Minimal Change and Maximal Coherence: A Basis for Belief Revision and Reasoning about Actions

Anand S. Rao
Australian AI Institute
Carlton, Vic-3053
Australia
Email: anand@aaii.oz.au

Norman Y. Foo
University of Sydney
Sydney, NSW-2006
Australia
Email: norman@basser.oz.au

Abstract

The study of *belief revision* and *reasoning about actions* have been two of the most active areas of research in AI. Both these areas involve reasoning about *change.* However very little work has been done in analyzing the principles common to both these areas. This paper presents a formal characterization of belief revision, based on the principles of *minimal change* and *maximal coherence.* This formal theory is then used to reason about actions. The resulting theory provides an elegant solution to the conceptual frame and ramification problems. It also facilitates reasoning in dynamic situations where the world changes during the execution of an action. The principles of minimal change and maximal coherence seem to unify belief revision and reasoning about actions and may form a fundamental core for reasoning about other dynamic processes that involve change.

1 Introduction

Belief revision is the process by which an agent revises her set of beliefs at the current instant of time, based on some input from the external world, to move into the next instant of time, possibly with a different set of beliefs. Thus the essential problem in belief revision is determining what beliefs the agent must acquire and what beliefs the agent must give up.

Reasoning about actions usually involves determining the state of the world after an action has been performed. Again this involves determining what beliefs *should change* and what beliefs *should not change* when moving from one state to another. The latter problem has long been known as the *frame problem* [McCarthy and Hayes, 1969]. Associated with this is the problem of specifying all the side-effects or ramifications of an action and making sure that the agent acquires (or gives up) beliefs about those side-effects as well.

Both belief revision and reasoning about actions, involve reasoning about change. However very little work has been done in analyzing the principles common to both these areas. Work on belief revision has for the most part concentrated on building efficient systems that perform belief revision [Doyle, 1979, de Kleer, 1986]. Work on reasoning about actions has addressed both the computational and foundational aspects [Brown, 1987, Georgeff and Lansky, 1987], but has been studied independently of belief revision. In this paper, we present a unified picture of both areas by formalizing the underlying principles of belief revision and reasoning about actions.

Minimal change and *maximal coherence* are two of the most important principles involved in reasoning about change. Minimization has been used in AI and philosophical logic in a variety of ways. In this paper, by *minimizing change* we mean *minimizing the acquisition of beliefs* or *minimizing the loss of beliefs* when an agent moves from one state to another. By *maximizing coherence* we mean retaining as many coherent beliefs as possible during a state change.

Theories of belief revision can be divided into two broad categories - the *coherence* and *foundational* theories of belief revision. A detailed analysis of both these theories has been carried out elsewhere [Rao and Foo, 1989]. In this paper we give a brief overview of the coherence theory of belief revision, which is based on the principles of minimal change and maximal coherence. The axiomatization of the theory is based on the work by Alchourron, Gardenfors and Makinson [1985] (henceforth called the AGM-theory) and the semantics is based on the possible-worlds formulation. This theory of belief revision is then used to describe a theory of actions which has a clear semantics and also solves the conceptual frame and ramification problems. [1] One of the advantages of having a uniform framework for belief revision and reasoning about actions, is that it facilitates reasoning in dynamic situations, where the world changes during the execution of an action. As the paper is part of the first author's thesis, in-

No claims are being made regarding the computational efficiency of our solution. We address only the conceptual problem in this paper.

terested readers are referred to the complete report [Rao, 1989] for more details.

2 Coherence Theory of Belief Revision

The underlying logic of belief revision is a modal logic of time-dependent beliefs. The formula BEL(a, t, ϕ) denotes that the agent a believes ϕ at time t. A branching time temporal logic is used with a standard first-order interpretation. Quantifying individual constants into the scope of modal operators is not allowed, but is allowed for time constants. The semantics for such a belief system can be given in terms of a slightly modified Kripke interpretation. The axiomatization is a KD45 axiomatization together with the axioms of transitivity and backwards linearity for time [Rao, 1989].

