

# IAI : Building Intelligent Agents

© John A. Bullinaria, 2005

1. Rational Agents
2. Intelligent Agents
3. Simple Reflex Agents
4. Reflex Agents with an Internal State
5. Goal Based Agents
6. Utility Based Agents
7. Machine Learning
8. The Agent Environments
9. Components of an AI Agent

## Rational Agents

“An *agent* is anything that can be viewed as *perceiving* its environment through *sensors* and *acting* upon that environment through *actuators*.” (Russell & Norvig, page 32)



It is no accident that this diagram is of exactly the same form as the diagram for the interaction of the human nervous system with its environment.

A *rational agent* is one that acts in a manner that causes it to be as successful as it can. What counts as successful obviously depends on what we require the agent to do. We need to determine an appropriate *performance measure* in each case. One has to be careful in choosing that measure, as there are often many conflicting criteria that need to be taken into account.

## Ideal Rational Agents

Being rational also depends on whatever the agent has perceived in the past, what the agent knows about its environment, and what actions it is able to perform.

An *ideal rational agent* is one that takes whatever action is expected to maximise its performance measure on the basis of the evidence provided by its perceptual history and whatever built-in knowledge it has.

Carrying out actions to obtain useful evidence for future actions is all part of rationality. For example, it is usually rational to want to learn!

If an agent's actions are based only on its built-in knowledge, and not on its own experience with its environment, then we say that the agent lacks autonomy. An *autonomous agent* will supplement its built-in knowledge with its own acquired (or learned) knowledge in order to act appropriately.

It is often a good AI strategy to build systems/agents that have enough knowledge to get them started, and then leave them to learn the rest.

# Intelligent Agents

Agents are comprised of an *architecture* (e.g. a computer) plus a *program* that runs on that architecture. In this module we are primarily interested in designing the programs.

In designing intelligent systems there are four main factors to consider:

- P **Percepts** – the inputs to our system
- A **Actions** – the outputs of our system
- G **Goals** – what the agent is expected to achieve
- E **Environment** – what the agent is interacting with

We shall consider four types of agent system of increasing sophistication:

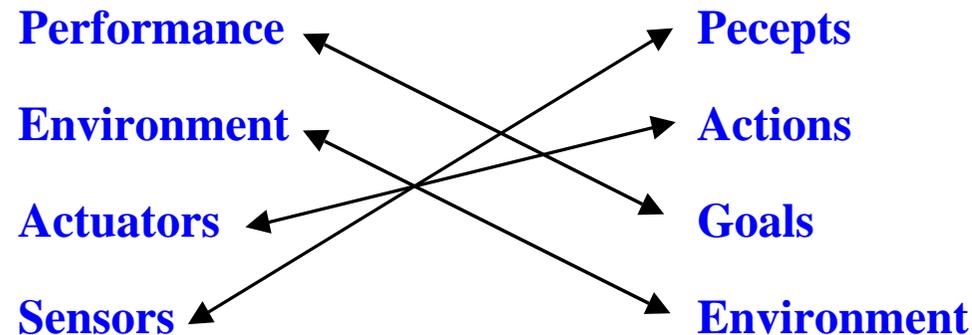
1. Simple Reflex Agents
2. Reflex Agents with an Internal State
3. Goal based agents
4. Utility based agents

## PAGE Descriptors vs. PEAS Descriptors

*PAGE Descriptors* are not the unique way of describing intelligent systems. One popular alternative involves *PEAS Descriptors*:

- P **Performance** – how we measure the system’s achievements
- E **Environment** – what the agent is interacting with
- A **Actuators** – what produces the outputs of the system
- S **Sensors** – what provides the inputs to the system

We can see that there is a clear and simple mapping between *PAGE* and *PEAS*:



Which descriptors we use makes little difference to the systems we build.

