Part Nine:
Practical Cognition / Planning

- **Given an agent capable of sophisticated epistemic cognition, how can it make use of that in practical cognition?**
- **We can regard practical cognition as having four components:**
  - goal selection
  - plan-construction
  - plan-selection
  - plan-execution
- **Although it is natural to think of these as components of practical cognition, most of the work will be carried out by epistemic cognition.**

1. Plan Construction

Means-End Reasoning

- To achieve a goal, we consider an action that would achieve it under some specified circumstances.
- Then we try to find a way of putting ourselves in those circumstances in order to achieve the goal by performing the action.
- Putting ourselves in those circumstances becomes a *subgoal*.
- The idea is to work backwards from the goal through subgoals until we arrive at subgoals that are already achieved.
- The resulting sequence of actions constitutes a *plan* for achieving the goal.

2. Goal-Regression Planning

- In this section I will formulate a somewhat more precise description of means-end reasoning.
- In section three I will argue that, assuming the basic correctness of the conventional theory, rational agents situated in a complex environment cannot in general perform means-end reasoning in quite the way AI planning theory proposes to implement it. In a sense to be explained, planning must be done defeasibly rather than by a planner that generates plans in a more mechanical (recursively enumerable) way.
- In section four I will argue that the conventional theory runs afoul of the Frame Problem and must be modified to accommodate a solution to the Frame Problem.
- In section five I will show how the conventional theory of means-end reasoning must be modified to accomplish this.

Planning-Conditionals

Planning proceeds in terms of causal knowledge of the form "performing action A under circumstances C is causally sufficient for achieving goal G". This is encoded by a *planning-conditional*:

(A / C) ➤ G

C is the *precondition* of the conditional
A is the *action*
G is the *goal*

Plans

- Plans are constructed out of *plan-steps*, which prescribe actions.
- Plan-steps cannot be identified with the actions they prescribe, because the same action may be prescribed by more than one step in a single plan.
- The plan-steps must be executed in a proper order.
- A *causal-link* is a triple ⟨$step_1$, *subgoal*, $step_2$⟩ recording the fact that $step_1$ is performed in order to achieve *subgoal*, which is a precondition for $step_2$.

- I will identify a plan with the ordered quadruple consisting of the set of its plan-steps, the ordering of the plan-steps, the causal-links describing its causal structure, and the goal of the plan.

Goal-Regression Planning

PROPOSE-NULL-PLAN

Given an interest in finding a plan for achieving *goal*, if *goal* is already true, propose a null-plan for *goal*.

GOAL-REGRESSION

Given an interest in finding a plan for achieving G, and given a planning-conditional (A / C) ➤ G, adopt an interest in finding a plan for achieving C. If a plan *subplan* is proposed for achieving C, construct a plan by (1) adding a new step to the end of *subplan* where the new step prescribes the action A, and (2) ordering the new step after all steps of *subplan*. Propose the new plan as a plan for achieving G.

Splitting Conjunctive Goals

- The subgoals generated by GOAL-REGRESSION will usually be conjunctions.
  - For example, if my goal is to light a fire, I may observe that I could do so by lighting a match provided I have a match and I have tinder.
  - GOAL-REGRESSION will thus generate the conjunctive subgoal *I have a match and I have tinder.*

Splitting Conjunctive Goals

- The subgoals generated by GOAL-REGRESSION will usually be conjunctions.
- We will generally be unable to make further progress in our plan construction by applying GOAL-REGRESSION once more to such a conjunctive subgoal (C_1 & C_2).
- To do so would require our having a planning-conditional of the form (A / C) ➤ (C_1 & C_2).
- It is rare that we will have a single planning-conditional like this that will achieve both conjuncts of a conjunctive subgoal.
- The most we can generally hope for is to have two separate planning-conditionals (A / C) ➤ C_1 and (A* & C*) ➤ C_2, which will allow us to construct separate subplans for the individual conjuncts.
- Given subplans for achieving each conjunct, we can then attempt to construct a plan for achieving the conjunction by merging the plans for the conjuncts.

Splitting Conjunctive Goals

Given two plans $plan_1$ and $plan_2$, let $plan_1 + plan_2$ be the plan that results from combining the plan-steps and ordering-constraints of each.

SPLIT-CONJUNCTIVE-GOAL

Given an interest in finding a plan for achieving a conjunctive goal (G_1 & G_2), adopt interest in finding plans $plan_1$ for G_1 and $plan_2$ for G_2. If such plans are proposed, propose $plan_1 + plan_2$ as a plan for (G_1 & G_2).

SPLIT-CONJUNCTIVE-GOAL-SAFELY

Given an interest in finding a plan for achieving a conjunctive goal (G_1 & G_2), adopt interest in finding plans $plan_1$ for G_1 and $plan_2$ for G_2. If such plans are proposed and do not destructively interfere with each other, propose $plan_1 + plan_2$ as a plan for (G_1 & G_2).

