

On Defaults in Action Theories

Hannes Strass and Michael Thielscher

Department of Computer Science
Dresden University of Technology
{hannes.strass,mit}@inf.tu-dresden.de

Abstract. We study the integration of two prominent fields of logic-based AI: action formalisms and non-monotonic reasoning. The resulting framework allows an agent employing an action theory as internal world model to make useful default assumptions. We show that the mechanism behaves properly in the sense that all intuitively possible conclusions can be drawn and no implausible inferences arise. In particular, it suffices to make default assumptions only once (in the initial state) to solve projection problems.

1 Introduction

This paper combines works of two important areas of logic-based artificial intelligence: we propose to enrich formalisms for reasoning about actions and change with default logic. The present work is not the first to join the two areas—non-monotonic logics have already been used by the reasoning about actions community in the past. After McCarthy and Hayes discovered the fundamental problem of determining the non-effects of actions, the frame problem [1], it was widely believed that non-monotonic reasoning were necessary to solve it. Then Hanks and McDermott gave an example of how straightforward use of non-monotonic logics in reasoning about actions and change can lead to counter-intuitive results [2]. When monotonic solutions to the frame problem were found [3, 4], non-monotonic reasoning again seemed to be obsolete.

In this paper, we argue that utilizing default logic still is of use when reasoning about actions. We will not use it to solve the frame problem—the solution to the frame problem we use here is monotonic and similar to the one of [4]—, but to make useful default assumptions about states. We consider action theories with deterministic actions where all effects are unconditional and a restricted form of default assumptions. The main reasoning task we are interested in is the projection problem, that is, given an initial state and a sequence of actions, the question whether a certain condition holds in the resulting state. As the main result of this paper, we show that restricting default application to the initial state not only guarantees a maximal set of states that are reachable, but also all possible inferences about these states.

The rest of the paper is organized as follows. The next section introduces the two areas this work is concerned with. Section 3 then combines the two fields and develops the main results. In the last section, we shortly sketch the limitations of our approach, outline directions for further work, and conclude.

2 Background

This section presents the formal underpinnings of the paper. In the first subsection we acquaint the reader with a unifying action calculus that we use to logically formalize action domains, and in the second subsection we recall the notions of default logic [5].

2.1 The Unifying Action Calculus

The stated objective of introducing the unifying action calculus (UAC) in [6] was to provide a universal framework for research in reasoning about actions. Since we want to formulate our results in a most general manner, we adopt the UAC for the present work.

The most notable generalization established by the UAC is its abstraction from the underlying time structure: it can be instantiated with formalisms using the time structure of situations (as the Situation Calculus [7] or the Fluent Calculus [4]), as well as with formalisms using a linear time structure (like the Event Calculus [8]).

The UAC uses only the sorts FLUENT, ACTION, and TIME along with the predicates $< : \text{TIME} \times \text{TIME}$ (denoting an ordering of time points), $Holds : \text{FLUENT} \times \text{TIME}$ (stating whether a fluent evaluates to true at a given time point), and $Poss : \text{ACTION} \times \text{TIME} \times \text{TIME}$ (indicating whether an action is applicable for particular starting and ending time points). Uniqueness-of-names is assumed for all (finitely many) functions into sorts FLUENT and ACTION.

The following definition introduces the most important types of formulas of the unifying action calculus: they allow to express properties of states and applicability conditions and effects of actions.

Definition 1. *Let \mathbf{s} be a sequence of variables of sort TIME.*

- A state formula $\Phi[\mathbf{s}]$ in \mathbf{s} is a first-order formula with free variables \mathbf{s} where
 - for each occurrence of $Holds(\varphi, s)$ in $\Phi[\mathbf{s}]$ we have $s \in \mathbf{s}$ and
 - predicate $Poss$ does not occur.

Let s, t be variables of sort TIME and A be a function into sort ACTION.

