

Steps Towards Commonsense-Driven Belief Revision in the Event Calculus

Nikoleta Tsampanaki, Giorgos Flouris, Theodore Patkos

Institute of Computer Science, FORTH, Greece

{tsaban, fgeo, patkos}@ics.forth.gr

Abstract

Recent extensions of the Event Calculus resulted in powerful formalisms, able to reason about a multitude of commonsense phenomena in causal domains, involving epistemic notions, functional fluents and probabilistic aspects, among others. Surprisingly, little attention has been paid to the problem of automatically revising (correcting) a Knowledge Base when an observation contradicts predictions regarding the world. Despite mature work on the related belief revision field, adapting such results for the case of action theories is non-trivial. This paper reports on ongoing work for addressing this problem by proposing a generic framework in the context of the Event Calculus, along with ASP encodings of the revision algorithm.

1 Introduction

Action languages are well-established logical theories for reasoning about the dynamics of changing worlds, aiming at “*formally characterizing the relationship between the knowledge, the perception and the action of autonomous agents*” [Van Harmelen *et al.*, 2007]. One of the most prominent action languages is the Event Calculus [Kowalski and Sergot, 1986; Miller and Shanahan, 2002], which incorporates certain useful features for representing causal and narrative information that differentiate it from other similar formalisms. The Event Calculus explicitly represents temporal knowledge, enabling reasoning about the effects of a narrative of events along a time line. It also relies on a non-monotonic treatment of events, in the sense that by default there are no unexpected effects or event occurrences.

Powerful extensions of the main formalism have been developed to accommodate, for instance, epistemic extensions [Miller *et al.*, 2013; Ma *et al.*, 2013; Patkos and Plexousakis, 2009], probabilistic uncertainty [Skarlatidis *et al.*, 2015; D’Asaro *et al.*, 2017] or knowledge derivations with non-binary-valued fluents [Miller *et al.*, 2013]. Moreover, progress in generalizing the stable model semantics used in Answer Set Programming (ASP) has opened the way for the reformulation of Event Calculus axiomatizations into logic programs that can be executed with efficient ASP solvers [Ferraris *et al.*, 2011]. This allowed for exploiting

state-of-the-art tools that outperform previous SAT- or logic programming-based solvers in almost all classes of problems related to practical applications [Lee and Palla, 2012].

However, to the best of our knowledge, little work has been done on supporting *belief change* in the Event Calculus, in cases when the new information contradicts the already inferred knowledge. Specifically, the existing non-epistemic extensions accommodate *belief update*, which concerns beliefs that change as the result of the realization that the world has changed through some action. The epistemic extensions further focus on modeling the notions of knowledge, thus supporting *belief expansion*, where newly acquired information can enrich the belief state of agents about aspects that were previously considered unknown. Yet, the ability to accommodate, through proper *revisions*, sensed information that contradicts existing beliefs is not supported. This problem is more general than belief expansion or even than diagnosis, as it not only intends to identify the reasons that explain the contradictions, but also to suggest proper modifications of the belief state of the agent under certain, potentially domain-dependent, criteria [Alchourron *et al.*, 1985].

In this paper, we present steps towards a formal method for accommodating belief revision on top of Event Calculus axiomatizations. We consider both the epistemic and non-epistemic case, relying on the possible-worlds representation to give formal semantics to an agent’s belief state. We formalize notions of commonsense revisions that take into consideration different knowledge states, such as factual (or observed), inferred and unknown beliefs. Finally, we present a methodology and an ASP encoding that can implement the formalism. The current framework is based on certain simplifying assumptions, such as deterministic domains and lack of state constraints (state axioms), which limit its broadness. Yet, this work can form the substrate for further extensions concerning a richer set of commonsense features, such as default beliefs, non-determinism, introspective belief changes, non-binary aspects etc, along with formal results showing that it is generic enough to be applied to different Event Calculus dialects.

Example: Consider the classical Yale Shooting scenario, where a loaded gun is fired against a living, walking turkey. An observer may believe that, after the shot, the turkey is dead. If future observations contradict her beliefs, e.g., by noticing that the turkey is still walking, the observer will need

to assess different potential revisions of her belief state: can it be that she was so mistaken and the shooter did not fire the gun in the first place? Or is it just that the initial, default belief about the gun being loaded was not accurate? Moreover, how would the revisions be affected if the initial state of the gun is unknown?

Although simplistic, this setting of the Yale Shooting scenario can be generalized to account for different levels of commonsense inferences, some of which may be domain-independent, e.g., revising aspects that were initially unknown rather than aspects that have already been observed, while others may be domain-dependent, e.g., considering certain observations as being less reliable than others. For such types of domains, we develop in the sequel a formal methodology for revising the belief state of an agent, taking into consideration commonsense and epistemic notions.

It should be noted that, even though the belief change literature has been used as a source of inspiration for addressing the problem, the related technical results cannot be directly applied, as they are based on assumptions that are not relevant for our setting (e.g., monotonicity of the underlying representation language). Thus, our approach leverages only the related methodologies, establishing connections among our ideas and the corresponding ideas from belief change.

The rest of the paper is structured as follows: in Section 2 we remind the basics about the Event Calculus that will be needed in the next sections. Sections 3 and 4 give the theoretical underpinnings of our methodology for the non-epistemic and the epistemic case respectively, describing the problem (and the proposed solution) in formal terms. In Section 5 we describe the implementation of the methodology in ASP, while in Section 6 we discuss related work. The paper concludes in Section 7 with remarks about our next steps.