Destructive Interference

Two plans $plan_1$ and $plan_2$ *destructively interfere* with each other iff $plan_1 + plan_2$ contains a causal-link ⟨$step_1$, *subgoal*, $step_2$⟩ and a subplan that establishes ~*subgoal*, where there is a consistent ordering of the plan-steps in which ~*subgoal* is established *between* $step_1$ and $step_2$.

3. R.E. Planning and Defeasible Planning

- An *r.e. planner* executes an effective computation, i.e., the set of pairs ⟨problem,solution⟩ that characterize the planner is recursively enumerable.
- An r.e. planner will only be possible if the set of destructive interferences is effectively computable.
- In order for destructive interference to be computable, it must be computable whether a particular condition (the negation of a precondition of one of the plan steps) is a consequence of an action under specifiable circumstances.

Autonomous Planning Agents

- Our objective here is the construction of autonomous rational agents capable of maneuvering through a complex, variable, and often uncooperative environment.
- It cannot be assumed that a planning agent has exactly the knowledge it needs to solve a planning problem.
- The principal function of epistemic cognition in an autonomous agent is to provide the information needed for practical cognition.
- As such, the course of epistemic cognition is driven by practical interests.
- Rather than coming to the planning problem equipped with all the knowledge required for its solution, the planning problem itself directs epistemic cognition, focusing epistemic endeavors on the pursuit of information that will be helpful in solving current planning problems.

Autonomous Planning Agents — contd.

- Paramount among this information is knowledge about what will happen if certain actions are taken under certain circumstances.
- Sometimes the agent already knows what will happen, but often it has to figure it out.
- At the very least this will require reasoning from current knowledge.
- In many cases it will require the empirical acquisition of new knowledge that cannot be obtained just by reasoning from what is already known.
- An essential characteristic of planning agents is that planning and epistemic cognition are interleaved.
- Unlike applied planning, it is impossible to require of a planning agent capable of functioning in realistically complex environments that it acquire all the requisite knowledge before beginning the plan search.

Autonomous Planning Agents — contd.

- Autonomous planning agents cannot rely on precompiled knowledge.
- They must engage in genuine reasoning about the consequences of actions, and we should not expect that reasoning to be any simpler than general epistemic reasoning.
- Realistically, epistemic reasoning must be defeasible, which makes the set of conclusions at best Δ_2.
- Even if we could construct an agent that did only first-order deductive reasoning, the set of conclusions is not effectively computable—it is recursively enumerable.
- This means that when the planning algorithm computes plans for the conjuncts of a conjunctive goal and then considers whether they can be merged without destructive interference, the reasoning required to find any particular destructive interference may take indefinitely long.
- If there is no destructive interference, there will be no point at which the planner can draw the conclusion that there is none simply on the grounds that none has been found.
- Thus the planning algorithm will bog down at this point and will never be able to produce the merged plan for the conjunctive goal.

Autonomous Planning Agents — contd.

- **Theorem: If the set of threats is not recursive, then the set of planning ⟨problem,solution⟩ pairs is not recursively enumerable.**
- **Corollary: A planner that insists upon ruling out threats before merging plans for the conjuncts of a conjunctive goal may never terminate.**

"The logical foundations of goal-regression planning in autonomous agents", *Artificial Intelligence*, 1998.

Autonomous Planning Agents — contd.

- If destructive interference is not computable, how can a planner get away with dividing conjunctive goals into separate conjuncts and planning for each conjunct separately?
- Humans assume defeasibly that the separate plans do not destructively interfere with one another, and so infer defeasibly that the merged plan is a good plan for the conjunctive goal.
- A rational agent operating in a realistically complex environment must make defeasible assumptions in the course of its planning, and then be prepared to change its planning decisions later if subsequent epistemic reasoning defeats those defeasible assumptions.
- In other words, the reasoning involved in planning must be a species of defeasible reasoning.
- Planning in autonomous agents cannot be done by an r.e. planner.

Autonomous Planning Agents — contd.

- The general way means-end reasoning must work is by performing goal regression, splitting conjunctive goals into their conjuncts and planning for them separately, and then merging the plans for the individual conjuncts into a combined plan for the conjunctive goal.
- The practical reasoner will infer defeasibly that the merged plan is a solution to the planning problem.
- A defeater for this defeasible inference consists of discovering that the plan contains destructive interference.
- Whenever a defeasible reasoner makes a defeasible inference, it must adopt interest in finding defeaters, so in this case the agent will adopt interest in finding destructive interference.
- Finding such interference should lead the agent to try various ways of repairing the plan to eliminate the interference, and then lead to a defeasible inference that the repaired plan is a solution to the planning problem.
- The tentative conclusion being adopted is that the plan will achieve its goal.
- Means-end reasoning becomes a form of epistemic reasoning to the effect that if a plan is executed (in any way consistent with the ordering) then it is defeasibly reasonable to expect the goal to be achieved.

Achieving Goals

A *linearization* of a partial-order plan is a linear plan that results from adding additional ordering constraints sufficient to linearly order the plan-steps.

A partial-order plan will achieve its goal iff every linearization of it will achieve its goal.

Let us take an action-sequence to be a linear sequence of actions.