- A precondition axiom is of the form

$$Poss(A(\mathbf{x}), s, t) \equiv \pi_A[s] \quad (1)$$

where $\pi_A[s]$ is a state formula in s with free variables among s, t, \mathbf{x} .

- An effect axiom is of the form

$$Poss(A(\mathbf{x}), s, t) \supset (\forall f)(Holds(f, t) \equiv (\gamma_A^+ \vee (Holds(f, s) \wedge \neg \gamma_A^-))) \quad (2)$$

where

$$\gamma_A^+ = \bigvee_{0 \leq i \leq n_A^+} f = \varphi_i \text{ and } \gamma_A^- = \bigvee_{0 \leq i \leq n_A^-} f = \psi_i$$

and the φ_i and ψ_i are terms of sort FLUENT with free variables among \mathbf{x} .

Readers may be curious as to why the predicate *Poss* carries two time arguments instead of just one: $Poss(a, s, t)$ is to be read as “action a is possible starting at time s and ending at time t .” The formulas γ_A^+ and γ_A^- enumerate the positive and negative effects of the action, respectively. This definition of effect axioms is a restricted version of the original definition of [6]—it only allows for deterministic actions with unconditional effects.

Definition 2. A (UAC) domain axiomatization consists of a finite set of foundational axioms Ω (that define the underlying time structure and do not contain the predicates *Holds* and *Poss*), a set Π of precondition axioms (1), and a set Υ of effect axioms (2); the latter two for all functions into sort ACTION.

A domain axiomatization is progressing, if

- $\Omega \models (\exists s : \text{TIME})(\forall t : \text{TIME})s \leq t$
- $\Omega \cup \Pi \models Poss(a, s, t) \supset s < t$

A domain axiomatization is sequential, if it is progressing and

$$\Omega \cup \Pi \models Poss(a, s, t) \wedge Poss(a', s', t') \supset (t < t' \supset t \leq s') \wedge (t = t' \supset (a = a' \wedge s = s'))$$

That is, a domain axiomatization is progressing if there exists a least time point and time always increases when applying an action. A sequential domain axiomatization furthermore requires that no two actions overlap.

Since we are mainly interested in the projection problem, our domain axiomatizations will usually include a set Σ_0 of state formulas in the least time point that characterize the initial state.

To illustrate the intended usage of the introduced notions, we make use of a variant of the example of [2], the Yale Shooting scenario.

Example 1. Consider the domain axiomatization $\Sigma = \Omega_{sit} \cup \Pi \cup \Upsilon \cup \Sigma_0$ ¹. The precondition axioms say that the action *Shoot* is possible if the gun is loaded and the actions *Load* and *Wait* are always possible.

$$\begin{aligned} \Pi = \{ & Poss(\text{Shoot}, s, t) \equiv (\text{Holds}(\text{Loaded}, s) \wedge t = \text{Do}(\text{Shoot}, s)), \\ & Poss(\text{Load}, s, t) \equiv t = \text{Do}(\text{Load}, s), Poss(\text{Wait}, s, t) \equiv t = \text{Do}(\text{Wait}, s) \} \end{aligned}$$

With these preconditions and foundational axioms Ω_{sit} , the domain axiomatization is sequential. The effect of shooting is that the turkey ceases to be alive, loading the gun causes it to be loaded, and waiting does not have any effect. All effect axioms in Υ are of the form (2), we state only the γ^\pm different from the empty disjunction: $\gamma_{\text{Load}}^+ = (f = \text{Loaded})$, $\gamma_{\text{Shoot}}^- = (f = \text{Alive})$. Finally, we state that the turkey is alive in the initial situation, $\Sigma_0 = \{\text{Holds}(\text{Alive}, S_0)\}$. We can now employ logical entailment to answer the question whether the turkey is still alive after applying the actions *Load*, *Wait*, and *Shoot*, respectively. With the notation $\text{Do}([a_1, \dots, a_n], s)$ as abbreviation for $\text{Do}(a_n, \text{Do}(\dots, \text{Do}(a_1, s) \dots))$, it is easy to see that $\Sigma \models \neg \text{Holds}(\text{Alive}, \text{Do}([\text{Load}, \text{Wait}, \text{Shoot}], S_0))$.