2 Preliminaries

Our account of change and causality is based on the discrete time Event Calculus axiomatized in [Mueller, 2015], while the modeling of possible worlds for representing epistemic notions is inspired by the epistemic extension of the Functional Event Calculus (\mathcal{EFEC}) [Miller *et al.*, 2013]. In particular, we consider a sort \mathcal{E} for events (variables e, e', e_1, \dots), a sort \mathcal{F} for fluents (f, f', f_1, \dots) and a sort \mathcal{T} for timepoints (t, t', t_1, \dots), which is restricted to the integers.¹ The key predicates are $HoldsAt() \subseteq \mathcal{F} \times \mathcal{T}$ denoting the truth value of a fluent at a particular timepoint, $Happens() \subseteq \mathcal{E} \times \mathcal{T}$ capturing the occurrence of events, $Initiates() \subseteq \mathcal{E} \times \mathcal{F} \times \mathcal{T}$ and $Terminates() \subseteq \mathcal{E} \times \mathcal{F} \times \mathcal{T}$, denoting that an event e causes fluent f to become true or false respectively in the next timepoint. The notions of cause, effect and inertia are captured in the \mathcal{DEC} domain independent set of axioms (see [Mueller, 2015]). As we restrict our considerations to deterministic domains for the time being, we do not axiomatize fluents that are released from the law inertia.

In order to support epistemic reasoning, we introduce two new sorts, in the style of \mathcal{EFEC} : a sort \mathcal{W} for representing

¹In the sequel, variables start with a small letter, and constants with a capital letter. Wherever not explicitly stated, variables are assumed to be universally quantified.

possible worlds (variables w, w', w_1, \dots) and a sort \mathcal{I} for instants (variables i, i', i_1, \dots). The idea is to represent time as a system of parallel lines, where each world is understood as an identifier for a possible time line. Finally, we assume that the constant W_a of sort \mathcal{W} signifies the actual world.

3 Revisions in the Non-Epistemic Case

3.1 Revision Setting and Principles

The representation of a dynamic domain requires the coupling of domain-independent and domain-dependent axioms with our knowledge about the world and the related narrative. Overall, we define a Knowledge Base as follows:

Definition 1 A Knowledge Base (KB) capturing a dynamic domain is defined as $\Phi = \mathcal{DEC} \wedge \Sigma \wedge \Gamma^0 \wedge \Delta \wedge \Omega$ where

- \mathcal{DEC} is the conjunction of the Discrete Event Calculus domain-independent axioms,
- Σ is the conjunction of the domain-dependent axioms,
- Γ^0 is the initial knowledge, i.e., a conjunction of ground $(\neg)HoldsAt(F_i, 0)$ axioms at timepoint 0,
- Δ is the narrative of actions, i.e., a conjunction of ground $Happens(E_i, T_j)$ axioms,
- Ω is a conjunction of unique name axioms.

Domain axioms in Δ can be partially defined and then minimized to address the Frame Problem and related issues. Γ^0 axioms cannot be partially defined, as we assume complete world knowledge initially (this assumption will be lifted in Section 4). We denote by $\Phi \models \phi$ the fact that Φ implies ϕ .

We assume that from time to time we observe some part of the world, i.e., we obtain the truth value of certain fluents. Our current assumption is that observations can only contain a conjunction of $(\neg)HoldsAt()$ statements. We denote by Γ^T an observation obtained at timepoint T .

Now let's turn our attention to the problem of revising a KB Φ with an observation Γ^T . We follow the *Principle of Consistency Maintenance* [Dalal, 1988], which requires that the result of revising Φ with Γ^T should be consistent. In addition, we make the standard assumption that is expressed by the *Principle of Primacy of New Information* [Dalal, 1988] (and formalized by the *postulate of success* in the AGM postulates [Alchourron *et al.*, 1985]), which states that the new observation should always be entailed after the revision.

The special characteristics of the Event Calculus force us to introduce two new principles. The first is the *Principle of Persistence of Background Knowledge*, which states that the revision process will only affect the initial knowledge (Γ^0) and/or the narrative (Δ). Thus, the domain-independent (\mathcal{DEC}) and domain-dependent (Σ) axioms, as well as the unique name axioms (Ω), should not be affected. This avoids issues associated to the problem of learning the domain from observations, which is not in the scope of this paper.

The second new requirement is the *Principle of Disallowing Proactive Change*, which, informally, states that we cannot use an observation referring to time T in order to add events in the narrative beyond that timepoint. Essentially, this

limits the direct effects of an observation (and the corresponding revision) to past timepoints, even though such effects may also have indirect ramifications related to the truth value of fluents in the future.

Finally, we adopt the *Principle of Minimal Change* [Katsuno and Mendelzon, 1991] (also known as the *Principle of Persistence of Prior Knowledge* [Dalal, 1988]), which states that the new KB should be as “close” as possible to the original KB; in other words, from all the possible change results (revision candidates) that satisfy the other principles, we should choose the one that retains “the most information”.

3.2 The Revision Operator

Following the above principles, we can formally define the set of *revision candidates* as follows:

Definition 2 Given a KB $\Phi = \mathcal{DEC} \wedge \Sigma \wedge \Gamma^0 \wedge \Delta \wedge \Omega$ and an observation Γ^T , a KB Φ' is a revision candidate of Φ with Γ^T iff:

- Φ' is of the form $\Phi' = \mathcal{DEC} \wedge \Sigma \wedge \Gamma^{T0} \wedge \Delta' \wedge \Omega$ (*Principle of Persistence of Background Knowledge*).
- No formula in $(\Delta \setminus \Delta') \cup (\Delta' \setminus \Delta)$ refers to timepoints $t > T$ (*Principle of Disallowing Proactive Change*).
- Φ' is a consistent KB (*Principle of Consistency*).
- $\Phi' \models \Gamma^T$ (*Principle of Primacy of New Information*).

The set of all revision candidates of Φ with Γ^T will be denoted by $RC(\Phi, \Gamma^T)$.

Note that Definition 2 imposes that the part $\mathcal{DEC} \wedge \Sigma \wedge \Omega$ of all revision candidates is identical to the corresponding part of the original KB (following the Principle of Persistence of Background Knowledge), and also formalizes all other principles (except from the Principle of Minimal Change). The latter is not considered because $RC(\Phi, \Gamma^T)$ is meant to represent all the *conceivably possible* revision results, not the optimal ones. The notion of minimal change is often subjective, context- and/or domain-dependent, so we chose to include it as a separately configurable component of our framework.

