Artificial intelligence
(CK0031/CK0248)

Francesco Corona

Department of Computer Science
Federal University of Ceará, Fortaleza

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Intelligent agents

The concept of **rational agent**: It is central to our approach to AI

- We try to make this notion more concrete

The concept of rationality can be applied to a variety of agents

- Operations in any imaginable environment

We use this concept to develop a set of design principles

⇝ Systems that can reasonably be called intelligent

1. We begin by examining **agents** and **environments**
2. Then, we discuss the **coupling** between them

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Intelligent agents (cont.)

The starting observation is that some agents behave *better* than others

- This leads to the idea of a rational agent

⇝ One that behaves as *well* as possible

How *well* an agent can behave depends on the **environment**

- As, some environments are more difficult than others

We start by giving a crude categorisation of the environments

- We show how their properties influence the agent design

We shall also describe a number of basic ‘skeleton’ agent designs

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agents and environment
Intelligent agents

Agents and environment

Definition

*An **agent** is anything equipped with two abilities*

1. ***Perceiving** its environment thru **sensors***
2. ***Acting** upon that environment through **actuators***

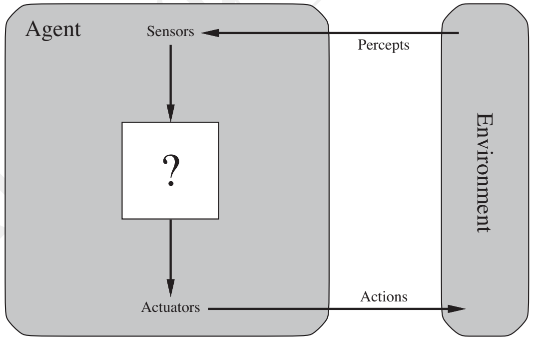

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agents and environment (cont.)

Example

Human agent

- Eyes, ears, and other organs for **sensors**
- Hands, legs, vocal tract, and so on for **actuators**

Robotic agent

- Cameras and infrared range finders for **sensors**
- Various motors for **actuators**

Software agent

- Keystrokes, file contents, incoming net packets as **sensory inputs**
- Displaying, writing files, outgoing net packets, as **acting outputs**

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agents and environment (cont.)

We use the term **percept** to refer to the agent's perceptual inputs

⇝ At any given instant

A **percept sequence** is the history of everything ever perceived

Remark

In general, an agent's choice of action, at any given instant, can depend on the entire percept sequence observed to date

- But it can NOT depend on anything that has NOT been perceived

Suppose that the specification of the agent's choice of action is fully known

- For ALL possible percept sequences

This specification says everything about the agent

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agents and environment (cont.)

Definition

*Mathematically, an agent's behaviour is described by the **agent function***

- *The agent function maps any percept sequence onto an action*

Imagine tabulating the agent function that describes any agent

- For most agents, this would be a very large table

Table is infinite, unless length of percept sequences is bounded

Given an agent to experiment with, we could construct this table

⇝ Try out all possible percept sequences

⇝ Record which actions the agent does in response

The table is an external characterisation of the agent

Definition

*Internally, the agent function will be implemented by an **agent program***

Agents and environment (cont.)

It is important to keep the idea of agent function and program distinct

- The agent function is an abstract mathematical description
- The agent program is a concrete implementation
- (The program runs in some physical system)

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agents and environment (cont.)

Example

The vacuum-cleaner and its world

It is a simple world: We can describe everything that happens

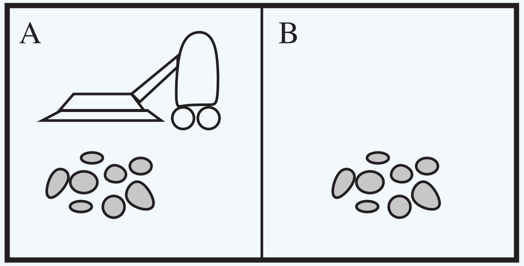

It is a made-up world, we can invent many variations

- This particular world has two locations
- Square A and square B

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agents and environment (cont.)

The perceptions of the agent

- ⇝ Which square it is in
- ⇝ Whether there is dirt in it

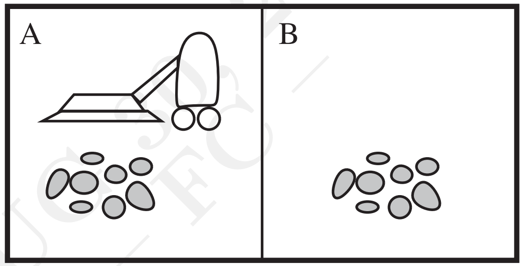

The available actions

- ⇝ Move left or move right
- ⇝ Suck up dirt or do nothing

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agents and environment (cont.)

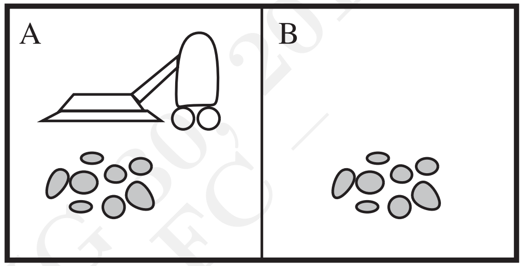

One very simple agent function

- If current square (location) is dirty (status), then suck (action);
- Otherwise, move to the other square (action)

Agents and environment (cont.)

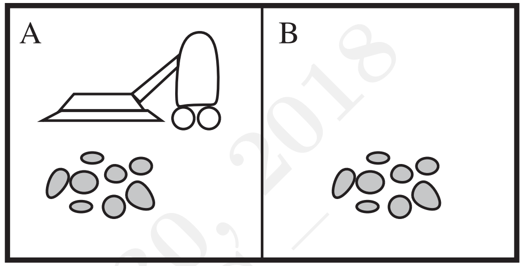

A partial tabulation of this agent function

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*
⋮	⋮
[A, *Clean*], [A, *Clean*], [A, *Clean*]	*Right*
[A, *Clean*], [A, *Clean*], [A, *Dirty*]	*Suck*
⋮	⋮

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agents and environment (cont.)

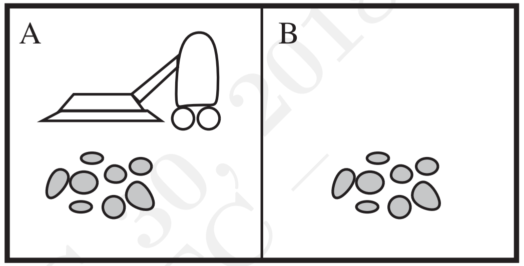

An agent program that implements the function

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*

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agents and environment (cont.)

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*
⋮	⋮
[A, *Clean*], [A, *Clean*], [A, *Clean*]	*Right*
[A, *Clean*], [A, *Clean*], [A, *Dirty*]	*Suck*
⋮	⋮

Various agents can be defined by filling the ‘action’ column in various ways

- What is the right way to fill out the table?

We are asking what makes an agent good/bad, intelligent/dumb?

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agents and environment (cont.)

Remark

The notion of agent is meant to be a tool for analysing systems

- The world is not divided into agents and non-agents
- It is not an absolute characterisation

Example

One can view a standard hand-held calculator as an agent

- The agent chooses the action of displaying ‘4’
- Given the percept sequence ‘2+2=’

Engineering can be seen as designing artefacts that interact with the world

- AI operates at the most interesting extreme of the spectrum
- The task environment requires non-trivial decision making
- Artefacts have significant computational resources

Rationality and good behaviour

Intelligent agents

Rationality and good behaviour

The **rational agent** is defined as one that **does the *right thing***

- Every entry in the table (agent function) is filled out *correctly*

Doing the *right thing* is better than doing the *wrong thing* (uh?)