¹ Ω_{sit} denotes the foundational axioms for situations. They have been omitted from the presentation due to a lack of space and can be found in [9].

2.2 Default Logic

Introduced in the seminal work by Reiter [5], default logic has become one of the most important formalisms for non-monotonic reasoning. The semantics for supernormal defaults used here is taken from [10], which is itself an enhancement of a notion developed in [11].

Definition 3. A supernormal default rule, or, for short, default, is a closed first-order formula. Any formulas with occurrences of free variables are taken as representatives of their ground instances.

For a set of closed formulas S , we say the default δ is active in S if both $\delta \notin S$ and $\neg\delta \notin S$.

A (supernormal) default theory is a pair (W, \mathcal{D}) , where W is a set of sentences and \mathcal{D} a set of default rules.

A default rule can thus also be seen as a hypothesis that we are willing to assume, but prepared to give up in case of contradiction. A default theory then adds a set of formulas, the *indefeasible knowledge*, that we are not willing to give up for any reason.

Definition 4. Let (W, \mathcal{D}) be a default theory where all default rules are supernormal and \prec be a total order on \mathcal{D} . Define $E_0 := Th(W)$ and for all $i \geq 0$,

$$E_{i+1} = \begin{cases} E_i & \text{if no default is active in } E_i \\ Th(E_i \cup \{\delta\}) & \text{otherwise, where } \delta \text{ is the } \prec\text{-} \\ & \text{minimal default active in } E_i. \end{cases}$$

Then the set $E := \bigcup_{i \geq 0} E_i$ is called the extension generated by \prec .

A set of formulas E is a preferred extension for (W, \mathcal{D}) if there exists a total order \prec that generates E . The set of all preferred extensions for a default theory (W, \mathcal{D}) is denoted by $Ex(W, \mathcal{D})$.

An extension for a default theory can be seen as a way of assuming as many defaults as possible without creating inconsistencies. It should be noted that, although the definition differs, our extensions are extensions in the sense of [5].

Based on extensions, one can define skeptical and credulous conclusions for default theories: skeptical conclusions are formulas that are contained in every extension, credulous conclusions are those that are contained in at least one extension.

Definition 5. Let (W, \mathcal{D}) be a supernormal default theory and Ψ be a first-order formula.

$$W \approx_{\mathcal{D}}^{skept} \Psi \stackrel{\text{def}}{=} \Psi \in \bigcap_{E \in Ex(W, \mathcal{D})} E, \quad W \approx_{\mathcal{D}}^{cred} \Psi \stackrel{\text{def}}{=} \Psi \in \bigcup_{E \in Ex(W, \mathcal{D})} E$$

3 Domain Axiomatizations with Defaults

The concepts established up to this point are now easily combined to the notion of a domain axiomatization with defaults, our main object of study. It is essentially a default theory having an action domain axiomatization as infeasible knowledge.

Definition 6. A domain axiomatization with defaults is a pair $(\Sigma, \mathcal{D}[s])$, where Σ is a UAC domain axiomatization and $\mathcal{D}[s]$ is a set of supernormal defaults of the form $\text{Holds}(\varphi, s)$ or $\neg\text{Holds}(\varphi, s)$ for a fluent φ .

We next define what it means for a time point to be reachable in an action domain. Intuitively, it means that there is a sequence of actions that leads to the time point when applied in sequence starting in the initial time point.

