Introducing Uncertainty

(It is not the world that is imperfect,
it is our knowledge of it)

R&N: Chap. 13

Slides from Jean-Claude Latombe at Stanford University
(used with permission)

- So far, we have assumed that:
  - World states are perfectly observable,
    $\rightarrow$ the current state is exactly known
  - Action representations are perfect,
    $\rightarrow$ states are exactly predicted

- We will now investigate how an agent can cope with imperfect information

Sources of Uncertainty

The Real World and its Representation

3x3 matrix filled with 1, 2, .., 8, and 'empty'

Agent's conceptualization
(→ representation language)

Real world

8-puzzle

The Real World and its Representation

Logic sentences using propositions like Block(A), On(A,B), Handempty, ... and connectives

Agent's conceptualization
(→ representation language)

Real world

Blocks world

Who provides the representation language?

- The agent's designer
- As of today, no practical techniques exist allowing an agent to autonomously abstract features of the real world into useful concepts and develop its own representation language using these concepts
- Inductive learning techniques are steps in this direction, but much more is needed
- The issues discussed in the following slides arise whether the representation language is provided by the agent's designer or developed over time by the agent

First Source of Uncertainty: The Representation Language

- There are many more states of the real world than can be expressed in the representation language
- So, any state represented in the language may correspond to many different states of the real world, which the agent can't represent distinguishably

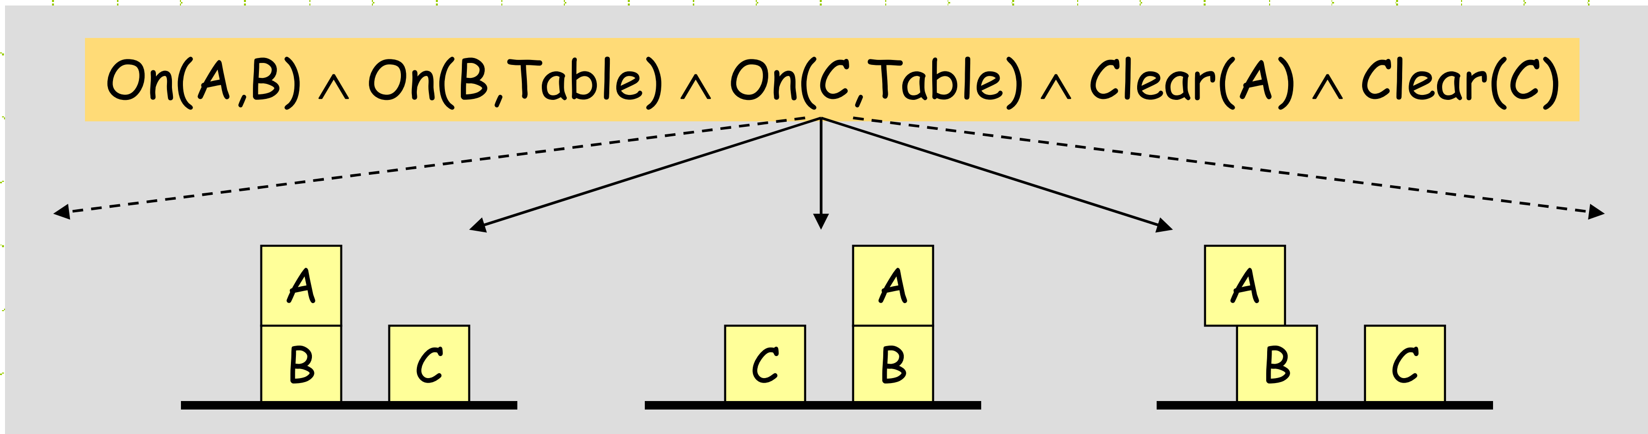

First Source of Uncertainty: The Representation Language

- 6 propositions On(x,y), where x, y = A, B, C and $x \neq y$
- 3 propositions On(x,Table), where x = A, B, C
- 3 propositions Clear(x), where x = A, B, C
- At most 2^{12} states can be distinguished in the language [in fact much fewer, because of state constraints such as $On(x,y) \rightarrow \neg On(y,x)$]
- But there are infinitely many states of the real world

