

AGM Revision of Beliefs about Action and Time

Marc van Zee*

University of Luxembourg
Luxembourg
marcvanee@gmail.com

Dragan Doder†

University of Luxembourg
Luxembourg
dragan.doder@uni.lu

Mehdi Dastani

Utrecht University
The Netherlands
M.M.Dastani@uu.nl

Leendert van der Torre

University of Luxembourg
Luxembourg
leendert@vandertorre.com

Abstract

The AGM theory of belief revision is based on propositional belief sets. In this paper we develop a logic for revision of temporal belief bases, containing expressions about temporal propositions (tomorrow it will rain), possibility (it may rain tomorrow), actions (the robot enters the room) and pre- and post-conditions of these actions. We prove the Katsuno-Mendelzon and the Darwiche-Pearl representation theorems by restricting the logic to formulas representing beliefs up to certain time. We illustrate our belief change model through several examples.

1 Introduction

Reasoning about the interplay between action and time is fundamental for the design and development of intelligent systems such as autonomous systems, robotic applications, and service agents [Wooldridge, 2000; Doherty *et al.*, 2009]. Analyzing the behavior of such systems, which are often specified in terms of actions' pre- and post-conditions together with some (behavioral/temporal) constraints, is not only indispensable for checking system properties (e.g., safety and liveness), there are also various applications that require automatic reasoning about the interplay between actions and time. For example, a calendar agent assisting its user in managing her calendar needs to reason about the actions/activities scheduled in time to detect possible conflicts of activities.

On the one hand, reasoning about actions and time has received an overwhelming attention in the last couple of decades [Reiter, 2001; Kvarnström, 2005; Mueller, 2010; Thielscher, 2001; Broersen, 2009], as well as in philosophy [Nuel Belnap, 2001]. On the other hand, the theory of

belief revision has been investigated in detail [Peppas, 2007; Gärdenfors, 2003; Van Benthem, 2011]. However, few attempts have been made to combine these two approaches, i.e. *revision* of theories that are defined in terms of action specifications with behavioral and temporal constraints.

Moreover, virtually all of the attempts to combine logics for reasoning about action and time with belief revision have been restricted purely to the syntactical level, mostly by showing that revision in the proposed logic satisfies the AGM postulates (e.g., [Shapiro *et al.*, 2011; Jin and Thielscher, 2004; Scherl, 2005; Scherl and Levesque, 2003], but see Section 6 for a more detailed discussion). Surprisingly, none of these attempts prove the well-known representations theorems linking revision to a total pre-order on models.

The aim of this paper is to develop a logic about action and time, such that this logic can be used within a more general framework of belief revision. It is this constraint of using the logic within a belief revision setting that drives the design of the logic. Therefore, since belief revision is originally defined for propositional logic, it is our methodology to stay as close to propositional logic as possible.

We first develop the logic Parameterized-time Action Logic (PAL). The language of this logic contains formulas to reason about preconditions, postconditions, and the execution of actions. Atoms in the language are parameterized with the state at which they are true. The main result of the paper is that we prove the Katsuno-Mendelzon (KM) representation theorem and the Darwiche-Pearl (DP) representation theorem in PAL. To this end, we define a revision operator that revises formulas up to a specific time point. We show that this leads to models of system behaviors which can be finitely generated, i.e. be characterized by a single formula. We illustrate our approach by examples and show how it can be applied to various models of time, possibility and action. For example, doing an action may imply its precondition, or it may be equivalent to it.

The rest of this paper is organised as follows. Section 2 is preliminary and introduces belief revision concepts, in Section 3 we introduce PAL, in Section 4 we study belief revision

*Funded by the National Research Fund (FNR), Luxembourg, RationalArchitecture project.

†Funded by the National Research Fund (FNR), Luxembourg, PRIMAT project.

sion in PAL, in Section 5 we provide some examples, and in Section 6 we discuss related work.

2 Preliminaries: Belief Revision

The AGM postulates [Alchourron *et al.*, 1985] formulate properties that should be satisfied by any (rational) revision operators defined on deductively closed sets of propositional formulas. Katsuno and Mendelzon (KM) [1991] represent a belief set B as a propositional formula ψ such that $B = \{\varphi \mid \psi \vdash \varphi\}$. They define the following six postulates for revision on ψ and show that these are equivalent to the eight AGM postulates:

- (R1) $\psi \circ \varphi$ implies φ
- (R2) If $\psi \wedge \varphi$ is satisfiable, then $\psi \circ \varphi \equiv \psi \wedge \varphi$
- (R3) If φ is satisfiable, then $\psi \circ \varphi$ is also satisfiable
- (R4) If $\psi \equiv \psi'$ and $\varphi \equiv \varphi'$, then $\psi \circ \varphi \equiv \psi' \circ \varphi'$
- (R5) $(\psi \circ \varphi) \wedge \varphi'$ implies $\psi \circ (\varphi \wedge \varphi')$
- (R6) If $(\psi \circ \varphi) \wedge \varphi'$ is satisfiable, then $\psi \circ (\varphi \wedge \varphi')$ implies $(\psi \circ \varphi) \wedge \varphi'$

Given a set \mathbb{I} of all interpretations over some propositional language, they define a faithful assignment as a function that assigns each ψ to a pre-order \leq_ψ on models satisfying the following three conditions:

1. If $I, I' \in \text{Mod}(\psi)$, then $I <_\psi I'$ does not hold.
2. If $I \in \text{Mod}(\psi)$ and $I' \notin \text{Mod}(\psi)$, then $I <_\psi I'$ holds.
3. If $\psi \equiv \phi$, then $\leq_\psi = \leq_\phi$.