## Examples of Intelligent Agents

<b>Agent Type</b>	<b>Percepts</b>	<b>Actions</b>	<b>Goals</b>	<b>Environment</b>
Medical diagnostic system	Symptoms, test results, patient's answers	Questions, test requests, treatments, referrals	Healthy patients, minimise costs	Patient, hospital, staff
Satellite image analysis system	Pixels of varying intensity and colour	Display a categorisation of the scene	Correct image categorisation	Images from orbiting satellite
Part-picking robot	Pixels of varying intensity and colour	Pick up parts and sort them into bins	Place parts into correct bins	Conveyor belt with parts, bins
Refinery controller	Temperature, pressure and chemical readings	Open and close valves, adjust temperature	Maximise purity, yield, safety	Refinery, staff
Interactive English tutor	Typed words	Display exercises, suggestions, corrections	Maximise student's exam results	Set of students, exam papers

## Simple Reflex Agents

Finding the Action appropriate for a given set of Percepts in a look-up table is clearly going to be impossible for all but the most trivial situations, because of the prohibitive number of entries such a table would require.

However, one can often summarise portions of the look-up table by noting commonly occurring input/output associations which can be written as *condition-action rules*:

**if** *{set of percepts}* **then** *{set of actions}*

These are sometimes also called *situation-action rules*, *productions*, or *if-then rules*. Generally, large sets of rules like these, e.g.

**if** *it is raining* **then** *put up umbrella*

may together produce actions that appear intelligent, and they can be implemented very efficiently. Humans have many learned responses and innate reflexes of this form. However, we shall see that their range of applicability is actually very narrow.

## Reflex Agents with an Internal State

One problem with Simple Reflex Agents is that their actions depend only on the current information provided by their sensors.

If a reflex agent were able to keep track of its previous states (i.e. maintain a description or “model” of its previous history in the world), and had knowledge of how the world evolves, then it would have more information about its current state, and hence be able to make better informed decisions about its actions.

Implicit in the knowledge about how the world evolves will be knowledge of how the agent’s own actions affect its state. It will need this so that it can update its own state without having to rely on sensory information.

We say that such reflex agents have an *internal state*. They will have access to information/predictions that are not currently available to their sensors. This will give them more scope for intelligent action than simple reflex agents agents.

## Goal Based Agents

Even with the increased knowledge of the current state of the world provided by an agent's internal state, the agent may still not have enough information to tell it what to do. The appropriate action for the agent will often depend on what its goals are, and so it must be provided with some *goal information*.

Given knowledge of how the world evolves, and of how its own actions will affect its state, an agent will be able to determine the consequences of all (or many) possible actions. It can then compare each of these against its goal to determine which action(s) achieves its goal, and hence which action(s) to take.

If a long sequence of actions is required to reach the goal, then *Search* and *Planning* are the sub-fields of AI that must be called into action. In simple reflex agents, such information would have to be pre-computed and built into the system by the designer.

Goal based agents may appear less efficient, but they are far more flexible. We can simply specify a new goal, rather than re-programming all the rules.

## Utility Based Agents

Goals alone are not enough to generate high quality behaviour. Often there are many sequences of actions that can result in the same goal being achieved. Given appropriate criteria, it may be possible to choose the 'best' sequence of actions from a number that all result in the goal being achieved.

We talk about the *utility* of particular states. Any utility based agent can be described as possessing a *utility function* that maps a state, or sequence of states, on to a real number that represents its utility or usefulness.

We can then use the utility to choose between alternative sequences of actions/states that lead to a given goal being achieved.

When there are conflicting goals, only some of which can be achieved, the utility is able to quantify the appropriate trade-offs.

When there are several goals, none of which can be achieved with certainty, the utility allows the likeliness of success to be weighed up against the importance of the goals.

# Machine Learning

We can see that percepts should not only be used for generating an agent's immediate actions, but also for improving its ability to act in the future, i.e. for learning.

In practice, *learning* can correspond to anything from trivial memorisation, to the creation of complete scientific theories.