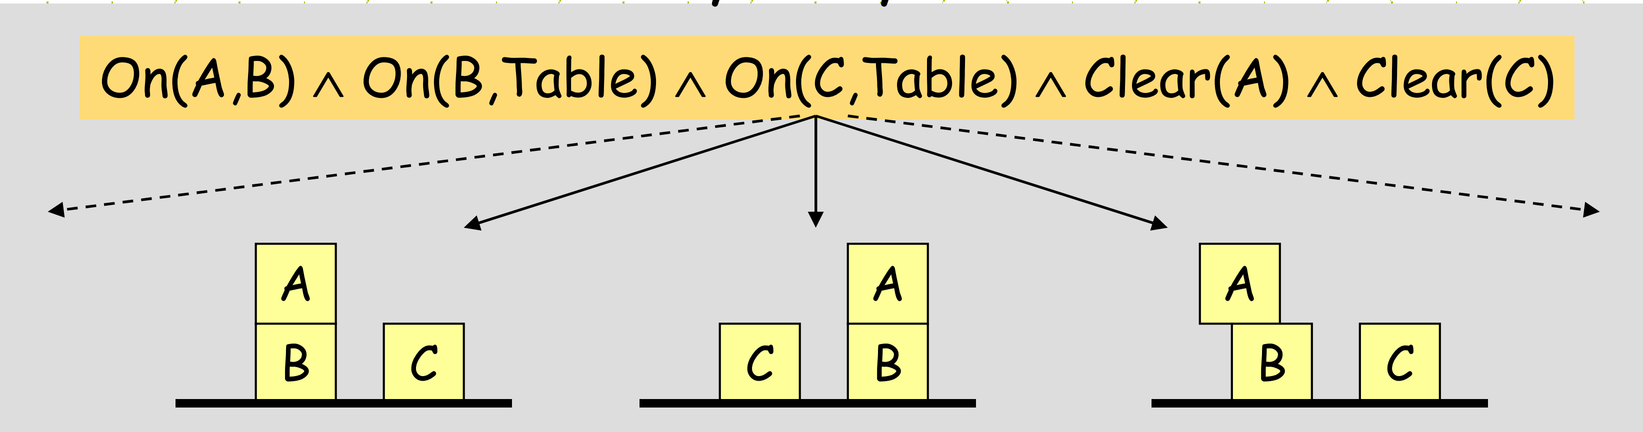

→ An action representation may be incorrect ...

Stack(C,A)

P = Holding(C) ∧ Block(C) ∧ Block(A) ∧ Clear(A)

D = Clear(A), Holding(C)

A = On(C,A), Clear(C), Handempty

is likely not to have the described effects in case 3 because the precondition is "incomplete"

