

AGM-Style Revision of Beliefs and Intentions

Marc van Zee and Dragan Doder
Computer Science and Communication
University of Luxembourg

marcvanee@gmail.com, dragan.doder@uni.lu

Abstract. We introduce a logic for temporal beliefs and intentions based on Shoham’s database perspective and we formalize his coherence conditions on beliefs and intentions. In order to do this we separate strong beliefs from weak beliefs. Strong beliefs are independent from intentions, while weak beliefs are obtained by adding intentions to strong beliefs and everything that follows from that. We provide AGM-style postulates for the revision of strong beliefs and intentions: strong belief revision may trigger intention revision, but intention revision may only trigger revision of weak beliefs. After revision, the strong beliefs are coherent with the intentions. We show in a representation theorem that a revision operator satisfying our postulates can be represented by a pre-order on interpretations of the beliefs, together with a selection function for the intentions.

1 Introduction

Recently there has been an increase in articles studying the dynamics of intentions in logic [8, 11, 16, 27, 15, 10]. Most of those papers take as a starting point the logical frameworks derived from Cohen and Levesque [6], which in turn formalize Bratman’s [4] planning theory of intention. In this paper, we take a different starting point, and study the revision of intentions from a *database perspective* [24]. The database perspective consists of a planner, a belief database and an intention database. Shoham [25] describes it as “(...) a generalization of the AGM scheme for belief revision, (...). In the AGM framework, the intelligent database is responsible for storing the planner’s beliefs and ensuring their consistency. In the enriched framework, there are two databases, one for beliefs and one for intentions, which are responsible for maintaining not only their individual consistency but also their mutual consistency.” (p.48) Shoham further developed these ideas with Jacob Banks, one of his PhD students, and behavioral economist Dan Ariely in the intelligent calendar application Timeful, which attracted over \$6.8 million in funding and was acquired by Google in 2015¹, who aim to integrate it into their Calendar applications. As Shoham [25] says himself: “The point of the story is there is a direct link between the original journal paper and the ultimate success of the company.” (p.47) Thus, it seems clear that his philosophical proposal has led to some success on the practical side. In this paper, we investigate whether his proposal can lead to interesting theoretical insights as well. More specifically, the aim of this paper is to develop a suitable formal theory for the belief and intention databases in the database perspective, to formalize the coherence conditions that Shoham puts on the databases, and to study belief and intention revision for this theory. Following Shoham’s proposal, our

methodology is to generalize AGM revision [1] for temporal beliefs and intentions and to prove a representation theorem.

In the area of intention revision and reconsideration, Grant *et al.* [10] combine intention revision with AGM-like postulates. There have also been a number of contributions applying AGM-style revision to action logics [23, 12, 21, 22, 3]. However, these proposals only characterize revision using a set of postulates, without proving representation theorems. There are also approaches that focus on the semantical level by postulating revision on a Kripke model [2].

In this paper, we first review two recent formalisations based on Shoham’s database perspective, namely the IPS (Icard-Pacuit-Shoham) framework [11] and our own previous work PAL (Parameterized-time Action Logic) [29], and we discuss the shortcomings of these logics. We extend PAL in order to formalize Shoham’s database perspective. An overview of our approach is displayed in Figure 1. We separate strong beliefs from weak beliefs. Strong beliefs are beliefs that are independent of the intentions of the agent, while weak beliefs are those beliefs that are obtained by adding the consequences of all intended actions to the strong beliefs, and everything that follows from that. A planner (outside of our agent) may add beliefs and intentions. The belief database consists of strong beliefs, and can only be updated by a strong belief formula. Revision of strong beliefs affects weak beliefs but, according to Shoham’s database perspective, they may only remove intentions. Intention revision may in turn trigger revision of the weak beliefs. The main result is to characterize this revision process correctly through postulates and to prove a representation theorem.²

The structure of this paper is as follows. Section 2 is preliminary and introduces Shoham’s database perspective, the IPS framework, and PAL. In Section 3 we formalize strong and weak beliefs, and we formalize Shoham’s coherence condition in Section 4. In Section 5 we study revision of beliefs and intentions, and in Section 6 we discuss related work.

2 Preliminaries: The Database Perspective

We start by introducing our running example that we will use frequently throughout the paper.

Example 1 (Running Example) *An agent located in Luxembourg is considering to attend the IJCAI conference in New York City, NY (USA) in July 2016 and the ECAI conference in The Hague (the Netherlands) in August 2016. Although it would like to attend both*

¹ <http://venturebeat.com/2015/05/04/google-acquires-scheduling-app-timeful-and-plans-to-integrate-it-into-google-apps/>

² **NOTE TO REVIEWERS:** We provide the full proofs of all theorems and propositions in this paper in a separate technical report: <http://www.marcvanzee.nl/proofs/ecai2016.pdf>.

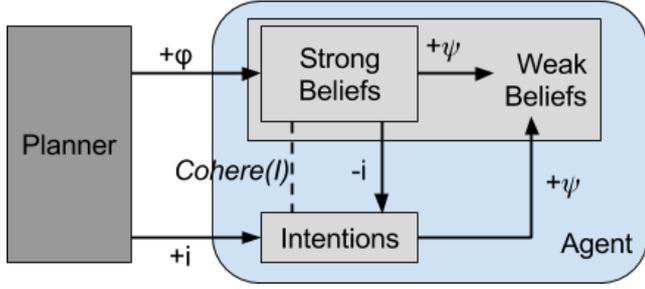


Figure 1: Our formalization of the database perspective. Strong beliefs are denoted by ϕ , weak beliefs by ψ , and intentions by i . An arrow indicates that a component can add (+) or remove (-) a formula from another component. The dashed line represents a coherence condition on strong beliefs and intentions. Weak beliefs are obtained from strong beliefs by adding the consequences of all intentions. The planner is considered a black box and can update strong beliefs and intentions. Strong beliefs may update intentions, but intentions may only update weak beliefs.

events, there is insufficient budget available for traveling. The agent thus believes that it is possible to attend IJCAI at time 0 (July 2016) and that it is possible to attend ECAI at time 1 (August 2016), but also believes that it is impossible to attend both conferences. If the agent decides to attend IJCAI, then it would like to combine this with a visit to a colleague in New York at time 2 (September 2016).

2.1 Shoham’s Database Perspective

Shoham’s database perspective [24] contains a planner (e.g., a STRIPS-like planner) that is itself engaged in some form of practical reasoning. In the course of planning, it may add actions to be taken at various times in the future to an intention database and add observations to a belief database. The intentions are *future-directed intentions* of the form (a, t) , meaning that action a will be executed at time t .³ The beliefs are also time-indexed, and are of the form p_t , meaning that p is true at time t . Shoham treats the planner as a “black box”: It provides the databases with input but its internal workings are unknown. Shoham proposes informal revision procedures for beliefs and intentions based on the following coherence conditions:

1. Beliefs must be internally consistent.
2. Intentions must be internally consistent.
 - (a) At most one action can be intended for a given time moment.
 - (b) If two intended actions immediately follow one another, the earlier cannot have postconditions that are inconsistent with the preconditions of the latter.
3. Intentions must be consistent with beliefs.
 - (a) If you intend to take an action you cannot believe that its preconditions do not hold.⁴
 - (b) If you intend to take an action, you believe that its postconditions hold.

³ The notion of intention here is rather restrictive and important characteristics of intentions are missing. See the conclusion for a discussion.

