Two Logical Theories of Plan Recognition

WAYNE WOBCKE, *School of Computer Science and Engineering,*
University of New South Wales, Sydney, NSW 2052 Australia.
E-mail: wobcke@cse.unsw.edu.au

Abstract

We present a logical approach to plan recognition that builds on Kautz's theory of keyhole plan recognition, defined as the problem of inferring descriptions of high-level plans from a set of descriptions of observed, typically low-level, actions. Kautz's approach is based on a formalization of the structural information contained in an abstraction/decomposition hierarchy of plan schemas: the inferred plans are selected by an observer agent from amongst those that can be generated from this hierarchy. In this paper, we present two logics for plan recognition. The first theory, like Kautz's, is monotonic, meaning that all possible plans compatible with the observations are treated as equally plausible by the observer. The second theory is based on rejecting this simplifying assumption, and is nonmonotonic. To develop this theory, we adapt ideas from belief revision and conditional logic to define plan recognition inference as a nonmonotonic consequence operation based on an observer's plausibility ordering on the possible plans. The logics incorporate a sound and complete theory of temporal intervals based on Allen's approach, enabling the plan recognition agent to handle temporal information relating multiple observations. An underlying motivation for our work is to provide an intuitive semantic interpretation of plans using situation semantics and to interpret plan recognition as inference over courses of events.

Keywords: Plan recognition, nonmonotonic reasoning, belief revision.

1 Plan recognition

The problem of plan recognition can be defined succinctly as 'to take as input a sequence of actions performed by an actor and to infer the goal pursued by the actor and also to organize the action sequence in terms of a plan structure' [46, p. 52]. Plan recognition is primarily a prediction task; the value to the observer is either to provide assistance to another agent (in a cooperative task) or to thwart another's plans (in adversarial scenarios). A major assumption is that the observed agent is rational, i.e. following a plan that the observer deems reasonable: there is no other basis upon which the observer can make predictions. In practice, both observer and observed agent are assumed to share common knowledge of a class of possible plans, usually represented as an abstraction/decomposition hierarchy relating complex abstract plans to their simpler concrete components.

One distinction drawn by Cohen, Perrault and Allen [13] is between *keyhole plan recognition* and *intended plan recognition*. In the former case, the observed agent may make no attempt to help or hinder the observer agent to recognize its plans, and may not even be aware that it is being observed (it is as if the observer looks through a keyhole). In the latter case, the observer and observed agents are typically collaborating on a joint task, and the observed agent intends the observer agent to recognize the plans it is executing. A special case of this is cooperative communication, specifically as construed by Grice [23], in which the recognition of intentions is paramount in guaranteeing successful communication. In the present paper, we deal exclusively with keyhole plan recognition, which formally is more general

J. Logic Computat., Vol. 12 No. 3, pp. 371–412 2002

in making no assumption about the cooperativeness of agents. However, this is not to deny the importance of intended plan recognition, which has been long studied in AI, c.f. Sidner and Israel [48]. Indeed the kind of intention recognition necessary for successful speech act recognition may be far more difficult to formalize than keyhole plan recognition, due to its reliance on non-syntactic features of communication such as gesture or intonation. Moreover, special attention may need to be paid to mismatches between the beliefs of the observer and the observed agent, c.f. Pollack [42].

There have been a number of attempts to formalize keyhole plan recognition in the AI literature. We broadly follow the approach of Kautz as taken in Kautz and Allen [28] and Kautz [26, 27], which is based on representing the information contained in a plan hierarchy using a first-order logic of events, then characterizing the plan recognition inferences using circumscription, McCarthy [40]. The intuition is that the plan hierarchy represents only information that is certain knowledge (from the observer's point of view) about the possible plans that can be constructed; for this to be used in plan recognition, the hierarchy needs to be 'completed' to enable stronger conclusions to be drawn—this essentially amounts to making the assumption that the hierarchy enables the construction of *all* possible plans that the observed agent could be following. Thus the intuition behind Kautz's approach is that the observer has incomplete information about the observed agent's planning knowledge, which through the use of circumscription, is turned into an assumption of complete knowledge.

Despite the use of circumscription, the logic of plan recognition that Kautz develops is monotonic, and his inference operation is defined using ordinary first-order deduction. This is because the result of performing the circumscriptions (for a given plan hierarchy) is representable as a first-order theory T, and the plan recognition inferences are defined as those formulae P which follow using first-order entailment from a set of observation formulae O together with T. This defines a monotonic inference operation, i.e. one for which, given any sets of sentences O and O' and any sentence P, $O \vdash P$ implies $O' \vdash P$ whenever $O \subseteq O'$. Intuitively, monotonicity here means that as the observer gathers more information, it need never retract any previous inferences. As noted by Kautz and Allen [28], this definition is based on the unrealistic assumption that all possible plans compatible with the observations are treated as equally plausible by the observer agent: there is no way for the observer to express the *a priori* plausibility of the various possible plans. Thus the inferences (and hence the predictive power) of a plan recognition system based on Kautz's theory will in general be too weak except in all but the simplest domains (such as help systems for Unix commands). Hence any general logical theory of plan recognition must be nonmonotonic.

In this paper, we propose an alternative approach to formalizing keyhole plan recognition. The basic motivation behind our approach is to define a semantics for plans using a version of situation semantics [7], then to characterize plan recognition as inference over situations. Situations (as opposed to first-order models or possible worlds) are suitable for modelling plans because they include a rich variety of semantic objects including individuals, objects, properties, relations, events and space-time locations, exactly the ingredients of plans. We construe observations of actions as descriptions of situations, i.e. that the observed situation includes the event being observed, then take plan recognition as a computational means of gaining information about the observed situation.

We present two logical theories of plan recognition. The first theory is monotonic, treating plan recognition as entailment over a class of plans that we call the *simple plans*. As in Kautz's theory, this sort of plan recognition can be characterized as deduction over a first-order language. The second theory is nonmonotonic. The basic idea is that some ordering on the simple plans is needed to capture the prior plausibility of these plans, analogous to the use of such orderings for modelling default reasoning, e.g. Shoham [47]. We can then characterize plan recognition as a nonmonotonic consequence operation defined as preferential entailment over the class of simple plans augmented with this plausibility ordering. Our work is thus based on established connections between orderings on models and the approach to nonmonotonic reasoning using expectations [19, 21]. This work derives from the theory of belief revision developed by Alchourrón, Gärdenfors and Makinson [2], which in turn is closely related to the work of Lewis [35] on the semantics of conditional logic. Indeed there is also a close intuitive connection between plan recognition inferences and indicative conditionals. Thus in this paper, the formal connections between plan recognition, nonmonotonic reasoning, belief revision and conditional logic are further strengthened.

The outline of this paper is as follows. We begin in Section 2 with a summary of Kautz's theory of plan recognition. Then in Section 3, we present a logic of entailment over situations, reconstruct Allen's logic of temporal reasoning for use in modelling courses of events, then show how to characterize plan recognition as deduction over theories corresponding to the class of simple plans. This work is our analogue of Kautz's theory of plan recognition. In Section 4, we show how to define a nonmonotonic consequence operation for plan recognition based on an ordered structure of situations that corresponds to a prior plausibility ordering on the class of simple plans. We construe an agent's epistemic state as a belief set together with a nonmonotonic consequence operation, and show how epistemic states should be revised in the light of new observations. This constitutes our nonmonotonic theory of plan recognition. We conclude with a discussion of related work, focusing on the role of context and epistemic state dynamics in plan recognition.

2 Kautz's theory of plan recognition

In Kautz's theory of plan recognition [28, 26, 27], the observer agent is assumed to start with a collection of plan schemas (action types). An action may have a number of subtypes (ways of specializing the action type) and may have a decomposition (a procedure to be followed in performing the action), together forming an intertwined abstraction/decomposition hierarchy. For example, a subtype of the action type for going to the city might be to take the train to the city, which in turn might have the decomposition of going to the train station, buying a ticket, etc. A sample hierarchy from the cooking domain, taken from [26], is shown in Figure 1 and will be used to illustrate the theory. Shaded arrows represent the abstraction relation (going from special to general), while thin black arrows represent the decomposition relation (going from plan schema to component steps). The decomposition arrows for a single action type are labelled with distinct function symbols distinguishing the different substeps in the plan: the names **s1**, **s2**, etc. carry no formal significance but are meant to be suggestive of first substep, second substep, etc.

Kautz makes a number of assumptions about the plan hierarchy. First, the plan hierarchy is assumed to form a finite acyclic graph, i.e. there is no path from any node to itself. Second, the hierarchy is assumed to include special action types *AnyEvent* and *End*: the specializations of *End* are those action types that are assumed to be 'self-motivating' and indicate the point at which plan recognition stops. Third, *End* and its specializations are exactly the set of schemas that do not appear in the decomposition of any action type.

The logical language used in Kautz's theory is first-order predicate calculus with a standard equality predicate. Kautz uses the 'reified' representation of events from Allen [4] in which

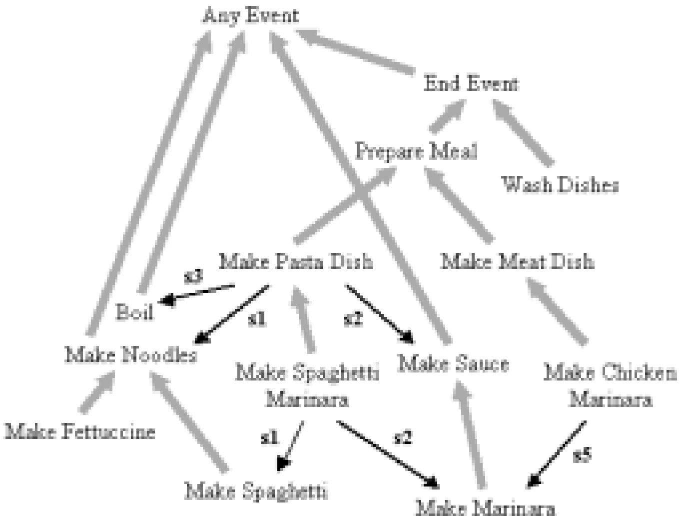

FIGURE 1. Cooking hierarchy

event tokens are represented by terms and event types by predicates, so that a formula such as $\alpha(x)$ denotes the proposition that an event token x is of type α.[1] The language is also assumed to include terms denoting temporal intervals, and the logic is assumed to include a sound and complete first-order theory of temporal relations on the intervals.

The basic knowledge of plans contained in the abstraction/decomposition hierarchy is represented as the set of abstraction and decomposition formulae.

DEFINITION 2.1
An action type α' *directly abstracts* an action type α if there is a single abstraction arrow from α to α'. The transitive closure of this relation is the *abstraction* relation.

Abstraction. If action type α' directly abstracts action type α:

(ABS) $\forall x(\alpha(x) \rightarrow \alpha'(x))$.

EXAMPLE 2.2
$\forall x(\mathit{PastaDish}(x) \rightarrow \mathit{PrepareMeal}(x))$

Decomposition. If action type α has decomposition $\alpha_1, \cdots, \alpha_n$ labelled $\mathbf{s_1}, \cdots, \mathbf{s_n}$ respectively:

(DEC) $\forall x(\alpha(x) \rightarrow (\alpha_1(\mathit{step}_1(x)) \wedge \cdots \wedge \alpha_n(\mathit{step}_n(x)) \wedge \kappa))$.

In such formulae, each symbol step_i denotes a function identifying the correspondingly labelled step $\mathbf{s_i}$ in the decomposition. The symbol κ denotes a conjunction of constraints on the decomposition, such as conditions on the types of role fillers, the preconditions and effects of the plan, and the temporal relationships between steps in the decomposition. These

---

[1] The earlier formulation of the theory in Kautz and Allen [28] differs in using an ‘occurs’ predicate. Here we follow the later presentations of Kautz [26, 27].

constraints are typically expressed using functions to identify role fillers and relations on temporal intervals. For example, the constraint that two substeps in a decomposition have the same agent and occur in temporal sequence might be represented using the formula $\mathit{agent}(\mathit{step}_1(x)) = \mathit{agent}(\mathit{step}_2(x)) \wedge \mathit{before}(\mathit{time}(\mathit{step}_1(x)), \mathit{time}(\mathit{step}_2(x)))$, where *agent* and *time* are functions denoting the agent and time of an action, and *before* denotes the temporal precedence relation.

EXAMPLE 2.3

$$\begin{aligned}&\forall x(\mathit{PastaDish}(x) \rightarrow\\ &\quad \mathit{Noodles}(s1(x)) \wedge \mathit{Sauce}(s2(x)) \wedge \mathit{Boil}(s3(x)) \wedge\\ &\quad agent(s1(x)) = agent(x) \wedge agent(s2(x)) = agent(x) \wedge agent(s3(x)) = agent(x) \wedge\\ &\quad during(time(s1(x)), time(x)) \wedge \mathit{before}(time(s1(x)), time(s3(x))))\end{aligned}$$

The knowledge in the plan hierarchy is assumed to be correct, but the central motivation of Kautz's theory is that this knowledge is, by itself, insufficient for plan recognition. For plan recognition, the theory representing the information in the plan hierarchy must be 'completed' using circumscription. The essential effect of performing the circumscriptions is that where the original hierarchy contains facts such as that a schema has a number of subtypes and occurs in a number of decompositions, the completed knowledge is that the action type has *only* these subtypes as possible subtypes, and occurs *only* in those decompositions. This additional information is not certain knowledge from the observer's point of view, because the observer agent cannot be sure that it is aware of all subtypes of an action type, nor all action types in which a given action type occurs. It is not necessary to review details of the formal circumscription operations here: what matters is that the additional assumptions that result are representable as a set of first-order formulae, their models can be characterized in terms of predicate minimization, and that plan recognition can be formalized as first-order deduction.

The following definition of compatibility is needed.

DEFINITION 2.4

Two action types are *compatible* if they have a common specialization under the (transitively closed) abstraction relation.

Note that as special cases of this definition, two action types are compatible if they are identical or if one is a specialization of the other.

Completing the abstraction hierarchy amounts to assuming that an action must be specialized by one of its *known* subtypes. The result of completing the abstraction hierarchy is the following set of formulae.

Exhaustiveness assumptions. If $\alpha_1, \cdots, \alpha_n$ are all the action types that α directly abstracts:

(EXA) $\forall x(\alpha(x) \rightarrow \alpha_1(x) \vee \cdots \vee \alpha_n(x))$.

EXAMPLE 2.5

$\forall x(\mathit{PrepareMeal}(x) \rightarrow (\mathit{PastaDish}(x) \vee \mathit{MeatDish}(x)))$

Disjointedness assumptions. If α_1 and α_2 are not compatible:

(DJA) $\forall x(\neg\alpha_1(x) \vee \neg\alpha_2(x))$.

EXAMPLE 2.6

$\forall x(\neg \mathit{PastaDish}(x) \vee \neg \mathit{MeatDish}(x))$

The class of models satisfying these formulae is characterized by a notion of A-closure.

DEFINITION 2.7
A model M in a set S is *minimal in a set of predicates* Π if there is no other model M' in S that agrees with M on the interpretation of all constants, functions and predicates not in Π, and for which the extension of each predicate $\pi \in \Pi$ is a subset of the extension of π in M.

DEFINITION 2.8
Let H_{NB} be the set of non-basic action types (those with at least one specialization), including the special action types *AnyEvent* and *End*. A model M is *closed under specialization* if M is minimal in H_{NB} among the models of (ABS).

Suppose an action in some model is of some basic type, and hence by (ABS), is of a number of non-basic types. The effect of closure under specialization is to ensure that in the model, the action is of only those non-basic types licensed by the (ABS) formulae.

THEOREM 2.9 (Kautz)
M is a model of (ABS) closed under specialization iff $M \models$ (ABS) $\cup$ (EXA).

DEFINITION 2.10
Let E be the set of all action types, including the special action types *AnyEvent* and *End*. A model M is *closed under abstraction* if M is minimal in $E - \{\mathit{AnyEvent}\}$ among the models of (ABS) that are closed under specialization.

Let H be the set of all (ABS) and (DEC) formulae, i.e. the certain knowledge represented in the plan hierarchy.

DEFINITION 2.11
A model M is an *A-closed model of H* if M is a model of H and a model of (ABS) that is closed under abstraction.

The effect of A-closure is to ensure that in any model, each action (object of type *AnyEvent*) is of a single basic type, and hence of only those types that abstract this basic type.

THEOREM 2.12 (Kautz)
M is an A-closed model of H iff $M \models H \cup$ (EXA) $\cup$ (DJA).

Completing the decomposition hierarchy essentially amounts to assuming that each action occurs as a component of another *known* action whenever possible. Thus this heuristic enforces a kind of coherence of a set of actions. To express this heuristic, Kautz utilizes the dummy *End* action type, which is defined so that its specializations are exactly those action types which are not components (i.e. do not occur directly in the decomposition) of any other action type. The result of performing the completion is the following set of formulae known as Component/Use assumptions.

Component/Use assumptions. Consider all the decomposition axioms derived from the plan hierarchy where an action type α_i compatible with an action type α occurs as a direct component of an action type α^j, as i and j vary. These formulae are all of the following form.

(DEC) $\forall x(\alpha^j(x) \rightarrow (\cdots \wedge \alpha_i(\mathit{step}_{i_j}(x)) \wedge \cdots \wedge \kappa_j))$.

Then the component/use assumptions are the formulae of the following form.

(CUA) $\forall x(\alpha(x) \rightarrow \mathit{End}(x) \vee (\cdots \vee \exists y(\alpha^j(y) \wedge (x = \mathit{step}_{i_j}(y))) \vee \cdots))$.

Each disjunct on the right-hand side, apart from $\mathit{End}(x)$, corresponds to an occurence in a (DEC) formula of $\alpha_i(x)$ where α_i is compatible with α. The (CUA) formula means that any action of type α is either an *End* action or it occurs as a component (the i_jth component) of an action of type α^j (because one of its specializations that is also a subtype of α_i is in the decomposition of α^j).

EXAMPLE 2.13

$$\forall x(\mathit{Noodles}(x) \rightarrow$$
$$\exists y(\mathit{PastaDish}(y) \wedge (x = \mathit{step}_1(y))) \vee \exists y(\mathit{SpaghettiMarinara}(y) \wedge (x = \mathit{step}_1(y))))$$

Note the importance of using compatibility of action types: an observation of *Noodles* does *not* lead to the prediction of *PastaDish* as it would by merely following the component arrow: the possibility that the noodles are spaghetti is also taken into account, in which case the observed plan would be *SpaghettiMarinara*; in fact, by an exhaustiveness assumption, the observed plan *must* be *SpaghettiMarinara.*

The models satisfying the (CUA) formulae are the class of *covering models*.

DEFINITION 2.14

A model M is a *covering model of* H if M is minimal in $E - \{\mathit{End}\}$ among the A-closed models of H.

In any model, if an action is of a type which has a decomposition, there are related actions whose types are those of the steps in the decomposition. In covering models, a partial converse is true: every action whose type is not a specialization of *End* occurs as a component of some other action (whose type either specializes *End* or specializes a schema which appears as the component of another action type, and so on, until *End* is reached).

THEOREM 2.15 (Kautz)

M is a covering model of H iff $M \models H \cup$ (EXA) $\cup$ (DJA) $\cup$ (CUA).

As Kautz states, using these formulae does not ensure that multiple observations are interpreted as being components of the *same* action whenever possible. To achieve this, it is necessary to minimize the number of *End* actions generated by a set of observations. The syntactic analogue of this is that the strongest consistent minimum cardinality assumption of the following form (the one with the fewest variables) must be assumed. Note that automatically determining this formula requires a check for consistency in first-order logic, which is generally not computable.

