


ARTIFICIAL INTELLIGENCE

Russell & Norvig

Chapter 2: Intelligent Agents

Agents


- “An **agent** is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators.”



Percepts

- **Percepts** are the agent's "perceptual inputs" at any instance.
- The **Percept History** is the complete sequence of everything the agent has perceived.
- *An agent's choice of action at any given instant can depend on the entire percept sequence observed to date, but not anything it hasn't perceived.*
- Agent function maps from percept history to actions


Vacuum cleaner world example



- Percepts: location and contents, e.g., [A,Dirty]
- Actions: *Left*, *Right*, *Suck*, *NoOp*

A vacuum-cleaner 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
...	



Rationality

- A **rational agent** is one that does the right thing.
 - Every entry in the table is filled out correctly.
- What is the right thing?
 - Approximation: the most *successful* agent.
 - *Measure of success?*
- Performance measure should be objective
 - E.g. the amount of dirt cleaned within a certain time.
 - E.g. how clean the floor is.
 - ...
- *Performance measure according to what is wanted in the environment instead of how the agents should behave.*

Rational Agent

- What is rational depends on:
 - The performance measure that defines the criterion of success
 - The agent's prior knowledge of the environment
 - The actions that the agent can perform
 - The agent's percept sequence to date

- Definition of a rational agent:

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

More rationality!

- Rationality \neq omniscience
 - An omniscient agent knows the actual outcome of its actions.
- Rationality \neq perfection
 - Rationality maximizes *expected* performance, while perfection maximizes *actual* performance.
- Components required for rationality
 - Exploration (information gathering)
 - Learning (go beyond *apriori* knowledge)
 - Autonomy (independent of prior knowledge)

Environments

- To design a rational agent we must specify its task environment
- PEAS description of the environment:
 - Performance
 - Environment
 - Actuators
 - Sensors

Automated Taxi Driver Example

- PEAS Environment:
 - *Performance measure*: Safe, fast, legal, comfortable trip, maximize profits
 - *Environment*: Roads, other traffic, pedestrians, customers
 - *Actuators*: Steering wheel, accelerator, brake, signal, horn, display
 - *Sensors*: Cameras, sonar, speedometer, GPS, odometer, engine sensors, keyboard

Agent: Medical diagnosis system

Performance Measure	Healthy patient, reduced cost
Environment	Patient, hospital staff
Actuators	Display of questions, tests, diagnoses, treatments, referrals
Sensors	Keyboard entry of symptoms, findings, patient's answers

Agent: Part-picking robot

Performance Measure	Percentage of parts in correct bins
Environment	Conveyor belt with parts, bins
Actuators	Jointed arm and hand
Sensors	Camera, joint angle sensors

Environment Types

- **Fully observable** (vs. partially observable): An agent's sensors give it access to the complete state of the environment at each point in time.
- **Deterministic** (vs. stochastic): The next state of the environment is completely determined by the current state and the action executed by the agent. (If the environment is deterministic except for the actions of other agents, then the environment is **strategic**)
- **Episodic** (vs. sequential): The agent's experience is divided into atomic "episodes" (each episode consists of the agent perceiving and then performing a single action), and the choice of action in each episode depends only on the episode itself.

Environment types, continued

- **Static** (vs. dynamic): The environment is unchanged while an agent is deliberating. (The environment is **semidynamic** if the environment itself does not change with the passage of time but the agent's performance score does)
- **Discrete** (vs. continuous): A limited number of distinct, clearly defined percepts and actions.

Example task environments

	Crossword puzzle	Chess w/ clock	Poker	Back-gammon	Taxi driving	Medical Diagnosis
Observable	Fully	Fully	Partial	Fully	Partial	Partial
Agents	Single	Multi	Multi	Multi	Multi	Single
Deterministic	Determ	Determ	Stoch	Stoch	Stoch	Stoch
Episodic	Seq	Seq	Seq	Seq	Seq	Seq
Static	Static	Semi	Static	Static	Dyn	Dyn
Discrete	Disc	Disc	Disc	Disc	Cont	Cont

Simple vs. Real Environments

- The simplest environment is
 - Fully observable, deterministic, episodic, static, discrete and single-agent
- Most real situations are:
 - Partially observable, stochastic, sequential, dynamic, continuous and multi-agent