⁴ Shoham notes here that “it is important to distinguish between the time of belief, and the time to which the belief refers. When an intention to act at time t_2 is added at time t_1 (with $t_1 < t_2$), then at time t_1 it is believed that right after the action is taken at t_2 its postconditions will hold.” [24, p.7]

Note that requirement 3a and 3b describe an asymmetry between pre- and postconditions: The postconditions are believed to be true after an intended action, but the preconditions may not. Therefore, we might think of the requirements as one of “optimistic” beliefs. According to Shoham [24]: “It is a good fit with how planners operate. Adopting an optimistic stance, they feel free to add intended actions so long as they are consistent with current beliefs.” (p. 7)

Shoham then “sketches” informal revision procedures, in which belief revision may trigger intention revision and visa versa, potentially leading to a long cascade of changes. Facts that are believed because they are postconditions of currently held intentions are annotated as such, because “if the intention is withdrawn then the belief in the postcondition can be eliminated as well.” (p. 8)

2.2 Icard *et al.* (IPS)

Icard *et al.* [11] develop a “formal semantical model to capture action, belief and intention, based on the ‘database perspective’” (p.1). They assume a set of atomic sentences $\text{Prop} = \{p, q, r, \dots\}$ and deterministic primitive actions $\text{Act} = \{a, b, c, \dots\}$. Entries in the belief database are represented by a language generated from:

$$\phi := p_t \mid pre(a)_t \mid post(a)_t \mid Do(a)_t \mid \Box\phi \mid \phi \wedge \phi \mid \neg\phi$$

with $p \in \text{Prop}$, $a \in \text{Act}$, and $t \in \mathbb{Z}$. p_t means that p is true at time t , $Do(a)_t$ means that the agent does action a at time t , and $pre(a)_t$ and $post(a)_t$ represent respectively the precondition and postcondition of action a at time t .

Icard *et al.* use a semantics of *appropriate paths*. They define $P = \mathcal{P}(\text{Prop} \cup \{pre(a), post(a) : a \in \text{Act}\})$, and a *path* $\pi : \mathbb{Z} \rightarrow (P \times \text{Act})$ as a mapping from a time point to a set of proposition-like formulas true at that time (denoted $\pi(t)_1$) and the next action a on the path (denoted $\pi(t)_2$). They define an equivalence relation $\pi \sim_t \pi'$, which means that π and π' represent the same situation up to t . Using this, they propose a notion of appropriateness:

Definition 1 (Appropriate Set of Paths) A set of paths Π is appropriate iff for all $\pi \in \Pi$:

- If $\pi(t)_2 = a$, then $post(a) \in \pi(t+1)_1$,
- If $pre(a) \in \pi(t)_1$, then there exists $\pi' \sim_t \pi$ s.t. $\pi'(t)_2 = a$.

The truth definition \models_{Π} is defined relative to an appropriate set of paths Π , and the modality is defined as follows:

$$\pi, t \models_{\Pi} \Box\phi, \text{ iff for all } \pi' \in \Pi, \text{ if } \pi \sim_t \pi' \text{ then } \pi', t \models \phi.$$

A model for a formula is an appropriate set of paths. They introduce an intention database $I = \{(a, t), \dots\}$ as a set of action-time pairs (a, t) and put the following coherence condition on their logic:

$$Cohere^*(I) = \Diamond \bigwedge_{(a,t) \in I} pre(a)_t.$$

This captures the intuition that an agent considers it possible to carry out all intended actions. They state that a pair (Π, I) is coherent if and only if there exists a path in Π in which $Cohere^*(I)$ is true. IPS distinguishes *intention-contingent*, or weak, beliefs from *non-contingent*, or strong, beliefs. Contingent beliefs B^I are obtained from a belief-intention database (B, I) as follows:

$$B^I = Cl(B \cup \{Do(a)_t : (a, t) \in I\}).$$

In order to switch from belief bases to an appropriate set of paths, Icard *et al.* introduce the functions ρ and β : “Given a set of formulas

B , we can consider the set of paths on which all formulas of B hold at time 0, denoted $\rho(B)$. Conversely, given a set of paths Π , we let $\beta(\Pi)$ be defined as the set of formulas valid at 0 in all paths in Π .³ (p.3)

The first issue with IPS is that the definition of non-contingent beliefs is problematic for coherence. We demonstrate this using our running example.

Example 2 (Non-contingent beliefs in IPS) *Suppose that the agent believes it has to go to the ECAI conference at time 1 because the event will be held in The Hague, which is very close to Luxembourg (where the agent is located) and thus is cheaper for traveling. However, the agent has the intention to go to IJCAI. We can formalize this in the IPS framework as a belief-intention base (B, I) with $\{post(ECAI)_2, \neg\Diamond(do(IJCAI)_0 \wedge post(ECAI)_2)\} \subseteq B$ and $(IJCAI, 0) \in I$. The coherence condition holds because $Cohere^*(I) = \Diamond(pre(IJCAI)_0)$ is consistent with B . However, the contingent beliefs B^I are inconsistent since $\neg\Diamond(do(IJCAI)_0 \wedge post(ECAI)_2) \wedge do(IJCAI)_0$ implies $\neg post(ECAI)_2$, but this is inconsistent with the initial belief $post(ECAI)_2$.*

The example above shows that contingent beliefs may become inconsistent if the agent has intentions that are conflicting with what it believes non-contingently. While the agent believes that it will not go to IJCAI, it still has the intention to do so. This is coherent according to $Cohere^*(I)$, since the agent doesn't believe that it is impossible to go to IJCAI. However, adding the consequences of the intention to go to IJCAI results in a conflict with the fact that the agent believes it will attend ECAI. The source of the problem is that non-contingent beliefs in the IPS framework such as $post(ECAI)_2$ are actually not non-contingent, because they are contingent on the actions that occur in the path (namely to attend ECAI in time 1). We solve this problem in Section 3 by requiring that strong beliefs are always prefixed by a modal operator. In this way, the beliefs are about possibility and necessity, but they are not beliefs about a specific future.

The second issue with IPS is that the definition of ρ is circular, and as a result it does not seem to be possible to apply it to all formulas of their logic. The IPS definition of ρ is $\rho(B) = \{\pi \mid \pi \models_{\Pi} B\}$. That is, the set of paths for a belief base B is those paths in which all formulas of B are true, given an appropriate set of paths Π . However, the function ρ should construct the appropriate set of paths. Therefore, the definition of ρ is circular. It seems that the function ρ only works for belief bases containing no modalities. We omit details for space constraints, but the construction of the canonical model in the proof of their representation theorem uses the function ρ to switch from a belief base to a set of paths (see Proof Sketch in the Appendix of Icard *et al.* [11]). Therefore, the representation theorem does not hold for all formulas of the logic, since it is not possible to apply the function ρ to all beliefs.

Summarizing, we recognize two shortcomings of the IPS framework as a formal basis for the database perspective: the definition of contingent beliefs is problematic, and the representation theorem does not hold for belief bases containing modalities.

2.3 Parameterized-time Action Logic

In our previous work, we develop Parameterized-time Action Logic (PAL) [29, 30] as an alternative to IPS. PAL differs syntactically from IPS in that the \Box -modality is indexed by a time-point, and semantically in that it uses a standard branching time logic semantics.

Definition 2 (PAL Language [30]) *Let $Act = \{a, b, c, \dots\}$ be a finite set of deterministic primitive actions, and $Prop = \{p, q, r, \dots\} \cup \{pre(a), post(a) \mid a \in Act\}$ be a finite set of propositions.⁵ The sets $Prop$ and Act are disjoint. The language \mathcal{L}_{PAL} is inductively defined by the following BNF grammar:*

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

with $\chi \in Prop, a \in Act$, and $t \in \mathbb{N}$. We abbreviate $\neg\Box_t\neg$ with \Diamond_t , and define $\perp \equiv p_0 \wedge \neg p_0$ and $\top \equiv \neg \perp$.