Minimum cardinality assumption. Assume the strongest formula of the following form.

$$(\mathrm{MCA}n)\ \ \forall x_1 \cdots \forall x_n(\mathit{End}(x_1) \wedge \cdots \wedge \mathit{End}(x_n) \rightarrow \bigvee_{i,j}(x_i = x_j)).$$

The formula with n variables means that there are no more than $n-1$ *End* actions.

EXAMPLE 2.16

$$\forall x_1 \forall x_2 \forall x_3(\mathit{End}(x_1) \wedge \mathit{End}(x_2) \wedge \mathit{End}(x_3) \rightarrow ((x_1 = x_2) \vee (x_1 = x_3) \vee (x_2 = x_3)))$$

Minimum cardinality models are characterized simply as those in whose extensions the *End* predicate is minimal.

DEFINITION 2.17

A model M has *minimum cardinality in a set* S *with respect to a predicate* π if there is no other model M' in S for which the size of the extension of π is less than the size of the extension of π in M.

THEOREM 2.18 (Kautz)
If (MCA_n), the strongest minimum cardinality assumption, exists, then M is a minimum cardinality model of H with respect to *End* iff $M \models H \cup (\mathrm{EXA}) \cup (\mathrm{DJA}) \cup (\mathrm{CUA}) \cup (\mathrm{MCA}_n)$.

DEFINITION 2.19
Let Γ be the collection of all (ABS), (DEC), (EXA), (DJA) and (CUA) formulae together with the strongest minimum cardinality assumption (MCA_n). A set of observation formulae O *entails* a prediction P iff every minimum cardinality model of O satisfies P.

COROLLARY 2.20 (Kautz)
Let Γ be the collection of all (ABS), (DEC), (EXA), (DJA) and (CUA) formulae together with the strongest minimum cardinality assumption (MCA_n). A set of observation formulae O entails a prediction P iff $\Gamma \cup O \vdash P$ using first-order deduction with the standard axioms of equality and the interval theory.

3 A monotonic logic for plan recognition

In this section, we develop a formal model of plan recognition as deduction over theories characterizing a class of plans that we call the ‘simple plans’. This approach is analogous to Kautz’s work, but is really only a preliminary step towards developing a nonmonotonic logic of plan recognition, which is done in Section 4. In contrast to Kautz’s theory, which uses first-order logic and circumscription, our work is based on providing a more direct semantics for plans based on situation semantics. Thus the first step is to define a monotonic logic, or more precisely, a class of monotonic consequence operations, that corresponds to entailment over sets of situations. The next step is to reconstruct Allen’s theory of temporal intervals to enable a richer class of situations — courses of events — to be represented. The final step is to define the class of simple plans for any given plan hierarchy and to show how entailment over this class can be represented as a monotonic consequence operation.

The reason for using situation semantics is that situations, as defined by Barwise and Perry [7], are portions of space-time consisting of individuals, objects, properties, relations, events and space-time locations, and as such are natural semantic counterparts of plans. In our usage of situation semantics, we follow the terminology of Devlin [15]. Situations resolve *issues*: an issue is an ascription of a property to an object, a claim that objects stand in some relation in some location or throughout some time period, or that an event occurs throughout a time period. If an issue is resolved in a situation, it is resolved with some *polarity*: either positive or negative. Issues with polarity are called *basic infons* (the basic units of information). Compound infons may be formed as logical combinations of basic and other infons using the operations of conjunction, disjunction and negation (note that we do not consider implication). If an issue is resolved positively in a situation, the situation *supports* the corresponding infon; if negatively, the situation *rejects* the infon.

Situations are contrasted with possible worlds and propositional models, although some authors, e.g. Cresswell [14], have argued that the notion of a possible world is broad enough to include situations. The main difference between situations and possible worlds (as understood by Kripke [30] or Lewis [36]) is that although every proposition must be assigned a truth value in every possible world, and of course this is the case in every propositional model, a situation may leave some issues unresolved. A formal difference is that although $A \vee \neg A$ holds in every possible world and every propositional model, there are situations in which $A \vee \neg A$ does not hold, namely those situations in which A is undefined. However, any inference that holds over the class of situations is valid over the class of possible worlds (and possible models), hence the logic of entailment over situations is strictly weaker (allows fewer inferences in general) than the propositional calculus.

3.1 Entailment over situations

We now define a consequence operation corresponding to entailment over situations, which we later use to formalize a deductive approach to plan recognition similar to Kautz's theory. Our modelling of situations comes from Devlin [15] and indirectly from Barwise and Perry [7]: a situation is modelled as a set of infons which are themselves issues with an assigned polarity. We follow Devlin in defining supports and rejects relations between situations and the class of compound infons: a situation supports a disjunction if it supports either disjunct and rejects a disjunction if it rejects both disjuncts; a situation supports a conjunction if it supports both conjuncts and rejects a conjunction if it rejects either conjunct; and a situation supports the negation of an infon if it rejects the infon, and rejects the negation of an infon if it supports the infon. Devlin's approach here borrows from that employed in partial logic (see the survey article by Blamey [10]).

We begin with the basic definitions, assuming a given collection of issues.

DEFINITION 3.1
A *basic infon* is an issue together with a polarity (positive or negative).

DEFINITION 3.2
The *infons* are defined as follows. A basic infon is an infon, and if A and B are infons then so are $\neg A$, $A \wedge B$ and $A \vee B$.

DEFINITION 3.3
The *dual* $A^{\perp}$ of a basic infon A is the issue of A with opposite polarity.

DEFINITION 3.4
A *situation* is a set of basic infons.

DEFINITION 3.5
A situation is *coherent* if it does not contain any basic infon and its dual.

DEFINITION 3.6
A coherent situation σ *supports (rejects)* an infon A, written $\sigma \models A$ ($\sigma \mathrel{=\!\!|} A$), under the following conditions.

$$\begin{array}{ll}
\sigma \models A & \text{if } A \in \sigma \text{ for a basic infon } A \\
\sigma \models \neg A & \text{if } \sigma \mathrel{=\!\!|} A \\
\sigma \models A \wedge B & \text{if } \sigma \models A \text{ and } \sigma \models B \\
\sigma \models A \vee B & \text{if } \sigma \models A \text{ or } \sigma \models B \\
 & \\
\sigma \mathrel{=\!\!|} A & \text{if } A^{\perp} \in \sigma \text{ for a basic infon } A \\
\sigma \mathrel{=\!\!|} \neg A & \text{if } \sigma \models A \\
\sigma \mathrel{=\!\!|} A \wedge B & \text{if } \sigma \mathrel{=\!\!|} A \text{ or } \sigma \mathrel{=\!\!|} B \\
\sigma \mathrel{=\!\!|} A \vee B & \text{if } \sigma \mathrel{=\!\!|} A \text{ and } \sigma \mathrel{=\!\!|} B
\end{array}$$

DEFINITION 3.7
A set of infons Γ *entails* an infon A with respect to a set of situations Σ, written $\Gamma \models_{\Sigma} A$, if every coherent situation in Σ that supports all the infons in Γ also supports A.

Situations provide the semantic grounding for plans; inference over situations provides a deductive account of plan recognition. This is formalized using the notion of a (monotonic) consequence operation, an operation, denoted $\vdash$, relating premiss sets (corresponding to observations) to inferred conclusions (corresponding to predictions). This approach derives from work in logic on Gentzen proof systems, which also forms the basis of the work of Gabbay [17] on nonmonotonic consequence operations, the approach we follow in Section 4 when our own nonmonotonic theory of plan recognition is developed.

DEFINITION 3.8
A *consequence operation* is a relation $\vdash$ between sets of infons and infons that satisfies the following conditions (where $\Gamma \vdash \Delta$ means that $\Gamma \vdash A$ for each $A \in \Delta$).

If $\Gamma \vdash A$ and $\Gamma \subseteq \Delta$ then $\Delta \vdash A$ (Monotonicity)
If $\Gamma \vdash \Delta$ and $\Gamma \cup \Delta \vdash A$ then $\Gamma \vdash A$ (Cut)
If $\Gamma \vdash A$ and $\Gamma \vdash B$ then $\Gamma \vdash A \wedge B$ (And)

DEFINITION 3.9
A consequence operation $\vdash$ is *compact* if whenever $\Gamma \vdash A$, there is some finite subset Γ' of Γ such that $\Gamma' \vdash A$.

DEFINITION 3.10
A *situated consequence operation* $\vdash$ is a compact consequence operation that satisfies the following postulates ($A \vdash B$ is used as an abbreviation for $\{A\} \vdash B$, and $A \dashv\vdash B$ as an abbreviation for $A \vdash B$ and $B \vdash A$ – say A and B are *provably equivalent under* $\vdash$).

$A \vdash A$ (Identity)
$A \vdash A \vee B$, $B \vdash A \vee B$ (Or-Introduction)
$A \wedge B \vdash A$, $A \wedge B \vdash B$ (And-Elimination)
$A \dashv\vdash \neg\neg A$ (Double Negation)
$A \wedge (B \vee C) \dashv\vdash (A \wedge B) \vee (A \wedge C)$ (And-Distribution)
$(A \vee B) \wedge (A \vee C) \dashv\vdash A \vee (B \wedge C)$ (Or-Distribution)
$\neg(A \wedge B) \dashv\vdash \neg A \vee \neg B$ (And-De Morgan)
$\neg(A \vee B) \dashv\vdash \neg A \wedge \neg B$ (Or-De Morgan)
$A \wedge (\neg A \vee B) \vdash B$ (Or-Elimination)
If $A \vdash B$ and $A \vdash C$ then $A \vdash B \wedge C$ (And-Introduction)
If $A \vdash C$ then $A \wedge B \vdash C$ (And-Premiss)
If $A \vdash C$ and $B \vdash C$ then $A \vee B \vdash C$ (Or-Premiss)

The names of the postulates come from inference rules in natural deduction systems for the propositional calculus. Conspicuously absent is the rule of contradiction, i.e. if $A \vdash$ *false* infer $\neg A$ (where *false* is some contradiction). This rule is invalid for entailment over situations since $B \wedge \neg B \vdash$ *false* but $\neg(B \wedge \neg B)$, which is equivalent to $B \vee \neg B$ using the De Morgan principles, is not valid (it fails to hold at those situations in which B is undefined).

PROPOSITION 3.11
Any situated consequence operation $\vdash$ satisfies the following properties.

If $\Gamma \dashv\vdash \Delta$ then $\Gamma \vdash A$ iff $\Delta \vdash A$ (Left Logical Equivalence)
If $\Gamma \vdash A$ and $A \vdash B$ then $\Gamma \vdash B$ (Right Weakening)
If $\Gamma \cup \{A\} \vdash C$ and $\Gamma \cup \{B\} \vdash C$ then $\Gamma \cup \{A \vee B\} \vdash C$ (Or)

(Left Logical Equivalence) means that provably equivalent sets of infons have identical consequences—this follows from (Monotonicity), (Cut) and compactness. (Right Weakening) allows any consequence of an inferred conclusion to be inferred: thus the consequences of a given set form a closed set under consequence. Finally, (Or) sanctions disjunction in the premisses.

Soundness and completeness of the class of consequence operations $\vdash$ with respect to entailment over situations is straightforward.

THEOREM 3.12
The class of situated consequence operations is sound and complete with respect to entailment over sets of situations. That is, (soundness) the consequence operation $\vdash_\Sigma$ defined as entailment over a set of situations Σ by $\Gamma \vdash_\Sigma A$ iff $\Gamma \models_\Sigma A$ is a situated consequence operation, and (completeness) given a situated consequence operation $\vdash$, there is a set of situations Σ such that $\Gamma \models_\Sigma A$ iff $\Gamma \vdash A$.

The proof makes use of a characterization of the coherent situations in terms of prime closed consistent sets of infons (whereas for propositional logic, possible worlds/models are characterized in terms of maximal consistent sets).

DEFINITION 3.13
A set of infons Γ is *closed* under a consequence operation $\vdash$ if $A \in \Gamma$ whenever $\Gamma \vdash A$.

DEFINITION 3.14
A set of infons Γ is *prime* if whenever Γ contains $A \vee B$, Γ contains A or Γ contains B.

3.2 Modelling courses of events

Our next task is to consider the temporally extended version of situations, known as *courses of events* in Barwise and Perry [7]. In this section, we first consider the logic of inference over formulae representing temporal constraints, adapting the approach of Allen [4]. We then show how this logic can be used to define consequence operations for reasoning with actions and events. In the next section, we show how this can be applied to plan recognition.

Our work follows that of Allen [4], who aimed to provide a general theory of action and time for use in planning and plan recognition. Allen's approach is based on an interval calculus in which intervals are first class entities, and basic propositions denote relationships between intervals: this approach has a natural affinity with the situation semantics treatment although based on a different technical framework (first-order logic). More specifically, we follow Allen [3], which described a system for representing and reasoning about temporal relationships between intervals. Allen's system has 13 primitive relations that together provide the expressive power for representing any ordering relationship between the start and end points of two intervals. Allen was not primarily concerned with logical questions: his focus was on using a constraint propagation algorithm to efficiently reason with temporal relations. This method is viable because a significant portion of the temporal logic can be represented using ternary constraints (analogous to the axioms expressed in the transitivity table, see below). The system was also limited in linguistic expressiveness in that the formulae representable in a constraint network (over which constraint propagation occurs) correspond to disjunctions of basic relationships between pairs of intervals, e.g. $(i \prec j) \vee (i \succ j)$.

As shown by Allen [3], the constraint propagation algorithm is incomplete. However, Ladkin and Maddux [32], in work cited in Ladkin [31], claim that Allen's logic is complete with respect to interval structures based on the real line. Primarily to demonstrate that there is no obstacle in incorporating Allen's theory into our own approach, we give an appropriate reconstruction of Allen's theory and prove the result below. The main technical issue is to ensure that Allen's system of temporal reasoning can be reconstructed using only the proposition-like language of infons, rather than the first-order language Allen used. In our language, infons represent basic temporal relationships holding between pairs of intervals denoted by primitive interval terms, and the models of the language are given by interval structures, defined below.

DEFINITION 3.15
An *interval constant* or *variable* is a primitive symbol that denotes an interval. An *interval term* is an interval constant or variable or an expression of the form $f(i)$ where f is a function symbol and i is an interval term.

DEFINITION 3.16
A *temporal infon* over a set of interval terms $\mathcal{I}$ is an infon all of whose basic constituents are of the form $i\ r\ j$ (with positive polarity) where i and j are in $\mathcal{I}$ and r is one of the relations $=$, $\prec$, during, overlaps, meets, starts and finishes, or their inverses (for all except equality), denoted $\succ$, during$'$, overlaps$'$, meets$'$, starts$'$ and finishes$'$.

DEFINITION 3.17
An *interval structure* over a set of interval terms $\mathcal{I}$ is a set ξ of finite closed intervals over the real line giving an interpretation for each interval term $i \in \mathcal{I}$ in which the interval interpreting i is identified by its start point i_x and end point i_y (with $i_x < i_y$).

DEFINITION 3.18
An interval structure ξ *supports (rejects)* a temporal infon A, written $\xi \models A$ ($\xi \mathrel{=\!\!|} A$), under the following conditions. First, the satisfaction conditions for temporal infons are as follows.

$\langle i_x, i_y \rangle = \langle j_x, j_y \rangle$	if $i_x = j_x$ and $i_y = j_y$
$\langle i_x, i_y \rangle \prec \langle j_x, j_y \rangle$	if $i_y < j_x$
$\langle i_x, i_y \rangle$ during $\langle j_x, j_y \rangle$	if $j_x < i_x$ and $i_y < j_y$
$\langle i_x, i_y \rangle$ overlaps $\langle j_x, j_y \rangle$	if $i_x < j_x$, $j_x < i_y$ and $i_y < j_y$
$\langle i_x, i_y \rangle$ meets $\langle j_x, j_y \rangle$	if $i_y = j_x$
$\langle i_x, i_y \rangle$ starts $\langle j_x, j_y \rangle$	if $i_x = j_x$ and $i_y < j_y$
$\langle i_x, i_y \rangle$ finishes $\langle j_x, j_y \rangle$	if $j_x < i_x$ and $i_y = j_y$

Second, the inverse relations are defined as follows: $i \succ j$ iff $j \prec i$ and $i\ r'\ j$ iff $j\ r\ i$ for r equal to during, overlaps, meets, starts and finishes. Finally, the satisfaction conditions for complex temporal infons are as for situations except that for a basic infon A, $\xi \mathrel{=\!\!|} A$ iff $\xi \not\models A$.

The intuitive meaning of the relations is as follows: $i = j$ means i and j are identical intervals, $i \prec j$ means i temporally precedes j, i during j means i is (properly) contained in j, i overlaps j means i starts before j and ends after j starts but before j ends (importantly, i overlaps j is different from j overlaps i), i meets j means the end point of i is the start point of j, i starts j means i forms a (proper) initial segment of j, and i finishes j means i forms a (proper) final segment of j. The relations are depicted in Figure 2, taken from Allen [5]. As can be seen in the figure, for an interval i, when $i_y < j_y$ there are five possibilities for the ordering of i_x and i_y with respect to j_x, giving the interval relations $\prec$, meets, overlaps, starts

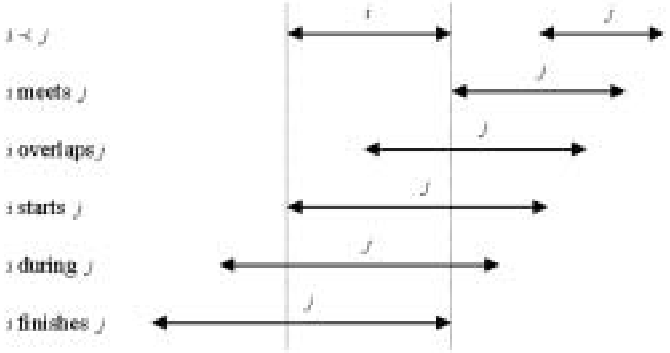

FIGURE 2. Allen's temporal relations

and during (and similarly when $j_y < i_y$ giving their inverse relations), and when $i_y = j_y$ there are three more possibilities, giving $=$ (not shown), finishes and the inverse of finishes. Thus it should be apparent from the figure that all possible orderings between the start and end points of two intervals generate an expressible interval relation.

Situated consequence operations based on the following set of axiom schemes $\Im$ are sound and complete with respect to the interval structures. We assume that an interval theory is based on a given countable set of interval terms $\mathcal{I}$: the infons of $\Im$ are the temporal infons over $\mathcal{I}$. Following Allen [3], we use the abbreviation $i\ (r_1 \cdots r_n)\ j$ to denote the finite disjunction $(i\ r_1\ j) \vee \cdots \vee (i\ r_n\ j)$. We also use the transitivity table from Allen [3], reproduced in Figure 3, for ease in expressing the 144 transitivity axiom schemes (based on the 12 temporal relations other than equality). Each entry in the table represents an axiom scheme of the form $(i\ r_1\ j) \wedge (j\ r_2\ k) \supset (i\ (\mathcal{T}(r_1, r_2))\ k)$, and the temporal relations are abbreviated as follows: d for during, o for overlaps, m for meets, s for starts, and f for finishes. We also use $A \equiv B$ as an abbreviation for $(A \supset B) \wedge (B \supset A)$. Finally, in the second equality axiom scheme, $A(j/i)$ is any temporal infon with j replacing any (or all) occurrences of i in A.