To formalize the Principle of Minimal Change, we will use the standard approach of introducing a *preference relation* \prec_{Φ}^T . The idea is that if $\Phi_1 \prec_{\Phi}^T \Phi_2$, Φ_1 is strictly more preferred than Φ_2 for the result of the revision of Φ with an observation at timepoint T . Note that this is different from the relations among interpretations [Katsuno and Mendelzon, 1991] and formulas [Gardenfors and Makinson, 1988] that have been used elsewhere for the same purpose. Establishing the connection among our preference relation and these works is part of our future work. For now, it suffices to assume that \prec_{Φ}^T is wellfounded (so that we can always find a minimal element in a non-empty set). Further properties (e.g., totality, transitivity) may improve algorithmic efficiency in identifying the optimal solution, but this is irrelevant for now.

We are now ready to define the revision operator. Intuitively, the idea is that we select those elements of $RC(\Phi, \Gamma^T)$ that are minimal with respect to \prec_{Φ}^T . In case multiple minimal elements exist, their disjunction is taken. It is also interesting to note that, in the special case when $RC(\Phi, \Gamma^T) = \emptyset$, we do not revise the KB; this can happen, e.g., when the

observation itself is inconsistent or when there is no way to satisfy the observation without changing background knowledge, such as the domain axioms. Formally:

Definition 3 The revision operator \bullet is a binary operator, defined as follows:

- $\Phi \bullet \Gamma^T = \Phi$ if $RC(\Phi, \Gamma^T) = \emptyset$.
- $\Phi \bullet \Gamma^T = \bigvee \{ \Phi' \mid \Phi' \in RC(\Phi, \Gamma^T) \text{ and there is no } \Phi'' \in RC(\Phi, \Gamma^T) \text{ such that } \Phi'' \prec_{\Phi}^T \Phi' \}$ otherwise.

3.3 Defining a Specific Preference Relation

To define the preference relation more precisely, we will leverage a *cost-based model* which assesses minimality based on the amount of information lost or modified from the original KB in order to accommodate the observation. In particular, considering two KBs Φ, Φ' the *cost* to move from Φ to Φ' will be defined on the basis of the formulas that can be inferred by one of these KBs but not the other. To formalise this, we first define the following sets:

Modified Knowledge $MK^T(\Phi, \Phi') = \{ HoldsAt(F, T') \mid T' \leq T \text{ and either } \Phi \models HoldsAt(F, T') \text{ and } \Phi' \models \neg HoldsAt(F, T'), \text{ or } \Phi \models \neg HoldsAt(F, T') \text{ and } \Phi' \models HoldsAt(F, T') \}$. This represents all *HoldsAt()* statements whose truth value was changed during the transition from Φ to Φ' , up to T .

New Events $NE^T(\Phi, \Phi') = \{ Happens(E, T') \mid T' \leq T, \Phi \models \neg Happens(E, T') \text{ and } \Phi' \models Happens(E, T') \}$. This represents all events that we had to add in the narrative of Φ to accommodate the observation, up to T .

Lost Events $LE^T(\Phi, \Phi') = \{ Happens(E, T') \mid T' \leq T, \Phi \models Happens(E, T') \text{ and } \Phi' \models \neg Happens(E, T') \}$. This represents all events that we had to retract from the narrative of Φ to accommodate the observation, up to T .

Note that the above definitions do not consider the consequences of changes for future timepoints, i.e., beyond a certain timepoint T . This will be used to ignore any future repercussions of our changes, considering only the changes up to the timepoint of the observation.

The *cost* between two KBs Φ, Φ' up to the timepoint T is defined as:

$$cost^T(\Phi, \Phi') = w_{MK} \cdot |MK^T(\Phi, \Phi')| + w_{NE} \cdot |NE^T(\Phi, \Phi')| + w_{LE} \cdot |LE^T(\Phi, \Phi')|,$$

where w_{MK}, w_{NE}, w_{LE} are the corresponding weights associated to each change (a parameter of our model).

Now, the \prec_{Φ}^T relation can be easily defined as follows: $\Phi_1 \prec_{\Phi}^T \Phi_2$ iff $cost^T(\Phi, \Phi_1) < cost^T(\Phi, \Phi_2)$. It is trivial to show that this relation is well-founded, as required by the definition.

3.4 Discussion on the preference function

Note that the above definition of the cost function corresponds to the non-epistemic version that our current implementation supports. Yet, we also consider various alternative options:

- Partitioning fluents and/or events into different “importance categories”, each with its own weight, w_1, \dots, w_n . This way, e.g., two different “new events” would

cost differently, depending on the weight of the involved event. For this case an aggregation function would be required, to aggregate w_i with the weight (w_{MK}, w_{NE}, w_{LE}) of the corresponding category of the formula whose weight is considered. This could help accommodate the case where default fluents can more easily be changed than non-default ones, or for cases where specific fluents cost more (in terms of epistemic effort or practical consequences) to change.

- Another future consideration would consider a degradation of the cost over time. This can support the intuition that it should be more expensive to change knowledge about past timepoints than knowledge about more recent timepoints. Again, this could be supported with an appropriate aggregation function, combining the weight coming from the type of change (w_{MK}, w_{NE}, w_{LE}) , the weight coming from the fluent or event (w_i) and the weight coming from the timepoint where the corresponding change occurred.
- Further, it is clear that the theory requires a qualitative, relative ordering of the different revision candidates. Even though this can be reduced to the comparison of the result of numerical functions, like the above, it is still possible to use a purely qualitative comparison method, or even hybrid methods combining qualitative and quantitative components. For this early version of the work, we chose the more simple quantitative approach, which is also more amenable to the implementation method chosen (see Section 5).

The above ideas can be neatly formalized by considering a weight function associating a “weight” to each possible formula ($HoldsAt()$, $Happens()$) and computing the cost as the total “weight” of the formulas implied by Φ but not Φ' (or vice-versa).

