

## **Plan Evaluation with Incomplete Action Descriptions**

Andrew Garland  
Neal Lesh

TR-2002-05 January 2002

### **Abstract**

This paper presents a framework that justifies an agent's goal-directed behavior, even in the absence of a provably correct plan. Most prior planning systems rely on a complete causal model and circumvent the *frame problem* by implicitly assuming that no unspecified relationships exist between actions and the world. In our approach, a domain modeler provides explicit circumscribed statements about which actions have been incompletely specified. Thus, an agent can minimize its dependence on implicit assumptions when selecting an action sequence to achieve its goals.

*Submitted to AAAI-02*

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Submitted January 2002.

# Plan Evaluation with Incomplete Action Descriptions

Andrew Garland and Neal Lesh  
Mitsubishi Electric Research Laboratories  
{garland,lesh}@merl.com

## Abstract

This paper presents a framework that justifies an agent’s goal-directed behavior, even in the absence of a provably correct plan. Most prior planning systems rely on a complete causal model and circumvent the *frame problem* by implicitly assuming that no unspecified relationships exist between actions and the world. In our approach, a domain modeler provides explicit circumscribed statements about which actions have been incompletely specified. Thus, an agent can minimize its dependence on implicit assumptions when selecting an action sequence to achieve its goals.

## 1 Introduction

Traditional work on planning makes the simplifying assumption that the planner is given a complete and correct description of the current state of the world, and a complete and correct model of actions it can perform. Past research on dealing with incomplete information has developed two kinds of approaches: *acquire* better information or *minimize* dependence on incomplete information.

Both approaches have been used to overcome incomplete state information. The acquisition approach involves plans with sensing actions (Etzioni *et al.*, 1992), while the minimization approach typically involves constructing plans that will succeed in all states of the world consistent with the known information (Smith & Weld, 1998). To date, however, only the acquisition approach has been used to overcome incompleteness in action models. In particular, machine learning techniques have been applied to the results of executing actions in order to acquire more accurate action descriptions (Gil, 1994; Wang, 1995; Oates & Cohen, 1996).

We describe novel techniques for minimizing dependence on incomplete information in a given action model. We present a tractable algorithm for identifying the *risks* of a plan, each of which represents a potential source of execution failure due to incompleteness of the action model. Our work is motivated by the observation that producing complete action descriptions is an exceedingly difficult and time-consuming task. In our approach, the domain experts who generate the action descriptions can provide additional information about the completeness of the model, similar to statements used to reason about incomplete states (Levesque, 1996; Golden, 1998; Babiak & Schmolze, 2000). For example, the experts can indicate which actions have been completely described, or that executing an action will not

change the truth value of a predicate. This information can be used to eliminate risks of candidate plans.

Essentially, we are revisiting the infamous *frame problem* (McCarthy & Hayes, 1969), which motivated the traditional completeness assumptions (Reiter, 1991). Without the completeness assumptions, it is necessary for the effects of an action description to include all of the facts whose truth value does *not* change as a result of executing it. The conceptual shift we make is to a willingness to execute plans with some risk of failure if the action model is incomplete. Our methods simply prefer to execute plans with fewer risks.

The advantage of our approach is to execute plans that achieve their goal more often. As a simple example, a plan composed of completely-modeled actions is more likely to succeed than one with incompletely-modeled actions. Additionally, our techniques will sometimes prefer certain orderings of actions over others, the inclusion of “extra” steps that reestablish conditions that could have been clobbered, and the choice of one action over another. We show that even in the case where the domain expert provides no information about the completeness of the model, our techniques can select plans that are more likely to succeed than if the incompleteness of the model is ignored.

In this paper, we present a linear-time algorithm for identifying risks and a flexible framework for evaluating plans based on risks. We prove that plans with certain *critical* risks can not be guaranteed to succeed. Finally, we describe two experimental case studies into the impact of our techniques on plan success.

## 2 Formulation

We assume a traditional planning representation. Each action description consists of a specification of its preconditions and effects as sets of fluent literals. We assume that for any planning domain, there exists a set of complete and correct action descriptions, which we will refer to as  $D_{true}$ . Executing an action in a state in which all of its preconditions are true will produce a new state in which all of its effects are true and all other fluents remain unchanged. Executing an action with any false preconditions has no effect on the state. A *plan* is a sequence of actions.

The set of action descriptions  $D$  given to the planner, however, will be a subset of the information in  $D_{true}$ . That is, each action description  $a_i$  in  $D$  may include only a subset of  $a_i$ ’s preconditions and effects in  $D_{true}$ .

