

Modular action language \mathcal{ALM}

DANIELA INCLEZAN

*Department of Computer Science and Software Engineering, Miami University
Oxford, OH 45056, USA
(e-mail: inclezd@miami0h.edu)*

MICHAEL GELFOND

*Department of Computer Science, Texas Tech University
Lubbock, TX 79409, USA
(e-mail: michael.gelfond@ttu.edu)*

submitted 1 December 2013; revised 15 October 2014; accepted 12 May 2015

Abstract

The paper introduces a new modular action language, \mathcal{ALM} , and illustrates the methodology of its use. It is based on the approach of Gelfond and Lifschitz (1993, *Journal of Logic Programming* 17, 2–4, 301–321; 1998, *Electronic Transactions on AI* 3, 16, 193–210) in which a high-level action language is used as a front end for a logic programming system description. The resulting logic programming representation is used to perform various computational tasks. The methodology based on existing action languages works well for small and even medium size systems, but is not meant to deal with larger systems that require *structuring of knowledge*. \mathcal{ALM} is meant to remedy this problem. Structuring of knowledge in \mathcal{ALM} is supported by the concepts of *module* (a formal description of a specific piece of knowledge packaged as a unit), *module hierarchy*, and *library*, and by the division of a system description of \mathcal{ALM} into two parts: *theory* and *structure*. A *theory* consists of one or more modules with a common theme, possibly organized into a module hierarchy based on a *dependency relation*. It contains declarations of sorts, attributes, and properties of the domain together with axioms describing them. *Structures* are used to describe the domain's objects. These features, together with the means for defining classes of a domain as special cases of previously defined ones, facilitate the stepwise development, testing, and readability of a knowledge base, as well as the creation of knowledge representation libraries.

KEYWORDS: logic programming, reasoning about actions and change, action language

1 Introduction

In this paper, we introduce a new modular action language, \mathcal{ALM} , and illustrate the principles of its use. Our work builds upon the methodology for representing knowledge about discrete dynamic systems introduced by Gelfond and Lifschitz (1993; 1998). In this approach, a system is viewed as a *transition diagram* whose nodes correspond to possible states of the system and whose arcs are labeled by actions. The diagram is defined by a *system description* – a collection of statements in a high-level *action language* expressing the direct and indirect effects of actions as

well as their executability conditions (see, for instance, action languages \mathcal{A} (Gelfond and Lifschitz 1993), \mathcal{B} (Gelfond and Lifschitz 1998); \mathcal{AL} (Turner 1997; Baral and Gelfond 2000); the non-modular extension of \mathcal{AL} with multi-valued fluents (Dovier *et al.* 2007); \mathcal{C} (Giunchiglia and Lifschitz 1998); $\mathcal{C}+$ (Giunchiglia *et al.* 2004a); \mathcal{K} (Eiter *et al.* 2004); \mathcal{D} (Strass and Thielscher 2012); \mathcal{E} (Kakas and Miller 1997); \mathcal{H} (Chintabathina *et al.* 2005; Chintabathina 2012)). Such languages allow concise representations of very large diagrams. In order to reason about the system, its action language description is often translated into a logic program under the answer set semantics (Gelfond and Lifschitz 1988; 1991). This allows for the use of Answer Set Programming (ASP) (Gelfond and Lifschitz 1991; Niemelä 1998; Marek and Truszczyński 1999) to perform complex reasoning tasks such as planning, diagnosis, etc. This methodology was successfully used in a number of interesting medium size applications, but does not seem to be fully adequate for applications requiring a larger body of knowledge about actions and their effects, step-wise design, and multiple use of, possibly previously designed, pieces of knowledge. (The phenomenon is of course well known in Computer Science. Similar considerations led to the early development of notions of subroutine and module in procedural programming. In logic programming, early solutions were based on the concepts of macro and template (Baral *et al.* 2006; Calimeri and Ianni 2006).) Just a few examples of domains that we consider large enough to benefit from the above-mentioned practices are: the Zoo World and Traffic World examples proposed by Sandewall (1999) and modeled in (Henschel and Thielscher 2000; Akman *et al.* 2004); the Monkey and Banana Problem by McCarthy (1963; 1968) and formalized in (Erdoğan and Lifschitz 2006; Erdoğan 2008); the Missionaries and Cannibals Problem by McCarthy (1998) represented in (Gustafsson and Kvarnström 2004; Erdoğan 2008).

