

Adding Modal Operators to the Action Language \mathcal{A}

Aaron Hunter

Simon Fraser University
Burnaby, B.C.
Canada V5A 1S6
amhunter@cs.sfu.ca

Abstract

The action language \mathcal{A} is a simple high-level language for describing transition systems. In this paper, we extend the action language \mathcal{A} by allowing a unary modal operator in the underlying propositional logic. The extended language requires very little new machinery, and it is suitable for describing transitions between Kripke structures. We consider some formal restrictions on action descriptions that preserve natural classes of Kripke structures, and we prove that the modal epistemic extension of \mathcal{A} naturally subsumes related approaches to reasoning about knowledge. We conclude with some plans for future work.

Introduction

The action language \mathcal{A} is a simple high-level language for reasoning about the effects of actions (Gelfond & Lifschitz 1993). The basic language is suitable only for simple action domains, but it has been extended several times to address a wide range of problems (Baral & Gelfond 1997; Baral, Gelfond, & Proveti 1997). In this paper, we suggest that it is possible to increase the representational power of \mathcal{A} without changing the action language itself. Instead, we look at extending the underlying propositional logic by adding modal operators. We consider the expressive power of the modal extension, and compare the framework with related work on epistemic extensions of \mathcal{A} .

Preliminaries

Notation and Conventions

We employ a standard set $\{\neg, \wedge, \vee\}$ of propositional connectives. Given a propositional signature F , a literal is an element of F or an element of F prefixed with the negation symbol. The complement of a literal f will be denoted by \bar{f} .

We assume the reader is familiar with propositional modal logic, as outlined in (Chellas 1980). To fix notation and terminology, we reiterate a few important definitions. We restrict attention to modal logics with a single unary modal operator. A Kripke structure is a triple $\mathcal{M} = \langle M, R, \pi \rangle$, where M is a non-empty set, R is a binary accessibility relation on M and π associates a subset of M with every atomic formula. The satisfaction relation $\mathcal{M}, m \models \phi$ is given by the standard recursive definition; we omit the mention of \mathcal{M} if it is clear from the context.

Let \mathcal{L} be a modal logic given by a set of axiom schemata, and let Π be a set of Kripke structures. We say that \mathcal{L} is *determined by* Π if the set of theorems of \mathcal{L} is identical to the set of formulas valid in Π . In practice, we will refer to a modal logic either by a set of axioms or by a set of Kripke structures, depending on which presentation is more convenient for the task at hand.

Many important modal logics are determined by placing natural restrictions on the accessibility relation. We mention three examples: KT , KB , and $K4$. These modal logics are determined by the classes of structures in which the accessibility relation is reflexive, symmetric, and transitive, respectively. If we combine all three restrictions then the accessibility relation must be an equivalence relation, and we have the modal logic $S5$.

The logic $S5$ is the standard modal epistemic logic. The intuition is that two worlds are related by the accessibility relation if they are indistinguishable to the underlying agent. Typically, in epistemic logic, we use the symbol K to denote the modal operator. In the general case, there may be several knowledge operators corresponding to different agents. For our purposes, it will be sufficient to restrict attention to the single-agent case. For a detailed discussion of epistemic logic, we refer the reader to (Fagin *et al.* 1995).

Action Language \mathcal{A}

We briefly review the syntax and semantics of the action language \mathcal{A} , as introduced in (Gelfond & Lifschitz 1993).

Let \mathbf{F} denote a fixed set of fluent symbols and let \mathbf{A} denote a fixed set of action symbols. We think of \mathbf{F} as a propositional signature, and it is understood that formulas and literals in \mathcal{A} are taken over the set \mathbf{F} . The syntax of \mathcal{A} is given by specifying a class of propositions. We restrict attention to the so-called effect propositions.

Definition 1 *An effect proposition of the language \mathcal{A} is an expression of the form*

$$A \text{ causes } L \text{ if } F$$

where $A \in \mathbf{A}$, L is a literal, and F is a conjunction of literals. A set of effect propositions is called an *action description*.

In order to give meaning to action descriptions, we need to associate a transition relation Φ_{AD} with every action de-

scription AD . The following definition describes how this is done.

Definition 2 Let AD be an action description, let s be an interpretation of \mathbf{F} and let A be an action symbol. Then $\Phi_{AD}(s, A, s')$ if

$$E(A, s) \subseteq s' \subseteq E(A, s) \cup s$$

where $E(A, s)$ is the set of literals L such that $(A \text{ causes } L \text{ if } F) \in AD$ and every literal in F holds in s .

Intuitively, the transition relation maps a pair (s, A) to a new interpretation s' that is exactly like s except that the fluents affected by A have changed values.

A great deal of work has been done on planning with action language \mathcal{A} and its extensions (Lifschitz 1999). The standard approach is to translate \mathcal{A} into logic programming in a manner that allows plans to be identified with answer sets.

