Agents and Environments

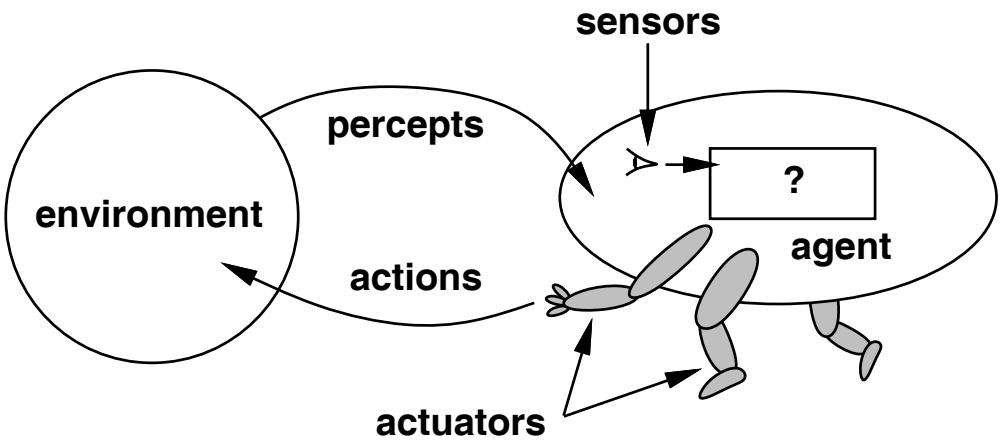

Agents include humans, robots, softbots, thermostats

The *agent function* maps from percept histories to actions

abstract mathematical description

The *agent program* runs on the physical architecture to produce the agent function.

concrete implementation

Vacuum-Cleaner World

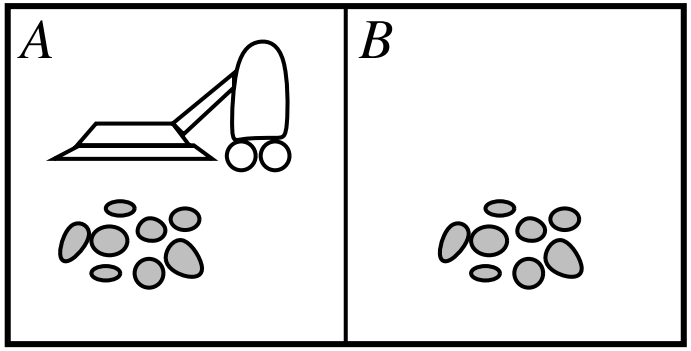

Percepts

current location and contents

e.g., [A, dirty]

Actions

Left, Right, Suck, NoOp

A Vacuum-Cleaner Agent

Agent Function

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

Can it be implemented in a small program?

Agent Program

What is the right program?

```
function REFLEX-VACUUM-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
```

Mike's House as an Agent

Percepts

- motion detectors
- light level sensors
- door/window status
- sound level

Actions

- light intensity
- furnace on/off
- space heaters on/off
- water heater on/off

Rationality

Fixed *performance measure* evaluates sequence of actions performed by agent in environment.

- e.g., number of squares cleaned up by time *T*
- e.g., number of squares cleaned up per time step
- e.g., large penalty at time *T* if any dirty squares remain

A rational agent chooses the action that maximizes the *expected value* of the performance measure given the percept sequence and *a priori* knowledge

- i.e., could not build a better agent with the same built-in knowledge
- does not imply that the agent is omniscient, clairvoyant, or successful
- does require exploration, learning, autonomy

Specifying the Task Environment

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

PEAS

- performance measure
- environment
- actuators
- sensors

E.g., automated taxi

- performance measure: safety, speed, legality, comfort, profit
- environment: US streets & freeways, traffic, pedestrians, weather variation
- actuators: steering direction, accelerator, brake, horn, speaker/display
- sensors: cameras, sonar, speedometer, accelerometer, GPS, engine status

Internet Shopping Agent

User provides agent with a specification of product

- e.g., Treos 600 phone
- e.g., cell phone under $400 with PDA functions

Agent finds lowest prices on web.

Agent bids on Ebay to get a lower price.

PEAS

- performance?
- environment?
- actuators?
- sensors?