Further, we shall treat the BEL operator as a self-belief operator as in Autoepistemic Logic (AEL) [Konolige, 1988, Moore, 1985]. This interpretation will be called an *autoepistemic Kripke interpretation* denoted by AKI. The attractive feature of such a time-dependent AE belief system is that the stable sets or possible worlds are uniquely determined by the first-order formulas and universal AE (UAE) beliefs (formulas of the form $\forall$t BEL(a, t, ϕ)) that they contain. This feature will be exploited in discussing the dynamics of the belief system.

There are three possible states (ignoring the inconsistent state) in which an agent could be at some time t and world w with respect to a formula ϕ received by the agent from the external world. These states are: (1) the agent believes in ϕ at t, (2) the agent believes in $\neg\phi$ at t, and (3) the agent does not believe in ϕ nor $\neg\phi$ at t and hence the agent is *uncommitted* about ϕ. By *dynamics* of belief systems or *belief revision* we mean the process by which an agent moves from one of the three states (1, 2 or 3) to any of the other two. Ignoring the trivial transitions of remaining in the same state we are left with six different transitions: (1-2) **(N-)Expansion:** Uncommitted state $\rightarrow$ Belief in ($\neg$) ϕ, (3-4) **(N-)Contraction:** Belief in ($\neg$) ϕ $\rightarrow$ Uncommitted state, (5-6) **(N-)Revision:** [2] Belief in (ϕ) $\neg\phi$ $\rightarrow$ Belief in ($\neg\phi$) ϕ. The terminology and approach is an extension of the AGM-theory.

Three *dynamic modal operators*, EXP, CON, and REV, are introduced to denote the above transitions. Thus EXP(a, t, ϕ, u) is read as 'the expansion by the agent a at time t with respect to ϕ results in time u', where ϕ is a first-order sentence and t and u are time points. In fact, u is the next time point to t. Note that there can be more than one next instant, as the underlying temporal logic is a branching time temporal logic. This makes the logic more expressive than situation calculus [McCarthy and Hayes, 1969] and is equivalent in expressive power to the extended situation calculus discussed by GeorgefF [1987].

The semantics of dynamic operators is based on *selection functions,* which select some possible worlds as being *closer* to the current world than the others. When the agent performs expansion or contraction, he is said to move into one of these closer worlds and designate these worlds as the worlds of the next time instant(s). The interpretation for this language and the semantics of the dynamic operations are formally defined as follows:

Definition: The *dynamic coherence interpretation* is a tuple, CI = <($\mathcal{E}$, C, AKI), where $\mathcal{E}$ and $\mathcal{C}$ are expansion and contraction functions that map W (worlds), AC (agent constants), TD (time points), and 2^W (set of worlds) to a $2^{T£)}$ (set of time points) and AKI is an autoepistemic Kripke interpretation.

Definition: A dynamic coherence interpretation CI, satisfies a well-formed formula ϕ at world w and time point [3] t (written as CI, w, $\models \phi$)v e n the following conditions,

1. CI, w, t $\models$ EXP(a, t, ϕ, u) iff u $\in \mathcal{E}$(w, a, t, $||\phi||^{CI}$), where ϕ is a first-order sentence.
2. CI, w, t $\models$ CON(a, t, ϕ, u) iff u $\in \mathcal{C}$(w, a, t, $||\phi||^{CI}$), where ϕ is a first-order sentence.
3. CI, w, t $\models$ REV(a, t, ϕ, u) iff CI, w, t $\models$ CON(a, t, $\neg\phi$, v) and CI, w, v $\models$ EXP(a, v, ϕ, u), where ϕ is a first-order sentence.

The notation $||\phi||^{CI}$ stands for all the worlds of the interpretation CI, which satisfy ϕ. More formally, $||\phi||^{CI}$ = {w | CI, w, t $\models \phi$ for all t $\in$ TD }. The unique properties of AEL allows us to restrict our attention to first-order sentences rather than arbitrary belief formulas [Rao, 1989].