The system $\Im$ consists of the following axiom schemes.

(PC) $A \vee \neg A$ for A a temporal infon
(Equality) $i = i$, $(i = j) \supset (j = i)$, $(i = j) \wedge (j = k) \supset (i = k)$
(Equality) $(i = j) \supset (A \equiv A(j/i))$
(Definition) $(i \prec j) \equiv (j \succ i)$
(Definition) $(i\ r\ j) \equiv (j\ r'\ i)$ for r = during, overlaps, meets, starts, finishes
(Transitivity) $(i\ r_1\ j) \wedge (j\ r_2\ k) \supset (i\ (\mathcal{T}(r_1, r_2))\ k)$ from the transitivity table
(Totality) $i\ (= \prec \succ\ d\, d'\ o\, o'\ m\, m'\ s\, s'\ f\ f')\ j$
(Exclusivity) $\neg(i\ r_1\ j) \vee \neg(i\ r_2\ j)$ for distinct relations r_1 and r_2.

Note that one difference between our reconstruction and Allen's original theory is that the above transitivity axioms can be *derived* in Allen's system from a first-order theory containing only the meets and equality relations as primitive. The axioms of this theory that are not representable in our language concern the existence of intervals, e.g. representing properties of time such as 'for any interval i there exists an interval j such that i meets j'. However, it is not necessary to use such properties in reasoning about relationships between given intervals. This is because the class of situated consequence operations whose infons are all temporal

	$\prec$	$\succ$	d	d'	o	o'	m	m'	s	s'	f	f'
$\prec$	$\prec$	all	$\prec\ d\ o\ m\ s$	$\prec$	$\prec$	$\prec\ d\ o\ m\ s$	$\prec$	$\prec\ d\ o\ m\ s$	$\prec$	$\prec$	$\prec\ d\ o\ m\ s$	$\prec$
$\succ$	all	$\succ$	$\succ\ d\ o'\ m'\ f$	$\succ$	$\succ\ d\ o'\ m'\ f$	$\succ$	$\succ\ d\ o'\ m'\ f$	$\succ$	$\succ\ d\ o'\ m'\ f$	$\succ$	$\succ$	$\succ$
d	$\prec$	$\succ$	d	all	$\prec\ d\ o\ m\ s$	$\succ\ d\ o'\ m'\ f$	$\prec$	$\succ$	d	$\succ\ d\ o'\ m'\ f$	d	$\prec\ d\ o\ m\ s$
d'	$\prec\ d'\ o\ m\ f'$	$\succ\ d'\ o'\ m'\ s'$	$d\ d'\ o\ o'\ s\ s'\ f\ f'\ =$	d'	$d'\ o\ f'$	$d'\ o'\ s'$	$d'\ o\ f'$	$d'\ o'\ s'$	$d'\ o\ f'$	d'	$d'\ o'\ s'$	d'
o	$\prec$	$\succ\ d'\ o'\ m'\ s'$	$d\ o\ s$	$\prec\ d'\ o\ m\ f'$	$\prec\ o\ m$	$d\ d'\ o\ o'\ s\ s'\ f\ f'\ =$	$\prec$	$d'\ o'\ s'$	o	$d'\ o\ f'$	$d\ o\ s$	$\prec\ o\ m$
o'	$\prec\ d'\ o\ m\ f'$	$\succ$	$d\ o'\ f$	$\succ\ d'\ o'\ m'\ s'$	$d\ d'\ o\ o'\ s\ s'\ f\ f'\ =$	$\succ\ o'\ m'$	$d'\ o\ f'$	$\succ$	$d\ o'\ f$	$\succ\ o'\ m'$	o'	$d'\ o'\ s'$
m	$\prec$	$\succ\ d'\ o'\ m'\ s'$	$d\ o\ s$	$\prec$	$\prec$	$d\ o\ s$	$\prec$	$f\ f'\ =$	m	m	$d\ o\ s$	$\prec$
m'	$\prec\ d'\ o\ m\ f'$	$\succ$	$d\ o'\ f$	$\succ$	$d\ o'\ f$	$\succ$	$s\ s'\ =$	$\succ$	$d\ o'\ f$	$\succ$	m'	m'
s	$\prec$	$\succ$	d	$\prec\ d'\ o\ m\ f'$	$\prec\ o\ m$	$d\ o'\ f$	$\prec$	m'	s	$s\ s'\ =$	d	$\prec\ o\ m$
s'	$\prec\ d'\ o\ m\ f'$	$\succ$	$d\ o'\ f$	d'	$d'\ o\ f'$	o'	$d'\ o\ f'$	m'	$s\ s'\ =$	s'	o'	d'
f	$\prec$	$\succ$	d	$\succ\ d'\ o'\ m'\ s'$	$d\ o\ s$	$\succ\ o'\ m'$	m	$\succ$	d	$\succ\ o'\ m'$	f	$f\ f'\ =$
f'	$\prec$	$\succ\ d'\ o'\ m'\ s'$	$d\ o\ s$	d'	o	$d'\ o'\ s'$	m	$d'\ o'\ s'$	o	d'	$f\ f'\ =$	f'

FIGURE 3. Allen's transitivity table: $(i\ r_1\ j) \wedge (j\ r_2\ k) \supset (i\ (\mathcal{T}(r_1, r_2))\ k)$

infons, and which contain the system $\Im$ as axioms, is sound and complete with respect to the class of interval structures.

DEFINITION 3.19
An *interval theory* is a closed set of temporal infons containing $\Im$ that is based on a countable set of interval terms $\mathcal{I}$.

THEOREM 3.20
The class of interval theories based on $\mathcal{I}$ is sound and complete with respect to the class of interval structures over $\mathcal{I}$.

As an aside, because temporal infons are propositional in nature (in the sense that any such infon is defined in any interval structure in which all interval terms occurring in the infon are interpreted), a deduction theorem for interval theories can be proven. Here $\vdash$ denotes a situated consequence operation.

PROPOSITION 3.21 (Deduction theorem)
If $\Gamma \cup \{A\} \vdash B$ then $\Gamma \vdash A \supset B$ for an interval theory Γ and temporal infons A and B.

We now consider the modelling of courses of events. Following Barwise and Perry [7], the basic infons in a course of events are either events or temporal constraints of the form modelled above. By an ‘event’ in the context of plan recognition, we mean the occurrence of an action instance over a period of time. Thus each basic event infon comprises an action type together with an interval of occurrence. The interval of occurrence in turn forms part of an interval structure associated with the situation. The purpose of the interval logic is to ensure the consistency and completeness of constraints between the intervals of different events. In this way, a course of events can be regarded not as a generalized situation, but as a special type of situation where the infons are of a particular sort (and where a situation consists of just the set of positive infons it supports).

DEFINITION 3.22
A *(positive) event infon* is an infon with positive polarity of the form $A(i)$ where A is an action type and i is an interval term. A *negative event infon* is the negation of a positive event infon. An *event infon* is a positive or negative event infon.

DEFINITION 3.23
A *course of events* is a pair $\langle \sigma, \xi \rangle$ where σ is a situation consisting of a set of event infons, and ξ is an interval structure over a set of interval terms that includes all interval terms occurring in an event infon in σ.

The last condition on courses of events simply ensures that each interval associated with an action in σ is interpreted in the interval structure ξ.

The satisfaction conditions for courses of events are inherited from those for situations and interval structures.

DEFINITION 3.24
A coherent course of events $\langle \sigma, \xi \rangle$ *supports (rejects)* a basic infon A, written $\langle \sigma, \xi \rangle \models A$ ($\langle \sigma, \xi \rangle \mathrel{=\!\!|} A$), under the following conditions.

$$\begin{array}{ll} \langle \sigma, \xi \rangle \models A & \text{if } A \in \sigma \text{ for an event infon } A \\ \langle \sigma, \xi \rangle \models A & \text{if } \xi \models A \text{ for a temporal infon } A \\[1em] \langle \sigma, \xi \rangle \mathrel{=\!\!|} A & \text{if } A^{\perp} \in \sigma \text{ for an event infon } A \\ \langle \sigma, \xi \rangle \mathrel{=\!\!|} A & \text{if } \xi \mathrel{=\!\!|} A \text{ for a temporal infon } A. \end{array}$$

The satisfaction conditions for complex infons are as for situations.

The class of situated consequence operations that include the axioms of interval structures is sound and complete with respect to the class of courses of events. Note that the second axiom scheme for equality in $\Im$ is here understood to apply to all infons, e.g. allowing the event infon $A(j)$ to be inferred from the event infon $A(i)$ and the temporal infon $i = j$.

COROLLARY 3.25
The class of situated consequence operations $\vdash$ for which $\vdash \Im$ is sound and complete with respect to entailment over sets of courses of events.

3.3 Plan recognition as deduction

In this section, we show how to define a consequence operation for use in plan recognition that implicitly represents the information contained in a plan hierarchy. This provides our deductive approach to plan recognition that is the analogue of Kautz's work [28, 26, 27]. The main technical task in this section is to define a language for representing actions, and to define and characterize a class of plans (the simple plans) which are the objects being recognized in the plan recognition process. The main technical difference between our work and Kautz's is, as noted above, our use of a proposition-like language in contrast to Kautz's use of first-order logic. We also make more precise the definition of a plan hierarchy used in our theory.

The notation we use for representing actions is a slight modification of the original notation for events in situation semantics developed in Barwise and Perry [7]. We first assume a finite collection of object types e.g. *person*, *food*, etc., arranged in an abstraction hierarchy, so that types may have subtypes, e.g. *pasta* is a subtype of *food*, etc. In this way, a plan hierarchy may contain role-specializations of an action type obtained by constraining the types of one or more of its objects. We represent an action occurrence using a predicate consisting of the action's type and a list of arguments, each consisting of an object variable with the type of the object.

DEFINITION 3.26
An *action predicate* is an expression of the form $\pi(x_1 : \tau_1, \cdots, x_n : \tau_n)$, where π is an action type and each x_i is a variable denoting an object constrained to be of type τ_i.

Our representation of plan schemas is inspired by standard techniques in use in hierarchical planning systems such as NOAH [45] and NONLIN [51]. Each plan schema can be considered a frame containing the following information:

- a unique name;
- a set of role variables each with a constraining object type;
- a (possibly null) parent action of which the schema is a *subtype*;
- a (possibly empty) expansion giving the *component* actions of the schema;
- a set of temporal constraints on the intervals of the component actions.

Our hierarchies do not contain any *AnyEvent* or *End* action types.

DEFINITION 3.27
A *plan hierarchy* is an abstraction relation and decomposition relation on a finite set of action types such that (i) both relations are irreflexive, (ii) the abstraction relation is transitive, (iii)

if A is a role-specialization of B then A is abstracted by B, (iv) no action type is abstracted by two distinct action types, and (v) the temporal constraints on each action decomposition are consistent interval theories.

DEFINITION 3.28
An action type A is a *specialization* of an action type B iff B abstracts A.

DEFINITION 3.29
A plan hierarchy is *coherent* if whenever a plan schema A has a component action B, any action A' that is a specialization of A has a component action B' that is a specialization of B or is equal to B.

We only consider coherent plan hierarchies in this paper. The purpose of the condition is to ensure that when a plan is elaborated, the specialized actions added to the plan are consistent with the more abstract actions already contained in the plan. Note that although the converse of this condition might also be sensible (i.e. if all actions in A's decomposition are abstracted by all actions in B's decomposition then A is a specialization of B), it is not essential for our theory. The cooking hierarchy of Figure 1 is coherent, and incidentally, also satisfies the converse condition.

What makes plan recognition a recognition task is that it is assumed that the executed plan being observed is an instance of one that the observer has ‘seen before’. We make this precise by defining the class of such plans, which we call the *simple plans*. We then show how inference over the courses of events corresponding to this class can be characterized as a monotonic consequence operation. We assume that the simple plans correspond to the possible executions of certain *primary* actions: these correspond to the ‘high-level’ plans in the hierarchy that are assumed to be self-motivating. Also, following Kautz's approach, we assume that the class of recognizable plans is drawn from the class of plans that the observer agent can construct from the plan hierarchy (the plan library).

In some sense, plan recognition is the inverse of planning. We view the simple plans as a set of plans that can be constructed by a planner with ‘minimal effort’. In the standard planning framework, a planning agent is considered as taking some goal Q and an initial plan consisting of one high-level action that achieves Q. The agent then constructs a plan to achieve Q by extending its initial plan in a number of ways, each defining an elaboration of the plan, stopping only when all actions in the plan are executable. A simple plan is a plan that, intuitively, the observer agent has seen before, and hence is represented ‘directly’ in the plan hierarchy. We define the class of simple plans constructively in terms of the operation of the planner, but limit the complexity of plans that can be constructed: essentially a simple plan is one that can be constructed from a plan hierarchy using only the operation of specialization of actions (disallowing more complex operations such as concatenation of partial plans). The construction of a simple plan starts with a primary action and finishes with a set of terminal actions.[2]

DEFINITION 3.30
A *primary action type* is a most abstract action type (in a given plan hierarchy) that is not a component of any plan schema.

DEFINITION 3.31
A *terminal action type* (in a given plan hierarchy) is an action type that has no subtypes or all of whose subtypes are components of other plan schemas.

---

[2] The name is chosen to reflect the close analogy between plan elaboration and sentence generation, and consequently between plan recognition and parsing.

The terminal action types in the cooking hierarchy from Figure 1 are *Boil*, *Fettuccine*, *Spaghetti*, *SpaghettiMarinara*, *ChickenMarinara*, *Sauce*, *Marinara* and *WashDishes*; *Sauce* is a terminal action type because all of its subtypes (here only *Marinara*) are components of other plan schemas. Because, as we shall define below, plan elaboration stops with a set of terminal actions, this definition goes against Kautz's assumption of complete knowledge of subtypes in the abstraction hierarchy. But there is some intuitive support for this condition: since all subtypes of the action are components of other plan schemas, they must be subtypes that are 'special' in some way, hence there must be other 'normal' ways of executing the action, so the abstraction hierarchy cannot be complete. In the cooking hierarchy, our approach gives intuitively correct results, for the assumption that the hierarchy is complete implies that the action of making fettuccine can never occur. To see this, suppose *Fettuccine*(x) is observed. Then with Kautz's theory, *Noodles*(x) can be inferred using (ABS), then *PastaDish*$(y) \vee$ *SpaghettiMarinara*(y) using (CUA), where $x = s_1(y)$, but then *SpaghettiMarinara*(y) follows using (EXA), giving *Marinara*$(s_2(y))$ and *Sauce*$(s_2(y))$ as well as *Spaghetti*(x) by (DEC), which contradicts *Fettuccine*(x) using (DJA). But if *Fettuccine* occurs in some plan, it must be a specialization of *PastaDish*, hence this plan must include an action of type *Sauce* that cannot be *Marinara*. The upshot is that there is information in the plan hierarchy itself that shows that the hierarchy is incomplete, hence it is irrational of the observer agent to make the complete knowledge assumptions required of Kautz's theory. Thus some weaker condition is needed, and this is the reason for the clause that means an action is terminal if there are possibly unknown ways of executing the action. Of course, in this example, the exhaustiveness assumption enabling *SpaghettiMarinara* to be inferred from *PastaDish* is also problematic, and blocking this inference is achieved through the definition of the simple plans.

DEFINITION 3.32

A *simple plan* (derived from a given plan hierarchy) is a set of event and temporal infons that can be constructed as follows. Initially, the plan P consists of a positive event infon $A(i_0)$ where A is a primary action and i_0 is an arbitrary interval constant, negative event infons $\neg A'(i_0)$ for the other primary actions A', and event infons $A'(s_j(i_0))$ for each action A' that abstracts A_j, the jth component of A. Then the plan is constructed by repeatedly elaborating a nonterminal action, stopping only when all such actions have been elaborated. To elaborate a nonterminal action A such that $A(i) \in P$, first it is decided whether to specialize A or each of A's nonterminal components. Then, for each of the chosen actions, one of its role-specializations or subtypes B that is not a component of any plan schema is chosen and the infons $B(i)$ and $\neg B'(i)$ are added to P for all other role-specializations or subtypes B'. In addition, whenever an event infon $C(i)$ is added to P where C has a nonempty decomposition, $C'(s_j(i))$ is added to P for all actions C' that abstract C_j, the jth component of C (and if A was chosen for specialization, any of A's components automatically specialized by this process are not considered further). Finally, any constraints on the intervals i and the $s_j(i)$ are added to P.

Thus a simple plan is maximally specific in the sense that all nonterminal actions in the plan are specialized as far as possible into a set of terminal actions. There are a number of subtleties in the above definition. One relates to the coherence condition on plan hierarchies, in that if an action A with decomposition $A_1, \cdots, A_n$ is specialized to a subtype B, any actions added to the plan as components of B are specializations of the A_i, so that our analogue of the (EXA) rule with A_i as antecedent is satisfied: this is why such A_i need not be considered further. Another subtlety relates to the condition that the subtype B cannot be a component of another plan schema, which means that B can only be in a simple plan if one of the schemas containing B as a component is also in the plan, so that our analogue of the (CUA) rule with B as antecedent is satisfied. A third subtlety relates to the condition that only the nonterminal actions are specialized. This handles the problem with *Sauce* in the cooking hierarchy as described above, since there is no sure way to specialize this action according to the knowledge in the plan hierarchy. What this means to the observer agent is that the action would be specialized by the planning agent in some unknown way (or possibly to *Marinara*, but then only if *SpaghettiMarinara* is the plan being followed). A fourth subtlety relates to the condition that either A or all of the nonterminal A_i are specialized, enabling a simple plan containing *Fettuccine* to be generated.

We now illustrate the idea of simple plans using the cooking hierarchy. The only primary action types are *PrepareMeal* and *WashDishes*. The simple plans correspond to the sets of terminal actions that form an elaboration of a single primary action. The following are all the simple plans derived from this hierarchy (omitting negative event infons and interval symbols for simplicity). The first simple plan below would not be recognized using Kautz's theory.

PrepareMeal, PastaDish, Noodles, Fettuccine, Sauce, Boil
PrepareMeal, PastaDish, SpaghettiMarinara, Noodles, Spaghetti, Sauce, Marinara, Boil
PrepareMeal, MeatDish, ChickenMarinara, Sauce, Marinara
WashDishes

However, the following sets of actions are *not* simple plans derived from this hierarchy.

PrepareMeal, PastaDish, Noodles, Sauce, Boil
PrepareMeal, PastaDish, Noodles, Spaghetti, Sauce, Boil
PrepareMeal, PastaDish, Noodles, Sauce, Marinara, Boil
PrepareMeal, PastaDish, Noodles, Spaghetti, Sauce, Marinara, Boil
PrepareMeal, PastaDish, Noodles, Fettuccine, Sauce, Marinara, Boil

Thus the notion of simple plans embodies some assumptions of complete knowledge and specificity, following Kautz's theory of plan recognition. The complete knowledge assumptions derive from the condition that a simple plan must contain a (known) specialization of an abstract action whenever possible. For example, an observed agent cannot make noodles *per se*; it must do this by making a particular (known) type of noodles, spaghetti or fettuccine. The specificity assumptions derive from the condition that an action A can only be specialized to an action that is not a component of another plan schema, so if any other subtype of A is in a simple plan, it must be because the plan schema of which it is a component is also in the plan. For example, the observed agent does not make spaghetti as part of some unknown type of pasta dish: it must be as a component of *SpaghettiMarinara*; however, in contrast, since there is no known special type of pasta dish that includes *Fettuccine*, the agent observed to be making fettuccine must be making some unknown type of pasta dish. Thus a specificity assumption can also be taken as an assumption of complete knowledge about the motivations for performing an action.