Motivating Example: Adding Knowledge

Many actions do not change the state of the world, but they do change the knowledge of an agent. For example, the card game Clue can be formalized as a game in which the state of the world never changes and every action simply alters the knowledge of the players (van Ditmarsch 2002). As a more practical example, cryptographic protocols can be seen as a sequence of knowledge producing actions (Burrows, Abadi, & Needham 1989). In order to formalize action domains of this nature, we need to extend \mathcal{A} with some formal representation of knowledge.

As far as we know, there have been no previously published modal epistemic extensions of \mathcal{A} . However, there have been extensions that address knowledge by introducing new propositions for representing the effects of sensing actions (Lobo, Mendez, & Taylor 2001; Son & Baral 2001). Basically these approaches focus on modelling dynamic knowledge about atomic facts. Lobo et. al. acknowledge that there are some situations in which a modal approach would be advantageous. For example, they suggest that a modal approach may provide a more natural framework for modelling situations in which introspective agents need to perform frequent checks on the current knowledge state.

We suggest that adding a modal knowledge operator to \mathcal{A} has some practical advantages over alternative approaches. In particular, by adding a modal operator, we obtain an action language that is immediately familiar and comprehensible to those with an elementary knowledge of modal logic. Moreover, using a modal knowledge operator is a natural way to represent nested knowledge in a multi-agent environment. Representing nested knowledge is essential for cryptographic protocol verification.

We remark that we do not attempt to give a complete account of knowledge change. Instead, we give a simple treatment of knowledge change at the normal level of abstraction for \mathcal{A} descriptions. We hope that the resulting language will be useful for high-level evaluation of related formalisms, similar to the evaluation demonstrated in (Kahramanoğullari & Thielscher 2003).

Overview

The main contribution of this paper is the introduction of a simple extension of \mathcal{A} that allows a unary modal operator in the underlying propositional logic. The basic framework can be seen as a first step towards a complete treatment of modal change in \mathcal{A} ; it provides a simple foundation that may be extended to reason about a wide range of modal action effects. Although the present system is very simple and relatively limited in the transition relations that it can represent, the epistemic variant is sufficiently expressive to describe planning domains with sensing actions and incomplete information.

In the next section, we give the basic syntax and semantics for the modal extension of the action language \mathcal{A} , and we look at a motivating example. In the following section, we look at the modal epistemic extension of \mathcal{A} and we demonstrate that it subsumes two existing approaches to the representation of epistemic action effects in \mathcal{A} . We conclude with some directions future work.

Adding Modal Operators

Syntax

Let \mathbf{A} be a fixed set of action symbols, let \mathbf{F} be a fixed set of fluent symbols, and let \mathcal{L} be a fixed modal logic with a single unary modal operator \square .

We want to extend \mathcal{A} minimally to allow modal action effects. We remark that action effects in \mathcal{A} are always literals; this restriction allows us to avoid dealing with disjunctive effects in the semantics. We will assure that disjunctive modal effects are also prevented.

Definition 3 A proposition of the action language $\mathcal{A}[\mathcal{L}]$ is an expression of the form

$$A \text{ causes } \phi \text{ if } PRE$$

where $A \in \mathbf{A}$, PRE is a formula, and ϕ is either a literal or a formula of the form $\square\psi$ for some ψ .

Notice that ψ need not be a literal; any formula can be the modal effect of an action. As in the standard case, an *action description* is a set of propositions.

Let \mathcal{L}_K denote the logic $S5$ with unary modal operator K ; we think of K as a modal knowledge operator. In this context, it is natural to think of propositions of the form

$$A \text{ causes } K\phi \text{ if } PRE$$

as descriptions of sensing action effects.

Semantics

As in \mathcal{A} , the semantics of $\mathcal{A}[\mathcal{L}]$ is given by associating a transition relation with each set of propositions. However, to deal with modal action effects, we need transition relations between Kripke structures. To facilitate the discussion, we adopt a functional notation. Hence, with each action description AD , we define a transition function Φ_{AD} . With this notation, $\Phi_{AD}(\mathcal{M}, A)$ denotes the Kripke structure that results when action A is executed in the structure \mathcal{M} .

The following definition associates a transition function Φ_{AD} with an action description AD .

Definition 4 Let AD be an action description in $\mathcal{A}[\mathcal{L}]$. Let $\mathcal{M} = \langle M, R, \pi \rangle$ be a Kripke structure for \mathcal{L} and let $A \in \mathcal{A}$. The Kripke structure $\Phi_{AD}(\mathcal{M}, A) = \langle M^*, R^*, \pi^* \rangle$ is defined as follows.

