What did you learn last week ?

- Groups and Agents
  - *Groups contains agents which are located in a*
- Geography
  - *and these situated agents have*
- Beliefs and Facts
  - *that lead them to engage in*
- Activities
  - *that are specified within an agent's*
- Workframes
  - *that are triggered by agent's beliefs that lead them to actions like communication, movement, primitive or composite which cause conclusions on the agent's beliefs or affect*
- Classes and Objects
  - *which could change facts in the world that can be detected by agents and change an agent's beliefs which triggers an agent's*
- Thoughtframes
  - *that can have conclusions that again change an agent's beliefs.*

What is an Agent?

An **intelligent agent** perceives its **environment** via **sensors** and acts **rationally** upon that environment with its **actuators**.

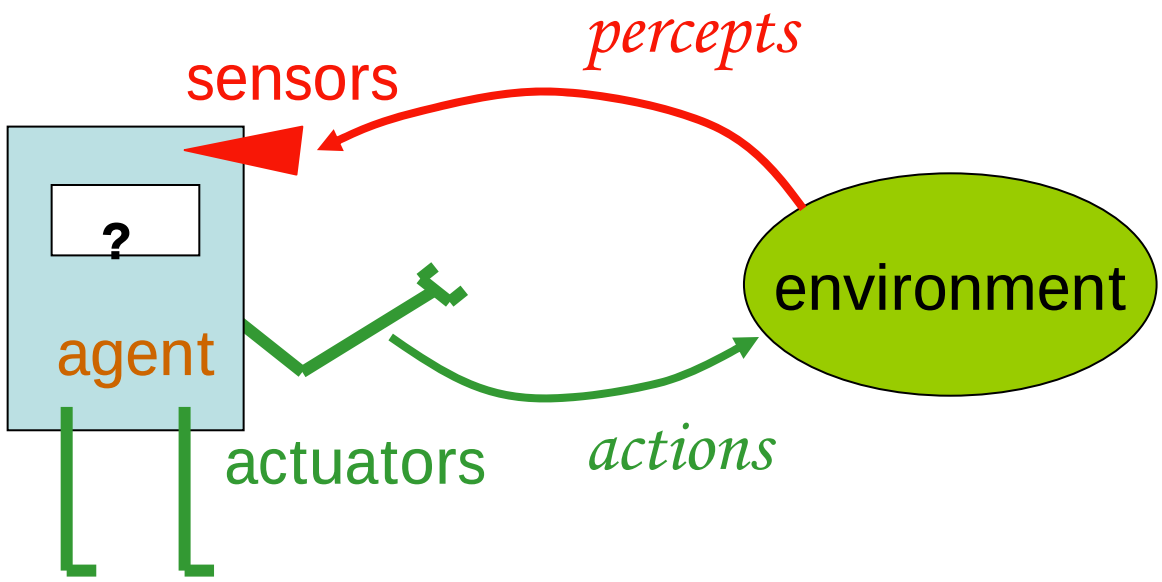

Agent Definition
Key Properties

An agent is a hardware or software system that is:

- **Situated**:

  i.e. embedded in some environment (which may be the physical world, a software environment, a community of agents) which they can:

  - sense (through physical sensors or message receipt or event detection giving partial info on environment state) and
  - act upon (via effectors, messages or event generation with possible non-deterministic outcomes)

- **Reactive :**

  i.e. responds in a timely fashion to messages, sensed data or detected events - so actively monitors state of its environment

- **Autonomous:**

  i.e. operates without the direct intervention of humans or other agents, with independent control over its actions and internal state

E.g. Humans

- Situated
- Sensors:
  - Eyes (vision), ears (hearing), skin (touch), tongue (taste), nose (olfaction), neuromuscular system (proprioception)
- Percepts:
  - At the lowest level – electrical signals
  - After preprocessing – objects in the visual field (location, textures, colors, …), auditory streams (pitch, loudness, direction), …
- Actuators: limbs, digits, eyes, tongue, …
- Actions: lift a finger, turn left, walk, run, carry an object
- (Often) Intelligent and Autonomous

Agent Definition
Other Possible properties

- Social:

  can interact with other agents and possibly humans using messages or actions that change the shared environment
- Pro-active:

  has one or more goals which it tries to achieve by communicating with other agents or acting on its environment
- Has a mentalistic model:

  agent has an internal architecture that can be understood in terms of mentalistic notions such as beliefs, desires, intentions and obligations

Rational agents

- An agent should strive to "do the right thing", based on what it can perceive and the actions it can perform.
- The ***right action*** is the one that will cause the agent to be **most successful**
- **Performance measure**: An **objective criterion for success** of an agent's behavior
- E.g., performance measure of a vacuum-cleaner agent could be amount of dirt cleaned up, amount of time taken, amount of electricity consumed, amount of noise generated, etc.

