

LECTURE 4: PRACTICAL REASONING AGENTS

An Introduction to Multiagent Systems

CIS 716.5, Spring 2010

What is Practical Reasoning?

- Practical reasoning is reasoning directed towards actions — the process of figuring out what to do:

Practical reasoning is a matter of weighing conflicting considerations for and against competing options, where the relevant considerations are provided by what the agent desires/values/cares about and what the agent believes.
(Bratman)

- Distinguish practical reasoning from *theoretical reasoning*. Theoretical reasoning is directed towards beliefs.

The Components of Practical Reasoning

- Human practical reasoning consists of two activities:
 - *deliberation*
deciding *what* state of affairs we want to achieve
— the outputs of deliberation are *intentions*;
 - *means-ends reasoning*
deciding *how* to achieve these states of affairs
— the outputs of means-ends reasoning are *plans*.

Intentions in Practical Reasoning

1. Intentions pose problems for agents, who need to determine ways of achieving them.

If I have an intention to ϕ , you would expect me to devote resources to deciding how to bring about ϕ .

2. Intentions provide a “filter” for adopting other intentions, which must not conflict.

If I have an intention to ϕ , you would not expect me to adopt an intention ψ that was incompatible with ϕ .

3. Agents track the success of their intentions, and are inclined to try again if their attempts fail.

If an agent's first attempt to achieve ϕ fails, then all other things being equal, it will try an alternative plan to achieve ϕ .

4. Agents believe their intentions are possible.

That is, they believe there is at least some way that the intentions could be brought about.

5. Agents do not believe they will not bring about their intentions.

It would not be rational of me to adopt an intention to ϕ if I believed I would fail with ϕ .

6. Under certain circumstances, agents believe they will bring about their intentions.

If I intend ϕ , then I believe that under “normal circumstances” I will succeed with ϕ .

7. Agents need not intend all the expected side effects of their intentions.

If I believe $\phi \Rightarrow \psi$ and I intend that ϕ , I do not necessarily intend ψ also. (Intentions are not closed under implication.)

This last problem is known as the *side effect* or *package deal* problem.


I may believe that going to the dentist involves pain, and I may also intend to go to the dentist — but this does not imply that I intend to suffer pain!

Intentions are Stronger than Desires

My desire to play basketball this afternoon is merely a potential influencer of my conduct this afternoon. It must vie with my other relevant desires [. . .] before it is settled what I will do. In contrast, once I intend to play basketball this afternoon, the matter is settled: I normally need not continue to weigh the pros and cons. When the afternoon arrives, I will normally just proceed to execute my intentions. (Bratman, 1990)


Means-ends Reasoning/Planning

- Planning is the design of a course of action that will achieve some desired goal.
- Basic idea is to give a planning system:
 - (representation of) goal/intention to achieve;
 - (representation of) actions it can perform; and
 - (representation of) the environment;and have it generate a *plan* to achieve the goal.
- This is *automatic programming*.



- How do we do this?
- Well, we have already seen one way when we saw how we might develop a simple reactive plan.
- Next we'll look at an approach that is more declarative.
- STRIPS, the Stanford Research Institute Problem Solver.

- Used in Shakey the robot:



Representations

- Question: How do we *represent*. . .
 - goal to be achieved;
 - state of environment;
 - actions available to agent;
 - plan itself.
- Answer: We use logic, or something that looks a lot like logic.


- We'll illustrate the techniques with reference to the *blocks world*.
- A simple (toy) world, in this case one where we consider toys:



- The blocks world contains a robot arm, 3 blocks (A, B and C) of equal size, and a table-top.



- The aim is to generate a plan for the robot arm to build towers out of blocks.
- For a formal description, we'll clean it up a bit:



- To represent this environment, need an *ontology*.

On(x, y) obj x on top of obj y
OnTable(x) obj x is on the table
Clear(x) nothing is on top of obj x
 Holding(x) arm is holding x

- Here is a representation of the blocks world described above:

Clear(A)

On(A, B)

OnTable(B)

OnTable(C)

- Use the *closed world assumption*
 - Anything not stated is assumed to be *false*.

- A *goal* is represented as a set of formulae.
- Here is a goal:

$$\{OnTable(A), OnTable(B), OnTable(C)\}$$

- *Actions* are represented as follows.

Each action has:

- a *name*
which may have arguments;
- a *pre-condition list*
list of facts which must be true for action to be executed;
- a *delete list*
list of facts that are no longer true after action is performed;
- an *add list*
list of facts made true by executing the action.

