Elaborating domain descriptions

Andreas Herzig **Laurent Perrussel** **Ivan Varzinczak***
IRIT – 118 route de Narbonne
31062 Toulouse Cedex – France
e-mail: {herzig,perrusse,ivan}@irit.fr

Abstract

In this work we address the problem of *elaborating* domain descriptions (alias action theories), in particular those that are expressed in dynamic logic. We define a general method based on contraction of formulas in a version of propositional dynamic logic with an incorporated solution to the frame problem. We present the semantics of our theory change and define syntactical operators for contracting a domain description. We establish soundness and completeness of the operators w.r.t. the semantics for descriptions that satisfy a principle of modularity that we have proposed elsewhere. We also investigate an example of changing non-modular domain descriptions.

Introduction

Suppose a situation where an agent has always believed that if the light switch is up, then there is light in the room. Suppose now that someday, she observes that even if the switch is in the upper position, the light is off. In such a case, the agent must change her beliefs about the relation between the propositions "the switch is up" and "the light is on". This example is an instance of the problem of changing propositional belief bases and is largely addressed in the literature about belief change (Gärdenfors 1988) and belief update (Katsuno & Mendelzon 1992).

Next, let our agent believe that whenever the switch is down, after toggling it, there is light in the room. This means that if the light is off, in every state of the world that follows the execution of toggling the switch, the room is lit up. Then, during a blackout, the agent toggles the switch and surprisingly the room is still dark.

Imagine now that the agent never worried about the relation between toggling the switch and the material it is made of, in the sense that she ever believed that just toggling the switch does not break it. Nevertheless, in a stressful day, she toggles the switch and then observes that she had broken it.

Completing the wayside cross our agent experiments in discovering the world's behavior, suppose she has believed that it is always possible to toggle the switch, provided some conditions like being close enough to it, having a free hand, the switch is not broken, etc, are satisfied. However, in a beautiful April fool's day, the agent discovers that someone has glued the switch and, consequently, it is no longer possible to toggle it.

The last three examples illustrate situations where changing the beliefs about the behavior of the action of toggling the switch is mandatory. In the first one, toggling the switch, once believed to be deterministic, has now to be seen as nondeterministic, or alternatively to have a different outcome in a specific context (e.g. if the power station is overloaded). In the second example, toggling the switch is known to have side-effects (ramifications) one was not aware of. In the last example, the executability of the action under concern is questioned in the light of new information showing a context that was not known to preclude its execution. Carrying out modifications is what we here call *elaborating* a domain description, which has to do with the principle of *elaboration tolerance* (McCarthy 1988).

Such cases of theory change are very important when one deals with logical descriptions of dynamic domains: it may always happen that one discovers that an action actually has a behavior that is different from that one has always believed it had.

Up to now, theory change has been studied mainly for knowledge bases in classical logics, both in terms of revision and update. Only in a few recent works it has been considered in the realm of modal logics, viz. in epistemic logic (Hansson 1999), and in action languages (Eiter *et al.* 2005). Recently, several works (Shapiro *et al.* 2000; Jin & Thielscher 2005) have investigated revision of beliefs about facts of the world. In our examples, this would concern e.g. the current status of the switch: the agent believes it is up, but is wrong about this and might subsequently be forced to revise his beliefs about the current state of affairs. Such belief revision operations do not modify the agent's beliefs about the action laws. In opposition to that, here we are interested exactly in such modifications. The aim of this work is to make a step toward that issue and propose a framework that deals with the contraction of action theories.

Dynamic logic, more specifically propositional dynamic logic (PDL (Harel 1984)), has been extensively used in reasoning about actions in the last years (Castilho, Gasquet, & Herzig 1999; Castilho, Herzig, & Varzinczak 2002; Zhang & Foo 2001; Foo & Zhang 2002; Zhang, Chopra, &

*Supported by a fellowship from the government of the FEDERATIVE REPUBLIC OF BRAZIL. Grant: CAPES BEX 1389/01-7.

Foo 2002). It has shown to be a viable alternative to situation calculus approaches because of its simplicity and existence of proof procedures for it. In this work we investigate the elaboration of domain descriptions encoded in a simplified version of such a logical formalism, viz. the multimodal logic K_n. We show how a theory expressed in terms of static laws, effect laws and executability laws is elaborated: usually, a law has to be changed due to its generality, i.e., the law is too strong and has to be weakened. It follows that elaborating an action theory means contracting it by static, effect or executability laws, before expanding the theory with more specific laws.

The present text is organized as follows: in the next section we define the logical framework we use throughout this work and show how action theories are encoded. Then we present our semantics of theory change and its syntactical counterpart. After that we establish soundness and completeness of our change operators w.r.t. the semantics, where completeness is conditioned by a notion of modularity that we have proposed in previous work. We then analyse an example of correcting a non-modular theory. Before concluding, we address related work on the field and discuss on how elaboration tolerant the framework here proposed is.

Background

Following the tradition in the reasoning about actions community, action theories are going to be collections of statements that have the particular form: "if *context*, then *effect* after *every execution* of *action*" (effect laws); and "if *precondition*, then *action executable*" (executability laws). Statements mentioning no action at all represent laws about the world (static laws). Besides that, statements of the form "if *context*, then *effect* after *some execution* of *action*" will be used as a causal notion to solve the frame and the ramification problems.

Logical preliminaries

Let $\mathfrak{Act} = \{a_1, a_2, \ldots\}$ be the set of all *atomic action constants* of a given domain. An example of atomic action is *toggle*. To each atomic action a there is associated a modal operator $[a]$.

$\mathfrak{Prop} = \{p_1, p_2, \ldots\}$ denotes the set of all *propositional constants*, also called *fluents* or *atoms*. Examples of those are *light* ("the light is on") and *up* ("the switch is up"). The set of all literals is $\mathfrak{Lit} = \{l_1, l_2, \ldots\}$, where each l_i is either p or $\neg p$, for some $p \in \mathfrak{Prop}$. If $l = \neg p$, then we identify $\neg l$ with p.

We use small Greek letters $\varphi, \psi, \ldots$ to denote *classical formulas*. They are recursively defined in the usual way:

$$\varphi ::= p \mid \top \mid \bot \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \rightarrow \varphi \mid \varphi \leftrightarrow \varphi$$

$\mathfrak{Fml}$ is the set of all classical formulas. An example of a classical formula is $up \rightarrow light$. By $val(\varphi)$ we denote the set of valuations making φ true. We view a valuation as a maximally-consistent set of literals. For $\mathfrak{Prop} = \{light, up\}$, there are four valuations: $\{light, up\}$, $\{light, \neg up\}$, $\{\neg light, up\}$ and $\{\neg light, \neg up\}$. Given a set of formulas Σ, by $lit(\Sigma)$ we denote the set of all literals appearing in formulas of Σ.