In order to focus on issues related to incomplete action models, we do not address uncertainty in the state in this paper. Instead, we assume that the planner is given a complete state of the world,  $s_0$ . However, this does not sidestep any fundamental issues because executing actions with incompletely modeled effects creates uncertainty about the state of the world.

In our discussion, we will often need to reason about whether a plan will be successful given the information contained in  $D$ . We define  $achieves(D, s_0, g, p)$  as returning true iff the goal  $g$  is true in the state that results from simulating the execution of plan  $p$  from state  $s_0$  assuming that  $D$  is correct and complete.

## 2.1 Locally closed-world statements

We now present a language for supplementing  $D$  with statements about the completeness of action descriptions. Our approach follows the use of locally closed-world (LCW) statements that are used to overcome incomplete state information. In that setting, an LCW statement expresses limited knowledge of the state, such as that all the files that exist in a particular directory are known by a software agent (Golden, 1998; Babaian & Schmolze, 2000).

In our work, an LCW statement expresses limited completeness of action descriptions. The statement  $DoesNotRelyOn(a, x)$  asserts that action  $a$  does not have precondition  $x$  or  $\neg x$  in  $D_{true}$ . We define  $CompletePreconditions(a)$  as a shorthand for the formula:

$$\forall x. x \notin preconditions(a) \supset DoesNotRelyOn(a, x)$$

which is equivalent to explicitly stating that the completeness of preconditions assumption is valid for  $a$ .

The statement  $DoesNotMakeTrue(a, x)$  asserts that  $a$  does not have effect  $x$  in  $D_{true}$ . The statement  $DoesNotMakeFalse(a, x)$  asserts that  $a$  does not have effect  $\neg x$  in  $D_{true}$ . We define  $CompleteEffects(a)$  as a shorthand for the formula:

$$\forall x. x \notin effects(a) \supset DoesNotMakeTrue(a, x) \wedge DoesNotMakeFalse(a, x)$$

which is equivalent to explicitly stating that the completeness of effects assumption is valid for  $a$ .

## 2.2 Plan selection

A *plan selection problem* is a 5-tuple  $(g, C, s_0, D, L)$ , where  $g$  is a goal,  $C$  is a set of candidate plans,  $s_0$  describes the initial state,  $D$  is a potentially incomplete action model, and  $L$  is a set of locally closed world statements. The ideal solution to this problem is to find the plan that will achieve its goal given the actual action descriptions. More precisely, the objective is to find  $c_i \in C$  so that  $achieves(D_{true}, s_0, g, c_i)$ . The quality of a plan selection algorithm can be measured by how frequently it chooses a plan that actually achieves  $g$ . Another measure is how often the plan selection algorithm indicates that no plan should be executed if, in fact, no plan in  $C$  will achieve  $g$  given  $D_{true}$ .

## 3 Algorithms and Analysis

We consider the following four types of risks, each of which is a potential source of execution failure, for a plan com-

posed of actions  $a_1, \dots, a_n$ . Without loss of generality, we encode the initial state as an action  $a_0$  with no preconditions and effects  $s_0$  and encode the goal  $g$  as action  $a_{n+1}$  which has preconditions  $g$  and no effects.

- $POSSCLOB(a_i, x) ::$  action  $a_i$  might have the effect  $\neg x$ , and there exists action  $a_j$ , for  $i < j$  that has precondition  $x$  in  $D$  and no action between  $a_i$  and  $a_j$  has effect  $x$  in  $D$ .
- $PRECOPEN(a_i) ::$  action  $a_i$  might have an unlisted precondition that will not be true when  $a_i$  is executed in  $a_0, \dots, a_n$ .
- $UNLISTEDEFFECT(a_i, x) ::$   $a_i$  might have effect  $x$  that is not listed in  $D$ .
- $PRECFALSE(a_i, x) ::$   $a_i$  has a precondition  $x$  in  $D$  which will be false when executed in  $a_0, \dots, a_n$ , according to  $D$ .

The first two types of risks correspond to relying on knowledge implied by completeness assumptions, i.e., that  $D$  is a good approximation to  $D_{true}$ . In contrast, the latter two risks rely on the incompleteness of the model in order to justify plans that would not necessarily succeed even if  $D = D_{true}$ . For example, an action with a  $PRECFALSE(a_i, x)$  risk cannot successfully execute unless a previous action has an  $x$  effect that is not in  $D$ , or the action which negates  $x$  has a precondition that is not in  $D$ . If no plan composed of the actions in  $D$  can achieve goal  $g$  then it is worth considering plans with other possible risks. For example, it seems preferable to execute a plan with one invented effect than one with five invented effects. However, these plans must be weighed against the alternative of doing nothing, which can be the best option if the chance of achieving a goal is small. Assessing these options would be more practical within a probabilistic framework.

