



EDA132: Applied Artificial Intelligence Agents (Chapter 2 of AIMA)

Jacek Malec

Dept. of Computer Science, Lund University, Sweden

January 19th, 2016




Plan for the 2nd hour

- What is an agent?
- PEAS (Performance measure, Environment, Actuators, Sensors)
- Agent architectures.
- Environments
- Multi-agent systems.



Agent



Agents include humans, robots, web-crawlers, thermostats, etc.


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

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

The *agent program* runs on a physical *architecture* to produce f .



The vacuum-cleaning world



Percepts: location and contents, e.g. $\langle A, \text{Dirty} \rangle$

Actions: *Left*, *Right*, *Suck*, *NoOp*



A vacuum-cleaning 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 vacuum-cleaning 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
...	...

```
function Reflex_Vacuum_Agent (location, status)
  if status == Dirty then return Suck
  if location == A then return Right
  if location == B then return Left
```



A vacuum-cleaning 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
...	...

```
function Reflex_Vacuum_Agent (location, status)
  if status == Dirty then return Suck
  if location == A then return Right
  if location == B then return Left
```

What is the *RIGHT* function?



Rationality

Fixed performance measure evaluates the environment sequence:

- one point per square cleaned up in time T ?
- one point per clean square per time step, minus one per move?
- penalize for $> k$ dirty squares?



Rationality

Fixed performance measure evaluates the environment sequence:

- one point per square cleaned up in time T ?
- one point per clean square per time step, minus one per move?
- penalize for $> k$ dirty squares?

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



Rationality

Fixed performance measure evaluates the environment sequence:

- one point per square cleaned up in time T ?
- one point per clean square per time step, minus one per move?
- penalize for $> k$ dirty squares?

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

Rational is not omniscient
as percepts may not supply all relevant information



Rationality

Fixed performance measure evaluates the environment sequence:

- one point per square cleaned up in time T ?
- one point per clean square per time step, minus one per move?
- penalize for $> k$ dirty squares?

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

Rational is not omniscient
as percepts may not supply all relevant information
Rational is not clairvoyant
as action outcomes may not be as expected



Rationality

Fixed performance measure evaluates the environment sequence:

- one point per square cleaned up in time T ?
- one point per clean square per time step, minus one per move?
- penalize for $> k$ dirty squares?

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

Rational is not omniscient
as percepts may not supply all relevant information
Rational is not clairvoyant
as action outcomes may not be as expected
Hence, rational is not necessarily successful



A rational agent

[Wooldridge, 2000]

An agent is said to be *rational* if it chooses to perform actions that are in its own best interests, given the beliefs it has about the world.

Properties of rational agents:

- Autonomy (they decide);
- Proactiveness (they try to achieve their goals);
- Reactivity (they react to changes in the environment);
- Social ability (they negotiate and cooperate with other agents).



PEAS

- PEAS: Performance measure, Environment, Actuators, Sensors
- Must first specify the setting for intelligent agent design
- Consider, e.g., the task of designing an automated taxi driver:
 - Performance measure
 - Environment
 - Actuators
 - Sensors



PEAS, example

AUTOMATED TAXI DRIVER:

- Performance measure: Safe, fast, legal, comfortable trip, maximize profits
- Environment: Roads, other traffic, pedestrians, customers
- Actuators: Steering, accelerator, brake, signal, horn
- Sensors: Cameras, radars, speedometer, GPS, odometer, engine sensors, car-human interface



Autonomous agents

Can make decisions on their own.

Why do they need to? Because of the following properties of real environments (cf. Russell and Norvig):

- the real world is inaccessible (partially observable);
- the real world is nondeterministic (stochastic, sometimes strategic);
- the real world is nonepisodic (sequential);
- the real world is dynamic (non-static);
- the real world is continuous (non-discrete).