⇝ What does it mean to do the *right thing*?

Age-old question to be answered in an age-old way

⇝ We consider the consequences

- (Of the agent's actions)

Rationality and good behaviour (cont.)

Consider an agent in an environment, it generates a sequence of actions

- The actions depend on the percepts it receives

Sequence of actions influences the environment, a sequence of its states

⇝ If the sequence of states is *desirable*, the agent performed well

This notion of desirability is captured by a **performance measure**

- It evaluates any given sequence of environment states

Rationality and good behaviour (cont.)

Various performance measure for different environment-task-agent setups

- The designer devise one appropriate to the circumstances
- (This is not as easy as it sounds)

Remark

As a general rule

'*It is better to design performance measures according to what one actually wants in the environment, rather than according to how one thinks the agent should behave*'

Rationality and good behaviour (cont.)

Example

Consider the vacuum-cleaner agent

We measure performance by the amount of dirt cleaned up in a shift

- A rational agent executes actions to maximise performance

The agent could maximise this performance measure in many ways

1. Cleaning up the dirt
2. Dumping it all on the floor
3. Cleaning it up again
4. And, so on

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Rationality and good behaviour (cont.)

A more suitable performance measure would reward the agent

⇝ A clean floor

Award one point for each clean square, at each time step

- Penalise for electricity consumed
- Penalise for noise generated
- ...

The notion of 'clean floor' is based on average cleanliness over time

The same average cleanliness can be achieved by different agents

- One agent does a mediocre job, but it functions all the time
- Another agent cleans energetically, but takes long breaks

■

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Rationality and good behaviour (cont.)

This is a deep philosophical question with profound implications

Which one is better?

- A reckless life of highs and lows
- A safe but mundane existence?

Which one is better?

- An economy where everyone lives in mild poverty (richness)
- An economy in which some live in plenty while others are very poor?

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Rationality

Rationality and good behaviour

Rationality

What is rational, for an agent, at any given time, depends on four things

1. The performance measure (the criterion of success)
2. The agent's prior knowledge of the environment
3. The actions the agent can perform
4. The agent's percept sequence to date

This leads to the definition of a rational agent

Definition

Rational agent

'For each possible percept sequence, a rational agent should select an action that is expected to maximise its performance measure, given the evidence by the percept sequence and whatever built-in knowledge the agent has'

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Rationality (cont.)

Example

Consider the vacuum-cleaner agent

It cleans a square if it is dirty and it moves to the other square if is not

- Is this a rational agent?
- That depends!

Things that we need to specify before we can answer

1. What is known about the environment
2. What are the sensors and actuators
3. What the performance measure is

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Rationality (cont.)

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*
⋮	⋮
[*A*, *Clean*], [*A*, *Clean*], [*A*, *Clean*]	*Right*
[*A*, *Clean*], [*A*, *Clean*], [*A*, *Dirty*]	*Suck*
⋮	⋮

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Rationality (cont.)

Let us make an assumption

Assumption

The performance measure awards 'one point, for each clean square,

- At each time step, over a 'lifetime' of 1K time steps'

- The 'geography/map' of the environment is known *a priori*
- Dirt distribution and initial location are unknown *a priori*

- The only actions available to the agent are `Left`, `Right`, and `Suck`
- The `Left` and `Right` actions move the agent left and right
- ⇝ When this would take the agent outside the environment, the agent remains where it is
- `Sucking` cleans the current square
- Clean squares stay clean

- The agent correctly perceives its location
- The agent perceives whether its location contains dirt

Rationality (cont.)

We safely claim that this **agent is rational**, under these circumstances

⇝ Its expected performance is at least as high as any other agent's

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Rationality (cont.)

Example

The vacuum-cleaner agent function is rational, under the assumptions

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*
⋮	⋮
[A, *Clean*], [A, *Clean*], [A, *Clean*]	*Right*
[A, *Clean*], [A, *Clean*], [A, *Dirty*]	*Suck*
⋮	⋮

To show rationality, show that it cleans at least as fast as any other agent

- For all possible actual environments
- (all dirt distributions and initial locations)

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Rationality (cont.)

This is trivially true when there is no dirt (00)

Dirt in the initial location and none in the other location (10)

- The world is clean after one step (no agent can do better)

No dirt in the initial location and dirt in the other location (01)

- The world is clean after two steps (no agent can do better)

Dirt in both locations (11)

- The world is clean after three steps (no agent can do better)

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Rationality (cont.)

The same agent is considered irrational, under different circumstances

For example,

- Once all dirt is cleaned up, the agent will oscillate back and forth
- Bad, if the performance measure penalises movements

A better agent for this case would done that does nothing

- Once it is sure that all the squares are clean

If clean squares can become dirty again, the agent should check

⇝ Re-clean them if needed, occasionally

The agent will need to explore the environment

- Rather than stick to squares A and B

If the geography of the environment is unknown

Omniscience, learning and autonomy

Rationality and good behaviour

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Omniscience, learning and autonomy

We must carefully distinguish between rationality and omniscience

- An omniscient agent knows the actual outcome of its actions
- ⇝ That is, the agent can act accordingly

Omniscience is impossible in reality

Example

I am walking around and I see an old friend across the street

There is no traffic nearby and I am not otherwise engaged

- So, being rational, I start to cross the street

Meanwhile, at 33K feet, a door falls off a passing airliner

- Before I make it to the other side I am flattened

Was I irrational to cross the street?

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Omniscience, learning and autonomy (cont.)

The example shows that rationality is not the same as perfection

Remark

- ⇝ **Rationality maximizes expected performance**
- ⇝ **Perfection maximizes actual performance**

Perfection should not be a requirement for agents

We cannot expect agents to do what turns out to be the best action

- It would be impossible to design agents
- Impossible to fulfil the specification

Our definition of rationality does not require omniscience

- ⇝ The rational choice depends only on the percepts
- The percept sequence, to date

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Omniscience, learning and autonomy (cont.)

Ensure that the agent is not allowed to engage in sub-intelligent activities

Example

A percept sequence will not tell that there is a truck approaching fast

- If an agent does not look both ways before crossing

Does our definition of rationality say that it is OK to cross? NO, it does not!

Not rational to cross the road, given the uninformative percept sequence

- ⇝ The risk of accident from crossing without looking is too big

Rationality suggests the ‘looking’ action, before start crossing the street

- ⇝ Looking helps maximise the expected performance

Omniscience, learning and autonomy (cont.)

Doing actions to modify future percepts is an important part of rationality

- It is called **information gathering**

Example

A second example of information gathering

- The **exploration** that must be undertaken by a vacuum-cleaner
- In an initially unknown environment

Our definition requires a rational agent to gather info

⇝ To learn from what it perceives

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Omniscience, learning and autonomy (cont.)

Remark

Agent's initial configuration could reflect prior knowledge of environment

- As the agent gains experience, this may be modified and augmented

Extreme case: The environment is completely known a priori

- The agent need not perceive or learn
- It simply acts correctly

Such agents are fragile

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Omniscience, learning and autonomy (cont.)

Example

The dung beetle

Digs its nest and lays its eggs, fetches a ball of dung to plug the entrance

- Suppose the ball of dung is removed from its grasp *en route*
- The beetle continues its task, pantomimes plugging the nest
- ... with the non-existent dung ball

It never notices the ball of dung is missing

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Omniscience, learning and autonomy (cont.)

Evolution has built an assumption into the beetle's behaviour

- When it is violated, unsuccessful behaviour results

Omniscience, learning and autonomy (cont.)