1. $M^* = M$
2. If $f \in \mathbf{F}$ and $m \in M$, then $m \in \pi^*(F)$ iff one of the following conditions holds:
 - $m \in \pi(f)$ and neither f nor \bar{f} occurs as a non-modal effect of a proposition in AD
 - f occurs as a non-modal effect of a proposition P in AD , and the precondition of P is satisfied by m in \mathcal{M} .
3. $R^*(m_1, m_2)$ holds iff the following both hold
 - $(m_1, m_2) \in R$
 - there is no rule in AD of the form

$A \text{ causes } \Box\phi \text{ if } PRE$

where $\mathcal{M}, m_1 \models PRE$ and $\mathcal{M}, m_2 \models \neg\phi$

Definition 4 is more intuitive if it is considered procedurally. Given a structure \mathcal{M} and an action A , we construct $\Phi_{AD}(\mathcal{M}, A)$ as follows:

1. The fluent values of each world are updated exactly as they would be in \mathcal{A} , looking only at the propositions with non-modal effects.
2. For each modal effect $\Box\phi$ and each world m , remove all edges Rm_1m_2 where $\mathcal{M}, m_1 \models PRE$ and $\mathcal{M}, m_2 \models \neg\phi$.

It is easy to see that this procedure gives the correct result.

Note that preconditions are always evaluated in the initial Kripke structure, rather than the successor structure. The accessibility relation is changed in the successor structure to guarantee that action effects will be true in all appropriate states, although it is possible that the preconditions will no longer be true. This is the natural extension of \mathcal{A} effects to a modal setting.

We remark that, in the transition between structures, edges are never added. Hence, in the case of knowledge, we can think of modal action effects as *refinements* to the agent's knowledge. Certainly it would be interesting to consider actions that reduce an agent's knowledge as well; such actions could be represented by action effects of the form $\neg\Box\phi$. We leave this problem for future work, and restrict our attention to simple refinements for the present paper.

The following example illustrates how to apply the basic definitions.

Example Consider the language $\mathcal{A}[\mathcal{L}_K]$. We represent a situation with a single agent inside a room with a window. Looking out the window allows the agent to determine if it is raining or not. Let AD denote the action description containing the following propositions:

LookOutWindow causes $K(Rain)$ if $Rain$

LookOutWindow causes $K(\neg Rain)$ if $\neg Rain$.

Informally, the first proposition says that looking out the window causes the agent to know it is raining, provided that

it is in fact raining. The second proposition makes the parallel assertion for non-raining worlds.

Suppose that $\mathcal{M} = \langle M, R, \pi \rangle$ is a structure where the accessibility relation is universal. We construct $\Phi_{AD}(\mathcal{M}, \text{LookOutWindow})$.

According to Definition 4, the set of worlds M remains unchanged. Moreover, since there are no non-modal action effects, the interpretation function π also remains unchanged. Hence, all that changes is the accessibility relation. Let $m \in M$ and suppose $m \in \pi(Rain)$. Due to the first proposition in AD , we need to remove all edges from m to worlds where it is not raining. So, in $\Phi_{AD}(\mathcal{M}, \text{LookOutWindow})$, the world m will be related to a world m' if and only if $m' \in \pi(Rain)$. Similarly, by the second proposition in AD , we remove all edges from non-raining worlds to raining worlds. The resulting accessibility relation is an equivalence relation with two equivalence classes that partition the worlds based on the value of $Rain$. This result is consistent with the intuitive interpretation of the accessibility relation as an indistinguishability relation; after looking out the window, the agent is able to distinguish raining worlds from non-raining worlds.

The preceding example highlights an interesting issue. In particular, one might observe that the effects of both propositions in AD are obtained by adding a \Box to the preconditions. By contrast, one might be interested in the interpretation of an action description AD' containing the single proposition

LookOutWindow causes $K(Rain)$.

This proposition asserts that looking out the window causes the agent to know it is raining, whether or not it is actually raining. Suppose that \mathcal{M} is an $S5$ structure containing a world m where it is not raining. Let $\mathcal{M}' = \Phi_{AD'}(\mathcal{M}, \text{LookOutWindow})$. Applying Definition 4, it is clear that m is a world in \mathcal{M}' but m is not related to itself in \mathcal{M}' . Informally, the edge (m, m) is removed in the transition between structures. Therefore \mathcal{M}' is not an $S5$ structure, because the accessibility relation is not reflexive.

Clearly this is a problem. The transition function between Kripke structures is intended to describe how the knowledge of an agent changes. Presumably, however, the fundamental nature of knowledge should not be changed by action execution. Does this mean that action descriptions like AD' are pathological? We suggest that the status of AD' depends on the modal logic of interest. For example, if we are interested in augmenting \mathcal{A} with an $S5$ modality, then we would like to assure that the transition functions defined by action descriptions always map equivalence relations to equivalence relations. Hence, for epistemic logic, we want to say that AD is an admissible action description, but AD' is not admissible because it does not preserve reflexivity. For some other modal logics, however, AD' may be perfectly acceptable. For example, in a modal logic of belief, we may allow action descriptions like AD' because preserving reflexivity would not be important.

In a general modal setting, we would like to ensure that action descriptions preserve all of the important structural

characteristics of the modality under consideration. Preservation properties of this sort are the topic of the next section.

Standard Modal Logics

Let Π be a class of Kripke structures, and let AD be an action description. We say that AD *preserves* Π if $\Phi_{AD}(\mathcal{M}, A) \in \Pi$ whenever $\mathcal{M} \in \Pi$.

