

INTELLIGENT AGENTS

CHAPTER 2

Reminders

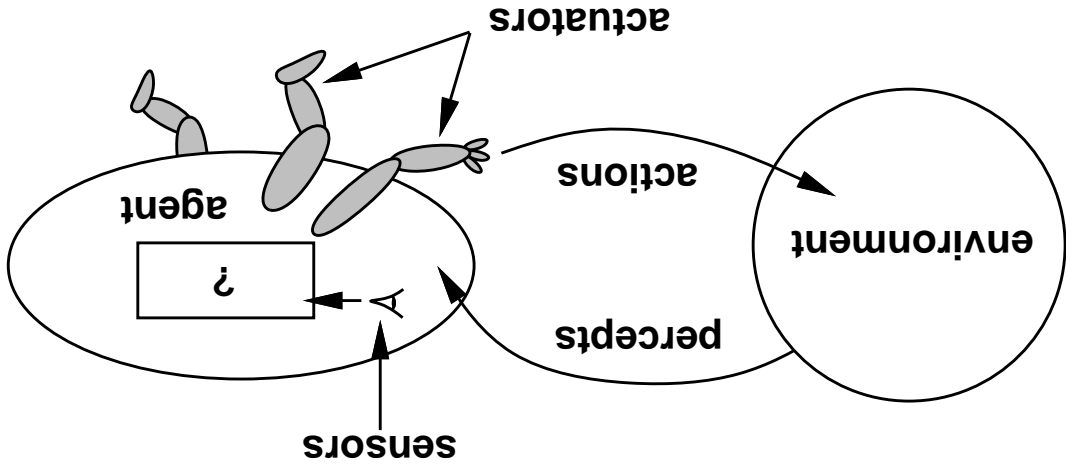
Assignment 0 (lisp refresher) due 1/28

Lisp/emacs/AIMA tutorial: 11-1 today and Monday, 271 Soda

Outline

- ◇ Agents and environments
- ◇ Rationality
- ◇ PEAS (Performance measure, Environment, Actuators, Sensors)
- ◇ Environment types
- ◇ Agent types

Agents and environments



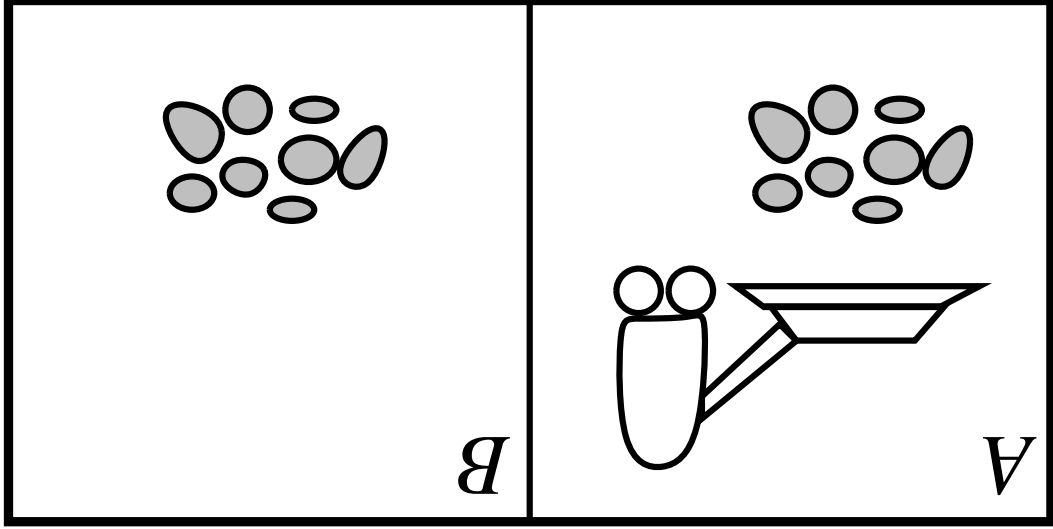
Agents include humans, robots, softbots, thermostats, etc.

The agent function maps from percept histories to actions:

$$f : P^* \rightarrow A$$

The agent program runs on the physical architecture to produce f

Vacuum-cleaner world



Percepts: location and contents, e.g., [A, Dirty]

Actions: *Left, Right, Suck, NoOp*

A vacuum-cleaner agent

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

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*

What is the **right** function?

Can it be implemented in a small agent program?

