

# LECTURE 3: DEDUCTIVE REASONING AGENTS

An Introduction to Multiagent Systems

CIS 716.5, Spring 2010

## Agent Architectures

- Pattie Maes (1991):

‘[A] particular methodology for building [agents]. It specifies how . . . the agent can be decomposed into the construction of a set of component modules and how these modules should be made to interact. The total set of modules and their interactions has to provide an answer to the question of how the sensor data and the current internal state of the agent determine the actions . . . and future internal state of the agent. An architecture encompasses techniques and algorithms that support this methodology.’

- Leslie Kaelbling (1991):

‘[A] specific collection of software (or hardware) modules, typically designated by boxes with arrows indicating the data and control flow among the modules. A more abstract view of an architecture is as a general methodology for designing particular modular decompositions for particular tasks.’

## Types of Agents

- 1956–present: *Symbolic Reasoning Agents*

Agents make decisions about what to do via *symbol manipulation*.

Its purest expression, proposes that agents use *explicit logical reasoning* in order to decide what to do.

- 1985–present: *Reactive Agents*

Problems with symbolic reasoning led to a reaction against this — led to the *reactive agents* movement, 1985–present.

- 1990–present: *Hybrid Agents*

*Hybrid* architectures attempt to combine the best of reasoning and reactive architectures.

## Symbolic Reasoning Agents


- The classical approach to building agents is to view them as a particular type of knowledge-based system, and bring all the associated methodologies of such systems to bear.
- This paradigm is known as *symbolic AI*.
- We define a deliberative agent or agent architecture to be one that:
  - contains an explicitly represented, symbolic model of the world;
  - makes decisions (for example about what actions to perform) via symbolic reasoning.

## Two Issues

- *The transduction problem:*  
that of translating the real world into an accurate, adequate symbolic description, in time for that description to be useful.  
... vision, speech understanding, learning.
- *The representation/reasoning problem:*  
that of how to symbolically represent information about complex real-world entities and processes, and how to get agents to reason with this information in time for the results to be useful.  
... knowledge representation, automated reasoning, automatic planning.
- Most researchers accept that neither problem is anywhere near solved.

## The transduction problem

- Identifying objects is hard.



## The representation/reasoning problem

- Underlying problem lies with the complexity of symbol manipulation algorithms.
- In general many (most) search-based symbol manipulation algorithms of interest are *highly intractable*.
- Hard to find *compact* representations.
- Because of these problems, some researchers have looked to alternative techniques for building agents; we look at these later.

## Deductive Reasoning Agents

- How can an agent decide what to do using theorem proving?
- Basic idea is to use logic to encode a theory stating the *best* action to perform in any given situation.
- Let:
  - $\rho$  be this theory (typically a set of rules);
  - $\Delta$  be a logical database that describes the current state of the world;
  - $Ac$  be the set of actions the agent can perform;
  - $\Delta \vdash_{\rho} \phi$  mean that  $\phi$  can be proved from  $\Delta$  using  $\rho$ .



- How does this fit into the abstract description we talked about last time?
- The perception function is as before:

$$see : E \rightarrow Per$$

of course, this is (much) easier said than done.

- The next state function revises the database  $\Delta$ :

$$next : \Delta \times Per \rightarrow \Delta$$

- And the action function, well a possible action function is on the next slide.


```
for each  $\alpha \in Ac$  do      /* try to find an action explicitly prescribed */
    if  $\Delta \vdash_{\rho} Do(\alpha)$  then
        return  $\alpha$ 
    end-if
end-for

for each  $\alpha \in Ac$  do      /* try to find an action not excluded */
    if  $\Delta \not\vdash_{\rho} \neg Do(\alpha)$  then
        return  $\alpha$ 
    end-if
end-for

return null                /* no action found */
```

## The Vacuum World

- Goal is for the robot to clear up all dirt.