Example

The sphex wasp

- The female sphex will dig a burrow, go out and sting a caterpillar
- Drag it to the burrow, enter the burrow again to check all is well
- Drag the caterpillar inside, and lay its egg

The caterpillar serves as a food source when the eggs hatch

Suppose the caterpillar is moved away while the sphex is doing the check

- It will revert to the 'first drag' step

It will continue the plan w/o modification, even after dozens interventions

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Omniscience, learning and autonomy (cont.)

The sphex innate plan failed (drama)

- The sphex is unable to learn this
- Thus, will not change the plan

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Omniscience, learning and autonomy (cont.)

Definition

Autonomy (or, lack thereof)

The extent that an agent relies on the prior knowledge of its designer

- *Rather than on its own percepts*

A truly rational agent should be autonomous, it should learn

⇝ To compensate for partial or incorrect prior knowledge

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Omniscience, learning and autonomy (cont.)

Example

A vacuum-cleaner that learns to foresee where and when dirt will appear

⇝ It will do better than one that does not

Practically, complete autonomy is seldom required from start

- With little/no experience, the agent has to act randomly
- Unless the designer gave some assistance

Omniscience, learning and autonomy (cont.)

Reasonable to provide an agent with initial knowledge and ability to learn

- After sufficient experience of its environment, the behaviour can become effectively *independent of prior knowledge*
- Incorporation of learning allows to design a single rational agent that will succeed in a variety of environments

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

The environment

Intelligent agents

AUG 30, 2018 _FC

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

The environment

AUG 30, 2018 _FC

Equipped with a definition of rationality, we start building rational agents

- First, we discuss **task environments**

'Problems' to which rational agents are 'solutions'

We may begin by showing how to specify a task environment

- We illustrate the process with a number of examples

We show that there exists a variety of task environments

- The flavour affects the design for the agent program

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Specification of task environments

The environment

AUG 30, 2018 _FC

Specification of task environments

When discussing rationality (vacuum-cleaner), we had to specify

- Environment
- Performance measure
- Actuators and sensors

We group them under the heading **task environment**

Definition

Task environment

PEAS**: **Performance**, **Environment**, **Actuators**, **Sensors

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Specification of task environments (cont.)

Remark

The design of an agent

First step must always be to specify the task environment

- The specification must be as complete as possible

The vacuum world was a simple example

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Specification of task environments (cont.)

Example

An automated taxi driver

This is a more complex problem

- '*We should point out, before the reader becomes alarmed, that a fully automated taxi is currently somewhat beyond the capabilities of existing technology*'

The full driving task is open-ended

- No limit to novel combinations of circumstances that can arise

The PEAS description for the taxi's task environment

Agent Type	Performance Measure	Environment	Actuators	Sensors
Taxi driver	Safe, fast, legal, comfortable trip, maximize profits	Roads, other traffic, pedestrians, customers	Steering, accelerator, brake, signal, horn, display	Cameras, sonar, speedometer, GPS, odometer, accelerometer, engine sensors, keyboard

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Specification of task environments (cont.)

Agent Type	Performance Measure	Environment	Actuators	Sensors
Taxi driver	Safe, fast, legal, comfortable trip, maximize profits	Roads, other traffic, pedestrians, customers	Steering, accelerator, brake, signal, horn, display	Cameras, sonar, speedometer, GPS, odometer, accelerometer, engine sensors, keyboard

Performance measure the automated driver should aspire to?

- Getting to the correct destination;
- Minimising fuel consumption and wear and tear;
- Minimising the trip time or cost;
- Minimising violations of traffic laws;
- Minimise disturbances to other drivers;
- Maximising safety and passenger comfort;
- Maximising profits

Some of these goals conflict, so tradeoffs will be required

Specification of task environments (cont.)

Agent Type	Performance Measure	Environment	Actuators	Sensors
Taxi driver	Safe, fast, legal, comfortable trip, maximize profits	Roads, other traffic, pedestrians, customers	Steering, accelerator, brake, signal, horn, display	Cameras, sonar, speedometer, GPS, odometer, accelerometer, engine sensors, keyboard

What is the driving **environment** that the taxi will face?

- A variety of roads, from rural and urban to 12-lane freeways
- Other traffic elements, pedestrians, animals and road works
- More traffic elements, police cars, puddles, and potholes
- The taxi must interact with potential/actual passengers

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Specification of task environments (cont.)

There are also some optional choices:

- Operate where snow is uncommon or where it is common
- Driving on the right, on the left, or both depending

Remark

The more restricted the environment, the easier the design problem

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Specification of task environments (cont.)

Agent Type	Performance Measure	Environment	Actuators	Sensors
Taxi driver	Safe, fast, legal, comfortable trip, maximize profits	Roads, other traffic, pedestrians, customers	Steering, accelerator, brake, signal, horn, display	Cameras, sonar, speedometer, GPS, odometer, accelerometer, engine sensors, keyboard

Actuators for an automated taxi are those available to humans

- Control over the engine through the accelerator
- Control over steering
- Braking

In addition, it will need output to a display or voice synthesiser

- To talk back to the passengers
- To communicate with other vehicles

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Specification of task environments (cont.)

Agent Type	Performance Measure	Environment	Actuators	Sensors
Taxi driver	Safe, fast, legal, comfortable trip, maximize profits	Roads, other traffic, pedestrians, customers	Steering, accelerator, brake, signal, horn, display	Cameras, sonar, speedometer, GPS, odometer, accelerometer, engine sensors, keyboard

What are the basic **sensors** for the taxi to have

- One or more controllable videocameras, to see the road;
- Infrared/sonar sensors, to detect distances to other cars and obstacles
- A speedometer, safety and to avoid speeding tickets
- An accelerometer, safety to control stability

Specification of task environments (cont.)

To determine the mechanical state of the vehicle

- Engine, fuel, and electrical system sensors

A global positioning system (GPS), to avoid getting lost

A keyboard/microphone for the passenger to request a destination

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Specification of task environments (cont.)

The set of basic PEAS elements vary with agent typology and task

Agent Type	Performance Measure	Environment	Actuators	Sensors
Medical diagnosis system	Healthy patient, reduced costs	Patient, hospital, staff	Display of questions, tests, diagnoses, treatments, referrals	Keyboard entry of symptoms, findings, patient's answers
Satellite image analysis system	Correct image categorization	Downlink from orbiting satellite	Display of scene categorization	Color pixel arrays
Part-picking robot	Percentage of parts in correct bins	Conveyor belt with parts; bins	Jointed arm and hand	Camera, joint angle sensors
Refinery controller	Purity, yield, safety	Refinery, operators	Valves, pumps, heaters, displays	Temperature, pressure, chemical sensors
Interactive English tutor	Student's score on test	Set of students, testing agency	Display of exercises, suggestions, corrections	Keyboard entry

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments

The environment

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments

The range of task environments that might arise in AI is vast

- We can identify a selection of a few dimensions
- To help categorise task environments

They determine the applicability of the main techniques of implementation

These dimensions determine, largely, the right agent design

- The definitions here are informal

Properties of task environments (cont.)

Definition

Fully observable v partially observable

If an agent's sensors give it access to the complete state of the environment at each point in time, then the task environment is fully observable

A task environment is effectively fully observable if sensors detect everything

- That are relevant to the choice of action
- Relevance depends on the performance measure

Remark

Fully observable environments are convenient

⇝ No need to maintain any internal state
⇝ No need to keep track of the world

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments (cont.)

An environment may be **partially observable**: Noisy/inaccurate sensors

- Also because parts of the state are missing from the sensor data

Example

- A vacuum-cleaner agent with only a local dirt sensor cannot tell whether there is dirt in other squares
- An automated taxi cannot see what other drivers are thinking

If the agent has no sensors, then the environment is called **unobservable**

- One might think that in such cases the agent's plight is hopeless
- The goals may still be achievable, sometimes with certainty

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments (cont.)

