

PATRICK J. HAYES*

The Frame Problem and Related Problems in Artificial Intelligence

Summary

The frame problem arises in attempts to formalise problem-solving processes involving interactions with a complex world. It concerns the difficulty of keeping track of the consequences of the performance of an action in, or more generally of the making of some alteration to, a representation of the world. The paper contains a survey of the problem, showing how it arises in several contexts and relating it to some traditional problems in philosophical logic. In the second part of the paper several suggested partial solutions to the problem are outlined and compared. This comparison necessitates an analysis of what is meant by a representation of a robot's environment. Different notions of representation give rise to different proposed solutions. It is argued that a theory of causal relationships is a necessity for any general solution. The significance of this, and the problem in general, for natural (human and animal) problem solving is discussed, and several desiderata for efficient representational schemes are outlined.

Introduction

We consider some problems which arise in attempting a logical analysis of the structure of a robot's beliefs.

A *robot* is an intelligent system equipped with sensory capabilities, operating in an environment similar to the everyday world inhabited by human robots.

* University of Edinburgh

By *belief* is meant any piece of information which is explicitly stored in the robot's memory. New beliefs are formed by (at least) two distinct processes: *thinking* and *observation*. The former involves operations which are purely internal to the belief system: the latter involves interacting with the *world*, that is, the external environment and, possibly, other aspects of the robot's own structure.

Beliefs will be represented by statements in a formal logical calculus, called the *belief calculus* L_b . The process of inferring new assertions from earlier ones by the *rules of inference* of the calculus will represent thinking (McCarthy, 1959, 1963; McCarthy and Hayes, 1969; Green, 1969; Hayes, 1971).

There are convincing reasons why L_b must include L_c — classical first-order logic. It has often been assumed that a moderately adequate belief logic can be obtained merely by adding *axioms* to L_c (a first-order theory); however I believe that it will certainly be necessary to add extra rules of inference to L_c , and extra syntactic richness to handle these extra rules.

One can show that, under very general conditions, logical calculi obey the *extension property*: If $S \vdash p$ and $S \subseteq S'$ then $S' \vdash p$. The importance of this is that if a belief p is added to a set S , then all thinking which was legal before, remains legal, so that the robot need not check it all out again.

Time and Change

For him to think about the real world, the robot's beliefs must handle *time*. This has two distinct but related aspects.

(a) There must be beliefs *about* time. For example, beliefs about causality.

(b) The robot lives *in* time: the world changes about him. His beliefs must accommodate in a rational way to this change.

Of these, the first has been very extensively investigated both in A.I. and philosophical logic, while the second has been largely ignored until very recently: it is more difficult. The first is solely concerned with thinking: the second involves observation.

The standard device for dealing with (a) is the introduction of *situation variables* (McCarthy, 1963; McCarthy and Hayes, 1969) or *possible*

worlds (Hintikka, 1967; Kripke, 1963). Symbols prone to change their denotations with the passage of time are enriched with an extra argument place which is filled with a term (often a variable) denoting a *situation* which one can think of intuitively as a time instant; although other readings are possible. In order to make statements about the relationships between situations, and the effects of actions, we also introduce terms denoting *events*, and the function R (read: *result*) which takes events and situations into new situations. Intuitively, " $R(e,s)$ " denotes the situation which results when the event e happens in the situation s . By "event" we mean a change in the world: "his switching on the light", "the explosion", "the death of Caesar". This is a minor technical simplification of the notation and terminology used in McCarthy and Hayes (1969) and Hayes (1971). Notice that all the machinery is defined within L_c . The situation calculus is a first-order theory.

Using situations, fairly useful axiomatisations can be obtained for a number of simple problems involving sequences of actions and events in fairly complicated worlds (Green, 1969; McCarthy and Hayes, 1969).