This inadequacy is due to the fact that most action languages, with some notable exceptions like *MAD* (Erdoğan and Lifschitz 2006; Lifschitz and Ren 2006; Desai and Singh 2007) and *TAL-C* (Gustafsson and Kvarnström 2004), have *no built-in features for supporting the description of a domain's ontology and its objects, and for structuring knowledge and creating knowledge-based libraries*. *ALM* is designed to address these problems. It is based on an earlier action language, \mathcal{AL} , introduced in Gelfond and Incezan (2009) where it is called \mathcal{AL}_d , which so far has been the authors' language of choice (see, for instance, (Gelfond and Kahl 2014)). However, the basic ideas presented in the paper can be used for defining versions of *ALM* based on other action languages.

ALM has constructs for representing *sorts* (i.e., classes, kinds, types, categories) of objects relevant to a given domain, their *attributes*,¹ and a subsort relation that can be viewed as a directed acyclic graph (DAG). We refer to this relation as a *sort hierarchy*. These constructs support a methodology of knowledge representation that starts with determining the sorts of objects in a domain and formulating the domain's causal laws and other axioms in terms of these sorts. The *specialization*

¹ Attributes are intrinsic properties of a sort of objects. In *ALM*, they are represented by *possibly* partial functions defined on elements of that sort.

construct of the language, which corresponds to the links of the sort hierarchy, allows to define new sorts (including various sorts of actions) in terms of other, previously defined sorts.

The definition of particular objects populating the sorts is usually given only when the domain knowledge is used to solve a particular task, e.g., predicting the effects of some particular sequences of actions, planning, diagnosis, etc.

It is worth noting that allowing definitions of actions as special cases of other, previously defined actions was one of the main goals of actions languages like \mathcal{ALM} and MAD . Such definitions are not allowed in traditional action languages. \mathcal{ALM} 's solution consists in allowing action sorts, which do not exist in MAD . We believe that the \mathcal{ALM} solution is simpler than the one in MAD , where special cases of actions are described using import statements (similar to bridge rules in $\mathcal{C}+$).

\mathcal{ALM} also facilitates the introduction of particular domain objects (including particular actions) that are defined as *instances* of the corresponding sorts. For example, an action $go(bob, london, paris)$ can be defined as an instance of action sort *move* with attributes *actor*, *origin*, and *destination* set to *bob*, *london*, and *paris* respectively; action $go(bob, paris)$ is another instance of the same sort in which the origin of the *move* is absent. Note that since axioms of the domain are formulated in terms of sorts and their attributes, they are applicable to both of these actions. This is very different from the traditional action language representation of objects as *terms*, which requires separate axioms for $go(bob, london, paris)$ and $go(bob, paris)$.

Structuring of knowledge in \mathcal{ALM} is supported by the concepts of *module*, *module hierarchy*, and *library*, and by the division of a system description of \mathcal{ALM} into two parts: *theory* and *structure*. *Theories* contain declarations of sorts, attributes, and properties of the domain together with axioms describing them, while *structures* are used to describe the domain's objects. Rather traditionally, \mathcal{ALM} views a module as a formal description of a specific piece of knowledge packaged as a unit. A theory consists of one or more modules with a common theme, possibly organized into a module hierarchy based on a *dependency relation*. Modules of a theory can be developed and tested independently, which facilitates the reuse of knowledge and step-wise development and refinement (Wirth 1971) of knowledge bases, and increases their *elaboration tolerance* (McCarthy 1998).

Theories describing recurrent knowledge may be stored in libraries and used in different applications. The *structure* part of an \mathcal{ALM} system description contains definitions of objects of the domain together with their sorts, values of their attributes, and *statics* – relations between objects that cannot be changed by actions. If a system description of \mathcal{ALM} satisfies some natural consistency requirements and provides complete information about the membership of its objects in the system's sorts then it describes the unique transition diagram containing all possible trajectories of the system. In this sense, \mathcal{ALM} is semantically similar to \mathcal{AL} . There are also some substantial differences. First, if no complete information about membership of objects in sorts is given, then the system description specifies the *collection* of transition diagrams corresponding to various possible placements of objects in the system's sorts. This has no analog in \mathcal{AL} . Second, in addition to the semantics of its system descriptions, \mathcal{ALM} provides semantics for its theories.