- Use 3 *domain predicates* in this exercise:

$In(x, y)$      agent is at  $(x, y)$

$Dirt(x, y)$     there is dirt at  $(x, y)$

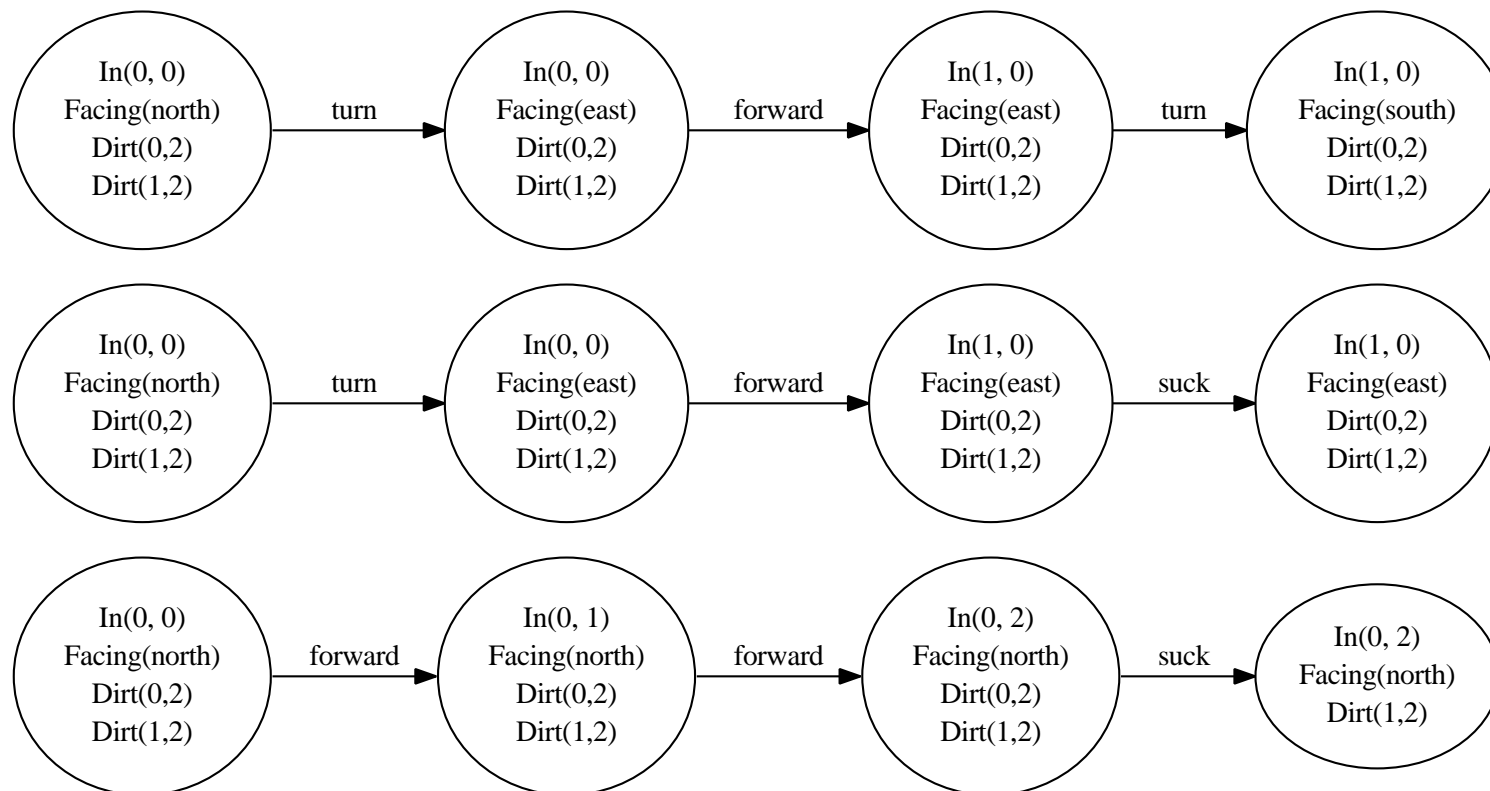
$Facing(d)$     the agent is facing direction  $d$

- Possible actions:

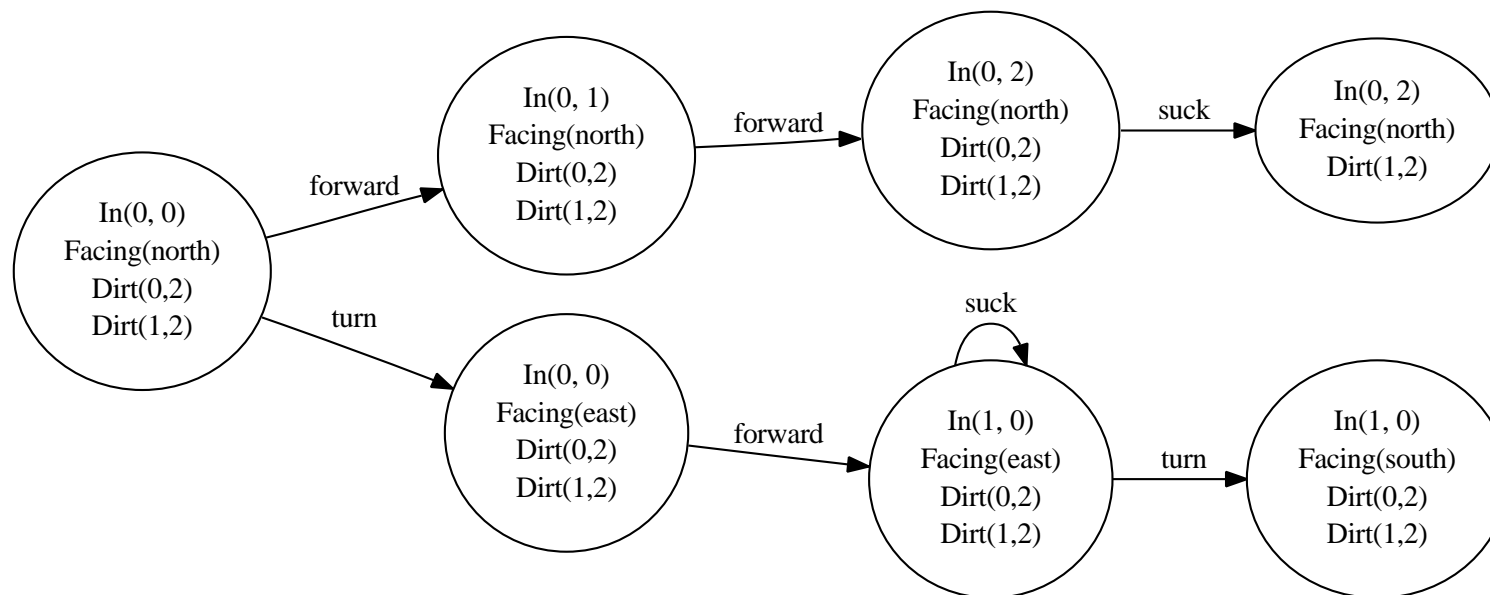
$$Ac = \{turn, forward, suck\}$$

NB: *turn* means “turn right”.

- With the system as depicted above, here are some possible ways that the system might run.



- If we only write down the distinct states, we get part of the statespace for the problem.



- Rules  $\rho$  for determining what to do:

$$In(0, 0) \wedge Facing(north) \wedge \neg Dirt(0, 0) \longrightarrow Do(forward)$$

$$In(0, 1) \wedge Facing(north) \wedge \neg Dirt(0, 1) \longrightarrow Do(forward)$$

$$In(0, 2) \wedge Facing(north) \wedge \neg Dirt(0, 2) \longrightarrow Do(turn)$$

$$In(0, 2) \wedge Facing(east) \longrightarrow Do(forward)$$

- ... and so on!
- Using these rules (+ other obvious ones), starting at  $(0, 0)$  the robot will clear up dirt.

- Problems:
  - how to convert video camera input to *Dirt*(0, 1)?
  - decision making assumes a *static* environment: *calculative* rationality.
  - decision making using first-order logic is *undecidable*!
- Even where we use *propositional* logic, decision making in the worst case means solving co-NP-complete problems.  
(NB: co-NP-complete = bad news!)
- Typical solutions:
  - weaken the logic;
  - use symbolic, non-logical representations;
  - shift the emphasis of reasoning from *run time* to *design time*.
- We now look at some examples of these approaches.