Definition

Single agent v multi-agent

The distinction between single- and multi-agent environments is intuitive

Example

An agent solving cross-words by itself is in a single-agent environment

- An agent playing chess is in a two-agent environment

There are some subtle issues to be considered

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments (cont.)

We started by describing how an entity may be viewed as an agent

- We have not explained which entities must be viewed as agents

Does agent A (taxi driver) have to treat object B (vehicle) as an agent?

- Or can it treat it merely as an object?
- (behaving according to the laws of physics)

Remark

The key distinction/question

Is the behaviour of agent B best described as maximising a performance measure, whose value depends on agent A's behaviour?

Properties of task environments (cont.)

Example

In chess, opponent entity B is trying to maximise its performance measure

- By the rules of chess, this minimises agent A's performance measure

Thus, chess is a competitive **multi-agent environment**

Example

In the taxi-driving environment, avoiding collisions maximises performance

- This is true for all the agents

Taxi-driving is a partially cooperative **multi-agent environment**

- The environment is also partially competitive
- (As only one car can occupy a parking space)

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments (cont.)

The design in multi-agent and in single-agent environments are different

- In some multi-agent environments, communication often emerges as a rational behaviour
- In some competitive environments, randomised behaviour is rational because it avoids the pitfalls of predictability

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments (cont.)

Definition

Deterministic *v* ***stochastic***

If the next state of the environment is completely determined by the current state and action executed by the agent, then we say the environment is ***deterministic***

- *Otherwise, we say it is* ***stochastic***

There is no uncertainty in fully observable, deterministic environments

- Partially observable environments could appear to be stochastic

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments (cont.)

Most real situations are seriously complex

- It is not possible to keep track of all the unobserved aspects
- For practical purposes, they must be treated as stochastic

Example

Taxi driving as we described it is stochastic

- One can never predict the behaviour of traffic (or else) exactly

The vacuum world as we described it is deterministic

- Variations include stochastic elements (randomly appearing dirt)
- (or an unreliable suction mechanism)

Properties of task environments (cont.)

Definition

*Environment is **uncertain** if it is not fully observable or not deterministic*

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments (cont.)

'Stochastic', in which uncertainty about outcomes quantified by probabilities

Non-deterministic environments

- Actions are characterised by their *possible* outcomes
- No probabilities are attached to them

Remark

Non-deterministic environment descriptions can be associated with performances that require success, for *all possible* outcomes of agent's actions

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments (cont.)

Definition

Episodic v sequential

*In **episodic** environments, agent's experience is divided into atomic episodes*

- *In each episode the agent receives a percept*
- *Then performs a single action*
- ⇝ *The next episode does not depend on the actions taken in previous episodes*

*In **sequential** environments, a decision could affect all future decisions*

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments (cont.)

Example

Many classification tasks are episodic

An agent that has to spot defective parts on an assembly lie

- Each decision based on current part, regardless of previous decisions
- The current decision does not affect whether the next part is faulty

Chess and taxi driving are sequential

- In both cases, short-term actions can have long-term consequences

Properties of task environments (cont.)

Remark

Episodic environments are simpler than sequential environments

- The agent does not need to think ahead

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments (cont.)

Definition

Static v dynamic

*If the environment can change while an agent is deliberating, then the environment is **dynamic** for that agent*

- *Otherwise, it is **static***

Static environments are easy to deal with

- No need to keep looking at the world while deciding on an action
- No need it worry about the passage of time

- Dynamic environments continuously ask what to do
- If no decision yet, that counts as deciding to do nothing

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments (cont.)

If the environment itself does not change with passage of time but agent's performance score does, then the environment is **semi-dynamic**

Example

- Taxi driving is dynamic: Other cars and the taxi itself keep moving while the driving algorithm ponders over what to do next
- Chess, when played with a clock, is semi-dynamic
- Crossword puzzles are static

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments (cont.)

Definition

Discrete v continuous

The discrete/continuous distinction applies to the state of the environment, to the way time is handled, and to the percepts and actions of the agent

Example

Chess has a finite number of distinct states (excluding the clock)

- Chess also has a discrete set of percepts and actions

Taxi driving is a continuous-state and continuous-time problem

- Speed and location of the taxi and of the other vehicles take a range of continuous values, and do so smoothly over time
- Taxi-driving actions are also continuous (steering angles, etc.)
- Input from digital cameras is discrete (strictly speaking), but it is treated as representing continuously varying intensities and locations

Properties of task environments (cont.)

Definition

Known v unknown

Strictly, this distinction refers not to the environment but to the agent's (or designer's) state of knowledge about the 'laws of physics' of the environment

- *In a **known** environment, outcomes (or outcome probabilities if the environment is stochastic) for all actions are given*
- *If the environment is **unknown**, the agent will have to learn how it works in order to make good decisions*

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments (cont.)

The distinction between known/unknown environments is not the same as the one between fully-observable/partially-observable environments

It is possible for a known environment to be partially observable

- In solitaire card games, I know the rules
- Still I am unable to see all the cards

An unknown environment can be fully observable

- In a new video game, the screen may show the entire game state
- Still I do not know what the buttons do until I try them

AUG 30, 2018 _ FC

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Properties of task environments (cont.)

Remark

The hardest case

- Partially observable, multi-agent, stochastic
- Sequential, dynamic, continuous, and unknown

Task Environment	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword puzzle	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with a clock	Fully	Multi	Deterministic	Sequential	Semi	Discrete
Poker	Partially	Multi	Stochastic	Sequential	Static	Discrete
Backgammon	Fully	Multi	Stochastic	Sequential	Static	Discrete
Taxi driving	Partially	Multi	Stochastic	Sequential	Dynamic	Continuous
Medical diagnosis	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Image analysis	Fully	Single	Deterministic	Episodic	Semi	Continuous
Part-picking robot	Partially	Single	Stochastic	Episodic	Dynamic	Continuous
Refinery controller	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Interactive English tutor	Partially	Multi	Stochastic	Sequential	Dynamic	Discrete

AUG 30, 2018 _ FC

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

The agents

Intelligent agents

AUG 30, 2018 _ FC

The agents

We introduced agents by describing them in terms of their behaviour

- The action performed after any given sequence of percepts

⇝ Now we talk about how their inside works

AI designs **agent programs** that implement the **agent function**

⇝ The **mapping from percepts to actions**

Assumption

We assume that the agent program operates on a computing device

- The device has physical sensors and actuators

This general setup is called the **architecture**

- **agent** = **architecture** + **program**

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

The agents (cont.)

- Architecture makes percepts from sensors available to the program
- Architecture feeds program's actions to the actuators

m The program is chosen appropriate for the architecture

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent programs

The agents

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent programs

The agent programs that we discuss all have the same skeleton:

- They take the current percept as input from the sensors
- They return an action to the actuators

Remark

Notice the difference between agent program and agent function

- The **agent program** takes the **current percept** as input
- The **agent function** takes the entire **percept history**

The agent program takes just the current percept as input

- Nothing more is available from the environment

An agent can have actions that depend on the entire percept sequence

- It will have to remember the percepts

Agent programs (cont.)

We describe the agent programs in simple pseudocode language

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent programs (cont.)

Example

The trivial agent program that keeps track of the percept sequence

- It uses it to index into a table of actions to decide what to do

```
function TABLE-DRIVEN-AGENT(percept) returns an action
  persistent: percepts, a sequence, initially empty
              table, a table of actions, indexed by percept sequences, initially fully specified

  append percept to the end of percepts
  action ← LOOKUP(percepts, table)
  return action
```

TABLE-DRIVEN-AGENT program is invoked for each new percept

- It retains the complete percept sequence in memory
- It returns an action each time

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent programs (cont.)