Each of these may contain *variables*.

- The *stack* action occurs when the robot arm places the object x it is holding is placed on top of object y .

$Stack(x, y)$
pre $Clear(y) \wedge Holding(x)$
del $Clear(y) \wedge Holding(x)$
add $ArmEmpty \wedge On(x, y)$

- The *unstack* action occurs when the robot arm picks an object x up from on top of another object y .

UnStack(x, y)
pre $On(x, y) \wedge Clear(x) \wedge ArmEmpty$
del $On(x, y) \wedge ArmEmpty$
add $Holding(x) \wedge Clear(y)$

Stack and UnStack are *inverses* of one-another.


- The *pickup* action occurs when the arm picks up an object x from the table.

Pickup(x)
pre $Clear(x) \wedge OnTable(x) \wedge ArmEmpty$
del $OnTable(x) \wedge ArmEmpty$
add $Holding(x)$

- The *putdown* action occurs when the arm places the object x onto the table.

PutDown(x)
pre *Holding*(x)
del *Holding*(x)
add *Holding*(x) \wedge *ArmEmpty*

- What is a plan?
A sequence (list) of actions, with variables replaced by constants.
- So, to get from:



- We need the set of actions:

Unstack(A)

Putdown(A)

Pickup(B)


Stack(B, C)

Pickup(A)

Stack(A, B)

- Creating the list of actions is not so hard in this case, but gets harder as the number of possible runs gets bigger.

- As here for example:



Formal representation

- Let's relate the STRIPS model back to the formal description of an agent we talked about before.
- This will help us to see how it fits into the overall picture.
- As before we assume that the agent has a set of actions A_C , and we will write individual actions as α_1, α_2 and so on.
- Now the actions have some structure, each one has preconditions, add list and delete list, for each $\alpha_i \in A_C$:

$$\alpha_i = \langle P_{\alpha_i}, D_{\alpha_i}, A_{\alpha_i} \rangle$$

- A plan is just a sequence of actions:

$$\pi = (\alpha_1, \dots, \alpha_n)$$

where each action is one of the actions from A_C

- A *planning problem* is then:

$$\langle B, Ac, i \rangle$$

where:

- B_0 is the set of beliefs the agent has about the world.
- Ac is the set of actions, and
- i is a goal (or intention)

- Since actions change the world, any rational agent will change its beliefs about the world as a result of carrying out actions.
- Thus, a plan π for a given planning problem will be associated with a sequence of sets of beliefs:

$$B_0 \xrightarrow{\alpha_1} B_1 \xrightarrow{\alpha_2} \dots \xrightarrow{\alpha_n} B_n$$

where for $1 \leq i \leq n$:

$$B_i = (B_{i-1} \setminus D_{\alpha_i}) \cup A_{\alpha_i}$$

- In other words at each step of the plan the beliefs are updated by removing the items in the delete list of the relevant action and adding the items in the add list.

- A plan π is said to be *acceptable* with respect to the problem $\langle B, Ac, i \rangle$ if and only if:

$$B_{i-1} \models P_{\alpha_i}$$

for all $1 \leq i \leq n$

- In other words, the pre-requisites for each action have to be true right before the action is carried out.
- (We say $B_{i-1} \models P_{\alpha_i}$ because the pre-conditions don't have to be in B_{i-1} , we just have to be able to prove the pre-conditions from B_{i-1} .)
- A plan π is *correct* if:
 1. it is acceptable; and
 2. $B_n \models i$
- In other words, it is correct if it is acceptable and the final state makes the goal true.

Implementing Practical Reasoning Agents

- A first pass at an implementation of a practical reasoning agent:

Agent Control Loop Version 1

1. while true
2. observe the world;
3. update internal world model;
4. deliberate about what intention
 to achieve next;
5. use means-ends reasoning to get
 a plan for the intention;
6. execute the plan
7. end while

- We will not be concerned with stages (2) or (3) except to say that these are related to the functions *see* and *next* from Lecture #2.
- *see* is as before:

$$see : E \rightarrow Per$$

but instead of *next*:

$$next : I \times Per \rightarrow I$$

we have:

$$brf : 2^{Bel} \times Per \rightarrow 2^{Bel}$$

- Note that $B \subseteq Bel$.