They show in a representation theorem that a revision operator \circ satisfies postulates (R1)-(R6) iff there exists a faithful assignment that maps each formula ψ to a total preorder \leq_ψ such that

$$\text{Mod}(\psi \circ \varphi) = \min(\text{Mod}(\varphi), \leq_\psi).$$

Darwiche and Pearl (DP) [1997] observe that the AGM postulates are too permissive to enforce plausible iterated revision. In order to remedy this, they suggest two changes to the KM postulates (R1)-(R6):

- Instead of performing revision on a propositional formula, perform revision on an abstract object called an *epistemic state* Ψ , which contains, in addition to the propositional beliefs $\text{Bel}(\Psi)$, the entire information needed for coherent reasoning. Formally, this means that each ψ in (R1)-(R6) is replaced with Ψ , where Ψ means $\text{Bel}(\Psi)$ whenever it is embedded in a propositional formula. So $\Psi \wedge \varphi$ means $\text{Bel}(\Psi) \wedge \varphi$, and $\Psi_1 \equiv \Psi_2$ means $\text{Bel}(\Psi_1) \equiv \text{Bel}(\Psi_2)$, while $\Psi \circ \varphi$ means an (abstract) epistemic state.
- Postulate (R4) is weakened as follows:
(R*4) If $\Psi = \Psi'$ and $\varphi \equiv \varphi'$, then $\Psi \circ \varphi \equiv \Psi' \circ \varphi'$

We refer to the DP postulates defined on epistemic states as (R*1)-(R*6). DP propose the following four additional postulates for a revision operator \circ on epistemic states.

- (C1) If $\varphi \models \varphi'$, then $(\Psi \circ \varphi') \circ \varphi \equiv \Psi \circ \varphi$.

- (C2) If $\varphi \models \neg\varphi'$, then $(\Psi \circ \varphi') \circ \varphi \equiv \Psi \circ \varphi$

- (C3) If $\Psi \circ \varphi \models \varphi'$, then $(\Psi \circ \varphi') \circ \varphi \models \varphi'$

- (C4) If $\Psi \circ \varphi \not\models \neg\varphi'$, then $(\Psi \circ \varphi') \circ \varphi \not\models \neg\varphi'$

They alter the KM definition of a faithful assignment to epistemic states and obtain a representation theorem similar to that of KM above. They show in a second representation theorem that a revision operator \circ_t satisfying their postulates (R*1)-(R*6) satisfies postulates (C1)-(C4) if and only if the operator and its corresponding faithful assignment satisfy:

- (CR1) If $m_1 \models \varphi$ and $m_2 \models \varphi$, then $m_1 \leq_\Psi m_2$ iff $m_1 \leq_{\Psi \circ \varphi} m_2$.
- (CR2) If $m_1 \not\models \varphi$ and $m_2 \not\models \varphi$, then $m_1 \leq_\Psi m_2$ iff $m_1 \leq_{\Psi \circ \varphi} m_2$.
- (CR3) If $m_1 \models \varphi$, $m_2 \not\models \varphi$ and $m_1 <_\Psi m_2$, then $m_1 <_{\Psi \circ \varphi} m_2$.
- (CR4) If $m_1 \models \varphi$, $m_2 \not\models \varphi$ and $m_1 \leq_\Psi m_2$, then $m_1 \leq_{\Psi \circ \varphi} m_2$.

All the above results are based on propositional logic. Logics for reasoning with beliefs about action specifications and time are generally not propositional, but use modalities or high-order logics (see Section 6 for a more detailed discussion). When moving from propositional logic to such more expressive logics, it is in general no longer possible to represent a belief set by a single formula ψ , because it can no longer be guaranteed that the number of models for the logic is finite. This may have different reasons: Time can be infinite, actions can be non-deterministic, or modalities may lead to infinite models. Consequently, the KM and DP representation theorem can no longer be obtained directly for such logics. In the following section, we develop a simple temporal logic and show in the section following it that we can reproduce both the KM and the DP representation theorems for a restricted fragment of this logic.

3 Parameterized-time Action Logic (PAL)

Our aim in this section is to develop a logical system that represents an agent's beliefs about the current moment and future moment and actions that may be performed. Beliefs are represented by the formal language \mathcal{L} .

Definition 1 (Language). Let $\text{Act} = \{a, b, c, \dots\}$ be a finite set of deterministic primitive actions, and $\text{Prop} = \{p, q, r, \dots\} \cup \{\text{pre}(a), \text{post}(a) \mid a \in \text{Act}\}$ be a finite set of propositions.¹ The sets Prop and Act are disjoint. The language \mathcal{L} is inductively defined by the following BNF grammar:

$$\varphi ::= \chi_t \mid \text{do}(a)_t \mid \Box_t \varphi \mid \varphi \wedge \varphi \mid \neg\varphi$$

with $\chi \in \text{Prop}$, $a \in \text{Act}$, and $t \in \mathbb{N}_0$. We abbreviate $\neg\Box_t \neg$ with \Diamond_t , and define $\perp \equiv p_0 \wedge \neg p_0$ and $\top \equiv \neg \perp$. $\text{Past}(t)$ is the set of all PAL formulas generated by boolean combinations of $p_{t'}$, $\text{pre}(a)_{t'}$, $\text{post}(a)_{t'}$, $\Box_{t'} \varphi$, and $\text{do}(a)_{t'-1}$ where $t' \leq t$ and φ is some PAL formula.

Intuitively, p_t means that the atomic formula p is true at time t , $\text{do}(a)_t$ means that action a is executed at time t . To every action and every time we associate formulas $\text{pre}(a)_t$ and $\text{post}(a)_{t+1}$, which are understood as the preconditions and postconditions of action a at time t . The modal operator

¹Throughout this paper we denote atomic propositions with χ .

\Box_t is interpreted as necessity, indexed with a time point t . The other boolean connectives are defined as usual.

The semantics of our logic is similar to CTL* [Reynolds, 2002], namely a tree structure containing nodes and edges connecting the nodes.² With each natural number $i \in \mathbb{N}_0$ we associate a set of states S_i such that all these sets are disjoint. We then define the accessibility relation between states such that it generates an infinite, single tree.