Axioms (AE1) - (AE5) and (AC1) - (AC4) of Appendix -I capture the proof-theoretic notion of *minimal change* and *maximal coherence* for expansion and contraction, respectively. The semantic conditions analogous to these axioms capture the semantic notion of *closest* possible worlds [Rao, 1989]. The description of these axioms and semantic conditions and a constructive procedure for $\mathcal{E}$ and $\mathcal{C}$ can be found elsewhere [Rao and Foo, 1989, Rao, 1989]. The class of models whose $\mathcal{E}$ ($\mathcal{C}$) selection function satisfies these semantic conditions are called $\mathcal{E}$-models ($\mathcal{C}$-models). We shall refer to the modal system KD45 together with the above axioms and inference rules as the *Coherence Modal System* or CS-modal system. We define the class of coherence $\mathcal{D}$-models as models whose $\mathcal{B}$ relation is a $\mathcal{B}$-model (i.e. $\mathcal{B}$ is serial, transitive and euclidean), $\mathcal{E}$ is a $\mathcal{E}$-model and $\mathcal{C}$ is a $\mathcal{C}$-model. The proof of the following theorem, which provides a sound and complete theory of belief revision, is given in the complete report [Rao, 1989].

[2] The word "belief revision" will be used in the more general sense to denote dynamics of belief systems and the word "revision" will be used in the more restricted sense as defined above.

[3] Temporal variables and constants are mapped to time points by a standard temporal assignment [Rao, 1989].

Theorem: The coherence modal system or the CS-modal system is sound and complete with respect to the class of coherence $\mathcal{D}$-models.

3 Reasoning about Actions

In this section we shall treat actions as a special type of dynamic operation and define it in terms of the other dynamic operators, EXP, CON and REV. Thus reasoning about actions is also based on the intuitive principles of minimal change and maximal coherence. As expansions and contractions can be viewed as the semantic counterparts of the *addlist* and *deletelist* of STRIPS [Fikes and Nilsson, 1971], the logic described here can be viewed as a semantic theory for STRIPS and the principles of minimal change and maximal coherence as a generalization of the STRIPS assumption.

In addition to predicate letters the logic of actions consists of *action letters,* which are used to form *action formulas.* The symbols ';' and '||' denote *sequential and parallel* actions. The modal operator ACT is used to denote the performance of an action. If the modal ACT formula is satisfiable/provable then the action is said to be successful; otherwise the action is said to have failed. This allows reasoning about both successful and failed actions which is crucial for reasoning about real world domains [Georgeff, 1987]. A *well-formed action formula* can either be a simple action formula or a sequential or parallel action formula. A simple action formula is denoted by $z(s_1, ...s_n)^a$ where z is an action letter, $s_1...s_n$ are first-order terms and a is the agent who is performing the action. The modal formula ACT(t, ψ, u) is read as 'at time t, action formula ψ is carried out to reach time point u'. The action formula ψ can be a combination of sequential or parallel action formulas.

Actions have usually been defined in terms *ot preconditions,* which must hold before the action is performed, and the *consequents,* which must hold after the action has been performed. This approach has been used in situation calculus [McCarthy and Hayes, 1969] and STRIPS [Fikes and Nilsson, 1971]. Thus, associated with every simple action formula *ifr,* are two formulas *p* and *c,* which are respectively the precondition and consequent of . Given the current world, a time point and a simple action formula, the action function *A* selects the closest possible world(s) at the next time point(s), that satisfy the following conditions:

- the precondition formula of the action is believed in the current world
- the consequent of the action must be believed at the next time instant, *with minimal change and maximal coherence with respect to the original world.*

The interesting part of the condition is the italicized part which helps in solving the frame and ramification problems. In the case of sequential action formulas, the action function $\mathcal{A}$ will select the closest worlds with respect to each and every action formula in that sequence. In the case of parallel action formulas, the action function $\mathcal{A}$ tries to select the same closest worlds with respect to all the action formulas executed in parallel. As the closest world is required to be consistent, parallel actions that interfere will be blocked by the action formula [Rao, 1989].

