

Part Ten Adopting Plans

- Plan construction produces plans that purport to achieve their goals, but *adopting* such a plan requires a further cognitive step.
 - Such a plan is not automatically adoptable:
 - » Its execution costs might be greater than the value of the goal achieved.
 - » It might interact adversely with other plans already adopted, increasing their execution costs or lowering the value of their goals.
 - Sometimes the impacted plan should be rejected.
 - Sometimes the new plan should be rejected.
- Nothing is certain.
 - We must discount the value of the goal by the probability of achieving it, and discount the execution costs by their probabilities.
- In other words, plans should be evaluated decision-theoretically.

Two Approaches to DT Planning

- Local Plans
 - Plans have the same structure as classical plans.
 - » Such plans aim at restricted goals like transporting an object to a location, acquiring a certain bit of information, recharging one's battery, etc.
 - » This makes the construction of such plans manageable by goal regression techniques.
 - Expected values are computed by assigning values to possible outcomes, discounting the values by the probabilities of the outcomes actually occurring.
- Global Plans
 - Plans aim to achieve all goals simultaneously.
 - Markov decision planners (MDP's) and partially observable Markov decision planners (POMDP's) build a state space whose nodes represent all possible states of the world.
 - Links between nodes correspond to actions that would move the world from one state to another with some specified probability.
 - Planners search for an optimal path through the space.

A Problem for Local Plans

- This represents too quick a generalization of classical decision theory to plans.
 - Local plans cannot be compared and chosen for execution just by comparing their expected values.
 - The basic difficulty is that classical decision theory concerned acts, which were taken to be unstructured entities that were logically independent of one another.
- Plans can embed one another as subplans, or they can have overlapping parts.
 - This makes it impossible to compare them just by comparing their expected values.
 - » For example, if one plan embeds another as a subplan, its having a higher expected value than its subplan does not automatically make it a better plan.
 - » It may have a higher expected value just because it achieves additional goals.
 - » If there are other ways to achieve those additional goals, it might be better to adopt the subplan together with some other plan for the other goals rather than adopting the single plan that aims to achieve all the goals.

Example

- A plan to run two errands on a single trip is often preferable to plans to run the errands separately, but it is not always preferable.
- Suppose I must pick up a ton of lead and a ton of gold from a single repository and deliver both to the same destination.
 - » Both will fit in my truck, and I could pick them up on a single trip, but doing so would risk damaging my springs.
- Then it might be better to deliver them on two separate trips.
 - » However, the plan to deliver them on a single trip, by virtue of achieving both goals, might have a higher expected value than any single plan with which it competes, e.g., the plan to deliver the gold without delivering the lead.
- What is better than adopting the plan to deliver them both on a single trip is adopting the two separate plans to deliver the gold on one trip and deliver the lead on another trip.
- So the plan with the higher expected value is not the best plan to choose.
- One should instead choose two other plans.

Example — continued

- It is always possible to construct a fourth plan that merges the plans for the two trips into a single plan prescribing both trips, and that plan will have a higher expected value than the plan to deliver both the gold and lead on a single trip.
- But do we want to require agents to always consider all the ways of combining their plans into larger plans before they decide which plans to adopt?
 - » That is computationally problematic.
- It is also subject to a logical difficulty.
 - » If we are allowed to merge plans arbitrarily into larger and more inclusive plans, then in a complex world there may be no optimal plans.
 - » For every plan, it may be possible to construct a preferable plan by merging the first plan with other plans for achieving other goals.
 - » If decision-theoretic planning requires the adoption of optimal plans, this will have the consequence that no plans will be adopted.