PAL uses a CTL*-like [20] tree semantics consisting of a tree $T = (S, R, v, act)$ where S is a set of states, R is an accessibility relation that is serial, linearly ordered in the past and connected, $v: S \rightarrow 2^{Prop}$ is a valuation function from states to sets of propositions, and $act: R \rightarrow Act$ is a function from accessibility relations to actions, such that actions are deterministic, i.e. if $act((s, s')) = act((s, s''))$, then $s' = s''$.

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 $(s_t, s_{t+1}) \in R$. The formula π_t refers to the t 'th state of the path π , so $v(\pi_t)$ and $act((\pi_t, \pi_{t+1}))$ refer respectively to the propositions true and the next action on path π at time t . For readability, $act((\pi_t, \pi_{t+1}))$ is abbreviated with $act(\pi_t)$.

Similarly to IPS, we define an equivalence relation on paths: 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 .

Definition 3 (Model [30]) *A model is a pair (T, π) where $T = (S, R, v, act)$ is a tree and π is a path in T , and for all $\pi \in T$ 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 $act(\pi'_t) = a$.*

Definition 4 (Truth definitions) *Let (T, π) be a model with $T = (S, R, v, act)$:*

$$\begin{aligned} T, \pi &\models \chi_t \text{ iff } \chi \in v(\pi_t) \text{ with } \chi \in Prop \\ T, \pi &\models do(a)_t \text{ iff } act(\pi_t) = a \\ T, \pi &\models \neg \varphi \text{ iff } T, \pi \not\models \varphi \\ T, \pi &\models \varphi \wedge \varphi' \text{ iff } T, \pi \models \varphi \text{ and } T, \pi \models \varphi' \\ T, \pi &\models \Box_t \varphi \text{ iff for all } \pi' \text{ in } T: \text{ if } \pi' \sim_t \pi, \text{ then } T, \pi' \models \varphi \end{aligned}$$

We axiomatize PAL and show that it is sound and strongly complete, i.e. $T \vdash \varphi$ iff $T \models \varphi$.

Furthermore, we characterize AGM belief revision in this logic by bounding all inputs and output of the revision process up to some time t . Using these constraints, we are able to represent a belief set B as a propositional formula ψ such that $B = \{\varphi \mid \psi \vdash \varphi\}$ and we prove the Katsuno and Mendelzon (KM) [13] representation theorem.

We now formalize the beliefs of our agent in the running example in PAL.

Example 3 (PAL model) *A possible partial PAL model (T, π_2) of the beliefs of our conference planning agent is shown in Figure 2, where the thick path represents the actual path. In the actual path, the agent believes it attends IJCAI at time 0 and visit a colleague at time 1. It also considers it possible to do nothing at time 0 and attend ECAI at time 1 in an alternative path. However, it does not consider it to be possible to attend both conferences.⁶ Some formulas that are true in Figure 2 are:*

⁵ Throughout this paper we denote atomic propositions with χ .

⁶ Note that $pre(nop) \equiv post(nop) \equiv \top$. They have been omitted from the figure for readability.

$$\begin{aligned}
T, \pi_3 & \models \neg do(visit)_1 \\
T, \pi_2 & \models post(visit)_2 \rightarrow \neg do(ECAI)_1 \\
T, \pi_1 & \models \diamond_0(do(IJCAI)_0 \wedge \neg do(visit)_1) \\
T, \pi_2 & \models \diamond_0 \square_1 do(ECAI)_1 \\
T, \pi_1 & \models \neg \diamond_0(do(IJCAI)_0 \wedge do(ECAI)_1)
\end{aligned}$$

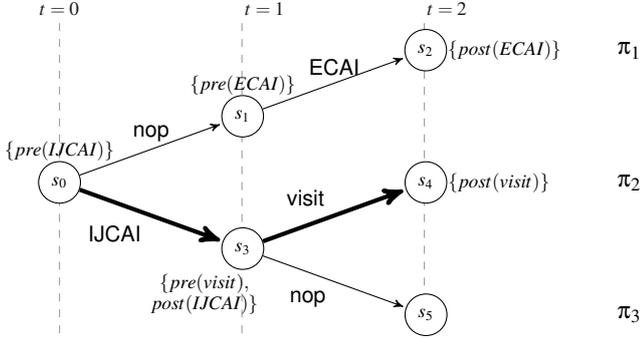


Figure 2: Example PAL Model (T, π_2) from $t = 0$ to $t = 2$.

In our previous work, we only consider the revision of PAL formulas. However, in the current approach we distinguish between strong and weak beliefs, and define the belief database as a set of strong beliefs (Figure 1). This means that our previous approach cannot be applied directly. We cannot apply it to strong beliefs, because it is not ensured that after revision we end up with strong beliefs again. We can also not apply it to the beliefs that result as consequence of intentions (weak beliefs), because revision of these beliefs should not affect strong beliefs. We explain the details of strong and weak beliefs in more detail in the next section.

3 Formalizing Strong and Weak Beliefs

Suppose the conference planning agent of our running example adopts the intention to attend IJCAI at time 0. When adopting this intentions, the agent will form new beliefs based on the success of this intentions, e.g., that it will be in New York at time 1. These further beliefs can be used in the course of further planning, for instance it may adopt the intention to visit a colleague in New York at time 1. If the agent then learns that its flight from Luxembourg to New York is canceled, it should drop the intention to attend IJCAI. Yet, by dropping this intention that was based on the now-dropped belief, other beliefs, including the belief that it will be in New York at time 1, should also be dropped, which may in turn force the intention to visit a colleague to be dropped. And so on. Thus, belief revision may trigger intention revision, which again can trigger belief revision, etc. However, intention revision should not change the “basic” beliefs of an agent. For instance, if the agent adopts the intention to attend IJCAI at time 0, and it does not annotate the beliefs based on this intention in any way, then the planner may believe that it is no longer necessary to attend IJCAI in order to be in New York at time 1, and thus decide to retract the intention. This is the so-called Little Nell paradox, and has been discussed extensively in the literature [17, 5, 28]. Shoham proposed as well to annotate postconditions of intentions in coherence condition 3b in order to separate them from ordinary beliefs (see Section 2.1).

In order to deal with these nuances, we separate strong and weak beliefs (this is the terminology used by van der Hoek *et al.* [28] as

well). The idea behind strong beliefs is that they represent the agent’s ideas about what is inevitable, no matter how it would act in the world. They thus represent the agent’s view on the current situation and the future within which it can plan its action. Formally, strong beliefs at some time t are formulas that start either with \diamond_t or \square_t .

Definition 5 (Strong Beliefs) *The set of all strong beliefs \mathbb{B}_t in time t for \mathcal{L}_{PAL} is inductively defined by the following BNF grammar:*

$$\varphi ::= \square_t \psi \mid \varphi \wedge \varphi \mid \neg \varphi,$$

where $\psi \in \mathcal{L}_{PAL}$ and $t \in \mathbb{N}$. A strong belief set B_t is a set of strong beliefs closed under consequence, i.e. $B_t = Cn(B')$ where $B' \subseteq \mathbb{B}_t$. In the remainder of this paper, we assume $t = 0$ and we simply write \mathbb{B} and B to abbreviate \mathbb{B}_0 and B_0 .

The weak beliefs $WB(B, I)$ are obtained from strong beliefs B by adding beliefs that are contingent on the intentions I . In other words, weak beliefs are all the strong beliefs and moreover the consequences of all intentions, and everything that follows from this.

