Title: Intelligent Agents
AIMA: Chapter 2

Introduction to Artificial Intelligence
CSCE 476-876, Spring 2004
URL: `www.cse.unl.edu/~choueiry/S04-476-876`

Berthe Y. Choueiry (Shu-we-ri)
`choueiry@cse.unl.edu, (402)472-5444`

Intelligent Agents

1. Agents and environments
2. Rationality
3. PEAS
   (Performance measure, Environment, Actuators, Sensors)
4. Types of environments
5. Types of Intelligent Agents

Agent

Anything that
- perceives its environment through sensors
- acts upon its environment through actuators

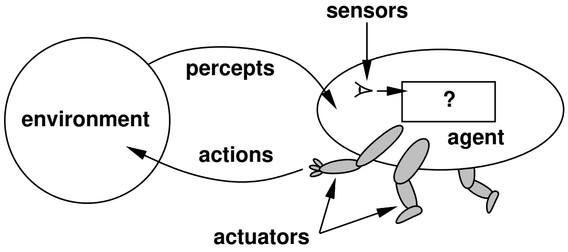

Agents include: Humans, robots, software, etc. Sensors? Actuators?

The **agent function** maps from percept sequences to actions:

$$f : \mathcal{P}^* \rightarrow \mathcal{A}$$

The **agent program** runs on the physical **architecture** to produce f

Vacuum-cleaner world

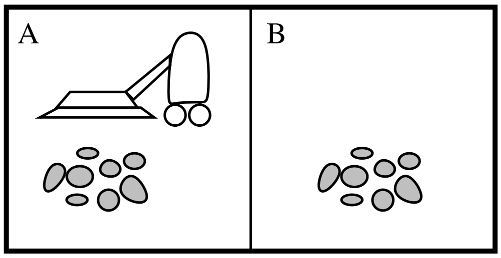

Percepts: locations and contents, e.g., $[A, dirty]$

Actions: $Left$, $Right$, $Suck$, $NoOp$

A Vacuum-cleaner Agent

Percept sequence	Action
$[A, Clean]$	Right
$[A, Dirty]$	Suck
$[B, Clean]$	Left
$[B, Dirty]$	Suck
$[A, Clean]$,$[A, Clean]$	Right
$\vdots$	
$[A, Clean]$,$[A, Clean]$,$[A, Clean]$	Right
$\vdots$	

Function Reflex-Vaccuum-Agent ($[location, status]$) **returns** an action
  if $status = Dirty$ then return $Suck$
  else if $location = A$ then return $Right$
  else if $location = B$ then return $Left$

Goal of AI

Build rational agents.

Rational = ?

What is "rational" depends on:

1. Performance measures (how, when)
2. The agents' prior knowledge of the environment
3. The actions the agent can perform
4. Percept sequence to date (history): everything agent has perceived so far

Rationality

Fixed **performance measure** evaluates the **environment sequence**
– one point per square cleaned up in time t?
– one point per clean square per time step, minus one per move?
– penalize for $> k$ dirty squares?

A **rational agent** chooses whichever action maximizes the **expected** value of the performance measure **given the percept sequence to date**

Rational $\neq$ omniscient, clairvoyant
Rationality maximizes expected performance
Perfection maximizes actual performance

Rational $\Longrightarrow$ exploration, learning, autonomy
After a suffficient experience of its environment, behavior of a rational agents becomes effectively undependent of prior knowledge.

PEAS

To design a rational agent, we must specify the **task environment**

Performance measure?

Environment?

Actuators?

Sensors?

Consider, e.g., the task of designing an automated taxi..

PEAS: Automated taxi

Performance measure: safety, destination, profits, legality, comfort, ...

Environment: US urban streets, freeways, traffic, pedestrians, stray animals, weather, ...

Actuators: steering, accelerator, brake, horn, speaker/display, ...

Sensors: video, accelerometers, gauges, engine sensors, keyboard, GPS, ...

Environment (1)

1. Fully Observable vs. Partially Observable
2. Deterministic vs. stochastic
3. Episodic vs. sequential
4. Static vs. dynamic
5. Discrete vs. continuous
6. Single agent vs. multiagent

Environment (2)

Fully/Partially Observable: sensors can detect all aspects of the world
Effectively fully observable: relevant aspects

Deterministic vs. stochastic: from the agent's view point
Next state determined by current state and agents' actions
Partially observable + deterministic appears stochastic

Episodic vs. sequential: Agent's experience divided into atomic episodes; subsequent episodes do not depend on actions in previous episodes

Environment (3)

Static vs. dynamic:
Dynamic: Environment changes while agent is deliberating
Semidynamic: environment static, performance scores dynamic

Discrete vs. continuous: Finite number of precepts, actions

Single agent vs. multiagent: B's behavior maximizes a performance measure whose value depends on A's behavior.
Cooperative, competitive, communication.

Chess? Taxi driving? hardest case?