- Problem: deliberation and means-ends reasoning processes are not instantaneous.
They have a *time cost*.
- Suppose that deliberation is *optimal* in that if it selects some intention to achieve, then this is the best thing for the agent.
- (It maximises expected utility.)
- So the agent has selects an intention to achieve that would have been optimal *at the time it observed the world*.
This is *calculative rationality*.
- The world *may change* in the meantime.
- Even if the agent can compute the right thing to do, it may not do the right thing.
- Optimality is hard.

- Let's make the algorithm more formal.

Agent Control Loop Version 2

```
1.  $B := B_0$ ; /* initial beliefs */
2. while true do
3.     get next percept  $\rho$ ;
4.      $B := brf(B, \rho)$ ;
5.      $I := deliberate(B)$ ;
6.      $\pi := plan(B, I)$ ;
7.      $execute(\pi)$ 
8. end while
```

where $I \subseteq Int$, $plan$ is exactly what we discussed above, and $execute$ is a function that executes each action in a plan.

- How might we implement these functions?

Deliberation

- How does an agent deliberate?
 - begin by trying to understand what the *options* available to you are;
 - *choose between them*, and *commit* to some.

Chosen options are then intentions.

- The *deliberate* function can be decomposed into two distinct functional components:
 - option generation; and
 - filtering.

- In *option generation*, the agent generates a set of possible alternatives.
- Represent option generation via a function, *options*:

$$options : 2^{Bel} \times 2^{Int} \rightarrow 2^{Des}$$

- This takes the agent's current beliefs and current intentions, and from them determines a set of options (= *desires*).
- In *filtering*, the agent chooses between competing alternatives, and commits to achieving them.
- In order to select between competing options, an agent uses a *filter* function.

$$filter : 2^{Bel} \times 2^{Des} \times 2^{Int} \rightarrow 2^{Int}$$

Agent Control Loop Version 3

1. $B := B_0$;
2. $I := I_0$;
3. while true do
4. get next percept ρ ;
5. $B := brf(B, \rho)$;
6. $D := options(B, I)$;
7. $I := filter(B, D, I)$;
8. $\pi := plan(B, I)$;
9. execute(π)
10. end while

- where $D \subseteq Des$.

Commitment Strategies

Some time in the not-so-distant future, you are having trouble with your new household robot. You say “Willie, bring me a beer.” The robot replies “OK boss.” Twenty minutes later, you screech “Willie, why didn’t you bring me that beer?” It answers “Well, I intended to get you the beer, but I decided to do something else.” Miffed, you send the wise guy back to the manufacturer, complaining about a lack of commitment. After retrofitting, Willie is returned, marked “Model C: The Committed Assistant.” Again, you ask Willie to bring you a beer. Again, it accedes, replying “Sure thing.” Then you ask: “What kind of beer did you buy?” It answers: “Genessee.” You say “Never mind.” One minute later, Willie trundles over with a Genessee in its gripper. [. . .] After still more tinkering, the manufacturer sends Willie back, promising no more problems with its commitments. So, being a somewhat trusting customer, you accept the rascal back into your household, but as a test, you ask it to bring you your last beer. [. . .] The robot gets the beer and starts towards you. As it approaches, it lifts its arm, wheels around, deliberately smashes the bottle, and trundles off. Back at the plant, when interrogated by customer service as to why it had abandoned its commitments, the robot replies that according to its specifications, it kept its commitments as long as required — commitments must be dropped when fulfilled or impossible to achieve. By smashing the bottle, the commitment became unachievable.

Degrees of Commitment

- *Blind commitment*

A blindly committed agent will continue to maintain an intention until it believes the intention has actually been achieved. Blind commitment is also sometimes referred to as *fanatical* commitment.

- *Single-minded commitment*

A single-minded agent will continue to maintain an intention until it believes that either the intention has been achieved, or else that it is no longer possible to achieve the intention.

- An agent has commitment both to *ends* (i.e., the state of affairs it wishes to bring about), and *means* (i.e., the mechanism via which the agent wishes to achieve the state of affairs).
- Currently, our agent control loop is overcommitted, both to means and ends.

Modification: *replan* if ever a plan goes wrong.