We denote complex formulas (with modal operators) by $\Phi, \Psi, \ldots$ They are recursively defined in the following way:

$$\Phi ::= \varphi \mid [a]\Phi \mid \neg\Phi \mid \Phi \wedge \Phi \mid \Phi \vee \Phi \mid \Phi \rightarrow \Phi \mid \Phi \leftrightarrow \Phi$$

$\langle a \rangle$ is the dual operator of $[a]$, defined as $\langle a \rangle\Phi =_{\text{def}} \neg[a]\neg\Phi$. An example of a complex formula is $\neg up \rightarrow [toggle]up$.

The semantics is that of multimodal logic K (Popkorn 1994).

Definition 1 A K_n-*model* is a tuple $\mathscr{M} = \langle W, R \rangle$ where W is a set of valuations, and R a function mapping action constants a to accessibility relations $R_a \subseteq W \times W$.

Definition 2 Given a K_n-model $\mathscr{M} = \langle W, R \rangle$,

- $\models_w^{\mathscr{M}} p$ (p is true at world w of model $\mathscr{M}$) if $p \in w$;
- $\models_w^{\mathscr{M}} [a]\Phi$ if for every w' such that $wR_a w'$, $\models_{w'}^{\mathscr{M}} \Phi$;
- truth conditions for the other connectives are as usual.

Definition 3 $\mathscr{M}$ is a model of Φ (noted $\models^{\mathscr{M}} \Phi$) if and only if for all $w \in W$, $\models_w^{\mathscr{M}} \Phi$. $\mathscr{M}$ is a model of a set of formulas Σ (noted $\models^{\mathscr{M}} \Sigma$) if and only if $\models^{\mathscr{M}} \Phi$ for every $\Phi \in \Sigma$. A formula Φ is a *consequence of the set of global axioms* Γ in the class of all K_n-models (noted $\Gamma \models_{\mathsf{K}_n} \Phi$) if and only if for every K_n-model $\mathscr{M}$, if $\models^{\mathscr{M}} \Gamma$, then $\models^{\mathscr{M}} \Phi$.

Describing the behavior of actions in K_n

Given a domain, we are interested in theories whose statements describe the behavior of actions. K_n allows for the representation of such statements, that we call *action laws*. Here we distinguish several types of them. The first kind of statement represents the *static laws*, which are formulas that must hold in every possible state of the world.

Definition 4 A *static law* is a formula $\varphi \in \mathfrak{Fml}$.

An example of a static law is $up \rightarrow light$, saying that if the switch is up, then the light is on. The set of all static laws of a domain is denoted by $\mathcal{S} \subseteq \mathfrak{Fml}$.

The second kind of action law we consider is given by the *effect laws*. These are formulas relating an action to its effects, which can be conditional.

Definition 5 An *effect law for action a* is of the form $\varphi \rightarrow [a]\psi$, where $\varphi, \psi \in \mathfrak{Fml}$.

The consequent ψ is the effect which always obtains when action a is executed in a state where the antecedent φ holds. If a is a nondeterministic action, then the consequent ψ is typically a disjunction. The set of effect laws of a domain is denoted by $\mathcal{E}$. An example of an effect law is $\neg up \rightarrow [toggle]light$, saying that whenever the switch is down, after toggling it, the room is lit up. If ψ is inconsistent, we have a special kind of effect law that we call an *inexecutability law*. For example, $broken \rightarrow [toggle]\bot$ expresses that *toggle* cannot be executed if the switch is broken.

Finally, we also define *executability laws*, which stipulate the context where an action is guaranteed to be executable. In K_n, the operator $\langle a \rangle$ is used to express executability. $\langle a \rangle\top$ thus reads "the execution of a is possible".

Definition 6 An *executability law for action a* is of the form $\varphi \to \langle a \rangle \top$, where $\varphi \in \mathfrak{Fml}$.

For instance, $\neg broken \to \langle toggle \rangle \top$ says that toggling can be executed whenever the switch is not broken. The set of all executability laws of a given domain is denoted by $\mathcal{X}$.

The rest of this work is devoted to the elaboration of action models and theories.

Models of contraction

When an action theory has to be changed, the basic operation is that of *contraction*. (In belief-base update (Winslett 1988; Katsuno & Mendelzon 1992) it has also been called *erasure*.) In this section we define its semantics.

In general we might contract by any formula Φ. Here we focus on contraction by one of the three kinds of laws. We therefore suppose that Φ is either φ, where φ is classical, or $\varphi \to [a]\psi$, or $\varphi \to \langle a \rangle \top$.

For the case of contracting static laws we resort to existing approaches in order to change the set of static laws. In the following, we consider any belief change operator such as Forbus' update method (Forbus 1989), or the possible models approach (Winslett 1988; 1995), or WSS (Herzig & Rifi 1999) or MPMA (Doherty, Łukaszewicz, & Madalinska-Bugaj 1998).

Contraction by φ corresponds to adding new possible worlds to W. Let $\ominus$ be a contraction operator for classical logic.

Definition 7 Let $\langle W, R \rangle$ be a K_n-model and φ a classical formula. The set of models resulting from contracting by φ is the singleton $\langle W, R \rangle^-_{\varphi} = \{\langle W', R \rangle\}$ such that $W' = W \ominus val(\varphi)$.

Observe that R should, a priori, change as well, otherwise contracting a classical formula may conflict with $\mathcal{X}$.[1] For instance, if $\neg\varphi \to \langle a \rangle \top \in \mathcal{X}$ and we contract by φ, the result may make $\mathcal{X}$ untrue. However, given the amount of information we have at hand, we think that whatever we do with R (adding or removing edges), we will always be able to find a counter-example to the intuitiveness of the operation, since it is domain dependent. For instance, adding edges for a deterministic action may render it nondeterministic. Deciding on what changes to carry out on R when contracting static laws depends on the user's intuition, and unfortunately this information cannot be generalized and established once for all. We opt for a priori doing nothing with R and postponing correction of executability laws.

Action theories being defined in terms of effect and executability laws, elaborating an action theory will mainly involve changes in these two sets of laws. Let us consider now both these cases.

Suppose the knowledge engineer acquires new information regarding the effect of action a. Then it means that the law under consideration is probably too strong, i.e., the expected effect may not occur and thus the law has to be weakened. Consider e.g. $\neg up \to [toggle]light$, and suppose it has to be weakened to the more specific $(\neg up \wedge \neg blackout) \to [toggle]light$.[2] In order to carry out such a weakening, first the designer has to contract the set of effect laws and second to expand the resulting set with the weakened law.

Contraction by $\varphi \to [a]\psi$ amounts to adding some 'counterexample' arrows from φ-worlds to $\neg\psi$-worlds. To ease such a task, we need a definition. Let $PI(\varphi)$ denote the set of prime implicates of φ.

Definition 8 Let $\varphi_1, \varphi_2 \in \mathfrak{Fml}$. $NewCons_{\varphi_1}(\varphi_2) = PI(\varphi_1 \wedge \varphi_2) \setminus PI(\varphi_1)$ computes the *new consequences* of φ_2 w.r.t. φ_1: the set of strongest clauses that follow from $\varphi_1 \wedge \varphi_2$, but do not follow from φ_1 alone (cf. e.g. (Inoue 1992)).