The table represents the agent function the agent program embodies

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*
⋮	⋮
[*A*, *Clean*], [*A*, *Clean*], [*A*, *Clean*]	*Right*
[*A*, *Clean*], [*A*, *Clean*], [*A*, *Dirty*]	*Suck*
⋮	⋮

To build a rational agent, is equivalent to construct the table

- The table contains the appropriate action
- For every possible percept sequence

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent programs (cont.)

A table-driven approach to agent construction is doomed to fail

- Let $\mathcal{P}$ be the set of all possible percepts
- Let T be the entire agent lifetime
- ($\mathcal{T}$, the total number of percepts)

The lookup table will contain $\sum_{t=1}^{T} |\mathcal{P}|^t$ entries

Example

Consider an automated taxi with visual input from a single cam

- Stuff comes in at the rate of $\sim$ 27MBs[1]
- One hour drive: A lookup table of $+10^{250,000,000,000}$ entries

A lookup table for chess (a tiny, well-behaved fragment of world)

- At least 10^{150} entries

[1] 30fps, 640×480 pixels with 24-bit colour information

Agent programs (cont.)

Because of the size of these tables

- No physical agent will have the space to store the table
- The designer would not have time to create the table
- No agent could learn all the right table entries from experience
- The designer has no guidance about how to fill the table entries in

Example

Despite all this, TABLE-DRIVEN-AGENT does do what we want

- It implements the desired agent function

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent programs (cont.)

A key challenge

- Find how to write programs that produce rational behaviour
- From a smallish program, to the extent possible

(Rather than from a vast table)

Examples show that this can be done successfully in many areas

Example

Tables of square roots used by engineers and school children prior to the 70s

- Now replaced by a five-line program on PCs

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent programs (cont.)

The four basic kinds of (intelligent) agent programs

- **Simple reflex agents**
- **Model-based reflex agents**
- **Goal-based agents**
- **Utility-based agents**

Each kind of agent program combines particular components

- This is how they generate actions

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Simple reflex agents

The agents

Simple reflex agents

Simple reflex agent: The simplest kind of agent

- These agents select actions on the basis of the current percept
- Ignore the rest of the percept history

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Simple reflex agents (cont.)

Example

The vacuum agent is a simple reflex agent

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*
⋮	⋮
[*A*, *Clean*], [*A*, *Clean*], [*A*, *Clean*]	*Right*
[*A*, *Clean*], [*A*, *Clean*], [*A*, *Dirty*]	*Suck*
⋮	⋮

```
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
```

Decision on current location and on whether that location contains dirt

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Simple reflex agents (cont.)

The agent program is small compared to the corresponding table

- The largest reduction comes from ignoring the percept history
- (This reduces the number of possibilities from 4^T to 4)

When current square is dirty, action does not depend on location

- (Another, smaller, reduction)

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Simple reflex agents (cont.)

Simple reflex behaviours occur even in more complex environments

Example

It is easy to imagine yourself as the driver of the automated taxi

- If the car in front brakes and its brake lights come on

⇝ Then, you should notice this and initiate braking

Processing is done on visual input to establish the condition

- '`The car in front is braking`'

This triggers established connection in the program to action

⇝ '`Initiate braking`'

We call such a connection a **condition-action rule**

```
if car-in-front-braking then initiate-braking
```

Simple reflex agents (cont.)

Humans have many such connections

- Some of which are learned (car braking)
- Some others are innate reflexes (blinking)

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Simple reflex agents (cont.)

A general and flexible approach to design simple reflex agents

1. First, build a general-purpose interpreter for condition-action rules
2. Then, create rule sets for specific task environments

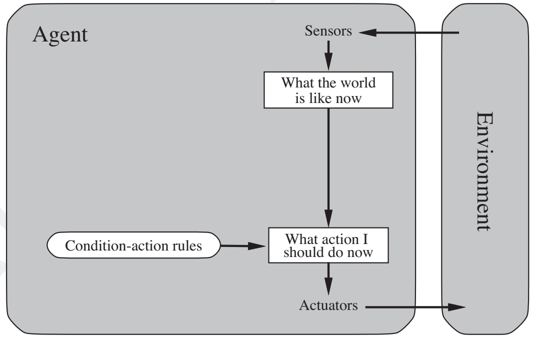

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Simple reflex agents (cont.)

The schematic of the general program shows condition-action rules

- The rules allow to connect percept to action

- Rectangles denote the current internal state of the decision process
- Ovals represent the background info used in the decision process

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Simple reflex agents (cont.)

The simple reflex agent acts according to the rule

- The rule condition matches the current state
- Current state is as defined by the percept

Code 1

```
function SIMPLE-REFLEX-AGENT(percept) returns an action
  persistent: rules, a set of condition–action rules

  state ← INTERPRET-INPUT(percept)
  rule ← RULE-MATCH(state, rules)
  action ← rule.ACTION
  return action
```

The **INTERPRET-INPUT** function

- Generates an abstracted description of current state from percept

The **RULE-MATCH** function

- Returns the first rule among those that matches the state description

Simple reflex agents (cont.)

Simple reflex agents have the admirable property of being simple

- But, they are also of limited intelligence

Remark

Situations in which the simple reflex agent works

If the correct decision can be made on the basis of only the current percept

- Only if the environment is fully observable

Even a tiny bit of un-observability can be cause of serious troubles

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Simple reflex agents (cont.)

Example

- The braking rule assumes condition `car-in-front-braking` to be determined from current percept, a single frame of video

This works if the car in front has a centrally mounted brake light

- Unfortunately, older models have different configurations
- (tail lights, brake lights, and turn-signal lights)

Not always possible to tell from a single image whether the car is braking

A simple reflex agent driving behind such a car would either

- Brake continuously and unnecessarily
- Never brake at all

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Simple reflex agents (cont.)

Example

Issues arise with a vacuum agent

A simple reflex agent that has lost its location sensor, only a dirt sensor

Such an agent has two possible percepts: `[Dirty]` and `[Clean]`

- It can `Suck` in response to `[Dirty]`

What should it do in response to `[Clean]`?

- Moving `Left` fails (forever) if it happens to start in square A
- Moving `Right` fails (forever) if it happens to start in square B

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Simple reflex agents (cont.)

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*
⋮	⋮
[*A*, *Clean*], [*A*, *Clean*], [*A*, *Clean*]	*Right*
[*A*, *Clean*], [*A*, *Clean*], [*A*, *Dirty*]	*Suck*
⋮	⋮

```
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
```

Simple reflex agents (cont.)

For simple reflex agents that operate in partially observable environments

⇝ Infinite loops are often unavoidable

Escape from infinite loops is possible, the agent must randomise its actions

Example

If the vacuum agent perceives [`Clean`]

- Flip a coin to choose between `Left` and `Right`
- The agent will reach the other square in an average of two steps

Then, if that square is dirty, the agent will clean it

- the task will be complete

Artificial intelligence

UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Simple reflex agents (cont.)

Randomised simple reflex agents can outperform deterministic equivalents

Randomised behaviour can be rational in some multi-agent environments

- In single-agent environments, randomisation is usually not rational

Remark

It is a useful trick that helps a simple reflex agent in some cases

- We can do better with sophisticated deterministic agents
- (in most situations)

Artificial intelligence

UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Model-based reflex agents

The agents

Artificial intelligence

UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Model-based reflex agents

The most effective way to handle partial observability

⇝ Keep track of the part of the world that can NOT be seen now

The agent should maintain some sort of **internal state**

- The internal state must depend on the percept history
- It reflects at least some of the unobserved aspects of current state

Model-based reflex agents (cont.)