A Second Problem for Local Plans

- Both the values of goals and the values of execution costs are typically a function of the circumstances under which a plan is executed, and that in turn will be strongly influenced by what other plans the agent adopts.
 - » To take a trivial example, if my goal is to eat a dish of vanilla ice cream, the value of that goal will be seriously diminished by my adopting a plan that calls for my eating a dill pickle first.
- And even more obviously, execution costs can be seriously affected by the agent's other plans.
 - » A plan to deliver a package will be much harder to execute if a prior plan first takes the agent to the other side of town.
- Most of the literature on decision-theoretic classical planning assumes that execution costs are constant values for each action in a plan, and goals have fixed values.
 - » But for an agent operating in a realistic environment, those assumptions aren't even good approximations.
 - » To compute an expected value for a plan, we must know what other plans the agent has adopted, and the right response to the construction of a new plan may be to adopt it, to reject it, or to adopt it and withdraw an earlier plan (in response to the new plan interacting negatively with that earlier plan).

Problems for Local Plans

- So there are two distinct problems for decision-theoretic classical planning.
- The first is that plans cannot be evaluated in isolation from one another, because they can affect each other's expected values.
- The second is that local plans cannot be chosen for adoption just because they have higher expected values than any competing plans the agent has constructed.

Problems for Global Plans

- The natural response to the problems encountered by local plans is to require plans to be global. Then they are automatically comparable.
- A generally recognized problem for Markov decision planning is that it is computationally infeasible in any but the simplest environments.
 - To estimate the complexity of the real world, it has been estimated that there are 10^{78} elementary particles.
 - If we take the state of a particle to be determined by four quantum states each having two possible values (a gross underestimation), each particle can be in 16 states, and so there are $16^{10^{78}}$ states of the universe.
 - This is a bigger number than we can write in the form $10^{...}$. The length of the exponent would be greater than the number of elementary particles in the universe.

Problems for Global Plans

- Clearly, an optimal policy cannot prescribe actions for all of these different states.
- It must abstract from the true complexity of the universe, making the assumption that most differences between states do not make any difference to how the agent should behave.
- Suppose we could confine our attention to just 300 two-valued variables.
 - » That is pretty unrealistic—it seems clear that many more than 300 parameters can make a difference to optimal behavior, and many of them are continuous-valued rather than two-valued.
- But even if we could confine our attention to 300 two-valued variables, an optimal policy would have to distinguish between 2^{300} states and prescribe behavior for each.
 - » 2^{300} is approximately equal to 10^{90} , which is twelve orders of magnitude larger than the number of elementary particles in the universe.
- Clearly, a real agent can not deal with policies that large, and even such policies would be woefully inadequate because in some cases they would fail to make crucial distinctions.

A Compromise Proposal Locally Global Planning

- MDP's and POMDP's get the logic of the planning problem right, but they do so at the expense of being impractical for realistic planning in complex environments.
- Assigning values to classical plans is a computationally more feasible alternative, but it does not get the logic right.
- Locally Global Planning:
use classical planning techniques to produce plans (but base them upon probabilistic connections rather than exceptionless causal connections), and then reason defeasibly about the expected value, not of individual plans, but of the whole package of plans that the agent has adopted at any one time.

Master Plans

- Although we do not want the agent to consider all the possible ways of merging its local plans into larger plans, we can consider the single plan that results from merging all of the agent's local plans into a master plan.
- I suggest that it is master plans that are the appropriate objects of decision-theoretic evaluation.
- Evaluating a plan in the context of the agent's other plans is just a matter of evaluating its contribution to the value of the master plan.
 - Two qualifications:
 - » adding a plan to the master plan may only increase the value of the master plan if we simultaneously delete other plans from the master plan.
 - » plans may have to be added in groups rather than individually.

Destructive Interference

- In classical planning, the only way plans can destructively interfere with each other is by undermining each other's causal-links.
 - A subplan undermines a causal link if it constitutes a plan for the negation of the subgoal between the time it is produced and the time it is used.
 - A new plan may make it impossible to successfully execute another plan if it introduces steps into the master plan which combine with other steps to form a subplan that undermines a causal-link of the earlier plan.
- In DT planning, the new plan may simply lower the probability that a causal-link will succeed.
- In DT planning, the new plan might also lower the values of the goals achieved by an earlier plan, or raise the execution costs.