Performance measure

- Objective: develop agents that perform well in their environments
- A performance measure indicates how successful an agent is
- Two aspects: **how** and **when**

How is performance assessed:

- Different performance measures will be suitable for different types of agents and environments
- **Contrast a trading agent with a vacuum cleaning agent**
- Objective performance measures are defined by us as **external observers** of a system

When is performance assessed:

- Important
- Continually, periodically or one-shot

Tasks for Agents

- We build agents in order to carry out *tasks* for us
- The task must be *specified* by us…
- But we want to tell agents what to do *without* telling them how to do it

Utility-based agents

- What if there are multiple alternative ways of achieving the same goal?
- Goals provide coarse distinction between “happy” and “unhappy” states.
- Utility-based agents have finer degrees of comparison between states.
- World Model + Goals + utility functions

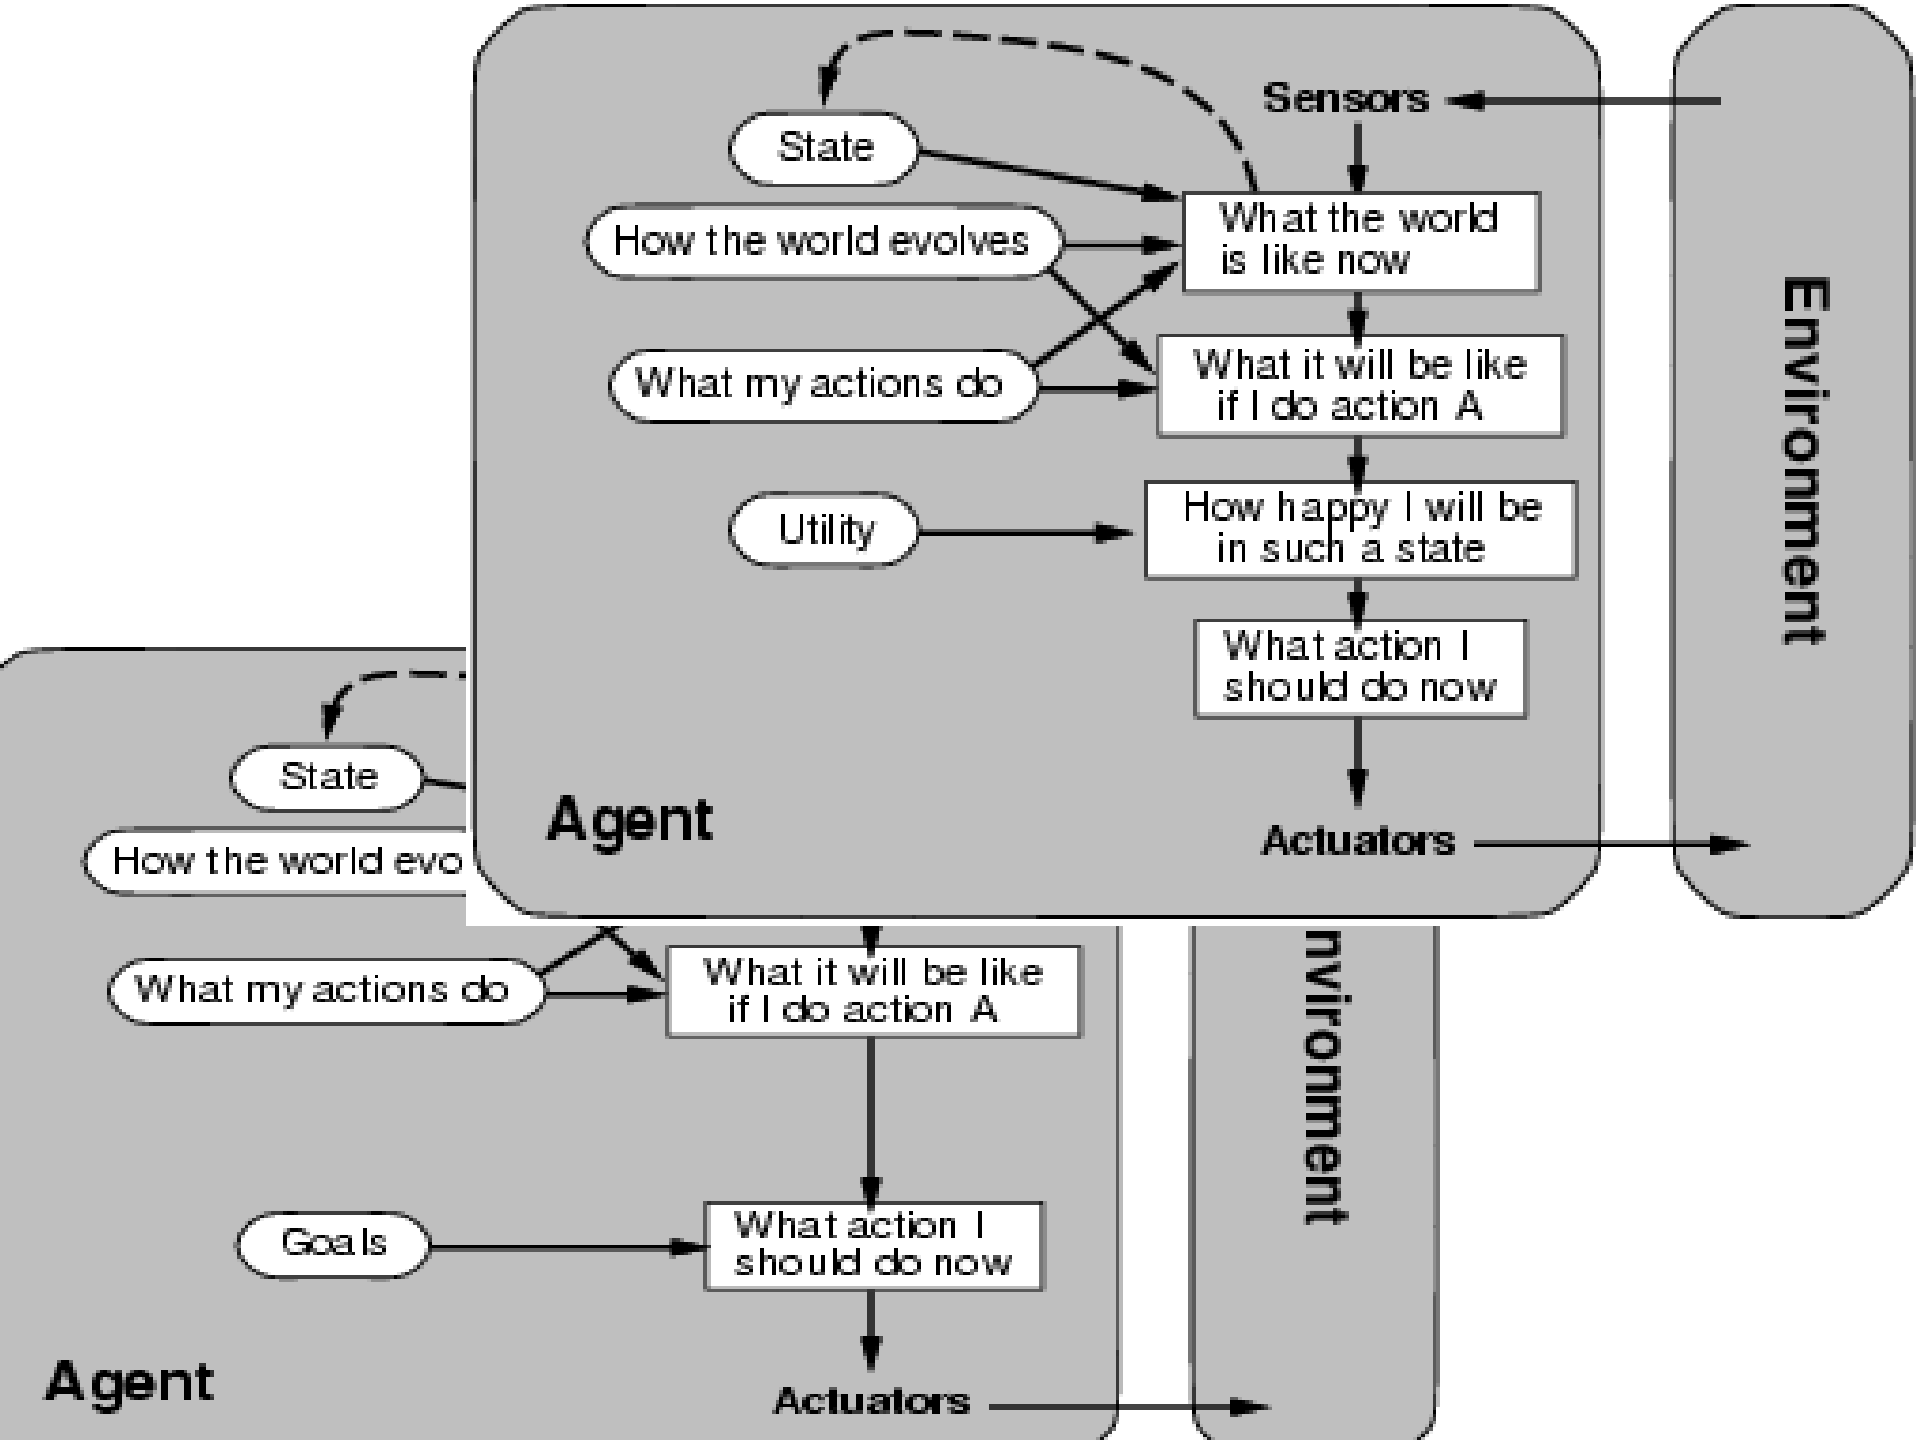
Sensors
State
How the world evolves
What the world is like now
What my actions do
What it will be like if I do action A
Utility
How happy I will be in such a state
What action I should do now
Actuators
Agent
Environment
State
How the world evo
What my actions do
What it will be like if I do action A
Goals
What action I should do now
Actuators
Agent
nvironment