The function `FINDRISKS`, shown in Figure 1, produces the set of risks of a plan  $c = a_1, \dots, a_n$ , a goal  $g$ , and a set of locally closed world statements  $L$ . The algorithm requires  $O(nm)$  time, where  $n$  is the length of the plan, and  $m$  is the number of literals in  $D$ . The pseudo-code is streamlined by assuming that  $c$  is a provably-correct plan in  $D$ ; therefore, there is no need to check for unlisted effects, false preconditions, or that  $g$  is achieved.<sup>1</sup> A reasonable heuristic for most domains is to assume that for any action  $a_i$  and literal  $x$  such that  $x$  is in  $PRECS(a_i)$  but is not in  $DELS(a_i)$ ,  $DoesNotMakeFalse(a_i, x) \in L$ .<sup>2</sup>

`FINDRISKS` requires one pass over the plan, working backwards from the last action. During processing, it keeps track of the set of literals that must eventually be achieved by an earlier action or be true in the initial state of the world. Each action  $a_i$  may provide support for some of these literals, namely the literals that are made true by the action. For each remaining unsupported literal  $x$ , the set of LCW statements  $L$  is checked to see if  $a_i$  might clobber  $x$ ; if so,  $POSSCLOB(a_i, x)$  is added to the result. Also,  $L$  is checked to see if  $a_i$ 's preconditions are completely modeled; if not,  $PRECOPEN(a_i)$  is added to the result. It might seem strange

<sup>1</sup>A more general version of this algorithm exists that efficiently finds these risks as well. When simulating the execution of the plan using  $D$ , if an action's preconditions are false, we treat it as having no effects and add a  $PRECFALSE$  risk to  $RiskSet$ .

<sup>2</sup>This can be done trivially by a pre-processor.

```

FINDRISKS ( $\langle a_1, \dots, a_n \rangle, g, L$ )  $\equiv$ 
  RiskSet  $\leftarrow \emptyset$ 
  Unsupported  $\leftarrow g$ 
  for  $i = n$  to 1
    Supported  $\leftarrow$  Unsupported  $\cap$  ADDS( $a_i$ )
    Supported  $\leftarrow$  Supported  $\setminus$  PRECS( $a_i$ )
    Unsupported  $\leftarrow$  Unsupported  $\setminus$  Supported
    for all  $x$  in Unsupported
      if DoesNotMakeFalse( $a_i, x$ )  $\notin L$ 
        RiskSet  $\leftarrow$  RiskSet  $\cup$  {POSSCLOB( $a_i, x$ )}
      if CompletePreconditions( $a_i$ )  $\notin L$ 
        RiskSet  $\leftarrow$  RiskSet  $\cup$  {PRECOPEN( $a_i$ )}
    Unsupported  $\leftarrow$  Unsupported  $\cup$  PRECS( $a_i$ )
  return RiskSet

```

Figure 1: Finding risks in a provably correct plan.

that  $\text{DELS}(a_i)$  are not examined and that  $\text{PRECS}(a_i)$  are not checked to be true; however, this is not needed since we assume that the input plan is provably correct in  $D$ .

### 3.1 Analysis

We now discuss the implications of risks. We show that a plan without risks cannot fail. We then describe a tractable method for identifying a subset of risks which imply the possibility of plan failure.

Roughly speaking, we say that it is possible for a plan to fail if there exists a possible world that is consistent with the information given the planner in which  $p$  does not achieve its goals. More precisely, we mean there exists some  $D'$  such that  $D \subseteq D'$ ,  $D'$  is consistent with the locally-closed world statements  $L$ , and  $\neg \text{achieves}(D', s_0, g, c)$ .

Determining if a plan will succeed involves reasoning over all possible states that might occur in the execution of the plan, which is exponential in the length of the the plan if there is uncertainty in the outcome of the actions. Under our assumption of a correct but incomplete model, any plan that can fail in  $D_{\text{true}}$  will have at least one risk. The following theorem states that a plan that has no risks and will succeed given action model  $D$  will also succeed given any action model that is a superset of  $D$ . As a corollary to our theorem, the plan will succeed given  $D_{\text{true}}$ .

**Theorem:** if  $D \subseteq D'$  and  $\text{FINDRISKS}(c, g, L)$  contains no risks and  $\text{achieves}(D, s_0, g, c)$  then  $\text{achieves}(D', s_0, g, c)$ .