Achieving Goals

A *linearization* of a partial-order plan is a linear plan that results from adding additional ordering constraints sufficient to linearly order the plan-steps.

Conventional goal-regression planning is based upon the following definitions.

Definition: A partial-order plan will achieve its goal iff every linearization of it will achieve its goal.

Let us take an action-sequence to be a linear sequence of actions.

(R1) Where *start-state* is a state of affairs and *conditionals* is a set of planning-conditionals, P is a result of $\langle A_1,...,A_n\rangle$ relative to *start-state* and *conditionals* iff either:
- (i) n = 0 and P is true in *start-state*; or
- (ii) n > 0 and *conditionals* contains a conditional (A_n / C) ➤ P such that C is a result of $\langle A_1,...,A_{n-1}\rangle$; or
- (iii) n > 0, P is a result of $\langle A_1,...,A_{n-1}\rangle$, and *conditionals* does not contain a conditional of the form (A_n / C) ➤ ~Q such that Q is either P or a conjunct of P and C is a result of $\langle A_1,...,A_{n-1}\rangle$; or
- (iv) n > 0 and P is a conjunction whose conjuncts are results of $\langle A_1,...,A_n\rangle$.

Soundness Assumption

A linear plan will achieve its goal relative to a state *start-state* iff its goal is a result of the sequence of actions prescribed by its plan-steps relative to *start-state* and the set of all true planning-conditionals.

- The Soundness Assumption provides the mathematical basis for a complete theory of means-end reasoning.
- The end result is a proof that when means-end reasoning is performed in accordance with these rules, the plans it produces will achieve their goals,

  and if there is a plan that will achieve a particular goal, some such plan will be found by following these rules of means-end reasoning.

Principles of Defeasible Planning

GOAL-REGRESSION

Given an interest in finding a plan for achieving G-at-t, adopt interest in finding a planning-conditional (A/C) ⇒ G. Given such a conditional, adopt interest in finding a plan for achieving C-at-t*. If it is concluded that a plan *subplan* will achieve C-at-t*, construct a plan by (1) adding a new step to the end of *subplan* where the new step prescribes the action A-at-t*, (2) adding a constraint (t* < t) to the ordering-constraints of *subplan*, and (3) adjusting the causal-links appropriately. Infer defeasibly that the new plan will achieve G-at-t.

SPLIT-CONJUNCTIVE-GOAL

Given an interest in finding a plan for achieving (G_1-at-t_1 & G_2-at-t_2), adopt interest in finding plans *$plan_1$* for G_1-at-t_1 and *$plan_2$* for G_2-at-t_2. Given such plans, infer defeasibly that the result of merging them will achieve (G_1-at-t_1 & G_2-at-t_2).

— a number of additional reason-schemas are also required —

"The logical foundations of goal-regression planning in autonomous agents", *Artificial Intelligence*, 1998.

Example — Pednault's Briefcase

- **(at-home briefcase)**
- **(at-home paycheck)**
- **(in-briefcase paycheck)**
- **(∀x)[((in-briefcase x) & (remove-from-briefcase x)) ⇒ ~(in-briefcase x)]**
- **[(at-home briefcase) & take-briefcase-to-office] ⇒ (at-office briefcase)**
- **(∀x)[((at-home briefcase) & (in-briefcase x) & take-briefcase-to-office) ⇒ ~(at-home x)]**

Goal: (at-home paycheck) & (at-office briefcase)