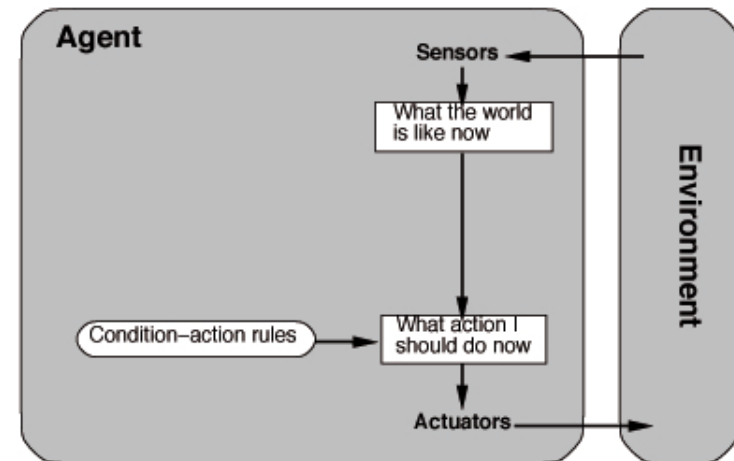


Agent taxonomy

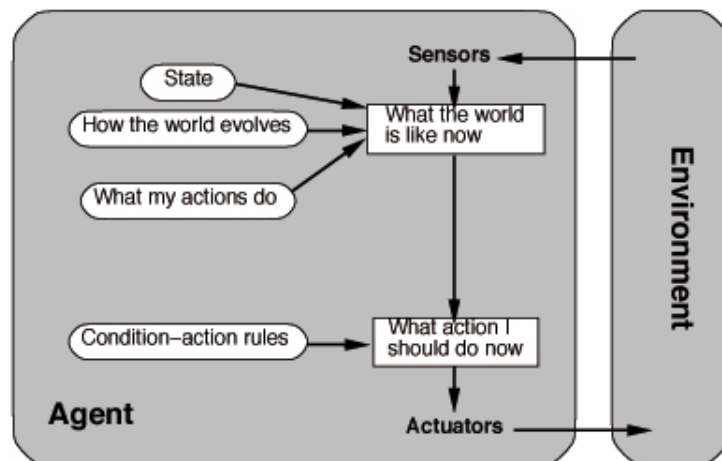
- simple reflex agents
- reflex agents with state
- goal-based agents
- utility-based agents
-
- learning agents - independent property from the list above