**Proof sketch:** Consider the first action  $a_1$  in  $c$ . Given that there are no  $\text{PRECFALSE}$  risks, all of  $a_1$ 's preconditions in  $D$  must be true in  $s_0$ . Since there are no  $\text{PRECOPEN}$  risks, all of  $a_1$ 's preconditions in  $D'$  must be true in  $s_0$  and thus  $a_1$  will execute successfully in  $s_0$ . A similar argument says that  $a_2$  will execute successfully as long as no effect of action  $a_1$  clobbered a precondition  $x$  of  $a_2$  that was true in  $s_0$ . However, if such a clobbering existed and  $x$  were in  $D$  then there would be a  $\text{POSSCLOB}(a_1, x)$  risk. If  $x$  were not in  $D$  then there would have to be a  $\text{PRECOPEN}(a_2)$  risk. Thus  $a_2$  will execute successfully. By induction, the entire plan will execute successfully, including the last action which can encode the goal, as described above.

A plan with risks may also succeed when executed in the world because the actions are, in fact, completely modeled or the omissions in the model do not adversely effect the plan or because the plan succeeds with all possible action descriptions that are consistent with  $D$ . We can, however, define a subset of risks which are *critical*. A *critical* risk is not guaranteed to cause plan failure, but it does guarantee the possibility of failure.

In order to identify the critical risks, we first identify a set of *vulnerable* conditions in  $p$ . Intuitively, a condition is vulnerable if has only one source of support in a plan. Formally, we define vulnerability recursively as follows. A conjunct of the goal is vulnerable in a plan iff it is established exactly once by an action in the plan or the initial state  $s_0$ . A precondition  $x$  of an action in a plan is vulnerable iff the action has an effect which establishes a vulnerable condition and  $x$  is established exactly once by either a prior action in the plan or by the initial state. The set of vulnerable conditions in a plan can be quickly computed in a simple pass backwards through the plan.

We can now define critical risks for the actions in a plan  $p = a_1, \dots, a_n$ . A  $\text{PRECOPEN}(a_i)$  risk is critical if  $a_i$  establishes a vulnerable condition. A  $\text{UNLISTEDEFFECT}(a_i, x)$  condition is critical if  $x$  is vulnerable. A  $\text{POSSCLOB}(a_i, x)$  is critical if  $x$  is vulnerable and  $a_i$  would execute successfully under the assumption that  $D$  is complete. A  $\text{PRECFALSE}(a_i, x)$  risk is critical if action  $a_i$  establishes a vulnerable condition. The critical risks can also be easily computed in a single pass through the plan.

The following theorem states that critical risks guarantee possible plan failure.

**Theorem:** If  $\text{FINDRISKS}(c, g, L)$  contains a critical risk, then there exists some  $D'$  such that  $D \subset D'$ ,  $D'$  is consistent with the LCW in  $L$ , and  $\neg \text{achieves}(D', s', g, c)$ .

**Proof sketch:** We briefly consider each type of risk. In each case, we show there exists a  $D'$  such that a vulnerable condition of the plan will not be established or will be clobbered, which causes plan failure. An action  $a_i$  with a critical  $\text{PRECOPEN}$  risk can fail because there exists an  $x$  (or  $\neg x$ ) that is a precondition of  $a_i$  in  $D'$  such that  $x$  will be false after  $a_{i-1}$  is executed. If a plan has a critical  $\text{POSSCLOB}(a_i, x)$  risk then it can fail if  $D'$  is identical to  $D$  except that action  $a_i$  has effect  $x$ , which will clobber a vulnerable condition. If  $D' = D$ , then a plan with a critical  $\text{UNLISTEDEFFECT}$  or  $\text{PRECFALSE}$  action will fail.

### 3.2 Plan Selection algorithm

We now present our solution to the plan selection problem. Below we list a variety of policies for preferring one plan to another. Each policy takes two candidate plans,  $c_1$  and  $c_2$ , and returns true if  $c_1$  is preferred to  $c_2$ . We use  $\text{risks}(c)$  as a shorthand for  $\text{FINDRISKS}(c, g, L)$ . We assume that a plan that achieves its goal with the given action descriptions is preferable to any plan that does not. Thus we assume that either both  $c_1$  and  $c_2$  achieve the goal given  $D$  or neither do.

- $RP_{\emptyset}(c_1, c_2) :: \text{risks}(c_1) = \emptyset$ , and  $\text{risks}(c_2)$  contains a critical risk.
- $RP_C(c_1, c_2) :: \text{risks}(c_1) \subset \text{risks}(c_2)$ .