Informally, a theory of \mathcal{ALM} can be viewed as a function taking as an input objects of the domain, their sort membership, and the values of static relations, and returning the corresponding transition diagram – a possible *model* of the theory. (This definition has some similarity with the notions of module developed for logic programs under the answer set semantics, e.g., (Oikarinen and Janhunnen 2006) and (Lierler and Truszczyński 2013). Accurate mathematical analysis of these similarities and their use for automatic reasoning in \mathcal{ALM} is a matter for future research.) The availability of a formal semantics clarifies the notion of an \mathcal{ALM} theory and allows us to define an entailment relation (T entails q if q is true in every model of T).

To accurately define the semantics of \mathcal{ALM} theories, we introduce the notion of a *basic action theory* (BAT) – a pair consisting of a specific type of sorted signature (which we call an *action signature*), and a set of axioms over this signature. An interpretation I of the signature of a BAT theory T defines: objects, their sort membership, and statics; while T can be viewed as a function that takes I as input and returns the transition diagram $T(I)$ defined by I . In a sense, $T(I)$ is very similar to system descriptions of \mathcal{AL} and other traditional action languages. The difference is in the forms of their signatures and axioms. As in \mathcal{AL} , the precise definition of states and transitions of $T(I)$ is given in terms of its translation into logic programs under the answer set semantics.

A system description D of \mathcal{ALM} can be viewed as a formal definition of a particular BAT theory T , and a class of its interpretations. The latter is given by the structure of D , the former by its theory. If the structure of D is complete, i.e., defines exactly one interpretation I , then D represents $T(I)$.

An earlier version of \mathcal{ALM} has been tested in the context of a real-life application, as part of our collaboration on Project Halo. Project Halo is a research effort by Vulcan Inc. aimed towards the development of a Digital Aristotle – “*an application containing large volumes of scientific knowledge and capable of applying sophisticated problem-solving methods to answer novel questions*” (Gunning *et al.* 2010). The Digital Aristotle uses the knowledge representation language called SILK (Semantic Inferencing on Large Knowledge) (Grosz *et al.* 2009), which is based on the well-founded semantics (Van Gelder *et al.* 1991) and transaction logic with defaults and argumentation theories (Fodor and Kifer 2011). Our first contribution to Project Halo consisted in creating an \mathcal{ALM} formalization of an important biological process, *cell division* (Incezan and Gelfond 2011). The use of \mathcal{ALM} allowed us to create libraries of knowledge and reuse information when representing the cell division domain. As a second step, we created a question answering system capable of answering complex temporal projection questions about this biological process (Incezan 2010). Our model of cell division represented in the higher level language \mathcal{ALM} served as a front end for the question answering system, which was implemented both in ASP and in the language of the Digital Aristotle.

Our language has evolved since our collaboration on Project Halo. The version of \mathcal{ALM} presented here differs from that described in previous papers (Gelfond and Incezan 2009; Incezan and Gelfond 2011) in various ways. We simplified and generalized the basic concepts of our language, as well as its syntax and semantics.

(We say more about the new features of \mathcal{ALM} in the conclusion section of the paper.) The reasoning in \mathcal{ALM} is based on the reduction of temporal projection, planning, diagnosis, etc. to the problem of computing the answer sets of logic programs (for a general description see, for instance, (Baral 2003)) by ASP solvers (see (Niemelä and Simons 1997), (Gebser *et al.* 2012), or (Leone *et al.* 2006)).

The rest of this paper is organized as follows: we first introduce the concept of BAT , which is a fundamental concept in this work. We then describe language \mathcal{ALM} and the methodology of \mathcal{ALM} 's use. We end with conclusions and future work. There are three online appendices containing the grammar of \mathcal{ALM} (pp. 1–3), the description of the use of \mathcal{ALM} in Digital Aristotle (pp. 4–10), and a comparison between \mathcal{ALM} and MAD (pp. 11–17).