- To write the algorithm down we need to refine our notion of plan execution.
- If π is a plan, then:
 - $empty(\pi)$ is true if there are no more actions in the plan.
 - $hd(\alpha)$ returns the first action in the plan.
 - $tail(\alpha)$ returns the plan minus the head of the plan.
 - $sound(\pi, I, B)$ means that π is a correct plan for I given B .
- Now we can say the following:

- The next version of the control loop:

Agent Control Loop Version 4

1. $B := B_0 ;$
2. $I := I_0 ;$
3. while true do
4. get next percept $\rho ;$
5. $B := brf(B, \rho) ;$
6. $D := options(B, I) ;$
7. $I := filter(B, D, I) ;$
8. $\pi := plan(B, I) ;$

⋮

Agent Control Loop Version 4 (cont)

:

```
9.      while not empty( $\pi$ ) do
10.          $\alpha := hd(\pi)$ ;
11.         execute( $\alpha$ );
12.          $\pi := tail(\pi)$ ;
13.         get next percept  $\rho$ ;
14.          $B := brf(B, \rho)$ ;
15.         if not sound( $\pi, I, B$ ) then
16.              $\pi := plan(B, I)$ 
17.         end-if
18.     end-while
19. end-while
```

- Still overcommitted to intentions: Never stops to consider whether or not its intentions are appropriate.
- Modification: stop to determine whether intentions have succeeded or whether they are impossible:

single-minded commitment.

- Next version:

Agent Control Loop Version 5

1. $B := B_0 ;$
2. $I := I_0 ;$
3. while true do
4. get next percept $\rho ;$
5. $B := brf(B, \rho) ;$
6. $D := options(B, I) ;$
7. $I := filter(B, D, I) ;$
8. $\pi := plan(B, I) ;$
- \vdots

```
⋮
9.   while not empty( $\pi$ )
      or succeeded( $I, B$ )
      or impossible( $I, B$ ) do
10.   $\alpha := hd(\pi) ;$ 
11.  execute( $\alpha$ ) ;
12.   $\pi := tail(\pi) ;$ 
13.  get next percept  $\rho ;$ 
14.   $B := brf(B, \rho) ;$ 
15.  if not sound( $\pi, I, B$ ) then
16.     $\pi := plan(B, I)$ 
17.  end-if
18.  end-while
19. end-while
```

Intention Reconsideration

- Our agent gets to reconsider its intentions once every time around the outer control loop, i.e., when:
 - it has completely executed a plan to achieve its current intentions; or
 - it believes it has achieved its current intentions; or
 - it believes its current intentions are no longer possible.
- This is limited in the way that it permits an agent to *reconsider* its intentions.
- Modification: Reconsider intentions after executing every action.

Agent Control Loop Version 6

⋮

9. while not (*empty*(π) or *succeeded*(I, B)
 or *impossible*(I, B)) do

10. $\alpha := hd(\pi);$

11. *execute*(α);

12. $\pi := tail(\pi);$

13. get next percept ρ ;

14. $B := brf(B, \rho);$

15. $D := options(B, I);$

16. $I := filter(B, D, I);$

17. if not *sound*(π, I, B) then

18. $\pi := plan(B, I)$

⋮

- But intention reconsideration is *costly*!

A dilemma:

- an agent that does not stop to reconsider its intentions sufficiently often will continue attempting to achieve its intentions even after it is clear that they cannot be achieved, or that there is no longer any reason for achieving them;
 - an agent that *constantly* reconsiders its intentions may spend insufficient time actually working to achieve them, and hence runs the risk of never actually achieving them.
- Solution: incorporate an explicit *meta-level control* component, that decides whether or not to reconsider.

Agent Control Loop Version 7

⋮

```
10.     while not (empty( $\pi$ ) or succeeded( $I, B$ )
           or impossible( $I, B$ )) do
11.          $\alpha := hd(\pi)$ ;
12.         execute( $\alpha$ );
13.          $\pi := tail(\pi)$ ;
14.         get next percept  $\rho$ ;
15.          $B := brf(B, \rho)$ ;
16.         if reconsider( $I, B$ ) then
17.              $D := options(B, I)$ ;
18.              $I := filter(B, D, I)$ ;
19.         if not sound( $\pi, I, B$ ) then  $\pi := plan(B, I)$ 
⋮
```

- The possible interactions between meta-level control and deliberation are:

Situation number	Chose to deliberate?	Changed intentions?	Would have changed intentions?	<i>reconsider(...)</i> optimal?
1	No	—	No	Yes
2	No	—	Yes	No
3	Yes	No	—	No
4	Yes	Yes	—	Yes


- An important assumption: cost of *reconsider(...)* is *much* less than the cost of the deliberation process itself.