Definition 5 Let \mathcal{L} be a modal logic determined by a class of structures Π . An \mathcal{L} -description is an action description for $\mathcal{A}[\mathcal{L}]$ that preserves Π .

We now provide restricted classes of action descriptions that preserve some natural systems of modal logic. The proofs are included at the end of the paper.

Proposition 1 Let AD be a set of propositions such that $PRE \models \phi$ for every rule in AD of the form

$$A \text{ causes } \Box\phi \text{ if } PRE.$$

Then AD is a *KT-description*.

Proposition 2 Let AD be a set of propositions such that, for every rule in AD of the form

$$A \text{ causes } \Box\phi \text{ if } PRE,$$

AD also contains a rule of the form

$$A \text{ causes } \Box\neg PRE \text{ if } \neg\phi.$$

Then AD is a *KB-description*.

Proposition 3 Let AD be a set of propositions such that, for every rule in AD of the form

$$A \text{ causes } \Box\phi \text{ if } PRE,$$

AD also contains a rule of the form

$$A \text{ causes } \Box\Box\phi \text{ if } PRE.$$

Then AD is a *K4-description*.

Note that the conditions of each proposition are sufficient, but not necessary. Hence, although any action description satisfying Propositions 1 - 3 is an *S5-description*, there are also many *S5-descriptions* that do not satisfy the given conditions. For example, there are certainly finite action descriptions that preserve *S5*, but every action description satisfying the condition in Proposition 3 must be infinite.

Informally, the class of action descriptions that are admissible in the action language $\mathcal{A}[\mathcal{L}]$ is the class of \mathcal{L} -descriptions. However, giving a constructive definition of this class for any interesting modal logic \mathcal{L} is a non-trivial problem. For some natural modal logics, it is clear that no simple syntactic characterization can be given. For example, specifying a useful class of descriptions that preserve seriality is difficult, due to the fact that we only allow refinements. As a result, the current framework has somewhat limited applicability to logics determined by non-reflexive, serial structures.

Related Formalisms

Epistemic Action Languages

Two epistemic extensions of \mathcal{A} have been proposed in the literature (Lobo, Mendez, & Taylor 2001; Son & Baral 2001). Originally, each of them was named \mathcal{A}_K . In order to reduce ambiguity, we refer to the extension of Lobo et. al. as \mathcal{A}_L and we refer to the extension of Baral et. al. as \mathcal{A}_B .

The Action Language \mathcal{A}_L

Assume that the action symbols in \mathbf{A} are partitioned into sensing actions and non-sensing actions. In \mathcal{A}_L , there are two kinds of propositions. First of all, if A is a non-sensing action, f is a literal, and PRE is a conjunction of literals then

$$A \text{ causes } f \text{ if } PRE$$

is a proposition of \mathcal{A}_L . If A is a sensing action, f is a fluent symbol, and PRE is a conjunction of literals, then

$$A \text{ causes to know } f \text{ if } PRE$$

is a proposition. In this proposition, f must be a fluent symbol because the intended interpretation is that the execution of A causes the agent to know the truth value of f . This contrasts with our modal effect propositions which assert that an action causes a certain modal formula to be true. Informally, \mathcal{A}_L is making a higher level assertion about the property f rather than a first-order assertion about the truth value of a formula.

We remark that non-deterministic action effects can also be represented in \mathcal{A}_L through a third propositional form. However, we will not consider non-deterministic effects in this paper.

Given an action description AD , we say that f is a *potential sensing effect* of A if AD contains a proposition of the form

$$A \text{ causes to know } f \text{ if } PRE.$$

The *knowledge precondition* of a fluent symbol f with respect to a sensing action A is the disjunction of all of the preconditions appearing in propositions involving the action A and the sensing effect f .

The semantics of \mathcal{A}_L uses the notion of a *situation*. A situation is a set of *states* and a state is an interpretation of the set of fluent symbols. A fluent f is true in a situation Σ if it is true in every state in Σ , it is false if it is false in every state in Σ and it is unknown otherwise. Truth or falsity in \mathcal{A}_L is understood to reflect the knowledge of an agent, and knowledge is understood to be correct but not necessarily complete.

The semantics of \mathcal{A}_L associates a transition relation Φ_{AD} with every action description AD . We give the definition for the special case where each action has at most one potential sensing effect f . Let Σ be a situation and let A be an action symbol. The triple (Σ, A, Σ^*) is in Φ_{AD} if and only if the following hold.

1. If A is non-sensing, then the interpretation associated with each world in Σ^* is the interpretation obtained by updating the worlds of Σ in accordance with the \mathcal{A} propositions in AD .

2. If A is sensing, and f is unknown with precondition P , then Σ^* satisfies one of the following three conditions

- (a) Σ^* is the set of states in Σ where P and f hold
- (b) Σ^* is the set of states in Σ where P and $\neg f$ hold
- (c) Σ^* is the set of states in Σ where $\neg P$ holds

Hence, given a pair (Σ, A) where A has a single potential sensing effect, there will generally be three possible successor situations. A set of situations is called an *epistemic state*. Hence, the semantics of \mathcal{A}_L actually maps a situation and an action to an epistemic state.