A ***plan interpretation***, PI, is a dynamic coherence interpretation CI, together with an action function $\mathcal{A}$ [4]. The modal operator ACT is given a semantics similar to the semantics of EXP and CON using the selection function $\mathcal{A}$. This is done elsewhere [Rao, 1989] and here we discuss only the axioms of simple and sequential actions. The axioms which follow determine how the selection function $\mathcal{A}$ is related to the other functions discussed in the previous section. The axioms and inference rules are listed in Appendix - I.

The **Axiom of Simple Actions (AA1)** states that if agent a performs a simple action at time t to result in time u, then the precondition of the action must have been believed by the agent at time t and the agent must have revised at t with respect to the consequent of the action. The **Axiom of Sequential Actions (AA2)** states that carrying out a sequence of actions is equivalent to carrying out the first action and then carrying out the next action in the resulting state. In addition to the above axioms we also need the inference rule (RA1) which states that equivalent actions have equivalent effects. The axiom of parallel actions and the soundness and completeness of the system are discussed in the complete report [Rao, 1989].

Example: Consider a situation occurring at time point t_1, in which blocks A and B are on the table. The situation at t_1 can be said to contain the following set of beliefs.

(L1) $\forall t\ \mathrm{BEL}(a_1, t, (on(b,l) \wedge b \neq d \wedge l \neq table) \rightarrow \neg on(d,l))$

(L2) $\forall t\ \mathrm{BEL}(a_1, t, (on(b,l) \wedge l \neq m) \rightarrow \neg on(b,m)$

$\mathrm{BEL}(a_1, t_1, on(A,table))$ $\neg\mathrm{BEL}(a_1, t_1, \neg on(A,table))$
$\mathrm{BEL}(a_1, t_1, on(B,table))$ $\neg\mathrm{BEL}(a_1, t_1, \neg on(B,table))$
$\mathrm{BEL}(a_1, t_1, \neg on(A,B))$ $\neg\mathrm{BEL}(a_1, t_1, on(A,B))$.
$\mathrm{BEL}(a_1, t_1, \neg on(B,A))$ $\neg\mathrm{BEL}(a_1, t_1, on(B,A))$.

Now, if the agent carries out the action $move(A, B)^{a_1}$, assuming that no qualification disables this action (i.e., ignoring the qualification problem), one would ideally like to *completely* determine the state of the world after the action has been performed without stating any frame axioms. This is in fact the crux of the frame and ramification problems.

[4] The selection function $\mathcal{A}$ can be described in terms of the selection functions $\mathcal{E}$ and $\mathcal{C}$ described earlier. This, together with the constructive procedure for the selection functions $\mathcal{E}$ and $\mathcal{C}$, makes the selection function $\mathcal{A}$ less mysterious.

If t_2 is the time point after the action $move(A, B)^{a_1}$, and the action has the usual preconditions and consequents [Lifschitz, 1987], then from the axiom of simple actions we have:

$ACT(t_1, move(A,B)^{a_1}, t_2) \equiv BEL(a_1, t_1, (\forall x \neg on(x,A)) \wedge (\forall x \neg on(x,B)) \wedge (B \neq A)) \wedge REV(a_1, t_1, on(A,B), t_2)$.

Using the domain closure axiom and unique names assumption it is easy to verify that the precondition is satisfied by the belief system at time t_1. If the revision formula is satisfiable at t_1, we not only have the successful completion of the action $move(A, B)^{a_1}$, but also a completely determined next state giving all the beliefs of the agent at that time point. As discussed earlier, revision with respect to on(A,B) is contraction with respect to ¬on(A,B) followed by expansion with respect to on(A,B). Let t_2' be the time point after contraction. Using the axioms of contraction and the beliefs at time t_1, one can derive the following beliefs for time t_2'. (The satisfaction of these beliefs can also be verified using the selection function $\mathcal{E}$ or the corresponding construction procedure.)