For example, the set of prime implicates of p_1 is just $\{p_1\}$, that of the formula $p_1 \wedge (\neg p_1 \vee p_2) \wedge (\neg p_1 \vee p_3 \vee p_4)$ is $\{p_1, p_2, p_3 \vee p_4\}$, hence we have that $NewCons_{p_1}((\neg p_1 \vee p_2) \wedge (\neg p_1 \vee p_3 \vee p_4)) = \{p_2, p_3 \vee p_4\}$.

Definition 9 Let $\langle W, R \rangle$ be a K_n-model and $\varphi \to [a]\psi$ an effect law. The models resulting from contracting by $\varphi \to [a]\psi$ is $\langle W, R \rangle^-_{\varphi \to [a]\psi} = \{\langle W, R \cup R'_a \rangle : R'_a \subseteq \{(w, w') : \models^{\langle W,R \rangle}_{w} \varphi, \models^{\langle W,R \rangle}_{w'} \neg\psi \text{ and } w' \setminus w \subseteq lit(NewCons_{\mathcal{S}}(\neg\psi))\}\}$.

In our context, $lit(NewCons_{\mathcal{S}}(\neg\psi))$ corresponds to all the ramifications that action a can produce.

Suppose now the knowledge engineer learns new information about the executability of a. This usually occurs when there are executability laws that are too strong, i.e., the condition in the theory guaranteeing the executability of a is too weak and has to be made more restrictive. Let e.g. $\langle toggle \rangle \top$ be the law to be contracted, and suppose it has to be weakened to the more specific $\neg broken \to \langle toggle \rangle \top$. To implement such a weakening, the designer has to first contract the set of executability laws and then to expand the resulting set with the weakened law.

Contraction by $\varphi \to \langle a \rangle \top$ corresponds to removing some arrows leaving worlds where φ holds. Removing such arrows has as consequence that a is no longer always executable in context φ.

Definition 10 Let $\langle W, R \rangle$ be a K_n-model and $\varphi \to \langle a \rangle \top$ an executability law. The set of models that result from the contraction by $\varphi \to \langle a \rangle \top$ is $\langle W, R \rangle^-_{\varphi \to \langle a \rangle \top} = \{\langle W, R' \rangle : R' = R \setminus R''_a, R''_a \subseteq \{(w, w') : wR_a w' \text{ and } \models^{\langle W,R \rangle}_{w} \varphi\}\}$.

In the next section we make a step toward syntactical operators that reflect the semantic foundations for contraction.

---

[1] We are indebted to the anonymous referees for pointing this out to us.

[2] The other possibility of weakening the law, i.e., replacing it by $\neg up \to [toggle](light \vee \neg light)$ looks silly. We were not able to find examples where changing the consequent could give a more intuitive result. In this sense, we prefer to always weaken a given law by strengthening its antecedent.

Contracting an action theory

Having established the semantics of action theory contraction, we can turn to its syntactical counterpart. Nevertheless, before doing that we have to consider an important issue. As the reader might have expected, the logical formalism of K_n alone does not solve the frame problem. For instance,

$$\left\{\begin{array}{c} up \rightarrow light, \\ \neg up \rightarrow [toggle]up, \\ up \rightarrow [toggle]\neg up, \\ \langle toggle \rangle \top \end{array}\right\} \not\models_{\mathsf{K}_n} broken \rightarrow [toggle]broken.$$

Thus, we need a consequence relation powerful enough to deal with the frame and ramification problems. This means that the deductive power of K_n has to be augmented in order to ensure that the relevant frame axioms follow from the theory. Following the logical framework developed in (Castilho, Gasquet, & Herzig 1999), we consider metalogical information given in the form of a dependence relation:

Definition 11 A *dependence relation* is a binary relation $\leadsto \subseteq \mathfrak{Act} \times \mathfrak{Lit}$.

The expression $a \leadsto l$ denotes that the execution of action a *may* change the truth value of the literal l. On the other hand, $\langle a, l \rangle \notin \leadsto$ (written $a \not\leadsto l$) means that l can never be caused by a. In our example we have $toggle \leadsto light$ and $toggle \leadsto \neg light$, which means that action *toggle* may cause a change in literals *light* and $\neg light$. We do not have $toggle \leadsto \neg broken$, for toggling the switch never repairs it.

We assume $\leadsto$ is finite.

Definition 12 A *model of a dependence relation* $\leadsto$ is a K_n-model $\mathscr{M}$ such that $\models^{\mathscr{M}} \{\neg l \rightarrow [a]\neg l \; : \; a \not\leadsto l\}$.

Given a dependence relation $\leadsto$, the associated consequence relation in the set of models for $\leadsto$ is noted $\models_{\leadsto}$. For our example we obtain

$$\left\{\begin{array}{c} up \rightarrow light, \\ \neg up \rightarrow [toggle]up, \\ up \rightarrow [toggle]\neg up, \\ \langle toggle \rangle \top \end{array}\right\} \models_{\leadsto} broken \rightarrow [toggle]broken.$$

We have $toggle \not\leadsto \neg broken$, i.e., $\neg broken$ is never caused by *toggle*. Therefore in all contexts where *broken* is true, after every execution of *toggle*, *broken* still remains true. The consequence of this independence is that the frame axiom $broken \rightarrow [toggle]broken$ is valid in the models of $\leadsto$.

Such a dependence-based approach has been shown (Demolombe, Herzig, & Varzinczak 2003) to subsume Reiter's solution to the frame problem (Reiter 1991) and moreover treats the ramification problem, even when actions with both indeterminate and indirect effects are involved (Castilho, Herzig, & Varzinczak 2002; Herzig & Varzinczak 2004a).

Definition 13 An *action theory* is a tuple of the form $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$.

In our example, the corresponding action theory is

$$\mathcal{S} = \{up \rightarrow light\}, \; \mathcal{E} = \left\{\begin{array}{c} \neg up \rightarrow [toggle]up, \\ up \rightarrow [toggle]\neg up \end{array}\right\}$$

$$\mathcal{X} = \{\langle toggle \rangle \top\}, \; \leadsto = \left\{\begin{array}{c} \langle toggle, light \rangle, \\ \langle toggle, \neg light \rangle, \\ \langle toggle, up \rangle, \\ \langle toggle, \neg up \rangle \end{array}\right\}$$

And we have $\mathcal{S}, \mathcal{E}, \mathcal{X} \models_{\leadsto} \neg up \rightarrow [toggle]light$. (For parsimony's sake, we write $\mathcal{S}, \mathcal{E}, \mathcal{X} \models_{\leadsto} \Phi$ instead of $\mathcal{S} \cup \mathcal{E} \cup \mathcal{X} \models_{\leadsto} \Phi$.)

Let $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ be an action theory and Φ a K_n-formula. $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle^-_{\Phi}$ is the action theory resulting from the contraction of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ by Φ.