Despite its early stage of development, the proposed \prec_{Φ}^T relation has several intuitively desirable formal properties. First, we show that “fewer” changes (with respect to the standard set-theoretic subset relation) are better than “more”:

Proposition 1 *Consider three KBs Φ, Φ_1, Φ_2 . Set $C^T(\Phi_i) = \{\phi \mid \Phi \models \phi, \Phi_i \not\models \phi \text{ and } \phi \text{ refers to a timepoint } t' \leq T\}$, for $i = 1, 2$. If $C^T(\Phi_1) \subset C^T(\Phi_2)$, then $\Phi_1 \prec_{\Phi}^T \Phi_2$.*

As a corollary, we get that not changing a KB is always cheaper than changing it, and this will happen whenever the observation does not contradict our expectations:

Proposition 2 $\Phi \prec_{\Phi}^T \Phi'$ for all Φ, Φ', T .

Proposition 3 If $\Phi \in RC(\Phi, \Gamma^T)$, then $\Phi \bullet \Gamma^T = \Phi$.

Proposition 4 If $\Phi \models \Gamma^T$, then $\Phi \bullet \Gamma^T = \Phi$.

Example (cont.) Returning to the Yale shooting example described before, the observer’s KB Φ_{Yale} can be described by the following axiomatization, stating that the gun is loaded at timepoint 0 and fired at timepoint 1.

$$Initiates(Load, Loaded, t) \quad (3.1)$$

$$HoldsAt(Loaded, t) \rightarrow Terminates(Shoot, Loaded, t) \quad (3.2)$$

$$HoldsAt(Loaded, t) \rightarrow Terminates(Shoot, Alive, t)$$

$$HoldsAt(Alive, 0) \quad (3.3)$$

$$HoldsAt(Loaded, 0) \quad (3.4)$$

$$Happens(Shoot, 1) \quad (3.5)$$

That is, $\Sigma = (3.1) \wedge (3.2) \wedge (3.3)$, $\Gamma^0 = (3.4) \wedge (3.5)$ and $\Delta = (3.6)$, whereas the remaining components of Φ_{Yale} follow from Definition 1. It can be shown that $\Phi_{Yale} \models \neg HoldsAt(Alive, 2)$.

Assume now that the observer receives information that contradicts her current inferences, e.g., $\Gamma^3 = HoldsAt(Alive, 3)$ (note that $\Phi_{Yale} \not\models HoldsAt(Alive, 3)$). A possible reaction to this observation would be that the observer was mistaken and the shooter did not fire the gun. That is, $\Delta' = \emptyset$ and $\Gamma'^0 = \Gamma^0$. So, a revision candidate of Φ_{Yale} , following Definition 2, would be $\Phi'_{Yale} = \mathcal{DEC} \wedge \Sigma \wedge \Gamma'^0 \wedge \Delta' \wedge \Omega$.

Another possible revision would be that the observer was mistaken and the gun was not loaded in the first place.

That is, $\Delta'' = \Delta$ and $\Gamma''^0 = (3.4) \wedge \neg HoldsAt(Loaded, 0)$. The revision candidate of Φ_{Yale} is now $\Phi''_{Yale} = \mathcal{DEC} \wedge \Sigma \wedge \Gamma''^0 \wedge \Delta'' \wedge \Omega$.

Consequently, $\Phi'_{Yale}, \Phi''_{Yale} \in RC(\Phi_{Yale}, \Gamma^3)$. Note that many other KBs are included in $RC(\Phi_{Yale}, \Gamma^3)$, but all of them would introduce more changes (with regards to the subset relation) than these two (see also Proposition 1), so they are not considered due to lack of space.

To find the \prec_{Φ}^3 -minimal element(s) of $RC(\Phi_{Yale}, \Gamma^3)$, we only need to compare Φ' with Φ'' . The weights associated with the cost function were set to $w_{MK} = 1, w_{NE} = 2, w_{LE} = 2$, so the corresponding costs are: $cost^3(\Phi_{Yale}, \Phi'_{Yale}) = 6$, $cost^3(\Phi_{Yale}, \Phi''_{Yale}) = 4$, so $\Phi''_{Yale} \prec_{\Phi}^3 \Phi'_{Yale}$ and Φ''_{Yale} is the revision result.

4 Revisions in the Epistemic Case

The high demands that are imposed on autonomous systems in real domains have led to variations of Event Calculus theories that can support reasoning with partial world knowledge. Such epistemic extensions can accommodate both known and unknown fluents, using a special type of “sense” actions to acquire new knowledge, which by definition only affect the belief state of the agent, causing no effect to the state of the domain. In this section, we discuss how revision of beliefs can be achieved when the sensed information contradicts existing knowledge. We rely on the approach introduced in the recent \mathcal{EFEC} dialect that implemented an adaptation of the possible worlds model to give formal semantics to belief predicates.

In \mathcal{EFEC} , the function $\langle \rangle: \mathcal{W} \times \mathcal{I} \rightarrow \mathcal{T}$ is introduced to map world/instant pairs to timepoints. Timepoint $\langle W, I \rangle$ represents instant I in possible world W , where:

$$\forall t \exists w, i. t = \langle w, i \rangle \quad (\text{DOX1})$$

The time lines believed to be accessible at any given moment are captured by the relation $K \subseteq \mathcal{W} \times \mathcal{W}$, which represents the accessibility relation between possible worlds, as in modal logics. As ordinary, we formally define belief of some fluent f at some timepoint as the fact that this fluent has the same truth value in all worlds that are accessible from the actual world:

$$Bel(f, \langle W_a, i \rangle) \equiv \quad (\text{DOX2})$$

$$\begin{aligned} \forall wK(w, W_a) &\rightarrow HoldsAt(f, \langle w, i \rangle) \\ BelNot(f, \langle W_a, i \rangle) &\equiv \end{aligned} \quad (DOX3)$$

$$\begin{aligned} \forall wK(w, W_a) &\rightarrow \neg HoldsAt(f, \langle w, i \rangle) \\ BelWh(f, \langle W_a, i \rangle) &\equiv \\ Bel(f, \langle W_a, i \rangle) \vee BelNot(f, \langle W_a, i \rangle) &\end{aligned} \quad (DOX4)$$

In contrast to \mathcal{EFEC} though, we do not consider K to be an equivalence relation.² Instead, in order to model belief rather than knowledge, we only assume that K is serial, which is equivalent to stating that the agent cannot believe contradictions (also known as the Consistency Axiom):

$$\forall i. Bel(f, \langle W_a, i \rangle) \rightarrow \neg BelNot(f, \langle W_a, i \rangle) \quad (DOX5)$$

$$\forall i. BelNot(f, \langle W_a, i \rangle) \rightarrow \neg Bel(f, \langle W_a, i \rangle) \quad (DOX6)$$

Notice that from the above axioms we do not assume that the actual world is accessible too (K is not reflexive). As a result, erroneous beliefs can still be inferred, requiring a revision mechanism whenever observations (that reflect W_a) do not comply with the agent's beliefs.