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

- percepts may not supply all relevant information

Rational \neq clairvoyant

- action outcomes may not be as expected

Hence, rational \neq successful

Rational \Leftrightarrow exploration, learning, autonomy

PEAS

To design a rational agent, we must specify the **task environment**
Consider, e.g., the task of designing an automated taxi:

Performance measure?

Environment?

Actuators?

Sensors?

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

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

- Performance measure? safety, destination, profits, legality, comfort, ...
- Environment? US streets/freeways, traffic, pedestrians, weather, ...
- Actuators? steering, accelerator, brake, horn, speaker/display, ...
- Sensors? video, accelerometers, gauges, engine sensors, keyboard, GPS, ...

PEAS

Internet shopping agent

Performance measure??

Environment??

Actuators??

Sensors??

Internet shopping agent

Performance measure? price, quality, appropriateness, efficiency

Environment? current and future WWW sites, vendors, shippers

Actuators? display to user, follow URL, fill in form

Sensors? HTML pages (text, graphics, scripts)

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

Environment types

	Solitaire Backgammon Internet shopping Taxi
Observable? Deterministic? Episodic? Static? Discrete? Single-agent?	Yes Yes No No

Environment types

						Observable??	Deterministic??	Episodic??	Static??	Discrete??	Single-agent??
Taxi	No	No	Partly	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Internet shopping	No	No	Partly	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Backgammon	No	No	Partly	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Solitaire	No	No	Partly	No	No	Yes	Yes	Yes	Yes	Yes	Yes

Environment types

						Observable?	Deterministic?	Episodic?	Static?	Discrete?	Single-agent??
Solitaire	Backgammon	Internet shopping	Taxi	Yes	Yes	No	Yes	No	No	No	No
				Yes	No	No	Partly	No	No	No	No

Environment types

					Observable?	Yes
					Deterministic?	Yes
					Episodic?	No
					Static?	No
					Discrete?	Yes
					Single-agent?	Yes
Taxi	Internet shopping	Backgammon	Solitaire			
No	No	Yes	Yes	No	Yes	No
No	Partly	No	No	Yes	No	No
No	No	No	No	No	No	No
No	Semi!	Semi!	Yes	Yes	Semi!	Semi!

Environment types

Environment types

		<u>Observable?</u>	Yes	Yes	No	No	No	No	No
		<u>Deterministic?</u>	Yes	No	No	Partly	No	No	No
		<u>Episodic?</u>	No	No	No	No	No	No	No
		<u>Static?</u>	Yes	Semi!	Semi!	Semi!	Semi!	No	No
		<u>Discrete?</u>	Yes	Yes	Yes	Yes	Yes	Yes	No
		<u>Single-agent?</u>	Yes	Yes	Yes	Yes	Yes	Yes	No
Solitaire	Backgammon	Internet shopping	Taxi						

