Intelligent Agents in Artificial Intelligence

Rushikesh B. Gholap[1] Abhishek P. Gosavi[2] Sabaji S. Gawade[3] Ashwini D. Padekar[4]
[1,2,3]Student [4]Assistant Professor
[1,2,3,4]Department of Computer Engineering
[1,2,3,4]MGM's College of Engineering & Technology, Kamothe, Navi Mumbai, India

Abstract— Since we are in era of AI (artificial Intelligence). Were each and everything is hyped with AI though it doesn't make valuable difference in performance. But it is in developing stage and to get better result we have to give it some time to get better .AI at initial stage is always in learning stage. This paper is all about how AI works in certain environment. This paper completely depends on working of AI mechanism and various algorithms to make most effective AI technology that adapt very well to given environment.

Key words: Artificial Intelligence

I. Introduction

A rational agent is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome. All the skills needed for the Turing Test also allow an agent to act rationally. Knowledge representation and reasoning enable agents to reach good decisions. We need to be able to generate comprehensible sentences in natural language to get by in a complex society. We need learning not only for erudition, but also because it improves our ability to generate effective behavior. The rational-agent approach has two advantages over the other approaches. First, it is more general than the "laws of thought" approach because correct inference is just one of several possible mechanisms for achieving rationality. Second, it is more amenable to scientific development than are approaches based on human behavior or human thought. The standard of rationality is mathematically well defined and completely general, and can be "unpacked" to generate agent designs that provably achieve it. Human behavior, on the other hand, is well adapted for one specific environment and is defined by, well, the sum total of all the things that humans do.

II. Agents & Environments

An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators. A human agent has eyes, ears, and other organs for sensors and hands, legs, vocal tract, and so on for actuators. A robotic agent might have cameras and infrared range finders for sensors and various motors for actuators. A software agent receives keystrokes, file contents, and network packets as sensory inputs and acts on the environment by displaying on the screen, writing files, and sending network packets.

We use the term percept to refer to the agent's perceptual inputs at any given instant. "An agent's percept sequence is the complete history of everything the agent has ever perceived." In general, an agent's choice of action at any given instant can depend on the entire percept sequence observed to date, but not on anything it hasn't perceived. By specifying the agent's choice of action for every possible percept sequence, we have said more or less everything there is to say about the agent. Mathematically speaking, we say that an agent's behavior is described by the agent function that maps any given percept sequence to an action.

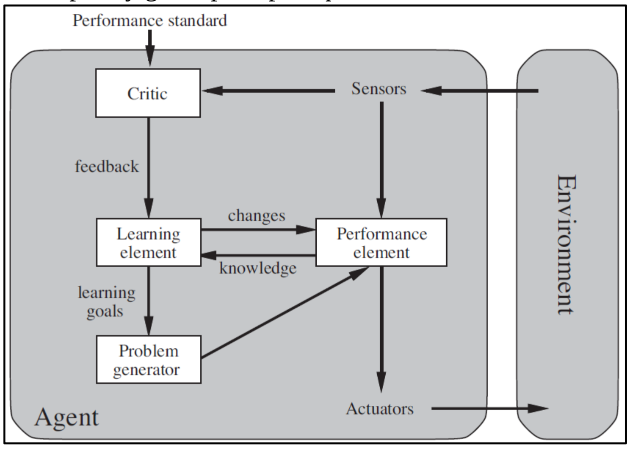

Fig. 1: Agents Interact with Environments through Sensors and Actuators

III. Good Behavior: The Concept of Rationality

A rational agent is one that does the right thing—conceptually speaking, every entry in the table for the agent function is filled out correctly. Obviously, doing the right thing is better than doing the wrong thing, but what does it mean to do the right thing? We answer this age-old question in an age-old way: by considering the consequences of the agent's behavior. When an agent is plunked down in an environment, it generates a sequence of actions according to the percepts it receives. This sequence of actions causes the environment to go through a sequence of states. If the sequence is desirable, then the agent has performed well. This notion of desirability is captured by a performance measure that evaluates any given sequence of environment states.

1) Rational at any given time depends on four things:
   - The performance measure that defines the criterion of success.
   - The agent's prior knowledge of the environment.
   - The actions that the agent can perform.
   - The agent's percept sequence to date.
2) This leads to a definition of a rational agent:

For each possible percept sequence, a rational agent should select an action that is expected to maximize its performance measure, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has.