Example

For the braking problem, internal state is not too extensive

- The previous frame from camera

Detect when two red lights at the vehicle edge go on or off simultaneously

Other driving tasks such as changing lanes

- The agent needs to keep track of where other cars are
- (if it cannot see them all at once)

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Model-based reflex agents (cont.)

The internal state information must be updated as time goes by

It requires two kinds of knowledge to be encoded

- We need info about how the world evolves independently of the agent[2]
- We need info about how the agent's own actions affect the world [3]

Knowledge about 'how the world works' is a **model** of the world

- Can be implemented as Boolean circuits
- Principled theories
- ...

An agent that uses such a model is a **model-based agent**

[2] For example, an overtaking car generally will be closer than it was before

[3] For example, when the agent steers clockwise the car turns to the right, or after driving for five minutes northbound one is about five miles norther

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Model-based reflex agents (cont.)

The structure of the model-based reflex agent with internal state

- It shows how current percept is combined with old internal state

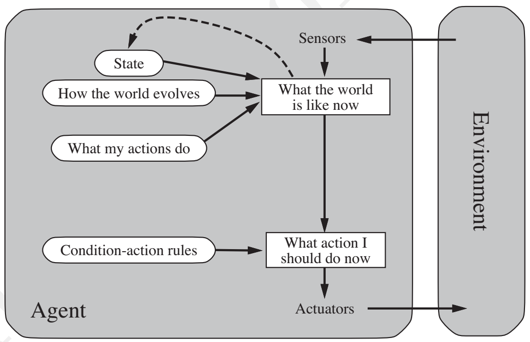

⇝ To generate the updated description of the current state

- Based on the agent's model of how the world works

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Model-based reflex agents (cont.)

A model-based reflex agent keeps track of the current state of the world

- It uses an internal model, then it chooses an action
- The choice is done in the same way as the reflex agent

Code 2

function MODEL-BASED-REFLEX-AGENT(*percept*) **returns** an action
  persistent: *state*, the agent's current conception of the world state
    model, a description of how the next state depends on current state and action
    rules, a set of condition–action rules
    action, the most recent action, initially none

  state ← UPDATE-STATE(*state*, *action*, *percept*, *model*)
  rule ← RULE-MATCH(*state*, *rules*)
  action ← *rule*.ACTION
  return *action*

Interesting part of agent program: Function **UPDATE-STATE**

- It is responsible for creating the new internal state description

Model-based reflex agents (cont.)

The details of how models and states are represented vary widely

- Depending on type of environment and technology used in design

Remark

Seldom possible to determine THE state of partially observable environment

- The box labeled ‘what the world is like now’
- It represents the agent’s ‘best guess(es)’

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Model-based reflex agents (cont.)

Example

The taxi may not be able to see around the truck stopped in front of it

- It can only guess about what may be causing the hold-up

Thus, uncertainty about the current state may be unavoidable

- Yet the agent still has to make a decision

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Model-based reflex agents (cont.)

The internal ‘state’ maintained by a model-based agent

- Not a ‘what the world is like now’ description

Example

A taxi driving back home may have a rule telling it to fill-up with gas

- Unless it has at least 50% tank

‘Driving back home’ may seem to be an aspect of the world state

- It is actually an aspect of the agent’s internal state

The taxi could be in exactly the same place at the same time

- But towards a different destination

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Goal-based agents

The agents

Goal-based agents

Knowing something about the state of the environment is not always enough

- It is still difficult to decide what to do

Example

At a road junction, the taxi can turn left, right, or go straight

- The correct decision depends on where it is trying to get to

The agent needs some **goal** information describing desirable situations

- As well as a current state description

Goal-based agents (cont.)

The agent program can combine scope with the model of the environment

- This allows to choose actions that achieve the goal

This can be achieved by using same info for the model-based reflex agent

Goal-based agents (cont.)

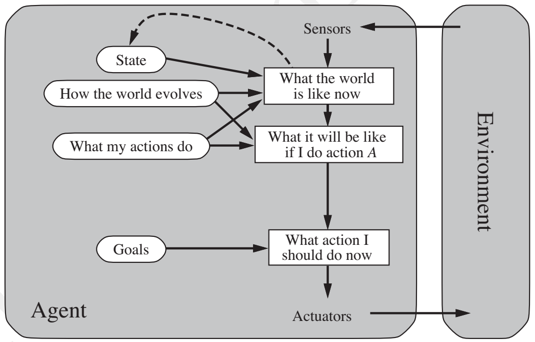

Goal-based agents (cont.)

Sometimes goal-based action selection is straightforward

- When goal satisfaction results straight from a single action

Sometimes goal-based action selection will be more tricky

- The agent may have to consider long sequences of twists and turns
- (To find a way to achieve the goal)

Definition

Search *and* ***planning*** *are the subfields of AI*

- *Finding action sequences that achieve the agent's goals*

Goal-based agents (cont.)

Decision making of this kind differs from condition-action rules

- It involves consideration of the future
- ('What will happen if I do such-and-such?')
- ('Will that make me happy?')

In reflex agent designs this explicit information is not represented

- Its built-in rules map directly from percepts to actions
- The reflex agent brakes when it sees brake lights

Example

A goal-based agent, in principle, could reason

If the car in front has its brake lights on, it will slow down

- The action that achieves the goal of not hitting other cars is to brake

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Goal-based agents (cont.)

The goal-based agent appears less efficient, but it is more flexible

- The knowledge that supports decisions is represented explicitly
- It can be modified

Example

It starts to rain

The agent can update its knowledge of how effectively its brakes will operate

This automatically causes all relevant behaviours to be altered

- To suit the new conditions

For the reflex agent, we must rewrite many condition-action rules

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Goal-based agents (cont.)

The goal-based agent's behaviour can easily be changed

- Wanna go to a different destination
- Specify that destination as the goal

Example

The reflex agent's rules for when to turn and when to go straight

- They will work only for a single destination
- They must all be replaced to go somewhere new

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Utility-based agents

The agents

Utility-based agents

Goals alone are not enough to generate high-quality behaviour

Example

Many action sequences will get the taxi to destination (the goal)

- Some are quicker, safer, more reliable, or cheaper than others

Goals provide a binary distinction between 'happy/unhappy' states

A more general performance measure should allow a comparison of states

- How happy they would make the agent

'Happy' does not sound very scientific, we use the term **utility** instead

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Utility-based agents (cont.)

Performance measures assign a score to any sequence of environment states

- It can easily distinguish between more and less desirable ones
- (ways of getting to the taxi's destination)

Definition

*A **utility function** is kind of an internalisation of the performance*

- *If internal utility function and external performance measure are in agreement, then an agent that chooses actions to maximise its utility will be rational, according to performance*

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Utility-based agents (cont.)

This is not the only way to be rational

- The rational agent program for the vacuum world
- (It has no idea what its utility function is)

Like goal-based agents, a utility-based agent has many advantages

- Flexibility and learning

A utility-based agent can make rational decisions when goals are inadequate

- When there are conflicting goals, only some of which can be achieved (for example, speed v safety), the utility function specifies the appropriate tradeoff
- When there are several goals, none of which can be achieved with certainty, utility provides a way in which the likelihood of success can be weighed against the importance of the goals

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Utility-based agents (cont.)

Partial observability and stochasticity are ubiquitous in the real world

$\rightsquigarrow$ Decision making under uncertainty

Technically,

- Rational utility-based agents maximise the **expected utility**
- Expected utility is the utility the agent expects to derive
- On average, given probabilities and utilities of each outcome

Goal-based agents (cont.)

