Rational Agents (Ch. 2)

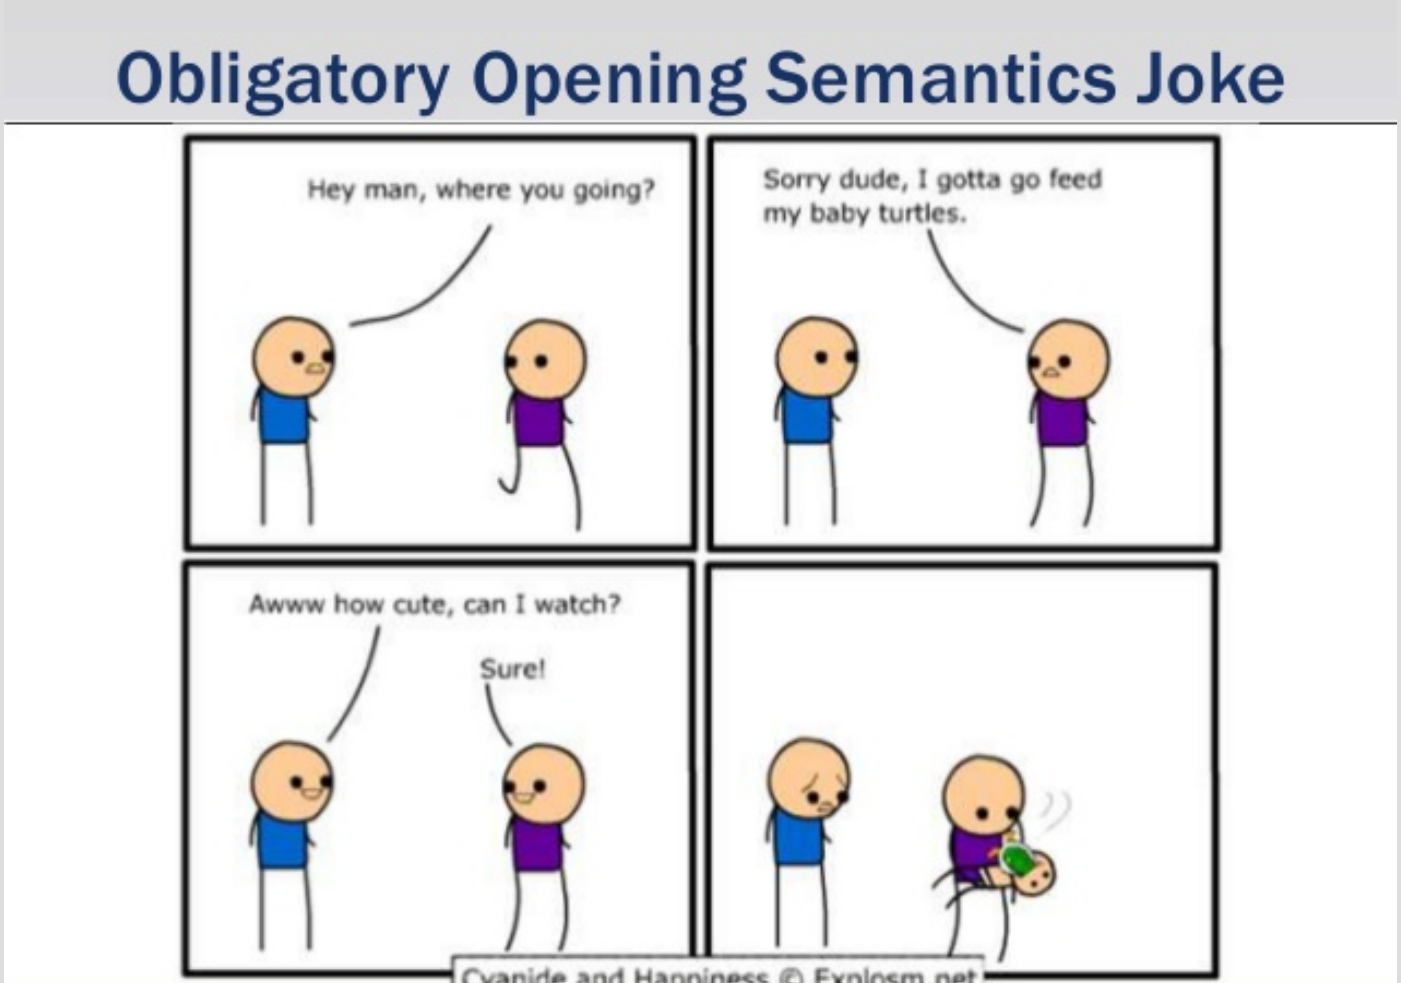

Rational agent

Remember vacuum problem?