The Frame Problem

Given a certain description of a situation s – that is, a collection of statements of the form $\phi \llbracket s \rrbracket$, where the fancy brackets mean that *every* situation in ϕ is an occurrence of 's' – we want to be able to infer as much as possible about $R(e,s)$. Of course, what we can infer will depend upon the properties of e . Thus we require assertions of the form:

$$\phi_1 \llbracket s \rrbracket \ \& \ \psi(e) \supset \phi_2 \llbracket R(e,s) \rrbracket \quad (1)$$

Such an assertion will be called a *law of motion*. The frame problem can be briefly stated as the problem of finding adequate collections of laws of motion.

Notice how easily human thinking seems to be able to handle such inferences. Suppose I am describing to a child how to build towers of bricks. I say "You can put the brick on top of this one onto some other one, if that one has not got anything else on it." The child *knows* that the other blocks will stay put during the move. But if I write the

corresponding law of motion:

$$(\text{on}(b_1, b_2, s) \ \& \ \forall z. \neg \text{on}(z, b_3, s)) \supset \text{on}(b_1, b_3, R(\text{move}(b_2, b_3), s)) \quad (2)$$

then nothing follows concerning the other blocks. What assertions could we write down which would capture the knowledge that the child has about the world?

One does not want to be obliged to give a law of motion for *every* aspect of the new situation. For instance, one feels that it is prolix to have a law of motion to the effect that if a block is not *moved*, then it stays where it is. And yet such laws – instances of (1) in which $\phi_1 = \phi_2$ – are necessary in first-order axiomatisations. They are called *frame axioms*. Their only function is to allow the robot to infer that an event does *not* affect an assertion. Such inferences are necessary: but one feels that they should follow from more general considerations than a case-by-case listing of axioms, especially as the number of frame axioms increases rapidly with the complexity of the problem. Raphael (1971) describes the difficulty thoroughly.

This phenomenon is to be expected. Logically, s and $R(e,s)$ are simply different entities. There is no *a priori* justification for inferring any properties of $R(e,s)$ from those of s . If it were usually the case that events made widespread and drastic alterations to the world (explosions, the Second Coming, etc.), then we could hardly expect anything better than the use of frame axioms to describe in detail, for each event, exactly what changes it brings about. Our expectation of a more general solution is based on the fact that the world is, fortunately for robots, fairly stable. Most events – especially those which are likely to be considered in planning – make only small local changes in the world, and are not expected to touch off long chains of cause and effect.

Frame Rules

We introduce some formalism in order to unify the subsequent discussions. Any general solution to the frame problem will be a method for allowing us to transfer properties from a situation s to its successor $R(e,s)$; and we expect such a licence to be sensitive to the form of the assertion, to what is known about the event e , and possibly to other facts.

Consider the rule scheme FR:

$$\chi, \phi[s], \psi(e) \vdash \phi[R(e, s)] \quad \text{(FR)}$$

provided $\aleph(e, \phi, \psi)$.

where \aleph is some condition on e, ϕ and ψ , expressed of course in the metalanguage. We will call such a rule a *frame rule*. The hope is that frame rules can be used to give a general mechanism for replacing the frame axioms, and also admit an efficient implementation, avoiding the search and relevancy problems which plague systems using axioms (Green, 1969; Raphael, 1971).

One must, when considering a frame rule, be cautious that it does not allow contradictions to be generated. Any addition of an inference rule to L_c , especially if not accompanied by extra syntax, brings the risk of inconsistency, and will, in any case, have dramatic effects on the metatheory of the calculus. For instance, the deduction theorem fails. Thus a careful investigation of each case is needed. In some cases, a frame rule has a sufficiently simple \aleph condition that it may be replaced by an *axiom scheme*, resulting in a more powerful logic in which the deduction theorem holds. This usually makes the metatheory easier and implementation more difficult.

Some Partial Solutions Using Frame Rules

The literature contains at least four suggestions for handling the problem which are describable by frame rules. In each case we need some extra syntactic machinery.