Constructing the Master Plan

- We can think of the master plan as the agent's current best approximation to a globally optimal policy.
 - It is "as global" as the agent's current planning has been able to manage, but it makes no attempt to discriminate between all possible states of affairs, and as such is small enough not to overwhelm the agent's memory capacity.

- Although the master plan is of manageable size from the perspective of information storage, it will still be a very large plan by the standards of current AI planning technology.

- It is not unreasonable to expect that at some given time a sophisticated autonomous agent will have adopted 1000 local plans with an average length of 10 steps each.
- That translates into a master plan of 10,000 steps. Furthermore, the local plans will typically deal with a very wide variety of goals and circumstances, and may draw collectively from 1000 different possible actions. The production of such a plan is several orders of magnitude beyond the capabilities of current automated planning algorithms.
- Weld (1999) observes that the current state of the art is represented by BLACKBOX (Kautz and Selman 1998), which can find a plan with 105 steps in a world with 10^{16} possible states.
- That is impressive when compared with previous planning technologies, but calculation reveals that $10^{16} = 2^{53}$, so this is still a world characterized by only 53 two-valued fluents.
- Computing an expected value for a 10,000 step plan is also an immense undertaking.
- It may be impossible to compute it exactly.

Constructing the Master Plan

- The master plan cannot be produced by combining all of the agent's goals into a single conjunctive goal and planning for that from scratch.
- But it can be produced by planning separately for the individual goals as they arise and then merging the resulting local plans into the composite master plan.
 - When they are merged, the agent must be on the lookout for both destructive and constructive interference, but what makes the planning manageable is that it can be assumed defeasibly that there is no interference until some is found.

Constructing the Master Plan

- It is suggestive that human planners seem to proceed in this way.
 - When solving local planning problems we do not look continuously at the big picture.
 - Rather, we find a plan that seems to work subject to local constraints, and then we worry later about how it fits in with the rest of our plans.
 - If we don't see a problem, we assume there is none, although we remain vigilant in case a problem later emerges.
- This is analogous to assuming defeasibly that subplans do not undermine each other when we merge them into a larger plan.

Evaluating the Master Plan

- The evaluation of the master plan can be made feasible by focusing on the evaluation of local plans.
- If the local plans are independent then the expected value of the master plan will be the sum of the expected values of the local plans it comprises.
 - The agent can assume defeasibly that the local plans it produces are independent, and evaluate the master plan accordingly.
 - It can then look for failures of independence, and when they are found, the local plans entering into the dependence can be merged together and evaluated directly, and then it can be assumed defeasibly that that is the value contributed by that set of local plans to the master plan.
 - If subsequent investigation turns up a larger set of dependencies, then the larger set can be merged and evaluated and it can again be assumed defeasibly that that evaluation constitutes the value contributed by that larger set of local plans.
 - In this way the agent never has to evaluate the entire master plan directly. It just evaluates relatively small plans, either local plans or composites of several local plans, and then sums the resulting values to get a defeasible estimate of the value of the master plan.

Locally Global Planning

- The proposal is then that direct plan construction and plan evaluation be confined to local plans or small composites of local plans.
- The master plan is constructed by merging the local plans and it is evaluated defeasibly by summing the values of the small plans.
- This makes the computational task of constructing plans and evaluating changes to the master plan relatively simple.
- This assumes that the agent has the computational tools required for constructing local plans and detecting interference between them.

Decision-Theoretic Goal-Regression

- Proposal — It is possible to perform feasible decision-theoretic planning by modifying conventional goal-regression planning in certain ways.
- Goal-regression planning can be performed by applying classical planning algorithms but appealing to probabilistic connections rather than exceptionless causal connections.
 - This is computationally easier than “probabilistic planning” (e.g., BURIDAN)
 - It isn't really the probability of the plan achieving its goals that is important — it is the expected value. The expected value can be high even with a low probability if the goals are sufficiently valuable.