Definition 2 (Tree). A tree is quadruple $T = (S, R, v, act)$ where $S = \bigcup_{n \in \mathbb{N}_0} S_n$ is a set of states, such that each S_t is the set of states at time t , $S_i \cap S_j = \emptyset$ for $i \neq j$; $R \subseteq \bigcup_{n \in \mathbb{N}_0} S_n \times S_{n+1}$ is an accessibility relation that is serial, linearly ordered in the past and connected (so S_0 is a singleton); $v : S \rightarrow 2^{Prop}$ is a valuation function from states to sets of propositions; and $act : R \rightarrow Act$ is a function assigning actions to elements of the accessibility relation, such that actions are deterministic, i.e. if $act((s, s')) = act((s, s''))$, then $s' = s''$.

Definition 3 (Path). Given a tree $T = (S, R, v, act)$, a path $\pi = (s_0, s_1, \dots)$ in T is a sequence of states such that $(s_t, s_{t+1}) \in R$. We write π_t to refer to the t 'th state of the path π and we thus write $v(\pi_t)$ and $act((\pi_t, \pi_{t+1}))$ to refer respectively to the propositions true and the next action on path π at time t . For readability, we abbreviate $act((\pi_t, \pi_{t+1}))$ with $act(\pi_t)$.

Intuitively, $v(\pi_t)$ are the propositions true at time t on path π , and $act(\pi_t)$ is the next action a on the path. There is a natural equivalence relation \sim_t on paths.

Definition 4 (Path equivalence). Two paths π and π' are equivalent up to time t , denoted $\pi \sim_t \pi'$, if and only if they contain the same states up to and including time t , i.e. $\pi \sim_t \pi'$ iff $(\forall t' \leq t). (v(\pi_{t'}) = v(\pi'_{t'}))$ and $(\forall t' < t). (act(\pi_{t'}) = act(\pi'_{t'}))$.

Formulas in PAL are evaluated on a path. Therefore, a model for a formula is pair consisting of a tree and a path in this tree. This, together with some additional constraints related to the pre- and post-conditions of actions, is our definition of a model.

Definition 5 (Model). A model is a pair (T, π) where $T = (S, R, v, act)$ is a tree and π is a path in T such that the following conditions hold:

1. If $act(\pi_t) = a$, then $post(a) \in v(\pi_{t+1})$,
2. If $pre(a) \in v(\pi_t)$, then there is some π' in T with $\pi \sim_t \pi'$ and $act(\pi'_t) = a$.

We denote models with m_1, m_2, \dots and sets of models with M_1, M_2, \dots . We denote the set of all models with \mathbb{M} .

Definition 6 (Truth definitions). Let $m = (T, \pi)$ be a model with $T = (S, R, v, act)$:

- $T, \pi \models \chi_t$ iff $\chi \in v(\pi_t)$ with $\chi \in Prop$
- $T, \pi \models do(a)_t$ iff $act(\pi_t) = a$
- $T, \pi \models \neg\phi$ iff $T, \pi \not\models \phi$

²A tree can equivalently be seen as an enfolded transition system, thereby representing all the possible runs through it. We choose to represent our semantics using trees because it simplifies the completeness proofs. See Reynolds [2002] for an overview of different kinds of semantics and conceptual underpinnings.

$T, \pi \models \phi \wedge \phi'$ iff $T, \pi \models \phi$ and $T, \pi \models \phi'$
 $T, \pi \models \Box_t \phi$ iff for all π' in T : if $\pi' \sim_t \pi$, then $T, \pi' \models \phi$

The logical notions of satisfiability, validity and semantics consequence are defined as usual.

Definition 7. The logic PAL consists of the following axiom schemas and rules:

- Propositional tautologies (PROP)
- $\Box_t(\phi \rightarrow \phi') \rightarrow (\Box_t \phi \rightarrow \Box_t \phi')$ (K)
- $\Box_t \phi \rightarrow \phi$ (T)
- $\Diamond_t \phi \rightarrow \Box_t \Diamond_t \phi$ (5)
- $\chi_t \rightarrow \Box_t \chi_t$, where $\chi \in Prop$ (A1)
- $\Diamond_t \chi_t \rightarrow \chi_t$, where $\chi \in Prop$ (A2)
- $\Box_t \phi \rightarrow \Box_{t+1} \phi$ (A3)
- $do(a)_t \rightarrow \Box_{t+1} do(a)_t$ (A4)
- $\Diamond_{t+1} do(a)_t \rightarrow do(a)_t$ (A5)
- $\bigvee_{a \in Act} do(a)_t$ (A6)
- $do(a)_t \rightarrow \bigwedge_{b \neq a} \neg do(b)_t$ (A7)
- $pre(a)_t \rightarrow \Diamond_t do(a)_t$ (A8)
- $do(a)_t \rightarrow post(a)_{t+1}$ (A9)
- $\Diamond_t (do(a)_t \wedge \alpha) \rightarrow \Box_t (do(a)_t \rightarrow \alpha)$ (A10)
where $\alpha \in Past(t+1)$
- From ϕ , infer $\Box_t \phi$ (NEC)
- From $\phi, \phi \rightarrow \phi'$, infer ϕ' (MP)

The relation \vdash is defined in the usual way with the restriction that (NEC) can be applied to theorems only.

Theorem 1 (Completeness Theorem). The logic PAL is sound and strongly complete, i.e. $T \vdash \phi$ iff $T \models \phi$.

