The STRIPS Assumption for Planning Under Uncertainty

Michael P. Wellman
AI Technology Office, Wright R&D Center
Wright-Patterson AFB, OH 45433
wellman@wrdc.af.mil

Abstract

The virtue of the STRIPS assumption for planning is that it bounds the information relevant to determining the effects of actions. Viewing the "assumption" as a statement about beliefs, we find that it does not actually assume anything about the world itself. We can characterize the assertion about beliefs in terms of probabilistic independence, thereby facilitating analysis of representations for planning under uncertainty. This interpretation separates the STRIPS assumption from other necessary features of a planning architecture, such as its model of persistence and its inferential policies. By isolating these factors, we can understand the role of dependence across a wide range of planners and action representations. Graphical models of dependence developed for probabilistic analysis provide a convenient tool for verifying the STRIPS assumption for a variety of planning systems. Investigation of a few representative systems reveals a Markovian event structure common to these planning models.

The Frame Problem and the STRIPS Assumption

The classic dilemma in representing and reasoning about the effects of actions is the *frame problem*, originally identified by McCarthy and Hayes [1969]. The frame problem has come to stand for a variety of computational and notational complexities arising from the apparent necessity of considering the possible change in status of every proposition for each action. Characterizations of the problem vary widely [Brown, 1987; Pylyshyn, 1987], proposed solutions even more so, but a kernel of consensus does seem to exist. AI researchers agree that part of the problem, at least, has to do with specifying the effects of actions without explicitly describing all ramifications and qualifications. In particular, we want to avoid a requirement for explicit *frame axioms* specifying the propositions *not* affected by each action.

Actual planners eschew frame axioms and restrict attention to propositions explicitly mentioned in their action specifications, a convention first applied by STRIPS [Fikes and Nilsson, 1971]. Waldinger has named this policy the "STRIPS assumption" [1977]. McDermott [1987] asserts that no program since STRIPS has been practically bothered by the frame problem, which is true if we define the problem narrowly as the need for frame axioms in the deductive planning approach. Nevertheless, building planners that perform well in complex dynamic environments is no easy task, in large part due to difficulties of representing and reasoning about change. As McDermott also points out, the frame problem *does* frustrate attempts at logical analyses of these programs and their environments, which should be relevant to the goal of designing better algorithms and representations. Much of the work on nonmonotonic logic is addressed to this issue.

Understanding the nature of the STRIPS assumption and the extent to which it circumvents the frame problem is a first step to understanding the larger issues in reasoning about actions. Previous discussions of the STRIPS assumption, including the original by Waldinger [1977], tended to encompass all of these issues, failing to distinguish the relatively small role played by this particular notational convention. While the broader views provide fuller accounts of the planners they address, their analyses are not transferable to planning frameworks that take significantly different approaches to representing and reasoning about change.

For example, Lifschitz's analysis [1986] focuses on conditions under which STRIPS's add/delete mechanism will be guaranteed to produce only valid plans. The analysis concludes essentially that STRIPS systems are sound as long as

1. the use of non-atomic sentences in operator descriptions and world models is restricted (in a precise manner described by Lifschitz), and
2. a kind of strong persistence holds, where no changes occur except as specified in add and delete lists.

These conditions (which clarify STRIPS significantly) apply to planning frameworks that adopt the same strong persistence model and forbid inference about the further consequences of specified effects. Many have been unwilling to accept these restrictions, and have worked on methods and semantic accounts of systems that go beyond them.

Research on the task of determining the implications of specified effects of actions (called the *ramification problem*), its counterpart for preconditions (the *qualification* problem), and development of models of persistence are important areas of investigation for AI planning. The point of this paper is that there is a separable aspect of the STRIPS assumption that is orthogonal to these issues, and therefore applicable across a variety of planning frameworks. Generally stated, the STRIPS assumption *per se* dictates that the effect of an action on the world model be completely determined by the direct effects explicit in its specification. By saying only that it is "completely determined," we permit the nature of the implicit effects to vary among planning systems.

I examine this interpretation below from the perspective of planning under uncertainty. Uncertainty provides further motivation for this view of the STRIPS assumption, and concepts from uncertain reasoning help to characterize it more precisely for application to existing planning frameworks.