We now characterize the simple plans (for a given plan hierarchy) as prime sets of infons that are closed under a set of postulates extending those for situated consequence operations. As with the interval theories, we assume that each such plan theory is based on a countable set of interval terms $\mathcal{I}$. In the following postulates, an interval variable may be replaced by any other interval term occurring in the set $\mathcal{I}$ (unless otherwise stated, terms such as i, i_1, etc. are interval variables).

The following postulates come directly from the meaning of the plan hierarchy: these are the postulates for abstraction (ABS), decomposition (DEC), and role-specialization (SPEC), which formalize the conditions on any valid plan. First, if an action token a in the plan is of type A, the plan also contains an event which is the instance a with the type of A's parent in the abstraction hierarchy and the same interval of occurrence.

(ABS) $A(i) \vdash A'(i)$, where A' directly abstracts A.

Second, if an expandable action is in a simple plan, all of its component actions are also in the plan, together with the associated constraints on the intervals of occurrence.

(DEC) $A(i) \vdash A_1(s_1(i)) \wedge \cdots \wedge A_n(s_n(i)) \wedge \kappa$, where the A_j are the actions in the expansion of A, the s_j are functions giving the jth step in the decomposition, and κ is a conjunction of temporal infons relating i and $s_1(i), \cdots, s_n(i)$.

Finally, any role-specialization A_r of an action A is a specialization of that action, thus if an instance of A_r occurs in a simple plan, it must also be an instance of A.

(SPEC) $A_r(i) \vdash A(i)$, for every role-specialization A_r of A.

The above postulates hold of any hierarchical plan. We now consider postulates that apply specifically to simple plans. We use XOR to mean n-ary exclusive or, i.e. $\mathsf{XOR}(\alpha_1, \cdots, \alpha_n)$ abbreviates $(\alpha_1 \vee \cdots \vee \alpha_n) \wedge \wedge_{i \neq j} \neg(\alpha_i \wedge \alpha_j)$. Thus a situation supports $\mathsf{XOR}(\alpha_1, \cdots, \alpha_n)$ iff it supports exactly one of the α_i and rejects all the others. First, exhaustiveness (EXH) says that every simple plan contains exactly one primary action over one associated interval of occurrence: we choose the constant denoting this interval arbitrarily.

(EXH) $\vdash \mathsf{XOR}(A_1(i_0), \cdots, A_n(i_0))$, where the A_j are the primary action types, and i_0 is an arbitrarily chosen interval constant.

(EXH) $A_j(i_1) \wedge A_j(i_2) \vdash i_1 = i_2$, where A_j is a primary action.

Second, the postulates for subtypes (SUB) state that an instance of exactly one subtype of a nonterminal action type or of each of its nonterminal components can be contained in a simple plan.

(SUB) $A(i) \vdash \mathsf{XOR}(A_1(i), \cdots, A_n(i), B_1^k(s_k(i)) \wedge \cdots \wedge B_{m_k}^k(s_k(i)))$, for a nonterminal action type A and each of A's nonterminal components B^k (A's kth step), where the A_j are all the subtypes of A and the B_l^k are the subtypes of B^k that are not components of any plan schema.

Finally, the postulates for components (COMP) state that if an action is included in a simple plan, at least one of the instances of the plan schemas containing it as a component must also be included in the plan—furthermore, the interval of occurrence of the schema must be correctly related to that of the component.

(COMP) $A(i) \vdash \cdots \vee (A_j(p(i)) \wedge (s_{k_j}(p(i)) = i)) \vee \cdots$, where the A_j are the most abstract plan schemas that include a specialization A' of A as the k_jth component action, and $p(i)$ gives the interval of occurrence of this plan instance.

We now show that sets of infons closed under the above postulates correspond to the simple plans, hence the courses of events satisfying the above postulates are the semantic analogues of the simple plans.

DEFINITION 3.33
The *postulates for plan recognition* $\mathcal{H}$ derived from a plan hierarchy H are the set of all (ABS), (DEC), (SPEC), (EXH), (SUB) and (COMP) postulates obtained from H.

DEFINITION 3.34
A *plan recognition inference operation* derived from a plan hierarchy H is a situated consequence operation $\vdash$ that satisfies the postulates for plan recognition derived from H and for which $\vdash \Im$ for some countable set of interval terms $\mathcal{I}$.

DEFINITION 3.35
A *plan theory* based on a plan recognition inference operation $\vdash$ derived from a plan hierarchy H is a prime consistent set of infons that is closed under $\vdash$.

DEFINITION 3.36
A course of events $\langle\sigma,\xi\rangle$ is a *realization* of a simple plan P based on a plan hierarchy H if (i) $\sigma \models E$ for all event infons $E \in P$, (ii) $\xi \models t$ for all temporal infons $t \in P$, and (iii) $\sigma \models A(i)$ for a positive event infon $A(i)$ only if $A(i) \in P$.

THEOREM 3.37
The plan theories based on the plan recognition inference operations derived from a plan hierarchy H correspond to the realizations of the simple plans based on H (up to renaming of interval symbols).

Note that in contrast to the first-order models used in Kautz's theory, there are typically very few negative event infons supported by a course of events corresponding to a simple plan: these come from the use of XOR in certain postulates. For example, if $A(i)$ is supported where A is a primary action, then by the first (EXH) postulate, $\neg A'(i)$ is supported for any other primary action A'. However, though Kautz stresses the importance of the similar disjointedness assumptions for plan recognition, there seem to be 'enough' negative event infons supported by such courses of events to force the desired conclusions to follow. Note also, however, that the above definition allows for a realization of a simple plan to support negative event infons that are not direct consequences of the event infons corresponding to actions in the plan and the conditions on simple plan construction, but the existence of such courses of events does not prevent any desired conclusion from being validly drawn.

COROLLARY 3.38
For any plan hierarchy H, the class of plan recognition inference operations derived from H is sound and complete with respect to entailment over the courses of events that are the realizations of the simple plans derived from H.

This means that if, for an observation A and prediction B, $A \vdash B$ where $\vdash$ is a plan recognition inference operation derived from a plan hierarchy H, B holds in every realization of every simple plan derived from H that supports A, and thus B is validly inferred from A. This is the basic result underpinning our monotonic theory of plan recognition.

To illustrate the theory, we now give the complete formalization of the cooking hierarchy (which, for us, does not include the *AnyEvent* and *End* action types). We assume the consequence operation is defined over a set of interval terms that includes $s_1(i), \cdots, s_5(i)$ and $p(i)$ whenever it includes an interval term i.

Abstraction

$\mathit{PastaDish}(i) \vdash \mathit{PrepareMeal}(i)$
$\mathit{MeatDish}(i) \vdash \mathit{PrepareMeal}(i)$
$\mathit{SpaghettiMarinara}(i) \vdash PastaDish(i)$
$\mathit{ChickenMarinara}(i) \vdash \mathit{MeatDish}(i)$
$\mathit{Spaghetti}(i) \vdash \mathit{Noodles}(i)$
$\mathit{Fettuccine}(i) \vdash \mathit{Noodles}(i)$

Decomposition

$\mathit{PastaDish}(i) \vdash \mathit{Noodles}(s_1(i)) \wedge \mathit{Sauce}(s_2(i)) \wedge \mathit{Boil}(s_3(i))$
$\mathit{SpaghettiMarinara}(i) \vdash \mathit{Spaghetti}(s_1(i)) \wedge \mathit{Marinara}(s_2(i)) \wedge \mathit{Boil}(s_3(i))$
$\mathit{ChickenMarinara}(i) \vdash \mathit{Marinara}(s_5(i))$

Exhaustiveness

$\vdash \mathsf{XOR}(\mathit{PrepareMeal}(i_0), \mathit{WashDishes}(i_0))$

Subtypes

$\mathit{PrepareMeal}(i) \vdash \mathsf{XOR}(\mathit{PastaDish}(i), \mathit{MeatDish}(i))$
$\mathit{PastaDish}(i) \vdash \mathsf{XOR}(\mathit{SpaghettiMarinara}(i), \mathit{Fettuccine}(s_1(i)))$
$\mathit{MeatDish}(i) \vdash \mathit{ChickenMarinara}(i)$
$\mathit{Noodles}(i) \vdash \mathsf{XOR}(\mathit{Spaghetti}(i), \mathit{Fettuccine}(i))$

Components

$\mathit{Noodles}(i) \vdash \mathit{PastaDish}(p(i)) \wedge (s_1(p(i)) = i)$
$\mathit{Spaghetti}(i) \vdash \mathit{SpaghettiMarinara}(p(i)) \wedge (s_1(p(i)) = i)$
$\mathit{Sauce}(i) \vdash (\mathit{PastaDish}(p(i)) \wedge (s_2(p(i)) = i)) \vee (\mathit{ChickenMarinara}(p(i)) \wedge (s_5(p(i)) = i))$
$\mathit{Marinara}(i) \vdash$
    $(\mathit{SpaghettiMarinara}(p(i)) \wedge (s_2(p(i)) = i)) \vee (\mathit{ChickenMarinara}(p(i)) \wedge (s_5(p(i)) = i))$
$\mathit{Boil}(i) \vdash \mathit{PastaDish}(p(i)) \wedge (s_3(p(i)) = i)$

We can now give a reconstruction of the example of plan recognition taken from Kautz [26]. In this example, the initial observation is of an agent making marinara sauce. We invent an arbitrary constant i_1 for the interval of occurrence of this event.

(1) $\mathit{Marinara}(i_1)$ by observation

It can be inferred that the agent is preparing a meal, as follows.

(2) $\mathit{SpaghettiMarinara}(p(i_1)) \vee \mathit{ChickenMarinara}(p(i_1))$ by (1) and (COMP)
(3) $\mathit{PastaDish}(p(i_1)) \vee \mathit{MeatDish}(p(i_1))$ by (2) and (ABS)
(4) $\mathit{PrepareMeal}(p(i_1))$ by (3) and (ABS)

The second observation is of an agent making fettuccine or spaghetti (we invent the constant i_2 for the interval of occurrence of this event). It can now be inferred that the agent will make spaghetti, as follows.

(5) $\mathit{Fettuccine}(i_2) \vee \mathit{Spaghetti}(i_2)$ by observation
(6) $\mathit{Noodles}(i_2)$ by (5) and (ABS)
(7) $\mathit{PastaDish}(p(i_2))$ by (6) and (COMP)
(8) $\mathit{PrepareMeal}(p(i_2))$ by (7) and (ABS)
(9) $p(i_1) = p(i_2)$ by (4), (8) and (EXH)

(10) *PastaDish*$(p(i_1))$ by (7) and (9)
(11) $\neg$*MeatDish*$(p(i_1))$ by (10) and (EXH)
(12) *SpaghettiMarinara*$(p(i_1)) \vee$ *MeatDish*$(p(i_1))$ by (2) and (ABS)
(13) *SpaghettiMarinara*$(p(i_1))$ by (11) and (12)
(14) *Spaghetti*$(s_1(p(i_1)))$ by (13) and (DEC)

This inference can be drawn because there is just one simple plan, *SpaghettiMarinara*, that contains both making marinara and making fettuccine or spaghetti, and this plan contains the action of making spaghetti.

The inferences allowed under our theory differ from those allowed using Kautz's approach. For example, Kautz's plan recognition theory derived from the cooking hierarchy includes the analogue of the following postulate as an exhaustiveness assumption.

PastaDish$(i) \vdash$ *SpaghettiMarinara*(i)

We have the following as a subtype postulate.

PastaDish$(i) \vdash \mathsf{XOR}($*SpaghettiMarinara*$(i),$ *Fettuccine*$(s_1(i)))$

In this case, the predictions made under our theory (from an observation *PastaDish*(i_1)) are weaker than those permitted under Kautz's, but our formulation seems intuitively correct: an agent making a pasta dish could be making fettuccine and hence a pasta dish of unknown type. Of course, this kind of inference occurs only when the 'observation' includes an abstract action. Kautz's theory also includes the analogue of the following postulate, which is not contained in our theory.

Sauce$(i) \vdash$ *Marinara*(i)

As described above, this leads to the conclusion that there can be no simple plan containing *Fettuccine*, since any such plan must include *Sauce* but not *Marinara*, directly conflicting with this postulate.

Kautz's theory handles multiple observations more directly than our theory, since we assume all observations relate to a *single* simple plan, whereas with Kautz's approach, the number of *End* actions is minimized, so that multiple plans may be recognized simultaneously. Using our theory, the existence of multiple concurrently executed simple plans is triggered by an observation that contradicts a prior prediction, but we think that this circumstance is better handled using extra-logical means than directly through the logic of plan recognition. It is certainly this aspect of Kautz's theory that leads to the technical complications in stating the minimum cardinality assumptions.

4 A nonmonotonic logic for plan recognition

From a logical point of view, one basic problem with Kautz's approach to plan recognition, as noted by Kautz himself, e.g. Kautz and Allen [28, p. 33], is that the resulting theory is monotonic. This means, in effect, that an observer treats all possible plans consistent with an observation as equally plausible. Thus stronger methods are needed to model the effect of the prior plausibility of plans on the formalization of plan recognition. In recent years, there have been various attempts to model the prior plausibility of plans using Bayesian probability theory [24, 43, 1], and using Dempster–Shafer theory [11, 9, 8]. From our perspective, these approaches are too specific in committing to a particular formalism. We aim to classify plan recognition inference as a particular kind of nonmonotonic consequence operation, which requires (only) determining the valid inference patterns underlying plan recognition: this is the approach to nonmonotonic reasoning initiated by Gabbay [17]. The approach exploits the connection between plan recognition and indicative conditionals, illustrated in the close correspondence between the inference of a prediction B from an observation A, and an indicative conditional 'if A then B'. Gabbay's approach thus borrows heavily from work on conditional logic [49, 35].

In this section, we present a modelling of a class of nonmonotonic consequence operations that are based on preferential entailment, following Shoham [47]. In Shoham's approach, B follows from A if B holds in all the 'most preferred' models of A, and thus the semantic framework requires an ordering on the elements of the models. While Shoham uses the term 'preference ordering' to describe these orderings, we use the term 'plausibility ordering' to indicate, in a more neutral way, an ordering representing the relative prior plausibility of two plans or situations. This section proceeds by first defining a class of nonmonotonic consequence operations that corresponds to preferential entailment over ordered structures of situations, then showing how to model the dynamics of an agent's epistemic states, where each epistemic state is a belief set together with a nonmonotonic consequence operation derived from a plausibility ordering on the simple plans. This constitutes our nonmonotonic theory of plan recognition. In order to develop the theory, we need to make the assumption of finite representability of the nonmonotonic consequence operation, meaning that our nonmonotonic theory of plan recognition is formally more restricted than our monotonic theory.

4.1 Preferential entailment over situations

The formal basis for preferential entailment is a plausibility ordering on some structure, here a structure of situations each modelling a simple plan. Under preferential entailment, an inference $A \mid\!\sim B$ holds in a structure of situations if B holds at all *minimal* A-supporting situations in the model. Thus the conditions on the ordering that determine minimality are crucial in reflecting properties of the consequence operation. The natural requirements on a plausibility ordering are reflexivity, antisymmetry and transitivity, so that the ordering forms a partial order, as in Shoham [47]. But in addition, we require a property known as almost-connectedness (see below), which in the context of the other conditions, is similar in power to totality, i.e. the condition that all pairs of situations are comparable. Some support for this strong condition follows from the intuition that if a situation supports $A \vee B$, it must support either A or B. To see this, suppose an agent observes $A \vee B$. The observed agent must be executing a plan that contains either A or B, and the corresponding minimal $A \vee B$-supporting situations also satisfy $A \vee B$ and hence either A or B (or both). Assuming that these minimal situations constitute *all* the minimal A-supporting situations (when satisfying A) and *all* the minimal B-supporting situations (when satisfying B), comparability ensues because, following Stalnaker [49], given any two situations σ_1 and σ_2 that are minimal A and B-supporting situations (respectively), a total pre-order of situations $\preceq$ can be defined by setting $\sigma_1 \preceq_\sigma \sigma_2$ iff σ_1 is a minimal $A \vee B$-supporting situation. It is also worth noting that comparability, while controversial in conditional logic, is built in to those approaches based on assigning a plausibility valuation, such as a prior probability or degree of belief, to each possible plan.

We now proceed to the technical definitions.

DEFINITION 4.1
An *almost-connected partial order* on a set S is a binary relation $\preceq$ that satisfies the following conditions (for all α, β, $\gamma \in S$):

(i) reflexive, i.e. $\alpha \preceq \alpha$;
(ii) antisymmetric, i.e. $\alpha \preceq \beta$ and $\beta \preceq \alpha$ implies $\alpha = \beta$;
(iii) transitive, i.e. $\alpha \preceq \beta$ and $\beta \preceq \gamma$ implies $\alpha = \gamma$;
(iv) almost-connected, i.e. $\alpha \preceq \beta$ implies $\alpha = \beta$, $\alpha \preceq \gamma$ or $\gamma \preceq \beta$.

Note that the condition of almost-connectedness implies that for any two elements α and β of S which are incomparable (i.e. $\alpha \not\preceq \beta$ and $\beta \not\preceq \alpha$), the set of elements γ such that $\alpha \preceq \gamma$ is the same as the set of elements γ such that $\beta \preceq \gamma$, and the set of elements γ such that $\gamma \preceq \alpha$ is the same as the set of elements γ such that $\gamma \preceq \beta$. Thus the ordering induces a set of equivalence classes in S: for each element of $s \in S$, there is an equivalence class containing s and the elements with which s is incomparable. This is already very similar to the 'sphere' semantics of Lewis [35] based on total pre-orders of possible worlds, and moreover, this connection can be made precise by noting that, given an almost-connected partial order $\preceq$, the ordering $\preceq'$ defined by $\alpha \preceq' \beta$ iff $\beta \not\preceq \alpha$ or $\alpha = \beta$ is a total pre-order.

The following provides the formal definition of an ordered structure of situations.

DEFINITION 4.2
A situation σ is *minimal* in a set of situations Σ ordered by $\preceq$ if $\sigma \in \Sigma$ and for all $\sigma' \in \Sigma$, $\sigma' \preceq \sigma$ implies $\sigma' = \sigma$.

DEFINITION 4.3
An ordering $\preceq$ on a set of situations Σ is *finitarily stoppered* if for any infon A, whenever some situation in Σ supports A, there is at least one and only a finite number of situations in Σ that support A and which are minimal amongst all A-supporting situations.

DEFINITION 4.4
An *ordered structure of situations* $\langle \Sigma, \preceq \rangle$ is a set Σ of finite coherent situations together with an ordering relation $\preceq$ on Σ that satisfies the following conditions:

(i) $\preceq$ is an almost-connected partial order on Σ;
(ii) $\preceq$ is finitarily stoppered.

Condition (i) represents our standard structural requirements on the ordering of elements in Σ. The clause of condition (ii) requiring the existence of minimal A-supporting situations in any model containing an A-supporting situation is known as 'finitarily stopperedness' in the terminology of Makinson [38]: it is a weakened version of the descending chain condition, and is related to the 'limit assumption' of Lewis [35]. Our definition is stronger in requiring only a finite number of such situations—this is the counterpart of the finite representability of the consequence operation. Note also that we do not require Σ to be nonempty.