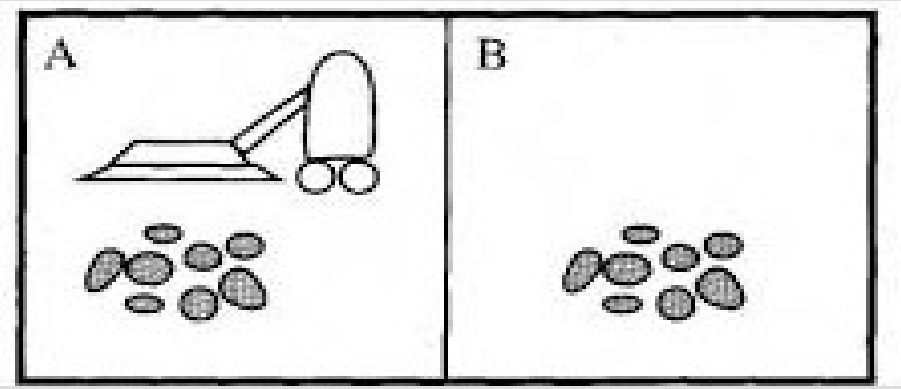

Agent program:

if [Dirty], return [Suck]

if at [room A], return [move right]

if at [room B], return [move left]

Agent models

Can also classify agents into four categories:

1. Simple reflex
2. Model-based reflex
3. Goal based
4. Utility based

Top is typically simpler and harder to adapt to similar problems, while bottom is more general representations

Agent models

1. Simple reflex = “plans” a single move using only current information

2. Model-based reflex = “plans” a single move using current and (some) past information

3. Goal based = plans multiple moves (until goal) using current and past information

4. Utility based = “goals” have different values

Agent models

What is the agent model of particles?