$\neg BEL(a_1, t_2', on(A,table))$ $\neg BEL(a_1, t_2', \neg on(A,table))$.
$BEL(a_1, t_2', on(B,table))$ $\neg BEL(a_1, t_2', \neg on(B,table))$.
$\neg BEL(a_1, t_2', \neg on(A,B))$ $\neg BEL(a_1, t_2', on(A,B))$.
$BEL(a_1, t_2', \neg on(B,A))$ $\neg BEL(a_1, t_2', on(B,A))$.

Contraction forces the agent to give up belief in ¬on(A,B). If the agent continues to believe on(A,table), then the UAE belief (L2) will force him to believe ¬on(A,B) as well. Hence the agent is forced to give up belief in on(A,table) as well.

Expanding at t_2' with respect to the formula on(A,B), yields the following beliefs at time t_2.

$\neg BEL(a_1, t_2, on(A,table))$ $BEL(a_1, t_2, \neg on(A,table))$.
$BEL(a_1, t_2, on(B,table))$ $\neg BEL(a_1, t_2, \neg on(B,table))$.
$\neg BEL(a_1, t_2, \neg on(A,B))$ $BEL(a_1, t_2, on(A,B))$.
$BEL(a_1, t_2, \neg on(B,A))$ $\neg BEL(a_1, t_2, on(B,A))$.

Expansion forces the agent to believe on(A,B). From the UAE belief (L2) the agent is also forced to believe ¬on(A,table). The primary effect of the action is to make the agent believe on(A,B) and the secondary effect or ramifications of the action is to make the agent believe ¬on(A,table). Thus the formula $ACT(t_1, move(A,B)^{a_1}, t_2)$ is satisfiable and the time point t_2 gives all the beliefs of the agent after the action has been performed. Thus our theory of actions provides an inferential solution and a model-theoretic solution to the frame and ramification problems based on the principles of minimal change and maximal coherence.

The example presented above is a very simple one and has been used mainly to illustrate the theory. More complicated, multiple-agent, dynamic-world problems are discussed in the complete report [Rao, 1989]. For example consider the situation where block B is an ice-block and it melts before the block A can be placed on top of it. As this is a change caused by nature, we essentially have a revision by the agent NATURE at time t_1 with respect to the formula-on(B, table). It is impossible to have a single time point where the move action by the agent and the above revision be satisfiable. As the agent has no control over the changes caused by NATURE, we can impose a partial order on the dynamic operations and conclude that the revision performed by the agent NATURE succeeds and the move action by the agent fails. Thus the unified theory of belief revision and reasoning about actions provides a framework for reasoning about dynamic, real-world situations, where the world changes *during* the performance of an action.

4 Comparison

AGM Theory

There are significant differences between the coherence theory of belief revision as outlined by Gardenfors et. al. [1985, 1988] and that of section 2. Firstly, AGM-theory carries out the dynamic analysis at a meta-level and hence lacks a model-theory and the soundness and completeness result as outlined in this paper. Also unlike the theory outlined in section 2, AGM-theory does not handle nested beliefs and does not exhibit static nonmonotonic properties [Rao, 1989]. Finally, AGM-theory d oes not consider reasoning about actions.

Lifschitz's Approach

Lifschitz [1987] introduces two predicates, success and affects, to define the notion of a successful action and the notion of what fluents are affected by an action. Using these, he defines the *law of change(LC)* and *law of inertia(LI)* By circumscribing the predicate causes, Lifschitz is able to solve the frame problem. This is hardly surprising because causes describes the dynamics of the belief system by specifying what changes when the agent moves from one state to another as a result of an action. The presence of (LC) and minimization of *causes* is essentially equivalent to minimizing the change caused by an action. Similarly, the presence of (LI) and minimization of *causes* is equivalent to maximizing the coherence between state changes. Also the definitions of causes and precond are analogous to the axiom of simple actions. Hence Lifschitz's theory essentially embodies the principles of minimal change and maximal coherence.