Proof Sketch. We prove the following formulation of completeness: each consistent set of formulas T has a model. We use the fact that each consistent set can be extended to a maximally consistent set (mcs) Σ , i.e. Σ is consistent and each proper superset of Σ is inconsistent. In the first step we extend T to a mcs T^* . We define a tree $M_{T^*} = (S, R, v, a)$:

1. $S = \bigcup_{t \in \mathbb{N}} S_t$ where $S_t = \{[\overline{T}]_t \mid \overline{T} \equiv_t T^*\}$
2. sRs' iff $(\exists \overline{T}^*, t \in \mathbb{N}). (s = [\overline{T}^*]_t \wedge s' = [\overline{T}^*]_{t+1})$
3. $p \in v(s)$ iff $(\exists \overline{T}^*, t \in \mathbb{N}). (s = [\overline{T}^*]_t \wedge p' \in \overline{T}^*)$.
4. $a = act((s, s'))$ iff $(\exists \overline{T}^*). (s = [\overline{T}^*]_t \wedge s' = [\overline{T}^*]_{t+1} \wedge do(a)_t \in \overline{T}^*)$.

where $[\overline{T}^*]_t = \{\overline{T} \mid T^* \equiv_t \overline{T}^*\}$ s.t. \equiv_t is an equivalence relation on mcs's defined by $T_1^* \equiv_t T_2^*$ iff $T_1^* \cap Past(t) = T_2^* \cap Past(t)$. If $\pi(T^*) = (s_0, s_1, \dots)$, where $s_t = [\overline{T}^*]_t$, then one can show that $(M_{T^*}, \pi(T^*))$ is a model. Finally, we prove that for each ϕ , $(M_{T^*}, \pi(T^*)) \models \phi$ iff $\phi \in T^*$, using induction on the depth of the proof. Consequently, $M_{T^*}, \pi(T^*) \models T$. \square

We shortly highlight the most important axiomatization considerations. First, note that exactly one action is executed in each time point (A6 and A7), and that actions are deterministic (A10). Semantically, this means that from a state, one can never reach two different successor states through the same action. The fact that postconditions of actions always hold on a path (A9), but that preconditions may not (A8), suggests that preconditions, unlike postconditions, need not be believed when an action is selected. We might therefore think

of our belief model as, in some sense, one of “optimistic” beliefs. Returning to the calendar agent, it may be well possible that a calendar contains the entry to possibly attend IJCAI in July 2015, i.e. $\Diamond_{Jan2015}do(attend_IJCAI)_{Jul2015}$, while the precondition is not satisfied yet, i.e. the paper is not accepted yet. And even if it would turn out that the paper is accepted, there may be no budget yet, or no flight booked. Therefore, we allow actions to be believed while preconditions are not satisfied (yet). We will return to these considerations in more detail in Section 5.

4 Belief Revision in PAL

Time in PAL is infinite in the future, so it is generally not possible to represent a PAL belief set closed under consequence by a single formula ψ , since this may potentially lead to an infinite conjunction. Therefore, we cannot prove the KM and DP representation theorems directly. In this section, we define a bounded revision function and we restrict the syntax and semantics of PAL up to a specific time point. We then prove two representation theorems.

Single-Step Revision We define a bounded revision function $*_t$ revising a set of PAL formulas B with another PAL formula ϕ , denoted $B *_t \phi$, where t is the maximal time point occurring in both B and ϕ . In order to do so, we first define the set of formulas $Form_t$ containing all PAL formulas with time points no larger than t , and Bel_t containing all sets of $Form_t$ formulas.

Definition 8 (AGM t -Bounded Revision Function). *Suppose some $t \in \mathbb{N}_0$. Let $\max_t(\phi)$ denote the maximal time point occurring in ϕ . Let $Form_t = \{\phi \in \mathcal{L} \mid \max_t(\phi) \leq t\}$ and $Bel_t = \{Cl(S) \mid S \subseteq Form_t\}$.*

An AGM t -bounded revision function $*_t : Bel_t \times Form_t \rightarrow Bel_t$ maps a deductively closed set from Bel_t and a formula from $Form_t$ to deductively closed set from Bel_t , and satisfies the AGM postulates [Alchourron et al., 1985].

We next define a bounded model as a pair consisting of a tree and a path, in which all paths in the tree are cut off up to a specific time point.

Definition 9 (t -Bounded Model). *Suppose some model $m = (T, \pi)$. A t -bounded path $\bar{\pi}[t]$ is defined from a path π in T as $\bar{\pi}[t] = (\pi_0, \dots, \pi_t)$. A t -bounded model $\bar{m}[t] = (\bar{T}[t], \bar{\pi}[t])$ is a pair where $\bar{T}[t] = \{\bar{\pi}[t] \mid \pi \in T\}$. $\bar{\mathbb{M}}[t]$ is the set of all t -bounded models. If the time point t is irrelevant or clear from the context, we may abbreviate $\bar{\pi}[t], \bar{T}[t], \bar{m}[t]$ and $\bar{\mathbb{M}}[t]$ with $\bar{\pi}, \bar{T}, \bar{m}$, and $\bar{\mathbb{M}}$, respectively.*

We will show in Lemma 4 below that it is possible to represent each $B \in Bel_t$ by a formula $\psi \in Form_t$ such that $Cl(B) = Cl(\psi)$. Using this lemma, we adapt the definition of a KM faithful assignment.

Definition 10 (t -Bounded Faithful Assignment). *A t -bounded faithful assignment is a function that maps each belief formula $\psi \in Form_t$ to a total preorder \leq^t_ψ on all models such that:*

1. If $m_1, m_2 \in Mod(\psi)$, then $m_1 \leq^t_\psi m_2$ and $m_2 \leq^t_\psi m_1$
2. If $m_1 \in Mod(\psi)$ and $m_2 \notin Mod(\psi)$, then $m_1 <^t_\psi m_2$

3. If $\psi \equiv \phi$, then $\leq^t_\psi = \leq^t_\phi$

4. If $\bar{m}_1 = \bar{m}_2$, then $m_1 \leq^t_\psi m_2$ and $m_2 \leq^t_\psi m_1$

Since we only consider beliefs up to some time t , we do not want to distinguish between models that are the same up to time t in the total pre-order \leq^t_ψ . This is essentially what condition (4) of the faithful assignment above ensures. The first three conditions are the same as those by KM.

Theorem 2 (Representation Theorem). *A t -bounded revision operator \circ_t satisfies postulates (R1)-(R6) iff there exists a t -bounded faithful assignment that maps each belief set ψ to a total preorder \leq^t_ψ such that*

$$Mod(\psi \circ_t \phi) = \min(Mod(\phi), \leq^t_\psi).$$

We will use the remainder of this subsection to prove the representation theorem above. We first show that the number of t -bounded model is finite.