Finally, we need to define a domain-independent axiom to ensure the existence of possible worlds in the initial state.

$$\forall f. \neg BelWh(f, \langle W_a, 0 \rangle) \rightarrow \quad (DOX7)$$

$$\exists w_1, w_2. K(w_1, W_a) \wedge K(w_2, W_a) \wedge$$

$$HoldsAt(f, \langle w_1, 0 \rangle) \wedge \neg HoldsAt(f, \langle w_2, 0 \rangle)$$

Notice that according to our assumption of never losing knowledge (there is no non-determinism), it is sufficient to preserve the number of possible worlds generated at the initial timepoint while reasoning, since there is no way of generating more worlds. We do not need to eliminate worlds either, in order to allow for reasoning about the past.

Based on the aforementioned formalization of belief, we extend the definition of a KB to accommodate lack of knowledge at the initial and at future timepoints:

Definition 4 An epistemic KB is defined as $e\Phi = \mathcal{DEC} \wedge \mathcal{DOX} \wedge \Sigma \wedge e\Gamma^0 \wedge \Delta \wedge \Omega$, where

- \mathcal{DEC} is the conjunction of the Discrete Event Calculus domain-independent axioms,
- \mathcal{DOX} is the conjunction of the epistemic axioms to support beliefs,
- Σ is the conjunction of the domain-dependent axioms,
- $e\Gamma^0$ is the initial beliefs, i.e., a conjunction of ground $Bel(F_i, \langle W_a, 0 \rangle)$, $BelNot(F_j, \langle W_a, 0 \rangle)$ axioms at timepoint 0,
- Δ is the narrative of actions, i.e., a conjunction of ground $Happens(E_i, \langle w, I_j \rangle)$ axioms,
- Ω is a conjunction of unique name axioms.

The definition for $e\Phi$ is more general than the one given for Φ . Specifically, if we assume complete world knowledge at timepoint 0, a single possible world is generated, making $e\Phi$ equivalent to Φ .

As before, Δ axioms can be partially defined and then minimized to address the Frame Problem and related issues. $e\Gamma^0$

²Note also that the current setting only permits belief statements that refer to the state of world fluent at the same time point. In other words, we do not represent the beliefs at some timepoint about the state of fluents at another timepoint. Such statements are supported in \mathcal{EFEC} and will be considered for future extensions of our axiomatization.

axioms can be partially defined as well; fluents that are unknown at time instant 0 generate the set of possible worlds according to axiom (DOX7).

The case of revision in the epistemic Event Calculus is virtually identical to the case of the non-epistemic one. Thus, whatever we discussed in Section 3 can be applied here as well. The main difference in the epistemic case is that we can now provide a more fine-grained preference relation, taking special provisions for the case where a fluent whose value was originally unknown, became known (true or false). To do so, an approach similar to the one used in Subsection 3.3 could be used, namely, identifying the sets of formulas of the form $Bel(\dots)$, $BelNot(\dots)$, $BelWh(\dots)$ that are implied/not implied by $e\Phi$, $e\Phi'$ respectively.

In particular, considering two epistemic KBs $e\Phi$, $e\Phi'$ the *cost* to move from $e\Phi$ to $e\Phi'$ will be defined on the basis of the formulas that can be inferred by one of these KBs but not the other. To formalise this, we first define the following sets:

Modified Knowledge $MK^T(e\Phi, e\Phi') = \{Bel(F, \langle W_a, T' \rangle) \mid T' \leq T \text{ and either } e\Phi \models Bel(F, \langle W_a, T' \rangle) \text{ and } e\Phi' \models BelNot(F, \langle W_a, T' \rangle), \text{ or } e\Phi \models BelNot(F, \langle W_a, T' \rangle) \text{ and } e\Phi' \models Bel(F, \langle W_a, T' \rangle)\}$. This represents all $Bel()$ or $BelNot()$ statements whose truth value was changed during the transition from $e\Phi$ to $e\Phi'$, up to T .

New Knowledge $NK^T(e\Phi, e\Phi') = \{Bel(F, \langle W_a, T' \rangle) \mid T' \leq T \text{ and either } e\Phi \models \neg BelWh(F, \langle W_a, T' \rangle) \text{ and } e\Phi' \models Bel(F, \langle W_a, T' \rangle), \text{ or } e\Phi \models \neg BelWh(F, \langle W_a, T' \rangle) \text{ and } e\Phi' \models BelNot(F, \langle W_a, T' \rangle)\}$. This represents all $\neg BelWh()$ statements whose unknown value was changed to known during the transition from $e\Phi$ to $e\Phi'$, up to T .

Lost Knowledge $LK^T(e\Phi, e\Phi') = \{\neg BelWh(F, \langle W_a, T' \rangle) \mid T' \leq T \text{ and either } e\Phi \models Bel(F, \langle W_a, T' \rangle) \text{ and } e\Phi' \models \neg BelWh(F, \langle W_a, T' \rangle), \text{ or } e\Phi \models BelNot(F, \langle W_a, T' \rangle) \text{ and } e\Phi' \models \neg BelWh(F, \langle W_a, T' \rangle)\}$. This represents all $Bel()$ or $BelNot()$ statements whose truth value was changed to unknown during the transition from $e\Phi$ to $e\Phi'$, up to T .