Environment (4)

Hardest case: patially observable, stochastic, sequential, dynamic, continuous, and multiagent

	Solitaire	Backgammon	Internet shopping	Taxi
Observable				
Deterministic				
Episodic				
Static				
Discrete				
Single-agent				

Answers depend on how you define/interpret the case

Episodic: chess tournament

Environment types

	Solitaire	Backgammon	Internet shopping	Taxi
Observable	Yes	Yes	No	No
Deterministic	Yes	No	Partly	No
Episodic	No	No	No	No
Static	Yes	Semi	Semi	No
Discrete	Yes	Yes	Yes	No
Single-agent	Yes	No	Yes (except auctions)	No

The environment type largely determines the agent design

The real world is (of course) partially observable, stochastic, sequential, dynamic, continuous, multi-agent

Types of Agents

Four, in order of increasing generality:

1. Simple reflex agents
2. Simple reflex agents with state
3. Goal-based agents
4. Utility-based agents
5. Learning agents

All these can be turned into learning agents.

Simple reflex agents

• Simple look-up table, mapping percepts to actions, is out of question (too large, too expensive to build)

• Many situations can be summarized by condition-action rules (humans: learned responses, innate reflexes)

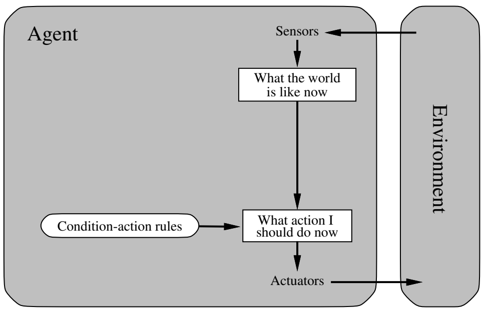

Rectangles: agent's internal state   Ovals: background information

Implementation: easy; Applicability: narrow

Simple reflex agents with state

- Sensory information alone is not sufficient
- Need to keep track of how the world evolves

(evolution: independently of agent, or caused by agent's actions)

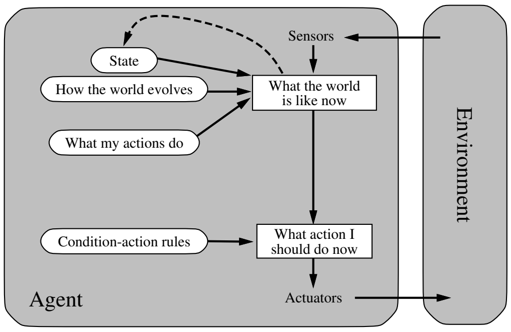

How the world evolved: model-bases agent

Goal-based agents

- State & actions don't tell <u>where</u> to go
- Need goals to build sequences of actions (planning)

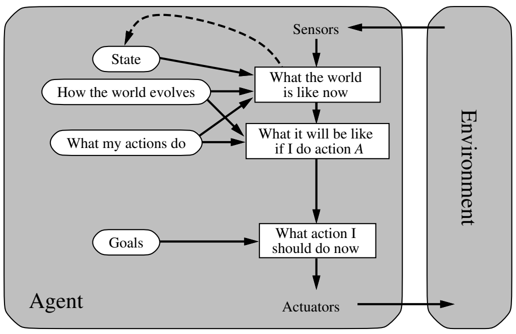

Goal-based: uses the same rules for different goals
Reflex: will need a complete set of rules for each goal

Utility-based agents

• Several action sequences to achieve some goal (binary process)

• Need to <u>select</u> among actions & sequences. Preferences.

• Utility: State → real number (express degree of satisfaction, specify trade-offs between conflicting goal)

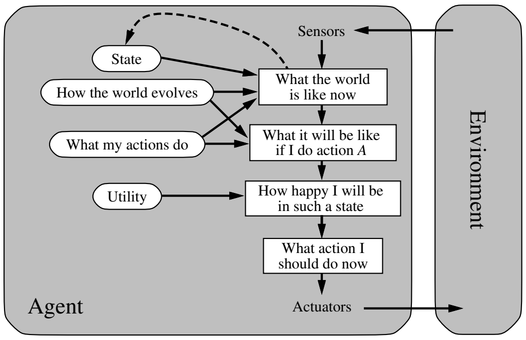

Learning agents

Agent operates in intially an unknown environment, and becomes more competent that its initial know knowledge alonge migh allow

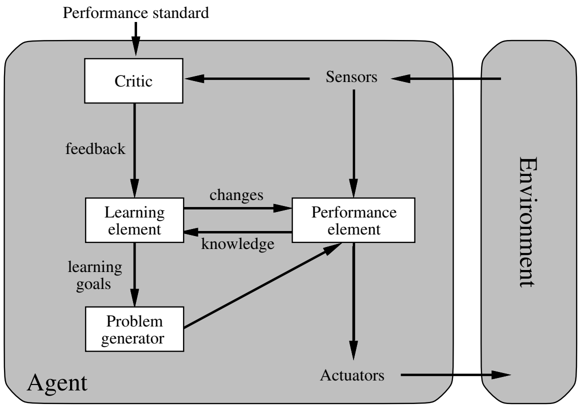

Learning: process of modification of each component of the agent to bring the components into closer agreement with the available feedback information, thus improving overall performance of the agent.