We illustrate the intuition behind the effects of sensing actions with an example.

Example Consider the proposition

Listen causes to know MusicOn if \neg EarPlugs.

If an agent executes the action *Listen*, there are 3 possible outcomes.

1. The agent learns that *MusicOn* is true.
2. The agent learns that *MusicOn* is false.
3. The agent does not learn the value of *MusicOn*.

The only way the third possibility can happen is if the agent is wearing ear plugs. Hence, if the agent listens and still does not know the value of *MusicOn*, then the agent is justified in concluding that *EarPlugs* is true.

In general, each action may have several potential sensing effects. We briefly outline how the definition above can be extended to handle multiple sensing effects. We say that a situation is (f, P) -admissible with respect to an action A if it satisfies the definition given above. Now suppose that A has n potential sensing effects f_1, \dots, f_n with corresponding knowledge preconditions P_1, \dots, P_n . In this case, $\Phi_{AD}(\Sigma, A, \Sigma^*)$ holds if and only if Σ^* is the intersection of n situations $\Sigma_1, \dots, \Sigma_n$ where each Σ_i is (f_i, P_i) -admissible with respect to A . We refer the reader to (Lobo, Mendez, & Taylor 2001) for the details.

From \mathcal{A}_L to $\mathcal{A}[\mathcal{L}_K]$

In this section, we translate \mathcal{A}_L into $\mathcal{A}[\mathcal{L}_K]$. To begin, we present the translation and give an intuitive explanation.

Definition 6 Let AD be an action description in \mathcal{A}_L . The $\mathcal{A}[\mathcal{L}_K]$ action description $\tau(AD)$ is obtained from AD as follows.

1. Every non-sensing proposition in AD is in $\tau(AD)$.
2. For each action A with potential sensing effect f and knowledge precondition P , $\tau(AD)$ contains the following propositions:

A causes $K(f \wedge P)$ if $f \wedge P$
 A causes $K(\neg f \wedge P)$ if $\neg f \wedge P$
 A causes $K\neg P$ if $\neg P$.

Suppose that AD is an action description involving an action A with a single potential sensing effect f with knowledge precondition P . If A is executed, then $\Phi_{\tau(AD)}$ maps a




Figure 1: Partitioning Accessibility

Kripke structure \mathcal{M} to a new structure \mathcal{M}' in which the accessibility relation is refined as illustrated in Figure 1. Each circled region represents the set of worlds in which the indicated formula is true. The edges of \mathcal{M} that go between the circled regions are removed in \mathcal{M}' . Clearly, the three circled regions together form a partition of the universe. This observation suggests that action descriptions in the image of τ will preserve equivalence relations. We formalize this claim in the following proposition.

Proposition 4 Let AD be a set of \mathcal{A}_L propositions. The set $\tau(AD)$ is an S5-description.

We will prove that $\tau(AD)$ is the translation of AD into $\mathcal{A}[\mathcal{L}_K]$. First, we illustrate that there is a natural way to turn an epistemic state E into a Kripke structure \mathcal{M}_E . For the moment, assume that the collection of situations in E are pairwise disjoint. We discuss this assumption below. Define $\mathcal{M}_E = \langle M, R, \pi \rangle$ as follows.

1. $M = \bigcup E$
2. $R(m_1, m_2)$ iff there is $\Sigma \in E$ such that $m_1, m_2 \in \Sigma$
3. for any fluent f , $m \in \pi(f)$ iff $f \in m$

Clearly R is an equivalence relation and, moreover, each $\Sigma \in E$ corresponds to the equivalence class $[s]$ generated by $s \in \Sigma$. If Σ is a situation, we write \mathcal{M}_Σ as an abbreviation for $\mathcal{M}_{\{\Sigma\}}$.

The assumption that the elements of E are pairwise disjoint is a simplifying assumption to assure that each state in each situation in E corresponds to a unique element in the universe of \mathcal{M}_E . Without this assumption, we can still define a natural structure representing E by using a universe of ordered pairs where one component is an interpretation s and the other component is a situation $\Sigma \in E$ containing s . However, for our purposes, it is sufficient to consider the restricted case described above.

The following result demonstrates the close relationship between \mathcal{A}_L and $\mathcal{A}[\mathcal{L}_K]$.

Theorem 1 Let AD be an \mathcal{A}_L action description, let Σ, Σ^* be non-empty situations, and let A be a sensing action in AD . Then $\Phi_{AD}(\Sigma, A, \Sigma^*)$ if and only if Σ^* is an equivalence class in $\Phi_{\tau(AD)}(\mathcal{M}_\Sigma, A)$.

Intuitively, this says that, under a natural translation between situations and structures, Φ_{AD} and $\Phi_{\tau(AD)}$ represent the same transition relation.