Planning under Uncertainty

An agent plans *under uncertainty* whenever it cannot flawlessly predict the state of the environment resulting from its actions. By this definition, uncertainty is a characteristic of the agent's knowledge rather than an inherent property of the environment. Given that we are never likely to achieve perfect prediction in realistic environments, all planning is actually performed under uncertainty; planning under *certainty* is an unrealizable—albeit often useful—idealization.

The frame problem arises in planning under uncertainty just as it does in the idealized framework. Planners must employ something like the STRIPS assumption to justify leaving non-effects of an action implicit in their omission from the action's specification. However, semantic accounts of the STRIPS assumption in classical planning (e.g., STRIPS itself [Lifschitz, 1986]) do not easily map over to the uncertain case. The conventional interpretation, that planners assume that relations in the world model are unchanged unless explicitly specified, cannot literally apply to planners that admit they have incomplete knowledge about the effects of their actions on the world.

Perhaps we could modify the interpretation to assume that changes of unspecified relations are unlikely rather than impossible. The problem of this approach is identifying a particular, well-motivated, likelihood assumption that is sufficiently general for domain-independent planning. As demonstrated by some work along these lines [Dean and Kanazawa, 1988; Hanks, 1990], defining such a convention is tantamount to adopting a model of persistence and probabilistic inference. Moreover, these persistence models tend to be more varied and complicated than those proposed for planning under certainty. These differences provide further motivation for a characterization of the STRIPS assumption that does not depend on a particular model of persistence.

The essential property of the STRIPS assumption that justifies implicit treatment of non-effects is the presumption that the information specified explicitly is sufficient to describe the agent's change in belief. In other words, once the direct effects are known, knowledge of the action itself is superfluous for purposes of prediction. Thus, the STRIPS assumption is fundamentally a statement that the agent's beliefs about changes in propositions not mentioned in an action's specification are independent of the action, given those effects explicitly specified. For planning under uncertainty, we can characterize beliefs in terms of probability distributions and use the concept of probabilistic independence to formalize this interpretation of the STRIPS assumption.

Probabilistic Independence

In a state of uncertainty, an agent's beliefs are representable by a probability distribution over possible situations (which is not to say that the agent's beliefs need be encoded as such in some data structure). We take situations to be assignments on a universe of variables describing the world, including such things as what actions are performed and their consequences. Belief states are then probability distributions over this space. Note that this framework avoids imposing a temporal ontology, which, while providing essential structure for the planning problem, would also detract from the generality of our analysis of the issue at hand.

To capture the meaning of the STRIPS assumption proposed above, we need a way to express the sufficiency of explicitly specified effects to describe the full impact of an action on the agent's beliefs about the world. For this purpose, the natural concept in probability theory is *conditional independence*. We say that random variables x and y are conditionally independent given z iff

$$\Pr(x|y,z) = \Pr(x|z) \quad (1)$$

for any possible values of the variables. In other words, once the value z is known, finding out the value of y has no effect on the agent's belief about x. In this case, y is superfluous information.

The STRIPS assumption is paraphrased by a schema for equation (1). Performance of an action is represented by y, z stands for the explicit effects of y plus the "background," and x represents "everything else." The independence assertion is that for a given background, knowing the explicit effects of y provides all the information useful for predicting its implicit effects, that is, everything else. Given its explicit effects, knowledge about the action's performance is redundant.

For a satisfactory interpretation, we need a more complete account of concepts like "background" and "everything else." To understand their role in planning systems, we investigate a class of representations for actions and events based on graphical models. Graphs provide a formal language for expressing (via adjacency) the locality of explicit effects in planning representations.

