REPRESENTATION OF ACTIONS THAT HAVE SIDE-EFFECTS

N.S. Sridharan and F. Hawrusik
Rutgers University
New Brunswick, NJ 08903

Systems that reason about actions, whether they do plan generation [1] or plan recognition [2] typically model the effects of actions that occur in a plan. Simple declarative schemata allow the specification of assertions to be added/deleted to model the primary effects of actions. Side-effects of actions are those that are conditional on properties of the state in which the action is taken. When representing actions which have side-effects, conventional wisdom suggests adopting a procedural representation for they allow detailed specification of side-effects. However the procedural representation hides this knowledge from other parts of the system, thereby hindering the system in reasoning about side-effects. We use a STRIPS like declarative schema for actions that has parameters, preconditions, assertional forms for goal and outcomes and investigate three methods of representing knowledge about side-effects and discuss how the system computes side-effects without running into severe combinatorics. The following discussion and examples deal with knowledge representation as implemented in the AIMDS system which forms the AI framework for the BELIEVER project.

Examples

Consider a normal input of the form "John walked from the office to the bus station" interpreted in a world model where "John is at the office" is true. The conclusions drawn include "John is at the bus station". This can be handled using an act schema with three variables P, FL and TL as shown below.

```
(Each WALK act has
    (agent      [a PERSON [refer: P]])
    (fromloc    [a LOCATION (refer: FL]])
    (toloc      [a LOCATION [refer: TL]])
    (goal       (PROPOSITION [P loc TL]))
    (precond    (PROPOSITION [P loc FL])))
```

We have extended the interpreter to deal with some simple cases of partial act instance descriptions and incomplete world models. The incomplete description "John walked to the bus station" can be filled in using the world model, so that the system now can conclude "John walked from the office". Similarly, if in the world model "John's location is unknown" then from the normal input "John walked from the office to the bus station" the system concludes "John was at the office before walking to the bus station". Dealing with side-effects requires additional knowledge. For example, consider the world model "John is at the office; John is holding a package; The package is at the office". To update the world model properly, the location of the package must be changed to the same location as John.

We explain three methods in which this knowledge about side-effects is represented.

(a) Conditional Outcomes: We associate with the WALK schema a set of conditional outcomes which include

```
[P inhand 0] => [0 loc TL]
```

with the interpretation "when the person P walks to TL and has in hand an object 0, the object O shifts location to TL". The interpreter that updates the world model then accounts for these conditional outcomes by testing each left side and effecting the changes prescribed in the right side. If the user attempts to be as complete as possible in writing down conditional sentences, he risks the possibility that most of these would be irrelevant to any particular instantiation of an act and the possibility that he might have missed some situations, Furthermore, the system is liable to waste effort in testing a combinatorially prohibitive number of conditional sentences. If a relevant conditional sentence is missing the system has no way of recognizing this to prompt the user.

(b) Consistency conditions on the world model! The discipline of stating consistency conditions on the world model introduced in [3] provides a second representation. A causal dependency of the physical world may be captured by writing a consistency condition such as,

```
For Person P and Object 0,
[P inhand 0] »> [(loeof P) = (loeof O)]
```

which is associated with the relation "inhand". This is deceptively similar to the conditional sentence shown in (a) but an important difference is that here it is NOT knowledge associated with WALK. When "John has in hand the package" is initially put into the world model, the system checks this consistency condition and makes note of

the depdendency of this assertion on the two supporting assertions "John is at the office" and "The package is at the office". When changing John's location as a result of the Walk act, the system recognizes that the support for "John has in hand the package" is being changed. At this point the system retrieves side-effect rules for updating "inhand" when "loc" is changed. This side effect rule specifies that the location of the package changes and that the person still has the object in hand.

It is worth noting that with a clean factoring out of causal dependency knowledge, (i) the system performs reduced search as a result of having precomputed the dependency between assertions in the world model; and (ii) if a side-effect rule is missing the system prompts the user for an appropriate rule, permitting knowledge acquisition in context; and (iii) the ability to follow causal dependencies is useful [SOPHIE] in answering "What if.." questions where a counterfactual denies a fact in the world model and is not a result of any specific action that the system knows about.

(c) Frames of Reference: Although the method TH) reduces the combinatorics of simulating a single action and makes the system more robust, it requires a painstaking updating of the world model each time a similar action occurs. For example, if John walks to several places in succession holding the package we want the system to save effort in updating the location of the package for each act. We introduce the notion of a frame of reference whereby at the time John is said to have the package in hand, the location of the package is changed from (Package loc Office) to (Package loc [locof John]) where [locof John] refers to John's location in the current world model. The actual location of the package will be computed when needed from its frame of reference - the current location of John. The WALK schema given before can be used in the usual manner to update only John's location and computation of the side-effect is carried out implicitly, and no effort is spent on this while updating the world model. The procedural attachments available in the AIMDS system provide a convenient mechanism for effecting such implicit side-effects. We use one procedure attached to "inhand" to shift the Object's location to the frame of reference of the person when "A inhand 0" is asserted; and another procedure attached to "inhand" gets activated when "A inhand 0" is removed and the activated procedure removes the object from its frame of reference and asserts an explicit location for it.

Thus while updating the world model for any one of a series of Walk acts there is no explicit computation of side-effects. Yet after these location transfers, the system determines the location of the object with ease. When John drops the package into a mailbox, the second procedure asserts that the package is now in the mailbox and John can walk away freely. This method is similar to the familiar technique of retaining only "primitive" assertions in the world model and computing all "derived" assertions whenever needed. However, the method presented is selective so that derived assertions not changed by an action are retained and need not be rederived.

Significance

We have presented three methods which complement each other, each suited well to a particular task. Method (a) makes explicit the side-effects, permitting reasoninq about side-effects; method (b) allows extending the action set and facilitates the acquisition of side-effect rules; and method (c) seems well suited for actually simulating the effects of an action. A topic of further research is into means for automatic re-representation of side-effect knowledge from one form into another, especially from method (b) into method (a) and method (c).

REFERENCES

Fikes, R.E. and Nilsson, N. "STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving", Journal of Artificial Intelligence, 3(1), 27-68, 1972.

Schmidt, C.F. and Sridharan, N.S. "Plan Recognition: A Hypothesize and Revise Paradiqm", Fifth "international Joint Conference on Artificial Intelligence, MIT, August 1977.

3. Srinivasan, C.V. "The Architecture of Coherent Information System: A General Problem Solving System", IEEE Transactions on Computers, April 1976.