Decision-Theoretic Interference

- Once a plan is constructed, an expected value can be computed. This computation is defeasible.
- If two plans are truly independent, they can be merged into a single plan and the expected value of that composite plan will be the sum of the expected values of the constituent plans.
- When the expected value of that composite plan is equal to the sum of the expected values of the constituent plans, let us say that they exhibit decision-theoretic independence.
- Decision-theoretic interference is the failure of decision-theoretic independence.

Decision-Theoretic Interference

- The literature on classical deterministic planning also recognizes failures of independence in its treatment of “threats” or “underminings”.
- This is regarded as an obstacle to merging the plans at all.
- But this is just the limiting case of the failure of decision-theoretic independence.
 - What is wrong with merging two plans when one undermines the other is that the undermined plan is then prevented from achieving its goal and hence from contributing its expected value to the expected value of the composite.
- So interference at the level of plan construction can be subsumed under decision-theoretic interference.

Decision-Theoretic Interference

- In adding local plans to the master plan, our defeasible assumption is one of decision-theoretic independence, so what is needed is tools for detecting decision-theoretic interference.
- The standard tools for detecting undermining will be a subspecies of these tools, but they must be generalized to handle the cases in which merging two plans raises or lowers the expected value of the composite rather than preventing one of the constituent plans from contributing anything at all to the value of the composite.

Probabilistic Undermining

- Classical undermining arises when executing the steps of one plan prevents a causal link of the second plan from working.
 - It does this by making the subgoal of the link false between the time it is produced and the time it is used.
- When the causal link records a merely probabilistic connection, a weaker variety of undermining can arise from the steps of the one plan lowering the probability of the subgoal being true when it is to be used.
 - This can happen in two ways.
 - » (1) It could change the probability of the subgoal being made true in the first place, or
 - » (2) it could change the probability of its remaining true until the causal link target is executed.
 - The latter is analogous to classical undermining, and includes it as the limiting case.
 - The former has no non-probabilistic analogue, because in classical planning it is assumed that actions are guaranteed to have their effects, regardless of circumstances, providing only that the preconditions are satisfied.

Probabilistic Undermining

The first kind of probabilistic undermining can be described more precisely as follows. The causal link is based upon a probability $\text{prob}(\text{subgoal} / \text{action} \ \& \ C)$, where C is a context established (or made probable) by the earlier parts of the plan. Undermining occurs when the second plan makes a context C^* probable and $\text{prob}(\text{subgoal} / \text{action} \ \& \ C \ \& \ C^*) \neq \text{prob}(\text{subgoal} / \text{action} \ \& \ C)$.

The second kind of probabilistic undermining is somewhat different. The plan relies upon a probability $\text{prob}(\text{subgoal-at-}t / \text{subgoal-at-}t_0)$. This may just be based upon a defeasible temporal projection, or it may be based upon more concrete probabilistic knowledge. Undermining occurs when there is a context C^* made probable by the undermining plan and an act A^* such that $\text{prob}(\text{subgoal-at-}t / A^* \ \& \ C^* \ \& \ \text{subgoal-at-}t_0) \neq \text{prob}(\text{subgoal-at-}t / \text{subgoal-at-}t_0)$.

Value-Undermining

- A third kind of undermining is *value-undermining*.
- Computing the expected value of a plan depends upon assumptions about the values of goals and side-effects. These are measured by conditional utilities $U(G/C)$ where C is a circumstance made probable by the plan. Value undermining occurs when the second plan makes C^* probable, where $U(G/C \& C^*) \neq U(G/C)$.

