of AI, the assumption was that the knowledge needed by an agent had to be built into it in some way.

b) Today, though, most AI systems make some use of machine learning - something we will discuss in the second half of the course. Learning can be used:

(1) To provide the knowledge the system needs to function initially. (E.g. systems that learn to classify images)

(2) To improve the performance of the system as it is used.

(3) Often both.

We will discuss this at some length in the second half of the course. In the first half, we will focus on "classical" AI

2. Mobile agents:

   a) General term for an agent that can move from place to place in accomplishing its goals

   b) May be mechanical mobility (e.g. a mobile robot)

   c) May be software mobility (e.g. an agent program that travels from computer to computer)

   Examples:

   ASK

   (1) Web-crawlers used by search engines

   (2) Various kinds of purchasing assistants that access and pool information from other sites (e.g. travel systems like Orbitz)

3. Interface agents: an agent that serves as a sort of personal assistant, carrying out assigned tasks (with some measure of autonomy) on behalf of a user

   Examples:

   ASK

   a) Email distribution systems

   b) Personal Calendar systems