Environment Types

aspects of environment	Solitaire	Back-gammon	Internet Shopping	Taxi
fully observable versus partially observable				
next state of environment is deterministic versus stochastic				
episodic or sequential task				
static versus dynamic environment				
discrete versus continuous (env., percepts, actions)				
single or multiple agents				

Environment Types

aspects of environment	Solitaire	Back-gammon	Internet Shopping	Taxi
fully observable versus partially observable	partially	fully	partially	partially
next state of environment is deterministic versus stochastic	determi-nistic	stochastic	semi determi-nistic	stochastic
episodic or sequential task	sequential	sequential	sequential	sequential
static versus dynamic environment	static	static	semi static	dynamic
discrete versus continuous (env., percepts, actions)	discrete	discrete	discrete	continuous
single or multiple agents	single	multiple	multiple (auctions)	multiple

Agent Types

Four basic types (in order of increasing generality)

- simple reflex agents
- reflex agents with state
- goal-based agents
- utility-based agents

Any of these can be made into a learning agent.

Simple Reflex Agents

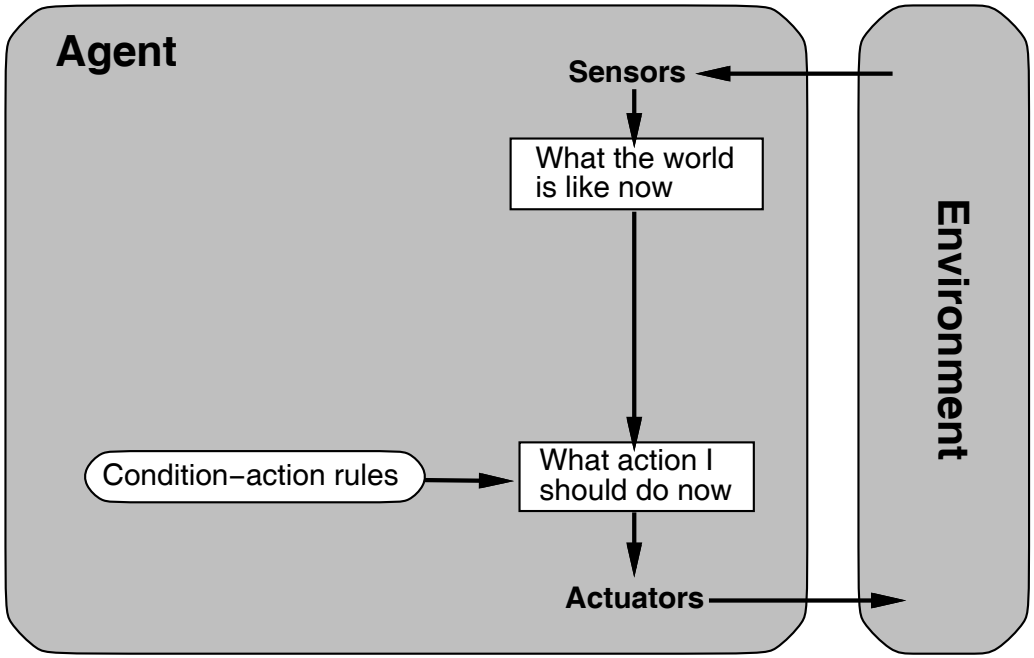

Reflex Agents With State

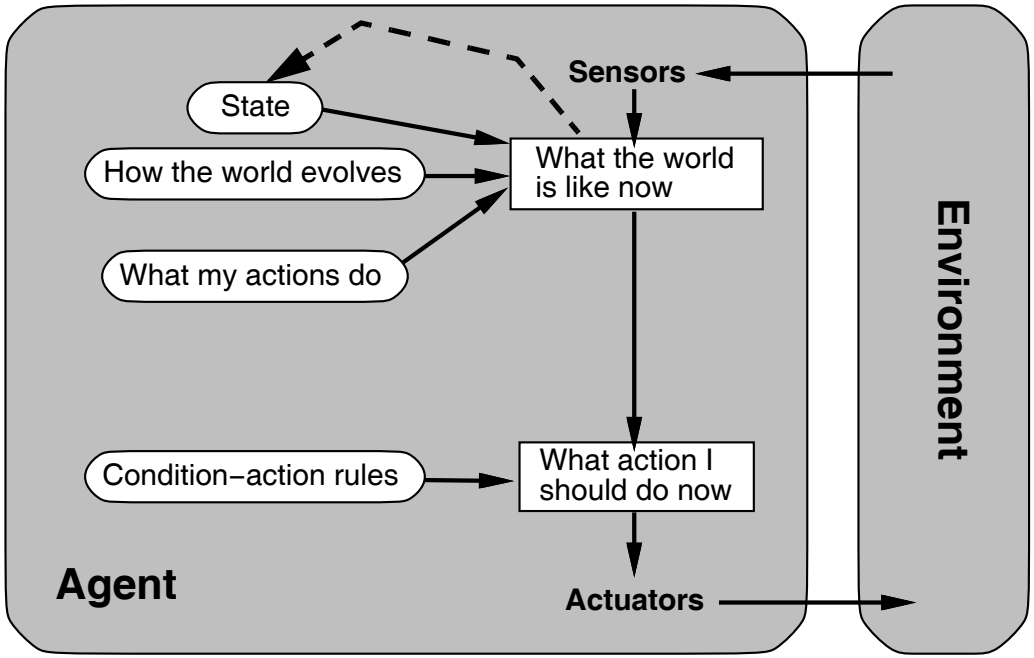

Goal-Based Agents

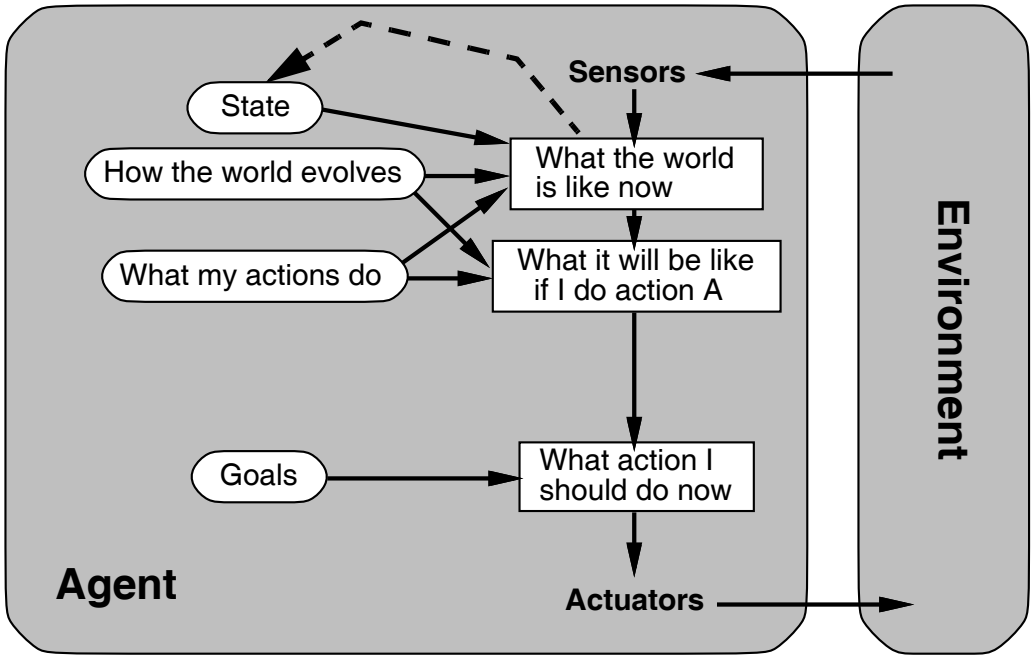

Utility-Based Agents