Sharing Steps

- Decision-theoretic interference occurs when the value of merging two plans is not the sum of the values of the plans.
- A fourth way in which this can happen is when the plans share steps, because in that case the execution costs of the merged plan will typically be less than the sum of the execution costs of the individual plans, due to something's having to be done only once.
 - This is the “good kind” of interference that improves the merged plan.
 - I suggest that the decision when to share steps between plans should usually be done at the level of the master plan, not at the level of the local plans, because it is a decision that must be based upon the expected values of the results.
 - The only exception to this occurs when resource constraints dictate that something cannot be done more than once, in which case the plan construction algorithm can determine that steps must be shared.

Maximizing vs. Satisficing

- Conventional decision theory would tell us to choose a master plan having a maximal expected value.
- That is at least computationally infeasible in complex domains.
- There may not even be a maximally good plan.
- We should instead *satisfice* — seek plans with positive expected values, and always maintain an interest in finding better plans.
 - A plan is defeasibly adoptable if it has a positive expected value, or if its addition to the master plan increases the value of the latter.
 - The adoption is defeated by finding another plan that can be added to the master plan in its place and will increase the value of the master plan further.
- So we are always on the lookout for better plans, but we are not searching for a single “best” plan.

Implementing DT Planning

- To incorporate decision-theoretic planning of the sort I am describing into an autonomous agent, we need procedures for finding decision-theoretic interference, and also procedures for fixing plans that exhibit it, i.e., making changes to them to eliminate undesirable interference.
- Of course, not all interference is undesirable, so we also need procedures for adding desirable interference.

Implementing DT Planning

- In general, whenever a plan with a positive expected value is found, the agent should adopt interest in finding ways to modify it to increase the expected value.
- These procedures can be modeled on conventional threat detection and plan repair techniques used in classical planning.
 - It can proceed by generalizing UNDERMINE-CAUSAL-LINKS, and supplementing ADD-ORDERING-CONSTRAINTS and CONFRONTATION.
 - The details remain to be worked out, but the idea seems relatively simple.

Hierarchical Planning

- A familiar idea from classical planning is that planning should be hierarchical — plan with “high-level” operators, and then expand them into low-level plans.
- This is usually defended on the grounds that it makes planning more efficient (Barrett and Weld 1994), but it is noteworthy that it can also produce plans with higher expected values.
- This is because we can be more confident of being able to perform a high level action (e.g., *drive across town*) in *some way or other* than we can in being able to perform it in any particular way.
 - If we build a detailed route into our plan, the probability of plan failure may be high, but if we just plan to drive across town some way or other, the probability of plan failure may be low.

Hierarchical Planning

- We make our plans more secure by planning hierarchically, where the high-level tasks are such that we can construct alternate plans when one plan fails or is apt to fail.
- We rely heavily upon being able to fix plans on the fly, and it is largely our decomposition plans that we are fixing or replacing.
- To accommodate this, the master plan must consist in part of high level plans, with links to lower level plans.
- To evaluate the adoption of decomposition plans for the performance of high level actions, it looks like we will have to evaluate the master plan on different levels—both by ignoring decomposition plans, and by adding them into the plan.

Conditional Planning

- Another way to improve expected values is to adopt conditional plans that prescribe different actions under different circumstances.
- If the success of a plan depends upon something (a contingency) being true that has only a certain probability of being true, we may be able to increase the expected value of the plan by adding a plan for what to do if the contingency fails.

Conditional Planning

- In the literature on conditional (or contingency) planning, it is generally assumed that a conditional plan must contain subplans for an exhaustive set of contingencies (see Peot and Smith 1992, and Pryor and Collins 1996).
- In fact, though, it can be useful to plan for a single contingency, or a small list of non-exhaustive contingencies.
- We need not assume that our plans are bullet-proof.
- We need only ensure that they have a reasonable expected-value, and then if they fail we shrug our shoulders and go on.
- On the other hand, we may also be able to increase the expected-value by adding additional subplans for further contingencies.