Contracting a theory by a static law φ amounts to using any existing contraction operator for classical logic. Let $\ominus$ be such an operator. Moreover, based on (Herzig & Varzinczak 2005b), we also need to guarantee that φ does not follow from $\mathcal{E}$, $\mathcal{X}$ and $\leadsto$. We define contraction of a domain description by a static law as follows:

Definition 14 $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle^-_{\varphi} = \langle \mathcal{S}^-, \mathcal{E}, \mathcal{X}^-, \leadsto \rangle$, where $\mathcal{S}^- = \mathcal{S} \ominus \varphi$ and $\mathcal{X}^- = \{(\varphi_i \wedge \varphi) \rightarrow \langle a \rangle \top \; : \; \varphi_i \rightarrow \langle a \rangle \top \in \mathcal{X}\}$.

We now consider the case of contracting an action theory by an executability law $\varphi \rightarrow \langle a \rangle \top$. For every executability in $\mathcal{X}$, we ensure that action a is executable only in contexts where $\neg\varphi$ is the case. The following operator does the job.

Definition 15 $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle^-_{\varphi \rightarrow \langle a \rangle \top} = \langle \mathcal{S}, \mathcal{E}, \mathcal{X}^-, \leadsto \rangle$, where $\mathcal{X}^- = \{(\varphi_i \wedge \neg\varphi) \rightarrow \langle a \rangle \top \; : \; \varphi_i \rightarrow \langle a \rangle \top \in \mathcal{X}\}$.

For instance, contracting $glued \rightarrow \langle toggle \rangle \top$ in our example would give us $\mathcal{X}^- = \{\neg glued \rightarrow \langle toggle \rangle \top\}$.

Finally, to contract a theory by $\varphi \rightarrow [a]\psi$, for every effect law in $\mathcal{E}$, we first ensure that a still has effect ψ whenever φ does not hold, second we enforce that a has no effect in context $\neg\varphi$ except on those literals that are consequences of $\neg\psi$. Combining this with the new dependence relation also linking a to literals involved by $\neg\psi$, we have that a may now produce $\neg\psi$ as outcome. In other words, the effect law has been contracted. The operator below formalizes this:

Definition 16 $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle^-_{\varphi \rightarrow [a]\psi} = \langle \mathcal{S}, \mathcal{E}^-, \mathcal{X}, \leadsto^- \rangle$, with $\leadsto^- = \leadsto \cup \{\langle a, l \rangle \; : \; l \in lit(NewCons_{\mathcal{S}}(\neg\psi))\}$ and $\mathcal{E}^- = \{(\varphi_i \wedge \neg\varphi) \rightarrow [a]\psi \; : \; \varphi_i \rightarrow [a]\psi \in \mathcal{E}\} \cup \{(\neg\varphi \wedge \neg l) \rightarrow [a]\neg l \; : \; \langle a, l \rangle \in (\leadsto^- \setminus \leadsto)\}$.

For instance, contracting the law $blackout \rightarrow [toggle]light$ from our theory would give us $\mathcal{E}^- = \{(\neg up \wedge \neg blackout) \rightarrow [toggle]up, (up \wedge \neg blackout) \rightarrow [toggle]\neg up\}$.

Results

In this section we present the main results that follow from our framework. These require the action theory under consideration to be modular (Herzig & Varzinczak 2005b). In our framework, an action theory is said to be modular if a formula of a given type entailed by the whole theory can also be derived solely from its respective module (the set of formulas of the same type) together with the static laws $\mathcal{S}$. As shown in (Herzig & Varzinczak 2005b), to make a domain description satisfy such a property it is enough to guarantee that there is no classical formula entailed by the theory that is not entailed by the static laws alone.

Definition 17 $\varphi \in \mathfrak{Fml}$ is an *implicit static law* of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ if and only if $\mathcal{S}, \mathcal{E}, \mathcal{X} \models_{\leadsto} \varphi$ and $\mathcal{S} \not\models \varphi$.

A theory is *modular* if it has no implicit static laws. Our concept of modularity of theories was originally defined in (Herzig & Varzinczak 2004b; 2005b), but similar notions have also been addressed in the literature (Cholvy 1999; Amir 2000a; Zhang, Chopra, & Foo 2002; Lang, Lin, & Marquis 2003; Herzig & Varzinczak 2005a). A modularity-based approach for narrative reasoning about actions is given in (Kakas, Michael, & Miller 2005).

To witness how implicit static laws can show up, consider the quite simple action theory below, depicting the walking turkey scenario (Thielscher 1995):

$$\mathcal{S} = \{walking \to alive\},\ \mathcal{E} = \left\{ \begin{array}{c} [tease]walking, \\ loaded \to [shoot]\neg alive \end{array} \right\}$$

$$\mathcal{X} = \{\ \langle tease \rangle \top, \langle shoot \rangle \top\ \},$$

$$\leadsto = \left\{ \begin{array}{c} \langle shoot, \neg loaded \rangle, \langle shoot, \neg alive \rangle, \\ \langle shoot, \neg walking \rangle, \langle tease, walking \rangle \end{array} \right\}$$

With this domain description we have $\mathcal{S}, \mathcal{E}, \mathcal{X} \models_{\leadsto} alive$: first, $\{walking \to alive, [tease]walking\} \models_{\leadsto} [tease]alive$, second $\models_{\leadsto} \neg alive \to [tease]\neg alive$ (from the independence $tease \not\leadsto alive$), and then $\mathcal{S}, \mathcal{E} \models_{\leadsto} \neg alive \to [tease]\bot$. As long as $\mathcal{S}, \mathcal{E}, \mathcal{X} \models_{\leadsto} \langle tease \rangle \top$, we must have $\mathcal{S}, \mathcal{E}, \mathcal{X} \models_{\leadsto} alive$. As $\mathcal{S} \not\models alive$, the formula *alive* is an implicit static law of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$.

Modular theories have several advantages (Herzig & Varzinczak 2004b;). For example, consistency of a modular action theory can be checked by just checking consistency of $\mathcal{S}$: if $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ is modular, then $\mathcal{S}, \mathcal{E}, \mathcal{X} \models_{\leadsto} \bot$ if and only if $\mathcal{S} \models \bot$. Deduction of an effect of a sequence of actions $a_1; \ldots; a_n$ (prediction) does not need to take into account the effect laws for actions other than $a_1, \ldots, a_n$. This applies in particular to plan validation when deciding whether $\langle a_1; \ldots; a_n \rangle \varphi$ is the case.

Throughout this work we have used multimodal logic K_n. For an assessment of the modularity principle in the Situation Calculus, see (Herzig & Varzinczak 2005a).

Here we establish that our operators are correct w.r.t. the semantics. Our first theorem establishes that the semantical contraction of the models of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ by Φ produces models of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle^-_\Phi$.

Theorem 1 Let $\langle W, R \rangle$ be a model of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$, and let Φ be a formula that has the form of one of the three laws. For all models $\mathscr{M}$, if $\mathscr{M} \in \langle W, R \rangle^-_\Phi$, then $\mathscr{M}$ is a model of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle^-_\Phi$.