Any rational agent behaves as if it has a utility function

- Of which it tries to maximise the expected value
- Agents with an explicit utility function can make rational decisions

General-purpose algos that do not depend on the specific utility function

Remark

The 'global' definition of rationality has been modified

- It become a 'local' constraint on rational-agent designs

Goal-based agents (cont.)

Model-based and utility-based agent use a model and an utility function

The utility function measures preferences among states of the world

⇝ Then they choose the action leading to best expected utility

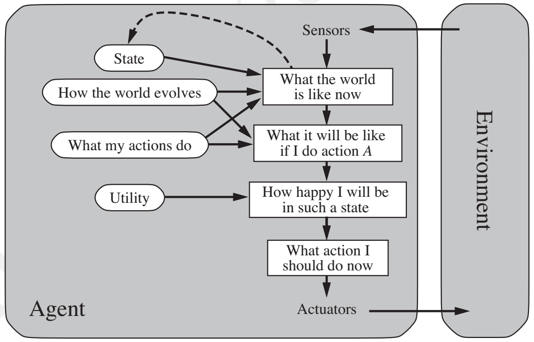

Expected utility is computed by averaging over all possible outcome states

- Weighted by the probability of the outcome

Utility-based agents (cont.)

Is it that simple? We build agents that maximise expected utility?

It's true that such agents would be intelligent, but it's not simple

A utility-based agent has to model and keep track of environment

- These tasks have involved research
- Perception, representation, reasoning, and learning

Choosing the utility-maximising course of action is also difficult

- It requires ingenious algorithms
- Perfect rationality may still be unachievable
- (computational complexity)

Learning agents

The agents

Learning agents

We described agent programs with various methods for selecting actions

- How do they come into being?

Turing (1950) considered the idea of programming machines by hand

- (One of his famous early papers)
- He estimates how much work this might take
- He concludes '*Some more expeditious method seems desirable*'

He proposed to build learning machines and then to teach them

Remark

This today is the preferred method for creating state-of-the-art systems

Learning allows the agent to operate in initially unknown environments

- Improve competence beyond the initial knowledge alone

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Learning agents (cont.)

A **learning agent** can be divided into four conceptual components

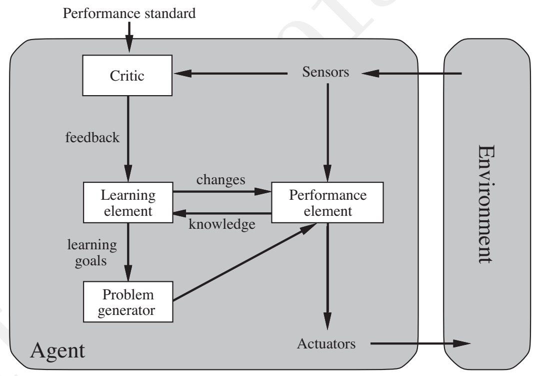

- **Learning element**, it is esponsible for making improvements
- **Performance element**, it is responsible for selecting external actions

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Learning agents (cont.)

The performance element is what we considered to be the agent

- It takes in percepts and decides on actions

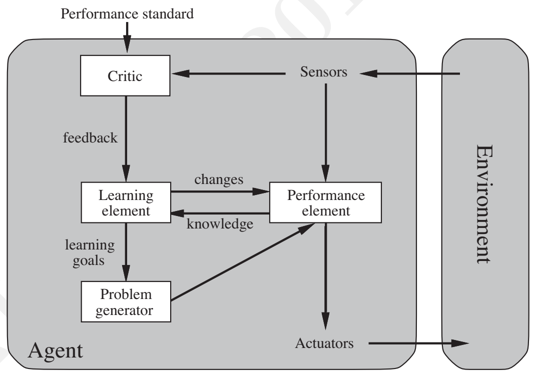

The learning element uses feedback from the **critic** element on how it is doing and determines how the performance element should be modified

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Learning agents (cont.)

The design of learning agents depends on the design of performance elements

Consider the design of an agent that learns a certain capability

- The first question is NOT 'How am I going to get it to learn this?'
- ⇝ 'What kind of performance element will my agent need to do this, once it has learned how?'

Given an agent design, learning mechanisms can be constructed

- We can improve every part of the agent

Learning agents (cont.)

Definition

The critic tells the learning element how well the agent is doing

- *With respect to a fixed performance standard*

The critic is necessary as percepts provide no indication of success

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Learning agents (cont.)

Example

A chess program could receive a percept indicating that it has checkmated its opponent, but it needs a performance standard to know this is good

- The percept itself does not say so

It is important that the performance standard be fixed

- Conceptually, one should think of it as being outside the agent
- The agent must not modify it to fit its own behaviour

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Learning agents (cont.)

The last component of a learning agent is the **problem generator**

⇝ Suggest actions leading to new/informative experiences

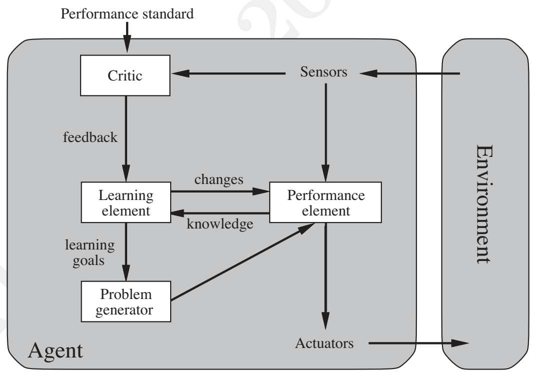

Artificial intelligence
UFC/DC CK0031/CK0248 2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Learning agents (cont.)

If the performance element had its way, it would keep doing best actions

- Given what it knows

If the agent is willing to explore, do suboptimal actions in the short run

- Then, it might discover much better actions for the long run

The problem generator's job is to suggest exploratory actions

Learning agents (cont.)

To make the design more concrete, we return to the automated taxi

The performance element

- Whatever collection of knowledge/procedures for selecting actions

The taxi drives using the performance element

- The critic observes the world
- It also passes information to the learning element

Learning agents (cont.)

Example

The taxi makes a quick left turn across three lanes of traffic

- The critic observes the language used by other drivers

From this experience, the learning element is able to formulate a rule

$\rightsquigarrow$ This was a bad action

The performance element is modified, by installation of the new rule

The generator may identify areas for improvement, suggest experiments

- Trying out the brakes on different road surfaces/conditions

Learning agents (cont.)

The learning element can make changes to the ‘knowledge’ components

- Simplest cases involve learning directly from percept sequence
- Observation of successive states can allow to learn
  $\rightsquigarrow$ ‘How the world evolves’
- Observation of the results of its actions can allow to learn
  $\rightsquigarrow$ ‘What my actions do’

Learning agents (cont.)

Example

If the taxi exerts a certain braking pressure when driving on a wet road

- Then, it will soon find out how much deceleration is achieved

These two tasks are more difficult if the environment is partially observable

Learning agents (cont.)

These forms of learning need no access to external performance standards

- The standard way: Make predictions that agree with experiment

The situation is more complex for a utility-based agent

- As they wish to learn an utility information

Example

The taxi-driving agent gets no tips from passengers who been shaken up

- The external performance standard must inform the agent that the loss of tips is a negative contribution to its overall performance
- The agent might be able to learn
- ⇝ Violent manoeuvres do not contribute to its own utility

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Learning agents (cont.)

Remark

The performance standard distinguishes part of the incoming percept

- A **reward** (or **penalty**)
- This part provides feedback on the quality of behaviour

Hard-wired performance standards such as pain/hunger in animals

- They can be understood in this way

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Learning agents (cont.)