However there are significant differences between our approach and that of Lifschitz. Firstly, Lifschitz's approach does not solve the ramification problem. Also the minimization process of Lifschitz is local [Hanks and McDermott, 1987] (i.e., with respect to a particular predicate) and relies on the user to specify the world in certain ways, whereas in the theory proposed here the minimization process is global (i.e., with respect to all the beliefs of the agent) and is independent of the syntax. Finally, the approach is incapable of handling dynamic situations. This is because the ontology is not powerful enough to represent or reason about changes caused by the external world.

Shoham's Approach

Based on the notion of preferential logics, Shoham [1987] introduces the *logic of chronological ignorance.* The underlying logic is a modal logic of temporal knowledge. A model M_2 of the above logic is *preferred over* another model M_1 if M_2 is *chronologically more ignorant* than M_1. The preferred models in the logic of chronological ignorance are the *chronologically maximally ignorant (c.m.i.)* models. By considering only the c.m.i. models of the given problem Shoham is able to solve the qualification problem (see the complete report [Rao, 1989] for more details) and extended prediction problem (a generalized form of the frame problem).

The notion of c.m.i. models can be simplified as follows - (1) prefer *maximally ignorant* models at each state (or time point) and (2) prefer *maximally ignorant* models when moving from one time point to another or during a state change. The first solves the qualification problem and the second solves the frame problem. Notice that in both cases one need not worry about a sequence of time points chronologically backward from the current time point. [5]

The net effect of revision is to believe as much of the original beliefs as possible and acquire as few beliefs as required. Acquiring as few beliefs as possible is the same as minimization of beliefs or maximization of non-beliefs (or ignorance in the case of knowledge). Thus (2) above, embodies the principle of minimal change.

The main differences between Shoham's approach and the one outlined in this paper are as follows. Firstly, his theory does not provide an axiomatic approach to the frame and ramification problems. Secondly, the principles of minimal change and maximal coherence are more general than the notion of chronologically maximizing ignorance. As a result our theory can be applied to database updates, belief revision, and reasoning about actions. However, it should be noted that Shoham defines a special class of theories called causal theories, which have an unique c.m.i model, and gives an algorithm for computing it. As mentioned before, this paper does not address the issue of efficient algorithms for computing contractions and revisions.

Reiter's Approach

In default logic [Reiter, 1980] one assumes by default that all actions and facts are normal. Actions which change the truth-value of certain facts are considered to be abnormal with respect to those facts. Thus given the default rules and the abnormality conditions the theory is supposed to predict the facts that will hold after the action has been performed. Unfortunately, this approach and other similar approaches using circumscription lead to counter-intuitive results as shown by Hanks and McDermott [1987].

[5]This fact was also noted by Lifschitz [1987].

In our opinion, the classical nonmonotonic theories like default logic and circumscription are not suitable for reasoning about change because they do not take into account the interaction with the environment. They are suitable only for static reasoning where the agent "jumps to conclusions" based on incomplete information]Rao, 1989]. Attempts to rectify the problem and make them suitable for dynamic reasoning, require extensive reformulation as done by Lifschitz and in any case has to incorporate some principles of change as discussed in this paper.

Possible-worlds Approach

Ginsberg and Smith [1987] first proposed the use of the possible-worlds model for solving the qualification, frame, and ramification problems. They provide a constructive procedure for computing the closest possible worlds. Their approach has certain problems [Winslett, 1988] as they minimize the formulas in the world and not the beliefs of the agent at a particular world as done in this paper. The significant difference between the possible-models approach (PMA) [Winslett, 1988] and our approach is the inability of the PMA to represent uncommitted states of the agent. This makes the notion of minimal change different in the two theories.

None of the above approaches are capable of handling dynamic situations [6] (the ice block example), where the state of the world changes during the execution of an action. This requires an explicit coupling of belief revision and reasoning about actions as carried out in this paper. Such situations will become more common when one starts addressing real-world problems.

5 Conclusion

This paper unifies the work in belief revision and reasoning about actions by analyzing the principles of minimal change and maximal coherence which forms a core part of both the areas. A sound and complete theory of belief revision is provided which is then used to represent and reason about actions. This theory of action solves the conceptual frame and ramification problems by utilizing the intuitive principles of minimal change and maximal coherence. It also provides a framework to reason about dynamic situations.