It remains to prove that the other way round, the models of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle^-_\Phi$ result from the semantical contraction of models of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ by Φ. This does not hold in general, as shown by the following example: suppose there is only one atom p and one action a, and consider the theory $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ such that $\mathcal{S} = \emptyset$, $\mathcal{E} = \{p \to [a]\bot\}$, $\mathcal{X} = \{\langle a \rangle \top\}$, and $\leadsto = \emptyset$. The only model of that action theory is $\mathscr{M} = \langle \{\{\neg p\}\}, \{(\{\neg p\}, \{\neg p\})\} \rangle$. By definition, $\mathscr{M}^-_{p \to \langle a \rangle \top} = \{\mathscr{M}\}$. On the other hand, $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle^-_{p \to \langle a \rangle \top} = \langle \emptyset, \{p \to [a]\bot\}, \{\neg p \to \langle a \rangle \top\}, \emptyset \rangle$. The contracted theory has *two* models: $\mathscr{M}$ and $\mathscr{M}' = \langle \{\{p\}, \{\neg p\}\}, (\{\neg p\}, \{\neg p\}) \rangle$. While $\neg p$ is valid in the contraction of the models of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$, it is invalid in the models of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle^-_{p \to \langle a \rangle \top}$.

Fortunately, we can establish a result for those action theories that are modular. The proof requires three lemmas. The first one says that for a modular theory we can restrict our attention to its 'big' models.

Lemma 1 Let $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ be modular. Then $\mathcal{S}, \mathcal{E}, \mathcal{X} \models_{\leadsto} \Phi$ if and only if $\models^{\langle W, R \rangle} \Phi$ for every model $\langle W, R \rangle$ of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ such that $W = val(\mathcal{S})$.

Note that the lemma does not hold for non-modular theories, as $\{\langle W, R \rangle : \langle W, R \rangle$ is a model of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ and $W = val(\mathcal{S})\}$ is empty then.

The second lemma says that modularity is preserved under contraction.

Lemma 2 Let $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ be modular, and let Φ be a formula of the form of one of the three laws. Then $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle^-_\Phi$ is modular.

The third one establishes the required link between the contraction operators and contraction of 'big' models.

Lemma 3 Let $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ be modular, and let Φ be a formula of the form of one of the three laws. If $\mathscr{M}' = \langle val(\mathcal{S}), R' \rangle$ is a model of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle^-_\Phi$, then there is a model $\mathscr{M}$ of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ such that $\mathscr{M}' \in \mathscr{M}^-_\Phi$.

Putting the three above lemmas together we get:

Theorem 2 Let $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ be modular, Φ be a formula of the form of one of the three laws, and $\langle \mathcal{S}^-, \mathcal{E}^-, \mathcal{X}^-, \leadsto^- \rangle$ be $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle^-_\Phi$. If it holds that $\mathcal{S}^-, \mathcal{E}^-, \mathcal{X}^- \models_{\leadsto^-} \Psi$, then for every model $\mathscr{M}$ of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ and every $\mathscr{M}' \in \mathscr{M}^-_\Phi$ it holds that $\models^{\mathscr{M}'} \Psi$.

Our two theorems together establish correctness of the operators:

Corollary 1 Let $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ be modular, Φ be a formula of the form of one of the three laws, and $\langle \mathcal{S}^-, \mathcal{E}^-, \mathcal{X}^-, \leadsto^- \rangle$ be $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle^-_\Phi$. Then $\mathcal{S}^-, \mathcal{E}^-, \mathcal{X}^- \models_{\leadsto^-} \Psi$ if and only if for every model $\mathscr{M}$ of $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ and every $\mathscr{M}' \in \mathscr{M}^-_\Phi$ it holds that $\models^{\mathscr{M}'} \Psi$.

We give a necessary condition for success of contraction:

Theorem 3 Let Φ be an effect or an executability law such that $\mathcal{S} \not\models_{\mathsf{K}_n} \Phi$. Let $\langle \mathcal{S}^-, \mathcal{E}^-, \mathcal{X}^-, \leadsto^- \rangle$ be $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle^-_\Phi$. If $\langle \mathcal{S}, \mathcal{E}, \mathcal{X}, \leadsto \rangle$ is modular, then $\mathcal{S}^-, \mathcal{E}^-, \mathcal{X}^- \not\models_{\leadsto^-} \Phi$.

Contracting implicit static laws

There can be many reasons why a theory should be changed. Following (Herzig & Varzinczak 2004b; 2005b;), here we focus on the case where it has some classical consequence φ the designer is not aware of.

If φ is taken as intuitive, then, normally, no change has to be done at all, unless we want to keep abide on the modularity principle and thus make φ explicit by adding it to $\mathcal{S}$. In the scenario example of last section, if the knowledge engineer's universe has immortal turkeys, then she would add the static law *alive* to $\mathcal{S}$.

The other way round, if φ is not intuitive, as long as φ is entailed by $\langle\mathcal{S},\mathcal{E},\mathcal{X},\leadsto\rangle$, the goal is to avoid such an entailment, i.e., what we want is $\mathcal{S}^-,\mathcal{E}^-,\mathcal{X}^- \not\models_{\leadsto^-} \varphi$, where $\langle\mathcal{S}^-,\mathcal{E}^-,\mathcal{X}^-,\leadsto^-\rangle$ is $\langle\mathcal{S},\mathcal{E},\mathcal{X},\leadsto\rangle^-_\Phi$. In the mentioned scenario, the knowledge engineer considers that having immortal turkeys is not reasonable and thus decides to change the domain description to $\langle\mathcal{S}^-,\mathcal{E}^-,\mathcal{X}^-,\leadsto^-\rangle$ so that $\mathcal{S}^-,\mathcal{E}^-,\mathcal{X}^- \not\models_{\leadsto^-}$ *alive*.

This means that action theories that are not modular need to be changed, too. Such a changing process is driven by the problematic part of the theory detected by the algorithms defined in (Herzig & Varzinczak 2004b) and improved in (Herzig & Varzinczak).

The algorithm works as follows: for each executability law $\varphi \to \langle a\rangle\top$ in the theory, construct from $\mathcal{E}$ and $\leadsto$ a set of inexecutabilities $\{\varphi_1 \to [a]\bot,\ldots,\varphi_n \to [a]\bot\}$ that potentially conflict with $\varphi \to \langle a\rangle\top$. For each i, $1 \leq i \leq n$, if $\varphi\wedge\varphi_i$ is satisfiable w.r.t. $\mathcal{S}$, mark $\neg(\varphi\wedge\varphi_i)$ as an implicit static law. Incrementally repeat this procedure (adding all the $\neg(\varphi\wedge\varphi_i)$ that were caught to $\mathcal{S}$) until no implicit static law is obtained.