Graphical Dependence Models

A *probabilistic network* (also called a *Bayesian* or *belief network* [Pearl, 1988] or *influence diagram*) is a directed acyclic graph (DAG) with nodes for the random variables connected by links indicating probabilistic relations. Associated with each node is a probability distribution for its variable given the possible values for its predecessors in the graph. Thus, a link from x to y indicates that y might depend probabilistically on x. Conversely, the absence of links restricts the dependencies that can be encoded in the network. The graphical condition for conditional independence in probabilistic networks is called *d-separation* [Geiger, 1990; Pearl *et al.*, 1989]. Two nodes x and y are d-separated by a set of nodes Z in a DAG iff for every *undirected* path between them either:

1. there is a node $z \in Z$ on the path with at least one of the incident edges leading out of z, or
2. there is a node z' on the path with both incident edges leading in, and neither z' nor any of its successors are in Z.

A dependency graph for which all d-separations are valid conditional independencies is called an *I-map*. Although any joint distribution can be represented graphically by some probabilistic network (which are all I-maps), the most efficient representations are those without superfluous links, called *minimal* I-maps.

We can characterize the independence condition underlying the STRIPS assumption in terms of these dependency graph concepts. Consider a probabilistic network with variables for all actions and events relevant to the planner. Every action node has an outgoing link exactly to those events explicitly represented as direct effects. Events may have arbitrary connections among themselves, as dictated by some world model (outside the scope of discussion here). Action nodes have no incoming links, reflecting our presumption that the planner has control over which actions are to be performed.

The STRIPS assumption is that the graph so constructed is an I-map. Let S_a be the set of event variables that action variable a directly affects, a's immediate successors in the dependency graph. By virtue of *I*-mapness, a is conditionally independent of any $e \notin S_a$ given e's predecessors (see, for example, [Wellman, 1990b, Lemma 4.1]). Each predecessor d of e, in turn, is either a direct effect of a or is conditionally independent given its own predecessors. Ultimately, the effect of a on e is completely determined by a's direct effects and e's relation to them. Note that we still need to describe the interaction, if any, between a and e in their joint effects.

The probabilistic STRIPS assumption does *not* require that a be conditionally independent of e given the direct effects S_a, or even by any subset of S_a. In Figure 1, for example, a and e are d-separated by $\{s, b\}$ but by no other variable set. The variable b is necessary for conditional independence of a and e even though b itself is unconditionally independent of a.

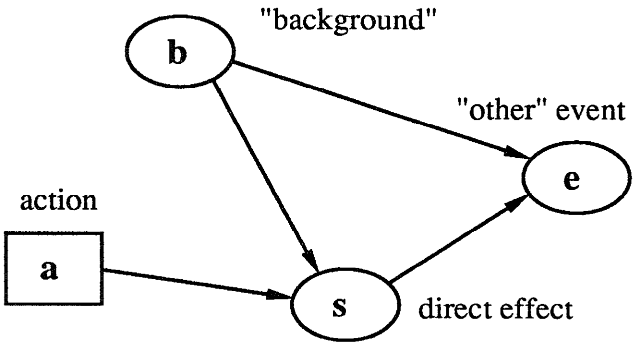

Figure 1: Action a is conditionally independent of e given $Z_a = \{s, b\}$ but not given any subset of its direct effects $S_a = \{s\}$.

If we enlarge the conditioning set to include predecessors of a's direct effects, however, we get another valid independence condition. Let $Z_a = S_a \cup B_a$, where B_a (the "background") is the set of variables that affect a's direct effects:

$$B_a = \bigcup_{s \in S_a} predecessors(s) - \{a\}.$$

The d-separation condition implies that a is conditionally independent of e given Z_a. In the graph of Figure 1, for example, the background $B_a = \{b\}$, and $Z_a = \{s, b\}$.

The dependency graph model permits us to formalize the STRIPS assumption in terms of probabilistic conditional independence. In particular, there must exist an I-map of variables in the world model where any variable e not specified as an effect of action a is not directly connected to a. Under this condition there may be a probabilistic dependency between a and e in some situations, but this can always be described in terms of a's and e's relations to S_a.

We can apply the graph construction to the informal statement of the independence condition given in the previous section. Filling in the terms, our statement is that the complete effects of an action a are fully specified by the direct effects, S_a, and the background, B_a. Everything else, e, is implicit in these variables. That is, e is conditionally independent of a given S_a and B_a. The fragment of Figure 1 can serve as a graphical schema for this pattern of relations, by interpreting the nodes as sets of variables and permitting the variables e to be connected via arbitrary paths to b and s.