- $RP_{<}(c_1, c_2) ::$  the number of each type of risk in  $risks(c_1)$  is less than that the number of that type of risk in  $risks(c_2)$ .
- $RP_w(c_1, c_2) ::$   $weighted(c_1) < weighted(c_2)$  where  $weighted(c)$  returns a real number by adding together the number of each type of risk multiplied by a pre-defined weight for that risk type.

Risks assessment can be incorporated with other methods for preferring plans. For example, most planners have an implicit preference for selecting the shortest plan that achieves the goal. These preference rules can either dominate the pre-existing preference method or be used only to break ties. Let  $TieBreak(RP, c_1, c_2)$  return true if either the pre-existing preference method prefers  $c_1$  to  $c_2$  or it ranks them equal and  $RP(c_1, c_2)$  is true.

The  $TieBreak(RP_0, c_1, c_2)$  policy is the most conservative use of the techniques discussed in this paper. It only uses our techniques in the case where two candidate plans are considered equal (ignoring the incompleteness issues) and one has no risk and the other has at least one critical risk. We present it as an extreme case in which reasoning about incompleteness clearly (if infrequently) improves plan selection. As shown above, if the action descriptions in  $D$  are correct but incomplete then a plan with critical risks might fail, while a plan without risks cannot fail. If there is no other basis for preferring one plan to another, it seems obviously beneficial to prefer the plan that cannot fail.

The other policies are more widely applicable but can prefer a plan that happens to fail to one that happens to succeed given  $D_{true}$ . This can happen with  $RP_{<}$  and  $RP_w$  if plan  $c_1$  has fewer risks than  $c_2$ , but all risks in  $c_1$  happen to correspond to real discrepancies between  $D$  and  $D_{true}$  while all the risks in  $c_2$  happen not to.

Somewhat surprisingly, however, even if  $c_1$  has a subset of the risks of  $c_2$ , it is possible that  $c_1$  will fail and  $c_2$  will succeed. One reason for this is that a  $POSSCLOB(a_i, x)$  risk can be *realized*, i.e.,  $a_i$  does have effect  $x$  in  $D_{true}$ , but then *corrected* by another action which has a  $\neg x$  effect which is also missing from  $D$ . Further, it is possible that the risk will be corrected in  $c_2$  but not in  $c_1$  if only  $c_2$  has the correcting action. If  $c_2$  has some other risks which are not realized, then it will succeed while  $c_1$  fails, even though  $c_1$  has a subset of  $c_2$ 's risks.

While the policies other than  $TieBreak(RP_0)$  can choose the inferior plan in certain cases, they are likely to improve plan selection, on average, under a reasonable set of assumptions about the distribution of planning problems. For example, for any given  $POSSCLOB$  risk, it seems much more likely that it will be realized than that it will be both realized and then corrected. If one plan has more  $POSSCLOB$  risks than another, then it is more likely to have a realized but uncorrected risk than the other, and therefore is more likely to fail.

## 4 Examples

We now describe three classes of planning problems for which it is beneficial to reason about the incompleteness of the action descriptions.

The first class of problems concerns the order of actions in a plan. Figure 2 shows a simple example designed to

illustrate how our risk analysis can prefer one ordering of plan actions over another. The figure shows a goal that can be achieved by executing two actions,  $a_1$  and  $a_2$ , in either order. If the action model is complete then both plans will achieve their goal.

Either plan can fail, however, if the model is incomplete. Candidate plan  $C_1 = [a_1, a_2]$  could fail if  $a_2$  has an effect  $\neg r$  which clobbers  $a_1$ 's effect. Similarly, plan  $C_2 = [a_2, a_1]$  could fail if  $a_1$  clobbers  $a_2$ 's effect. However,  $C_2$  could also fail if  $a_2$  has effect  $\neg p$  which would clobber  $a_1$ 's known precondition. Thus,  $C_2$  has more  $POSSCLOB$  risks than  $C_1$  and both the  $RP_{<}$  and  $RP_w$  policies would prefer  $C_1$  to  $C_2$ .

The preference for  $C_1$  over  $C_2$  is justified by imagining what we could add to  $D$  in order to create a  $D_{true}$  in which the plans would fail. For any combination of additional conditions and effects that would cause  $C_2$  to fail, there is a corresponding modification that would cause  $C_1$  to fail. However, in order to "match" the risk introduced by adding  $\neg p$  as an effect of  $C_2$ , we have to add both a precondition and an effect to  $C_1$ . Thus, in the absence of any LCW information,  $C_1$  seems the safer choice. On the other hand, if the given LCW eliminates the risks of plan  $C_2$ , then it becomes the better plan to execute.