For an example of the execution of the algorithm, consider $\langle\mathcal{S},\mathcal{E},\mathcal{X},\leadsto\rangle$ as above. For the action *tease*, we have the executability $\langle tease\rangle\top$. Now, from $\mathcal{E}$, and $\leadsto$ we try to build an inexecutability for *tease*. We take $[tease]walking$ and compute then all indirect effects of *tease* w.r.t. $\mathcal{S}$. From $walking \to alive$, we get that *alive* is an indirect effect of *tease*, giving us $[tease]alive$. But $\langle tease, alive\rangle \notin\leadsto$, which means the frame axiom $\neg alive \to [tease]\neg alive$ holds. Together with $[tease]alive$, this gives us the inexecutability $\neg alive \to [tease]\bot$. As $\mathcal{S}\cup\{\top,\neg alive\}$ is satisfiable ($\top$ is the antecedent of the executability $\langle tease\rangle\top$), we get the implicit static law *alive*. For this example no other inexecutability for *tease* can be derived, so the computation stops.

It seems that in general implicit static laws are not intuitive. Therefore their contraction is more likely to happen than their addition.[3] In the example above, the action theory has to be contracted by *alive*.[4] In order to contract the action theory, the designer has several choices:

[3]In all the examples in which we have found implicit static laws that are intuitive they are so evident that the only explanation for not having them explicitly stated is that they have been forgotten by the theory's designer.

[4]Here the change operation is a revision-based operation rather than an update-based operation since we mainly "fix" the theory.

1) Contract the set $\mathcal{S}$. (In this case, such an operation is not enough, since *alive* is a consequence of the rest of the theory.)

2) Weaken the effect law $[tease]walking$ to $alive \to [tease]walking$, since the original effect law is too strong. This means that in a first stage the designer has to contract the theory and in a second one expand the effect laws with the weaker law. The designer will usually choose this option if she focuses on the preconditions of the effects of actions.

3) Weaken the executability law $\langle tease\rangle\top$ by rephrasing it as $alive \to \langle tease\rangle\top$: first the executability is contracted and then the weaker one is added to the resulting set of executability laws. The designer will choose this option if she focuses on preconditions for action execution.

The analysis of this example shows that the choice of what change has to be carried out is up to the knowledge engineer. Such a task can get more complicated when ramifications are involved. To witness, suppose our scenario has been formalized as follows: $\mathcal{S} = \{walking \to alive\}$, $\mathcal{E} = \{[shoot]\neg alive\}$, $\mathcal{X} = \{\langle shoot\rangle\top\}$, and $\leadsto = \{\langle shoot, \neg alive\rangle\}$. From this action theory we can derive the inexecutability $walking \to [shoot]\bot$ and thus the implicit static law $\neg walking$. In this case we have to change the theory by contracting the frame axiom $walking \to [shoot]walking$ (which amounts to adding the missing indirect dependence $shoot \leadsto \neg walking$).

Elaboration tolerance

The principle of elaboration tolerance has been proposed by McCarthy (McCarthy 1988). Roughly, it states that the effort required to add new information to a given representation (new laws or entities) should be proportional to the complexity of the information being added, i.e., it should not require the complete reconstruction of the old theory (Shanahan 1997).

Since then many formalisms in the reasoning about actions field claim, in a more or less tacit way, to satisfy such a principle. However, for all this time there has been a lack of good formal criteria allowing for the evaluation of theory change difficulty and, consequently, comparisons between different frameworks are carried out in a subjective way.

The proposal by Amir (Amir 2000b) made the first steps in formally answering what difficulty of changing a theory means by formalizing one aspect of elaboration tolerance. The basic idea is as follows: let $\mathcal{T}_0$ be the original theory and let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two equivalent (and different) theories such that each one results from $\mathcal{T}_0$ by the application of some sequence of operations (additions and/or deletions of formulas). The resulting theory whose transformation from $\mathcal{T}_0$ has the shortest length (number of operations) is taken as the most elaboration tolerant.

Nevertheless, in the referred work only addition/deletion of *axioms* is considered, i.e., changes in the logical language or contraction of consequences of the theory not explicitly stated in the original set of axioms are not taken into account. This means that even the formal setting given in (Amir 2000b) is not enough to evaluate the complexity of theory change in a broad sense. Hence the community still needs formal criteria that allow for the comparison between more complex changes carried out by frameworks like ours, for example.

Of course, how elaboration tolerant a given update/revision method is strongly depends on its underlying formalism for reasoning about actions, i.e., its logical background, the solution to the frame problem it implements, the hypothesis it relies on, etc. In what follows we discuss how the dependence-based approach here used behaves when expansion is considered. Most of the comments concerning consequences of expansion can also be stated for contraction. We do that with respect to some of the qualitative criteria given in (McCarthy 1998). In all that follows we suppose that the resulting theory is consistent.

Adding effect laws In the dependence-based framework, adding the new effect law $\varphi \rightarrow [a]\psi$ to the theory demands a change in the dependence module $\leadsto$. In that case, the maximum number of statements added to $\leadsto$ is $|\{l : l \in lit(NewCons_{\mathcal{S}}(\psi))\}|$ (dependences for all indirect effects have to be stated, too). This is due to the explanation closure nature of the reasoning behind dependence (for more details, see (Castilho, Gasquet, & Herzig 1999)). Because of this, according to Shanahan (Shanahan 1997), explanation closure approaches are not elaboration tolerant when dealing with the ramification problem. In order to achieve that, the framework should have a mechanism behaving like circumscription that automatically deals with ramifications. This raises the question: "if we had an automatic (or even semi-automatic) procedure to do the job of generating the indirect dependences, could we say the framework is elaboration tolerant?". We think we can answer positively to such a question, and, supported by Reiter (Reiter 2001), we are working on a semi-automatic procedure for generating the dependence relation from a set of effect laws.

Adding executability laws Such a task demands only a change in the set $\mathcal{X}$ of executabilities, possibly introducing implicit static laws as a side effect.

Adding static laws Besides expanding the set $\mathcal{S}$, adding new (indirect) dependences may be required (see above).

Adding frame axioms If the frame axiom $\neg l \rightarrow [a]\neg l$ has to be valid in the resulting theory, expunging the dependence $a \leadsto l$ should do the job.

Adding a new action name Without loss of generality we can assume the action in question was already in the language. In that case, we expect just to add effect or executability laws for it. For the former, at most $|\mathfrak{Lit}|$ dependences will be added to $\leadsto$. (We point out nevertheless that the requirement made in (McCarthy 1998) that the addition of an action irrelevant for a given plan in the old theory should not preclude it in the resulting theory is too strong. Indeed, it is not difficult to imagine a new action forcing an implicit static law from which an inexecutability for some action in the plan can be derived. The same holds for the item below.)

Adding a new fluent name In the same way, we can suppose the fluent was already in the language. Such a task amounts thus to one or more of the above expansions. There will be at most $2 \times |\mathfrak{Act}|$ new elements added to $\leadsto$.

Related work

Following (Li & Pereira 1996; Liberatore 2000), Eiter *et al.* (Eiter *et al.* 2005) have investigated update of action domain descriptions. They define a version of action theory update in an action language and give complexity results showing how hard such a task can be.

Update of action descriptions in their sense is always relative to some conditions (interpreted as knowledge possibly obtained from earlier observations and that should be kept). This characterizes a constraint-based update. In the example they give, change must be carried out preserving the assumption that pushing the button of the remote control is always executable. Actually, the method is more subtle, as new effect laws are added constrained by *the addition* of viz. an executability law for the new action under concern. In the example, the constraint (executability of push) was not in the original action description and must figure in the updated theory.