Applications

The conditional independence interpretation is a valuable tool for studying specific planning systems and validating their use of the STRIPS assumption. A practical prerequisite for applying these results is identifying the relevant background context, B_a, for the various planners. Note that while planners adopt different policies regarding *how* the implicit effects are derived from the explicit effects and background, validity of the STRIPS assumption does not depend on these policies.

In the following sections I illustrate the application of the independence concepts by analyzing aspects of three planning systems. The planners examined differ in their probabilistic or deterministic representations for the effects of actions, as well as the type of temporal structure imposed on the planning environment.

SUDO-Planner

SUDO-PLANNER [Wellman, 1990a] uses *qualitative probabilistic networks* (QPNs) [Wellman, 1990b], abstractions of the models described above, for representing and reasoning about the effects of actions. When introducing actions and events of interest, the planner modifies the structure of the existing network to preserve the model's validity. One class of constructs appearing in SUDO-PLANNER's knowledge base, called *Markov influences*, specify the effect of an action on an event variable and its dependence on the previous value of that variable.

For example, consider a QPN for a medical therapy problem that includes a variable for the extent of a patient's coronary artery disease (CAD). One action considered by the planner is a coronary artery bypass graft (CABG): bypass surgery to alleviate the coronary disease. The effect of CABG is to decrease CAD (in a precise probabilistic sense [Wellman, 1990b]). Furthermore, the Markov influence specifies that the decrease is greater for patients who have more severe CAD initially. This relationship refers to the variable CAD at two distinct points in time—before and after CABG—and thus cannot be captured by simply adding CABG to the network. Instead, SUDO-PLANNER modifies the QPN by splitting CAD into two variables, CAD_1 and CAD_2. Figure 2 diagrams the result of this *mitosis* process. CABG negatively influences CAD_2, which is otherwise positively related to its value before surgery, CAD_1. The boxed minus sign indicates the synergistic interaction of CABG with CAD_1. Predecessors of the original CAD variable are connected to CAD_1, while its successors before processing the Markov influence are transferred to CAD_2.

This process has direct implications for conditional independence (which indeed was the reason for calling

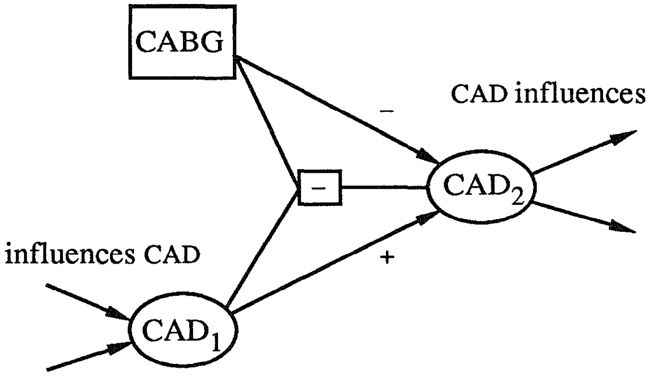

Figure 2: The Markov influence of CABG on CAD.

them *Markov* influences). Specifically, influencers of the original variable CAD are independent of CAD_2 given CAD_1, and CAD's original influences cannot depend on CAD_1 given CAD_2. (The reason is that any path between influences and influencers that circumvents the CAD_i variables must include at least one node with both incident edges leading in.) These conditions in turn imply that any variable in the network is independent of CABG given CAD_1 and CAD_2.

More generally, suppose the action a is defined exclusively by Markov influences on a set of event variables E. The STRIPS assumption dictates that the effects of a be completely captured by these influences. The corresponding independence condition is that any other event be conditionally independent of a, given $Z_a = S_a \cup B_a$, where the direct effects $S_a = E_2$, the second halves of the split event variables, and the background $B_a = E_1$, the first halves produced by SUDO-PLANNER's variable mitosis process.

Markov Influence Diagrams

Kanazawa and Dean [1989] propose a framework for planning under uncertainty based on "causal models," influence diagrams with the Markov property and some other features inessential for our purposes. In a Markov influence diagram, there is a node corresponding to every proposition of interest at every distinguished instant of time. The Markov property is enforced by permitting nodes at time t to depend only on nodes from time $t-1$. If this convention applies to actions as well, then any event at time t is d-separated from actions at time $t' < t$ by the action's direct effects (all at time $t'+1$), plus the events of time t'.