Acknowledgement: We would like to thank Mike Georgeff, Dave Wilkins and the anonymous referees' for their helpful comments. The first author was supported by the Sydney University Postgraduate Scholarship and the Research Foundation for Information Technology. Thanks are also due to IBM Corp. for their generous support during our stay at IBM T.J. Watson Research Center and IBM Thornwood respectively.

[6]More recently, Lifschitz and Rabinov [1989] have proposed a circumscriptive formalism to deal with dynamic situations. But their solution still suffers from some of the other drawbacks mentioned earlier and is incapable of representing multiple-agent domains.

References

[Alchourron et ai, 1985] Alchourron C, P. Gardenfors, and D. Makinson. On the logic of theory change. Journal of Symbolic Logic, 50:510 530, 1985.

[Brown, 1987] F. M. Brown, editor. Workshop on the Frame Problem in Artificial Intelligence. Morgan Kaufmann, Los Altos, 1987.

[de Kleer, 1986] J. de Kleer. An assumption-based TMS. Artificial Intelligence, 28:127-162, 1986.

[Doyle, 1979] J. Doyle. A Truth Maintenance System. Artificial Intelligence, 12:231-272, 1979.

[Kikes and Nilsson, 1971] R. Fikes and N. J. Nilsson. Strips: A new approach to the application of theorem proving to problem solving. Artificial Intelligence, 2:189-208, 1971.

[Gardenfors, 1988] P. Gardenfors. Knowledge in Flux: Modeling the Dynamics of Epistemic States. Bradford Book, MIT Press, Cambridge, Mass., 1988.

[Georgeff and Lansky, 1987] M. P. GeorgefF and A. L. Lansky, editors. Workshop on Reasoning about Actions and Plans. Morgan Kaufmann Publishers, San Mateo, CA, 1987.

[Georgeff, 1987] M. P. Georgeff. Actions, processes, and causality. In M. P. Georgeff and A. L. Lansky, editors, Workshop on Reasoning about Actions and Plans, pages 99-122, Los Altos, CA, 1987. Morgan Kaufmann Publishers.

[Ginsberg and Smith, 1987] M. L. Ginsberg and D. E. Smith. Reasoning about action 1: A possible worlds approach. In F. M. Brown, editor, The Frame Problem in Artificial Intelligence. Morgan Kaufmann, Los Altos, 1987.

[Hanks and McDermott, 1987] S. Hanks and D. McDermott. Nonmonotonic logic and temporal projection. Artificial Intelligence, 33:379—412, 1987.

[Konolige, 1988] K. Konolige. On the relation between default and autoepistemic logic. Artificial Intelligence, 35(3):343-382, 1988.

[Lifschitz and Rabinov, 1989] V. Lifschitz and A. Rabinov. Miracles in formal theories of action. Artificial Intelligence, 38(2), 1989.

[Lifschitz, 1987] V. Lifschitz. Formal theories of action. In F. M. Brown, editor, The Frame Problem in Artificial Intelligence, Los Altos, 1987. Morgan Kaufmann.

[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, volume 4, pages 463-502. American elsevier, New York, 1969.

[Moore, 1985] R. C. Moore. Semantical considerations on nonmonotonic logic. Artificial Intelligence, 25, 1985.

[Rao and Foo, 1989] A. S. Rao and N. Y. Foo. Formal theories of belief revision. In Proceedings of the First International Conference on Knowledge Representation and Reasoning, Toronto, Canada, 1989.

[Rao, 1989] A. S. Rao. Dynamics of belief systems: A philosophical, logical and AI perspective. Technical Note 2, Australian AI Institute, Carlton, Australia, 1989.

[Reiter, 1980] R. Reiter. A logic for default reasoning. Artificial Intelligence, 13:81-132,1980.

[SI loharn, 1987] Y. Shoharn. Reasoning About Change: Time and Causation from the Standpoint of Artificial Intelligence. The MIT Press, Cambridge, Mass., 1987.