Goal states

- One possible way to measure how well an agent is doing is to check that it has achieved its goal
- There may be a number of different action sequences that will enable an agent to satisfy its goal
- **A good performance measure should allow the comparison of different world states or sequences of states**

Problem Solving Agent : Example

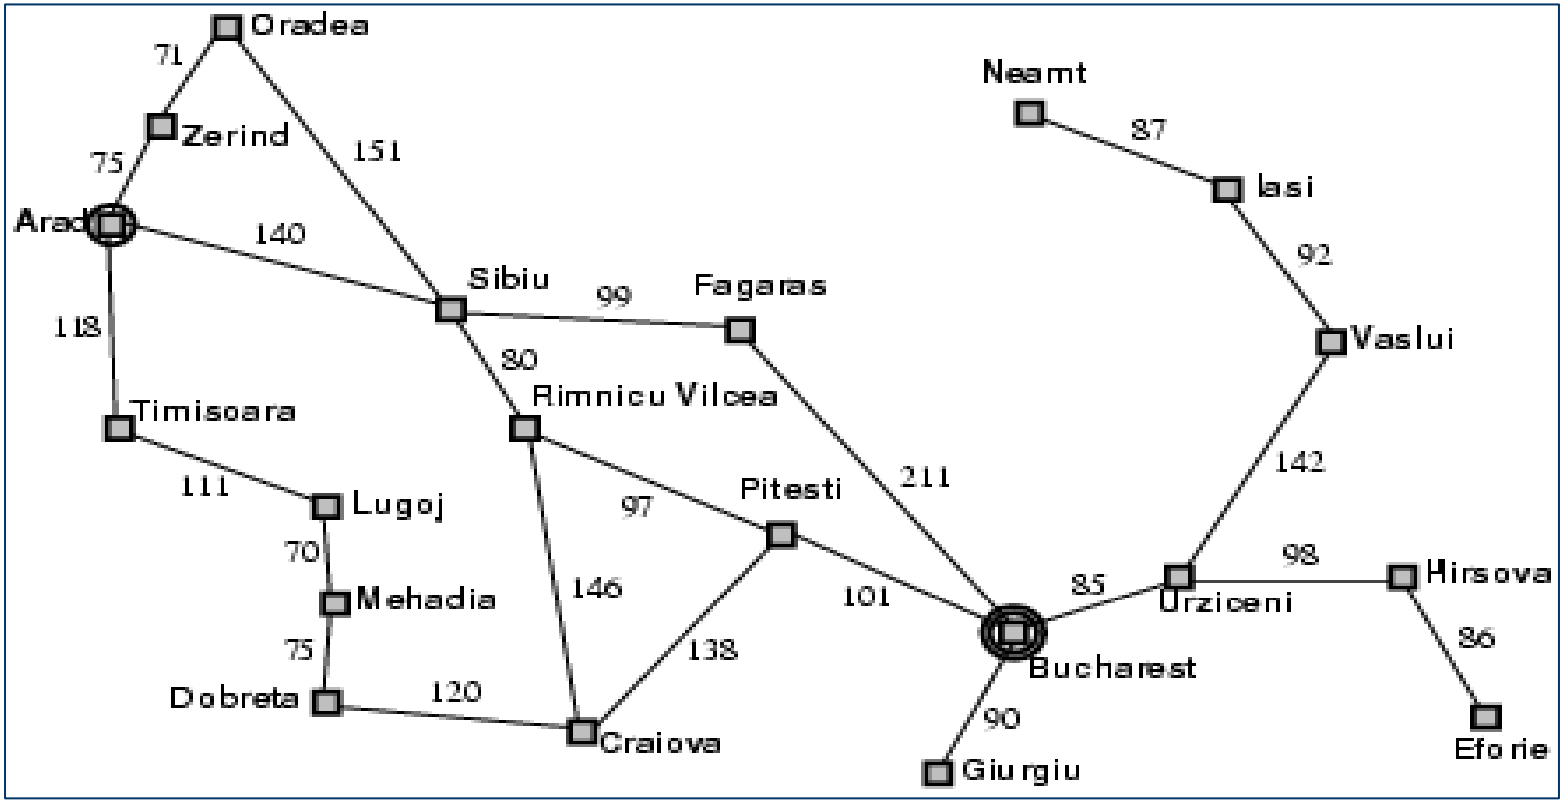

A Simplified Road Map of Part of Romania

Problem Solving Agent : Example

- On holiday in Romania; currently in Arad

- Formulate goal:
  - be in Bucharest
- Formulate problem:
  - states: various cities
  - actions: drive between cities
- Find solution:
  - sequence of cities, e.g., Arad, Sibiu, Fagaras, Bucharest

Rationality

Rationality depends on...

- The performance measure for success, usually taken as **utility**
- What the agent has perceived so far
- What the agent knows about the environment
- The actions the agent can perform

An **ideal rational agent:** *for each possible percept sequence, it acts to maximize its expected utility, on the basis of its knowledge and the evidence from the percept sequence*

Preferences and utilities

- Agents need to be able to express preferences over different goal states
- Each state *s* can be associated with a utility *u*(*s*) for each agent
- The utility is a real number which indicates the desirability of the state for the agent
- For two states *s* and *s'*
  - agent *i* prefers *s* to *s'* if and only if *u*(*s*)>*u*(*s'*)
  - is indifferent between the two states if and only if *u*(*s*)=*u*(*s'*)
- **The agent's objective is to bring about states of the environment that maximise its utility**

Utility-Based Agents

- Goals provide crude binary distinction between “happy” and “un happy”
  - If one state is preferred over the other, then it has higher utility for the agent

  utility-function (state) = real number (degree of happiness)

- Complete specification of utility-function allows rational decisions in the following circumstances
  - Taking decision when in presence of Conflicting goals
  - When there are several goals that the agent can aim for.