From \mathcal{A}_B to $\mathcal{A}[\mathcal{L}_K]$

Our summary of \mathcal{A}_B and the ensuing translation will be somewhat brief, because the details are very close to the translation we have just seen. The syntax of \mathcal{A}_B introduces a new set of propositions of the form

A determines f .

The intended interpretation of such a proposition is that an agent will know the value of f after executing A .

The semantics of \mathcal{A}_B is based on pairs $\langle s, \Sigma \rangle$, where s is a state and Σ is a set of states containing s . The state s represents the actual world, and Σ represents those worlds that are believed to be possible. For simplicity, we assume that all actions are either sensing actions or non-sensing actions; strictly speaking this distinction is not required.

We associate a transition function Φ_{AD} with every \mathcal{A}_B action description. If A is a non-sensing action, $\Phi_{AD}(\langle s, \Sigma \rangle, A)$ is obtained by updating each world in Σ in accordance with the semantics of A . If A is a sensing action, and

A determines f

is in AD , then $\Phi_{AD}(\langle s, \Sigma \rangle, A)$ is obtained by removing from Σ each world that differs from s in the interpretation of f .

The translation from \mathcal{A}_B is essentially identical to that from \mathcal{A}_L , except that there are no knowledge preconditions. Given AD , we construct $\sigma(AD)$ by replacing each proposition of the form

A determines f

with two propositions

A causes Kf if f
 A causes $K\neg f$ if $\neg f$

We have the following result.

Theorem 2 *Let AD be a set of \mathcal{A}_B propositions, let s be a state, let Σ be a set of states containing s , and let A be an action symbol that only occurs in propositions of the form*

A determines f .

Then $\Phi_{AD}(\langle s, \Sigma \rangle, A) = \langle s, [s] \rangle$ where $[s]$ denotes the equivalence class of s in $\Phi_{\sigma(AD)}(\mathcal{M}_\Sigma, A)$.

We remark that there is one important difference between \mathcal{A}_B and $\mathcal{A}[\mathcal{L}_K]$. Namely, the semantics of \mathcal{A}_B incorporates a distinguished state representing the actual state of the world. This difference can be eliminated by considering Kripke structures of the form $\langle M, R, \pi, a \rangle$, where a represents the actual world. The semantics of $\mathcal{A}[\mathcal{L}_K]$ is unchanged, except to state that a remains constant in every transition.

We have seen that both \mathcal{A}_B and (deterministic) \mathcal{A}_L can be naturally embedded in the language $\mathcal{A}[\mathcal{L}_K]$. As a result,

a simplified fragment of $\mathcal{A}[\mathcal{L}_K]$ can be translated into logic programming for planning (Son, Huy, & Baral 2004). Similarly, a restricted portion of the language can be translated into epistemic logic programming (Lobo, Mendez, & Taylor 2001).

Discussion

Future Work

There are three directions in which we would like to extend the present work. First of all, we would like to formally address propositions that add edges to the accessibility relation. Such propositions are useful for describing non-deterministic action effects. Moreover, some combination of adding and removing edges is required to give natural descriptions that preserve seriality. This is an important concern for the representation of some natural modal logics. For example, if we would like to represent change in the context of deontic logic, we need to be able to preserve seriality.

The second extension we would like to consider would allow multiple agents, each with their own individual knowledge. Some work has already been done on the treatment of epistemic action effects in a multi-agent environment (van Ditmarsch 2002). This is a difficult problem, because a single action may affect the knowledge of each agent differently. For example, van Ditmarsch formally analyzes the different epistemic effects brought about by whispering some information to another agent versus announcing the same information. We would like to be able to give a compact treatment of this kind of action effect in the action language framework.

Third, we would like to be able to implement a planner for a less restricted class of modal action languages. In particular, we would like to allow some limited nesting of modal operators in our descriptions. Such nesting is required to address the representation of simple knowledge games, and it is also required for the verification of communication protocols. We would be interested in demonstrating the practical utility of modal action languages by solving some realistic verification problems.

Finally, we remark that there is nothing special about the action language \mathcal{A} for our approach to modal action effects. It simply provides a semantic framework with which we update states in a Kripke structure. Clearly, many action languages could be used for this purpose.

Conclusion

There are many cases where a notion of necessity is useful for representing the state of the world. By combining a modal logic with the action language \mathcal{A} , we can create a simple tool for representing and reasoning about change in such an environment. In this paper, we have provided a simple extension of the action language framework for reasoning about transitions between Kripke structures. The paradigmatic example has been the representation of changes in the knowledge of a single agent. We have seen that the modal approach naturally subsumes existing approaches to reasoning about epistemic action effects, and it requires little for-

mal machinery on top of elementary modal logic and the action language \mathcal{A} .

Proofs

Proposition 1 *Let AD be a set of propositions such that $PRE \models \phi$ for every rule in AD of the form*

$$A \text{ causes } \Box\phi \text{ if } PRE.$$

Then AD is a KT -description.