We can now define a notion of preferential entailment for ordered structures of situations, by adapting the definition from Shoham [47], which has been used for formalizing non-monotonic consequence operations by Kraus, Lehmann and Magidor [29], Lehmann and Magidor [33], Freund [16] and Gärdenfors and Makinson [21]. Essentially, A preferentially entails B if B holds in all situations minimal amongst those supporting A. Note that if there are no minimal A-supporting situations, A preferentially entails B for any infon B.

DEFINITION 4.5
Given an ordered structure of situations $\langle \Sigma, \preceq \rangle$, let $[A]$ denote the set of situations in Σ that

support A. For a situation $\sigma \in \Sigma$, the set of *minimal A-supporting situations in* Σ, *min*(A), is defined as $\{\sigma'$: σ' is minimal in $[A]$ ordered by $\preceq\}$.

DEFINITION 4.6
An infon A *preferentially entails* an infon B with respect to an ordered structure of situations $\mathcal{M} = \langle \Sigma, \preceq \rangle$, written $A \mathrel{|\!\!\approx}_{\mathcal{M}} B$, if every minimal A-supporting situation in Σ supports B.

We now characterize preferential entailment as a set of properties of the consequence operation so generated. The precise definitions are as follows.

DEFINITION 4.7
An *almost-rational consequence operation* $\mid\!\sim$ based on a monotonic operation $\vdash$ is a consequence operation on infons that satisfies the following postulates.

If $A \vdash B$ then $A \mid\!\sim B$ (Supraclassicality)
If $A \mid\!\sim B \wedge \neg B$ then $A \vdash B \wedge \neg B$ (Consistency Preservation)
If $A \mid\!\sim B$ and $B \vdash C$ then $A \mid\!\sim C$ (Right Weakening)
If $A \mid\!\sim B$ and $B \mid\!\sim A$ then $A \mid\!\sim C$ implies $B \mid\!\sim C$ (Reciprocity)
If $A \mid\!\sim B$ and $A \mid\!\sim C$ then $A \wedge B \mid\!\sim C$ (Cautious Monotony)
If $A \mid\!\sim B$ and $A \wedge B \mid\!\sim C$ then $A \mid\!\sim C$ (Cut)
If $A \mid\!\sim B$ and $A \mid\!\sim C$ then $A \mid\!\sim B \wedge C$ (And)
If $A \mid\!\sim C$ and $B \mid\!\sim C$ then $A \vee B \mid\!\sim C$ (Or)
If $A \vee B \not\mid\!\sim B$ and $A \vee B \mid\!\sim C$ then $A \mid\!\sim C$ (Almost Rational Monotony)

The names of the postulates come from the survey article by Makinson [38]. The condition (Supraclassicality) refers to the fact that $\mid\!\sim$ includes all inferences based on the 'classical' consequence operation $\vdash$, although where in earlier work this has always been propositional logic, here $\vdash$ refers to a situated consequence operation. The postulate (Cautious Monotony) is a weakened form of monotonicity analogous to a commonly valid axiom of conditional logics, e.g. Stalnaker [49] and Lewis [35]; its formulation here is due to Gabbay [17]. The postulates (Right Weakening), (Cut), (And) and (Or) are analogues of those for situated consequence operations. The final postulate, which we call (Almost Rational Monotony), is a version of the postulate of (Rational Monotony) used by Lehmann and Magidor [33] to characterize total pre-orders of propositional models, and as such, corresponds to the (CV) axiom in Lewis's conditional logic. However, our postulate is considerably weaker than (Rational Monotony), primarily due to the weaker properties of consequence operations for entailment over situations as compared to the propositional calculus. However, almost-rational consequence operations do satisfy (Disjunctive Rationality), another postulate introduced by Kraus, Lehmann and Magidor [29] that is weaker than (Rational Monotony). The additional rules are defined as follows.

If $A \not\mid\!\sim \neg B$ and $A \mid\!\sim C$ then $A \wedge B \mid\!\sim C$ (Rational Monotony)
If $A \vee B \mid\!\sim C$ then $A \mid\!\sim C$ or $B \mid\!\sim C$ (Disjunctive Rationality)

PROPOSITION 4.8
Any almost-rational consequence operation $\mid\!\sim$ satisfies the following properties.

$A \mid\!\sim A$ (Reflexivity)
If $A \dashv\vdash B$ then $A \mid\!\sim C$ iff $B \mid\!\sim C$ (Left Logical Equivalence)
If $A \vee B \mid\!\sim C$ then $A \mid\!\sim C$ or $B \mid\!\sim C$ (Disjunctive Rationality)

Soundness and completeness of the class of almost-rational consequence operations can be proven. The completeness proof involves showing that for any such consequence operation, there is an ordered structure of situations such that preferential entailment in the structure corresponds to the inference operation. The proof works by determining a set of situations for each set of nonmonotonic consequences of an infon, i.e. each set of the form $\{B : A \mathrel{|\!\sim} B\}$ for some A. This method works, however, only when each such set is logically equivalent to a finite set of infons: this is the reason for the finitary constraints on the consequence operation, condition (ii) above.

DEFINITION 4.9
Two sets of infons Γ and Δ are *logically equivalent*, denoted $\Gamma \equiv \Delta$, if Γ and Δ are provably equivalent under every situated consequence operation (or equivalently, are provably equivalent under the weakest situated consequence operation, i.e. that satisfying only the postulates in Definition 3.10 and their implications).

DEFINITION 4.10
A consequence operation on infons $\mathrel{|\!\sim}$ based on a monotonic operation $\vdash$ is *finitely representable* if for each set of infons Γ and infon A, the sets $\{B : \Gamma \vdash B\}$ and $\{B : A \mathrel{|\!\sim} B\}$ are logically equivalent to finite sets of infons.

THEOREM 4.11
The class of finitely representable almost-rational consequence operations is sound and complete with respect to preferential entailment over ordered structures of situations. That is, (soundness) the consequence operation $\mathrel{|\!\sim}_{\mathcal{M}}$ defined by preferential entailment from an ordered structure of situations $\mathcal{M}$ by $A \mathrel{|\!\sim}_{\mathcal{M}} B$ iff $A \mathrel{|\!\!\approx}_{\mathcal{M}} B$ is a finitely representable almost-rational consequence operation based on some situated consequence operation $\vdash$, and (completeness) given a finitely representable almost-rational consequence operation $\mathrel{|\!\sim}$ based on a situated consequence operation $\vdash$, there is an ordered structure of situations $\mathcal{M} = \langle \Sigma, \preceq \rangle$ such that $\Gamma \models_{\Sigma} A$ iff $\Gamma \vdash A$ and $A \mathrel{|\!\!\approx}_{\mathcal{M}} B$ iff $A \mathrel{|\!\sim} B$.

4.2 Plan recognition as nonmonotonic inference

In this section, we formalize a view of plan recognition as nonmonotonic inference based on a plausibility ordering on the simple plans. The first step is to put together the characterization of the ordered structures of situations in terms of almost-rational consequence operations with the characterization of the simple plans in terms of a subclass of situated consequence operations (those satisfying the postulates $\mathcal{H}$ derived from a given plan hierarchy H and supporting the interval axioms $\Im$). This gives the following result, which forms the basis of our nonmonotonic theory of plan recognition. Note that in order to satisfy the finitary constraints on almost-rational consequence operations, the set of interval terms $\mathcal{I}$ occurring in the temporal infons in a plan recognition inference operation must be finite.

COROLLARY 4.12
For any plan hierarchy H, the class of finitely representable almost-rational consequence operations based on a plan recognition inference operation derived from H is sound and complete with respect to preferential entailment over the ordered structures of courses of events that are the realizations of the simple plans derived from H.

Using a nonmonotonic theory of plan recognition (based on some plausibility ordering), the observer agent can make stronger predictions than with a monotonic theory, but this introduces the converse problem of retracting predictions that conflict with prior observations.

Thus what is needed to complete the theory of plan recognition is an account of the dynamics of the observer agent's epistemic states, each epistemic state consisting of a set of beliefs representing the agent's current observations and predictions, and a nonmonotonic consequence operation derived from the current plausibility ordering on simple plans. In this way, viewing plan recognition as nonmonotonic inference involves a connection to belief revision, e.g. as construed by Gärdenfors [18]. Moreover, to account for multiple observations, iterated revisions must be allowed. These issues are the subject of the remainder of this section.

To define a revision operation for a plan recognition agent, we assume that there is an epistemological difference between the beliefs the observer agent has as the result of observation and those held as the result of prediction. We can partially capture this difference by assuming that all the agent's observations are veridical, hence cannot be retracted on the basis of further observations, but that predictions are subject to revision on the acquisition of additional information. We take the initial belief set of the observer agent to represent the assumption that the observed agent is performing one of the most plausible simple plans (so the agent has beliefs that do not follow purely from the postulates derived from the plan hierarchy). To model this formally, we define an epistemic state of the observer agent to consist of a set of beliefs K closed under a situated consequence operation $\vdash$, together with a nonmonotonic consequence operation $\mid\!\sim$ based on $\vdash$ representing the revisions the agent would perform on the basis of new observations, i.e. $A \mid\!\sim B$ iff the agent would predict B given the observation A (and, of course, does so if A is actually observed—this last condition is known as the Ramsey test in the literature on belief revision). Our work thus follows established connections between belief revision and nonmonotonic reasoning [39], especially that showing connections between belief revision operations and expectations [19, 21].

DEFINITION 4.13
A *belief system* $\langle \mathcal{K}, \vdash \rangle$ is a set of belief sets $\mathcal{K}$ together with a monotonic consequence operation $\vdash$ such that each belief set in $\mathcal{K}$ is a set of infons closed under $\vdash$.

DEFINITION 4.14
An *epistemic state* $\langle K, \mid\!\sim \rangle$ in a belief system $\langle \mathcal{K}, \vdash \rangle$ consists of a belief set $K \in \mathcal{K}$ and an almost-rational consequence operation $\mid\!\sim$ such that $K = \{A : \mathit{true} \mid\!\sim A\}$, where *true* is a special infon supported by all situations.

Here $\mathcal{K}$ is typically the class of theories determined by sets of simple plans derived from a plan hierarchy H, and $\vdash$ is a plan recognition inference operation derived from H. The initial belief set is the set of infons holding at all minimal situations in a plausibility ordering on the simple plans, and to accept an observation A, the belief set is revised to consist of the set of infons holding in $\mathit{min}(A)$. Moreover, since observations are taken to be veridical, it is straightforward to define a revision operation on plausibility orderings, making the belief revision operation iterative: essentially, all simple plans not supporting a given observation are discarded, giving a plausibility ordering restricted to the situations supporting A in which the agent's beliefs are those infons holding in all minimal situations in the new ordering. This is related to the definition of conditionalization in work on update in Bayesian probability theory, applied in a belief revision context by Rott [44] — see the discussion in Nayak [41].

DEFINITION 4.15
The *revision* of an ordered structure of situations $\mathcal{M} = \langle \Sigma, \preceq \rangle$ by an infon A is the ordered structure of situations $\mathcal{M}^*_A = \langle \Sigma', \preceq' \rangle$ defined by setting $\Sigma' = \Sigma \cap [A]$ and $\sigma_1 \preceq' \sigma_2$ iff $\sigma_1 \preceq \sigma_2$ when σ_1 and σ_2 are contained in Σ' (i.e. $\preceq'$ is $\preceq$ restricted to the set Σ').

The above definition induces an iterative revision operation on epistemic states: to revise a belief set K and consequence operation $\mid\!\sim$ by A, the new belief set K^*_A is the set of infons supported by all situations in *min*(*true*) in Σ', and the new consequence operation $\mid\!\sim_A$ is that determined by the ordering $\preceq'$ on the elements of Σ'; this is related to the operation $\mid\!\sim$ associated with K by the condition $B \mid\!\sim_A C$ iff $A \wedge B \mid\!\sim C$ (conditionalization, condition (C) below). It is also clear that K^*_A so defined satisfies the Ramsey test.

DEFINITION 4.16
A *revision operation on epistemic states* is a function from epistemic states $\langle K, \mid\!\sim \rangle$ and infons A to epistemic states $\langle K^*_A, \mid\!\sim_A \rangle$ that satisfies the following postulates.

(RT) $B \in K^*_A$ iff $A \mid\!\sim B$.
(C) $B \mid\!\sim_A C$ iff $A \wedge B \mid\!\sim C$.

The class of revision operations on epistemic states in finitely representable belief systems characterizes the class of revision operations on ordered structures of situations.

DEFINITION 4.17
A belief system $\langle \mathcal{K}, \vdash \rangle$ is finitely representable if every belief set in $\mathcal{K}$ is logically equivalent to a finite set of infons.

DEFINITION 4.18
A *revision operation on ordered structures of situations* based on a situated consequence operation $\vdash$ is a revision function defined for every model $\mathcal{M}$ whose entailment relation is based on $\vdash$.

COROLLARY 4.19
The class of revision operations on a finitely representable belief system $\langle \mathcal{K}, \vdash \rangle$ is sound and complete with respect to the class of revision operations on ordered structures of situations based on $\vdash$.

An alternative formulation of epistemic states and their revision operations can be given in terms of belief revision functions. Here, an epistemic state is defined to consist of a belief set K and a belief revision function $*$ (with respect to K) mapping infons to belief sets; the belief revision function corresponds to a nonmonotonic consequence operation using the Ramsey test. The revision operation on epistemic states produces a revised belief set K^*_A and a revised belief revision operation $*'$ (varying with A). The postulates satisfied by the revision operation on epistemic states are as follows: (K*1)–(K*7) correspond to the conditions on almost-rational consequence operations, while (K*8) corresponds to conditionalization.

DEFINITION 4.20
An *almost-rational* belief revision operation on epistemic states is an operation from epistemic states $\langle K, * \rangle$ and infons A to epistemic states $\langle K^*_A, *' \rangle$ that satisfies the following postulates for all belief sets K and all infons A and B.

(K*1) K^*_A is a belief set
(K*2) $A \in K^*_A$
(K*3) If $K \vdash A$ then $K^*_A = K$
(K*4) If $B \in K^*_A$ then $K^*_{A \wedge B} = K^*_A$
(K*5) If $A \dashv\vdash B$ then $K^*_A = K^*_B$
(K*6) $K^*_A \cap K^*_B \subseteq K^*_{A \vee B}$
(K*7) If $B \notin K^*_{A \vee B}$ then $K^*_{A \vee B} \subseteq K^*_A$
(K*8) $(K^*_A)^{*'}_B = K^*_{A \wedge B}$.

COROLLARY 4.21
Suppose each epistemic state $\langle K, *\rangle$ is defined from an epistemic state $\langle K, \mid\!\sim\rangle$ using the Ramsey test. Then the belief revision operation from epistemic states $\langle K, *\rangle$ and infons A to epistemic states $\langle K^*_A, *'\rangle$ is an almost-rational belief revision operation.

Let us conclude with an example from Kautz [26], based again on the cooking hierarchy from Figure 1. Suppose the simple plans derived from this hierarchy are ordered as follows (again omitting negative event infons and interval symbols).

PrepareMeal, *MeatDish*, *ChickenMarinara*, *Sauce*, *Marinara*
$\prec$
PrepareMeal, *PastaDish*, *SpaghettiMarinara*, *Noodles*, *Spaghetti*, *Sauce*, *Marinara*, *Boil*
PrepareMeal, *PastaDish*, *Noodles*, *Fettuccine*, *Sauce*, *Boil*
$\prec$
WashDishes

The observer agent initially believes that the observed agent is executing the *ChickenMarinara* plan (the *a priori* most plausible simple plan), and hence predicts the actions *Marinara*, *Sauce*, *MeatDish* and *PrepareMeal*. On observing the action *Fettuccine*, it revises its beliefs to predict the new actions *Noodles*, *Boil* and *PastaDish* (the new actions in the most plausible simple plans containing *Fettuccine*), while dropping the beliefs in *ChickenMarinara*, *Marinara* and *MeatDish*, and retaining the beliefs in *Sauce* and *PrepareMeal*. The consequence operation so defined is clearly nonmonotonic, e.g. *Sauce* $\mid\!\sim$ *ChickenMarinara* but *Sauce* $\wedge$ *Fettuccine* $\not\mid\!\sim$ *ChickenMarinara*.

5 Related work

The formulations of plan recognition in this paper address the issues of prior information about the plausibility of plans (as encoded in a preference ordering) and, in a more limited way, context (this ordering can be relative to the agent's prior beliefs). Both themes have been the subject of related work both in theoretical research on formalizing plan recognition and in practical systems that carry out plan recognition. The requirements expected of practical plan recognition systems serve to highlight the limitations of the present logical analysis of the problem, and suggest further directions in which this work could be strengthened.

Although it has long been recognized that prior information about the plausibility of various possible plans is essential in any practical plan recognition system, e.g. as noted by Kautz and Allen [28], only more recently have there been any attempts to model this precisely. Apart from work inspired by belief revision [52, 25], two techniques that have been used to directly model the uncertainty of plan recognition predictions are Dempster-Shafer theory and Bayesian networks. Dempster-Shafer theory has been used by Carberry [11], Bauer *et al.* [9] and Bauer [8], while Bayesian networks have been used by Huber, Durfee and Wellman [24], Pynadath and Wellman [43] and Albrecht, Zukerman and Nicholson [1]. Bayesian networks provide a convenient way of representing the prior likelihoods of plans and goals, and provide a method of learning these likelihoods from observation—the focus of the work of Albrecht, Zukerman and Nicholson [1].

The systems of Huber, Durfee and Wellman [24] and Pynadath and Wellman [43] are also interesting in that they operate in dynamic domains, whereas theories based on logic, even though handling multiple observations, are typically restricted to static domains. The difficult issue with dynamic domains is that the observer agent is forced to take into account the observed agent's actions in response to changes in the world, which may include attempts to recover from failures or unforeseen events in addition to executing an originally intended plan. A further issue considered by Pynadath and Wellman is the effect of context on plan recognition, 'context' here referring to the dynamically changing situation in which the agent is embedded. For example, in the traffic monitoring domain of this system, features such as the presence or absence of cars in a lane adjacent to the observed agent affect the plausibility of a plan such as *overtaking*. However, in their system, such contextual features are simply encoded as additional input nodes to the Bayesian network, an approach that could suffer from computational inefficiency in more complex domains.

Existing work on plan recognition in dynamic domains is based on the simplifying (but possibly realistic) assumption that observations are made in the same temporal order as that of the sequence of actions, whereas this assumption is typically not made with other previous work: e.g. for 'story understanding' applications, such an assumption would be wrong. By exploiting such an assumption, Lesh and Etzioni [34] have developed an efficient method for quickly determining whether a goal is consistent with a sequence of observations. Their goal recognizer is monotonic and makes no use of prior information about the relative plausibility of the various goals. With simple proposition-like actions, such techniques seem more closely related to those based on probabilistic finite state machines or Markov processes than to previous algorithms for plan recognition, and one avenue for future work is to explore this connection further.

Attention to context is always important for performing realistic plan recognition. In general, there has been a recent trend in plan recognition systems towards more sophisticated models of context. In most early work on keyhole plan recognition (e.g. up to and including that of Carberry [11]), the context of the plan recognition system is taken as the partial plan formed on the basis of previous observations, while approaches to intended plan recognition have always emphasized a richer context such as the beliefs, goals and intentions of the speakers in a conversation [37]. As mentioned above, the system of Pynadath and Wellman [43] takes as part of the context some dynamically changing features of the environment; thus the system is more able to 'track' an agent's actions and goals over time [50]. This notion of context is taken a step further in the plan recognition component of the TraumAID system [22] in which the context includes prior actions taken by the observed agents (physicians in an emergency room treating a trauma patient) in response to prior observations (information about the patient): these actions both generate new information (beliefs of the observed agents) and influence the appropriateness of future actions. Thus in this domain, not only is prior knowledge of the plausibility of various plans essential, but this plausibility assessment must itself be dynamically updated during the plan recognition process.