Definition 6 (Weak Beliefs) *Given a pair (B, I) , the weak beliefs are defined as:*

$$WB(B, I) = Cn(B \cup \{do(a)_t \mid (a, t) \in I\}).$$

Let us formalize these notions in our running example.

Example 4 (Strong and weak beliefs) *Let (B, I) be such that the beliefs up to time 2 are represented by the model in Figure 2⁷, and let $I = \{(IJCAI, 0), (visit, 1)\}$. Some examples of strong beliefs are $\diamond_0 do(IJCAI)_0$ (“it is possible at time 0 to attend IJCAI at time 0”), $\diamond_0 do(ECAI)_1$ (“it is possible at time 0 to attend ECAI at time 1”), and $\neg \diamond_0 (post(ECAI)_2 \wedge post(IJCAI)_1)$ (“It is not possible that the postconditions of attending ECAI and IJCAI are both true, respectively at time 2 and 1”). Some examples of weak beliefs are $do(IJCAI)_1$, $post(ijcai)_2$, and $post(visit)_2 \rightarrow do(IJCAI)_0$.*

The reader may already have noted that, semantically, strong beliefs are independent of the specific path on which they are true. For instance, returning to Example 3, the belief $\neg \diamond_0 (do(IJCAI)_0 \wedge do(ECAI)_1)$ is true in path π_1 , but it is also true in all other paths of the model. This seems to indicate that the models of a set of strong beliefs are rather sets of trees instead of sets of tree-path pairs. It captures the intuition that strong beliefs are not dependent on a specific future, but are about possibility and necessity. In Section 5, we will make this property more precise and use it in order to characterize revision of strong beliefs.

4 Formalizing Shoham’s Coherence Conditions

In this section we formalize Shoham’s coherence condition in our framework. We first demonstrate that $Cohere^*(I)$ of IPS (Section 2.2) is too permissive because it allows models in which intentions are not jointly executable.

Example 5 (Coherence in IPS) *Suppose the agent of the running example has two intentions: $I = \{(IJCAI, 0), (ECAI, 1)\}$. Intuitively,*

⁷ In other words, all models are the same up to $t = 2$. This single model property simplifies the exposition of our framework and we assume this throughout the paper.

the agent's intentions do not cohere with its beliefs, because it believes it cannot execute them both due to insufficient budget. However, according to $\text{Cohere}^*(I)$ the agent is coherent because the preconditions of all intentions hold on some path (namely the actual path of Figure 2). This is because while the precondition to attend IJCAI is true at time 0, the agent only executes this action in a path different from the path in which it attends ECAI as well.

Thus, $\text{Cohere}^*(I)$ does not fulfill Shoham's coherence condition 2b (Section 2.1). Although the preconditions of the intended actions are true on a path, the intentions are not jointly executable because the postcondition of the first action is inconsistent with the precondition of the second. More specifically, the problem is that it is not possible to define the precondition of a set of actions in terms of preconditions of individual actions, because it cannot be ensured that all the intentions are fulfilled on the same path as well. Therefore, in order to formalize a coherence condition in PAL, we extend the language with preconditions of finite action sequences, which ensures that after executing the first action, the precondition for the remaining actions are still true. We modify the language, the definition of a model, the axiomatization, and we show that the new axiomatization is sound and strongly complete. We call the new logic PAL-P (Parameterized-time Action Logic with extended Preconditions).

Definition 7 (PAL-P Language) The language \mathcal{L} is obtained from \mathcal{L}_{PAL} by adding $\{pre(a, b, \dots)_t \mid \{a, b, \dots\} \subseteq \text{Act}, t \in \mathbb{N}\}$ to the set of propositions.

We also extend the definition of a model accordingly.

Definition 8 (PAL-P Model) A model is a pair (T, π) with $T = (S, R, v, \text{act})$ such that for all $\pi \in T$ the following holds:

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

We refer to models of PAL-P with m_1, m_2, \dots , we refer to sets of models with M_1, M_2, \dots , and we refer to the set of all models with \mathbb{M} .

Definition 9 The logic PAL-P consists of the all the axiom schemas and rules of PAL [30] (Def. 7), and the following two:

$$pre(\dots, a, b)_t \rightarrow pre(\dots, a)_t \quad (A11)$$

$$(pre(a, b, \dots)_t \wedge do(a)_t) \rightarrow pre(b, \dots)_{t+1} \quad (A12)$$

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

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

Note that it is not directly possible in PAL-P to express preconditions for actions that do not occur directly after each other. In order to do so, we simply make a disjunction over all possible action combinations in the time points in between the actions. Thus, if for instance $\text{Act} = \{a, b\}$ and $I = \{(a, 1), (b, 3)\}$, then $\text{Cohere}(I) = \diamond_0 \bigvee_{x \in \text{Act}} pre(a, x, b)_1 = \diamond_0 (pre(a, a, b)_1 \vee pre(a, b, b)_1)$.⁸

⁸ Our construction of preconditions over action sequences may lead to a coherence condition involving a big disjunction. This is a drawback in terms of computational complexity. Alternatively, one may explicitly denote the time of each precondition, e.g. $pre(a, b)_{(t_1, t_2)}$. We chose the former since it is conceptually closer to the original syntax, but the latter can be implemented straightforwardly.

Definition 10 (Coherence) Given an intention database $I = \{(b_{t_1}, t_1), \dots, (b_{t_n}, t_n)\}$ with $t_1 < \dots < t_n$, let

$$\text{Cohere}(I) = \diamond_0 \bigvee_{\substack{a_k \in \text{Act}: k \notin \{t_1, \dots, t_n\} \\ a_k = b_k: k \in \{t_1, \dots, t_n\}}} pre(a_{t_1}, a_{t_1+1}, \dots, a_{t_n})_{t_1}.$$

For a given set of models M , we say that (M, I) is coherent iff there exists some $m \in M$ with $m \models \text{Cohere}(I)$. For a given agent $A = (B, I)$, we say that the A is coherent iff B is consistent with $\text{Cohere}(I)$, i.e., $B \not\vdash \neg \text{Cohere}(I)$.

Naturally, if a set of intentions is coherent with a set of strong beliefs, then its subset is coherent as well.

Lemma 1 if $I' \subseteq I$, then $\text{Cohere}(I) \vdash \text{Cohere}(I')$.

The following proposition was originally proposed by Icard *et al.*, and holds in our framework.

Proposition 1 If (B, I) is coherent, then $WB(B, I)$ is consistent.

Let us demonstrate the difference between $\text{Cohere}^*(I)$ of IPS and our $\text{Cohere}^*(I)$ using the running example. It turns out the above proposition does not hold for IPS.

Example 6 (Comparing the coherence conditions) Recall from Example 5 that the model of Figure 2 is coherent with intentions $I = \{(IJCAI, 0), (ECAI, 1)\}$ according to $\text{Cohere}^*(I)$ of the IPS framework. However, the weak beliefs of this agent are inconsistent because the agent believes it is impossible to execute both intentions. Thus, Proposition 1 does not hold here.

On the other hand, the model is not coherent according to $\text{Cohere}(I)$, because the agent does not have the possibility to jointly execute both intentions (i.e. the preconditions for both actions together is false). Thus, none of the models satisfy $\diamond_0 pre(IJCAI, ECAI)_0$, even though they satisfy $\diamond_0 pre(IJCAI)_0 \wedge \diamond_0 pre(ECAI)_1$. Thus, we see that in this case Proposition 1 holds.