Definition 7. Let Σ be a domain axiomatization and $\mathcal{D}[s]$ be a set of defaults.

$$\text{Reach}(r) \stackrel{\text{def}}{=} (\forall R)((\forall s)(\text{Init}(s) \supset R(s)) \wedge (\forall a, s, t)(R(s) \wedge \text{Poss}(a, s, t) \supset R(t))) \supset R(r) \quad (3)$$

$$\text{Init}(t) \stackrel{\text{def}}{=} \neg(\exists s)s < t \quad (4)$$

A time point τ is called

- finitely reachable in Σ if $\Sigma \models \text{Reach}(\tau)$;
- finitely, credulously reachable in $(\Sigma, \mathcal{D}[\sigma])^2$, if σ is finitely reachable in Σ and for some extension E for $(\Sigma, \mathcal{D}[\sigma])$ we have $E \models \text{Reach}(\tau)$;
- finitely, weakly skeptically reachable in $(\Sigma, \mathcal{D}[\sigma])$, if σ is finitely reachable in Σ and for all extensions E for $(\Sigma, \mathcal{D}[\sigma])$, we have $E \models \text{Reach}(\tau)$;
- finitely, strongly skeptically reachable in $(\Sigma, \mathcal{D}[\sigma])$, if σ is finitely reachable in Σ and there exist ground actions $\alpha_1, \dots, \alpha_n$ and time points τ_0, \dots, τ_n such that $\Sigma \approx_{\mathcal{D}[\sigma]}^{\text{skept}} \text{Poss}(\alpha_i, \tau_{i-1}, \tau_i)$ for all $1 \leq i \leq n$, $\tau_0 = \sigma$, and $\tau_n = \tau$.

With situations as underlying time structure, weak and strong skeptical reachability coincide. This is because the foundational axioms for situations [9] entail that situations have unique predecessors.

Example 1 (continued). We add a fluent **Broken** that indicates if the gun does not function properly. Shooting is now only possible if the gun is loaded *and* not broken:

$$\text{Poss}(\text{Shoot}, s, t) \equiv (\text{Holds}(\text{Loaded}, s) \wedge \neg\text{Holds}(\text{Broken}, s) \wedge t = \text{Do}(\text{Shoot}, s))$$

Unless there is information to the contrary, it should be assumed that the gun has no defects. This is expressed by the set of defaults $\mathcal{D}[s] = \{\neg\text{Holds}(\text{Broken}, s)\}$. Without the default assumption, it cannot be concluded that the action **Shoot** is

² By $\mathcal{D}[\sigma]$ we denote the set of defaults in $\mathcal{D}[s]$ where s has been instantiated by the term σ .

possible after performing **Load** and **Wait** since it cannot be inferred that the gun is not broken. Using the abbreviations $S_1 = Do(\text{Load}, S_0)$, $S_2 = Do(\text{Wait}, S_1)$, and $S_3 = Do(\text{Shoot}, S_2)$, we illustrate how the non-monotonic entailment relation defined earlier enables us to use the default rule to draw the desired conclusion:

$$\begin{aligned}\Sigma &\approx_{\mathcal{D}[S_0]}^{\text{skpt}} \neg \text{Holds}(\text{Broken}, S_2), \\ \Sigma &\approx_{\mathcal{D}[S_0]}^{\text{skpt}} \text{Poss}(\text{Shoot}, S_2, S_3), \text{ and} \\ \Sigma &\approx_{\mathcal{D}[S_0]}^{\text{skpt}} \neg \text{Holds}(\text{Alive}, S_3).\end{aligned}$$

The default conclusion that the gun works correctly, drawn in S_0 , carries over to S_2 and allows to conclude applicability of **Shoot** in S_2 and its effects on S_3 .

In the example just seen, default reasoning could be restricted to the initial situation. As it turns out, this is sufficient for the type of action domain considered here: effect axiom (2) never “removes” information about fluents and thus never makes more defaults active after executing an action. This observation is formalized by the following lemma. It essentially says that to reason about a time point in which an action ends, it makes no difference whether we apply the defaults to the resulting time point or to the time point when the action starts. This holds of course only due to the restricted nature of effect axiom (2).