We conventionally classify learning into three increasingly difficult classes:

**Supervised Learning** – learning with a teacher (e.g. the system is told what outputs it should produce for each of a set of inputs)

**Reinforcement learning** – learning with limited feedback (e.g. the system must produce outputs for a set of inputs and is only told whether they are good or bad)

**Unsupervised learning** – learning with no help (e.g. the system is given a set of inputs and is expected to make some kind of sense of them on its own)

Machine learning systems can be set up to do all three types of learning.

# Learning Agents

The most powerful agents are able to learn by actively exploring and experimenting with their environment.

A general learning agent has four basic components:

1. **The Performance Element** – which takes in percepts and decides on appropriate actions in the same way as a non-learning agent.
2. **The Critic** – which uses a fixed standard of performance to tell the learning element how well the agent is doing.
3. **The Learning Element** – that receives information from the critic and makes appropriate improvements to the performance element.
4. **The Problem Generator** – that suggests actions that will lead to new and informative experiences (e.g. as in carrying out experiments).

Note that not all learning agents will need a problem generator – a teacher, or the agent's normal mode of operation, may provide sufficient feedback for learning.

# The Agent Environments

We have talked a lot about intelligent agents, but little about the environments with which they interact. There are five particularly useful dimensions along which one can categorize the properties of an agent's environment:

**Accessible vs. Inaccessible** – An environment is *accessible* to an agent if the agent's sensory apparatus gives it access to the complete state of the environment. The agent will not then need to maintain an internal state to keep track of the world.

**Deterministic vs. Non-deterministic** – The environment is *deterministic* if its next state is determined completely by its current state and the actions of its agents. A deterministic but inaccessible environment may *appear* non-deterministic to the agent. Most real environments are so complex that, in practice, they have to be treated as non-deterministic.

**Discrete vs. Continuous** – A *discrete* environment has a limited/finite number of distinct, clearly defined percepts and actions.

**Episodic vs. Non-episodic** – In an *episodic* environment the agent’s experience can be divided into “episodes” consisting of the agent perceiving and then producing actions that depend only on that episode. Subsequent episodes do not depend on previous episodes, and so the agent can limit how far it needs to think ahead.

**Static vs. Dynamic** – A *static* environment does not change while an agent is deliberating. The agent will then have no need to keep checking the world while it is deciding on an action, nor worry about the passage of time.

It is clear that different types of environment will require different types of agents to deal with them effectively.

The hardest environments to cope with will be inaccessible, non-deterministic, continuous, non-episodic, and dynamic.

Often we run an *environment program* for our agents to interact with. This gives each agent its percepts, receives an action from each agent, and then updates the environment. It may also keep track of the performance measures for each agent.

## Components of an AI Agent

We now have a good idea of the components that need to be built into an AI agent:

1. A means to infer properties of the world from its percepts.
2. Information about the way the world evolves.
3. Information about what will happen as a result of its possible actions.
4. Utility information indicating the desirability of possible world states and the actions that lead to them.
5. Goals that describe the classes of states whose achievement maximises the agent's utility.
6. A mapping from the above forms of knowledge to its actions.
7. An active learning system that will improve the agents ability to perform well.

We shall be returning to these issues throughout the remainder of this module.

# Overview and Reading

1. The concepts of agent and rational agent were introduced.
2. The four PAGE descriptors are Percepts, Actions, Goals, and Environment.
3. Increasingly sophisticated agent types are: Simple reflex agents, Reflex agents with an internal state, Goal based agents, and Utility based agents.
4. We looked at the Machine Learning components that can lead to improved agent performance.
5. Five principal distinctions concerning agent environments are: Accessibility, Determinism, Discreteness, Episodicness, and Staticness.
6. We ended by summarising the main components required of an AI agent.

## Reading

1. Russell & Norvig: Sections 2.1, 2.2, 2.3, 2.4, 2.5. [This lecture ties in closely with this material – the other books do not cover it very well.]