- In situation (1), the agent did not choose to deliberate, and as a consequence, did not choose to change intentions. Moreover, if it *had* chosen to deliberate, it would not have changed intentions. In this situation, the *reconsider(...)* function is behaving optimally.
- In situation (2), the agent did not choose to deliberate, but if it had done so, it *would* have changed intentions. In this situation, the *reconsider(...)* function is not behaving optimally.
- In situation (3), the agent chose to deliberate, but did not change intentions. In this situation, the *reconsider(...)* function is not behaving optimally.
- In situation (4), the agent chose to deliberate, and did change intentions. In this situation, the *reconsider(...)* function is behaving optimally.

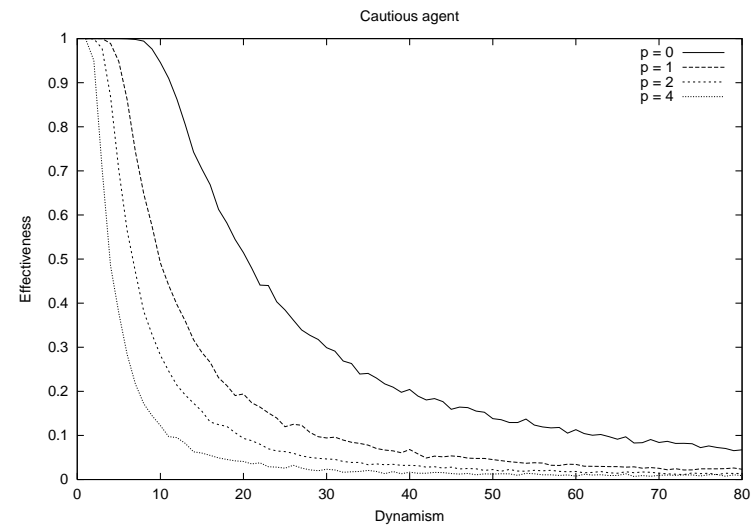
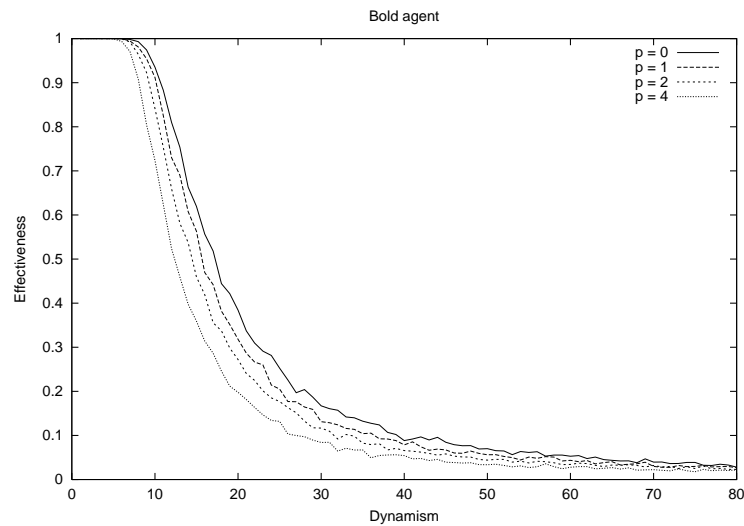
Optimal Intention Reconsideration

- Kinny and Georgeff's experimentally investigated effectiveness of intention reconsideration strategies.
- Two different types of reconsideration strategy were used:
 - *bold* agents
never pause to reconsider intentions, and
 - *cautious* agents
stop to reconsider after every action.
- *Dynamism* in the environment is represented by the *rate of world change*, γ .
- Experiments in the Tileworld.

- Tileworld revisited:




- Experiments replicated by Martijn Schut:



- So:
 - If γ is low (i.e., the environment does not change quickly), then bold agents do well compared to cautious ones. This is because cautious ones waste time reconsidering their commitments while bold agents are busy working towards — and achieving — their intentions.
 - If γ is high (i.e., the environment changes frequently), then cautious agents can outperform bold agents. This is because they are able to recognize when intentions are doomed, and also to take advantage of serendipitous situations and new opportunities when they arise.
 - When planning costs are high, this advantage can be eroded.

PRS

- We now make the discussion even more concrete by introducing an actual agent architecture: the PRS.
- In the PRS, each agent is equipped with a *plan library*, representing that agent's *procedural knowledge*: knowledge about the mechanisms that can be used by the agent in order to realise its intentions.
- The options available to an agent are directly determined by the plans an agent has: an agent with no plans has no options.
- In addition, PRS agents have explicit representations of beliefs, desires, and intentions, as above.