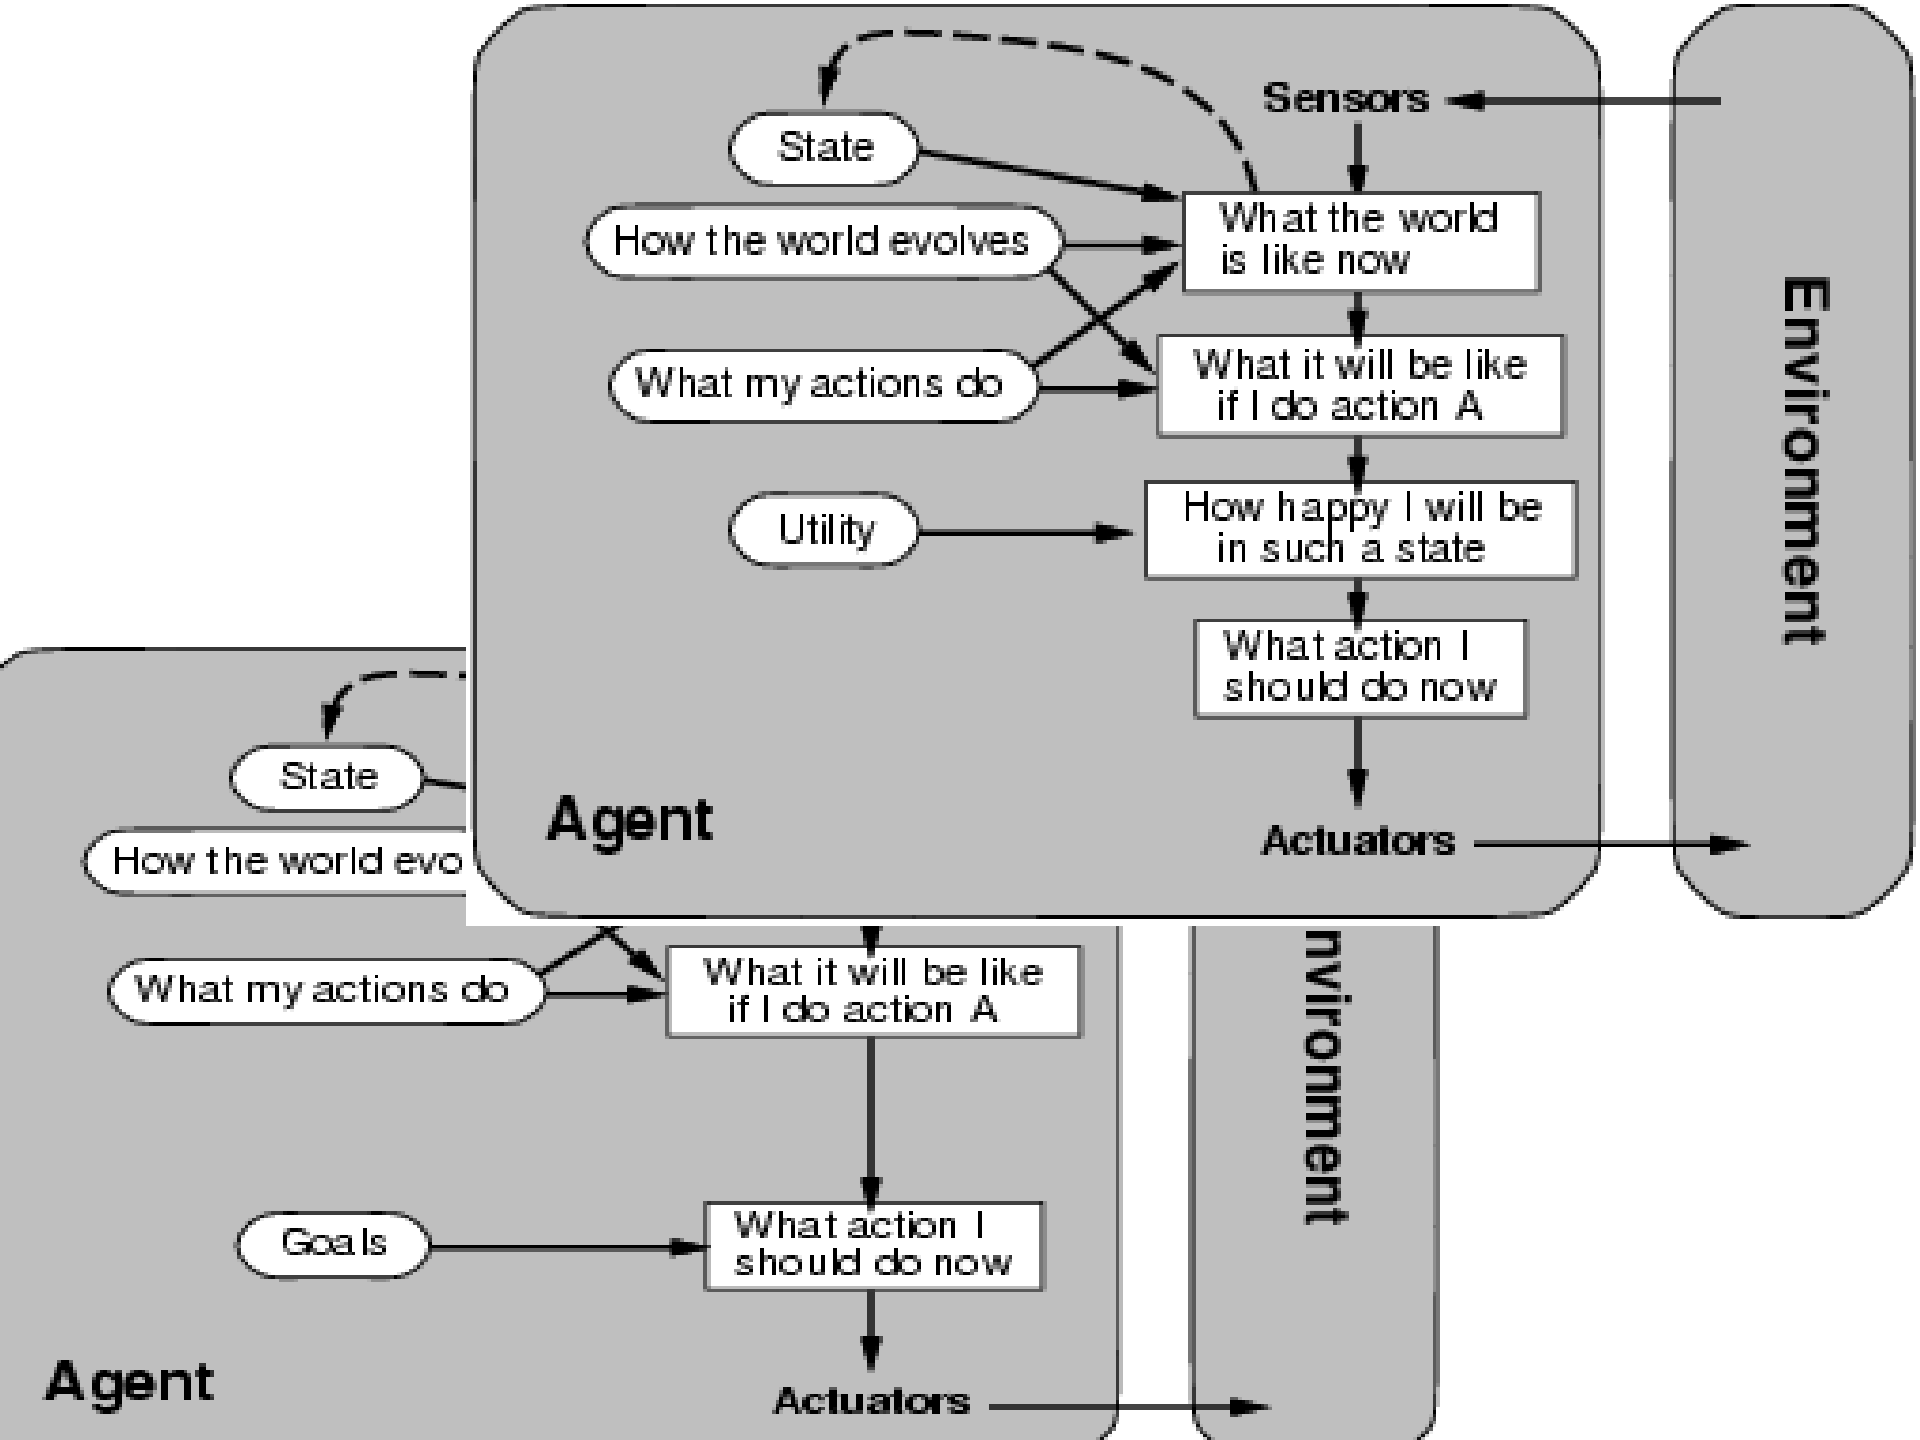
Sensors
State
How the world evolves
What the world is like now
What my actions do
What it will be like if I do action A
Utility
How happy I will be in such a state
What action I should do now
Actuators
Agent
Environment
State
How the world evo
What my actions do
What it will be like if I do action A
Goals
What action I should do now
Actuators
Agent
nvironment