Example — Pednault's Briefcase

Example — Pednault's Briefcase

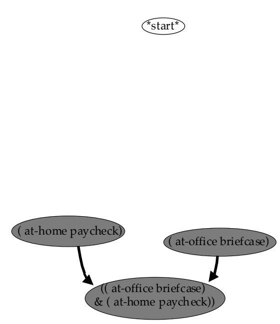

Example — Pednault's Briefcase

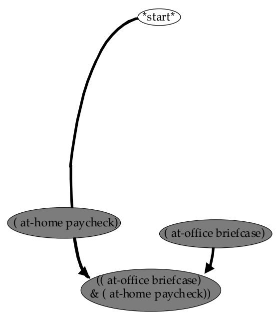

Example — Pednault's Briefcase

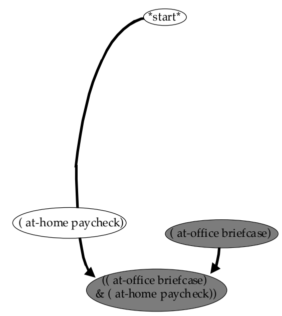

Example — Pednault's Briefcase

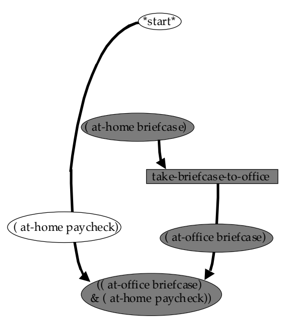

Example — Pednault's Briefcase

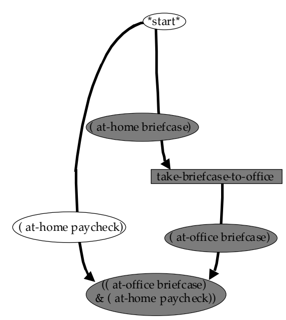

Example — Pednault's Briefcase

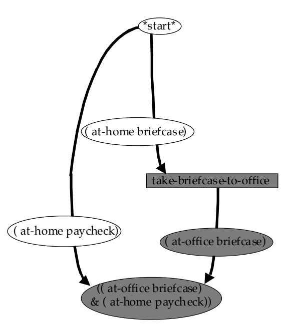

Example — Pednault's Briefcase

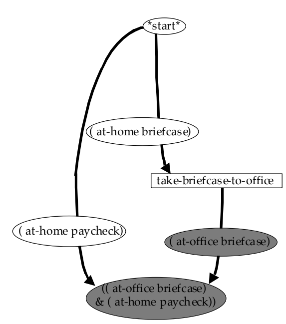

Example — Pednault's Briefcase

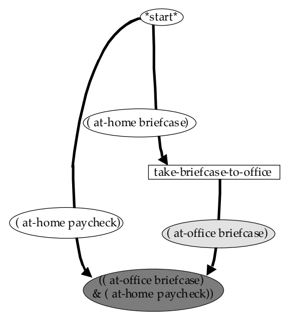

Example — Pednault's Briefcase

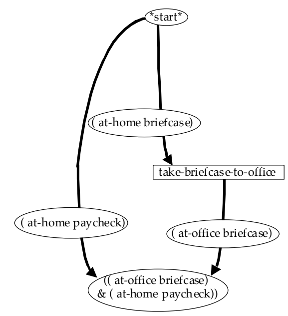

Example — Pednault's Briefcase

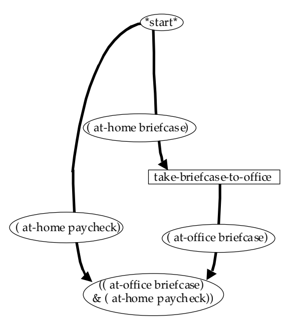

Example — Pednault's Briefcase

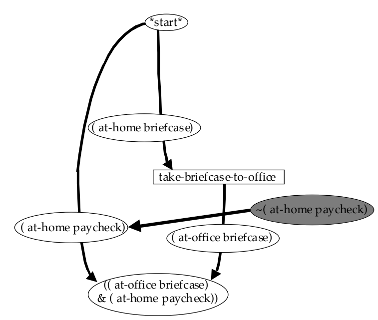

Example — Pednault's Briefcase

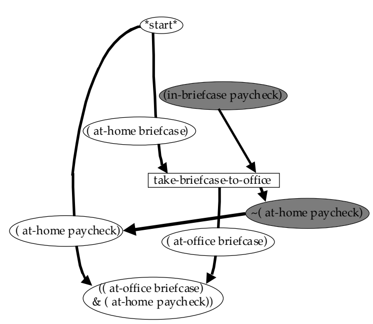

Example — Pednault's Briefcase

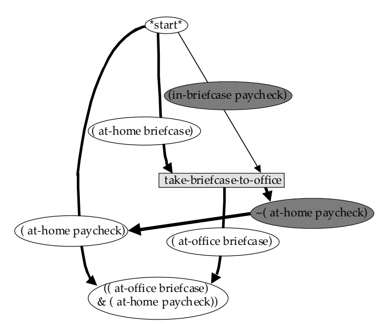

Example — Pednault's Briefcase

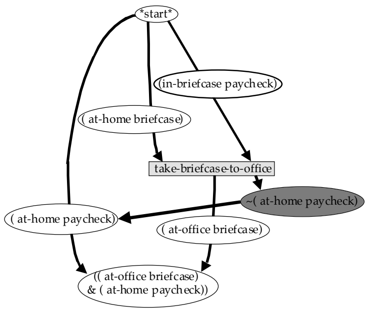

Example — Pednault's Briefcase

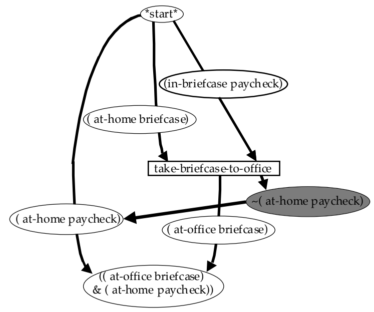

Example — Pednault's Briefcase

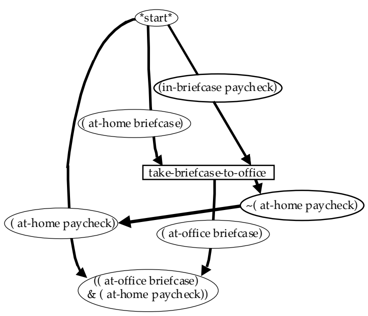

Example — Pednault's Briefcase

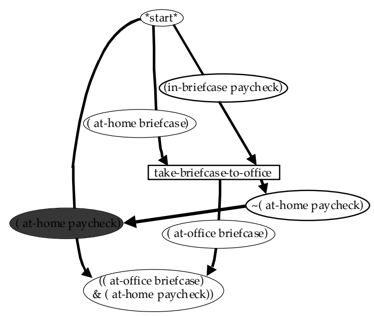

Example — Pednault's Briefcase

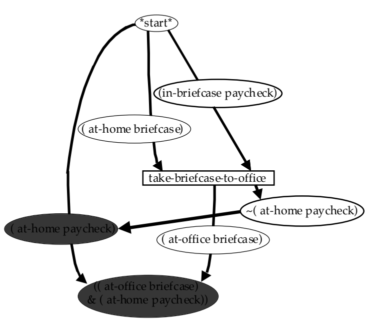

Example — Pednault's Briefcase

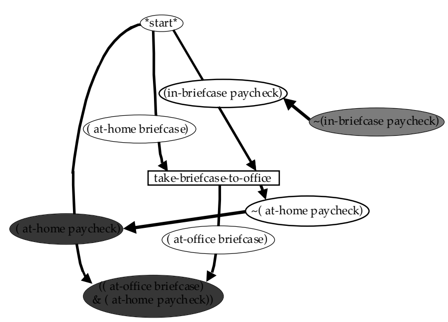

Example — Pednault's Briefcase

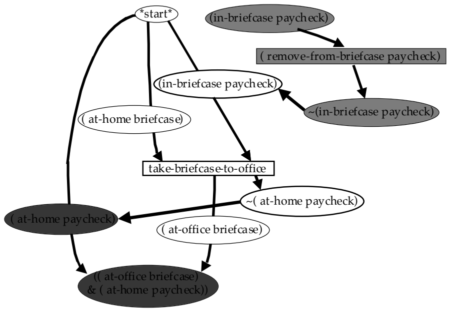