## Agent-oriented programming

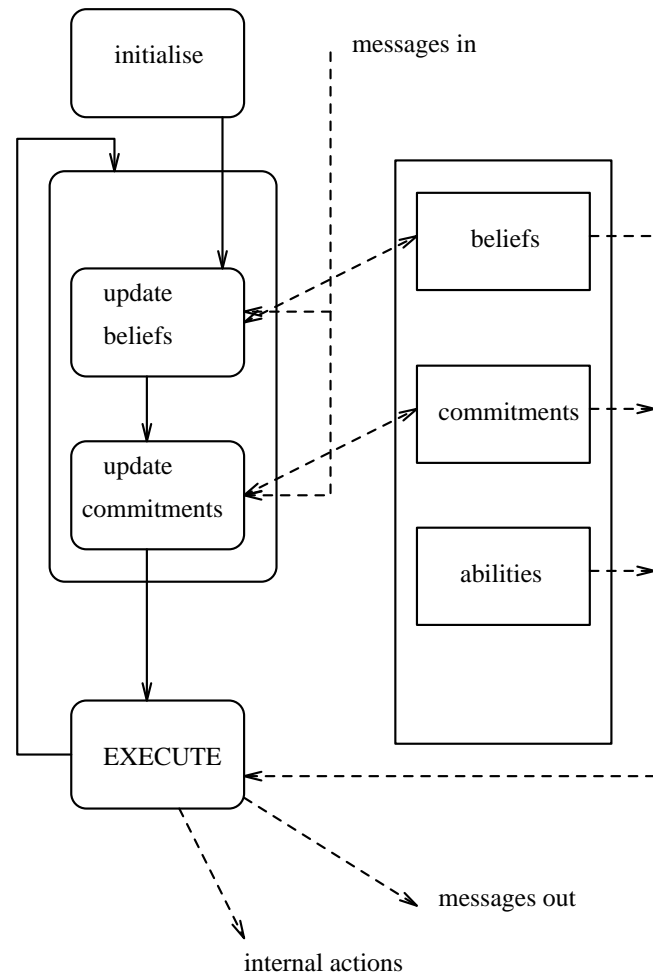
- Yoav Shoham introduced “agent-oriented programming” in 1990: “new programming paradigm, based on a societal view of computation”.
- The key idea:  
*directly programming agents in terms of intentional notions like belief, commitment, and intention.*

## Agent0

- AGENT0 is implemented as an extension to LISP.  
Each agent in AGENT0 has 4 components:
  - a set of capabilities (things the agent can do);
  - a set of initial beliefs;
  - a set of initial commitments (things the agent will do); and
  - a set of *commitment rules*.
- The key component, which determines how the agent acts, is the commitment rule set.

- Each commitment rule contains
  - a *message condition*;
  - a *mental condition*; and
  - an action.
- On each ‘decision cycle’ ...
  - The message condition is matched against the messages the agent has received;
  - The mental condition is matched against the beliefs of the agent.
  - If the rule fires, then the agent becomes committed to the action (the action gets added to the agents commitment set).

- Actions may be
  - *private*:  
an internally executed computation, or
  - *communicative*:  
sending messages.
- Messages are constrained to be one of three types:
  - “requests” to commit to action;
  - “unrequests” to refrain from actions;
  - “informs” which pass on information.



- A commitment rule:

```
COMMIT(  
  ( agent, REQUEST, DO(time, action)  
  ), ;;; msg condition  
  ( B,  
    [now, Friend agent] AND  
    CAN(self, action) AND  
    NOT [time, CMT(self, anyaction)]  
  ), ;;; mental condition  
  self,  
  DO(time, action)  
)
```

- This rule may be paraphrased as follows:  
if I receive a message from *agent* which requests me to do *action* at *time*, and I believe that:
  - *agent* is currently a friend;
  - I can do the action;
  - at *time*, I am not committed to doing any other action,then commit to doing *action* at *time*.

## PLACA

- A more refined implementation was developed by Thomas, for her 1993 doctoral thesis.
- Her Planning Communicating Agents (PLACA) language was intended to address one severe drawback to AGENT0: the inability of agents to plan, and communicate requests for action via high-level goals.
- Agents in PLACA are programmed in much the same way as in AGENT0, in terms of *mental change* rules.



- An example mental change rule:

```
((self ?agent REQUEST (?t (xeroxed ?x)))  
  (AND (CAN-ACHIEVE (?t xeroxed ?x))  
        (NOT (BEL (*now* shelving)))  
        (NOT (BEL (*now* (vip ?agent)))))  
  ((ADOPT (INTEND (5pm (xeroxed ?x))))))  
  ((?agent self INFORM  
    (*now* (INTEND (5pm (xeroxed ?x)))))))
```

- Paraphrased:

if someone asks you to xerox something, and you can, and you don't believe that they're a VIP, or that you're supposed to be shelving books, then

- Adopt the intention to xerox it by 5pm, and
- Inform them of your newly adopted intention.

## Concurrent METATEM

- Concurrent METATEM is a multi-agent language in which each agent is programmed by giving it a *temporal logic* specification of the behaviour it should exhibit.
- (Note, though, that the behavior it is capturing is reactive.)
- These specifications are executed directly in order to generate the behaviour of the agent.
- Temporal logic is classical logic augmented by *modal operators* for describing how the truth of propositions changes over time.
- Think of the world as being a number of discrete states.
- There is a single past history, but a number of possible futures — all the possible ways that the world might develop.

- For example. . .