Finally, an intuition promising a competing formalization of plan recognition is based on treating plan recognition as explanation (or abduction) [12, 6, 9]. The motivation is the idea that the purpose of plan recognition is not only to predict an agent's actions, but also to explain them. Thus the plan recognition agent is conceived of as searching for plausible explanations of its observations. However, in practice, heuristics are needed for determining the 'best explanations' (possibly restricting explanations to a special syntactic form or to a particular class of sentences). Thus to the extent that criteria for 'good' explanations are vague, such approaches may be more closely related to those based on nonmonotonic reasoning than at first appears.

6 Conclusion

We presented two logical theories of plan recognition. The first theory is a reconstruction of Kautz's theory using situation semantics to define the semantic analogues of plans. Plan recognition is treated as entailment over the class of courses of events modelling the simple plans, a class of plans characterized constructively, whose members the observer agent is aiming to recognize. This theory, as is Kautz's, is simplified in adopting the assumption that all simple plans are treated as equally plausible by the observer agent. The second theory of plan recognition is nonmonotonic, and is based on rejecting this simplifying assumption. The idea is that a plausibility ordering on courses of events corresponding to the simple plans generates a nonmonotonic consequence operation of a type which we have characterized as an 'almost-rational' consequence operation. We have then shown how to define the revision of an agent's epistemic state construed as a belief set together with a nonmonotonic consequence operation, allowing changes in the observer agent's state arising from a sequence of observations to be modelled. The conditions on ordered structures of situations assume that all simple plans are comparable under the plausibility ordering; thus our approach to plan recognition can be considered an abstraction of those based on Bayesian inference and Dempster–Shafer theory.

Acknowledgements

This work started as Ph.D. research in the Department of Computer Science at the University of Essex: thanks to my supervisor, Raymond Turner, for introducing me to situation semantics, and to Sam Steel, who taught me about planning research. This paper also draws heavily on my work on belief revision carried out in the Knowledge Systems Group at the University of Sydney, and thanks are due to Norman Foo for coordinating this group. My knowledge of situation semantics has been greatly enhanced by discussions with Greg Restall and Lawrence Cavedon; thanks also to David Israel and John Perry for hosting my sabbatical at CSLI. A much earlier version of this paper was presented at the Dagstuhl Seminar on Deductive Approaches to Plan Generation and Plan Recognition held at Schloß Dagstuhl in October, 1993. Thanks to DFKI Saarbrücken for sponsoring this workshop and to the participants for much stimulating discussion. This research was supported by an Australian Research Council Small Grant and the Symbolic Reasoning Systems Project of the Australian National University.

References

[1] D. W. Albrecht, I. Zukerman, and A. E. Nicholson. Bayesian models for keyhole plan recognition in an adventure game. *User Modeling and User-Adapted Interaction*, **8**, 5–47, 1998.

[2] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic*, **50**, 510–530, 1985.

[3] J. F. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, **26**, 832–843, 1983.

[4] J. F. Allen. Towards a general theory of action and time. *Artificial Intelligence*, **23**, 123–154, 1984.

[5] J. F. Allen. Temporal reasoning and planning. In *Reasoning About Plans*, J. F. Allen, H. A. Kautz, R. N. Pelavin and J. D. Tenenberg, eds. Morgan Kaufmann, San Mateo, CA, 1991.

[6] D. E. Appelt and M. E. Pollack. Weighted abduction for plan ascription. *User Modeling and User-Adapted Interaction*, **2**, 1–25, 1992.

[7] J. Barwise and J. Perry. *Situations and Attitudes.* MIT Press, Cambridge, MA, 1983.

[8] M. Bauer. A Dempster–Shafer approach to modeling agent preferences for plan recognition. *User Modeling and User-Adapted Interaction*, **5**, 317–348, 1995.

[9] M. Bauer, S. Biundo, D. Dengler, J. Koehler, and G. Paul. PHI—a logic-based tool for intelligent help systems. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pp. 460–466, 1993.

[10] S. Blamey. Partial logic. In *Handbook of Philosophical Logic: Volume 3*, D. M. Gabbay and F. Guenthner, eds. Reidel, Dordrecht, 1986.

[11] S. Carberry. Incorporating default inferences into plan recognition. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90)*, pp. 471–478, 1990.

[12] E. Charniak and R. P. Goldman. A Bayesian model of plan recognition. *Artificial Intelligence*, **64**, 53–79, 1992.

[13] P. R. Cohen, C. R. Perrault, and J. F. Allen. Beyond question answering. In *Strategies for Natural Language Processing*, W. G. Lehnert and M. H. Ringle, eds. Lawrence Erlbaum Associates, Hillsdale, NJ, 1982.

[14] M. J. Cresswell. *Semantical Essays.* Kluwer Academic Publishers, Dordrecht, 1988.

[15] K. Devlin. *Logic and Information.* Cambridge University Press, Cambridge, 1991.

[16] M. Freund. Injective models and disjunctive relations. *Journal of Logic and Computation*, **3**, 231–247, 1993.

[17] D. M. Gabbay. Theoretical foundations for non-monotonic reasoning in expert systems. In *Proceedings of the NATO Advanced Study Institute on Logics and Models of Concurrent Systems*, pp. 439–457, 1985.

[18] P. Gärdenfors. *Knowledge in Flux.* MIT Press, Cambridge, MA, 1988.

[19] P. Gärdenfors. Nonmonotonic inferences based on expectations: a preliminary report. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR'91)*, pp. 585–590, 1991.

[20] P. Gärdenfors and D. Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge (TARK 1988)*, pp. 83–95, 1988.

[21] P. Gärdenfors and D. Makinson. Nonmonotonic inference based on expectations. *Artificial Intelligence*, **65**, 197–245, 1994.

[22] A. S. Gertner and B. L. Webber. A bias towards relevance: recognizing plans where goal minimization fails. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI-96)*, pp. 1133–1138, 1996.

[23] H. P. Grice. Meaning. *Philosophical Review*, **66**, 377–388, 1957.

[24] M. J. Huber, E. H. Durfee, and M. P. Wellman. The automated mapping of plans for plan recognition. In *Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence*, pp. 344–351, 1994.

[25] P. Jachowicz and R. G. Goebel. Describing plan recognition as nonmonotonic reasoning and belief revision. In *Advanced Topics in Artificial Intelligence*, A. Sattar, ed. Springer-Verlag, Berlin, 1997.

[26] H. A. Kautz. A circumscriptive theory of plan recognition. In *Intentions in Communication*, P. R. Cohen, J. Morgan and M. E. Pollack, eds. MIT Press, Cambridge, MA, 1990.

[27] H. A. Kautz. A formal theory of plan recognition and its implementation. In *Reasoning About Plans*, J. F. Allen, H. A. Kautz, R. N. Pelavin and J. D. Tenenberg, eds. Morgan Kaufmann, San Mateo, CA, 1991.

[28] H. A. Kautz and J. F. Allen. Generalized plan recognition. In *Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI-86)*, pp. 32–37, 1986.

[29] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, **44**, 167–207, 1990.

[30] S. A. Kripke. Naming and necessity. In *Semantics for Natural Language*, D. Davidson and G. Harman, eds. Reidel, Dordrecht, 1972.

[31] P. B. Ladkin. Satisfying first-order constraints about time intervals. In *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI-88)*, pp. 512–517, 1988.

[32] P. B. Ladkin and R. D. Maddux. Representation and reasoning with convex time intervals. Technical Report KES.U.88.2, Kestrel Institute, 1988.

[33] D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, **55**, 1–60, 1992.

[34] N. Lesh and O. Etzioni. A sound and fast goal recognizer. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pp. 1704–1710, 1995.

[35] D. K. Lewis. *Counterfactuals.* Blackwell, Oxford, 1973.

[36] D. K. Lewis. *On the Plurality of Worlds.* Blackwell, Oxford, 1986.

[37] D. J. Litman and J. F. Allen. A plan recognition model for subdialogues in conversations. *Cognitive Science*, **11**, 163–200, 1987.

[38] D. Makinson. General patterns in nonmonotonic reasoning. In *Handbook of Logic in Artificial Intelligence and Logic Programming: Volume 3*, D. M. Gabbay, C. J. Hogger and J. A. Robinson, eds. Oxford University Press, Oxford, 1994.

[39] D. Makinson and P. Gärdenfors. Relations between the logic of theory change and nonmonotonic logic. In *The Logic of Theory Change*, A. Fuhrmann and M. Morreau, eds. Springer-Verlag, Berlin, 1991.

[40] J. M. McCarthy. Circumscription—a form of non-monotonic reasoning. *Artificial Intelligence*, **13**, 27–39, 1980.

[41] A. C. Nayak. Iterated belief change based on epistemic entrenchment. *Erkenntnis*, **41**, 353–390, 1994.

[42] M. E. Pollack. A model of plan inference that distinguishes between the beliefs of actors and observers. In *Proceedings of the Twenty-Fourth Annual Meeting of the Association for Computational Linguistics*, pp. 207–214, 1986.

[43] D. V. Pynadath and M. P. Wellman. Accounting for context in plan recognition, with application to traffic monitoring. In *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*, pp. 472–481, 1995.

[44] H. Rott. Two methods of constructing contractions and revisions of knowledge systems. *Journal of Philosophical Logic*, **20**, 149–173, 1991.

[45] E. D. Sacerdoti. The nonlinear nature of plans. In *Proceedings of the Fourth International Joint Conference on Artificial Intelligence*, pp. 206–214, 1975.

[46] C. F. Schmidt, N. S. Sridharan, and J. L. Goodson. The plan recognition problem: an intersection of psychology and artificial intelligence. *Artificial Intelligence*, **11**, 45–83, 1978.

[47] Y. Shoham. *Reasoning about Change*, MIT Press, Cambridge, MA, 1988.

[48] C. L. Sidner and D. J. Israel. Recognizing intended meaning and speakers' plans. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pp. 203–208, 1981.

[49] R. C. Stalnaker. A theory of Conditionals. In *Causation and Conditionals*, Sosa, E., ed. Oxford University Press, Oxford, 1968.

[50] M. Tambe and P. S. Rosenbloom. RESC: an approach for real-time, dynamic agent tracking. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pp. 103–110, 1995.

[51] A. Tate. Generating project networks. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pp. 888–893, 1977.

[52] W. R. Wobcke. Plan recognition as belief revision. In *Proceedings of the Second World Conference on the Fundamentals of Artificial Intelligence*, pp. 481–492, 1995.

Appendix

A Proofs

PROPOSITION 3.11
Any situated consequence operation $\vdash$ satisfies the following properties.

If $\Gamma \dashv\vdash \Delta$ then $\Gamma \vdash A$ iff $\Delta \vdash A$ (Left Logical Equivalence)
If $\Gamma \vdash A$ and $A \vdash B$ then $\Gamma \vdash B$ (Right Weakening)
If $\Gamma \cup \{A\} \vdash C$ and $\Gamma \cup \{B\} \vdash C$ then $\Gamma \cup \{A \vee B\} \vdash C$ (Or)

PROOF. For (Left Logical Equivalence), suppose $\Gamma \vdash A$ and $\Delta \vdash \Gamma$. Then since $\Gamma \subseteq \Delta \cup \Gamma$, $\Delta \cup \Gamma \vdash A$ by (Monotonicity), so since $\Delta \vdash \Gamma$, $\Delta \vdash A$ by (Cut). Similarly, if $\Delta \vdash A$ and $\Gamma \vdash \Delta$ then $\Gamma \vdash A$ using (Monotonicity) and (Cut). (Right Weakening) follows from (Cut) and (Monotonicity) with $\Delta = \{A\}$. For (Or), by compactness, there are finite subsets Γ_1 and Γ_2 of Γ such that $\Gamma_1 \cup \{A\} \vdash C$ and $\Gamma_2 \cup \{B\} \vdash C$. Let γ_1 (γ_2) be the conjunction of infons in Γ_1 (Γ_2). We show that $\Gamma_1 \dashv\vdash \gamma_1$ (similarly, $\Gamma_2 \dashv\vdash \gamma_2$): first, (Identity) gives $\gamma \vdash \gamma$ for each $\gamma \in \Gamma_1$, so by (Monotonicity), $\Gamma_1 \vdash \gamma$ for each such γ, but then by (And), $\Gamma_1 \vdash \gamma_1$; second, (And-Elimination) shows that $\gamma_1 \vdash \gamma$ for each infon $\gamma \in \Gamma_1$, so by definition, $\gamma_1 \vdash \Gamma_1$. By similar reasoning, $\Gamma_1 \cup \{A\} \dashv\vdash \gamma_1 \wedge A$ and $\Gamma_2 \cup \{B\} \dashv\vdash \gamma_2 \wedge B$. Now by (Left Logical Equivalence), $\gamma_1 \wedge A \vdash C$, and $\gamma_2 \wedge B \vdash C$, so by (And-Premise), $\gamma_1 \wedge \gamma_2 \wedge A \vdash C$ and $\gamma_1 \wedge \gamma_2 \wedge B \vdash C$. Hence by (Or-Premise), (And-Distribution) and (Left Logical Equivalence), $\gamma_1 \wedge \gamma_2 \wedge (A \vee B) \vdash C$. But again by reasoning similar to the above, $\Gamma_1 \cup \Gamma_2 \cup \{A \vee B\}$ is provably equivalent to $\gamma_1 \wedge \gamma_2 \wedge (A \vee B)$, so by (Left Logical Equivalence), $\Gamma_1 \cup \Gamma_2 \cup \{A \vee B\} \vdash C$, so $\Gamma \cup \{A \vee B\} \vdash C$ by (Monotonicity) and (Cut), using the fact that $\Gamma \vdash \Gamma_1 \cup \Gamma_2$, which holds by (Identity) and (Monotonicity). ∎

The following series of definitions and lemmas lead up to the characterization result for $\vdash$ in terms of entailment over situations. The main intermediate results concern a characterization of the coherent situations using prime closed consistent sets of infons.

LEMMA A.1
Any infon A is provably equivalent to a disjunctive normal form.

PROOF. The usual inductive proof for propositional calculus can easily be adapted, the postulates for $\vdash$ ensuring the provable equivalence of an infon and its disjunctive normal form. ■

LEMMA A.2
If Γ is a consistent set of infons then there exists a prime closed consistent set of infons containing Γ.

PROOF. As in the proof of Lindenbaum's lemma, start with an enumeration of the infons $A_1, A_2, \cdots$ in increasing order of complexity. Now given a consistent set of infons Γ, define an infinite sequence of sets of infons $\Gamma_0, \Gamma_1, \cdots$, as follows. First, define $\Gamma_0 = \Gamma$. At step i, if $\Gamma_i \cup \{A_i\} \nvdash$ *false* (where *false* is an arbitrary contradiction), define $\Gamma_{i+1} = Cn(\Gamma_i \cup \{A_i\})$ (the smallest closed set of sentences containing $\Gamma_i \cup \{A_i\}$), otherwise define $\Gamma_{i+1} = \Gamma_i$. We show that Γ_∞ is prime, closed and consistent. Clearly the construction guarantees that Γ_∞ is closed and consistent, for these properties hold at each stage of the process. For primeness, suppose that Γ_∞ contains $B_1 \vee B_2$ but neither B_1 nor B_2. Then at step i when $B_1 \vee B_2$ was considered, $\Gamma_i \cup \{B_1 \vee B_2\} \nvdash$ *false*. But $\Gamma_i \cup \{B_1\} \vdash$ *false* and $\Gamma_i \cup \{B_2\} \vdash$ *false* since B_1 and B_2 were inconsistent with subsets of Γ_i at steps before i and are therefore both inconsistent with Γ_i. So by disjunction in the premisses, $\Gamma_i \cup \{B_1 \vee B_2\} \vdash$ *false*, a contradiction. ■

COROLLARY A.3
If Γ is a consistent set of infons and A is an infon such that $\Gamma \nvdash A$ then there exists a prime closed consistent set of infons containing Γ that does not contain A.

PROOF. First note that this does not follow directly from Lemma A.2 since, given such a set Γ and infon A, $\neg A$ may not be consistent with Γ (e.g. in the case where A is $B \vee \neg B$), so Lemma A.2 cannot be applied to the set $\Gamma \cup \{\neg A\}$. However, the construction of Lemma A.2 can be generalized in the following way: where in Lemma A.2 the test at each stage is for whether $\Gamma_i \cup \{A_i\} \nvdash$ *false*, we replace this with the test for whether $\Gamma_i \cup \{A_i\} \nvdash A$. At the end of the process, the set Γ_∞ is clearly closed and does not contain A, and, using the same reasoning as in the proof of Lemma A.2, this set is also prime. ■

THEOREM 3.12
The class of situated consequence operations is sound and complete with respect to entailment over sets of situations. That is, (soundness) the consequence operation $\vdash_\Sigma$ defined as entailment over a set of situations Σ by $\Gamma \vdash_\Sigma A$ iff $\Gamma \models_\Sigma A$ is a situated consequence operation, and (completeness) given a situated consequence operation $\vdash$, there is a set of situations Σ such that $\Gamma \models_\Sigma A$ iff $\Gamma \vdash A$.

PROOF. (Soundness) Given a set of situations Σ, let the consequence operation $\vdash_\Sigma$ be defined by $\Gamma \vdash_\Sigma A$ iff $\Gamma \models_\Sigma A$. Clearly each of the postulates for situated consequence operations is satisfied as a result of the satisfaction conditions for situations. As an example, consider (And-Introduction): if $A \vdash B$ and $A \vdash C$ then $A \vdash B \wedge C$. Let $\sigma \in \Sigma$ be any coherent situation supporting A and suppose $A \vdash_\Sigma B$ and $A \vdash_\Sigma C$. By hypothesis, $\sigma \models B$ and $\sigma \models C$, hence $\sigma \models B \wedge C$. Thus by definition, $A \vdash_\Sigma B \wedge C$.

(Completeness) Given a situated consequence operation $\vdash$, let $\mathcal{G}$ be the set of prime closed consistent (i.e. closed under $\vdash$) sets of infons over the language of $\vdash$. For each such set $\Gamma \in \mathcal{G}$, define a situation σ_Γ that consists of the set of basic infons $\{\langle a, + \rangle : a \in \Gamma\} \cup \{\langle a, - \rangle : \neg a \in \Gamma\}$ where a is a basic infon in the language of $\vdash$. Define Σ to be the set of all such σ_Γ, i.e. $\Sigma = \{\sigma_\Gamma : \Gamma \in \mathcal{G}\}$.

We show by induction on infons A in disjunctive normal form that $A \in \Gamma$ iff $\sigma_\Gamma \models A$ for any prime closed consistent set of infons Γ. If A is a basic infon or the negation of a basic infon then $A \in \Gamma$ iff $\sigma_\Gamma \models A$ by definition. If A is of the form $A_1 \wedge A_2$ then $A \in \Gamma$ iff $A_1 \in \Gamma$ and $A_2 \in \Gamma$ iff $\sigma_\Gamma \models A_1$ and $\sigma_\Gamma \models A_2$ iff $\sigma_\Gamma \models A$. Finally, if A is of the form $A_1 \vee A_2$ then $A \in \Gamma$ iff $A_1 \in \Gamma$ or $A_2 \in \Gamma$ (using primeness) iff $\sigma_\Gamma \models A_1$ or $\sigma_\Gamma \models A_2$ iff $\sigma_\Gamma \models A$.