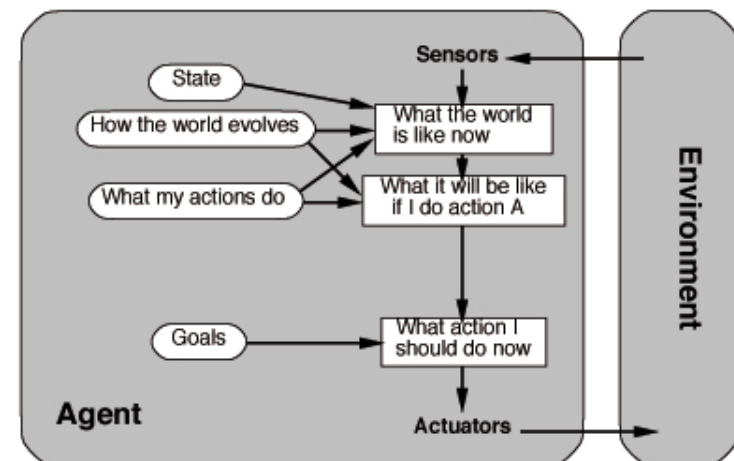
Simple reflex agent



Reflex agent with state

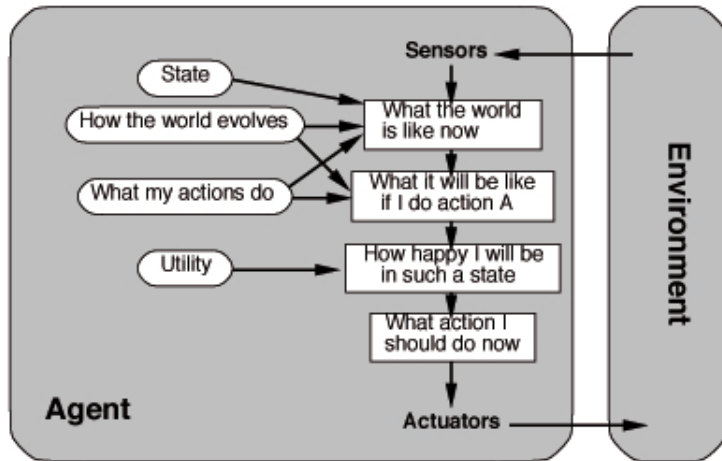


Goal-based agent

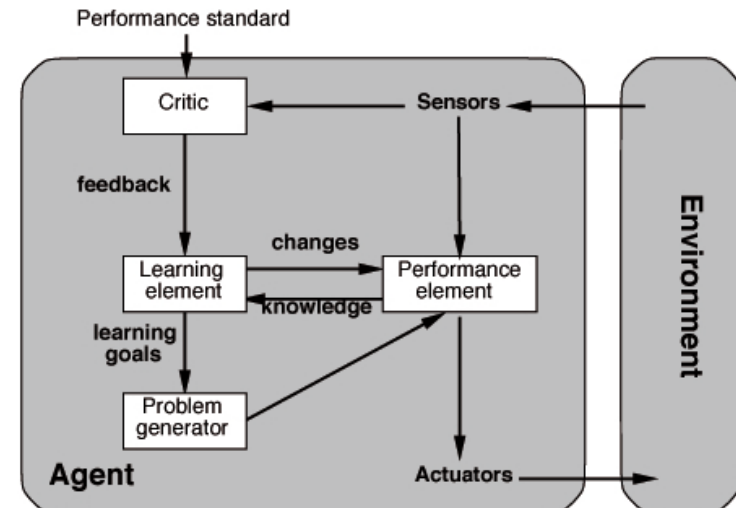




Utility-based agent



Learning agent



A bit more on rationality

Rationality is a very powerful assumption.


It allows us to compute things we wouldn't otherwise be able to dream of!

40 first years of AI were based solely on this assumption.

What do you think about?



Rodney Brooks, 1985





Subsumption

- horizontal vs. vertical decomposition
- a system is more than a sum of its parts (emergent intelligence)
- each behaviour can sense the environment and generate a physical action



Physical Grounding Hypothesis

- situatedness
"the world is its own best model"
- embodiment
- intelligence
"intelligence is determined by the dynamics of interaction with the world"
- emergence
"intelligence is in the eye of the observer"



Multi-agent systems

Interesting for a number of reasons:

- performance: many agents may do the job faster, with less effort
Sometimes only many agents can do the job (if they are heterogenous or if the deadline is hard)
- reliability, robustness: when one agent fails, the rest may do the job
- adaptivity: agents exposed to different environmental conditions can learn appropriately (and even communicate the results to others)

Note special case of faults: communication faults not occurring in a single-agent case



Interaction, Coordination, Cooperation

- *Interaction*: common resources
 - antagonistic (incompatibility of goals)
 - non-antagonistic
- *Coordination*: planning for use of common resources
- *Cooperation*: planning for maximisation of utility
 - *eusocial* behaviour (innate, McFarland)
 - cooperative behaviour (selfish agents maximising personal utility)



Summary

Agents interact with *environments* through *actuators and sensors*



Summary

Agents interact with *environments* through *actuators and sensors*

The *agent function* describes what the agent does in all circumstances



Summary

Agents interact with *environments* through *actuators and sensors*

The *agent function* describes what the agent does in all circumstances

The *performance measure* evaluates the environment sequence



Summary

Agents interact with *environments* through *actuators and sensors*

The *agent function* describes what the agent does in all circumstances

The *performance measure* evaluates the environment sequence

A *perfectly rational* agent maximizes expected performance



Summary

Agents interact with *environments* through *actuators and sensors*

The *agent function* describes what the agent does in all circumstances

The *performance measure* evaluates the environment sequence

A *perfectly rational* agent maximizes expected performance

Agent programs implement (some) agent functions



Summary

Agents interact with *environments* through *actuators and sensors*

The *agent function* describes what the agent does in all circumstances

The *performance measure* evaluates the environment sequence

A *perfectly rational* agent maximizes expected performance

Agent programs implement (some) agent functions

PEAS descriptions define task environments



Summary

Agents interact with *environments* through *actuators and sensors*

The *agent function* describes what the agent does in all circumstances

The *performance measure* evaluates the environment sequence

A *perfectly rational* agent maximizes expected performance

Agent programs implement (some) agent functions

PEAS descriptions define task environments

Environments are categorized along several dimensions:
observable? deterministic? episodic? static? discrete?
single-agent?



Summary

Agents interact with *environments* through *actuators and sensors*

The *agent function* describes what the agent does in all circumstances

The *performance measure* evaluates the environment sequence

A *perfectly rational* agent maximizes expected performance

Agent programs implement (some) agent functions

PEAS descriptions define task environments

Environments are categorized along several dimensions:
observable? deterministic? episodic? static? discrete?
single-agent?

Several basic agent architectures exist:
reflex, reflex with state, goal-based, utility-based