New Events $NE^T(\Phi, \Phi') = \{Happens(E, \langle W_a, T' \rangle) \mid T' \leq T, \Phi \models \neg Happens(E, \langle W_a, T' \rangle) \text{ and } \Phi' \models Happens(E, \langle W_a, T' \rangle)\}$. This represents all events that we had to add in the narrative of Φ to accommodate the observation, up to T .

Lost Events $LE^T(\Phi, \Phi') = \{Happens(E, \langle W_a, T' \rangle) \mid T' \leq T, \Phi \models Happens(E, \langle W_a, T' \rangle) \text{ and } \Phi' \models \neg Happens(E, \langle W_a, T' \rangle)\}$. This represents all events that we had to retract from the narrative of Φ to accommodate the observation, up to T .

The *cost* between two KBs $e\Phi$, $e\Phi'$ up to the timepoint T is defined as:

$$\begin{aligned} cost^T(e\Phi, e\Phi') &= w_{MK} \cdot |MK^T(e\Phi, e\Phi')| + w_{NK} \cdot \\ &|NK^T(e\Phi, e\Phi')| + w_{LK} \cdot |LK^T(e\Phi, e\Phi')| + w_{NE} \cdot \\ &|NE^T(e\Phi, e\Phi')| + w_{LE} \cdot |LE^T(e\Phi, e\Phi')|, \end{aligned}$$

where $w_{MK}, w_{NK}, w_{LK}, w_{NE}, w_{LE}$ are the corresponding weights associated to each change (a parameter of our model).

Now, the \prec_{Φ}^T relation can be easily defined as follows:
 $e\Phi_1 \prec_{\Phi}^T e\Phi_2$ iff $cost^T(e\Phi, e\Phi_1) < cost^T(e\Phi, e\Phi_2)$. It is trivial to show that this relation is well-founded, as required by the definition.

5 Implementation

The proposed framework was implemented for both the non-epistemic and the epistemic case, using the architecture shown in Figure 1.³ The figure shows the loop of steps performed whenever a new observation arrives, along with the corresponding input/output modules (rulesets). There are two main reasoning steps, interconnected via two Java programs: the first reasoning step generates one answer set in the non-epistemic case and multiple answer sets in the epistemic case, each denoting a possible world (see Section 4), and the second step generates cost-optimal revisions. Reasoning is performed by implementing the Event Calculus axiomatizations as ASP rules and by utilizing the Clingo reasoner [Gebser *et al.*, 2011]. Similarly, the domain-dependent axioms, the epistemic axioms, the main program, the meta-program, the new observations and the current beliefs are all implemented in ASP. The current beliefs module contains the running information about the world state and the narrative of actions. A Java parser intervenes between the two reasoning steps to transform the information contained in the generated answer sets, i.e., the possible worlds, into the agent’s belief predicates. These are introduced in ASP form in the revision meta-program, which provides the revision results, based on the DOX axioms described in Section 4. Finally, a Java parser parses all other modules needed for automating the process and for connecting the reasoning to the outside world.

The most important module is the revision meta-program, which implements the revision algorithm. In case of inconsistency, it takes as input the new observation and the result from the Java parser, which expresses the running belief state. The meta-program computes the revision using the cost-optimal

³<http://www.ics.forth.gr/isl/CS17BelRevPaper/>

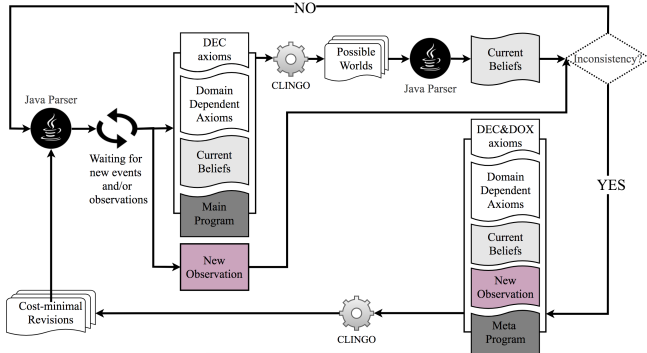


Figure 1: The reasoning loop for revising the KB of an agent.

revisions generator, via a logic program. Roughly, this program generates combinations of fluents in the initial state, as well as combinations of event occurrences at each timepoint, for every possible world, aiming at keeping only the combinations that lead to a consistent KB and are consistent with the new observation (revision candidates). The cost associated with each revision candidate is calculated, based on the cost function described in Section 4; this is implemented with ASP rules that penalize each truth value or event that is different from the output of the first reasoning step. Finally, an optimization statement filters out all non-optimal revision candidates (answer sets). At the end, the program returns the disjunction of the optimal candidates.

Returning to the Yale shooting example described before, an output of the program in the epistemic case is discussed next. Namely, the initial beliefs are that the turkey is alive at timepoint 0 and a shot happens at timepoint 1, but we do not know whether the gun is loaded or not at timepoint 0. Thus, we do not know whether the turkey is alive or not at timepoint 3. Assume now that we receive a new contradicting information that the turkey is alive at timepoint 3. We present the optimal revision that accommodates this observation into our knowledge base in Figure 2, based on the following weights: $w_{MK} = 1, w_{NK} = 1, w_{LK} = 2, w_{NE} = 2, w_{LE} = 2$. More specifically, the optimal revision is to assume that we were mistaken and the gun was not loaded in the first place. Thus, we have the least possible cumulative cost, as we gain knowledge on the state of the turkey at timepoints 2 and 3, as well as, on the state of the gun at timepoints 0 and 1. Had we accepted the revision that the shooter did not fire the gun, we would lose knowledge and event happenings at various timepoints, and as a result, the cost would be greater in total.