Finally, we show that $\Gamma \models_\Sigma A$ iff $\Gamma \vdash A$. First suppose $\Gamma \vdash A$. Then every prime closed consistent set containing Γ contains A, and hence every coherent situation in Σ that supports Γ supports A. That is, $\Gamma \models_\Sigma A$. Conversely, if $\Gamma \nvdash A$ then by Corollary A.3, there is some prime closed consistent set Γ' containing Γ that does not contain A, so the coherent situation $\sigma_{\Gamma'} \in \Sigma$ supports Γ but not A, hence $\Gamma \not\models_\Sigma A$. ■

THEOREM 3.20
The class of interval theories based on $\mathcal{I}$ is sound and complete with respect to the class of interval structures over $\mathcal{I}$.

PROOF. Fix notation as follows: for an interval i (assuming a fixed interval structure ξ), let the start point of i be denoted i_x and let the end point of i be denoted i_y.

As usual, soundness is easier than completeness. It is straightforward to check that all the axiom schemes whose interval terms are drawn from a set $\mathcal{I}$ are valid in every interval structure which contains an interpretation for each interval term in $\mathcal{I}$.

For completeness, we start with a maximal consistent interval theory Γ over a set of interval terms $\mathcal{I}$, and show how to construct an interval structure ξ such that ξ supports A iff A is contained in Γ. As the basis of the construction, we show that an interval theory generates a strict total ordering on a set of points—the start and end points of the intervals in the theory; we *assume* that such an ordering on a countable number of points can be mapped to the real line so as to preserve the ordering relation.

Given a maximal consistent interval theory Γ over $\mathcal{I}$, consider the start and end points i_x and i_y of the intervals in $\mathcal{I}$. Define an equivalence relation $=$ on such points by defining $i_x = j_x$ iff Γ contains $i = j$, i starts j or j starts i, $i_x = j_y$ iff j meets i, $i_y = j_x$ iff i meets j, and $i_y = j_y$ iff $i = j$, i finishes j or j finishes i. Now define an ordering $<$ on points by setting $i_x < j_x$ iff Γ contains $i < j$, j during i, i overlaps j, i meets j or j finishes i, $i_x < j_y$ iff Γ does not contain $i \succ j$ or j meets i, $i_y < j_x$ iff Γ contains $i < j$, and $i_y < j_y$ iff Γ contains $i < j$, i during j, i overlaps j, i meets j or i starts j.

The first step is to verify that $=$ is an equivalence relation on points such that whenever $p = q$, $p < r$ iff $q < r$ and $r < p$ iff $r < q$. This follows from the axioms of equality together with the transitivity axiom schemes. The second step is to verify that the relation $<$ is a strict total ordering. Here transitivity again follows from the transitivity axiom schemes, whereas irreflexivity follows from (Exclusivity) and the first equality axiom, and totality follows from the axiom scheme of the same name. The final step is to check that ξ supports A iff A is contained in Γ. It suffices to do this for the basic infons, i.e. those of the form $i\ r\ j$ where r is an interval relation. First, it is straightforward to check that for each such infon A, if Γ contains A then ξ supports A. Second, it is also easy to check that if ξ supports $i\ r\ j$ then ξ rejects $i\ r'\ j$ for any other relation r'. But by (Totality) and (Exclusivity), Γ contains only one such infon for any pair of intervals i and j; thus if Γ does not contain A, ξ rejects A. $\blacksquare$

PROPOSITION 3.21 (Deduction theorem)
If $\Gamma \cup \{A\} \vdash B$ then $\Gamma \vdash A \supset B$ for an interval theory Γ and temporal infons A and B.

PROOF. Let Γ be an interval theory and A and B temporal infons. First note that if $A \vdash B$ then $A \vee \neg A \vdash B \vee \neg A$ using (Identity), (Right Weakening) and (Or). But for any temporal infon A we have $\vdash A \vee \neg A$. Thus $\vdash A \supset B$ whenever $A \vdash B$.

For the Deduction theorem, if $\Gamma \vdash A \supset B$ then $\Gamma \cup \{A\} \vdash B$ by (Identity), (And), (Or-Elimination) and (Right Weakening). Now suppose $\Gamma \cup \{A\} \vdash B$. Then since $\vdash$ is compact, there is some finite subset Γ' of Γ such that $\Gamma' \cup \{A\} \vdash B$, so there is a temporal infon γ such that $\gamma \wedge A \vdash B$. By the above observation, $\vdash (\gamma \wedge A) \supset B$, i.e. $\vdash \neg\gamma \vee \neg A \vee B$, i.e. $\vdash \gamma \supset (A \supset B)$, so $\gamma \vdash A \supset B$ by the other half of the Deduction theorem. It follows that $\Gamma \vdash A \supset B$, as required. $\blacksquare$

COROLLARY 3.25
The class of situated consequence operations $\vdash$ for which $\vdash \Im$ is sound and complete with respect to entailment over sets of courses of events.

PROOF. (Soundness) This follows from the soundness parts of Theorems 3.12 and 3.20. More precisely, Theorem 3.20 ensures that the axioms of $\Im$ are valid in every course of events, while Theorem 3.12 ensures that the postulates for consequence operations are sound with respect to entailment over courses of events, since the satisfaction conditions for event infons and complex infons in courses of events are the same as for situations.

(Completeness) This follows from the completeness parts of Theorems 3.12 and 3.20. More precisely, if $\Gamma \nvdash A$, let Γ' be a prime closed consistent set of infons containing Γ but not A. The construction of Theorem 3.12 gives a situation supporting all and only the event infons in Γ', while the definitions from Theorem 3.20 provide an interval structure supporting all and only the temporal infons in Γ'. Thus the course of events $\langle \sigma, \xi \rangle$ supports all and only the infons in Γ', hence supports Γ but not A, i.e. $\Gamma \not\models A$. $\blacksquare$

THEOREM 3.37
The plan theories based on a plan recognition inference operation derived from a plan hierarchy H correspond to the realizations of the simple plans based on H (up to renaming of interval symbols).

PROOF. We first show that any realization of a simple plan P gives rise to a plan theory based on a plan recognition inference operation derived from H. Let $\mathcal{P}$ be the class of all realizations of simple plans based on H, and let $\vdash$

be the consequence operation defined by entailment over $\mathcal{P}$, which, by Theorem 3.12, is a situated consequence operation. Let $\epsilon = \langle \sigma, \xi \rangle$ be a realization of a simple plan P, and let T be the set of infons supported by ϵ. It is clear that T is prime and consistent since T derives from a course of events, and that T is closed under the postulates for situated consequence operations. We now show that T is closed under the postulates for plan recognition derived from H. Since P is a simple plan, it contains exactly one primary action A; rename the interval constant associated with A to agree with that in the (EXH) postulate so that (EXH) is satisfied. T is also clearly closed under (ABS), (DEC), (SPEC) and (SUB) using the coherence of the plan hierarchy, and additionally for (SPEC) and (SUB), the fact that no action type is a specialization of two distinct action types. We show closure under (COMP) by induction on the elaboration construction, i.e. that (COMP) is satisfied for the plans P in each stage of the construction of the simple plan from the primary action in P whose existence is guaranteed by (EXH). Now an action A is in P if (i) A is a specialization of the primary action in P, (ii) A is a subtype of an action in P such that A is not a component of any plan schema, or (iii) A abstracts an action A' that is a component of a plan schema in P. In case (i), the base case of the induction, (COMP) is satisfied trivially because the specializations of a primary action are components of no plan schemas. In case (ii), (COMP) is satisfied by the induction hypothesis since the addition of A to the plan results in the antecedents of no further (COMP) postulates being satisfied. Case (iii) also follows from the formulation of the (COMP) postulates. Finally, by Theorem 3.20, T contains $\Im$. Thus T is a plan theory based on a plan inference operation derived from H.

We next show that any plan theory T based on a plan recognition inference operation derived from H corresponds to a realization of a simple plan P based on H. Define a course of events $\langle \sigma, \xi \rangle$ as follows: set σ to be the set of positive or negative event infons E that are contained in T, and set ξ to be any interval structure based on the temporal infons in T (this exists because the temporal constraints on each action expansion are consistent). Since the (EXH), (SPEC) and (SUB) postulates reflect properties of the plan elaboration procedure, and (ABS) and (DEC) ensure that σ contains all parents and expansions of actions in T, σ supports all event infons E in some simple plan P. Also, by definition, ξ supports all temporal infons in T and hence all temporal infons in P. Clearly $\langle \sigma, \xi \rangle$ is coherent.

It remains to show that σ supports a positive event infon E only if $E \in P$. So, for a contradiction, suppose T contains some other positive event infon B. We associate with B a set of positive event infons $\mathcal{B}$, none of which are contained in P, but one of which is contained in T, then derive a contradiction by showing that none of the infons in $\mathcal{B}$ are, in fact, in T. Start with $\mathcal{B}$ as the set containing just B. Now B is (i) an event infon corresponding to the primary action in P but with a different interval variable, or an event infon corresponding to (ii) a primary action not in P, (iii) a specialization of an action in P, or (iv) an action such that neither itself nor any of its specializations are components of an action in P. If case (i)–(iii) applies, stop. In case (iv), there is a (COMP) postulate whose antecedent is B and consequent is $B_1 \vee \cdots \vee B_n$; replace B in $\mathcal{B}$ by the set of all such B_i. Again, the primeness of T ensures that at least one of the infons in $\mathcal{B}$ is contained in T. Now continually repeat this process for each of the different actions in $\mathcal{B}$, and for the new $\mathcal{B}$ that results from this set, building up a set of infons $\mathcal{B}$, one of which must be in T, but none of which correspond to actions in P. Eventually, the process must terminate with a set of infons that are either the primary action in P but with a different interval variable, a primary action not in P or a specialization of an action in P (termination is guaranteed because the hierarchy is finite). However, the (EXH), (SUB) and (SPEC) postulates ensure that such infons are in T iff the corresponding actions are in the plan P, giving a contradiction. Thus $\langle \sigma, \xi \rangle$ is a realization of P. ■

COROLLARY 3.38
For any plan hierarchy H, the class of plan recognition inference operations derived from H is sound and complete with respect to entailment over the courses of events that are the realizations of the simple plans derived from H.

PROOF. (Soundness) Let $\vdash$ be a plan recognition inference operation. If $\Gamma \vdash A$ then every plan theory that contains Γ contains A. So by Theorem 3.37, every course of events realizing a simple plan based on H that supports Γ supports A.

(Completeness) Again, let $\vdash$ be a plan recognition inference operation. If $\Gamma \not\vdash A$, let Γ' be a prime closed consistent set of infons containing Γ but not A, as guaranteed by Lemma A.2. The constructions of Theorems 3.12 and 3.20 give a course of events $\langle \sigma, \xi \rangle$ supporting all and only the infons in Γ', hence which does not support A. It suffices to show that $\langle \sigma, \xi \rangle$ is a realization of a simple plan derived from H. Now since $\vdash$ satisfies the postulates for plan recognition derived from H, Γ' is a plan theory, so by Theorem 3.37, corresponds to a simple plan P. Since σ supports all and only the event infons in Γ' and ξ supports all the temporal infons in Γ', $\langle \sigma, \xi \rangle$ is a realization of P by definition. ■

PROPOSITION 4.8
Any almost-rational consequence operation $\mid\!\sim$ satisfies the following properties.

$A \mathrel{|\!\sim} A$ (Reflexivity)
If $A \dashv\vdash B$ then $A \mathrel{|\!\sim} C$ iff $B \mathrel{|\!\sim} C$ (Left Logical Equivalence)
If $A \vee B \mathrel{|\!\sim} C$ then $A \mathrel{|\!\sim} C$ or $B \mathrel{|\!\sim} C$ (Disjunctive Rationality)

Proof. (Reflexivity) follows from the corresponding property of $\vdash$ together with (Supraclassicality). Similarly (Left Logical Equivalence) follows from (Supraclassicality) and (Reciprocity). (Disjunctive Rationality) follows from (Cautious Monotony), (Reciprocity) and (Almost Rational Monotony) as follows. If $A \vee B \mathrel{|\!\sim} A$ then $A \mathrel{|\!\sim} C$ follows from $A \vee B \mathrel{|\!\sim} C$ using (Cautious Monotony) and (Reciprocity), but if $A \vee B \not\mathrel{|\!\sim} A$ then $B \mathrel{|\!\sim} C$ follows from $A \vee B \mathrel{|\!\sim} C$ using (Almost Rational Monotony). ■

The next series of definitions lead up to the characterization of compact almost-rational consequence operations using ordered structures of situations. The initial results concern a characterization of the situations using prime infons. In the following lemmas, $\vdash_0$ denotes the weakest situated consequence operation, which exists because the intersection of all situated consequence operations is itself a situated consequence operation, and hence is the weakest such operation.

Definition A.4
A *literal* is an infon of the form A or $\neg A$ where A is a basic infon.

Definition A.5
An infon p is *prime* if it is a finite conjunction of literals $l_1 \wedge \cdots \wedge l_n$.

Lemma A.6
If p and q are consistent prime infons, then $p \vdash_0 q$ iff p includes every literal occurring as a conjunct in q.

Proof. The 'only if' part follows from the completeness part of Theorem 3.12. Suppose there is a literal l occurring as a conjunct in q that does not occur in p. It is then easy to construct a situation supporting p in which l is rejected, hence there is some model of p which does not support q, so $p \not\vdash_0 q$ by completeness. The 'if' part follows from (And-Elimination) and (And). ■

Lemma A.7
If p, q, and r are consistent prime infons, then $p \vdash_0 q \vee r$ implies $p \vdash_0 q$ or $p \vdash_0 r$.

Proof. First suppose that q and r are literals. We show the converse, i.e. if $p \not\vdash_0 q$ and $p \not\vdash_0 r$ then $p \not\vdash_0 q \vee r$. Suppose $p \not\vdash_0 q$ and $p \not\vdash_0 r$. The result then follows from Lemma A.6 and the completeness part of Theorem 3.12, since both q and r must be literals not contained in p, so there is a situation supporting p but rejecting both q and r, so that by completeness, $p \not\vdash_0 q \vee r$.

Now suppose q and r are arbitrary consistent prime infons. We show a contradiction. So suppose $p \vdash_0 q \vee r$ but $p \not\vdash_0 q$ and $p \not\vdash_0 r$. Since $p \vdash_0 q \vee r$, by (Or-Distribution) and (And-Elimination), $p \vdash_0 q' \vee r'$ for any literals q' in q and r' in r, and hence by the above case, $p \vdash_0 q'$ or $p \vdash_0 r'$ for any such q' and r'. But since $p \not\vdash_0 q$, by Lemma A.6, there is some literal q_i occurring as a conjunct in q that does not occur in p, hence $p \not\vdash_0 q_i$, and similarly, there is some literal r_j occurring as a conjunct in r such that $p \not\vdash_0 r_j$, giving a contradiction. ■

Definition A.8
A *prime decomposition* of an infon A is a finite disjunction of prime infons $p_1 \vee \cdots \vee p_n$ logically equivalent to A such that A is not logically equivalent to $p_1 \vee \cdots [p_i] \cdots \vee p_n$ for any i, where this infon is p with the prime p_i omitted.

Lemma A.9
Any consistent infon A has a prime decomposition unique up to logical equivalence.

Proof. Lemma A.1 ensures the existence of logically equivalent disjunctive normal forms. That is, any consistent infon A is logically equivalent to a disjunction of primes $p_1 \vee \cdots \vee p_n$. A prime decomposition can then be found by removing any prime p_i from this disjunction whenever there is some other prime p_j in the disjunction such that $p_i \vdash_0 p_j$ (this also has the effect of removing any inconsistent primes from the decomposition). Clearly the logical equivalence to A is preserved.

For uniqueness up to logical equivalence, suppose that A has two prime decompositions $P = p_1 \vee \cdots \vee p_n$ and $Q = q_1 \vee \cdots \vee q_m$. Then $p_i \vdash_0 q_1 \vee \cdots \vee q_m$ so $p_i \vdash_0 q_1$ or $\cdots$ or $p_i \vdash_0 q_m$ for each i by Lemma A.7. By suitably reordering the q_j, we can assume that $p_i \vdash_0 q_i$ for each i. Thus $m \leq n$ since otherwise Q contains a redundant prime. Similarly, for each i, $q_i \vdash_0 p_j$ for some j, so $q_i \vdash_0 p_j \vdash_0 q_j$, so $i = j$ (since Q is a prime decomposition), and $m = n$. Hence P is logically equivalent to Q. ■

The final series of lemmas concern properties of the sets of nonmonotonic consequences of an infon. Following Freund [16], the notation A^* is used to denote an infon provably equivalent to the set of nonmonotonic consequences of an infon A, which we suppose always to exist. But although prime decompositions are unique up to logical equivalence, they are not necessarily unique up to provable equivalence under the situated consequence operation $\vdash$ on which $\mathrel{|\!\sim}$ is based; however, using a special class of prime infons which are also *saturated*, the prime decomposition of A^* for any infon A is a disjunction of saturated prime infons that is unique up to provable equivalence under $\vdash$.

LEMMA A.10
For any infons A and B, if $A \mathrel{|\!\sim} B$ and $B \mathrel{|\!\sim} A$ then $A^* \equiv B^*$.

PROOF. It follows directly from (Reciprocity) that $A \mathrel{|\!\sim} C$ iff $B \mathrel{|\!\sim} C$. Thus the sets of nonmonotonic consequences of A and B are logically equivalent. Since each is also logically equivalent to its representing infon A^* and B^*, the two infons are also logically equivalent, i.e. $A^* \equiv B^*$. ■

LEMMA A.11
For any infon A, $A^* \equiv A^{**}$.

PROOF. Since both $A \mathrel{|\!\sim} A^*$ and $A^* \mathrel{|\!\sim} A$ (since $A^* \vdash_0 A$), $A^* \equiv A^{**}$ by Lemma A.10. ■

LEMMA A.12
For any infons A and B, $(A \vee B)^* \vdash_0 A^* \vee B^*$.

PROOF. Clearly $A \mathrel{|\!\sim} A^*$ and $B \mathrel{|\!\sim} B^*$, and since $A^* \vdash_0 A^* \vee B^*$ by (Or), $A \mathrel{|\!\sim} A^* \vee B^*$ using (Right Weakening). Similarly $B \mathrel{|\!\sim} A^* \vee B^*$. Hence by (Or), $A \vee B \mathrel{|\!\sim} A^* \vee B^*$. Thus by definition, $A^* \vee B^*$ is contained in the set of nonmonotonic consequences of $A \vee B$, hence $(A \vee B)^* \vdash_0 A^* \vee B^*$ since $(A \vee B)^*$ is logically equivalent to this set. ■

DEFINITION A.13
A prime infon p is *saturated* with respect to a finitely representable nonmonotonic consequence operation if $p^* \equiv p$.

LEMMA A.14
If $p_1 \vee \cdots \vee p_n$ is a prime decomposition of A^* then each p_i is saturated.