[Winslett, 1988] M. Winslett. Reasoning about action using a possible models approach. In AAAI-88, pages 89-93, St. Paul, Minnesota, 1988.

Appendix - I

Axioms for Expansion

(AE1) $\mathrm{EXP}(a,t,\phi,u) \rightarrow \mathrm{BEL}(a,u,\phi)$

(AE2) $\mathrm{EXP}(a,t,\phi,u) \rightarrow (\mathrm{BEL}(a,t,\alpha) \rightarrow \mathrm{BEL}(a,u,\alpha))$

(AE3) $\mathrm{BEL}(a,t,\phi) \wedge \mathrm{EXP}(a,t,\phi,u) \rightarrow (\mathrm{BEL}(a,t,\alpha) \equiv \mathrm{BEL}(a,u,\alpha))$

(AE4) $(\mathrm{BEL}(a,t,\alpha) \rightarrow \mathrm{BEL}(a, v,\alpha)) \wedge \mathrm{EXP}(a,t,\phi,u) \wedge \mathrm{EXP}(a,v,\phi,y) \rightarrow (\mathrm{BEL}(a,u,\beta) \rightarrow \mathrm{BEL}(a,y,\beta))$

(AE5) $\mathrm{BEL}(a,u,\phi) \rightarrow \mathrm{BEL}(a,t,\phi) \vee (\mathrm{EXP}(a,t,\alpha,u) \wedge \mathrm{BEL}(a,u, \mathrm{BEL}(a,u,\alpha) \supset \phi))$

(RE1) From $\vdash \phi_1 \equiv \phi_2$ infer $\vdash \mathrm{EXP}(a, t,\phi_1,u) \equiv \mathrm{EXP}(a,t,\phi_2,u)$

Axioms for Contraction

(AC1) $\mathrm{CON}(a,t,\phi,u) \rightarrow \neg\mathrm{BEL}(a,u,\phi)$

(AC2) $\mathrm{BEL}(a,u,\phi) \rightarrow \mathrm{BEL}(a,t,\phi) \vee (\mathrm{CON}(a,t,\alpha,u) \wedge \mathrm{BEL}(a,u, \neg\mathrm{BEL}(a,u,\alpha) \supset \phi))$

(AC3) $\neg\mathrm{BEL}(a,t,\phi) \wedge \mathrm{CON}(a,t,\phi,u) \rightarrow (\neg\mathrm{BEL}(a,t,\alpha) \equiv \neg\mathrm{BEL}(a,u,\alpha))$

(AC4) $\mathrm{CON}(a,t,\phi,u) \wedge \mathrm{EXP}(a,u,\phi,v) \rightarrow (\mathrm{BEL}(a,t,\alpha) \rightarrow \mathrm{BEL}(a,v,\alpha))$

(RC1) From $\vdash \phi_1 \equiv \phi_2$ infer $\vdash \mathrm{CON}(a,t,\phi_1,u) \equiv \mathrm{CON}(a,t,\phi_2,u)$

Axiom of Revision

(AR1) $\mathrm{REV}(a,t,\phi,u) \equiv \mathrm{CON}(a,t,\neg\phi,v) \wedge \mathrm{EXP}(a,v,\phi,u)$

Axioms for Actions

(AA1) $\mathrm{ACT}(t,\gamma^a,u) \equiv \mathrm{BEL}(a,t,p^{\gamma^a}) \wedge \mathrm{REV}(a,t,c^{\gamma^a},u)$

(AA2) $\mathrm{ACT}(t, {\gamma_1}^{a_1} ; {\gamma_2}^{a_2}, u) \equiv \mathrm{ACT}(t, {\gamma_1}^{a_1}, x) \wedge \mathrm{ACT}(x, {\gamma_2}^{a_2}, u)$

(RA1) From $\vdash p^{\psi_1} \equiv p^{\psi_2}$ and $\vdash c^{\psi_1} \equiv c^{\psi_2}$ infer $\vdash \mathrm{ACT}(t, \psi_1, u) \equiv \mathrm{ACT}(t, \psi_2, u)$