The second class of problems we consider are ones in which our plan selection policies will sometimes prefer plans with what seem like extra steps, i.e., steps that are not required to make the plan execute successfully given  $D$ . Figure 3 shows a simple example of how this could happen. If a precondition  $x$  of an action  $a_i$  is established by the initial state or by an early action then it has many chances to be clobbered. By adding in a "cleanup" step just before  $a_i$  to reestablish  $x$ , these potential clobbering risks are removed.

init	plan $C_1$	goal	init	plan $C_2$	goal
{p}	$p \xrightarrow{a_1} r \quad \xrightarrow{a_2} q$	{r,q}	{p}	$\xrightarrow{a_2} q \quad p \xrightarrow{a_1} r$	{r,q}
<u>Risks</u>			<u>Risks</u>		
	PRECOPEN( $a_1$ )			PRECOPEN( $a_1$ )	
	PRECOPEN( $a_2$ )			PRECOPEN( $a_2$ )	
	POSSCLOB( $a_2, r$ )			POSSCLOB( $a_1, q$ )	
				POSSCLOB( $a_2, p$ )	

Figure 2: **Action order.** These two plans differ only in the order of the actions, but they have different numbers of risks.

init	plan $C_1$	goal
{p}	$\xrightarrow{a_1} q \quad q \xrightarrow{a_2} r \quad r \xrightarrow{a_3} s \quad s, p \xrightarrow{a_4} t$	{t}
<u>Risks</u>		
	PRECOPEN( $a_1$ )	PRECOPEN( $a_2$ )
	PRECOPEN( $a_3$ )	PRECOPEN( $a_4$ )
	POSSCLOB( $a_1, p$ )	POSSCLOB( $a_2, p$ )
	POSSCLOB( $a_3, p$ )	

  

init	plan $C_2$	goal
{p}	$\xrightarrow{a_1} q \quad q \xrightarrow{a_2} r \quad r \xrightarrow{a_3} s \quad \xrightarrow{a_5} p \quad s, p \xrightarrow{a_4} t$	{t}
<u>Risks</u>		
	PRECOPEN( $a_1$ )	PRECOPEN( $a_2$ )
	PRECOPEN( $a_3$ )	PRECOPEN( $a_4$ )
	PRECOPEN( $a_5$ )	

Figure 3: **Additional steps.** The extra step  $a_5$  in plan  $C_2$  reduces the number of possible clobberings of action  $a_4$ .

Run	Percentiles of the number of risks distribution										Avg # of risks	Dev. of risks	Min # of risks	Max # of risks	Number of plans
	$t_{10}$	$t_{20}$	$t_{30}$	$t_{40}$	$t_{50}$	$t_{60}$	$t_{70}$	$t_{80}$	$t_{90}$	$t_{100}$					
AO	87.6	86.2	84.7	84.0	83.1	82.3	81.4	80.3	78.9	76.7	49.4	4.3	42.0	74.0	195254
OC	85.4	83.0	81.4	80.9	79.8	79.0	77.9	77.1	76.3	74.7	72.0	7.8	55.0	115.0	341469

Table 1: Impact of risk assessment on likelihood of successfully executing plans.

<u>init</u>	<u>plan <math>C_1</math></u>	<u>goal</u>	<u>init</u>	<u>plan <math>C_2</math></u>	<u>goal</u>		
{w}	$\vec{a}_1 \vec{r}, q$	$w, q \vec{a}_2 \vec{s}$	{r,s}	{w}	$w \vec{a}_3 \vec{p}$	$p \vec{a}_4 \vec{r}, s$	{r,s}
<u>Risks</u>			<u>Risks</u>				
	PRECOPEN( $a_1$ )			PRECOPEN( $a_3$ )			
	PRECOPEN( $a_2$ )			PRECOPEN( $a_4$ )			
	POSSCLOB( $a_1, w$ )						
	POSSCLOB( $a_2, r$ )						

Figure 4: **Operator choice.** The second plan is preferred because there are possible clobberings in the first.

However, the added step may also introduce risks, such as the PRECOPEN risk in Figure 3. If POSSCLOB risks are weighted more than half the weight of PRECOPEN risks, then the  $RP_w$  will prefer plan  $C_2$  in this example. Additionally,  $C_2$  can be preferred to  $C_1$  by any of the other selection policies if LCW excludes enough risks in  $C_2$ .