PROOF. Since $A^* \equiv ((A^* \wedge p_1) \vee \cdots \vee (A^* \wedge p_n))$, by Lemmas A.10, A.11 and A.12, $A^* \vdash_0 A^{**} \vdash_0 ((A^* \wedge p_1) \vee \cdots \vee (A^* \wedge p_n))^* \vdash_0 (A^* \wedge p_1)^* \vee \cdots \vee (A^* \wedge p_n)^* \vdash_0 p_1^* \vee \cdots \vee p_n^*$, i.e. $p_1 \vee \cdots \vee p_n \vdash_0 p_1^* \vee \cdots \vee p_n^*$. So for each i, $p_i \vdash_0 p_1^* \vee \cdots \vee p_n^*$ and so $p_i \vdash_0 p_j^*$ and hence $p_i \vdash_0 p_j$ for some j (since for each j, $p_j^* \vdash_0 p_j$). But $p_i \nvdash_0 p_j$ unless $i = j$. This means that for each i, $p_i \vdash_0 p_i^*$, and since also $p_i^* \vdash_0 p_i$, it follows that $p_i^* \equiv p_i$. ■

COROLLARY A.15
For any infon A, if A^* is consistent, the prime decomposition of A^* is unique up to provable equivalence under $\vdash$.

PROOF. Lemma A.9 ensures that A^* is logically equivalent to a disjunction of prime infons. Lemma A.14 ensures that each such prime infon is saturated. That is, for each such prime infon p, $p \equiv p^*$, so by (Supraclassicality), $p \equiv p^+$, where p^+ is the set $\{q : p \vdash q\}$. Consequently, the analogues of Lemma A.6 and so of Lemmas A.7 and A.9 hold with $\vdash$ replacing $\vdash_0$. That is, the prime decomposition of A^* is unique up to provable equivalence under $\vdash$. ■

To show the completeness of the class of nonmonotonic consequence operations, we need to construct, for each operation in the class, an ordered structure of situations over which preferential entailment exactly reflects the consequence operation. A preliminary step, following standard methods in conditional logic, is to define an ordering $\leq$ on infons reflecting their relative degree of plausibility. We define this relation over classes of infons considered equivalent if their sets of nonmonotonic consequences are provably equivalent.

DEFINITION A.16
Given a nonmonotonic consequence operation $\mathrel{|\!\sim}$ based on a situated consequence operation $\vdash$, the equivalence relation $\approx$ on infons is defined by setting $A \approx B$ iff $A^* \dashv\vdash B^*$.

DEFINITION A.17
Given a nonmonotonic consequence operation $\mathrel{|\!\sim}$, the ordering $\leq$ on equivalence classes of infons under $\approx$ is defined by setting $A \leq B$ iff $A \vee B \mathrel{|\!\sim} A$.

This is well defined as follows. First, if $A \approx C$ and $A \leq B$, then $C^* \dashv\vdash A^* \dashv\vdash (A \vee B)^* \dashv\vdash (C \vee B)^*$ so $C \leq B$. Second, if $B \approx D$ and $A \leq B$, then $(A \vee B)^* \dashv\vdash (A \vee D)^*$ so $A \leq D$.

LEMMA A.18
For an almost-rational consequence operation, the ordering $\leq$ on equivalence classes of infons is an almost-connected partial order.

PROOF. Following Gärdenfors and Makinson [20]. Reflexivity follows from the reflexivity of $\mid\!\sim$, i.e. since $A \mid\!\sim A$, $A \leq A$ using (Left Logical Equivalence). For antisymmetry, suppose $A \leq B$ and $B \leq A$, i.e. $A \vee B \mid\!\sim A$ and $A \vee B \mid\!\sim B$. Hence $A^* = (A \vee B)^* = B^*$ using (Left Logical Equivalence), (Cut) and (Cautious Monotony), and so $A \approx B$. For transitivity, suppose $A \leq B$ and $B \leq C$. This means that $B \vee C \mid\!\sim B$, so $A \vee B \vee C \mid\!\sim A \vee B$ by (Right Weakening) and (Or), hence $A \vee B \vee C \mid\!\sim A$ by (Left Logical Equivalence) and (Cut), since $A \leq B$ means $A \vee B \mid\!\sim A$. But $A \vee B \mid\!\sim A$ also implies $A \vee B \vee C \mid\!\sim A \vee C$ by (Right Weakening) and (Or), so since $A \vee B \vee C \mid\!\sim A$, $A \vee C \mid\!\sim A$ by (Left Logical Equivalence) and (Cautious Monotony), i.e. $A \leq C$. For almost connectedness, suppose $A \leq C$ but $A \not\leq B$ and $B \not\leq C$. First, $A \vee C \mid\!\sim A$ implies $A \vee B \vee C \mid\!\sim A \vee B$ by (Right Weakening) and (Or). But $A \vee B \not\mid\!\sim A$ implies $A \vee B \vee C \not\mid\!\sim A$ by (Left Logical Equivalence) and (Cautious Monotony), and similarly, $B \vee C \not\mid\!\sim B$ implies $A \vee B \vee C \not\mid\!\sim B$. These three conclusions contradict (Disjunctive Rationality). ■

LEMMA A.19
If $p_1, \cdots, p_n$ are saturated prime infons and p is a distinct saturated prime infon such that p^* is not provably equivalent to $(p_1 \vee \cdots \vee p_n)^*$, then if $p_1 \vee \cdots \vee p_n \leq p$ then $p_i \leq p$ for some i.

PROOF. We prove the converse. Suppose $p_i \not\leq p$ for each i. Then since $\leq$ is antisymmetric and p^* is not provably equivalent to p_i^* for each i, $p \leq p_i$ for each i. Thus $p \vee p_i \mid\!\sim p$ for each i and so by (Or), $p \vee (p_1 \vee \cdots \vee p_n) \mid\!\sim p$, so $p \leq p_1 \vee \cdots \vee p_n$. Since p^* is not provably equivalent to $(p_1 \vee \cdots \vee p_n)^*$, $p_1 \vee \cdots \vee p_n \not\leq p$ by the antisymmetry of $\leq$, as required. ■

THEOREM 4.11
The class of finitely representable almost-rational consequence operations is sound and complete with respect to preferential entailment over ordered structures of situations. That is, (soundness) the consequence operation $\mid\!\sim_{\mathcal{M}}$ defined by preferential entailment from an ordered structure of situations $\mathcal{M}$ by $A \mid\!\sim_{\mathcal{M}} B$ iff $A \mid\!\approx_{\mathcal{M}} B$ is a finitely representable almost-rational consequence operation based on some situated consequence operation $\vdash$, and (completeness) given a finitely representable almost-rational consequence operation $\mid\!\sim$ based on a situated consequence operation $\vdash$, there is an ordered structure of situations $\mathcal{M} = \langle \Sigma, \preceq \rangle$ such that $\Gamma \models_\Sigma A$ iff $\Gamma \vdash A$ and $A \mid\!\approx_{\mathcal{M}} B$ iff $A \mid\!\sim B$.

PROOF. (Soundness) Given an ordered structure of situations $\mathcal{M} = \langle \Sigma, \preceq \rangle$, let the consequence operation $\mid\!\sim_{\mathcal{M}}$ be defined by setting $A \mid\!\sim_{\mathcal{M}} B$ iff $A \mid\!\approx_{\mathcal{M}} B$, and let the situated consequence operation $\vdash_\Sigma$ be defined by setting $\Gamma \vdash_\Sigma A$ iff $\Gamma \models_\Sigma A$. Then by Theorem 3.12, $\vdash_\Sigma$ satisfies the postulates for situated consequence operations. Also each of the postulates for almost-rational consequence relations is satisfied as a result of the satisfaction conditions for ordered structures of situations. This includes (Supraclassicality), ensuring that $\mid\!\sim_{\mathcal{M}}$ is based on $\vdash_\Sigma$. It is also clear that $\mid\!\sim_{\mathcal{M}}$ is finitely representable, since each situation in Σ is finite and $\preceq$ is finitarily stoppered, hence for any set of infons Γ and infon A, the sets $\{B : \Gamma \vdash B\}$ and $\{B : A \mid\!\sim B\}$ are logically equivalent to finite disjunctions of finite conjunctions of infons.

(Completeness) Given an almost-rational consequence operation $\mid\!\sim$ based on a situated consequence operation $\vdash$, we need to define an ordered structure of situations $\mathcal{M} = \langle \Sigma, \preceq \rangle$ such that $\Gamma \vdash_\Sigma A$ iff $\Gamma \vdash A$ and $A \mid\!\approx_{\mathcal{M}} B$ iff $A \mid\!\sim B$. Let $\mathcal{P}$ be the set of consistent saturated prime infons, i.e. the set of consistent infons p for which $p \equiv p^*$. For each such infon $p \in \mathcal{P}$, define a coherent situation σ_p that consists of the set of basic infons $\{\langle a, + \rangle : p \vdash a\} \cup \{\langle a, - \rangle : p \vdash \neg a\}$ where a is a basic infon occurring in p. Define $\Sigma_{\mathcal{P}}$ to be the set of all such σ_p, i.e. $\Sigma_{\mathcal{P}} = \{\sigma_p : p \in \mathcal{P}\}$. In addition, define a set of situations Σ' as follows: whenever $A \nvdash B$ for some consistent infons A and B, but $p \vdash B$ for all consistent saturated prime infons $p \in \mathcal{P}$ for which $p \vdash A$, Σ' contains a situation that supports A but not B; such situations exist by Corollary A.3. Define Σ to be $\Sigma_{\mathcal{P}} \cup \Sigma'$. Clearly each situation in Σ is finite.

We verify that $\Gamma \models_\Sigma A$ iff $\Gamma \vdash A$ for any set of infons Γ and infon A. First, if $\Gamma \vdash A$ then whenever $p \vdash \Gamma$, $p \vdash A$ by (Monotonicity) and (Cut), so every situation $\sigma_p \in \Sigma$ that supports Γ supports A, and similarly by the construction of Corollary A.3, every situation in Σ' that supports Γ supports A, hence $\Gamma \models_\Sigma A$ by definition. Conversely, if $\Gamma \nvdash A$, by finite representability, the set of consequences of Γ is logically equivalent to some infon

B such that $B \nvdash A$. Now either there is a consistent saturated prime infon p for which $p \vdash B$ but $p \nvdash A$, with $\sigma_p \in \Sigma_{\mathcal{P}}$, or else by construction, there is a situation $\sigma \in \Sigma'$ such that $\sigma \models B$ but $\sigma \not\models A$. Thus $\Gamma \not\models_\Sigma A$, as required.

We now need to define the ordering $\preceq$ on Σ and show that it is an almost-connected partial order. The ordering $\leq$ on equivalence classes of infons enables us to simply set, for two situations σ_p and σ_q in $\Sigma_{\mathcal{P}}$, $\sigma_p \preceq \sigma_q$ iff $p \leq q$. It follows from Lemma A.18 that this defines an almost-connected partial order on $\Sigma_{\mathcal{P}}$. The ordering of situations in $\Sigma_{\mathcal{P}}$ with respect to those in Σ' is defined by setting $\sigma_p \preceq \sigma$ but $\sigma \not\preceq \sigma_p$ for all situations $\sigma_p \in \Sigma_{\mathcal{P}}$ and $\sigma \in \Sigma'$; the situations in Σ' are defined to be incomparable with respect to one another. Thus the situations in Σ' define a single cluster of situations that are all maximal in the $\preceq$ ordering, making this ordering on Σ an almost-connected partial order.

We need to show that $\preceq$ is finitarily stoppered. First note that by (Consistency Preservation), for any infon A, a minimal A-supporting situation cannot be contained in Σ', since if $A \mid\!\sim$ *false* then $A \vdash$ *false* and there are no A-supporting situations in Σ. So given an infon A, suppose $\sigma_p \models A$ for some $\sigma_p \in \Sigma$. We need to determine a situation σ° such that $\sigma^\circ \models A$ and for any situation σ' such that $\sigma' \models A$, $\sigma' \preceq \sigma^\circ$ implies $\sigma' = \sigma^\circ$. Since $p^* \dashv\vdash p$, $p \mid\!\sim A$ so that $p \vee A \mid\!\sim A$, i.e. $A \leq p$. Also for any infon B, if $A \mid\!\sim B$ then $A \vee B \mid\!\sim B$ so $B \leq A$. Thus $A^* \leq A$ and so by transitivity, $A^* \leq p$. Now let $p_1 \vee \cdots \vee p_n$ be a prime decomposition of A^*. Hence $p_1 \vee \cdots \vee p_n \leq p$ and so by Lemma A.19, either $p \approx A$ or $p_i \leq p$ for some i. But when $p \approx A$, A^* is itself a prime p_1 and hence also in this case, $p_i \leq p$ for some i ($i = 1$) by definition. The desired situation σ° is then σ_{p_i}. To check the required property, whenever there is some situation σ_q such that $\sigma_q \models A$ and $\sigma_q \preceq \sigma^\circ$, $q \leq p_i$ by definition. Note that since $\sigma_q \models A$, $A^* \leq q$ as above; note also that $q \not\approx p_j$ for $j \neq i$, since if $q \approx p_j$ ($j \neq i$), it follows that $p_j \leq p_i$, which is impossible when both are primes in the prime decomposition of A^*. Now if $q \approx p_i$ or the prime decomposition of A^* consists of just one prime p_i (so that $p_i \leq q$ since $A^* \leq q$), then $\sigma^\circ = \sigma_{p_i} \leq \sigma_q$, as required. Otherwise, since $A^* \leq q$, i.e. $p_1 \vee \cdots \vee p_n \leq q$, by Lemma A.19 again, $p_j \leq q$ for some j, so $p_j \leq p_i$ using transitivity since $q \leq p_i$. But since p_j and p_i are both primes in the prime decomposition of A^*, $p_j \leq p_i$ only if $j = i$. Hence $p_i \leq q$, i.e. $\sigma^\circ = \sigma_{p_i} \leq \sigma_q$, as required. Clearly also the number of such minimal A-supporting situations is finite.

Finally, we need to show that $A \mid\!\approx_{\mathcal{M}} B$ iff $A \mid\!\sim B$. First suppose that $A \mid\!\sim B$. If A^* is consistent, let $p_1 \vee \cdots \vee p_n$ be a prime decomposition of A^* and consider a situation $\sigma_p \in \mathit{min}(A)$. Then by the argument showing that $\preceq$ is stoppered, $\sigma_p = \sigma_{p_i}$ for some i. Hence $p \vdash A^* \vdash B$, so $A \mid\!\approx_{\mathcal{M}} B$. But if A^* is inconsistent, by construction, there are no situations in Σ supporting A, so $A \mid\!\sim_{\mathcal{M}} B$ and $A \mid\!\approx_{\mathcal{M}} B$ for all B, as required. Conversely, suppose $A \mid\!\not\sim B$. Then $A^* \nvdash B$, so A^* is consistent, and by (Or), for some prime p_i in the prime decomposition of A^*, $p_i \nvdash B$, so $p_i^* \nvdash B$, i.e. $p_i \mid\!\not\sim B$. Since by construction a situation corresponding to p_i is contained in Σ and the above argument that $\preceq$ is stoppered shows that this is a minimal A-supporting situation, it follows by definition that $A \mid\!\not\approx_{\mathcal{M}} B$. ∎

COROLLARY 4.12
For any plan hierarchy H, the class of finitely representable almost-rational consequence operations based on a plan recognition inference operation derived from H is sound and complete with respect to preferential entailment over the ordered structures of courses of events that are the realizations of the simple plans derived from H.

PROOF. (Soundness) By Theorem 4.11, the consequence operation $\mid\!\sim_{\mathcal{M}}$ defined by preferential entailment from an ordered structure of situations $\mathcal{M} = \langle \Sigma, \preceq \rangle$ by setting $A \mid\!\sim_{\mathcal{M}} B$ iff $A \mid\!\approx_{\mathcal{M}} B$ is a finitely representable almost-rational consequence operation based on $\vdash_\Sigma$, the situated consequence operation defined by entailment over Σ as in Theorem 3.12. Now if every situation in Σ is a course of events that is a realization of a simple plan based on H, it follows by Theorem 3.37 that $\vdash$ is a plan recognition inference operation.

(Completeness) If $\mid\!\sim$ is a finitely representable almost-rational consequence operation based on a plan recognition inference operation $\vdash$, then by Theorem 4.11, there is an ordered structure of situations $\mathcal{M} = \langle \Sigma, \preceq \rangle$ such that $A \mid\!\approx_{\mathcal{M}} B$ iff $A \mid\!\sim B$ and $\Gamma \models_\Sigma A$ iff $\Gamma \vdash A$. This construction uses prime infons p such that $p \equiv p^*$. Since each such p is provably equivalent to a prime consistent set of infons closed under the postulates for plan recognition derived from H, it follows that each such p is provably equivalent to a plan theory, hence by Theorem 3.37, the situations in the model $\mathcal{M}$ are all courses of events that are realizations of a simple plan based on H. Since the simple plans are all representable by the finite number of event and temporal infons they contain, it follows that the situations in $\mathcal{M}$ are all finite, assuming here that the set of interval terms occurring in $\mid\!\sim$ and $\vdash$ is also finite. ∎

COROLLARY 4.19
The class of revision operations on a finitely representable belief system $\langle \mathcal{K}, \vdash \rangle$ is sound and complete with respect to the class of revision operations on ordered structures of situations based on $\vdash$.

PROOF. (Soundness) Given a revision operation on ordered structures of situations based on a situated consequence operation $\vdash$, for any such $\mathcal{M} = \langle \Sigma, \preceq \rangle$, by Theorem 4.11, the consequence operation $\mid\!\sim_{\mathcal{M}}$ defined by setting $A \mid\!\sim_{\mathcal{M}} B$ iff $A \mid\!\approx_{\mathcal{M}} B$ is a finitely representable almost-rational consequence operation based on $\vdash$, which determines an epistemic state $\langle K, \mid\!\sim_{\mathcal{M}} \rangle$, where K consists of the set of infons holding at *min*(*true*) in $\mathcal{M}$. The belief revision operation on epistemic states defined using the Ramsey test and conditionalization corresponds to the revision operation on ordered structures of situations by definition.

(Completeness) Given a revision operation on a finitely representable belief system $\langle \mathcal{K}, \vdash \rangle$, for any epistemic state $\langle \mathcal{K}, \mid\!\sim \rangle$, by Theorem 4.11, the consequence operation $\mid\!\sim$ determines an ordered structure of situations $\mathcal{M} = \langle \Sigma, \preceq \rangle$ such that $A \mid\!\approx_{\mathcal{M}} B$ iff $A \mid\!\sim B$ iff $B \in K^*_A$ (using the Ramsey test). Note here that for each such belief set K, (Consistency Preservation) holds assuming that $\mid\!\sim$ is based on the restriction of $\vdash$ defined so that $A \vdash$ *false* for any infon A for which $K^*_A = K_\perp$. Finally, it again follows that for each such epistemic state, since the consequence operation $\mid\!\sim_A$ satisfies condition (C), the ordered structure of situations $\mathcal{M}^*_A$ is equal to $\langle \Sigma', \preceq' \rangle$ where $\Sigma' = \Sigma \cap [A]$ and $\preceq'$ is $\preceq$ restricted to Σ'. ∎

COROLLARY 4.21

Suppose each epistemic state $\langle K, * \rangle$ is defined from an epistemic state $\langle K, \mid\!\sim \rangle$ using the Ramsey test. Then the belief revision operation from epistemic states $\langle K, * \rangle$ and infons A to epistemic states $\langle K^*_A, *' \rangle$ is an almost-rational belief revision operation.

PROOF. That the belief revision operation satisfies postulates (K*1)–(K*7) is straightforward, these being translations of the postulates for almost-rational consequence operations. Similarly, (K*8) is a translation of condition (C). ∎

Received May 19, 2000