Proof Let AD be a set of propositions satisfying the premise, let A be an action symbol, let $\mathcal{M} = \langle M, R, \pi \rangle$ with R reflexive, and let $\Phi_{AD}(\mathcal{M}, A) = \langle M, R^*, \pi^* \rangle$. Suppose that $(m, m) \notin R^*$ for some $m \in M$. Since R is reflexive, we know that $(m, m) \in R$. Informally, this means that the edge (m, m) is removed in the transition between structures. Hence, there must be some proposition

$$A \text{ causes } \Box\phi \text{ if } PRE$$

in AD such that $\mathcal{M}, m \models PRE$ and $\mathcal{M}, m \models \neg\phi$. This contradicts our assumption that ϕ is a logical consequence of PRE .

Proposition 2 *Let AD be a set of propositions such that, for every rule in AD of the form*

$$A \text{ causes } \Box\phi \text{ if } PRE,$$

AD also contains a rule of the form

$$A \text{ causes } \Box\neg PRE \text{ if } \neg\phi.$$

Then AD is a KB -description.

Proof Let AD be a set of propositions satisfying the premise, let A be an action symbol, let $\mathcal{M} = \langle M, R, \pi \rangle$ with R symmetric, and let $\Phi_{AD}(\mathcal{M}, A) = \langle M, R^*, \pi^* \rangle$. Suppose that $(m, n) \in R^*$ and $(n, m) \notin R^*$ for some m, n . Since $(n, m) \notin R^*$, there must be some proposition of the form

$$A \text{ causes } \Box\phi \text{ if } PRE$$

in AD , where $\mathcal{M}, n \models PRE$ and $\mathcal{M}, m \models \neg\phi$. By assumption, AD also contains

$$A \text{ causes } \Box\neg PRE \text{ if } \neg\phi.$$

Since $\mathcal{M}, m \models \neg\phi$ and $\mathcal{M}, n \models PRE$, it follows from Definition 4 that $(m, n) \notin R^*$, which is a contradiction. Hence R^* is symmetric.

Proposition 3 *Let AD be a set of propositions such that, for every rule in AD of the form*

$$A \text{ causes } \Box\phi \text{ if } PRE,$$

AD also contains a rule of the form

$$A \text{ causes } \Box\Box\phi \text{ if } PRE.$$

Then AD is a $K4$ -description.

Proof Let AD be a set of propositions satisfying the premise, let A be an action symbol, let $\mathcal{M} = \langle M, R, \pi \rangle$ with R transitive, and let $\Phi_{AD}(\mathcal{M}, A) = \langle M, R^*, \pi^* \rangle$. Assume that $(m, n) \in R^*$ and $(n, p) \in R^*$. Now suppose $(m, p) \notin R^*$, so there is some proposition

$$A \text{ causes } \Box\phi \text{ if } PRE$$

in AD such that $\mathcal{M}, m \models PRE$ and $\mathcal{M}, p \models \neg\phi$. By assumption, AD also contains

$$A \text{ causes } \Box\Box\phi \text{ if } PRE.$$

But then, since $\mathcal{M}, m \models PRE$ and $(m, n) \in R^*$, it follows that $\mathcal{M}, n \models \Box\phi$. Then, since $(n, p) \in R^*$, we must have $\mathcal{M}, p \models \phi$. This is a contradiction, hence R^* is transitive.

Proposition 4 *Let AD be a set of \mathcal{A}_L propositions. The set $\tau(AD)$ is an $S5$ -description.*

Proof Let $\mathcal{M} = \langle M, R, \pi \rangle$. Let A be a sensing action with n potential sensing effects f_1, \dots, f_n with knowledge preconditions P_1, \dots, P_n . Let $\Phi_{\tau(AD)}(\mathcal{M}, A) = \langle M, R^*, \pi \rangle$. We remark that π has remained unchanged because A is a sensing action. We prove that R^* is an equivalence.

By Proposition 1, R^* is reflexive. Moreover, it is straightforward to modify the proof of Proposition 2 to prove that R^* is symmetric. All that remains is to show that R^* is transitive. Suppose that R^*mn and R^*np , but not R^*mp . There are three possible cases to consider.

1. $m \models f_i \wedge P_i$ and $p \models \neg(f_i \wedge P_i)$ for some i
2. $m \models \neg f_i \wedge P_i$ and $p \models \neg(\neg f_i \wedge P_i)$ for some i
3. $m \models \neg P_i$ and $p \models P_i$ for some i

Suppose the first case holds. Since R^*mn , it must be the case that $n \not\models \neg(f_i \wedge P_i)$. Since R^*np , it must be the case that $n \models \neg(f_i \wedge P_i)$. Hence, the first case is not possible. The other two cases lead to similar contradictions. Therefore R^* is transitive.

Theorem 1 *Let AD be an \mathcal{A}_L action description, let Σ, Σ^* be non-empty situations, and let A be a sensing action in AD . Then $\Phi_{AD}(\Sigma, A, \Sigma^*)$ if and only if Σ^* is an equivalence class in $\Phi_{\tau(AD)}(\mathcal{M}_\Sigma, A)$.*

Proof Let A be an action with n potential sensing effects f_1, \dots, f_n with knowledge preconditions P_1, \dots, P_n .