Lemma 1. *For each $t \in \mathbb{N}_0$, $\bar{\mathbb{M}}[t]$ is finite.*

Proof. Suppose some $t \in \mathbb{N}_0$. Since actions are deterministic and there are finitely many actions in our logic, each state has a finite number of successor states. Moreover, since there are finitely many propositions in our language, the number of possible valuations of the states is finite as well. Therefore, the number of models in $\bar{\mathbb{M}}[t]$ is finite. \square

The following lemma obtains a correspondence between semantic consequence of two models equivalent up to t . The proof is by induction on the depth of the formula.

Lemma 2. *For each $\phi \in Form_t$ and models $m_1, m_2 \in \bar{\mathbb{M}}$, if $\bar{m}_1[t] = \bar{m}_2[t]$, then $m_1 \models \phi$ iff $m_2 \models \phi$.*

Let $Ext(\bar{m})$ be the set of all possible extensions of the bounded model \bar{m} to models, i.e. $Ext(\bar{m}) = \{m' \in \bar{\mathbb{M}} \mid \bar{m}'[t] = \bar{m}\}$. We next show that we can represent each $Ext(\bar{m})$ by a single formula.

Lemma 3. *For all $t \in \mathbb{N}_0$ and $\bar{m} \in \bar{\mathbb{M}}[t]$, there exists a formula $form(\bar{m}) \in Form_t$ such that $Mod(form(\bar{m})) = Ext(\bar{m})$.*

Proof. For $t \in \mathbb{N}_0$ and $m = (T, \pi)$ we define:

$$\alpha_{\bar{\pi}} = \bigwedge_{n=0}^t \left(\bigwedge_{\chi \in v(\bar{\pi}_n)} \chi_n \wedge \bigwedge_{\chi \notin v(\bar{\pi}_n)} \neg \chi_n \wedge \bigwedge_{act(\bar{\pi}_n)=a} do(a)_n \wedge \bigwedge_{act(\bar{\pi}_n) \neq a} \neg do(a)_n \right).$$

Then we define:

$$form(\bar{m}) = \alpha_{\bar{\pi}} \wedge \bigwedge_{\bar{\pi}' \in \bar{T}} \Diamond_0 \alpha_{\bar{\pi}'} \wedge \bigwedge_{\bar{\pi}' \notin \bar{T}} \neg \Diamond_0 \alpha_{\bar{\pi}'}$$

It follows directly from our construction that m is a model of $form(\bar{m})$. By Lemma 2 for each $m' \in Ext(\bar{m}) : m' \models form(\bar{m})$. What remains to show is that if $m' = (T', \pi')$ and $\bar{m} \neq \bar{m}'$, then $m' \not\models form(\bar{m})$. Since $\bar{m}' \neq \bar{m}$ we have (1) $\bar{T} \neq \bar{T}'$ or (2) $\bar{\pi} \neq \bar{\pi}'$. In case (1), if there exists some path $\bar{\pi}'' \in \bar{T}$ s.t. $\bar{\pi}'' \notin \bar{T}'$, then the formula $\Diamond_0 \alpha_{\bar{\pi}''}$ is true in m but not in m' . But then $form(\bar{m})$ is also not true in

m' , hence $m' \not\models \text{form}(\bar{m})$. If there exists some path $\bar{\pi}'' \in \bar{T}'$ and $\bar{\pi}'' \notin \bar{T}$, then $\neg \diamond_0 \alpha_{\bar{\pi}''}$ is true in m but not in m' , so $m' \not\models \text{form}(\bar{m})$ follows as well. In case (2) $\alpha_{\bar{\pi}} \neq \alpha_{\bar{\pi}'}$ and thus $m' \not\models \text{form}(\bar{m})$. \square

We can now show that it is indeed possible to represent a belief base in Bel_t as a single formula.

Lemma 4. *For each $B \in Bel_t$, there exists a formula ψ such that $B = Cl(\psi)$.*

Proof. Suppose an arbitrary $B \in Bel_t$. It follows from Lemma 1 that the number of t -bounded models for B is finite. Therefore, using Lemma 3 we define $\psi = \bigvee_{\bar{m}[t] \models B} \text{form}(\bar{m})$. So, the set of t -bounded models of ψ is the union of the set of t -bounded models of B . Therefore, $\bar{m}[t] \models B$ iff $\bar{m}[t] \models \psi$. By Lemma 2, $\bar{m} \models B$ iff $\bar{m} \models \psi$, so $B \models \phi$ iff $\psi \models \phi$. Finally, by the completeness theorem we can conclude $B \vdash \phi$ iff $\psi \vdash \phi$. \square

We next provide a proof sketch of Theorem 2.

Proof Sketch (Theorem 2). (\Rightarrow): We define a t -bounded faithful assignment \leq_{ψ}^t as: $m_1 \leq_{\psi}^t m_2$ iff $m_1 \in Mod(\psi)$ or $m_1 \in Mod(\psi \circ_t \text{form}(\bar{m}) \vee \text{form}(\bar{m}'))$. Let us prove the new condition (4) of Def. 10. Let m and m' be such that $\bar{m} = \bar{m}'$. By Lemm 3, $Mod(\text{form}(\bar{m})) = Ext(\bar{m}) = Ext(\bar{m}') = Mod(\text{form}(\bar{m}'))$. Hence, $\text{form}(\bar{m}) \equiv \text{form}(\bar{m}')$, so $\text{form}(\bar{m}, \bar{m}') \equiv \text{form}(\bar{m})$. By (R4): $Mod(\psi \circ_t \text{form}(\bar{m})) = Mod(\psi \circ_t \text{form}(\bar{m}, \bar{m}'))$. By (R1), $m \in Mod(\psi \circ_t \text{form}(\bar{m}))$, so $m \in Mod(\psi \circ_t \text{form}(\bar{m}, \bar{m}'))$. Hence, by the definition of \leq_{ψ}^t : $m \leq_{\psi}^t m'$. We can prove $m' \leq_{\psi}^t m$ similarly.