We now discuss Shoham's coherence conditions. Condition 1 is a direct consequence of Proposition 1. Condition 2a follows from our notion of coherence: if $a \neq b$, $\{(a, t), (b, t)\} \subseteq I$, and B is some set of strong beliefs, then $WB(B, I) \vdash do(a)_t \wedge do(b)_t$. On the other hand, $do(a)_t \rightarrow \neg do(b)_t$ is an of PAL-P (Axiom A7, Def. 7 [29]), so the weak beliefs are inconsistent. By Proposition 1, B is not coherent with I . We show that condition 2b holds in the following proposition.

Proposition 2 (Coherence Condition 2b) If $\{(a, t), (b, t+1)\} \subseteq I$ and B is a set of strong beliefs such that (B, I) coherent, then $\{pre(b)_{t+1}, post(a)_{t+1}\}$ is consistent with B .

Condition 3a follows from our coherence condition: If (B, I) is coherent, i.e. $B \not\vdash \neg \text{Cohere}(I)$, then it follows by Lemma 1 that for all $(a, t) \in I$, $B \not\vdash \neg \diamond_0 pre(a)_t$ (since $\text{Cohere}(\{(a, t)\}) = \diamond_0 pre(a)_t$). Finally, condition 3b is formalized using our notion of weak beliefs and axiom A9 of PAL-P: $do(a)_t \rightarrow post(a)_t$ (Def. 7 [29]).

5 Belief and Intention Revision

In this section we formalize the revision procedures on an agent. That is, we formalize all the arrows from Figure 1. Recall that we distinguish between the revision of strong beliefs, the revision of intentions, and the revision of weak beliefs, which is a consequence of both. The revision of strong beliefs may trigger intention revision, while intention revision only triggers the revision of weak beliefs.

5.1 Revision Postulates

Following KM, we fix a way of representing a belief set B consisting of strong beliefs by a propositional formula ψ such that $B = \{\varphi \mid \psi \vdash \varphi\}$. Since intentions and beliefs that have been added by a planner are naturally bounded up to some time point t , we define a bounded revision function and we restrict the syntax and semantics of PAL-P up to a specific time point. As a consequence, it is then possible to obtain the single formula ψ for a set of strong beliefs B (Corollary 1). The difficulty of obtaining this result is that when revising a belief database that is bounded up to some time t with a strong belief, we have to ensure that the resulting belief database is also bounded up to t , and that it remains a strong belief. We first define some notation that we use in the rest of this paper.

Definition 11 An agent is a pair (ψ, I) consisting of a belief formula ψ , and an intention base I . \mathbb{A} denotes the set of all agents, \mathbb{B} denotes the set of all strong beliefs, \mathbb{I} denotes the set of all intentions, and \mathbb{D} denotes the set of all intention databases. We denote $\mathbb{A}, \mathbb{B}, \mathbb{I}$, and \mathbb{D} bounded up to t with respectively $\mathbb{A}^t, \mathbb{B}^t, \mathbb{I}^t$, and \mathbb{D}^t . However, if the restriction is clear from context, we may omit the superscript notation. We define ε as a special “empty” intention.

We now define a bounded revision function $*_t$ revising an agent (ψ, I) with a tuple (φ, i) consisting of a strong belief φ and an intention i , denoted $(\psi, I) *_t (\varphi, i)$, where t is the maximal time point occurring in ψ, I, φ , and i .

Definition 12 (Agent Revision Function) An Agent revision function $*_t : \mathbb{A} \times (\mathbb{B} \times \mathbb{I}) \rightarrow \mathbb{A}$ maps an agent, a strong belief formula, and an intention—all bounded up to t —to an agent bounded up to t such that if,

$$(\psi, I) *_t (\varphi, i) = (\psi', I'),$$

$$(\psi_2, I_2) *_t (\varphi_2, i_2) = (\psi'_2, I'_2),$$

then following postulates hold:

(P1) ψ' implies φ .

(P2) If $\psi \wedge \varphi$ is satisfiable, then $\psi' \equiv \psi \wedge \varphi$.

(P3) If φ is satisfiable, then ψ' is also satisfiable.

(P4) If $\psi \equiv \psi_2$ and $\varphi \equiv \varphi_2$ then $\psi' \equiv \psi'_2$.

(P5) If $\psi \equiv \psi_2$ and $\varphi_2 \equiv \varphi \wedge \varphi'$ then $\psi' \wedge \varphi'$ implies ψ'_2 .

(P6) If $\psi \equiv \psi_2$, $\varphi_2 \equiv \varphi \wedge \varphi'$, and $\psi' \wedge \varphi'$ is satisfiable, then ψ'_2 implies $\psi' \wedge \varphi'$.

(P7) (ψ', I') is coherent.

(P8) If $(\psi', \{i\})$ is coherent, then $i \in I'$.

(P9) If $(\psi', I \cup \{i\})$ is coherent, then $I \cup \{i\} \subseteq I'$.

(P10) $I' \subseteq I \cup \{i\}$.

(P11) If $I = I_2$, $i = i_2$, and $\psi' \equiv \psi'_2$, then $I' = I'_2$.

(P12) For all I'' with $I' \subset I'' \subseteq I \cup \{i\}$: (ψ', I'') is not coherent.

Postulates (P1)-(P6) are simply the KM postulates in our setting, which are equivalent to the AGM postulates [13]. They also state that the revision of strong beliefs does not depend on the intentions. Postulates (P7)-(P10) also appear in IPS. Postulate (P7) states that the outcome of a revision should be coherent. Postulate (P8) states that the new intention i take precedence over all other current intentions; if possible, it should be added, even if all current intentions have to be discarded. Postulate (P9) and (P10) together state that if it is possible to simply add the intention, then this is the only change that is made. Postulate (P11) states that if we revise with the same i but with a different belief, and we end up with the same belief in both cases, then we also end up with the same intentions. Finally, (P12) states that we do not discard intentions unnecessarily. This last postulate

is comparable to the *parsimony requirement* introduced by Grant *et al.* [10].

We now discuss our revision function in some more detail, starting with a simple example.

Example 7 (Adding an intention) Suppose we have an agent $A = (\psi, I)$ such that all models of the strong beliefs B are the same as the partial model in Figure 2 up to $t = 2$, and suppose that $I = \{(IJCAI, 0), (\text{visit}, 1)\}$. That is, that agent has the intention to attend IJCAI at time 0 and then visit a colleague at time 1. Suppose now that the agent changes its intention to attend ECAI at time 1. Formally: $(\psi, I) *_2 (\top, i) = (\psi', I')$ with $i = (ECAI, 1)$. Since $(\psi, \{i\})$ is coherent but (ψ, I') is not, from (P8) and (P9) we obtain $i \in I'$. Furthermore, from (P10) we have that $I' \subseteq \{(IJCAI, 0), (\text{visit}, 1), (ECAI, 1)\}$. Combining this gives $I' = \{(ECAI, 1)\}$ as the only coherent outcome. Thus, the agent no longer intends to attend IJCAI and to visit the colleague. Note that, although the strong beliefs didn't change after revising with the new intention, the weak beliefs did change. For example, $\text{post}(IJCAI)_1 \in \text{WB}(\psi, I) \setminus \text{WB}(\psi, I')$ and $\text{post}(ECAI)_2 \in \text{WB}(\psi, I') \setminus \text{WB}(\psi, I)$.