Utility-based agents

Utility functions map states to a measure of the utility of the states, often real numbers.

They are used to:

- Select between conflicting goals
- Select between alternative ways of achieving a goal
- Deal with cases of multiple goals, none of which can be achieved with certainty – weighing up likelihood of success against importance of goal.

Example

- The agent can move North, South, East and West
- Bumping onto a wall leaves the position unchanged
- Fully observable and stochastic: every action to the intended direction succeeds with p=0.8, but with probability p=0.2 the agent moves at right angles towards the intended direction
- Only the utilities of the terminal states are known

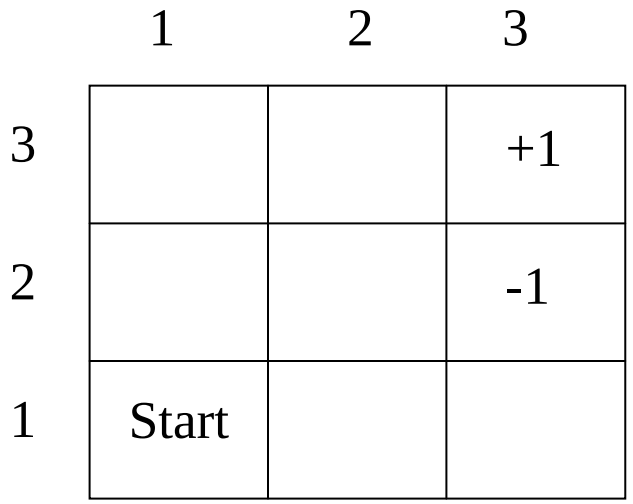

Deterministic vs stochastic environments

Deterministic environments

- The next state is completely determined by the current state and the actions performed by the agent
- The outcome of an agent's actions is uniquely defined, no need to stop and reconsider

Most environments are stochastic

- There is a random element that decides how the world changes
- Limited sphere of influence: the effects of an agent's actions are not known in advance
- An agent's actions may even fail
- Stochasticity complicates agent design

The utilities have been calculated using the value iteration method

	1	2	3
3	0.87	0.93	+1
2	0.82	0.78	-1
1	0.76	0.72	0.49

(a)

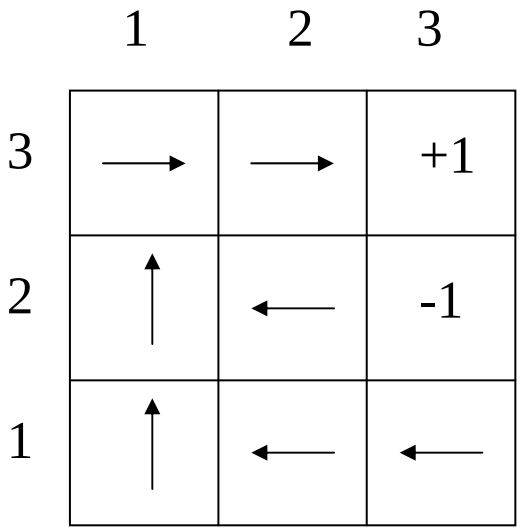

(b)

Utility Functions over States

- One possibility: associate *utilities* with individual states — the task of the agent is then to bring about states that maximize utility
- A task specification is a function

$$u : E \rightarrow \square$$

which associates a real number with every environment state

A complete specification of the utility function allows rational decisions when:

- there are conflicting goals, only some of which can be accomplished; the utility function indicates the appropriate tradeoff;
- there are several goals that the agent can endeavour to achieve, but none of which can be achieved with certainty; the utility provides a way in which the likelihood of success can be evaluated against the importance of the goals

Utility-Based Agents (Cont.)

- Preferred world state has higher utility for agent = quality of being useful
- Examples
  - quicker, safer, more reliable ways to get where going;
  - price comparison shopping
  - bidding on items in an auction
  - evaluating bids in an auction
- Utility function: state ==> U(state) = measure of happiness
- Search (goal-based) vs. games (utilities).

A Utility-Based Agent

function Utility-Based-Agent(*percept*)
static: a set of probabilistic beliefs about the state of the world

Update-Probs-for-Current-State(*percept*,*old-action*)
Update-Probs-for-Actions(*state*, *actions*)
Select-Action-with-Highest-Utility(*probs*)
return *action*