Conditional Planning

- A fundamental issue in conditional planning is when to do it. That is, for which contingencies should we plan?
- Existing conditional planners do not address this issue, taking it as part of the specification of the planning problem by the user.
- However, an autonomous planning agent must be able to figure this out for itself.
- The solution to this problem must appeal to decision-theoretic considerations.
- Roughly, we should plan for a contingency when we have reason to think that doing so may produce a plan with a reasonable expected-value.
- This will turn in part on the likelihood of the contingency coming true, and prior knowledge of what is apt to happen if it does come true.

Information Gathering

- To execute a conditional plan, we must know whether a contingency is true.
 - It is initially tempting to think that the antecedents of conditionals in conditional plans are the contingencies themselves.
 - This accords with the way we think about cases where it is obvious to us whether the contingencies hold (e.g., it is raining).
 - But as Pryor and Collins (1996) observe, sometimes information gathering will be complex, and we must make sure that the steps involved in gathering information don't conflict with the plan itself.
- This suggests that the plan should incorporate the information gathering steps as part of it.

Information Gathering

- The incorporation of information-gathering can be accommodated by making the antecedents of the conditionals epistemic:
 - “If you know that P, do A”, “If you know that ~P, do B”.
 - We could also have a condition “If you don't know that P, do C”, or “If you are uncertain whether P, do C”.
 - The latter could not be captured by taking the antecedents to be non-epistemic.
- The cases in which it will be obvious whether the contingencies hold can be accommodated by employing high level operators like “Observe whether P”.
 - Then we do not have to plan explicitly for information gathering, leaving that planning until the time of execution.
 - So there is an important interaction between hierarchical planning and conditional planning here.

Information Gathering

- An independent argument for including knowledge acquisition steps is that at the time when you want to know whether P, it may be too late to take the steps that are required for getting that knowledge.
 - For instance, you might have to ask a question of a friend before he leaves, or use the phone before you do something that makes it inaccessible.
- So you will need to incorporate these steps into the plan.
- However, in cases in which you expect it to be easy to acquire the knowledge when you need it, you can leave the details undetermined.

Information Gathering

- The inclusion of information acquisition steps interacts in important ways with plan execution.
 - It is tempting to suppose that in executing a plan the agent should monitor the course of execution and verify that contingencies hold and subgoals are achieved before continuing the execution.
 - However, that will not always be possible.
 - » For example, one step of a plan might involve calling a friend and asking him to do something.
 - » I may have no way of verifying that he does it.
 - » The best I can do is continue plan execution on the assumption that the subgoal has been achieved.
 - » If it hasn't, then I incur execution costs that do not actually contribute to the achievement of the goal.
 - What this illustrates is that in computing the expected value of a plan, one must take account of the possibility that plan execution will continue, incurring execution costs, even though the plan has already failed.
- This suggests that monitoring steps should be explicitly included in the plan when they are to be employed, because whether we monitor plan execution can affect the expected value of the plan.

Plan Execution

- The master plan also provides a useful database for plan execution.
- Execution requires epistemic monitoring to
 1. ensure that things are going as planned
 2. determine when or whether the antecedents of conditional steps are satisfied.
- We must update the master plan as steps are performed, and update its evaluation.
- It is defeasibly reasonable to expect the master plan to remain adoptable as steps are performed.
 - However, new information may alter this defeasible expectation, forcing the abandonment or modification of constituent plans in the master plan.

Realistic Planning

- Realistic planning in autonomous agents operating in complex and partially unpredictable environments must be
 - continuous
 - decision-theoretic
 - hierarchical
 - conditional
 - classical.

Realistic Planning

- The key to my proposed solution to this nexus of problems is that plan search should proceed classically, but using probabilistic connections rather than exceptionless causal connections.
- The kind of classical planning involved should be both hierarchical and conditional.
- The plans produced must be evaluated decision-theoretically, but that is done by incorporating them into the master plan and using defeasible rules for evaluating it.
- The master plan will also provide the data structure for keeping track of decomposition plans, and updating the planning task as new goals are adopted and plan steps are performed in the course of executing plans.