Example PRS (JAM) System

GOALS:

```
ACHIEVE blocks_stacked;
```

FACTS:

```
// Block1 on Block2 initially so need to clear Block2 before stacking.
```

```
FACT ON "Block1" "Block2";  
FACT ON "Block2" "Table";  
FACT ON "Block3" "Table";  
FACT CLEAR "Block1";  
FACT CLEAR "Block3";  
FACT CLEAR "Table";  
FACT initialized "False";
```

```
Plan: {
NAME: "Top-level plan"
DOCUMENTATION:
    "Establish Block1 on Block2 on Block3."
GOAL:
    ACHIEVE blocks_stacked;
CONTEXT:
BODY:
    EXECUTE print "Goal is Block1 on Block2 on Block2 on Table.\n";
    EXECUTE print "World Model at start is:\n";
    EXECUTE printWorldModel;
    EXECUTE print "ACHIEVEing Block3 on Table.\n";
    ACHIEVE ON "Block3" "Table";
    EXECUTE print "ACHIEVEing Block2 on Block3.\n";
    ACHIEVE ON "Block2" "Block3";
    EXECUTE print "ACHIEVEing Block1 on Block2.\n";
    ACHIEVE ON "Block1" "Block2";
    EXECUTE print "World Model at end is:\n";
    EXECUTE printWorldModel;
}
```

```
Plan: {
NAME: "Stack blocks that are already clear"
GOAL:
    ACHIEVE ON $OBJ1 $OBJ2;
CONTEXT:
BODY:
    EXECUTE print "Making sure " $OBJ1 " is clear\n";
    ACHIEVE CLEAR $OBJ1;
    EXECUTE print "Making sure " $OBJ2 " is clear.\n";
    ACHIEVE CLEAR $OBJ2;
    EXECUTE print "Moving " $OBJ1 " on top of " $OBJ2 " .\n";
    PERFORM move $OBJ1 $OBJ2;
UTILITY: 10;

FAILURE:
    EXECUTE print "\n\nStack blocks failed!\n\n";
}
```

```
Plan: {
NAME: "Clear a block"
GOAL:
    ACHIEVE CLEAR $OBJ;
CONTEXT:
    FACT ON $OBJ2 $OBJ;
BODY:
    EXECUTE print "Clearing " $OBJ2 " from on top of " $OBJ "\n";
    EXECUTE print "Moving " $OBJ2 " to table.\n";
    ACHIEVE ON $OBJ2 "Table";

EFFECTS:
    EXECUTE print "CLEAR: Retracting ON " $OBJ2 " " $OBJ "\n";
    RETRACT ON $OBJ1 $OBJ;

FAILURE:
    EXECUTE print "\n\nClearing block " $OBJ " failed!\n\n";
}
```

```
Plan: {
NAME: "Move a block onto another object"
GOAL:
    PERFORM move $OBJ1 $OBJ2;
CONTEXT:
    FACT CLEAR $OBJ1;
    FACT CLEAR $OBJ2;
BODY:
    EXECUTE print "Performing low-level move action"
    EXECUTE print " of " $OBJ1 " to " $OBJ2 ".\n";

EFFECTS:
    WHEN : TEST (!= $OBJ2 "Table") {
EXECUTE print "    Retracting CLEAR " $OBJ2 "\n";
RETRACT CLEAR $OBJ2;
    };
    FACT ON $OBJ1 $OBJ3;
EXECUTE print "    move: Retracting ON " $OBJ1 " " $OBJ3 "\n";
RETRACT ON $OBJ1 $OBJ3;
EXECUTE print "    move: Asserting CLEAR " $OBJ3 "\n";
ASSERT CLEAR $OBJ3;
EXECUTE print "    move: Asserting ON " $OBJ1 " " $OBJ2 "\n\n";
ASSERT ON $OBJ1 $OBJ2;

FAILURE:
    EXECUTE print "\n\nMove failed!\n\n";
$}
```

Summary

- This lecture has covered a lot of ground on practical reasoning.
- We started by discussing what practical reasoning was, and how it relates to intentions.
- We then looked at planning (how an agent achieves its desires) and how deliberation and means-ends reasoning fit into the basic agent control loop.
- We then refined the agent control loop, considering commitment and intention reconsideration.
- Finally, we looked at some implemented systems.