2 Basic action theories

In this section, we give the definition of a fundamental concept of \mathcal{ALM} called BAT . A BAT consists of a collection of axioms over a so-called *action signature* – a special type of sorted signature providing suitable vocabulary for representing knowledge about dynamic domains. Sorted signatures needed for our purpose are somewhat atypical. They allow partial functions and contain means for describing a hierarchy of sorts and attributes of their elements. We start with the precise definition of sorted signatures and their interpretations.

2.1 Sorted signatures and their interpretations

By *sorted signature*, we mean a tuple

$$\Sigma = \langle \mathcal{C}, \mathcal{O}, \mathcal{H}, \mathcal{F} \rangle$$

where \mathcal{C} , \mathcal{O} , and \mathcal{F} are sets of strings over some fixed alphabet. The strings are used to name *sorts*, *objects*, and (possibly partial) *functions* respectively. Each function symbol $f \in \mathcal{F}$ is assigned a positive integer n (called f 's arity), sorts c_0, \dots, c_n for its parameters, and sort c for its values. We refer to c as the *range* of f and use the standard mathematical notation $f : c_0 \times \dots \times c_n \rightarrow c$ for this assignment.

Finally, \mathcal{H} is a *sort hierarchy* – a DAG with two types of nodes: *sort nodes* labeled by sort names from \mathcal{C} , and *object nodes* labeled by object names from \mathcal{O} . Whenever convenient we identify nodes of the hierarchy with their labels. A link from sort c_1 to sort c_2 , denoted by $\langle c_1, c_2 \rangle$, indicates that elements of sort c_1 are also elements of sort c_2 . We refer to c_2 as a *parent* of c_1 . A link from object o to a sort c , denoted by $\langle o, c \rangle$, indicates that object o is of sort c . For simplicity, we assume that the graph has exactly one sink node, which corresponds to the sort containing all the elements of the hierarchy. A triple $\langle \mathcal{C}, \mathcal{O}, \mathcal{H} \rangle$ will be sometimes referred to as an *ontology*.

Sorts, object constants, and functions of a sorted signature are normally partitioned into *user-defined*, *pre-defined*, and *special*.

The collection of *pre-defined* symbols may include names for some commonly used sorts and functions, such as: sorts *booleans* and *integers*; a sort $[m..n]$ for every pair of natural numbers m and n such that $m < n$, denoting the set of natural

numbers in the closed interval $[m, n]$; standard object constants *true*, *false*, 0, 1, 2, etc., denoting elements of these sorts; standard arithmetic functions and relations $+$, $-$, $*$, $/$, *mod*, $<$, \leq , etc. (The list is not exhaustive. When needed, we may introduce other similar symbols.) All these symbols are pre-interpreted, i.e., come with their usual mathematical interpretations.

The collection of *special* symbols consists of

- Sorts and function symbols pertinent to sort hierarchies of sorted signatures:
 - Sort *nodes* denoting the collection of sorts labeling the sort nodes of \mathcal{H} . This sort is never used as a label of a node in \mathcal{H} .
 - Sort *object_constants* denoting the collection of constants labeling the object nodes of \mathcal{H} . This sort is never used as a label of a node in \mathcal{H} .
 - Sort *universe* denoting the collection of elements of sorts from \mathcal{H} .
 - Function symbol *link* : $nodes \times nodes \rightarrow booleans$ where *link*(c_1, c_2) returns *true* iff \mathcal{H} contains a link from sort c_1 to sort c_2 .
 - Function symbol *is_a* : $universe \times nodes \rightarrow booleans$ where *is_a*(x, c) returns *true* if c is a source node of \mathcal{H} (i.e., c has no subsorts in \mathcal{H}) and object x from the universe is of the sort denoted by c .
 - Function symbol *instance* : $universe \times nodes \rightarrow booleans$ denoting the membership relation between objects of the universe and the sorts of the domain. This function will be later defined in terms of function *is_a*.
 - Function symbols *subsort* : $nodes \times nodes \rightarrow booleans$, *has_child*, *has_parent*, *sink*, *source* : $nodes \rightarrow booleans$ describing properties of sorts of \mathcal{H} and their members. All these functions (with their self-explanatory meaning) will be later defined in terms of function *link*.
- Function symbol *dom_f* : $c_0 \times \dots \times c_n \rightarrow booleans$ (read as *domain of f*) for every user-defined function symbol $f : c_0 \times \dots \times c_n \rightarrow c$ with $n > 0$.