They describe domains of actions in a fragment of the action language $\mathcal{C}$ (Gelfond & Lifschitz 1998). However they do not specify which fragment, so it is not clear whether the claimed advantages $\mathcal{C}$ has over $\mathcal{A}$ really transfer to their framework. At one hand, their approach deals with indirect effects, but they do not talk about updating a theory by a law with a nondeterministic action. Anyway, except for concurrency, their account can be translated into ours, as shown in (Castilho, Gasquet, & Herzig 1999).

Eiter *et al.* consider an action theory $\mathcal{T}$ as comprising two main components: $\mathcal{T}_u$, the part of the theory that must remain unchanged, and $\mathcal{T}_m$, the part concerning the statements that are allowed to change. The crucial information to the associated solution to the frame problem is always in $\mathcal{T}_u$.

Given an action theory $\mathcal{T} = \mathcal{T}_u \cup \mathcal{T}_m$, $((\mathcal{T}_u \cup \mathcal{T}_m), \mathcal{T}', C)$ is the problem of updating $\mathcal{T}$ by $\mathcal{T}' \subseteq \mathcal{S} \cup \mathcal{E}$ warranting the result satisfies all constraints in $C \subseteq \mathcal{S} \cup \mathcal{X}$.

Even though they do not explicitly state postulates for their kind of theory update, they establish conditions for the update operator to be successful. Basically, they claim for consistency of the resulting theory; maintenance of the new knowledge and the invariable part of the description; satisfaction of the constraints in C; and minimal change.

In some examples that they develop, the illustrated "partial solution" does not satisfy C due to the existence of implicit laws (cf. Example 1, where there is an implicit inexecutability law). To achieve a solution, while keeping C, some other laws must be dropped (in the example, the agent gives up a static law).[5]

Just to see the link between update by subsumed laws and addition of implicit static laws, we note that Proposition 1 in the referred work is the same as Theorem 14 in (Herzig & Varzinczak 2005b): every implicit static law in Herzig and Varzinczak's sense is trivially a subsumed law in Eiter *et al.*'s sense.

[5]This does not mean however that the updated theory will necessarily contain no implicit law.

With their method we can also contract by a static and an effect law. Contraction of executabilities are not explicitly addressed, and weakening (replacing a law by a weaker one) is left as future work.

A main difference between the approach in (Eiter *et al.* 2005) and ours is that we do not need to add new fluents at every elaboration stage: we still work on the same set of fluents, refining their behavior w.r.t. an action a. In Eiter *et al.*'s proposal an update forces changing all the variable rules appearing in the action theory by adding to each one a new update fluent. This is a constraint when elaborating action theories.

Concluding remarks

In this work we have presented a general method for changing a domain description (alias action theory) given any formula we want to contract.

We have defined a semantics for theory contraction and also presented its syntactical counterpart through contraction operators. Soundness and completeness of such operators with respect to the semantics have been established (Corollary 1).

We have also shown that modularity is a necessary condition for a contraction to be successful (Theorem 3). This gives further evidence that our modularity notion is fruitful.

We have analysed an example of contraction of a non-modular theory by an implicit static law that is unintended.

Because of forcing formulas to be explicitly stated in their respective modules (and thus possibly making them inferable in independently different ways), intuitively modularity could be seen to diminish elaboration tolerance. For instance, when contracting a classical formula φ from a non-modular theory, it seems reasonable to expect not to change the set of static laws $\mathcal{S}$, while the theory being modular surely forces changing such a module. However it is not difficult to conceive non-modular theories in which contraction of a formula φ may demand a change in $\mathcal{S}$ as well. To witness, suppose $\mathcal{S} = \{\varphi_1 \to \varphi_2\}$ in an action theory from whose dynamic part we (implicitly) infer $\neg\varphi_2$. In this case, a contraction of $\neg\varphi_1$ keeping $\neg\varphi_2$ would necessarily ask for a change in $\mathcal{S}$. We point out nevertheless that in both cases (modular and non-modular) the extra work in changing other modules stays in the mechanical level, i.e., in the machinery that carries out the modification, and does not augment in a significant way the amount of work the knowledge engineer is expected to do.

What is the status of the AGM-postulates for contraction in our framework? First, contraction of static laws satisfies all the postulates, as soon as the underlying classical contraction operation $\ominus$ satisfies all of them.

In the general case, however, our constructions do not satisfy the central postulate of preservation $\langle\mathcal{S},\mathcal{E},\mathcal{X},\leadsto\rangle^-_\Phi = \langle\mathcal{S},\mathcal{E},\mathcal{X},\leadsto\rangle$ if $\mathcal{S},\mathcal{E},\mathcal{X} \not\models_{\leadsto} \Phi$. Indeed, suppose we have a language with only one atom p, and a model $\mathscr{M}$ with two worlds $w = \{p\}$ and $w' = \{\neg p\}$ such that wR_aw', $w'R_aw$, and $w'R_aw'$. Then $\models^{\mathscr{M}} p \to [a]\neg p$ and $\not\models^{\mathscr{M}} [a]\neg p$, i.e., $\mathscr{M}$ is a model of the effect law $p \to [a]\neg p$, but not of $[a]\neg p$. Now the contraction $\mathscr{M}^-_{[a]\neg p}$ yields the model $\mathscr{M}'$ such that $R_a = W \times W$. Then $\not\models^{\mathscr{M}'} p \to [a]\neg p$, i.e., the effect law $p \to [a]\neg p$ is not preserved. Our contraction operation thus behaves rather like an update operation.

Now let us focus on the other postulates. Since our operator has a behavior which is close to the update postulate, we focus on the following basic erasure postulates introduced in (Katsuno & Mendelzon 1991). Let $Cn(\mathcal{T})$ be the set of all logical consequences of a theory $\mathcal{T}$.

KM1 $Cn(\langle\mathcal{S},\mathcal{E},\mathcal{X},\leadsto\rangle^-_\Phi) \subseteq Cn(\langle\mathcal{S},\mathcal{E},\mathcal{X},\leadsto\rangle)$

Postulate **KM1** does not always hold because it is possible to make the formula $\varphi \to [a]\bot$ valid in the resulting theory by removing elements of R_a (cf. Definition 10).

KM2 $\Phi \notin Cn(\langle\mathcal{S},\mathcal{E},\mathcal{X},\leadsto\rangle^-_\Phi)$

Under the condition that $\langle\mathcal{S},\mathcal{E},\mathcal{X},\leadsto\rangle$ is modular, Postulate **KM2** is satisfied (cf. Theorem 3).

KM3 If $Cn(\langle\mathcal{S}_1,\mathcal{E}_1,\mathcal{X}_1,\leadsto\rangle) = Cn(\langle\mathcal{S}_2,\mathcal{E}_2,\mathcal{X}_2,\leadsto\rangle)$ and $\models_{\mathsf{K}_n} \Phi_1 \leftrightarrow \Phi_2$, then $Cn(\langle\mathcal{S}_1,\mathcal{E}_1,\mathcal{X}_1,\leadsto\rangle^-_{\Phi_2}) = Cn(\langle\mathcal{S}_2,\mathcal{E}_2,\mathcal{X}_2,\leadsto\rangle^-_{\Phi_1})$.