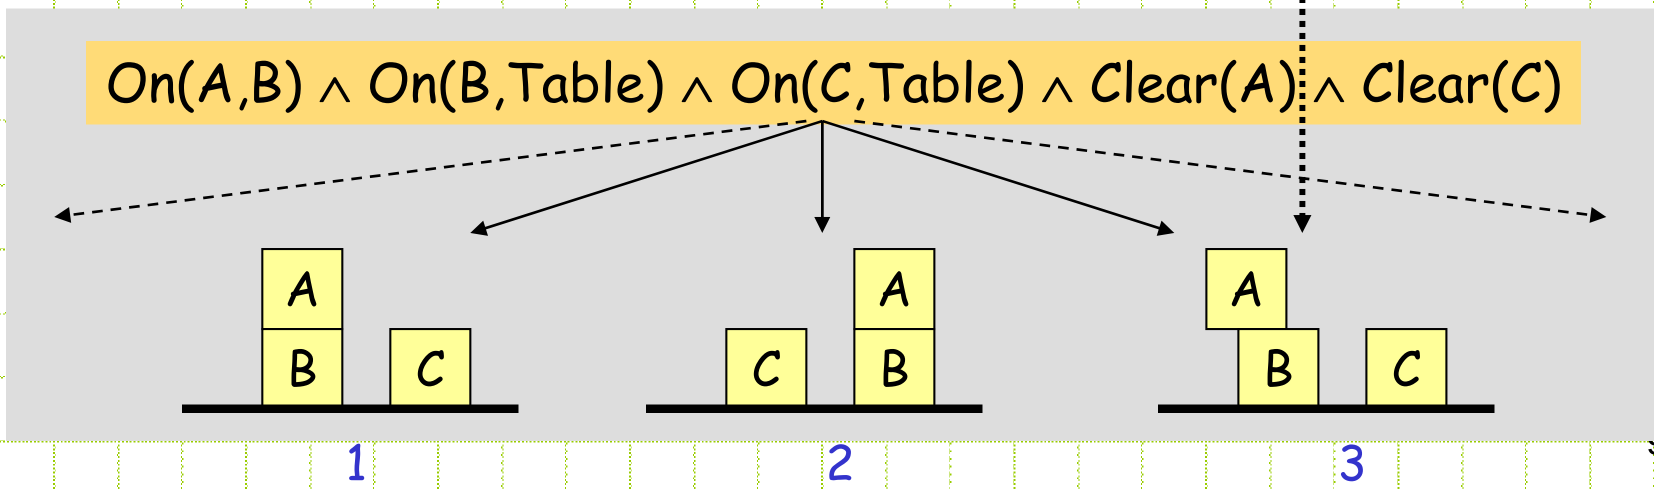

… or may describe several alternative effects

Stack(C,A)

P = Holding(C) ∧ Block(C) ∧ Block(A) ∧ Clear(A)
[If On(A,x) ∧ (x ≠ Table)]

E_1 { D = Clear(A), Holding(C)
A = On(C,A), Clear(C), Handempty

OR

E_2 { D = Holding(C), On(A,x)
A = On(C,Table), Clear(C), Handempty, On(A,Table), Clear(A), Clear(x)

On(A,B) ∧ On(B,Table) ∧ On(C,Table) ∧ Clear(A) ∧ Clear(C)