Frames

Following McCarthy and Hayes (1969), one assumes a finite number of monadic second-order predicates P_i . If $\vdash P_i(h)$ for a non-logical symbol h (predicate, function or individual constant) then we say that h is in the i th *block* of the frame. The frame rule is

$$P_{i_1}(h_1), \dots, P_{i_n}(h_n), \phi[s], P_j(e) \vdash \phi[R(e, s)] \quad (3)$$

where h_1, \dots, h_n are all the non logical symbols which occur *crucially* in ϕ , and $i_k \neq j, 1 \leq k \leq n$. Here *crucial* is some syntactic relation between

h and ϕ ; different relations give different logics, with a stronger or weaker frame rule.

Causal connection

We assume (Hayes 1971) that there is a 3-place predicate $\rightarrow(x, y, s)$ (read: x is connected to y in situation s) which has the intuitive meaning that if x is not connected to y , then any change to y does not affect x . It seems reasonable that \rightarrow should be a partial ordering on its first two arguments (reflexive and transitive). The frame rule is:

$$\phi[s], \neg \rightarrow(h_1, e, s), \dots, \neg \rightarrow(h_n, e, s) \vdash \phi[R(e, s)] \quad (4)$$

where (i) ϕ is an atom or the negation of an atom; (ii) h_1, \dots, h_n are all the terms which occur *crucially* in ϕ .

If we insisted only that $\neg \rightarrow(h_i, e, s)$ is not provable (rather than $\neg \rightarrow(h_i, e, s)$ is provable) then the rule is much stronger but no longer obeys the extension property. This is analogous to PLANNER's method below.

MICRO-PLANNER

The problem solving language MICRO-PLANNER (Sussman and Winograd, 1969) uses a subset of predicate calculus enriched with notations which control the system's search for proofs. We will ignore the latter aspect for the present and describe the underlying formalism. Its chief peculiarity is that it has no negation, and is therefore not troubled by the need for consistency.

Following MICRO-PLANNER we introduce the new unary propositional connective *therase*. Intuitively, *therase* ϕ will mean that ϕ is "erased". We also introduce the notion of a *transition*: an expression $\langle e: \phi_1, \dots, \phi_n \rangle$. This means intuitively "erase ϕ_1, \dots, ϕ_n in passing from s to $R(e, s)$ ". The frame rule is:

$$\chi, \phi[s], \langle e: \phi_1, \dots, \phi_n \rangle \vdash \phi[R(e, s)] \quad (5)$$

where (i) ϕ is an atom; (ii) ϕ contains no variables (other than s); (iii) χ , *therase* ϕ_1, \dots , *therase* $\phi_n \nmid$ *therase* $\phi[s]$. Notice the negated inference in (iii).

STRIPS

The problem-solving system STRIPS (Fikes and Nilsson, 1971) uses the full predicate calculus enriched with special notations ("operator descriptions") describing events, and ways of declaring certain predicates to be *primitive*. We can use transitions to describe this also. The frame rule is:

$$\phi[s], \langle e: \phi_1, \dots, \phi_n \rangle \vdash \phi[R(e, s)] \quad (6)$$

where (i) ϕ is an atom or the negation of an atom; (ii) ϕ contains no variables (other than s); (iii) the predicate symbol in ϕ is *primitive*; (iv) $\phi[s]$ is not an instance of any ϕ_i , $1 \leq i \leq n$. Notice the similarity to (5). *Primitive* can be axiomatised by the use of a monadic second-order predicate, as in (1) above.

These four rules have widely divergent logical properties. Rule (3) is replaceable by an axiom scheme, and is thus rather elementary. It is also very easy to implement efficiently (theorem-proving cognoscenti may be worried by the higher-order expressions, but these are harmless since they contain no variables). Variations are possible, e.g., we might have disjointness axioms for the P_i and require $\neg P_j(h_k)$ rather than $P_{ik}(h_k)$: this would be closely similar to a special case of (4).

Retaining consistency in the presence of (3) requires in non-trivial problems that the P_i classification be rather coarse. (For instance, *no* change in position *ever* affects the colour of things, so predicates of location *could* be classed apart from predicates of colour.) Thus frames, although useful, do not completely solve the problem.