The revision function $*_t$ takes a tuple (φ, i) as input, and the postulates (P1)-(P7) ensure that revision of strong beliefs occurs prior to the revision of intentions. Therefore, it may seem plausible that revising with (φ, i) is the same as first revising with (φ, ε) and then with (\top, i) . In other words, the following postulate seems to follow:

$$\begin{aligned} &\text{If } (\psi, I) *_t (\varphi, i) = (\psi', I') \\ &\text{and } ((\psi, I) *_t (\varphi, \varepsilon)) *_t (\top, i) = (\psi'', I''), \quad (\text{P13}^*) \\ &\text{then } \psi' \equiv \psi'' \text{ and } I' \equiv I''. \end{aligned}$$

However, this property does not follow from (P1)-(P12), and we show in the following example that adding the postulate would in fact conflict with the maximality postulate (P12).

Example 8 (Joint vs separate revision) Suppose some agent $A = (\psi, I)$ with beliefs up to $t = 2$ corresponding to the model on the left of Figure 3. The agent has the possible actions to go to the dentist (d) or to stay at work (w), and then go eating (e) or go to the movies (m). Before revision, the agent has intentions $I = \{(d, 0), (e, 1)\}$ (left image of Figure 3, intentions shown as bold lines). It then revises with the belief that it cannot go eating after going to the dentist (φ) and as a result with the intention to go to the movie at time 1 ($i = (m, 1)$). The resulting strong beliefs after revising with φ are shown on the right of Figure 3.

Let us first analyze joint revision. That is, $(\psi, I) *_2 (\varphi, i) = (\psi', I')$. Both $(\psi', \{(d, 0), (m, 1)\})$ and $(\psi', \{(m, 1)\})$ are coherent, so by the maximality postulate (P12), $I' = \{(d, 0), (m, 1)\}$. Hence, the agent intends to go to the dentist and go to the movie.

For separate revision, let $(\psi, I) *_2 (\varphi, \varepsilon) = (\psi', \bar{I})$ and $(\psi', \bar{I}) *_2 (\top, i) = (\psi', \bar{I}')$ (note that ψ' is the same as for joint revision, by (P4)). Now, since $(\psi', \{(d, 0)\})$ and $(\psi', \{(e, 1)\})$ are both coherent, we either have $\bar{I} = \{(d, 0)\}$ or $\bar{I} = \{(e, 1)\}$. Suppose that $\bar{I} = \{(e, 1)\}$. In that case, since $(\psi', \{(e, 1), (m, 1)\})$ is incoherent, we obtain $\bar{I}' = \{(m, 1)\}$ by the postulates (P8) and (P10). Thus, (P13*) doesn't hold.

In separate revision, the agent has to choose whether to go eating or to go to the dentist after revising beliefs. When it chooses to go eating, it has to drop this intention again when deciding to go to the movies, since these two intentions conflict. In joint revision, this is not the case since the agent can compare going to the movies with both possibilities and choose the biggest set of intentions.

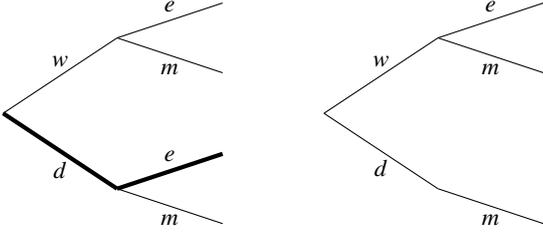


Figure 3: Left: Partial model of strong beliefs ψ of agent (ψ, I) with $I = \{(d, 0), (e, 1)\}$ (bold lines). Right: Revised strong beliefs of agent after learning its no possible to eat (e) after the dentist (d).

The following proposition states that revising with nothing doesn't change a coherent agent. That is, the new beliefs are equivalent to the one prior to revision and the intention database is unchanged.

Proposition 3 *Suppose an agent (ψ, I) is coherent, and $(\psi, I) *_t (\top, \varepsilon) = (\psi', I')$. Then $\psi \equiv \psi'$ and $I = I'$.*

5.2 Representation Theorem

We next characterize all revision schemes satisfying (P1)-(P12) in terms of minimal change with respect to an ordering among interpretations and a selection function accommodating new intentions while restoring coherence. We bound models of strong beliefs up to t , which means that all the paths in the model are “cut off” at t . This ensures finitely many non-equivalent formulas for some belief set B . A t -bounded model $m^t = (T^t, \pi^t)$ is a model containing a tree T in which all paths, including π , have length t . Strong beliefs are about possibility and necessity, and they are independent of a specific path. Therefore, if a single path in a tree is a model of a strong belief, then all paths in this tree are models of this strong belief. Formally, a set of models of a strong belief M_{SB} satisfies the following condition:

$$\text{If } (T, \pi) \in M_{SB}, \text{ then } (T, \pi') \in M_{SB} \text{ for all } \pi' \in T.$$

A set of t -bounded models of a strong belief M_{SB}^t contains only t restricted models of a strong belief. We write \mathbb{M}_{SB}^t to denote the set of all sets of t -bounded models of strong beliefs. We now show that we can represent a set of models of strong beliefs by a single formula.

Lemma 2 *Let $\text{Ext}(M_{SB}^t)$ be the set of all possible extensions of a set of bounded model of strong beliefs M_{SB}^t to models, i.e. $\text{Ext}(M_{SB}^t) = \{m \in \mathbb{M} \mid m^t \in M_{SB}^t\}$. Given a set of t -bounded models of strong beliefs M_{SB}^t , there exists a strong belief formula $\text{form}(M_{SB}^t)$ such that $\text{Mod}(\text{form}(M_{SB}^t)) = \text{Ext}(M_{SB}^t)$.*

Corollary 1 *Given a t -bounded strong belief set B , there exists a formula ψ such that $B = \{\varphi \mid \psi \vdash \varphi\}$.*

Proof Sketch. For a given belief set B , we can show that there exists a set of t -bounded models of a strong belief M_{SB}^t s.t. $\text{Ext}(M_{SB}^t) = \text{Mod}(B)$. If $\psi = \text{form}(M_{SB}^t)$, then $\text{Mod}(\psi) = \text{Mod}(B)$, and by the completeness theorem, $B = \text{Cl}(\psi)$. \square

Given an intention database I , we define a selection function γ_t^I that tries to accommodate a new intention based on strong beliefs. The selection function specifies preferences on which intention an agent would like to keep in the presence of the new beliefs.

Definition 13 (Selection Function) *Given an intention database I , a selection function $\gamma_t^I : \mathbb{M}_{SB} \times \mathbb{I} \rightarrow \mathbb{D}$ maps a set of models of a strong belief and an intention to an updated intention database—all bounded up to t —such that if $\gamma_t^I(M^t, \{i\}) = I'$, then:*

1. (M^t, I') is coherent.
2. If $(M^t, \{i\})$ is coherent, then $i \in I'$.
3. If $(M^t, I \cup \{i\})$ is coherent, then $I \cup \{i\} \subseteq I'$.
4. $I' \subseteq I \cup \{i\}$.
5. For all I'' with $I' \subset I'' \subseteq I \cup \{i\}$: (M^t, I'') is not coherent.

The five conditions on the selection function are in direct correspondence with postulates (P7)-(P10), (P12) of the agent revision function $*_t$. Note that postulate (P11) doesn't have a corresponding condition in the definition above but is represented by the fact that the selection function takes the revised beliefs as input. That is, intention revision occurs after belief revision.

The following proposition states that a selection function does not change intentions unnecessarily. That is, if an intention is already in the intention database, or if it's empty, the intention database remains unchanged.

Proposition 4 *Given some coherent pair (M, I) , if $i = \varepsilon$ or $i \in I$, then $\gamma_t^I(M, i) = I$.*

KM define a faithful assignment from a belief formula to a pre-order over models. Since we are also considering intentions, we extend this definition such that it also maps intentions databases to selection functions.