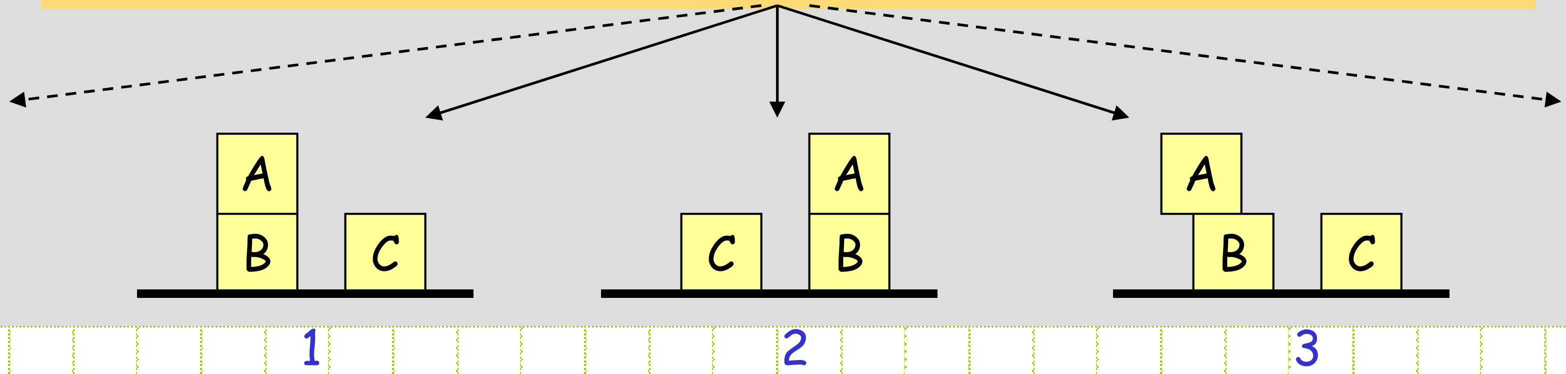

Observation of the Real World

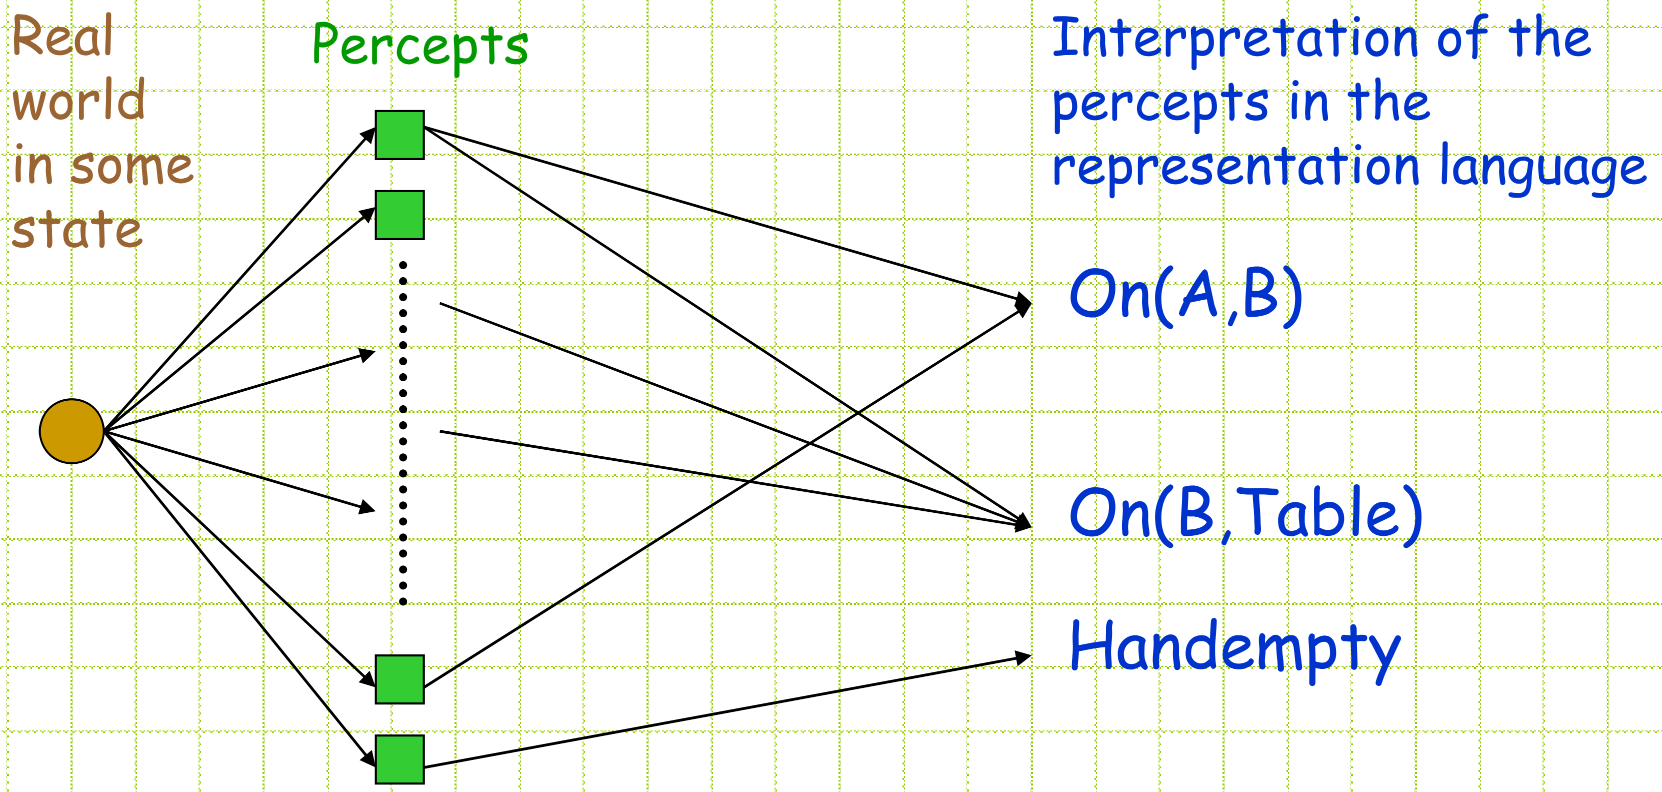

Percepts can be user's inputs, sensory data (e.g., image pixels), information received from other agents, ...

Second source of Uncertainty: Imperfect Observation of the World

Observation of the world can be:

- Partial, e.g., a vision sensor can't see through obstacles (lack of percepts)

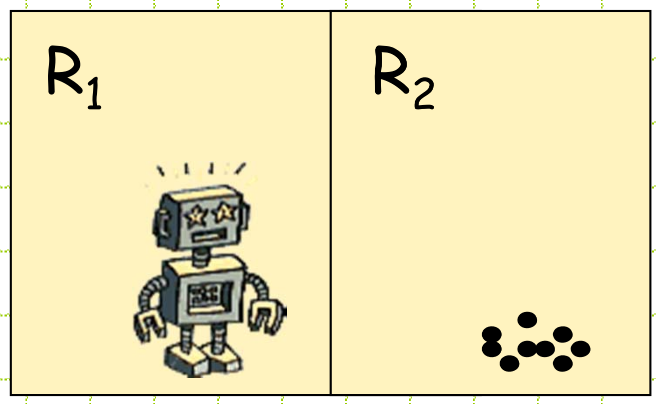

The robot may not know whether there is dust in room R2

Second source of Uncertainty: Imperfect Observation of the World

Observation of the world can be:

- Partial, e.g., a vision sensor can't see through obstacles
- Ambiguous, e.g., percepts have multiple possible interpretations

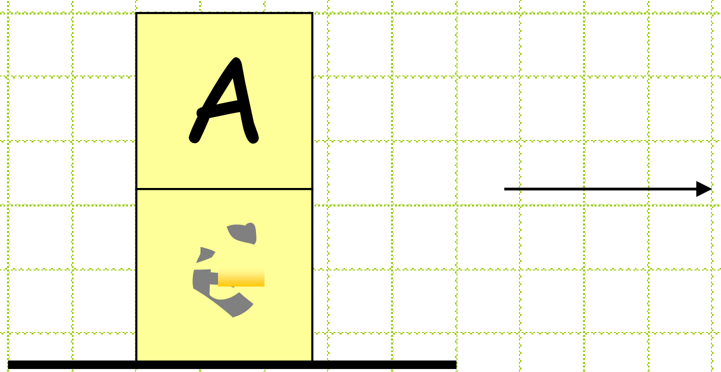

$On(A,B) \vee On(A,C)$

Second source of Uncertainty: Imperfect Observation of the World

Observation of the world can be:

- Partial, e.g., a vision sensor can't see through obstacles
- Ambiguous, e.g., percepts have multiple possible interpretations
- Incorrect

Third Source of Uncertainty: Ignorance, Laziness, Efficiency

- An action may have a long list of preconditions, e.g.:
  Drive-Car:
  P = Have(Keys) $\wedge$ $\neg$Empty(Gas-Tank) $\wedge$ Battery-Ok $\wedge$
  Ignition-Ok $\wedge$ $\neg$Flat-Tires $\wedge$ $\neg$Stolen(Car) ...
- The agent's designer may **ignore** some preconditions ... or by **laziness** or for **efficiency**, may not want to include all of them in the action representation
- The result is a representation that is either incorrect – executing the action may not have the described effects – or that describes several alternative effects

Representation of Uncertainty

- Many models of uncertainty
- We will consider two important models:
  - Non-deterministic model:
    Uncertainty is represented by a set of possible values, e.g., a set of possible worlds, a set of possible effects, ...
  - Probabilistic model:
    Uncertainty is represented by a probabilistic distribution over a set of possible values