Note that $\mathcal{M}_\Sigma = \langle \Sigma, R, \pi \rangle$, where R is universal and $p \in \pi(s)$ iff $p \in s$. Let $\Phi_{\tau(AD)}(\mathcal{M}_\Sigma, A) = \langle \Sigma, R^*, \pi \rangle$.

The relation R^* is obtained by making the following changes to R , for each $i \leq n$:

1. remove edges $(m, n) \in R$ where $m \models (f_i \wedge P_i)$ and $n \models \neg(f_i \wedge P_i)$
2. remove edges $(m, n) \in R$ where $m \models (\neg f_i \wedge P_i)$ and $n \models \neg(\neg f_i \wedge P_i)$
3. remove edges $(m, n) \in R$ where $m \models \neg P_i$ and $n \models P_i$

Hence R^*mn if and only if, for each i , one of these conditions holds:

1. $m, n \models f_i \wedge P_i$

2. $m, n \models \neg f_i \wedge P_i$

3. $m, n \models \neg P_i$

So Σ^* is an equivalence class in R^* if and only if $\Sigma^* = \bigcap_{i=1}^n \Sigma_i$ where each Σ_i is a situation satisfying exactly one of the following conditions:

1. Σ_i is the set of states in Σ where f_i and P_i hold
2. Σ_i is the set of states in Σ where $\neg f_i$ and P_i hold
3. Σ_i is the set of states in Σ where $\neg P_i$ holds.

By definition, this holds if and only if $\Phi_{AD}(\Sigma, A, \Sigma^*)$.

Theorem 2 Let AD be a set of \mathcal{A}_B propositions, let s be a state, let Σ be a set of states containing s , and let A be an action symbol that only occurs in propositions of the form

A determines f .

Then $\Phi_{AD}(\langle s, \Sigma \rangle, A) = \langle s, [s] \rangle$ where $[s]$ denotes the equivalence class of s in $\Phi_{\sigma(AD)}(\mathcal{M}_\Sigma, A)$.

Proof Let f_1, \dots, f_n be the set of fluent symbols which occur in AD in propositions of the form

A determines f_i .

Hence $\Phi_{AD}(\langle s, \Sigma \rangle, A) = \langle s, \Sigma^* \rangle$ where Σ^* is obtained from Σ by removing all states that do not agree with s on the value of some f_i .

Recall that the accessibility relation in \mathcal{M}_Σ is universal. Hence, the accessibility relation in $\Phi_{\sigma(AD)}(\mathcal{M}_\Sigma, A)$ is obtained by starting with a universal relation and then removing every edge from s to any state that assigns a different truth value to some f_i . The edges from s that remain in $\Phi_{\sigma(AD)}(\mathcal{M}_\Sigma, A)$ are precisely the edges to states that agree with s on the truth value of every f_i . Therefore $\Sigma^* = [s]$.

References

- Baral, C., and Gelfond, M. 1997. Reasoning about effects of concurrent actions. *Journal of Logic Programming* 31(1-3):85–117.
- Baral, C.; Gelfond, M.; and Proveti, A. 1997. Representing actions: Laws, observations and hypothesis. *Journal of Logic Programming* 31(1-3):201–243.
- Burrows, M.; Abadi, M.; and Needham, R. 1989. A logic of authentication. Technical Report 39, Digital Systems Research Center.
- Chellas, B. 1980. *Modal Logic: An Introduction*. Cambridge University Press.
- Fagin, R.; Halpern, J.; Moses, Y.; and Vardi, M. 1995. *Reasoning About Knowledge*. MIT Press.
- Gelfond, M., and Lifschitz, V. 1993. Representing action and change by logic programs. *Journal of Logic Programming* 17:301–321.
- Kahramanoğlu, O., and Thielscher, M. 2003. A formal assessment result for fluent calculus using the action description language \mathcal{A}_k . In Kruse, R., ed., *Proceedings of the German Annual Conference on Artificial Intelligence (KI)*, volume 2821 of *Lecture Notes in Artificial Intelligence*, 209–223. Springer-Verlag.

Lifschitz, V. 1999. Action languages, answer sets and planning. In Apt, K.; Marek, V.; Truszczyński, M.; and Warren, D., eds., *The Logic Programming Paradigm: A 25-Year Perspective*. Springer-Verlag.

Lobo, J.; Mendez, G.; and Taylor, S. 2001. Knowledge and the action description language a. *Theory and Practice of Logic Programming* 1(2):129–184.

Son, T., and Baral, C. 2001. Formalizing sensing actions: A transition function based approach. *Artificial Intelligence* 125(1-2):19–91.

Son, T.; Huy, P.; and Baral, C. 2004. Planning with sensing actions and incomplete information using logic programming. In *Proceedings of LPNMR7*.

van Ditmarsch, H. 2002. Descriptions of game actions. *Journal of Logic, Language and Information (JoLLI)* 11:349–365.