Definition 14 (Faithful assignment) *A faithful assignment is a function that assigns to each strong belief formula $\psi \in \mathbb{B}^t$ a total pre-order \leq_{ψ}^t over \mathbb{M} and to each intention database $I \in \mathbb{D}^t$ a selection function γ_t^I and satisfies the following conditions:*

1. If $m_1, m_2 \in \text{Mod}(\psi)$, then $m_1 \leq_{\psi}^t m_2$ and $m_2 \leq_{\psi}^t m_1$.
2. If $m_1 \in \text{Mod}(\psi)$ and $m_2 \notin \text{Mod}(\psi)$, then $m_1 < m_2$.
3. If $\psi \equiv \phi$, then $\leq_{\psi}^t = \leq_{\phi}^t$.
4. If $T^t = T_2^t$, then $(T, \pi) \leq_{\psi}^t (T_2, \pi_2)$ and $(T_2, \pi_2) \leq_{\psi}^t (T, \pi)$.

Conditions 1 to 3 on the faithful assignment are the same as those of KM. Condition 4 ensures that we do not distinguish between models in the total pre-order \leq_{ψ}^t whose trees are the same up to time t . This is essentially what is represented in the revision function by bounding the all input of the revision function $*_t$ up to t . Moreover, \leq_{ψ}^t does not distinguish between models obtained by selecting two different paths from the same tree. This corresponds to the fact that we are using strong belief formulas in the revision, which do not distinguish between different paths in the same tree as well.

Theorem 2 (Representation Theorem) *An agent revision operator $*_t$ satisfies postulates (P1)-(P12) iff there exists a faithful assignment that maps each ψ to a total pre-order \leq_{ψ}^t and each I to a selection function γ_t^I such that if $(\psi, I) *_t (\varphi, i) = (\psi', I')$, then:*

1. $\text{Mod}(\psi') = \min(\text{Mod}(\varphi), \leq_{\psi}^t)$
2. $I' = \gamma_t^I(\text{Mod}(\psi'), i)$

Finally, it turns out to be straightforward to formulate the DP postulates for iterated revision in our framework for the strong beliefs and to prove their representation theorem. Due to space constraints we have omitted the results, but they can be found in a separate technical report.⁹

⁹ See <http://www.marcvanzee.nl/proofs/ecai2016.pdf>

6 Related Work

Grant *et al.* [10] develop AGM-style postulates for belief, intention, and goal revision. They provide a detailed analysis and propose different reconsideration strategies, but restrict themselves to a syntactic analysis. Much effort in combining AGM revision with action logics (e.g., the Event Calculus [18], Temporal Action Logics [14], extensions to the Fluent Calculus [26], and extensions to the Situation Calculus (see [19, Ch.2] for an overview)) concentrates on extending these action theories to incorporate *sensing* or *knowledge-producing actions*. Shapiro *et al.* [23] extend 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 [12], and is further developed by Scherl [21] and Scherl and Levesque [22] by taking into account the frame problem as well. However, none of these approaches prove representation theorems linking revision to a total pre-order on models. Baral and Zhang [2] model belief updates on the basis of semantics of modal logic S5 and show that their knowledge update operator satisfies all the KM postulates. Bonanno [3] combines temporal logic with AGM belief revision by extending 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 prove representation theorems. The concept of strong beliefs has been discussed extensively in the literature, for instance in the story of *Little Nell* [17] or a paradox found in *knowledge-based programs* [9] (see van der Hoek *et al.* [28] for a detailed discussion).

7 Conclusion

We develop a logical theory for reasoning about temporal beliefs and intentions based on Shoham's database perspective. We propose postulates for revision of strong beliefs and intentions, and prove a representation theorem relating the postulates to our formal model. Icard *et al.* prove a comparable representation theorem, but we show in this paper that it does not hold in general. In their proof, they use a canonical model construction, while we proof our representation theorem using standard techniques from belief revision. It remains an open problem whether a canonical model construction is possible when proving the representation theorem of Icard *et al.*

For future work, we aim to make the role of the planner more explicit. Currently, our agent only receives updates from the planner, but allowing the agent to do planning tasks itself would allow it to, for instance, replace intentions instead of merely discarding them. This paves the road to develop a richer notion of intentions. For this, we see the use of hierarchical planning as an interesting approach, in which intentions can be regarded as partial plans. If the databases take over part of the planning, then well-known problems such as the frame problem become more stringent: Once a fact is established (for example, as a postcondition of an intention), it persists until it explicitly contradicts postconditions established by future intentions. Existing action logics (e.g., the Event Calculus or the Fluent Calculus) and database approaches (e.g., TMMS [7]) have dealt with these problems in detail, so comparing and possibly enriching them with our formalism seems both useful and relevant future work.

Proof Sketches

Proposition 5 (Coherence Condition 2b) *If $\{(a,t), (b,t+1)\} \subseteq I$ and B is a set of strong beliefs such that (B,I) coherent, then $\{pre(b)_{t+1}, post(a)_{t+1}\}$ is consistent with B .*

Proof Sketch. Let $I = \{(a,t), (b,t+1)\}$ and let B be a set of strong beliefs whose set of models is M . We need to show that consistency of $\{pre(b)_{t+1}, post(a)_{t+1}\}$ with B follows from our coherence condition. Note that the coherence formula is $Cohere(I) = \diamond_0 pre(a,b)_t$. By the axiom (A12) (Definition 4) and the Deduction theorem we have $pre(a,b)_t \wedge do(a)_t \vdash pre(b)_{t+1}$. Using the axiom (A9): $do(a)_t \rightarrow post(a)_{t+1}$ from PAL (Def. 7 [29]) and Deduction theorem we obtain $do(a)_t \vdash post(a)_{t+1}$. Consequently,

$$pre(a,b)_t \wedge do(a)_t \vdash pre(b)_{t+1} \wedge post(a)_{t+1}.$$

Thus, in order to prove that $\{pre(b)_{t+1}, post(a)_{t+1}\}$ is consistent with B , it is sufficient to show that $pre(a,b)_t \wedge do(a)_t$ is consistent with B . If (B,I) is coherent, then there is a model $m = (T, \pi) \in M$ such that $m \models Cohere(I)$, so there is $\pi' \in T$ such that $(T, \pi') \models pre(a,b)_t$. By Definition 3.2, then there exist $\pi'' \in T$ such that $\pi' \sim_t \pi''$ and $act(\pi'') = a$. Then $(T, \pi'') \models pre(a,b)_t \wedge do(a)_t$. Since M is a set of models of strong beliefs, we obtain $(T, \pi'') \in M$, i.e., (T, π'') is also a model of B . Then B is consistent with $pre(a,b)_t \wedge do(a)_t$, by Completeness theorem. \square

Proposition 6 *If an agent $A = (B,I)$ is coherent, then $WB(B,I)$ is consistent.*

Proof Sketch. Using axioms A8, A11, and A12, one can show that $\{pre(a_0, \dots, a_m)_t\} \vdash \diamond_t (do(a_0)_t \wedge \diamond_{t+1} (do(a_1)_{t+1} \wedge \diamond_{t+2} (\dots)))$. By taking the contrapositive of A3, $pre(a_0, \dots, a_m)_t$ implies $\diamond_t \bigwedge_{k=0}^m do(a_k)_{t+k}$. Therefore, $Cohere(I)$ (Def. 10) implies

$$\diamond_0 \bigvee_{\substack{a_k \in Act: k \notin \{t_1, \dots, t_k\} \\ a_k = b_k: k \in \{t_1, \dots, t_n\}}} \diamond_{t_1} (do(a_{t_1})_{t_1} \wedge do(a_{t_1+1})_{t_1+1} \wedge \dots \wedge do(a_{t_n})_{t_n})$$