(\Leftarrow): We show that there exists some $\psi \circ_t \phi \subseteq Form_t$ s.t. $Mod(\psi \circ_t \phi) = \min(Mod(\phi), \leq_{\psi}^t)$. Suppose that ϕ is consistent ($\phi \equiv \perp$ is trivial). If $m \in \min(Mod(\phi), \leq_{\psi}^t)$, by condition (4) we obtain $Ext(\phi) \subseteq \min(Mod(\alpha), \leq_{\psi}^t)$. Thus, there exists a set $S \subseteq \bar{M}[t]$ s.t. $\min(Mod(\phi), \leq_{\psi}^t) = \bigcup_{\bar{m} \in S} Ext(\bar{m})$. If $\psi \circ_t \phi$ is the formula $\bigvee_{\bar{m} \in S} \text{form}(\bar{m})$, then $Mod(\psi \circ_t \phi) = \bigcup_{\bar{m} \in S} Ext(\bar{m})$ by Lemma 3. So there exists such $\psi \circ_t \phi$. For the other parts of the proof we modify KM straightforwardly. \square

Iterated Revision Recall from Section 2 that Darwiche and Pearl propose to revise epistemic states instead of belief bases. When switching from belief revision on a belief state to belief revision on an epistemic state the definition of a faithful assignment should be adopted accordingly. We will do this now for our setting. Recall that Ψ stands for $Bel(\Psi)$ whenever it is embedded in a propositional formula.

Definition 11 (t -bounded faithful assignment on epistemic states). *A t -bounded faithful assignment on epistemic states maps each epistemic state Ψ with $Bel(\Psi)$ from Bel_t , to a total pre-order \leq_{Ψ}^t on all models such that:*

1. If $m_1, m_2 \models \Psi$, then $m_1 \leq_{\Psi}^t m_2$ and $m_2 \leq_{\Psi}^t m_1$
2. If $m_1 \models \Psi$ and $m_1 \not\models \Psi$, then $m_1 <_{\Psi}^t m_2$
3. If $\Psi = \Phi$, then $\leq_{\Psi} = \leq_{\Phi}$
4. If $\bar{m}_1 = \bar{m}_2$, then $m_1 \leq_{\Psi}^t m_2$ and $m_2 \leq_{\Psi}^t m_1$

The next representation theorem is similar to Theorem 2, but it uses the DP postulates on epistemic states and the t -bounded faithful assignment on epistemic states that we defined above. Recall that we refer to the DP postulates as (R*1)-(R*6).

Theorem 3 (Representation Theorem). *A t -bounded revision operator \circ_t satisfies postulates (R*1)-(R*6) precisely when there exists a t -bounded faithful assignment on epistemic states that maps each epistemic state Ψ to a total pre-order \leq_{Ψ}^t such that*

$$Mod(\Psi \circ_t \phi) = \min(Mod(\phi), \leq_{\Psi}^t)$$

Proof Sketch. Straightforward modification of the proof of Theorem 2. \square

Theorem 4. *Suppose that a t -bounded revision operator on epistemic states satisfies postulates (R*1)-(R*6). The operator satisfies postulates (C1)-(C4) iff the operator and its corresponding faithful assignment satisfy:*

- CR1 If $m_1 \models \phi$ and $m_2 \models \phi$, then $m_1 \leq_{\Psi}^t m_2$ iff $m_1 \leq_{\Psi \circ \phi}^t m_2$.
- CR2 If $m_1 \not\models \phi$ and $m_2 \not\models \phi$, then $m_1 \leq_{\Psi}^t m_2$ iff $m_1 \leq_{\Psi \circ \phi}^t m_2$.
- CR3 If $m_1 \models \phi$, $m_2 \not\models \phi$ and $m_1 <_{\Psi}^t m_2$, then $m_1 <_{\Psi \circ \phi}^t m_2$.
- CR4 If $m_1 \models \phi$, $m_2 \not\models \phi$ and $m_1 \leq_{\Psi}^t m_2$, then $m_1 \leq_{\Psi \circ \phi}^t m_2$.

Proof Sketch. (\Leftarrow): Identical to DP.

(\Rightarrow): We prove (C1) \Rightarrow (CR1), the proofs for the other postulates are similar. Suppose that (C1) holds and $m, m' \models \phi$. Let $\alpha = \text{form}(\bar{m}, \bar{m}')$. By Lemma Lemma 3, $Mod(\alpha) = Ext(\bar{m}) \cup Ext(\bar{m}')$. If $m'' \in Ext(\bar{m})$, then $\bar{m} = \bar{m}'$, so by Lemma Lemma 2 we obtain $m'' \models \phi$. Similarly, if $m'' \in Ext(\bar{m}')$, then $m'' \models \phi$, so $\alpha \models \phi$. By (C1), $(\Psi \circ_t \phi) \circ_t \alpha \equiv \Psi \circ_t \alpha$. In other words, $\min(Mod(\alpha), \leq_{\Psi \circ \phi}^t) = \min(Mod(\alpha), \leq_{\Psi}^t)$. Consequently, $m \leq_{\Psi \circ \phi}^t m'$ iff $m \leq_{\Psi}^t m'$. \square

5 Examples

We start with a simple example illustrating revision using a pre-order over models in our temporal logic.