Figure 3 depicts the generic structure of a Markov influence diagram. Note that all links relate an action or event to an event at the next time point. It is possible to relax this restriction—for example, by adding auxiliary atemporal variables—as long as the regularity of temporal states is retained. See, for instance, the variant scheme described by Berzuini et al. [1989].

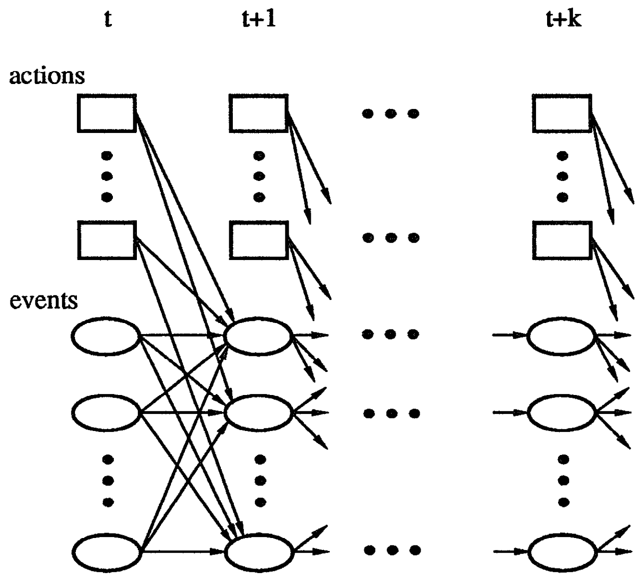

Figure 3: A schematic view of Markov influence diagrams. Actions and events at time t need not be connected to every event at $t+1$.

With respect to our statement of the STRIPS assumption, the background is (conservatively) the state of the world at the time of the action, that is, the actions and events appearing in the same column of Figure 3. The direct effects are the events from the next time point with links from the action node. In schemes for probabilistic temporal projection [Dean and Kanazawa, 1988; Hanks, 1990], these links are typically specified by a set of *causal rules* associating actions and events with their possible consequences.

STRIPS

Instantiated propositionally for a finite world, we find that a STRIPS model is actually a degenerate kind of Markov influence diagram. The Markov property follows from the linearity of the situation calculus framework [McCarthy and Hayes, 1969]. The propositions at a situation s are deterministic functions of those of the previous situation. And under Lifschitz's soundness conditions [1986], the relation of functional dependence is defined by mention in add and delete lists. The persistence model of STRIPS is that for a given proposition this function is the identity in the absence of an action performed at s affecting that proposition. Thus, the background required for any proposition is only its value in the previous situation.

In terms of dependency graphs, the STRIPS model is constructed as follows. Let a and e denote action and event types, with nodes a_t and e_t for every type in every situation t.

$$\begin{aligned} S_a &= \{e \mid e \text{ in } a\text{'s add or delete list}\}, \\ \textit{predecessors}(e_t) &= e_{t-1} \cup \{a_{t-1} \mid e \in S_a\}. \end{aligned}$$

Most STRIPS-like systems do not specify what happens when actions are performed in situations where their preconditions do not hold. To represent "context-dependent effects" [Pednault, 1988], we need only to add events mentioned in preconditions to the background of affected events.

For deterministic variables, there is a stronger graphical criterion for conditional independence, called *D-separation* (note uppercase) [Pearl *et al.*, 1989]. Although the independence condition for STRIPS's simple graph structure is trivial, the more powerful criterion might be useful for analyzing STRIPS-like systems that permit logical and perhaps probabilistic relations among propositions.

Summary

The interpretation presented here provides a new perspective on the STRIPS assumption, constraining the semantics of a planner's knowledge base of actions and events. Essentially, it mandates that the implicit consequences of an action be completely specified by its direct effects. Although described and motivated in terms of probabilistic conditional independence, the interpretation has implications for planning systems regardless of whether they employ probabilistic representations. Moreover, it is sufficiently general to capture the principle behind the STRIPS assumption for planning systems with action-event representations considerably more expressive than that of STRIPS.