Rule (4) is also replaceable by an axiom scheme, and the restriction to literals can be eliminated, with some resultant complication in the rule. Also, there is a corresponding model theory and a completeness result (Hayes, 1971), so that one can gain an intuition as to what (4) *means*. Retaining consistency with (4) requires some care in making logical definitions.

Rules (5) and (6) have a different character. Notice that (6) is almost a special case of (5): that in which *therase* $\phi \vdash$ *therase* ψ if ψ is not primitive or ψ is an instance of ϕ . The importance of this is that instantiation, and probably primitiveness also, are *decidable*, and conditions (iii) and (iv) in (6) are effectively determined solely by examining the transition, whereas condition (iii) in (5) is in general not decida-

ble and in any case requires an examination of all of χ : in applications, the whole set of beliefs. MICRO-PLANNER uses its ability to control the theorem-proving process to partly compensate for both of these problems, but with a more expressive language they would become harder to handle. Notice also that (5) does not satisfy the extension property, while (6) does, provided we allow at most one transition to be unconditionally asserted for each event.

Maintaining "consistency" with (5) is a matter of the axiom-writer's art. There seem to be no general guidelines. Maintaining consistency with (6) seems to be largely a matter of judicious choice of *primitive* vocabulary. There is no articulated model theory underlying (5) or (6). They are regarded more as syntactic tools — analogous to evaluation rules for a high-level programming language — than as descriptive assertions.

A (Very) Simple Example: Toy Bricks

$$\neg \text{above}(x, x, s) \quad (\text{A1})$$

$$x = \text{Table} \vee \text{above}(x, \text{Table}, s) \quad (\text{A2})$$

$$\text{above}(x, y, s) \equiv \text{on}(x, y, s) \vee \exists z. \text{on}(z, y, s) \ \& \ \text{above}(x, z, s) \quad (\text{A3})$$

$$\text{free}(x, s) \equiv \forall y \neg \text{on}(y, x, s) \quad (\text{A4})$$

To enable activity to occur we will have events *move*(x, y): the brick x is put on top of the brick y . Laws of motion we might consider include:

$$\text{free}(x, s) \ \& \ x \neq y \ \supset \ \text{on}(x, y, R(\text{move}(x, y), s)) \quad (\text{A5})$$

$$\text{free}(x, s) \ \& \ w \neq x \ \& \ \text{on}(w, z, s) \ \supset \ \text{on}(w, z, R(\text{move}(x, y), s)) \quad (\text{A6})$$

$$\text{free}(x, s) \ \& \ w \neq x \ \& \ \text{above}(w, z, s) \ \supset \ \text{above}(w, z, R(\text{move}(x, y), s)) \quad (\text{A7})$$

$$\text{free}(x, s) \ \& \ w \neq y \ \& \ \text{free}(w, s) \ \supset \ \text{free}(w, R(\text{move}(x, y), s)) \quad (\text{A8})$$

Of these, (A6–A8) are frame axioms. (In fact, (A7) and (A8) are redundant, since they can, with some difficulty, be derived from (A6) and (A3), (A4) respectively.) (A5) assumes somewhat idealistically that there is always enough space on y to put a new brick.

Rule (3) cannot be used in any intuitively satisfactory way to replace A6–A8.

Rule (4) can be used. We need only to specify when bricks are connected to events:

$$\rightarrow (x, \text{move}(y, z), s) \equiv x = y \vee \text{above}(x, y, s) \quad (\text{A9})$$

Using (A9) and (A3), (A4), it is not hard to show that

$$\begin{aligned} & \text{free}(x, s) \ \& \ w \neq x \ \& \ \text{on}(w, z, s). \supset. \neg \rightarrow (w, \text{move}(x, y), s) \ \& \\ & \neg \rightarrow (z, \text{move}(x, y), s) \end{aligned}$$

and thus, we can infer $\text{on} \{w, z, R[\text{move}(x, y), s]\}$ by rule (7). (A7) and (A8) are similar but simpler. (One should remark also that (A4) is an example of an illegal definition, in the presence of (4), since it suppresses a variable which the rule needs to be aware of. It is easy to fix this up in various ways.)