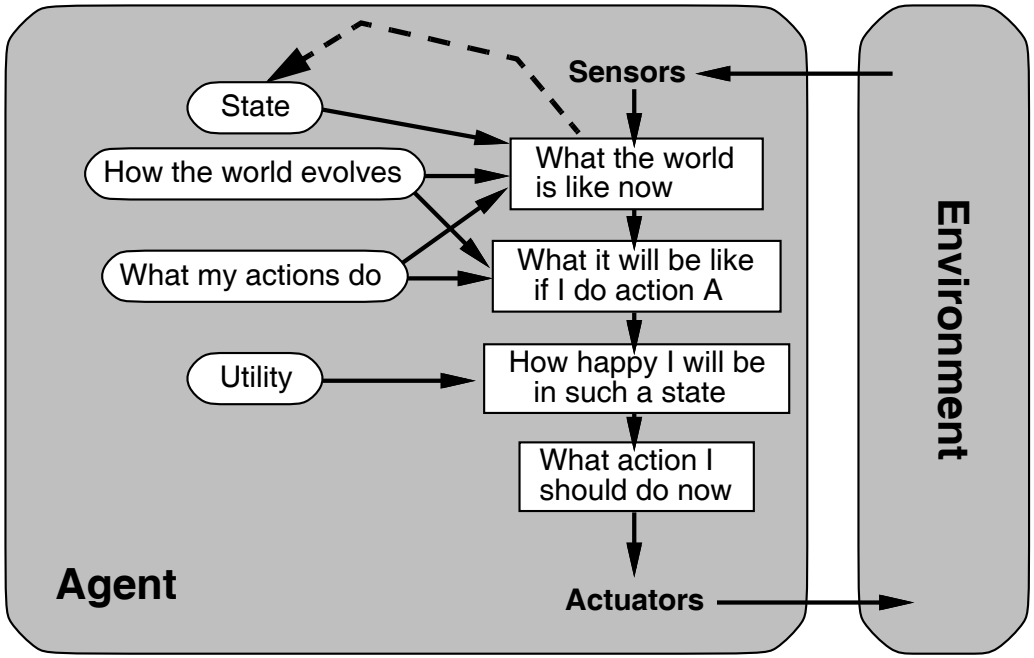

Learning Agents

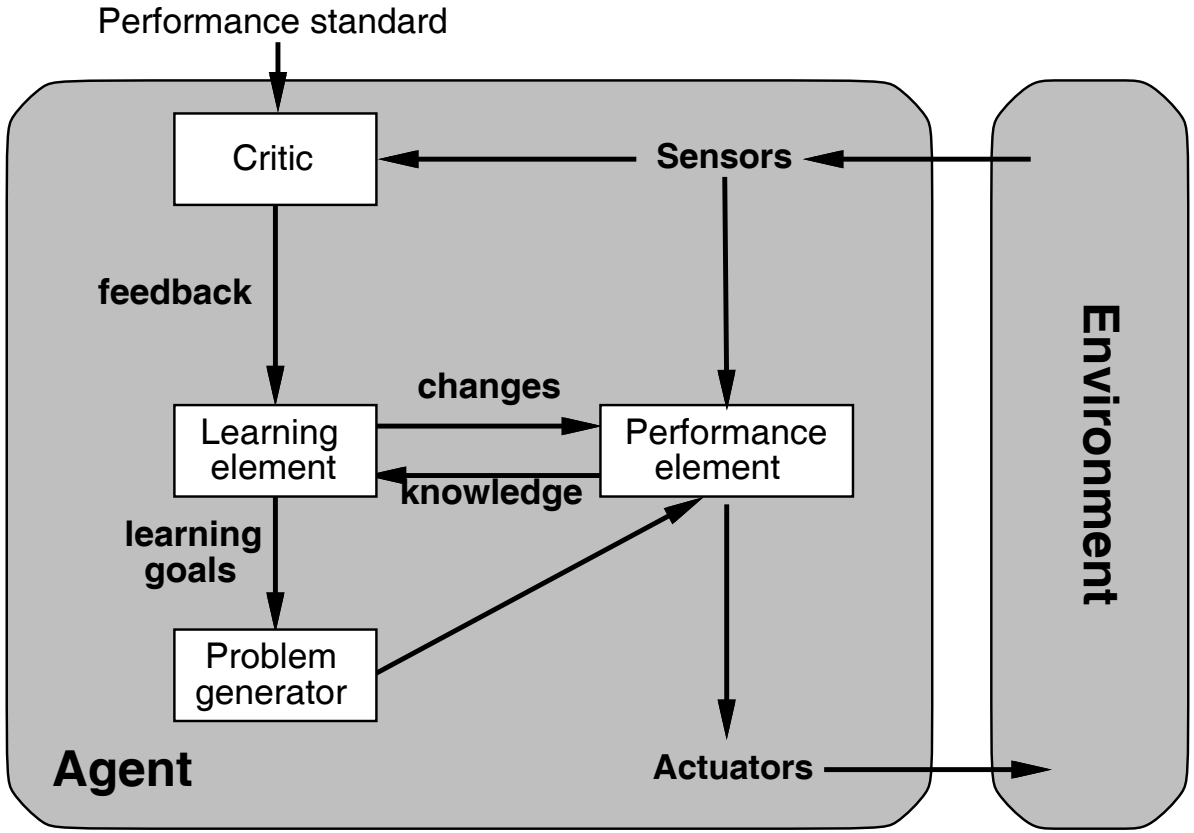

Wumpus World

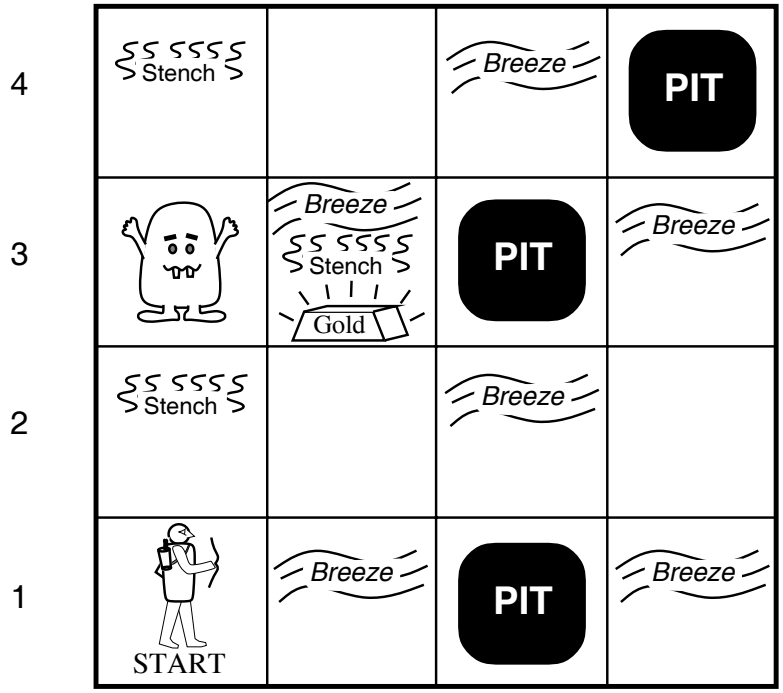

Cave consisting of rooms connected by passageways

Wumpus lurks in cave, eats anyone entering room

Agent can shoot wumpus

Some rooms contain botomless pits

One room contains pot of gold

Wumpus World

Performance measure

+1000 for picking up gold

-1000 for falling into a pit or being eaten by wumpus

-10 for shooting an arrow

-1 otherwise

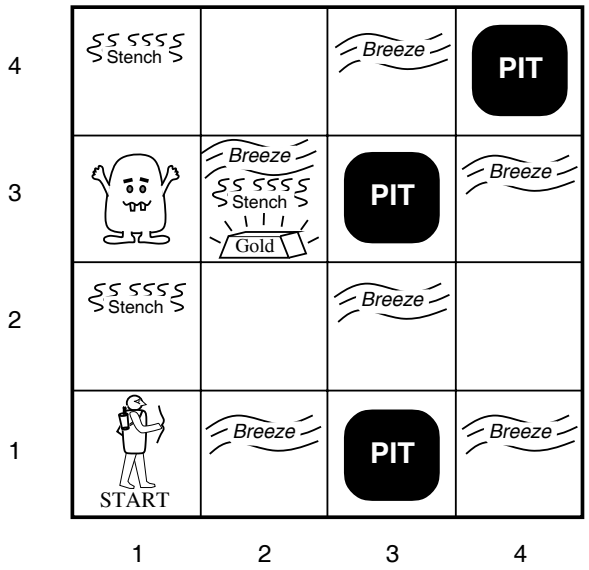

Environment

4x4 grid of rooms

agent always starts at [1,1] facing *right*