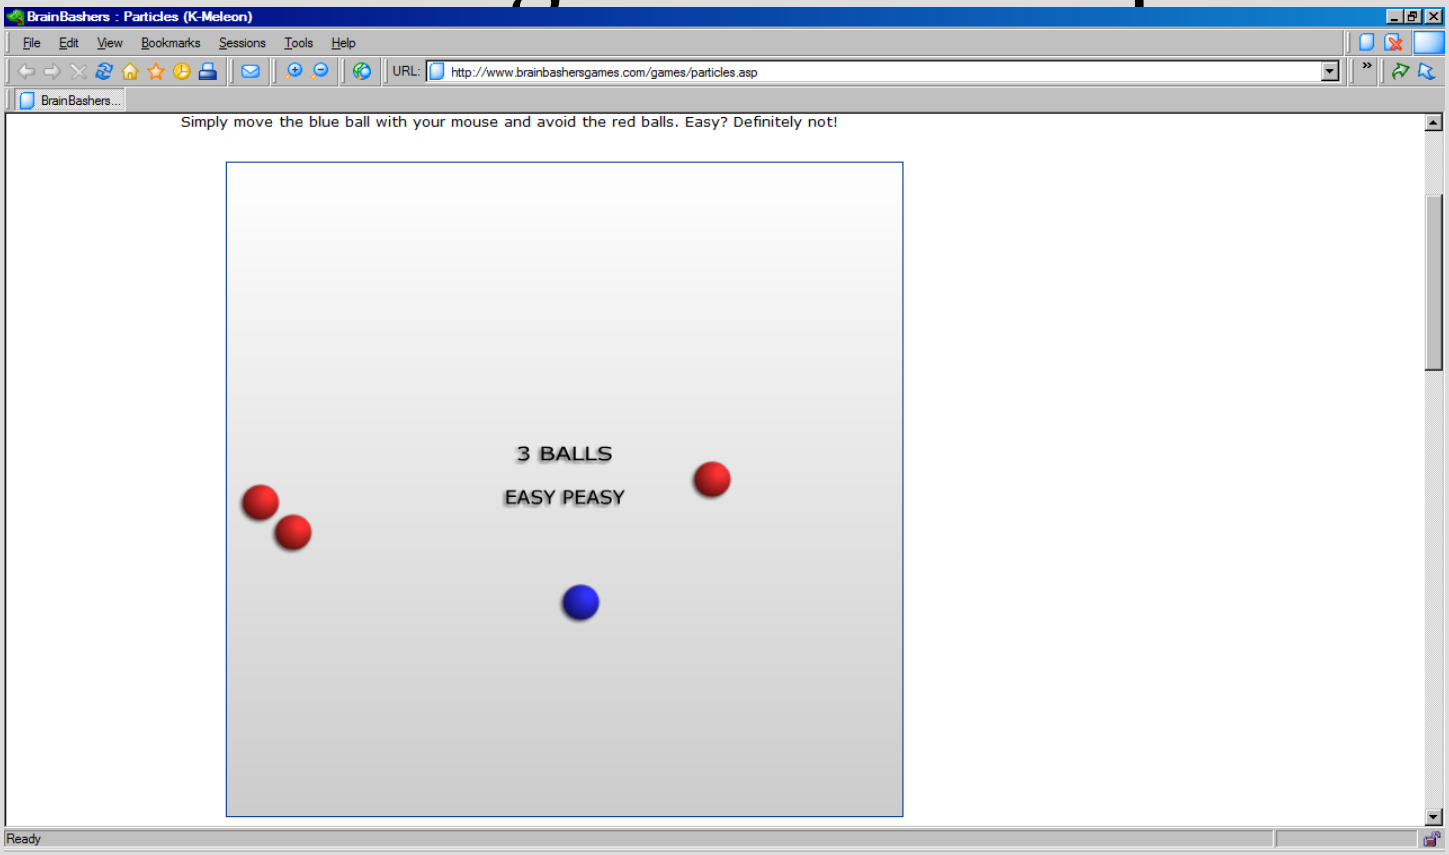

Think of a way to improve the agent and describe what model it is now

Environment classification

Environments can be further classified on the following characteristics:(right side harder)

1. Fully vs. partially observable
2. Single vs. multi-agent
3. Deterministic vs. stochastic
4. Episodic vs. sequential
5. Static vs. dynamic
6. Discrete vs. continuous
7. Known vs. unknown

Environment classification

In a fully observable environment, agents can see every part.

Agents can only see part of the environment if it is partially observable

Full ↗

Partial →

Environment classification

If your agent is the only one, the environment is a single agent environment

More than one is a multi-agent environment (possibly cooperative or competitive)

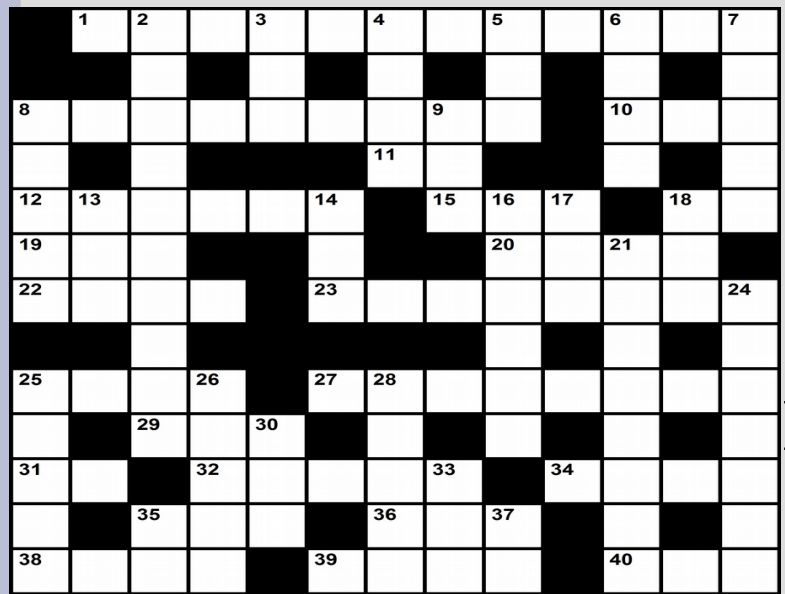

←single

multi→

Environment classification

If your state+action has a known effect in the environment, it is deterministic

If actions have a distribution (probability) of possible effects, it is stochastic