-----Current Beliefs-----				
-----FLUENTS-----				
Fluent\Timepoint	0	1	2	3
loaded	~BelWh	~BelWh	BelNot	BelNot
alive	Bel	Bel	~BelWh	~BelWh
-----EVENTS-----				
Event\Timepoint	0	1	2	3
load	-	-	-	-
shoot	-	Happ	-	-
-----Revised Beliefs-----				
-----Optimal Revision No1-----				
-----FLUENTS-----				
Fluent\Timepoint	0	1	2	3
loaded	BelNot	BelNot	BelNot	BelNot
alive	Bel	Bel	Bel	Bel
-----EVENTS-----				
Event\Timepoint	0	1	2	3
load	-	-	-	-
shoot	-	Happ	-	-

Figure 2: The output of the Java agent, divided in two sections. The upper section in blue represents the original belief state, and the lower section in green represents the revised belief state of the agent.

6 Related Work

Belief change (also known as *belief revision*) is a mature field of study dealing with the adaptation of a KB in the face of new information [Alchourron *et al.*, 1985]. Traditionally, two types of change have been considered: *revision* and *update* [Katsuno and Mendelzon, 1991]. Revision deals with cases where the new information is some observation or refinement of our knowledge about the world, whereas update deals with cases where the adaptation is dictated by some action or event that changed the world itself. The case of belief update is inherently captured by the semantics and reasoning of Event Calculus, where one can explicitly declare events, as well as the effects and preconditions of such events. However, studies regarding the revision of action theories when our observations of the actual world are inconsistent with the theory’s predictions regarding the world’s state are limited.

Most works in the classical belief change literature are dealing with the so-called *classical logics* [Flouris *et al.*, 2006], which have certain nice properties, both in terms of semantics (monotonicity, compactness, inclusion of the classical tautological implication, etc) and in terms of syntax (closed with respect to the usual operators \wedge, \vee, \neg , etc). Extensions of these theories to apply for ontological languages [Flouris *et al.*, 2006; Qi and Du, 2009], or compact and monotonic logics in general [Ribeiro *et al.*, 2013] have been considered as well. However, to the best of our knowledge, no such study exists for non-monotonic formalisms, partly because many non-monotonic formalisms (most notably, defeasible logic, default logic and paraconsistent logics) have inherent ways to reason under inconsistency without trivializing inference. Thus, technical results from the related literature are not directly applicable in our setting.

Studies that account for epistemic considerations of the Event Calculus are more closely related to ours. More specifically, the \mathcal{EFEC} variant introduced in [Miller *et al.*, 2013; Ma *et al.*, 2013] is the first to rely on the possible worlds semantics to reason about knowledge. \mathcal{EFEC} supports a multitude of features, such as reasoning about the future and past, or dealing with non-determinism and concurrency. Our work utilizes the same underlying structures to formalize the treatment of epistemic notions, extending them with the ability to revise contradicting knowledge, although it is currently significantly limited in the set of supporting features. In [Patkos and Plexousakis, 2009], a different epistemic extension of discrete-time Event Calculus theories is presented, using a deduction-oriented model of knowledge.

Beyond the Event Calculus, possible-worlds based epistemic extensions for reasoning about actions and knowledge have been developed in the context of other calculi. The first approach that inspired this direction of research is owed to [Moore, 1985], who presented a Kripke-like formulation of epistemic notions of modal logic in action languages by reifying possible worlds as situations. [Scherl and Levesque, 2003] adapted this framework in the Situation Calculus, using possible situations to specify how the mental state of an agent should change with ordinary and sense actions, providing also a solution to the frame problem for knowledge. Other studies introduced further fea-

tures: [Thielscher, 2000] adapted the model in the context of the Fluent Calculus, [Scherl, 2003] covered concurrent actions, while [Kelly and Pearce, 2008] introduced epistemic modalities for groups of agents. Non-possible-worlds based epistemic action frameworks include [Morgenstern, 1987; Demolombe and Pozos-Parra, 2000; Son and Baral, 2001; Petrick and Levesque, 2002; Vassos and Levesque, 2007; Liu and Lakemeyer, 2009]. In all these frameworks, knowledge is assumed to be always correct and observations that contradict inferred knowledge will lead to inconsistency.

The ability to deal with belief changes has lately started to gain interest within other action languages, as in [Shapiro *et al.*, 2011] and [Schwering *et al.*, 2015] in the Situation Calculus, but without taking time into account. [Van Zee *et al.*, 2015] developed a new action formalism for revision of temporal belief bases; even though related to our work, [Van Zee *et al.*, 2015] do not directly address the problem of revising theories in Event Calculus, but instead define a new logic of action that is closer to propositional logic, thereby allowing technical results from the belief change literature to be directly applicable in their framework.

7 Conclusion

This paper reported on early work towards a formal framework for changing Event Calculus theories in the face of new (and potentially unexpected) observations. Our framework is necessary in the cases where an intelligent agent observes, or otherwise becomes aware of, information that contradicts what was expected by the underlying theory. Even though the rich technical results from the belief change literature are not generally applicable to our setting, we leveraged on some key ideas and adapted them for our purposes. Our approach was based on a set of principles and a preference relation that models the well-known Principle of Minimal Change.

We are currently working on extending the framework with more features and providing a more efficient implementation, along with a more generic preference relation. We are working to accommodate default knowledge, irrelevant fluents, degradation of the cost over time and other domain-specific features. Our implementation will be extended to allow easy parameterization and customization of the preference relation to be used, even at run-time, in order to experiment with the behaviour of different preferences and preference families.

Also, we are planning on establishing stronger connections with existing results from belief change (e.g., satisfaction of certain postulates, or connections between our preference relation and various selection functions or orderings that have been used in other contexts), thereby more thoroughly understanding the properties of the proposed framework. Further, even though our theoretical framework is generic enough to support more complex flavours of action theories and \mathcal{DEC} , our implementation will need to be significantly extended to support different Event Calculus dialects and a richer set of commonsense features such as non-determinism, state constraints, and introspective belief changes.