The third and final class of problems arises when there are alternative actions with similar effects. The issues that arise in the ordering of actions and in adding additional steps also come up in the choice of actions to achieve the same goal or subgoal of a plan. The most obvious role of our techniques would be to prefer to use operators that are completely modeled over ones that are not. Additionally, Figure 4 shows two plans that look equally correct if the prospect of missing effects and preconditions is ignored, but one plan has more risks than the other.

## 5 Experiments

In order to measure how useful risk assessment is for plan selection, we conducted two experiments. We implemented a modified version of the Fast Forward (FF) planning system (Hoffmann & Nebel, 2001) that exploits augmented domain descriptions to find the risks in each generated plan. The domain descriptions allow some preconditions and effects to be marked as “hidden”, i.e. knowledge that is in  $D_{true}$  but not in  $D$ . Thus, our version of FF can generate plans using  $D$ , and then evaluate them using  $D_{true}$ .

Our first attempt to evaluate our plan-selection framework was to slightly modify the domain descriptions for the planning domains that FF comes with. However, these results were skewed by the fact that standard planning domains are highly crafted, and small differences between  $D$  and  $D_{true}$  led to success percentages of 100% or well under 50%. Neither case seemed to suitably represent situations where a domain model might be usable, yet incomplete.

Each of our two experiments are based on a suite of pseudo-randomly generated triples  $\langle g, D, D_{true} \rangle$ , where  $g$  is a planning goal that can be solved in  $D$  but may not be executable in  $D_{true}$ . There are precondition / effect dependencies in both  $D$  and  $D_{true}$  that reduce the number of plans produced by FF. Some additional dependencies are added to  $D_{true}$  but only partially added to  $D$  so that between 65% and 85% of the solutions in  $D$  will be executable. For these

experiments, add effects were never hidden (since then  $D$  produces no plan), but preconditions and delete effects were hidden at random. Also, each action was given, on average, the same number of preconditions to avoid introducing a bias.

For measuring the impact of risk assessment in action ordering decisions, the triples  $\langle g, D, D_{true} \rangle$  were generated so that  $g$  is always achieved by executing 7 actions exactly once. The question is which ordering of these 7 actions is most likely to be successfully executed.

For measuring the impact of risk assessment in operator selection decisions, the triples  $\langle g, D, D_{true} \rangle$  were generated so that  $g$  is always achieved by a sequence of length 10. For each position  $i$  in the plan, there are two possible choices  $a_{i,0}$  and  $a_{i,1}$ . These two actions share a common effect and a common precondition (which is the common effect of both  $a_{(i-1),0}$  and  $a_{(i-1),1}$ ). The question is which combination of choices is most likely to be successfully executed.

Table 1 gives statistics showing how risks and success percentages are related for the two experiments. The first row presents results for the action ordering (abbreviated “AO”) and the second presents results for operator choice (“OC”). Within each row, there are 10 columns that show the likelihood of successfully executing a plan drawn at random from different subsets of the set of generated plans. The final five columns show general statistics about the distribution of risks in the set of generated plans.

Within each row, the success percentages in the columns labelled  $t_{10}, \dots, t_{100}$  are derived from different subsets of the set of generated plans. For  $t_k$ , the subset contains all plans whose number of risks are in the lowest  $k$ th percentile of distribution. For example,  $t_{50}$  includes all plans with fewer risks than the median number of risks, and  $t_{100}$  includes all plans.

The results show what a significant impact risk assessment can have. For both runs, generating two plans and preferring the one with fewer risks will, on average, increase the chance of success for roughly 75% to 80%. Generating more candidates continues to provide benefits, as selecting a plan from  $t_{10}$  increase the likelihood of success to over 85%.

## 6 Related work

Much previous work has addressed the problem of planning with incomplete state information and non-deterministic or conditional effects (e.g. Kushmerick, Hanks, & Weld (1995); Smith & Weld (1998)). This is similar to the problem of planning with incomplete action models in the sense that both problems are concerned with uncertainty about the effects of actions. One important difference, however, is that non-deterministic planning systems demand even more elaborate action descriptions than traditional planners. For example, the action descriptions are required to describe all the different possible sets of effects that an action might have. In contrast, our techniques can improve planning even

without any additional information, and we provide a framework in which any additional effort to produce LCW statements can be factored into the planning process. Further, our techniques are designed to exploit various types of information, even statements about which fluents an action does *not* effect.

Additionally, the objective of work on non-deterministic planning is usually to generate plans that are guaranteed to succeed, or are guaranteed to succeed with some probability. As a result, even assessing a probabilistic or conditional plan to determine if it will succeed requires exponential computation win the length of the plan. In contrast, our methods simply prefer to execute plans with fewer risks. Further, our techniques are linear in the length of the plan and the size of the state, though we do require the planner to generate multiple plans to choose from.