Rule (5) can also be used, but we must ensure that *therase* does a sufficiently thorough job. Various approaches are possible. The following seems to be most in the spirit of MICRO-PLANNER. In its terms, *on* and *above* statements will be in the data-base, but *free* statements will not. The necessary axioms will be:

$$\text{therase free}(x, s) \quad (\text{A10})$$

$$\text{therase on}(x, y, s) \ \& \ \text{above}(y, z, s) \supset \text{therase above}(x, z, s) \quad (\text{A11})$$

$$\text{free}(x, s) \supset (\text{move}(x, y): \text{on}(x, z, s)) \quad (\text{A12})$$

To infer statements $\text{free}[x, R(e, s)]$, we must first generate enough $\text{on}[x, y, R(e, s)]$ statements by rule (5), and then use (A4), since by (A10), rule (5) never makes such an inference directly. (We could omit (A10) and replace (A12) by

$$\text{free}(x, s) \supset (\text{move}(x, y): \text{on}(x, z, s), \text{free}(y, s)) \quad (\text{A13})$$

This would, in MICRO-PLANNER terms, be a decision to keep *free* assertions in the data base.)

Notice that MICRO-PLANNER has no negation and hence no need to *therase* such assertions as $\neg \text{on}(x, y, s)$. If it had negation we would replace (A12) by

$$\text{free}(x, s) \supset (\text{move}(x, y): \text{on}(x, z, s), \neg \text{on}(x, y, s)) \quad (\text{A14})$$

and add

$$\text{therase } \neg \text{on}(x, y, s) \ \& \ \text{above}(y, z, s) \supset \text{therase } \neg \text{above}(x, z, s) \quad (\text{A15})$$

Notice the close relations between (A3), (A11) and (A15).

Rule (6) can be used similarly to (5), but we are no longer able to use axioms such as (A11) and (A15). The solution which seems closest in spirit to STRIPS is to declare that *on* is primitive but that *above* and *free* are not, and then simply use (A14). The "world model" (Fikes and Nilsson, 1971) would then consist of a collection of atoms $\text{on}(a, b)$, or their negations, and the system would rederive *above* and *free* assertions when needed. This is very similar to MICRO-PLANNER's "data-base", and we could have used rule (5) in an exactly similar fashion.

Implementing Frame Rules

Some ingenuity with list structures enables one to store assertions in such a way that

- (i) Given s , one can easily find all assertions $\phi[s]$.
- (ii) Each symbol denoting a situation is stored only once.
- (iii) The relationships between s and $R(e, s)$, etc., are stored efficiently and are easily retrieved.
- (iv) To apply a frame rule to s , one need only:
 - (a) Create a new cell pointing to s .
 - (b) Move two pointers.
 - (c) Check each $\phi[s]$ for condition \aleph : if it holds, move one pointer.

In the case of a rule like (5) or the variation to (4), where \aleph is a *negative* condition (\neq), we need only examine those $\phi[s]$ for which the condition *fails*, resulting in greater savings.

Space does not permit a description of the method, but MICRO-PLANNER and STRIPS use related ideas. (The authors of these systems seem to confuse to some extent their particular implementations with the logical description of the frame rules, even to the extent of claiming that a logical description is impossible.)

Consistency and Counterfactuals

Frame rules can be efficiently implemented and, in their various ways, allow the replacement of frame axioms by more systematic ma-

chinery. But there is a constant danger, in constructing larger axiomatizations, of introducing inconsistency. An alternative approach avoids this by transferring properties ϕ from s to $R(e,s)$ as long as it is consistent to do so, rather than according to some fixed-in-advance rule.