The main advantage of this approach is that it distinguishes the concept of belief dependency from the model of persistence of events in the world. It does not obviate the need for such a persistence theory, though it renders the issue orthogonal to the STRIPS assumption *per se*.

Examination of a variety of planning systems indicates that the dependency graph is a useful analytical tool for investigating the structure of relations among actions and events. When the analysis reveals regularity in this structure (as should be expected for reasonable planning architectures), general d-separation patterns can be derived, yielding constraints on the impact of actions on the agent's beliefs about the world. We can then exploit these constraints to design more efficient action representations and belief revision mechanisms.

It is not surprising that all the analyses of planning systems above appeal to some sort of Markov property. Any temporal structure on a pattern of conditional independence constitutes a Markovian form of model. This suggests that the theory of Markov models may be a good place to search for structured patterns of uncertain relationships among events over time.

The most important limitation of the analysis is that dependency graphs are an inherently propositional representation. Application to planning systems with quantified constructs (any nontrivial action and event representation) requires some instantiation mechanism. A potential solution approach is to apply the first-order axioms of conditional independence directly. This technique might be beneficial even for the propositional case, as the axiomatic theory may be stronger than the graphical [Pearl *et al.*, 1989].

References

[Berzuini *et al.*, 1989] Carlo Berzuini, Riccardo Bellazzi, and Silvana Quaglini. Temporal reasoning with probabilities. In *Proceedings of the Workshop on Uncertainty in Artificial Intelligence*, pages 14–21, Windsor, ON, 1989.

[Brown, 1987] Frank M. Brown, editor. *The Frame Problem in Artificial Intelligence: Proceedings of the 1987 Workshop*. Morgan Kaufmann, 1987.

[Dean and Kanazawa, 1988] Thomas Dean and Keiji Kanazawa. Probabilistic temporal reasoning. In *Proceedings of the National Conference on Artificial Intelligence*, pages 524–528, 1988.

[Fikes and Nilsson, 1971] Richard E. Fikes and Nils J. Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208, 1971.

[Geiger, 1990] Dan Geiger. Graphoids: A qualitative framework for probabilistic inference. Technical Report R-142, UCLA Computer Science Department, January 1990.

[Hanks, 1990] Steven John Hanks. Projecting plans for uncertain worlds. Technical Report 756, Yale University Department of Computer Science, January 1990.

[Kanazawa and Dean, 1989] Keiji Kanazawa and Thomas Dean. A model for projection and action. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 985–990, 1989.

[Lifschitz, 1986] Vladimir Lifschitz. On the semantics of STRIPS. In Michael P. Georgeff and Amy L. Lansky, editors, *Reasoning about Actions and Plans: Proceedings of the 1986 Workshop*, pages 1–9. Morgan Kaufmann, 1986.

[McCarthy and Hayes, 1969] J. McCarthy and P. J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie, editors, *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.

[McDermott, 1987] Drew McDermott. AI, logic, and the frame problem. In Brown [1987], pages 105–118.

[Pearl *et al.*, 1989] Judea Pearl, Dan Geiger, and Thomas Verma. Conditional independence and its representations. *Kybernetika*, 25:33–44, 1989.

[Pearl, 1988] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.

[Pednault, 1988] Edwin P. D. Pednault. Extending conventional planning techniques to handle actions with context-dependent effects. In *Proceedings of the National Conference on Artificial Intelligence*, pages 55–59. AAAI, 1988.

[Pylyshyn, 1987] Zenon W. Pylyshyn, editor. *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*. Ablex, 1987.

[Waldinger, 1977] Richard Waldinger. Achieving several goals simultaneously. In E. Elcock and D. Michie, editors, *Machine Intelligence 8*, pages 94–136. Edinburgh University Press, 1977.

[Wellman, 1990a] Michael P. Wellman. *Formulation of Tradeoffs in Planning Under Uncertainty*. Pitman and Morgan Kaufmann, 1990.

[Wellman, 1990b] Michael P. Wellman. Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence*, 1990.