← deterministic

stochastic →

Environment classification

An episodic environment is where the previous action does not effect the next observation (i.e. can be broken into independent events)

If there is the next action depends on the previous, the environment is sequential

← episodic

sequential →

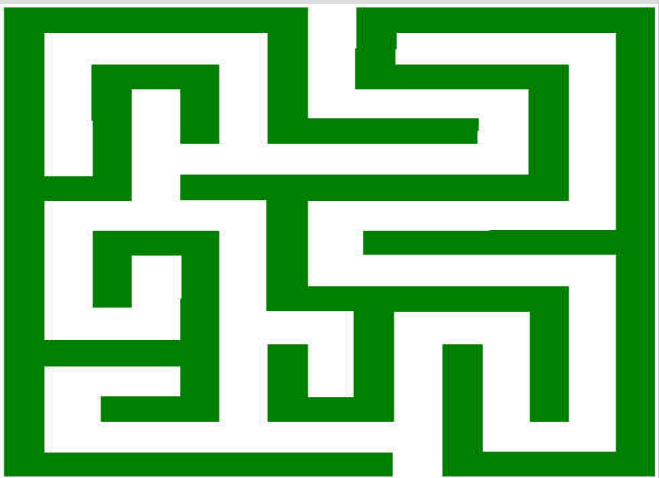

Environment classification

If the environment only changes when you make an action, it is static

a dynamic environment can change while your agent is thinking or observing

static

dynamic

Environment classification

Discrete = separate/distinct (events)
Continuous = fluid transition (between events)

This classification can applies: agent's percept and actions, environment's time and states

discrete (state)

continuous (state)

Environment classification

Known = agent's actions have known effects on the environment

Unknown = the actions have an initially unknown effect on the environment (can learn)

know how to stop

do not know how to stop

Environment classification

1. Fully vs. partially observable = how much can you see?
2. Single vs. multi-agent
   = do you need to worry about others interacting?
3. Deterministic vs. stochastic
   = do you know (exactly) the outcomes of actions?
4. Episodic vs. sequential
   = do your past choices effect the future?
5. Static vs. dynamic = do you have time to think?
6. Discrete vs. continuous
   = are you restricted on where you can be?
7. Known vs. unknown
   = do you know the rules of the game?

Environment classification

Some of these classifications are associated with the state, while others with the actions

State: Actions:

1. Fully vs. partially observable
2. Single vs. multi-agent
3. Deterministic vs. stochastic
4. Episodic vs. sequential
5. Static vs. dynamic
6. Discrete vs. continuous
7. Known vs. unknown

Environment classification

Pick a game/hobby/sport/pastime/whatever and describe both the PEAS and whether the environment/agent is:

1. Fully vs. partially observable
2. Single vs. multi-agent
3. Deterministic vs. stochastic
4. Episodic vs. sequential
5. Static vs. dynamic
6. Discrete vs. continuous
7. Known vs. unknown

Environment classification

Agent type	Performance	Environment	Actuators	Sensors
Particles	time alive	boarder, red balls	move mouse	screen-shot

Fully observable, single agent, deterministic, sequential (halfway episodic), dynamic, continuous (time, state, action, and percept), known (to me!)

State structure

An atomic state has no sub-parts and acts as a simple unique identifier

An example is an elevator:
Elevator = agent (actions = up/down)
Floor = state

In this example, when someone requests the elevator on floor 7, the only information the agent has is what floor it currently is on

State structure

Another example of an atomic representation is simple path finding:
If we start at Koffman, how would you get to Keller's CS office?

Go E. -> Cross N @ Ford & Amundson -> Walk to E. KHKH -> K. Stairs -> CS office

The words above hold no special meaning other than differentiating from each other

State structure

A factored state has a fixed number of variables/attributes associated with it

You can then reason on how these associated values change between states to solve problem

Can always “un-factor” and enumerate all possibilities to go back to atomic states, but might be too exponential or lose efficiency

State structure

Structured states simply describe objects and their relationship to others

Suppose we have 3 blocks: A, B and C
We could describe: A on top of B, C next to B

A factored representation would have to enumerate all possible configurations of A, B and C to be as representative

State structure

We will start using structured approaches when we deal with logic:

Summer implies Warm
Warm implies T-Shirt

The current state might be:
!Summer (¬Summer)
but the states have intrinsic relations between each other (not just actions)