Prior work has addressed the complementary problem of improving action models between planning episodes. One approach has been to develop knowledge acquisition systems that help domain experts convey their knowledge to a computer. For example, Tecuci *et al.* (Tecuci *et al.*, 1999) present techniques for producing hierarchical if-then task reduction rules by demonstration and discussion from a human expert. A second approach is to develop techniques for improving action descriptions based on the observed results of executing actions (Gil, 1994; Wang, 1995; Oates & Cohen, 1996). Our techniques are complementary since our methods are designed to improve planning when there are incomplete action descriptions. One possible synergy between these two lines of research would be to develop techniques for automatically learning LCW or helping to elicit it from the domain experts.

## 7 Conclusions and Future Work

This work is motivated by our belief that the primary obstacle to the practical application of planning technology is the difficulty of obtaining task models. This problem should be addressed both by developing better techniques for eliciting models, and by adapting planning algorithms to use models that are less difficult to elicit. In this paper, we pursue the latter approach.

Our future work includes incorporating other types of knowledge about how to do things into a planning process that reasons about incomplete action models. As mentioned above, if an important effect or precondition is missing from the action descriptions given to a planner, it can have little chance of achieving its goal. Indeed, once we forgo the completeness, the set of valid plans can no longer be deduced from the action descriptions. Thus, it becomes especially useful for human experts to provide additional information about how to accomplish goals and subgoals by, for example, providing a library of pre-defined goal-decomposition rules (also called recipes and hierarchical task networks). We are particularly interested in using LCW information to allow more flexible use of goal-decomposition knowledge.

Another line of future research is to integrate the techniques presented here into a system that integrates (re)planning and execution. The planning system will keep track of what actions it has executed and what state information it has sensed at different time points. Our techniques

for assessing risks can be applied to the execution history in order to determine which fluents it should most rely on. If the preconditions of an executed action had many risks associated with it, then the planner should be wary of using that action's presumed effects to support the preconditions of future actions.

## References

- Babaian, T., and Schmolze, J. G. 2000. PSIPlan: Open World Planning with  $\psi$ -forms. In *Proc. 5th Int. Conf. on AI Planning Systems*, 292–307.
- Etzioni, O.; Hanks, S.; Weld, D.; Draper, D.; Lesh, N.; and Williamson, M. 1992. An approach to planning with incomplete information. In *Proc. 3rd Int. Conf. on Principles of Knowledge Representation and Reasoning*, 115–125.
- Gil, Y. 1994. Learning by Experimentation: Incremental Refinement of Incomplete Planning Domains. In *Eleventh Intl Conf on Machine Learning*, 87–95.
- Golden, K. 1998. Leap before you Look: Information Gathering in the PUCINI planner. In *Proc. 4th Int. Conf. on AI Planning Systems*, 70–77.
- Hoffmann, J., and Nebel, B. 2001. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research* 14:253–302.
- Kushmerick, N.; Hanks, S.; and Weld, D. 1995. An Algorithm for Probabilistic Planning. *Artificial Intelligence* 76:239–286.
- Levesque, H. J. 1996. What is planning in the presence of sensing. In *Proc. 13th Nat. Conf. AI*, 1139–1146.
- McCarthy, J., and Hayes, P. J. 1969. Some philosophical problems from the standpoint of Artificial Intelligence. In Meltzer, B., and Michie, D., eds., *Machine Intelligence 4*. Edinburgh University Press. 463–502.
- Oates, T., and Cohen, P. 1996. Searching for planning operators with context-dependent and probabilistic effects. In *Proc. 13th Nat. Conf. AI*, 863–868.
- Reiter, R. 1991. The Frame Problem in the Situation Calculus: A Simple Solution (Sometimes) and a Completeness Result for Goal Regression. In Lifschitz, V., ed., *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*. Academic Press. 359–380.
- Smith, D., and Weld, D. 1998. Conformant Graphplan. In *Proc. 15th Nat. Conf. AI*, 889–896.
- Tecuci, G.; Boicu, M.; Wright, K.; Lee, S.; Marcu, D.; and Bowman, M. 1999. An integrated shell and methodology for rapid development of knowledge-based agents. In *Proc. 16th Nat. Conf. AI*, 250–257.
- Wang, X. 1995. Learning by observation and practice: an incremental approach for planning operator acquisition. In *Proc. 12th Int. Conf. on Machine Learning*, 549–557.