Example — Pednault's Briefcase

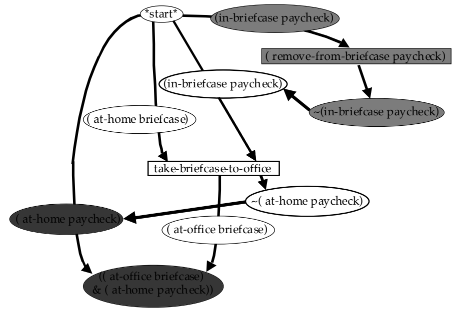

Example — Pednault's Briefcase

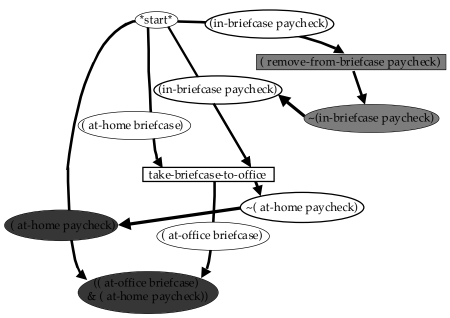

Example — Pednault's Briefcase

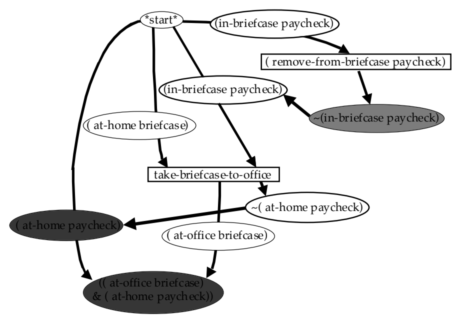

Example — Pednault's Briefcase

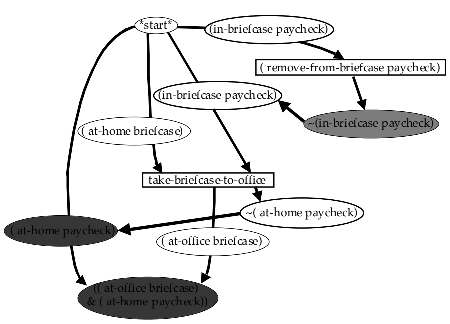

Example — Pednault's Briefcase

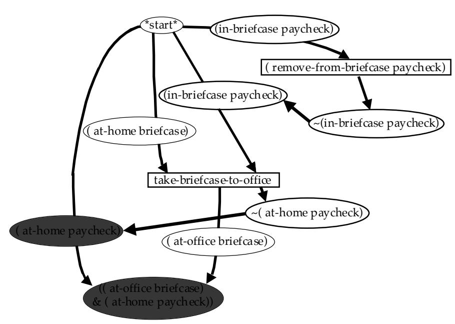

Example — Pednault's Briefcase

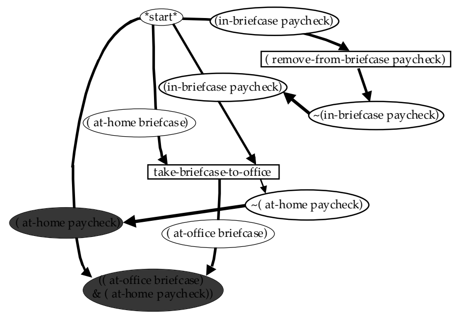

Example — Pednault's Briefcase

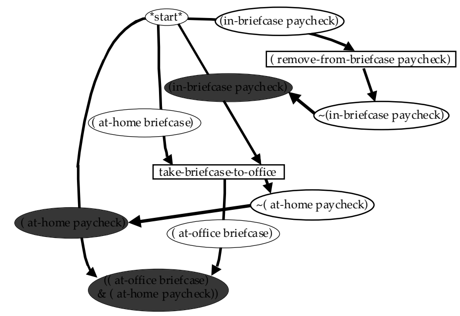

Example — Pednault's Briefcase

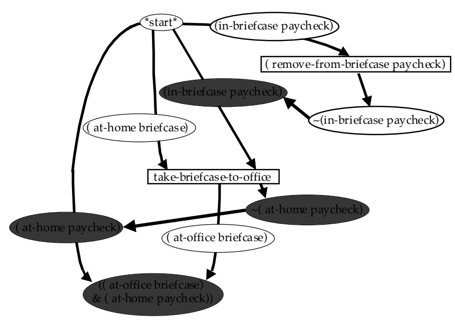

Example — Pednault's Briefcase

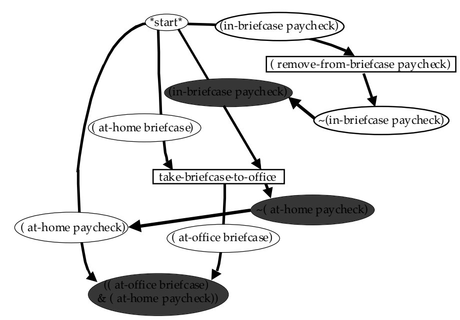

Example — Pednault's Briefcase

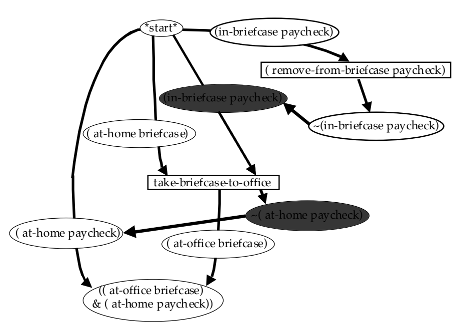

Example — Pednault's Briefcase

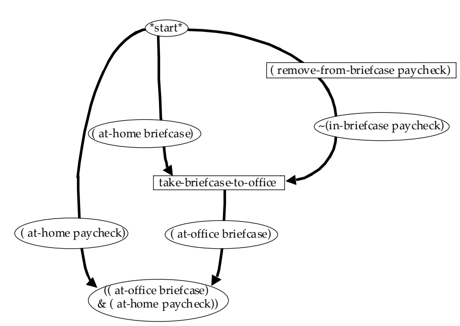

Defeasible Planning

- A defeasible planner can perform essentially the same plan search as a conventional goal-regression planner.
- The difference is that if the set of potential threats is not recursive, the conventional planner will not return a plan, but the defeasible planner can return a solution tentatively.
- When the agent must take action, it can then decide what to do by using the plans it has adopted defeasibly.
- This is the way any realistic agent must operate. It cannot wait for all reasoning to terminate.

4. Planning and the Frame Problem

The conventional theory of goal-regression planning depends heavily on the following:

Soundness Assumption