Example: Belief State

- In the presence of non-deterministic sensory uncertainty, an agent belief state represents all the states of the world that it thinks are possible at a given time or at a given stage of reasoning

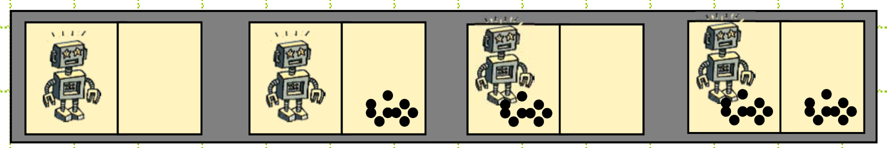

- In the probabilistic model of uncertainty, a probability is associated with each state to measure its likelihood to be the actual state

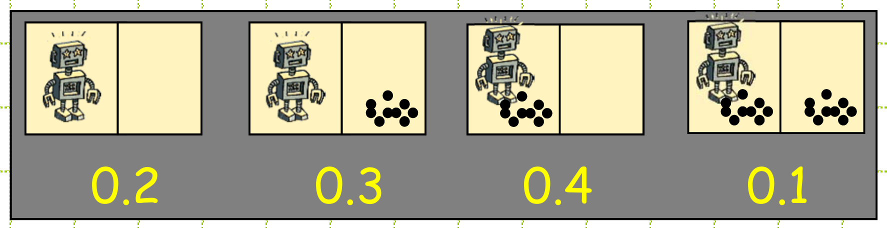

What do probabilities mean?

- Probabilities have a natural frequency interpretation
- The agent believes that if it was able to return many times to a situation where it has the same belief state, then the actual states in this situation would occur at a relative frequency defined by the probabilistic distribution

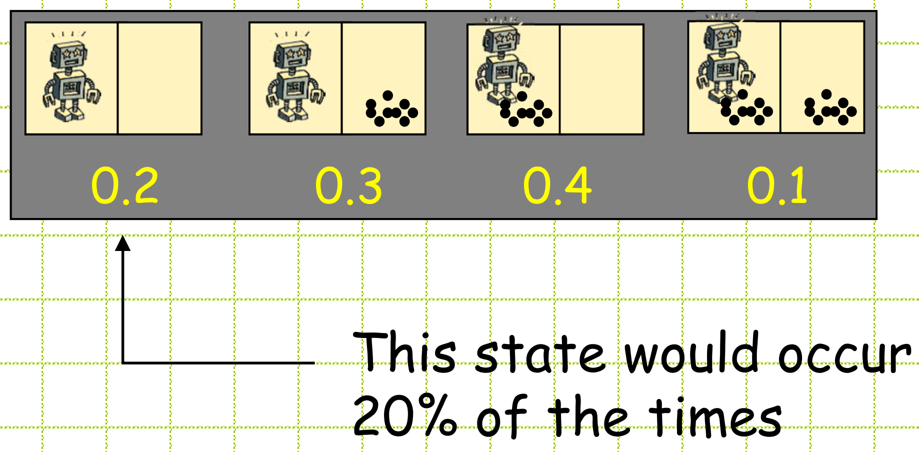

Example

- Consider a world where a dentist agent D meets a new patient P
- D is interested in only one thing: whether P has a cavity, which D models using the proposition Cavity
- Before making any observation, D's belief state is:

Cavity	$\neg$ Cavity
p	1-p

- This means that D believes that a fraction p of patients have cavities

Where do probabilities come from?

- Frequencies observed in the past, e.g., by the agent, its designer, or others
- Symmetries, e.g.:
  - If I roll a dice, each of the 6 outcomes has probability 1/6
- Subjectivism, e.g.:
  - If I drive on Highway 280 at 120mph, I will get a speeding ticket with probability 0.6
  - Principle of indifference: If there is no knowledge to consider one possibility more probable than another, give them the same probability