Consequently, $Cohere(I)$ implies $\diamond_0 \diamond_{t_1} \bigwedge_{k=1}^n do(b_k)_{t_k}$, and by (A3) this implies $\diamond_0 \bigwedge_{(a,t) \in I} do(a)_t$. Therefore, if (B,I) is coherent, then the set $B \cup \{\diamond_0 \bigwedge_{(a,t) \in I} do(a)_t\}$ is consistent. By the fact that B is a strong belief set, $B \cup \{\bigwedge_{(a,t) \in I} do(a)_t\}$ is consistent, i.e. $WB(B,I)$ is consistent. \square

Theorem 3 (Representation Theorem) *An agent revision operator $*_t$ satisfies postulates (P1)-(P12) iff there exists a faithful assignment that maps each ψ to a total pre-order \leq_ψ^t and each I to a selection function γ_I^t such that if $(\psi, I) *_t (\varphi, i) = (\psi', I')$, then:*

1. $Mod(\psi') = \min(Mod(\varphi), \leq_\psi^t)$
2. $I' = \gamma_I^t(Mod(\psi'), i)$

Proof Sketch. We only sketch the proof of " \Rightarrow ": Suppose that some agent revision operator $*_t$ satisfies postulates (P1)-(P12). Given models m_1 and m_2 , let $(\psi, \emptyset) *_t (form(m_1^t) \vee form(m_2^t), \varepsilon) = (\psi', \emptyset)$. We define \leq_ψ^t by $m_1 \leq_\psi^t m_2$ iff $m_1 \models \psi$ or $m_1 \models \psi'$. We also define γ_I^t by $\gamma_I^t(M^t_{SB}, i) = I'$, where $(form(M^t_{SB}), I) *_t (\top, i) = (\psi_2, I')$ (note that $\psi_2 \equiv form(M^t_{SB})$).

Let us prove condition 4 of Definition 9. For $m_1 = (T, \pi)$ and $m_2 = (T_2, \pi_2)$, let ψ' be as above. Since $\psi, \psi' \in \mathbb{B}^t$ and $T^t = T_2^t$, we have $m_1 \models \psi$ iff $m_2 \models \psi$ and $m_1 \models \psi'$ iff $m_2 \models \psi'$, so $m_1 \leq_\psi^t m_2$ and $m_2 \leq_\psi^t m_1$.

Following KM, one can show that conditions 1 to 3 from Definition 9 hold, and furthermore that $Mod(\psi') = \min(Mod(\varphi), \leq_\psi^t)$. We now prove $I' = \gamma_I^t(Mod(\psi'), i)$. By our definition of γ_I^t we have that $(\psi', I) *_t (\top, i) = (\psi_2, \gamma_I^t(Mod(\psi'), i))$ (recall that $\psi' \equiv \psi_2$). Since $(\psi, I) *_t (\varphi, i) = (\psi', I')$, by (P11) we obtain that $I' = \gamma_I^t(Mod(\psi'), i)$. Using postulate (P7)-(P10) and (P12) we can prove that γ_I^t is a selection function. \square

REFERENCES

- [1] Carlos E. Alchourron, Peter Gärdenfors, 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).
- [2] Chitta Baral and Yan Zhang, 'Knowledge updates: Semantics and complexity issues', *Artificial Intelligence*, **164**(1), 209–243, (2005).
- [3] Giacomo Bonanno, 'Axiomatic characterization of the AGM theory of belief revision in a temporal logic', *Artificial Intelligence*, **171**(2), 144–160, (2007).
- [4] Michael E. Bratman, *Intention, plans, and practical reason*, Harvard University Press, Cambridge, MA, 1987.
- [5] Philip R Cohen and Hector J Levesque, 'Intention is choice with commitment', *Artificial Intelligence*, **42**(2-3), 213–261, (1990).
- [6] Philip R. Cohen and Hector J. Levesque, 'Teamwork', *Notûs*, **25**(4), 487–512, (1991).
- [7] Thomas L Dean and Drew V McDermott, 'Temporal data base management', *Artificial Intelligence*, **32**(1), 1–55, (1987).
- [8] Hans Ditmarsch, Tiago Lima, and Emiliano Lorini, *Intention Change via Local Assignments*, 136–151, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
- [9] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi, *Reasoning about Knowledge*, MIT Press, 1995.
- [10] John Grant, Sarit Kraus, Donald Perlis, and Michael Wooldridge, 'Postulates for revising BDI structures', *Synthese*, **175**(1), 39–62, (2010).
- [11] Thomas Icard, Eric Pacuit, and Yoav Shoham, 'Joint revision of belief and intention', *Proc. of the 12th International Conference on Knowledge Representation*, 572–574, (2010).
- [12] Yi Jin and Michael Thielscher, 'Representing beliefs in the fluent calculus.', in *ECAI*, pp. 823–827. IOS Press, (2004).
- [13] Hirofumi Katsuno and Alberto O. Mendelzon, 'Propositional knowledge base revision and minimal change', *Artificial Intelligence*, **52**(3), 263–294, (dec 1991).
- [14] Jonas Kvarnström, *TALplanner and other extensions to Temporal Action Logic*, Ph.D. dissertation, Linköpings universitet, 2005.
- [15] 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*, pp. 198–211. Springer, (2009).
- [16] Emiliano Lorini and Andreas Herzig, 'A logic of intention and attempt', *Synthese*, **163**(1), 45–77, (2008).
- [17] Drew McDermott, 'A temporal logic for reasoning about processes and plans', *Cognitive science*, **6**(2), 101–155, (1982).
- [18] Erik T Mueller, *Commonsense reasoning*, Morgan Kaufmann, 2010.
- [19] Theodore Patkos, *A formal theory for reasoning about action, knowledge and time*, Ph.D. dissertation, University of Crete-Heraklion, 2010.
- [20] M. Reynolds, 'An axiomatization of full computation tree logic', *Journal of Symbolic Logic*, **66**(3), 1011–1057, (2002).
- [21] Richard B Scherl, 'Action, belief change and the frame problem: A fluent calculus approach', in *Proceedings of the Sixth workshop on Non-monotonic Reasoning, Action, and Change at IJCAI*, (2005).
- [22] Richard B Scherl and Hector J Levesque, 'Knowledge, action, and the frame problem', *Artificial Intelligence*, **144**(1), 1–39, (2003).
- [23] Steven Shapiro, Maurice Pagnucco, Yves Lesprance, and Hector J. Levesque, 'Iterated belief change in the situation calculus.', *Artificial Intelligence*, **175**(1), 165–192, (2011).
- [24] Yoav Shoham, 'Logical theories of intention and the database perspective', *Journal of Philosophical Logic*, **38**(6), 633–647, (2009).
- [25] Yoav Shoham, 'Why knowledge representation matters', *Commun. ACM*, **59**(1), 47–49, (January 2016).
- [26] Michael Thielscher, 'The concurrent, continuous fluent calculus', *Studia Logica*, **67**(3), 315–331, (2001).
- [27] Wiebe van der Hoek, Wojciech Jamroga, and Michael Wooldridge, 'Towards a theory of intention revision', *Synthese*, **155**(2), 265–290, (2007).
- [28] Wiebe Van der Hoek and Michael Wooldridge, 'Towards a logic of rational agency', *Logic Journal of IGPL*, **11**(2), 135–159, (2003).
- [29] 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*, (2015).
- [30] Marc van Zee, Dragan Doder, Mehdi Dastani, and Leendert van der Torre, 'AGM Revision of Beliefs about Action and Time', in *Proceedings of the International Joint Conference on Artificial Intelligence*, (2015).