References

- [Alchourron *et al.*, 1985] C. Alchourron, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [Dalal, 1988] M. Dalal. Investigations into a theory of knowledge base revision: Preliminary report. In *AAAI-88*, pages 475–479, 1988.
- [D’Asaro *et al.*, 2017] F. A. D’Asaro, A. Bikakis, L. Dickens, and R. Miller. Foundations for a probabilistic event calculus. In *LPNMR-17*, pages 57–63, 2017.
- [Demolombe and Pozos-Parra, 2000] R. Demolombe and M. P. Pozos-Parra. A simple and tractable extension of situation calculus to epistemic logic. In *ISMIS-00*, pages 515–524, 2000.
- [Ferraris *et al.*, 2011] P. Ferraris, J. Lee, and V. Lifschitz. Stable models and circumscription. *Artificial Intelligence*, 175(1):236–263, 2011.
- [Flouris *et al.*, 2006] G. Flouris, D. Plexousakis, and G. Antoniou. On generalizing the AGM postulates. In *STAIRS-06*, pages 132–143, 2006.
- [Gärdenfors and Makinson, 1988] P. Gärdenfors and D. Makinson. Revisions of knowledge systems using epistemic entrenchment. In *TARK-88*, pages 83–95, 1988.
- [Gebser *et al.*, 2011] Martin Gebser, Benjamin Kaufmann, Roland Kaminski, Max Ostrowski, Torsten Schaub, and Marius Schneider. Potassco: The Potsdam answer set solving collection. *AI Communications*, 24(2):107–124, 2011.
- [Katsuno and Mendelzon, 1991] H. Katsuno and A.O. Mendelzon. On the difference between updating a knowledge base and revising it. In *KR-91*, 1991.
- [Kelly and Pearce, 2008] R.F. Kelly and A.R. Pearce. Complex epistemic modalities in the situation calculus. In *KR-08*, pages 611–620, 2008.
- [Kowalski and Sergot, 1986] R. Kowalski and M. Sergot. A logic-based calculus of events. *New Generation Computing*, 4(1):67–95, 1986.
- [Lee and Palla, 2012] J. Lee and R. Palla. Reformulating the situation calculus and the event calculus in the general theory of stable models and in answer set programming. *JAIR*, 43(1):571–620, 2012.
- [Liu and Lakemeyer, 2009] Y. Liu and G. Lakemeyer. On first-order definability and computability of progression for local-effect actions and beyond. In *IJCAI-09*, 2009.
- [Ma *et al.*, 2013] J. Ma, R. Miller, L. Morgenstern, and T. Patkos. An Epistemic Event Calculus for ASP-based reasoning about knowledge of the past, present and future. In *LPAR-13*, pages 75–87, 2013.
- [Miller and Shanahan, 2002] R. Miller and M. Shanahan. Some alternative formulations of the event calculus. *Computational Logic: Logic Programming and Beyond, Essays in Honour of R. Kowalski Part 2*, 2408(1):452–490, 2002.
- [Miller *et al.*, 2013] R. Miller, L. Morgenstern, and T. Patkos. Reasoning about knowledge and action in an epistemic event calculus. In *Commonsense-13*, 2013.
- [Moore, 1985] R. C. Moore. A formal theory of knowledge and action. In *Formal Theories of the Commonsense World*, pages 319–358. J. Hobbs, R. Moore (Eds.), 1985.
- [Morgenstern, 1987] L. Morgenstern. Knowledge preconditions for actions and plans. In *IJCAI-87*, 1987.
- [Mueller, 2015] E.T. Mueller. *Commonsense Reasoning: An Event Calculus Based Approach*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2nd edition, 2015.
- [Patkos and Plexousakis, 2009] T. Patkos and D. Plexousakis. Reasoning with knowledge, action and time in dynamic and uncertain domains. In *IJCAI-09*, 2009.
- [Petrick and Levesque, 2002] R. Petrick and H. Levesque. Knowledge equivalence in combined action theories. In *KR-02*, pages 303–314, 2002.
- [Qi and Du, 2009] G. Qi and J. Du. Model-based revision operators for terminologies in Description Logics. In *IJCAI-09*, pages 891–897, 2009.
- [Ribeiro *et al.*, 2013] M.M. Ribeiro, R. Wassermann, G. Flouris, and G. Antoniou. Minimal change: Relevance and recovery revisited. 201:59–80, 2013.
- [Scherl and Levesque, 2003] R. Scherl and H. Levesque. Knowledge, action, and the frame problem. *Artificial Intelligence*, 144(1-2):1–39, 2003.
- [Scherl, 2003] R.B. Scherl. Reasoning about the interaction of knowledge, time and concurrent actions in the situation calculus. In *IJCAI-03*, pages 1091–1096, 2003.
- [Schwering *et al.*, 2015] C. Schwering, G. Lakemeyer, and M. Pagnucco. Belief revision and progression of knowledge bases in the epistemic situation calculus. In *IJCAI-15*, 2015.
- [Shapiro *et al.*, 2011] S. Shapiro, M. Pagnucco, Y. Lespérance, and H.J. Levesque. Iterated belief change in the situation calculus. *Artificial Intelligence*, 175(1):165–192, 2011.
- [Skarlatidis *et al.*, 2015] A. Skarlatidis, A. Artikis, J. Filippou, and G. Paliouras. A probabilistic logic programming event calculus. *TPLP*, 15:213–245, 2015.
- [Son and Baral, 2001] T. C. Son and C. Baral. Formalizing sensing actions – a transition function based approach. *Artificial Intelligence*, 125(1-2):19–91, 2001.
- [Thielscher, 2000] M. Thielscher. Representing the knowledge of a robot. In *KR-00*, pages 109–120, 2000.
- [Van Harmelen *et al.*, 2007] F. Van Harmelen, V. Lifschitz, and B. Porter. *Handbook of Knowledge Representation*. Elsevier Science, San Diego, USA, 2007.
- [Van Zee *et al.*, 2015] M. Van Zee, D. Doder, M. Dastani, and L. Van Der Torre. AGM revision of beliefs about action and time. In *IJCAI-15*, pages 3250–3256, 2015.
- [Vassos and Levesque, 2007] S. Vassos and H. Levesque. Progression of situation calculus action theories with incomplete information. In *IJCAI-07*, 2007.