In summary, we consider agents with a variety of components

- Components can be represented in many ways
- They are (in) the agent program

There appears to be variety among learning methods

- There is a single unifying theme

Learning is the process of modification of each component of the agent

- Bring components into agreement with feedback information
- ⇝ Thereby improving the overall performance of the agent

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2

Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent components

The agents

Agent components

We have described agent programs (in high-level terms)

Various components, whose function is to answer questions

- 'What is the world like now?'
- 'What action should I do now?'
- 'What do my actions do?'

'How do these components work?'

- How the components can represent the environment

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent components (cont.)

Representations along an axis of increasing complexity, expressive power

- **Atomic**, **factored**, and **structured**

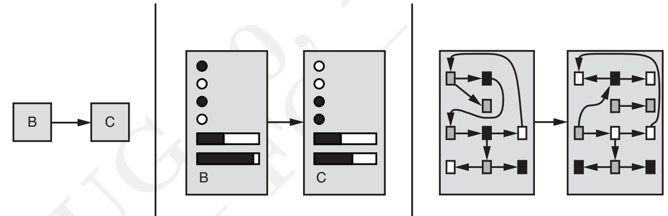

(a) Atomic (b) Factored (b) Structured

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent components (cont.)

Consider the agent component dealing with 'What my actions do?'

- This component describes the changes that might occur in the environment as the result of taking an action

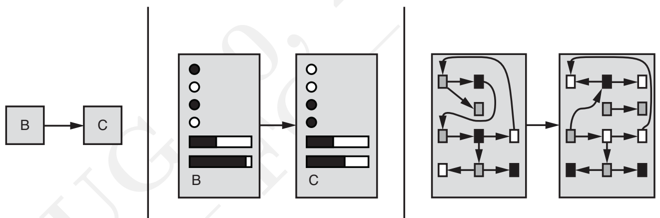

(a) Atomic (b) Factored (b) Structured

Three ways of schematically depicting states and transitions

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent components (cont.)

Definition

*In an **atomic representation** each state of the world is indivisible*

- *It has no internal structure*

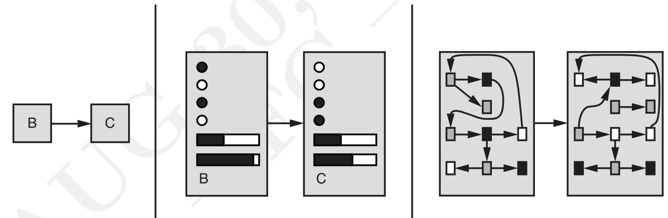

(a) Atomic (b) Factored (b) Structured

Agent components (cont.)

Example

The problem of finding a route from one end of a country to the other

⇝ Via some sequence of cities

Sufficient to reduce the state of the world to the city the agent is in

⇝ For the purposes of solving this problem

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent components (cont.)

A single atom, a 'black box'

The discernible property is being identical to or different from another box

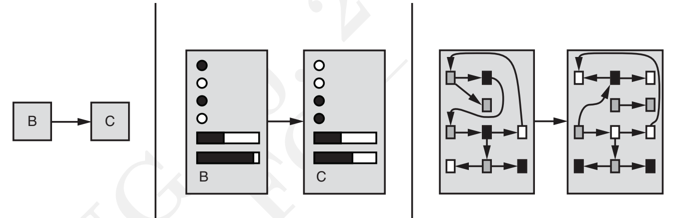

(a) Atomic (b) Factored (b) Structured

Algorithms underlying search and game-playing, hidden Markov models, and Markov decision processes work with atomic representations

- They treat representations as if they were atomic

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent components (cont.)

Now consider a higher-fidelity description for the same problem

Example

We need to be concerned with more than atomic location

We need to pay attention to things

- How much gas is in the tank
- Current GPS coordinates
- Whether or not the oil warning light is working
- How much spare change we have for toll crossings
- What station is on the radio
- ... and so on

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent components (cont.)

Definition

*A **factored representation** splits up each state into a fixed set of variables or attributes, each of which can have a value*

Agent components (cont.)

Different atomic states have nothing in common (different black boxes)

Different factored states can share some attributes (same GPS location)

- Other states cannot (lots of gas or no gas)

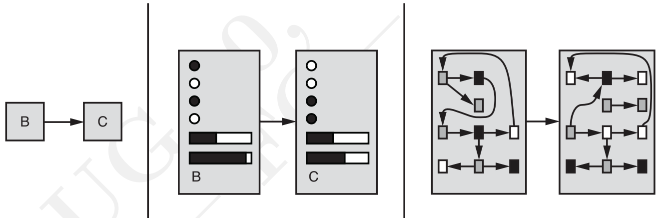

(a) Atomic (b) Factored (b) Structured

It is easier to work out how to turn one state into another

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent components (cont.)

With factored representations, we can also represent uncertainty

Many important areas of AI are based on factored representations

- Constraint satisfaction algorithms
- Propositional logic
- Planning
- Bayesian networks
- Machine learning algorithms

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent components (cont.)

At times, we need to understand the world as having things in it

- Things are related to each other
- Not just variables with values

Example

1. There is a truck ahead of us
2. It is reversing into the driveway of a dairy farm
3. A cow has got loose
4. The cow is blocking the way

A factored representation is unlikely to be have a true/false type attribute

`Truck-Ahead-Backing-Into-Dairy-Farm-Driveway-Blocked-By-Loose-Cow`

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent components (cont.)

In a **structured representation**

Objects (cows, trucks, ...) and their relationships can be described explicitly

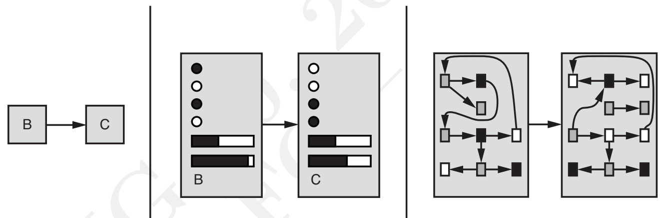

(a) Atomic (b) Factored (b) Structured

A state includes objects, each of which may have

- Relations with other objects
- Attributes of its own

Agent components (cont.)

Structured representations underlie relational databases, first-order logic, probability models, knowledge-based learning and natural language

- Almost everything that humans express in natural language
- Objects and their relationships

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent components (cont.)

The axis of atomic ⇒ factored ⇒ structured representations

⇝ The axis of increasing **expressiveness**

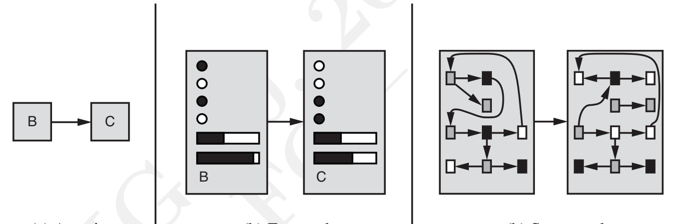

(a) Atomic (b) Factored (b) Structured

Roughly speaking, a more expressive representation can capture everything a less expressive one can (at least as concisely), plus some more

- Often, the more expressive language is much more concise

Artificial intelligence
UFC/DC
CK0031/CK0248
2018.2
Agents and environment
Rationality and good behaviour
Rationality
Omniscience, learning and autonomy
The environment
Specification of task environments
Properties of task environments
The agents
Agent programs
Simple reflex agents
Model-based reflex agents
Goal-based agents
Utility-based agents
Learning agents
Agent components

Agent components (cont.)

Example

The rules of chess

⇝ A page or two

- A structured-representation language, such as first-order logic

⇝ Thousands of pages

- A factored-representation language, such as propositional logic

Reasoning/learning gets more complex with power of representation

Remark

How to benefit from expressive representations while avoiding drawbacks

⇝ Need to operate at all points along the axis simultaneously