Terms of a sorted signature are defined as usual:

- A variable and an object constant is a term.
- If $f : c_0 \times \dots \times c_n \rightarrow c$ is a function symbol and t_0, \dots, t_n are terms then $f(t_0, \dots, t_n)$ is a term.

Expressions of the form

$$t_1 = t_2 \quad \text{and} \quad t_1 \neq t_2 \tag{1}$$

are called *literals*. Positive literals are also referred to as *atoms*. (For simplicity of presentation, we use standard shorthands and write t and $\neg t$ instead of $t = true$ and $t = false$, respectively; $3 \leq 5$ instead of $\leq (3, 5)$; etc.) Terms and literals not containing variables are called *ground*. Our notion of an interpretation of a sorted signature is slightly different from the traditional one.

Definition 1 (Interpretation)

An interpretation \mathcal{I} of Σ consists of

- A non-empty set $|\mathcal{I}|$ of strings called the *universe* of \mathcal{I} .

- An assignment that maps
 - every user-defined sort c of \mathcal{H} into a subset $\mathcal{I}(c)$ of $|\mathcal{I}|$ and user-defined object constant o into an element from $|\mathcal{I}|$;
 - every user-defined function symbol $f : c_0 \times \dots \times c_n \rightarrow c$ of Σ into a (possibly partial) function $\mathcal{I}(f) : \mathcal{I}(c_0) \times \dots \times \mathcal{I}(c_n) \rightarrow \mathcal{I}(c)$;
 - the special function is_a into function $\mathcal{I}(is_a)$ such that:
 - for every $x \in |\mathcal{I}|$ and every sort c of \mathcal{H} , $\mathcal{I}(is_a)(x, c)$ is *true* iff c is a source node of \mathcal{H} and $x \in \mathcal{I}(c)$ and
 - for every object o and sort c of \mathcal{H} , $\mathcal{I}(is_a)(\mathcal{I}(o), c)$ is *true* iff $\langle o, c \rangle \in \mathcal{H}$;
 - the special function $link$ into function $\mathcal{I}(link)$ such that for every two sort nodes c_1, c_2 , $\mathcal{I}(link)(c_1, c_2)$ is *true* iff $\langle c_1, c_2 \rangle \in \mathcal{H}$;
 - the special function dom_f for user-defined function $f : c_0 \times \dots \times c_n \rightarrow c$ into function $\mathcal{I}(dom_f)$ such that for every $\bar{x} \in \mathcal{I}(c_0) \times \dots \times \mathcal{I}(c_n)$, $\mathcal{I}(dom_f)(\bar{x})$ is *true* iff \bar{x} belongs to the domain of $\mathcal{I}(f)$.
- On pre-defined symbols, \mathcal{I} is identified with the symbols' standard interpretations.

An interpretation \mathcal{I} of Σ can be naturally extended to ground terms: if \mathcal{I} is defined on terms t_1, \dots, t_n and $\mathcal{I}(f)$ is defined on the tuple $(\mathcal{I}(t_1), \dots, \mathcal{I}(t_n))$ then

$$\mathcal{I}(f(t_1, \dots, t_n)) =_{def} \mathcal{I}(f)(\mathcal{I}(t_1), \dots, \mathcal{I}(t_n)).$$

Otherwise, $\mathcal{I}(f(t_1, \dots, t_n))$ is undefined.

Finally, we say that an atom $t_1 = t_2$ is

- *true* in \mathcal{I} if both $\mathcal{I}(t_1)$ and $\mathcal{I}(t_2)$ are defined and have the same value;
- *false* in \mathcal{I} if both $\mathcal{I}(t_1)$ and $\mathcal{I}(t_2)$ are defined and have different values; and
- *undefined* in \mathcal{I} otherwise.