A linear plan will achieve its goal relative to a state *start-state* iff its goal is a result of the sequence of actions prescribed by its plan-steps relative to *start-state* and the set of all true planning-conditionals.

(R1) Where *start-state* is a state of affairs and *conditionals* is a set of planning-conditionals, P is a result of $\langle A_1,...,A_n\rangle$ relative to *start-state* and *conditionals* iff either:
- (i) n = 0 and P is true in *start-state*; or
- (ii) n > 0 and *conditionals* contains a conditional (A_n & C) ➤ P such that C is a result of $\langle A_1,...,A_{n-1}\rangle$; or
- (iii) n > 0, P is a result of $\langle A_1,...,A_{n-1}\rangle$, and *conditionals* does not contain a conditional of the form (A_n & C) ➤ ~Q such that Q is either P or a conjunct of P and C is a result of $\langle A_1,...,A_{n-1}\rangle$; or
- (iv) n > 0 and P is a conjunction whose conjuncts are results of $\langle A_1,...,A_n\rangle$.

Soundness Assumption

A linear plan will achieve its goal relative to a state *start-state* iff its goal is a result of the sequence of actions prescribed by its plan-steps relative to *start-state* and the set of all true planning-conditionals.

What the soundness assumption seems to be saying is:

Necessarily, a linear plan will achieve a goal G when executed from a start-state iff G is a result of the sequence of actions prescribed by its plan-steps relative to the start-state and the set of all true planning-conditionals.

Interpreting (R1) in that light, what is the justification for clause (iii)?