The environment type largely determines the agent design

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

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

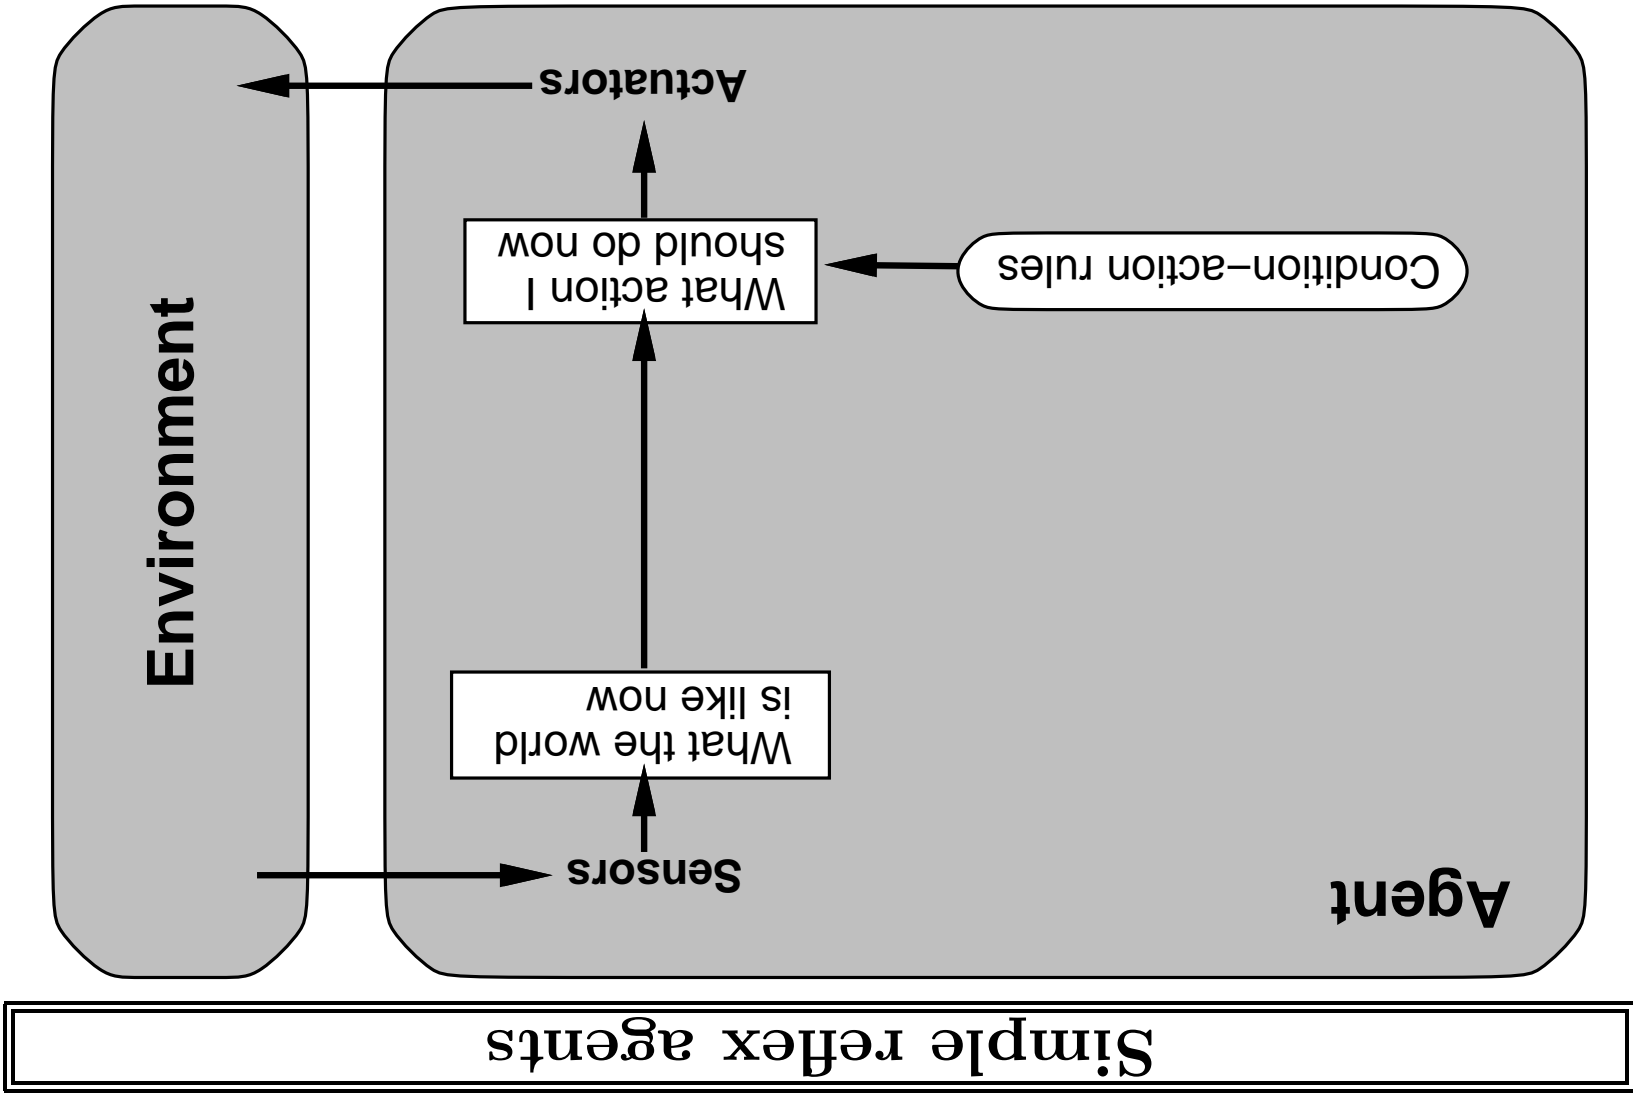
Environment types

Agent types

Four basic types in order of increasing generality:

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

All these can be turned into learning agents



```

(defun make-reflex-vacuum-agent-program ()
  #'(lambda (percept)
    (let ((location (first percept)) (status (second percept)))
      (cond ((eq status 'dirty) 'Suck)
            ((eq location 'A) 'Right)
            ((eq location 'B) 'Left))))))

(setq joe (make-agent :name 'joe :body (make-agent-body)
                     :program (make-reflex-vacuum-agent-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

```

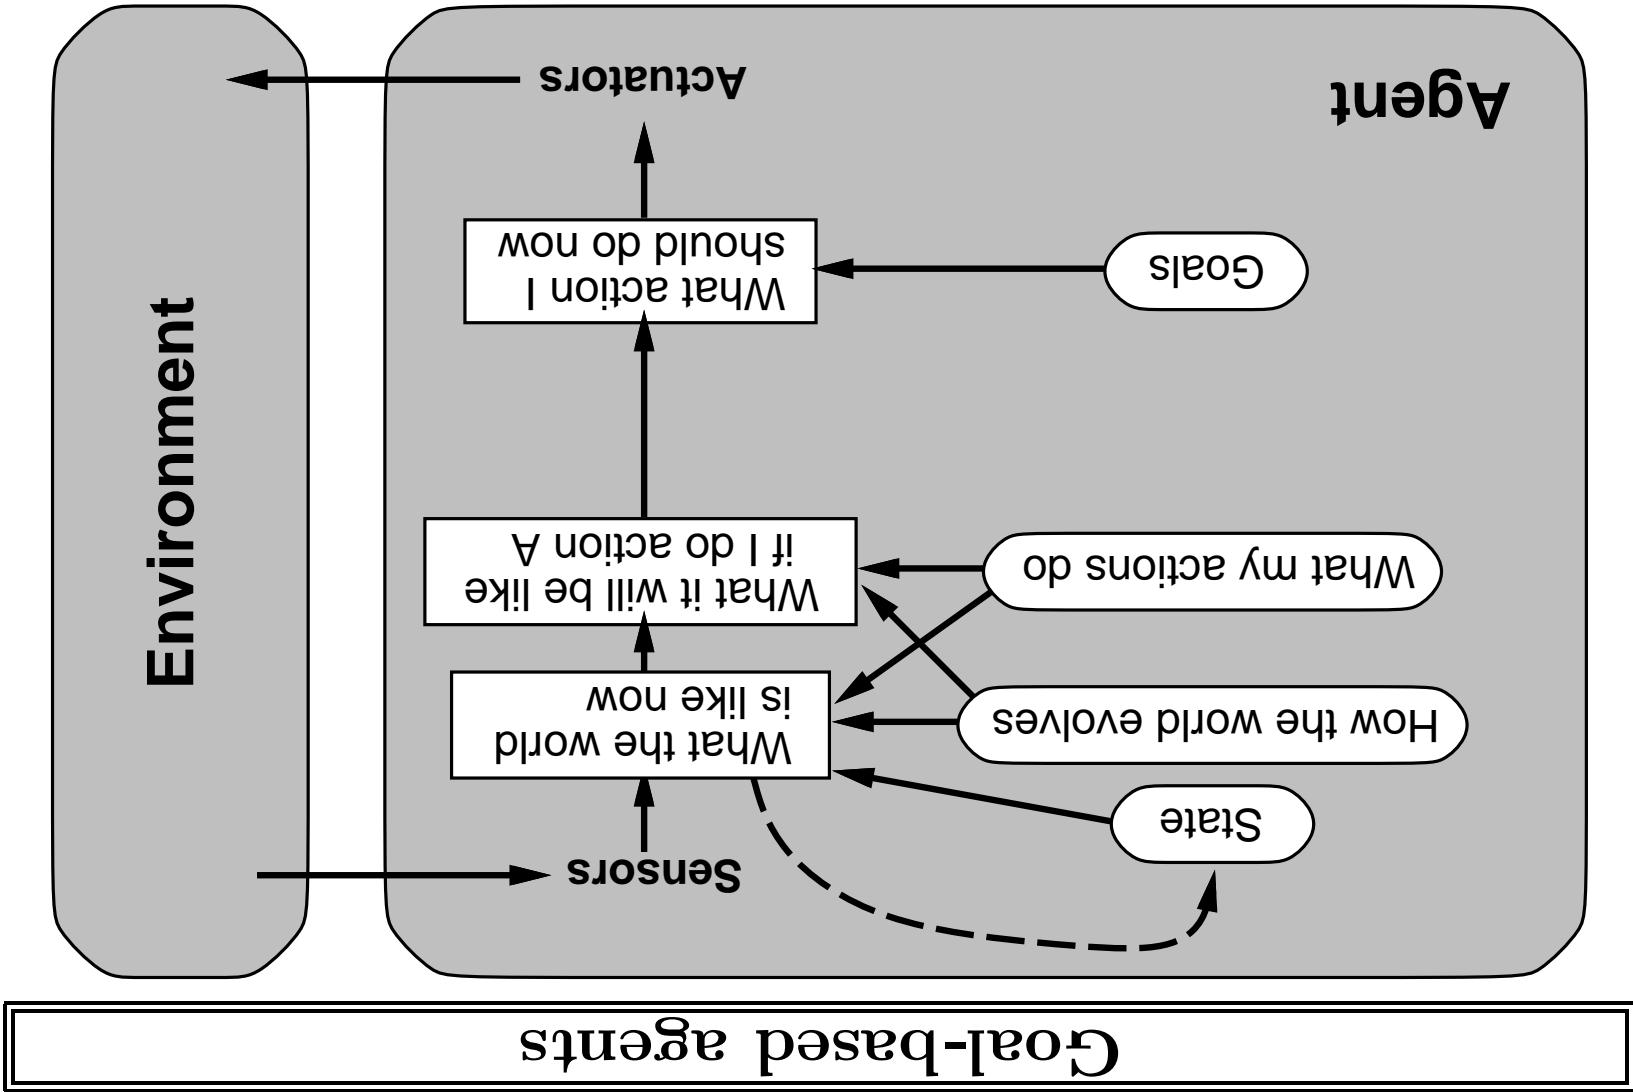
Example



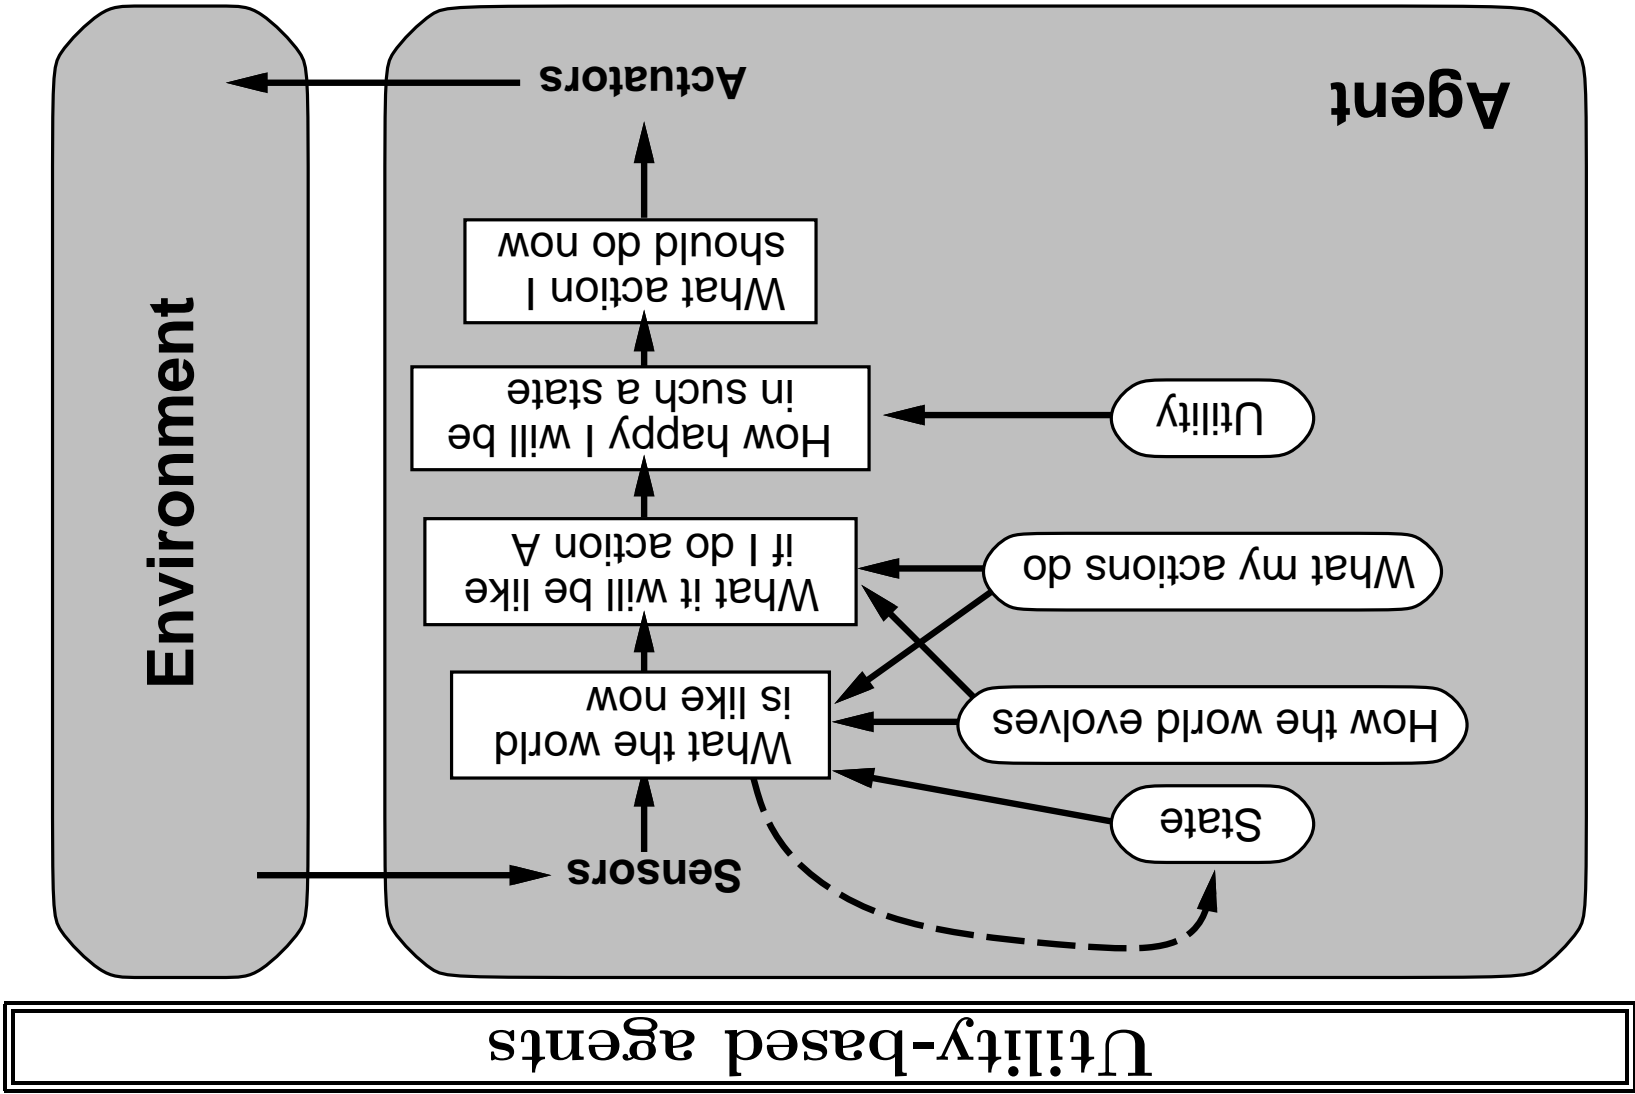
```
(defun make-reflex-vacuum-agent-with-state-program ()
  (let ((last-A infinity) (last-B infinity))
    #'(lambda (percept)
      (let ((location (first percept)) (status (second percept)))
        (incf last-A) (incf last-B)
        (cond
          ((eq status 'dirty)
           (if (eq location 'A) (setq last-A 0) (setq last-B 0))
           'Suck)
          ((eq location 'A) (if (> last-B 3) 'Right 'NoOp))
          ((eq location 'B) (if (> last-A 3) 'Left 'NoOp)))))))
```

```
function REFLEX-VACUUM-AGENT([location,status]) returns an action
static: last-A, last-B, numbers, initially ∞
if status = Dirty then ...
```

Example



Goal-based agents



Summary

Agents interact with environments through actuators and sensors

The agent function describes what the agent does in all circumstances

The performance measure evaluates the environment sequence

A perfectly rational agent maximizes expected performance

Agent programs implement (some) agent functions

PEAS descriptions define task environments

Environments are categorized along several dimensions:

observable? deterministic? episodic? static? discrete? single-agent?

Several basic agent architectures exist:

reflex, reflex with state, goal-based, utility-based