- important(agents)

- means “it is now, and will always be true that agents are important”

- ◇ important(ConcurrentMetateM)

- means “sometime in the future, ConcurrentMetateM will be important”

- ◆ important(Prolog)

- means “sometime in the past it was true that Prolog was important”

- More examples...

$(\neg \text{friends}(\text{us})) \mathcal{U} \text{apologise}(\text{you})$

means “we are not friends until you apologise”

$\bigcirc \text{apologise}(\text{you})$

means “tomorrow (in the next state), you apologise”.

- And a last couple of examples . . .

$$\odot \text{apologise}(\text{you}) \Rightarrow \bigcirc \text{friends}(\text{us})$$

means “if you apologised yesterday, then tomorrow we will be friends”.

The operator  $\odot$  indicates the previous state in time.

$$\text{friends}(\text{us}) \mathcal{S} \text{apologise}(\text{you})$$

means “we have been friends since you apologised”

- The root of the MetateM concept is Gabbay's *separation theorem*:

Any arbitrary temporal logic formula can be rewritten in a logically equivalent *past*  $\Rightarrow$  *future* form.

- Just like our last example.
- MetateM program is a collection of

*past*  $\Rightarrow$  *future*

rules.

- Execution proceeds by a process of continually matching rules against a “history”, and *firing* those rules whose antecedents are satisfied.
- The instantiated future-time consequents become *commitments* which must subsequently be satisfied.

- An example MetateM program: the resource controller...

$$\bigcirc \text{ask}(x) \Rightarrow \diamond \text{give}(x)$$

$$\text{give}(x) \wedge \text{give}(y) \Rightarrow (x=y)$$

- First rule ensure that an ‘ask’ is eventually followed by a ‘give’.
- Second rule ensures that only one ‘give’ is ever performed at any one time.



- A Concurrent MetateM system contains a number of agents (objects), each object has 3 attributes:
  - a name;
  - an interface;
  - a MetateM program.

- An agent's interface contains two sets:
  - messages the agent will *accept*,
  - messages the agent may *send*.
- For example, a 'stack' object's interface:

stack(pop, push)[popped, stackfull]

{pop, push} = messages received

{popped, stackfull} = messages sent

## Snow White & The Dwarves

- To illustrate the language Concurrent MetateM in more detail, here are some example programs. . .
- These are taken from a paper by Michael Fisher.
- Snow White has some sweets (resources), which she will give to the Dwarves (resource consumers).
- She will only give to one dwarf at a time.
- She will always eventually give to a dwarf that asks.

- Here is Snow White, written in Concurrent MetateM:

Snow-White(ask)[give]:

$$\begin{aligned} & \bigcirc \text{ask}(x) \Rightarrow \Diamond \text{give}(x) \\ & \text{give}(x) \wedge \text{give}(y) \Rightarrow (x = y) \end{aligned}$$

- The dwarf 'eager' asks for a sweet initially, and then whenever he has just received one, asks again.

eager(give)[ask]:

start  $\Rightarrow$  ask(eager)

⦿ give(eager)  $\Rightarrow$  ask(eager)

- Some dwarves are even less polite: 'greedy' just asks every time.

greedy(give)[ask]:

start  $\Rightarrow$   $\square$  ask(greedy)

- Fortunately, some have better manners; 'courteous' only asks when 'eager' and 'greedy' have eaten.

courteous(give)[ask]:

$$((\neg \text{ask}(\text{courteous}) \mathcal{S} \text{give}(\text{eager})) \wedge \\ (\neg \text{ask}(\text{courteous}) \mathcal{S} \text{give}(\text{greedy}))) \Rightarrow \\ \text{ask}(\text{courteous})$$

- And finally, 'shy' will only ask for a sweet when no-one else has just asked.

shy(give)[ask]:

- $start \Rightarrow \Diamond \text{ask}(\text{shy})$
- $\bigcirc \text{ask}(x) \Rightarrow \neg \text{ask}(\text{shy})$
- $\bigcirc \text{give}(\text{shy}) \Rightarrow \Diamond \text{ask}(\text{shy})$



## Summary

- This lecture has looked at how one might build an agent in symbolic logic.
- We discussed several ways of doing this:
  - Using simple propositional logic
  - Using the logic-inspired approach of Agent0 and PLACA.
  - In the executable temporal logic of METATEM.
- Note that none of agents we specified are very complex, nor did we specify behaviors that were much more than simply reactive.
- Papers describing more complex specifications can be found on the webpage.