Example 1. *Suppose a setting with two actions, i.e. $Act = \{a, b\}$ and no propositions other than the pre- and postconditions for these actions, i.e. $Prop = \{pre(a), post(a), pre(b), post(b)\}$. Suppose that Ψ is equivalent to*

$$do(a)_0 \wedge pre(a)_0 \wedge \neg pre(b)_0 \wedge \neg post(a)_0 \wedge \neg post(b)_0 \wedge \neg \diamond_0 pre(a)_1 \wedge \neg \diamond_0 pre(b)_1 \wedge \neg post(b)_1.$$

The three 1-bounded models of Ψ are $\{\bar{m}_1, \bar{m}_2, \bar{m}_3\}$. See Figure 1, note that formulas that are false in a state have been omitted. If we define \leq_{Ψ}^t as $m_1 \leq_{\Psi}^t m_2$ iff $m_1 \in Mod(\Psi)$ or $m_1 \in Mod(\Psi \circ_t (\text{form}(\bar{m}_1) \vee \text{form}(\bar{m}_2)))$, then these three models are minimal. Suppose we for instance revise with $\neg \diamond_0 post(b)_1$. Then, bounded models \bar{m}_2 and \bar{m}_3 are removed, i.e. $Mod[1](\Psi \circ_t \neg \diamond_0 post(b)_1) = \{\bar{m}_1\}$, while revising with $\diamond_0 do(b)_0$ will remove model \bar{m}_1 , since in this model it is not possible that action b is executed.




Figure 1: Example 1-bounded models, where bold lines indicate actual paths.

Our next example considers the relation between preconditions and actions in the axiomatization of PAL.

Example 2. *The calendar agent believes that IJCAI is attended in July 2015 because a paper has been submitted. Now, the agent learns in March that the paper to IJCAI is rejected, which causes the agent to believe that it will be impossible to attend IJCAI. However, the agent upholds its belief that IJCAI is attended in July. Let t_1 and t_2 as before.*

$$\Psi \equiv do(a)_{t_2}$$

$$\Psi \circ_{t_2} \square_{t_1} \neg pre(a)_{t_2} \equiv do(a)_{t_2} \wedge \neg pre(a)_{t_2}$$

As we already briefly mentioned in Section 3, in our current model preconditions are sufficient conditions of actions to be possible, but they are not necessary. Such a model can be used in planning domains in which it, at the moment of planning, may not be clear whether preconditions are true [Shoham, 2009]. While such optimistic beliefs allow an agent to reason about its actions while not being committed to accept the preconditions of these action, this assumption may in some cases be considered rather weak. We can straightforwardly strengthen this assumption by replacing axiom (A8) with the following:

$$pre(a)_t \leftrightarrow \diamond_t do(a)_t \quad (A8')$$

Completeness is preserved by changing the condition (2) of Definition 5 from an “if” to an “if and only if”. Indeed, it can be verified that in the previous example adopting the belief that the precondition to attend IJCAI is impossible will cause the agent to drop its belief that it will attend IJCAI.

6 Related Work

Many logical systems have been developed for reasoning about the pre and postconditions of actions with explicit time points, such as the Event Calculus [Mueller, 2010], Temporal Action Logics [Kvarnström, 2005], extensions to the Fluent Calculus [Thielscher, 2001], and extensions to the Situation Calculus [Papadakis and Plexousakis, 2003] (see [Patkos, 2010, Ch.2] for an overview). Much effort in this field concentrates on extending these action theories to incorporate *sensing* or *knowledge-producing actions*, i.e., actions whose

effects change the mental state of an agent instead of the world. Shapiro [2011] extends the Situation Calculus to reason about beliefs rather than knowledge by introducing a modality B and shows that both the AGM postulates and the DP postulates are satisfied in this framework. A similar approach concerning the Fluent Calculus has been formalized by Jin and Thielscher [2004], and is further developed by Scherl [2005] and Scherl and Levesque [2003] by taking into account the frame problem as well. Although these approaches all obtain correspondences with the AGM or DP postulates on a syntactical level, none of them prove representations theorems linking revision to a total pre-order on models.

Baral and Zhang [2005] model belief updates on the basis of semantics of modal logic S5 by specifying an update according to the minimal change on both the agent’s actual world and knowledge. They show that their knowledge update operator satisfies all the KM postulates. Bonanno [2007] combines temporal logic with AGM belief revision as well, and extends a temporal logic with a belief operator and an information operator. Both these approaches do not take action or time into account and do not provide a representation theorem.

Semantically, PAL is close to CTL* [Reynolds, 2002], dating back to Prior [1967] with three important differences. First, PAL contains time-indexed modalities, which allows one to express statements such as “It is possible in February that I will attend IJCAI in July”. Secondly, PAL allows for explicit reasoning about pre- and postconditions of deterministic actions (axiom A10). Thirdly, PAL only contains one type of modality, while CTL* contains a next operator and an until operator as well. The last two reasons arguably make PAL less expressive than CTL*, but in return the axiomatization is straightforward, we are able to obtain strong completeness, and it is possible to prove both representation theorems. It seems that this is not possible for CTL* in general.

The logic PAL is also closely related to the logic developed by Icard *et al.* [2010], who study the joint revision of beliefs and intentions using AGM-like postulates. However, we show in a recent paper that one of their axioms is not sound and that their logic is noncompact, so no axiomatization using their syntax and semantics (including theirs) is strongly complete [van Zee *et al.*, 2015]. Moreover, we focus on revision of beliefs up to time t and iterated revision.

7 Summary and Outlook

We present a temporal logic for reasoning about beliefs and action in time. To apply the Katsuno-Mendelzon and Darwiche-Pearl representation theorems to belief sets of this logic, we restrict ourselves to beliefs generated by formulas which represent beliefs up to certain time. Using this restriction, we prove both representation theorems for our logic.

The main challenge of this paper is to define revision in a temporal logic with infinite time. While much research has focused on extending existing logics for reasoning about action and change with the incorporation of new beliefs, the semantic counterpart in the form of a total pre-order on models is not taken into account. This paper has tried to take a first

step into this direction, but there are many directions from here that should be explored:

The frame problem is one of the most fundamental problems in reasoning about action and change. The challenge is how to specify the non-effects of actions succinctly. For instance, if a proposition such as “The table is red” is true at some time t , and no action occurs that affects the truth value of this proposition, then it seems plausible that the table is still red at time $t + 1$. Currently, our framework does not contain so-called “frame axioms”, and it is an interesting question how the AGM and DP postulates can be extended to account for such behavior as well.