Suppose we have a set χ of general laws which are to hold in every situation, and a description of – a set of assertions about – the situation s : $\{\phi_1[s], \dots, \phi_n[s]\}$. Using laws of motion we will directly infer certain properties ψ_1, \dots, ψ_m of $R(e,s)$: the set of these constitutes a partial description of $R(e,s)$. To compute a more adequate one, we add assertions $\phi_i[R(e,s)]$ in some order, *checking at each stage for consistency with χ* ; if a $\phi_i[R(e,s)]$ makes the set inconsistent, it is rejected. This continues until no more ϕ_i can be added. In this way we compute a maximal consistent subset (MCS) of the inconsistent set

$$\chi \cup \{\psi_1, \dots, \psi_m, \phi_1[R(e,s)], \dots, \phi_n[R(e,s)]\}.$$

There are two big problems: (1) Consistency is not a decidable or even semi-decidable property. Thus for practicality one has to accept a large restriction on the expressive power of the language. (2) There are in general many different MCSs of an inconsistent set, and so we must have ways of choosing an appropriate one. In terms of the procedure outlined above, we need a good ordering on the ϕ_i .

This procedure is closely similar to one described by Rescher (1964) to provide an analysis of counterfactual reasonings (“If I had struck this match yesterday, it would have lit”, when in fact I didn’t.). Rescher is aware of the first problem but gives no solution. His major contribution is to the second problem, which he solves by the use of *modal categories*: a hierarchical classification of assertions into grades of law-likeness. One never adds $\phi_i[R(e,s)]$ unless every ϕ_j with a lower classification has already been tested. This machinery is especially interesting as in (Simon and Rescher, 1966) it is linked to Simon’s theory of causality (Simon, 1953). One puts ϕ_i in a lower category than ϕ_j just in case ϕ_i causes ϕ_j (or $\neg\phi_j$), more or less. Space does not permit a complete description of this interesting material which is fully covered in the references cited. In spite of its appeal, the first problem is still unsolved.

In unpublished work at Stanford, Jack Buchanan has independently worked out another version of the procedure. The first problem is handled by accepting a drastic restriction on the language. Every ϕ_i is an atom or the negation of an atom – c.f. frame rules (7), (8) and (9) –

and, more seriously, χ contains only assertions of the form $t_1 \neq t_2$ or of the form $P(t_1, \dots, t, \dots, t_n)$ and $P(t_1, \dots, u, \dots, t_n) \supset t = u$. Under these constraints, consistency is decidable and can even be computed quite efficiently. Moreover, MCSs are unique, so the second problem evaporates. However, it is not clear whether non-trivial problems can be reasonably stated in such a restricted vocabulary.

Conclusions

In the long run, I believe that a mixture of frame rules and consistency-based methods will be required for non-trivial problems, corresponding respectively to the “strategic” and “tactical” aspects of computing descriptions of new situations. In the short term we need to know more about the properties of both procedures.

One outstanding defect of present approaches is the lack of a clear model theory. Formal systems for handling the frame problem are beginning to proliferate, but a clear *semantic* theory is far from sight. Even to begin such a project would seem to require deep insight into our presystematic intuitions about the physical world.

Observations and the Qualification Problem

We have so far been entirely concerned with thinking. The situation calculus is a belief calculus for beliefs *about* time. Observations – interactions with the real world – introduce new problems. We must now consider the second aspect of time (b.p.).

Almost any general belief about the result of his own actions may be contradicted by the robot’s observations. He may conclude that he can drive to the airport; only to find a flat tire. A human immediately says, “Ah, now I cannot go”. Simply *adding* a new belief (“the tire is flat”) renders an earlier conclusion false, though it was a valid conclusion from the earlier set of beliefs, *all of which are still present*. Thus we do *not* assume that the robot had concluded “If my tires are OK, then I can get to the airport” since there are no end of different things which might go wrong, and he cannot be expected to hedge his conclusions round with thousands of qualifications (McCarthy and Hayes, 1969).

Clearly this implies that the belief logic does not obey the extension property *for observations*: to expect otherwise would be to hope for omnipotence. However, we are little nearer any positive ideas for handling the inferences correctly.