(R1) Where *start-state* is a state of affairs and *conditionals* is a set of planning-conditionals, P is a result of $\langle A_1,...,A_n\rangle$ relative to *start-state* and *conditionals* iff either:
- (i) n = 0 and P is true in *start-state*; or
- (ii) n > 0 and *conditionals* contains a conditional (A_n & C) ➤ P such that C is a result of $\langle A_1,...,A_{n-1}\rangle$; or
- (iii) n > 0, P is a result of $\langle A_1,...,A_{n-1}\rangle$, and *conditionals* does not contain a conditional of the form (A_n & C) ➤ ~Q such that Q is either P or a conjunct of P and C is a result of $\langle A_1,... A_{n-1}\rangle$; or
- (iv) n > 0 and P is a conjunction whose conjuncts are results of $\langle A_1,...,A_n\rangle$.

Clause (iii) is often said to presuppose:

The STRIPS assumption

Nothing changes in the world unless it does so as a result of executing a step of the plan.

This is silly!

We have a *limited* expectation, not that nothing will change, but that the particular subgoals established by initial steps of the plan will not change unless executing later steps of the plan causes them to change.

Providing the logical foundations for such a defeasible expectation is just the Frame Problem. Such foundations were proposed earlier.

So construed, clause (iii) is based upon temporal projection. As such,we have only a defeasible expectation that subgoals will remain true.

(R1) should be about "expectable results".

(R1) Where *start-state* is a state of affairs and *conditionals* is a set of planning-conditionals, P is an expectable result of $\langle A_1,...,A_n\rangle$ relative to *start-state* and *conditionals* iff either:
- (i) n = 0 and P is true in *start-state*; or
- (ii) n > 0 and *conditionals* contains a conditional (A_n & C) ➤ P such that C is an expectable result of $\langle A_1,...,A_{n-1}\rangle$; or
- (iii) n > 0, P is an expectable result of $\langle A_1,...,A_{n-1}\rangle$, and *conditionals* does not contain a conditional of the form (A_n & C) ➤ ~Q such that Q is either P or a conjunct of P and C is an expectable result of $\langle A_1,...,A_{n-1}\rangle$; or
- (iv) n > 0 and P is a conjunction whose conjuncts are expectable results of $\langle A_1,...,A_n\rangle$.

Clause (i) tells us to expect that if P is true initially it will remain true. This is an instance of TEMPORAL-PROJECTION and accordingly requires addition of a temporal-projectibility constraint:

(i) n = 0 and P is true in *start-state* and temporally-projectible.

PROPOSE-NULL-PLAN requires a similar constraint.

(R1) Where *start-state* is a state of affairs and *conditionals* is a set of planning-conditionals, P is an expectable result of $\langle A_1,...,A_n\rangle$ relative to *start-state* and *conditionals* iff either:

(i) n = 0 and P is true in *start-state*; or

(ii) n > 0 and *conditionals* contains a conditional (A_n & C) ➤ P such that C is an expectable result of $\langle A_1,...,A_{n-1}\rangle$; or

(iii) n > 0, P is an expectable result of $\langle A_1,...,A_{n-1}\rangle$, and *conditionals* does not contain a conditional of the form (A_n & C) ➤ ~Q such that Q is either P or a conjunct of P and C is an expectable result of $\langle A_1,...,A_{n-1}\rangle$; or

(iv) n > 0 and P is a conjunction whose conjuncts are expectable results of $\langle A_1,...,A_n\rangle$.

Clause (iv) requires no modification.
SPLIT-CONJUNCTIVE-GOAL requires no modification

(R1) Where *start-state* is a state of affairs and *conditionals* is a set of planning-conditionals, P is an expectable result of $\langle A_1,...,A_n\rangle$ relative to *start-state* and *conditionals* iff either:

(i) n = 0 and P is true in *start-state*; or

(ii) n > 0 and *conditionals* contains a conditional (A_n & C) ➤ P such that C is an expectable result of $\langle A_1,...,A_{n-1}\rangle$; or

(iii) n > 0, P is an expectable result of $\langle A_1,...,A_{n-1}\rangle$, and *conditionals* does not contain a conditional of the form (A_n & C) ➤ ~Q such that Q is either P or a conjunct of P and C is an expectable result of $\langle A_1,...,A_{n-1}\rangle$; or

(iv) n > 0 and P is a conjunction whose conjuncts are expectable results of $\langle A_1,...,A_n\rangle$.

Clause (ii) describes causal inferences of the sort to which the Frame Problem is relevant. If we can expect *C-at-t_n* to be achieved by executing the sequence of actions $\langle A_1,...,A_{n-1}\rangle$ then given the conditional (A_n & C) ➤ P we can infer that for some time t^* between t_n and t_{n+1}, and *P-at-t^** will be made true by performing *A_n-at-t_n*. If *P* is temporally-projectible, we can then infer by temporal-projection that *P* will still be true at t_{n+1}. For this reasoning to work, a projectibility constraint must be added to clause (ii):

(ii) *n* > 0, *P* is temporally-projectibile, and *conditionals* contains a conditional (A_n & C) ➤ P such that *C* is an expectable-result of $\langle A_1,...,A_{n-1}\rangle$.