IV. Specifying the Task Environment

In our discussion of the rationality of the simple vacuum-cleaner agent, we had to specify the performance measure, the environment, and the agent's actuators and sensors. We group all these under the heading of the task environment. For the acronymically minded, we call this the PEAS (Performance, Environment, Actuators, Sensors) description. In designing an agent, the first step must always be to specify the task environment as fully as possible. The vacuum world is a simple example; let us consider a more complex problem: an automated taxi driver. We should point out, before the reader becomes alarmed, that a fully automated taxi is currently somewhat beyond the capabilities of existing technology. The full driving task is extremely open-ended.

1) Example

– PEAS description of the task environment for an automated taxi.
– Agent Type: Taxi driver.
– Performance Measure: Safe, fast, legal, comfortable trip, maximize profits.
– Environment: Roads, other traffic, pedestrians, customers
– Actuators: Steering, accelerator, brake, signal, horn, display.
– Sensors: Cameras, sonar, speedometer, GPS, odometer, accelerometer, engine sensors, keyboard.

2) Properties of Task Environments

The range of task environments that might arise in AI is obviously vast. We can, however, identify a fairly small number of dimensions along which task environments can be categorized. These dimensions determine, to a large extent, the appropriate agent design and the applicability of each of the principal families of techniques for agent implementation each kind of environment.

– Fully observable vs. partially observable
– Single agent vs. multi-agent
– Static vs. dynamic
– Discrete vs. continuous
– Known vs. unknown

V. THE STRUCTURE OF AGENTS

So far we have talked about agents by describing behavior—the action that is performed after any given sequence of percepts. Now we must bite the bullet and talk about how the insides work. The job of AI is to design an agent program that implements the agent function— the mapping from percepts to actions. We assume this program will run on some sort of computing device with physical sensors and actuators—we call this the architecture:

agent = architecture + program

A. *Agent Programs*

The agent programs that we design in this book all have the same skeleton: they take the current percept as input from the sensors and return an action to the actuators. Notice the difference between the agent program, which takes the current percept as input, and the agent function, which takes the entire percept history. The agent program takes just the current percept as input because nothing more is available from the environment; if the agent's actions need to depend on the entire percept sequence, the agent will have to remember the percepts.

B. *Simple Reflex Agents*

The simplest kind of agent is the simple reflex agent. These agents select actions on the basis of the current percept, ignoring the rest of the percept history.

C. *Model-Based Reflex Agents*

The most effective way to handle partial observability is for the agent to keep track of the part of the world it can't see now. That is, the agent should maintain some sort of internal state that depends on the percept history and thereby reflects at least some of the unobserved aspects of the current state.

D. *Goal-Based Agents*

Knowing something about the current state of the environment is not always enough to decide what to do. For example, at a road junction, the taxi can turn left, turn right, or go straight on. The correct decision depends on where the taxi is trying to get to. In other words, as well as a current state description, the GOAL agent needs some sort of goal information that describes situations that are desirable for example, being at the passenger's destination.

E. *Utility-Based Agents*

Goals alone are not enough to generate high-quality behavior in most environments. For example, many action sequences will get the taxi to its destination (thereby achieving the goal) but some are quicker, safer, more reliable, or cheaper than others. Goals just provide a crude binary distinction between "happy" and "unhappy" states. A more general performance measure should allow a comparison of different world states according to exactly how happy they would make the agent. Because "happy" does not sound very scientific, economists and computer scientists use the term utility instead.

REFERENCES

[1] Russell, Stuart J. ; Norvig, Peter (2003), Artificial Intelligence: A Modern Approach (2nd ed.), Upper Saddle River, New Jersey: Prentice Hall, ISBN 0-13-790395-2 , Chapter . 2
[2] Stan Franklin and Art Graesser (1996); Is it an Agent, or just a Program?: A Taxonomy for Autonomous Agents ; Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages, Springer-Verlag, 1996
[3] N. Kasabov, Introduction: Hybrid intelligent adaptive systems. International Journal of Intelligent Systems, Vol.6, (1998) 453–454.
[4] Weiss, G. (2013). Multiagent systems (2nd ed.). Cambridge, MA: The MIT Press.