Lemma 1. *Let $(\Sigma, \mathcal{D}[s])$ be a domain axiomatization with defaults, α be a ground action such that $\Sigma \models \text{Poss}(\alpha, \sigma, \tau)$ for some $\sigma, \tau : \text{TIME}$, and let $\Psi[\tau]$ be a state formula in τ . Then*

$$\Sigma \approx_{\mathcal{D}[\sigma]}^{\text{skpt}} \Psi[\tau] \text{ iff } \Sigma \approx_{\mathcal{D}[\tau]}^{\text{skpt}} \Psi[\tau]$$

The next theorem says that all *local* conclusions about a finitely reachable time point σ (that is, all conclusions about σ using defaults from $\mathcal{D}[\sigma]$) are exactly the conclusions about σ that we can draw by instantiating the defaults only with the least time point.

Theorem 1. *Let $(\Sigma, \mathcal{D}[s])$ be a progressing domain axiomatization with defaults, λ its least time point, $\sigma : \text{TIME}$ be finitely reachable in Σ , and $\Psi[\sigma]$ be a state formula. Then*

$$\Sigma \approx_{\mathcal{D}[\sigma]}^{\text{skpt}} \Psi[\sigma] \text{ iff } \Sigma \approx_{\mathcal{D}[\lambda]}^{\text{skpt}} \Psi[\sigma]$$

It thus remains to show that local defaults are indeed exhaustive with respect to local conclusions. The next lemma takes a step into this direction: it states that action application does not increase default knowledge about past time points.

Lemma 2. *Let $(\Sigma, \mathcal{D}[s])$ be a domain axiomatization with defaults, α be a ground action such that $\Sigma \models \text{Poss}(\alpha, \sigma, \tau)$ for some $\sigma, \tau : \text{TIME}$, and let $\Psi[\rho]$ be a state formula in $\rho : \text{TIME}$ where $\rho \leq \sigma$. Then*

$$\Sigma \approx_{\mathcal{D}[\tau]}^{\text{skpt}} \Psi[\rho] \text{ implies } \Sigma \approx_{\mathcal{D}[\sigma]}^{\text{skpt}} \Psi[\rho]$$

The converse of the lemma does not hold, since an action effect might preclude a default conclusion about the past. Using the above lemma and simple induction on the length of action sequences, one can establish the following.

Theorem 2. *Let $(\Sigma, \mathcal{D}[s])$ be a progressing domain axiomatization with defaults, let $\Psi[s]$ be a state formula, $\sigma < \tau$ be time points, and σ be finitely reachable in Σ . Then*

$$\Sigma \models_{\mathcal{D}[\tau]}^{\text{skept}} \Psi[\sigma] \text{ implies } \Sigma \models_{\mathcal{D}[\sigma]}^{\text{skept}} \Psi[\sigma]$$

The next theorem, our first main result, now combines Theorems 1 and 2 and tells us that default instantiation to the least time point subsumes default instantiation in any time point in the future of the time point we want to reason about.

Theorem 3. *Let $(\Sigma, \mathcal{D}[s])$ be a progressing domain axiomatization with defaults, λ be its least time point, $\Psi[s]$ be a state formula, and $\sigma < \tau$ be terms of sort TIME where σ is finitely reachable in Σ . Then*

$$\Sigma \models_{\mathcal{D}[\tau]}^{\text{skept}} \Psi[\sigma] \text{ implies } \Sigma \models_{\mathcal{D}[\lambda]}^{\text{skept}} \Psi[\sigma]$$

Proof. $\Sigma \models_{\mathcal{D}[\tau]}^{\text{skept}} \Psi[\sigma]$ implies $\Sigma \models_{\mathcal{D}[\sigma]}^{\text{skept}} \Psi[\sigma]$ by Theorem 2. By Theorem 1, this is the case iff $\Sigma \models_{\mathcal{D}[\lambda]}^{\text{skept}} \Psi[\sigma]$.