GOAL-REGRESSION is based directly on clause (ii), so it must contain a corresponding constraint, and it becomes defeasible:

GOAL-REGRESSION

Given an interest in finding a plan for achieving *G-at-t*, if *G is temporally-projectible* , adopt interest in finding planning-conditionals (*A* & *C*) ➤ *G* having *G* as their consequent. Given such a conditional, adopt an interest in finding a plan for achieving *C-at-t**. If it is concluded that a plan *subplan* will achieve *C-at-t**, construct a plan by (1) adding a new step to the end of *subplan* where the new step prescribes the action *A-at-t**, (2) adding the constraint ($t^* < t$) to the ordering-constraints of *subplan*, and (3) adjusting the causal-links appropriately. Infer defeasibly that the new plan will achieve *G-at-t*.

(iii) n > 0, P is an expectable-result of $\langle A_1,...,A_{n-1}\rangle$, and *conditionals* does not contain a conditional of the form (A_n & C) ➤ ~Q such that Q is either P or a conjunct of P and C is an expectable-result of $\langle A_1,...,A_{n-1}\rangle$

- Clause (iii) is, in effect, a statement of TEMPORAL-PROJECTION applied to the expectable-results of an action-sequence, together with the statement of a defeater for the application of TEMPORAL-PROJECTION .
- For TEMPORAL-PROJECTION to be applicable, we must require that P be temporally-projectible.
- Given that constraint, if it is defeasibly reasonable to expect P to be true after executing $A_1,...,A_{n-1}$, then it is defeasibly reasonable to expect P to remain true after executing A_n as well.
- A defeater for this defeasible expectation consists of having a reason for thinking that P will not be true.
- Given that it is defeasibly reasonable to expect C to be true after executing $A_1,...,A_{n-1}$, it follows in accordance with the preceding discussion of the Frame Problem that, given the conditional (A_n & C) ➤ ~Q , it is defeasibly reasonable to expect Q to become false after executing A_n.

(iii) n > 0, P is a temporally-projectible expectable-result of $\langle A_1,...,A_{n-1}\rangle$, and *conditionals* does not contain a conditional of the form (A_n & C) ➤ ~Q such that Q is either P or a conjunct of P and C is an expectable-result of $\langle A_1,...,A_{n-1}\rangle$

Building the temporal-projetibility constraints into (R1) yields:

(R2) Where *start-state* is a state of affairs and *conditionals* is a set of planning-conditionals, P is an expectable-result of $\langle A_1,...,A_n\rangle$ relative to *start-state* and *conditionals* iff either:

(i) n = 0 and P is temporally-projectible and true in *start-state* ; or

(ii) n > 0, P is temporally-projectible and *conditionals* contains a conditional (A_n & C) ➤ P such that C is an expectable-result of $\langle A_1,...,A_{n-1}\rangle$; or

(iii) n > 0, P is a temporally-projectible expectable-result of $\langle A_1,...,A_{n-1}\rangle$, and *conditionals* does not contain a conditional of the form (A_n & C) ➤ ~Q such that Q is either P or a conjunct of P and C is an expectable-result of $\langle A_1,...,A_{n-1}\rangle$; or

(iv) n > 0 and P is a conjunction whose conjuncts are expectable-results of $\langle A_1,...,A_n\rangle$.

Soundness Assumption

Executing a linear plan can be defeasibly expected to achieve a goal G relative to a state *start-state* iff G is an expectable-result of the sequence of actions prescribed by the plan-steps of the plan relative to *start-state* and the set of all true planning-conditionals.

A Semantics for Planning

The definition (R2) of "expectable-result" and the modified soundness assumption constitutes a semantics, relative to which we can prove the soundness and completeness of a set of rules for goal-regression planning.

6. Conclusions

- The theory of means-end reasoning developed here draws heavily on conventional AI planning theory, but there are also important differences.
- A planning agent embedded in a complex environment must interleave planning with epistemic cognition aimed at providing information needed for planning, and this makes r.e. logically impossible. Instead, planning must be done defeasibly.
- To accomplish this I have proposed taking means-end reasoning to be a species of epistemic cognition whose purpose is to generate defeasibly reasonable conclusions of the form "Plan p would achieve its goal if the prescribed plan-steps were executed in any order consistent with its ordering-constraints".

6. Conclusions

- The conventional approach adopts the definition (R1) of "result", and then attempts to prove the soundness and completeness of a planning algorithm.
- Once it is recognized that planning is based upon defeasible expectations rather than objectively determinate results of actions, it becomes apparent that no such definition of "result" is possible.
- The semantics of planning must instead be based upon the epistemic concept of an "expectable-result".

6. Conclusions

- It has been argued here that means-end reasoning presupposes a certain kind of solution to the Frame Problem.
- This requires the imposition of temporal-projectibility constraints in the definition of "expectable-result" and on the subgoals generated by GOAL-REGRESSION.

Inspect OSCAR planner