gold and wumpus rooms chosen at random with all rooms other than [1,1] being equiprobable; gold and wumpus cannot be in the same room

each room has probability .2 of containing a pit (excluding [1,1] and room containing gold)

Wumpus World

Sensors

Is there a stench in the current room?

Is there a breeze in the current room?

Is there a glitter?

Did the agent perceive a bump (i.e., walked into a wall)?

Was a woeful scream heard (i.e., wumpus was killed)?

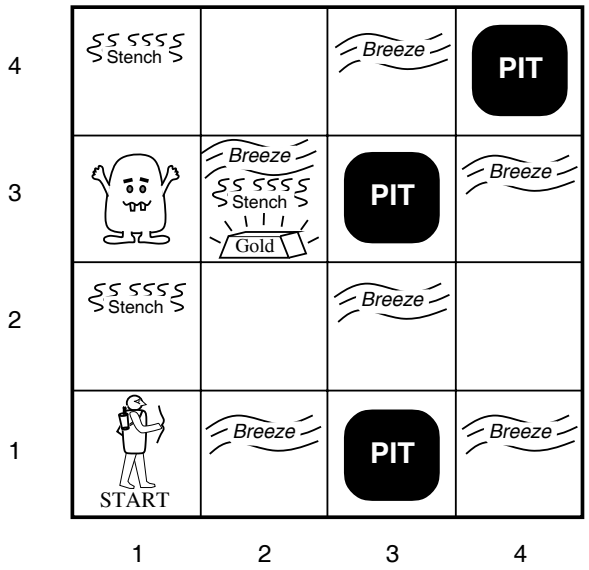

Actuators

move forward to adjacent room (but agent cannot walk off grid)

turn left 90 degrees

turn right 90 degrees

grab object in room (picks up gold if gold is in room)

fire an arrow in current direction (kills wumpus if in adjacent room)