Other mental attitudes, such as intentions, goals, desires and preferences, we have left out completely. This is not because we assume that they are unimportant, but because it was our goal to focus on belief revision in a temporal setting. However, returning to our calendar assistant, there is nothing keeping us from treating the appointments that a user makes as “intentions”, and given the recent work in the revision of intentions [Shoham, 2009; van der Hoek *et al.*, 2007; Icard *et al.*, 2010; Lorini *et al.*, 2009] this certainly is a direction worth exploring.

References

- [Alchourron *et al.*, 1985] Carlos E. Alchourron, Peter Gardenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 06 1985.
- [Baral and Zhang, 2005] Chitta Baral and Yan Zhang. Knowledge updates: Semantics and complexity issues. *Artificial Intelligence*, 164(1):209–243, 2005.
- [Bonanno, 2007] Giacomo Bonanno. Axiomatic characterization of the AGM theory of belief revision in a temporal logic. *Artificial Intelligence*, 171(2):144–160, 2007.
- [Broersen, 2009] Jan Broersen. A complete stit logic for knowledge and action, and some of its applications. In *Declarative Agent Languages and Technologies VI*, pages 47–59. Springer, 2009.
- [Darwiche and Pearl, 1997] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89(12):1 – 29, 1997.
- [Doherty *et al.*, 2009] Patrick Doherty, Jonas Kvarnström, and Fredrik Heintz. A Temporal Logic-based Planning and Execution Monitoring Framework for Unmanned Aircraft Systems. *Autonomous Agents and Multi-Agent Systems*, 19(3):332–377, 2009.
- [Gärdenfors, 2003] Peter Gärdenfors. *Belief revision*, volume 29. Cambridge University Press, 2003.
- [Icard *et al.*, 2010] Thomas Icard, Eric Pacuit, and Yoav Shoham. Joint revision of belief and intention. *Proc. of the 12th International Conference on Knowledge Representation*, pages 572–574, 2010.
- [Jin and Thielscher, 2004] Yi Jin and Michael Thielscher. Representing beliefs in the fluent calculus. In Ramon Lopez de Mntaras and Lorenza Saitta, editors, *ECAI*, pages 823–827. IOS Press, 2004.
- [Katsuno and Mendelzon, 1991] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, dec 1991.
- [Kvarnström, 2005] Jonas Kvarnström. *TALplanner and other extensions to Temporal Action Logic*. PhD thesis, Linköpings universitet, 2005.
- [Lorini *et al.*, 2009] Emiliano Lorini, Mehdi Dastani, Hans P. van Ditmarsch, Andreas Herzig, and John-Jules Ch. Meyer. Intentions and assignments. In *LORI*, volume 5834 of *Lecture Notes in Computer Science*, pages 198–211. Springer, 2009.
- [Mueller, 2010] Erik T Mueller. *Commonsense reasoning*. Morgan Kaufmann, 2010.
- [Nuel Belnap, 2001] Ming Xu Nuel Belnap, Michael Perloff. *Facing the Future : Agents and Choices in Our Indeterminist World: Agents and Choices in Our Indeterminist World*. Oxford University Press, USA, 2001.
- [Papadakis and Plexousakis, 2003] Nikos Papadakis and Dimitris Plexousakis. Actions with duration and constraints: The ramification problem in temporal databases. *International Journal on Artificial Intelligence Tools*, 12(3):315–353, 2003.
- [Patkos, 2010] Theodore Patkos. *A formal theory for reasoning about action, knowledge and time*. PhD thesis, University of Crete-Heraklion, 2010.
- [Peppas, 2007] P. Peppas. *Handbook of Knowledge Representation*, chapter Belief Revision. Elsevier, 2007.
- [Prior, 1967] Arthur N Prior. *Past, present and future*, volume 154. Clarendon Press Oxford, 1967.
- [Reiter, 2001] Raymond Reiter. *Knowledge in Action. Logical Foundations for Specifying and Implementing Dynamical Systems*. 2001.
- [Reynolds, 2002] M. Reynolds. An axiomatization of full computation tree logic. *Journal of Symbolic Logic*, 66(3):1011–1057, 2002.
- [Scherl and Levesque, 2003] Richard B Scherl and Hector J Levesque. Knowledge, action, and the frame problem. *Artificial Intelligence*, 144(1):1–39, 2003.
- [Scherl, 2005] Richard B Scherl. Action, belief change and the frame problem: A fluent calculus approach. In *Proceedings of the Sixth workshop on Nonmonotonic Reasoning, Action, and Change at IJCAI*, 2005.
- [Shapiro *et al.*, 2011] Steven Shapiro, Maurice Pagnucco, Yves Lesprance, and Hector J. Levesque. Iterated belief change in the situation calculus. *Artificial Intelligence*, 175(1):165–192, 2011.
- [Shoham, 2009] Yoav Shoham. Logical theories of intention and the database perspective. *Journal of Philosophical Logic*, 2009.
- [Thielscher, 2001] Michael Thielscher. The concurrent, continuous fluent calculus. *Studia Logica*, 67(3):315–331, 2001.
- [Van Benthem, 2011] Johan Van Benthem. *Logical dynamics of information and interaction*. Cambridge University Press, 2011.
- [van der Hoek *et al.*, 2007] Wiebe van der Hoek, Wojciech Jamroga, and Michael Wooldridge. Towards a theory of intention revision. *Synthese*, 155(2):265–290, February 2007.
- [van Zee *et al.*, 2015] Marc van Zee, Mehdi Dastani, Dragan Doder, and Leendert van der Torre. Consistency conditions for beliefs and intentions. In *Twelfth International Symposium on Logical Formalizations of Commonsense Reasoning*, March 2015.
- [Wooldridge, 2000] M.J. Wooldridge. *Reasoning about Rational Agents*. MIT Press, 2000.