John McCarthy recently pointed out to me that MICRO-PLANNER has a facility (called THNOT) which apparently solves the problem nicely. I will translate this into a slightly different notation.

We introduce a new unary propositional connective *proved*, which is supposed to mean "can be proved from the current collection of beliefs". Then we can write axioms like the following:

$$\text{flat (tire)} \supset \text{kaput (car)}$$

$$\neg \text{proved kaput (car)} \supset \text{at (robot, airport, R[drive(airport), s])} \quad (\text{A17})$$

from which *at(robot, airport, ...)* should be concluded *until* we add:

$$\text{flat (tire)} \quad (\text{A18})$$

at which point the $\neg \text{proved} \dots$ becomes false. ($\neg \text{proved}$ is PLANNER's THNOT).

To make this work we could try the following rules of inference.

$$\phi \vdash \text{proved } \phi \quad (\text{P1})$$

$$\chi \vdash \neg \text{proved } \phi \quad (\text{P2})$$

where $\chi \not\vdash \phi$.

(P2) fails the extension property, as expected. (It also has the difficulties of effectiveness which worry frame rule (5), but we will ignore these.)

Unfortunately, (P1) and (P2) are *inconsistent*. Suppose $\chi \not\vdash \phi$, but that ϕ is consistent with χ . Then by (P2), $\neg \text{proved } \phi$. But if we now add ϕ (an observation: the flat tire), then by (P1) *proved* ϕ : an overt contradiction. MICRO-PLANNER avoids this by denying (P1) and treating " ϕ and $\neg \text{proved } \phi$ " as consistent. But this is a counsel of despair, since it clearly is not, according to the intuitive meanings.

The logical answer is to somehow make *proved* refer to the set χ of antecedents. The direct approach to this requires extremely cumbersome notation and a very strong logic which partly contains its own metatheory, thus coming close to Gödel inconsistency. Fortunately we do not need to *describe* sets χ of assertions, but only to *refer* to them,

and this can be done with a very weak notation, similar to situation variables.

Assume that every belief is decorated with a constant symbol called the *index*: we will write it as a superscript. Indices denote the robot's internal belief states just as situation terms denote external situations. Observations are analogous to events. Assertions *proved* ϕ now have an extra index which identifies the state of belief at the time the inference was tested. The above rules of inference become:

$$\phi^s \vdash \text{proved}^s \phi^s \quad (\text{P1}')$$

$$\chi \vdash \neg \text{proved}^s \phi^s \quad (\text{P2}')$$

where $\chi \not\vdash \phi^s$ and every member of χ has index s .

In applications we now insist that:

(i) in applying P2', χ contains *all* beliefs with index s ;

(ii) whenever an *observation* is added to the beliefs, every index s is replaced by a new one s' , *except* those on *proved* assertions.

This is just enough to avoid inconsistency; it clearly does not involve any Gödel-ish difficulties; and (ii) can be very efficiently implemented by frame-rule methods (see section Implementing Frame Rules). Indeed, more complex versions of (ii) which allow for direct contradiction between beliefs and observations can be similarly implemented.

The logic of these indices is trivial, but extensions have some interest. For instance, if we identify indices with situation terms, then expressions of the form $\phi[s]^s$ become legal, with the intuitive meaning " ϕ is true *now*".

Seen this way, the qualification problem is closely linked with the frame problem, and one expects progress in either area to help with the other.

Acknowledgements

This work was supported in part by the Advanced Research Projects Agency of the Office of the Secretary of Defense (SD-183), and in part by the Science Research Council.

I am also grateful for comment, criticism and contributions from Jack Buchanan, Richard Fikes, Malcolm Newey, Nils Nilsson, John

Rulifson, Richard Waldinger, Richard Weyrauch and from John McCarthy, to whom I am also grateful for the invitation to visit the Stanford Artificial Intelligence Project, where this paper was written. Most of all, I thank my wife, Jackie, for improving my English; controlling my verbosity, and typing innumerable drafts of the manuscript.