What this theorem misses out, however, are time points that are not finitely reachable in Σ only, but where some action application along the way depends crucially on a default conclusion. To illustrate this, recall Example 1: the situation $Do([\text{Load, Wait, Shoot}], S_0)$ is not reachable in Σ , because the necessary precondition that the gun is not broken cannot be inferred without the respective default.

The following theorem, our second main result, now assures sufficiency of instantiation with the least time point also for time points that are only reachable by default.

Theorem 4. *Let $(\Sigma, \mathcal{D}[s])$ be a progressing domain axiomatization with defaults, λ its least time point, σ be a time point that is finitely reachable in Σ , $\Psi[s]$ be a state formula, and τ be a time point that is finitely, strongly skeptically reachable in $(\Sigma, \mathcal{D}[\sigma])$. Then*

1. τ is finitely, strongly skeptically reachable in $(\Sigma, \mathcal{D}[\lambda])$, and
2. $\Sigma \models_{\mathcal{D}[\sigma]}^{\text{skept}} \Psi[\tau]$ iff $\Sigma \models_{\mathcal{D}[\lambda]}^{\text{skept}} \Psi[\tau]$.

An immediate consequence of this result is that instantiation in the least time point also provides a “maximal” number of reachable time points: default instantiation with a later time point might potentially prevent actions in the least time point, which in turn might render yet another time point unreachable.

4 Conclusions and Future Work

We have presented an enrichment of action theories with a well-known non-monotonic logic, Raymond Reiter’s default logic. To the best of our knowledge, this is the first time this field is explored in the logic-based AI community. The approach has been shown to behave well (although no proofs could be included due to space limitations)—by the restrictions made in the definitions, defaults persist over time and it thus suffices to apply them only once (namely to the initial state).

With respect to further generalizations of our proposal, we remark that both allowing for disjunctive defaults and allowing for conditional effects causes unintuitive conclusions via the employed solution to the frame problem. Future research in this topic will therefore be devoted to generalizing both the defaults (from supernormal to normal) and the considered actions (from deterministic with unconditional effects to non-deterministic with conditional effects).

References

1. McCarthy, J., Hayes, P.J.: Some Philosophical Problems from the Standpoint of Artificial Intelligence. In: Machine Intelligence, Edinburgh University Press (1969) 463–502
2. Hanks, S., McDermott, D.: Nonmonotonic Logic and Temporal Projection. *Artificial Intelligence* **33**(3) (1987) 379–412
3. Reiter, R.: The Frame Problem in the Situation Calculus: A Simple Solution (Sometimes) and a Completeness Result for Goal Regression. In: Artificial Intelligence and Mathematical Theory of Computation – Papers in Honor of John McCarthy, Academic Press (1991) 359–380
4. Thielscher, M.: From Situation Calculus to Fluent Calculus: State Update Axioms as a Solution to the Inferential Frame Problem. *Artificial Intelligence* **111**(1–2) (1999) 277–299
5. Reiter, R.: A Logic for Default Reasoning. *Artificial Intelligence* **13** (1980) 81–132
6. Thielscher, M.: A Unifying Action Calculus. *Artificial Intelligence* **to appear** (2009)
7. McCarthy, J.: Situations and Actions and Causal Laws (1963) Stanford Artificial Intelligence Project: Memo 2.
8. Kowalski, R.A., Sergot, M.J.: A Logic-based Calculus of Events. *New Generation Computing* **4**(1) (1986) 67–95
9. Pirri, F., Reiter, R.: Some Contributions to the Metatheory of the Situation Calculus. *Journal of the ACM* **46**(3) (1999) 325–361
10. Brewka, G., Eiter, T.: Prioritizing Default Logic: Abridged Report. In: Festschrift on the occasion of Prof. Dr. W. Bibel’s 60th birthday, Kluwer (1999)
11. Brewka, G.: Preferred Subtheories: An Extended Logical Framework for Default Reasoning. In: Proceedings of the Eleventh International Conference on Artificial Intelligence. (1989) 1043–1048