Theorem 4 If $\langle\mathcal{S}_1,\mathcal{E}_1,\mathcal{X}_1,\leadsto\rangle$ and $\langle\mathcal{S}_2,\mathcal{E}_2,\mathcal{X}_2,\leadsto\rangle$ are modular and the propositional contraction operator $\ominus$ satisfies Postulate **KM3**, then Postulate **KM3** is satisfied for every $\Phi_1, \Phi_2 \in \mathfrak{Fml}$.

Here we have presented the case for contraction, but our definitions can be extended to revision, too. Our results can also be generalized to the case where learning new actions or fluents is involved. This means in general that more than one simple formula should be added to the belief base and must fit together with the rest of the theory with as little side-effects as possible. We are currently defining algorithms based on our operators to achieve that.

Acknowledgments

We are grateful to the anonymous referees for many useful comments on an earlier version of this paper.

References

Amir, E. 2000a. (De)composition of situation calculus theories. In *Proc. AAAI'2000*, 456–463. AAAI Press/MIT Press.

Amir, E. 2000b. Toward a formalization of elaboration tolerance: Adding and deleting axioms. In *Frontiers of Belief Revision*. Kluwer.

Castilho, M. A.; Gasquet, O.; and Herzig, A. 1999. Formalizing action and change in modal logic I: the frame problem. *J. of Logic and Computation* 9(5):701–735.

Castilho, M. A.; Herzig, A.; and Varzinczak, I. J. 2002. It depends on the context! a decidable logic of actions and plans based on a ternary dependence relation. In *NMR'02*, 343–348.

Cholvy, L. 1999. Checking regulation consistency by using SOL-resolution. In *Proc. Int. Conf. on AI and Law*, 73–79.

Demolombe, R.; Herzig, A.; and Varzinczak, I. 2003. Regression in modal logic. *J. of Applied Non-Classical Logics (JANCL)* 13(2):165–185.

Doherty, P.; Łukaszewicz, W.; and Madalinska-Bugaj, E. 1998. The PMA and relativizing change for action update. In *Proc. KR'98*, 258–269. Morgan Kaufmann.

Eiter, T.; Erdem, E.; Fink, M.; and Senko, J. 2005. Updating action domain descriptions. In *Proc. IJCAI'05*, 418–423. Morgan Kaufmann.

Foo, N. Y., and Zhang, D. 2002. Dealing with the ramification problem in the extended propositional dynamic logic. In *Advances in Modal Logic*, volume 3. World Scientific. 173–191.

Forbus, K. D. 1989. Introducing actions into qualitative simulation. In *Proc. IJCAI'89*, 1273–1278. Morgan Kaufmann.

Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press.

Gelfond, M., and Lifschitz, V. 1998. Action languages. *ETAI* 2(3–4):193–210.

Hansson, S. O. 1999. *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Kluwer.

Harel, D. 1984. Dynamic logic. In *Handbook of Philosophical Logic*, volume II. D. Reidel, Dordrecht. 497–604.

Herzig, A., and Rifi, O. 1999. Propositional belief base update and minimal change. *Artificial Intelligence* 115(1):107–138.

Herzig, A., and Varzinczak, I. Metatheory of actions: beyond consistency. To appear.

Herzig, A., and Varzinczak, I. 2004a. An assessment of actions with indeterminate and indirect effects in some causal approaches. Technical Report 2004–08–R, Institut de recherche en informatique de Toulouse (IRIT), Université Paul Sabatier.

Herzig, A., and Varzinczak, I. 2004b. Domain descriptions should be modular. In *Proc. ECAI'04*, 348–352. IOS Press.

Herzig, A., and Varzinczak, I. 2005a. Cohesion, coupling and the meta-theory of actions. In *Proc. IJCAI'05*, 442–447. Morgan Kaufmann.

Herzig, A., and Varzinczak, I. 2005b. On the modularity of theories. In *Advances in Modal Logic*, volume 5. King's College Publications. 93–109.

Inoue, K. 1992. Linear resolution for consequence finding. *Artificial Intelligence* 56(2–3):301–353.

Jin, Y., and Thielscher, M. 2005. Iterated belief revision, revised. In *Proc. IJCAI'05*, 478–483. Morgan Kaufmann.

Kakas, A.; Michael, L.; and Miller, R. 2005. $\mathcal{M}$odular-$\mathcal{E}$: an elaboration tolerant approach to the ramification and qualification problems. In *Proc. 8th Intl. Conf. Logic Programming and Nonmonotonic Reasoning*, 211–226. Springer-Verlag.

Katsuno, H., and Mendelzon, A. O. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52(3):263–294.

Katsuno, H., and Mendelzon, A. O. 1992. On the difference between updating a knowledge base and revising it. In Gärdenfors, P., ed., *Belief revision*. Cambridge University Press. 183–203.

Lang, J.; Lin, F.; and Marquis, P. 2003. Causal theories of action – a computational core. In *Proc. IJCAI'03*, 1073–1078. Morgan Kaufmann.

Li, R., and Pereira, L. 1996. What is believed is what is explained. In *Proc. AAAI'96*, 550–555. AAAI Press/MIT Press.

Liberatore, P. 2000. A framework for belief update. In *Proc. JELIA'2000*, 361–375.

McCarthy, J. 1988. *Mathematical logic in artificial intelligence*. Daedalus.

McCarthy, J. 1998. Elaboration tolerance. In *Proc. Common Sense'98*.

Popkorn, S. 1994. *First Steps in Modal Logic*. Cambridge University Press.

Reiter, R. 1991. 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. 359–380.

Reiter, R. 2001. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. Cambridge, MA: MIT Press.

Shanahan, M. 1997. *Solving the frame problem: a mathematical investigation of the common sense law of inertia*. Cambridge, MA: MIT Press.

Shapiro, S.; Pagnucco, M.; Lespérance, Y.; and Levesque, H. J. 2000. Iterated belief change in the situation calculus. In *Proc. KR'2000*, 527–538. Morgan Kaufmann.

Thielscher, M. 1995. Computing ramifications by post-processing. In *Proc. IJCAI'95*, 1994–2000. Morgan Kaufmann.

Winslett, M.-A. 1988. Reasoning about action using a possible models approach. In *Proc. AAAI'88*, 89–93. Morgan Kaufmann.

Winslett, M.-A. 1995. Updating logical databases. In *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 4. Oxford University Press. 133–174.

Zhang, D., and Foo, N. Y. 2001. EPDL: A logic for causal reasoning. In *Proc. IJCAI'01*, 131–138. Morgan Kaufmann.

Zhang, D.; Chopra, S.; and Foo, N. Y. 2002. Consistency of action descriptions. In *PRICAI'02, Topics in Artificial Intelligence*. Springer-Verlag.