Similarly, a literal $t_1 \neq t_2$ is *true* in \mathcal{I} if $t_1 = t_2$ is *false* in \mathcal{I} ; it is *false* in \mathcal{I} if $t_1 = t_2$ is *true* in \mathcal{I} ; and *undefined* otherwise.

Note that every interpretation \mathcal{I} can be uniquely represented by the collection of atoms that are true in this interpretation. For instance, for every sort c of \mathcal{H} , $\mathcal{I}(c)$ can be represented as the set $\{instance(o, c) : \mathcal{I}(o) \in \mathcal{I}(c)\}$; for a unary function f , $\mathcal{I}(f)$ can be viewed as the set $\{f(x) = y : \mathcal{I}(x) \in \mathcal{I}(dom_f) \text{ and } \mathcal{I}(f)(\mathcal{I}(x)) = \mathcal{I}(y)\}$, etc.

2.2 Action signature and axioms of a BAT

Since \mathcal{ALM} is a language for specifying properties of actions, in what follows we limit ourselves to *action signatures* – sorted signatures that

- contain a special sort *actions* and
- have their user-defined and special function symbols divided into three disjoint categories: *attributes*, *statics*, and *fluents*. Attributes describe intrinsic properties

of objects of a given sort; statics and fluents describe relations between objects. Values of attributes and statics are constants – they cannot be changed by actions. The values of fluents can. Both statics and fluents are further divided into *basic* and *defined*. The latter are *total boolean functions* that can be defined in terms of the former. They are used primarily for the brevity of representation.

A literal (atom) in which f is an attribute is called an *attribute literal (atom)*. Similarly for *static* and *fluent* literals that are, in turn, divided into *basic* and *defined*. We assume that all special functions of an action signature, except dom_f , are defined statics; dom_f is a basic fluent when f is a basic fluent and a defined static otherwise. Since the semantics of \mathcal{ALM} will be defined in terms of a version of ASP with function symbols, ASP{f} (Balduccini 2013), which does not allow terms with nested user-defined functions, we *limit atoms of an action signature to those constructed from terms with at most one user-defined function symbol*. (This is not a serious limitation and can easily be avoided by viewing nested terms as shorthands.)

We can now define the syntax and informal semantics of *statements* of a \mathcal{BAT} over a fixed action signature Σ . Variables in these statements are universally quantified.

Definition 2 (Statements of a BAT)

- A *dynamic causal law* is an expression of the form

$$\text{occurs}(a) \text{ causes } f(\bar{x}) = o \text{ if } \text{instance}(a, c), \text{cond} \quad (2)$$

where a and o are variables or object constants, f is a basic fluent, c is the sort *actions* or a subsort of it, and cond is a collection of literals. The law says that *an occurrence of an action a of the sort c in a state satisfying property cond causes the value of $f(\bar{x})$ to become o in any resulting state*.

- A *state constraint* is an expression of the form

$$f(\bar{x}) = o \text{ if } \text{cond} \quad (3)$$

where o is a variable or an object constant, f is any function except a defined function, and cond is a collection of literals. The law says that *the value of $f(\bar{x})$ in any state satisfying condition cond must be o* . Additionally, $f(\bar{x}) = o$ can also be replaced by the object constant *false*, in which case the law says that *there is no state satisfying condition cond* .

- The *definition* of a defined function p is an expression of the form

$$\begin{aligned} p(\bar{t}_1) & \text{ if } \text{cond}_1 \\ \dots & \\ p(\bar{t}_k) & \text{ if } \text{cond}_k \end{aligned} \quad (4)$$

where \bar{t} s are sequences of terms, and $\text{cond}_1, \dots, \text{cond}_k$ are collections of literals. Moreover, if p is a static then $\text{cond}_1, \dots, \text{cond}_k$ cannot contain fluent literals. Statements of the definition will be often referred to as its *clauses*. The statement says that for every \bar{Y} , $p(\bar{Y})$ is true in a state σ iff there is $1 \leq m \leq k$ such that statements cond_m and $\bar{t}_m = \bar{Y}$ are true in σ .

- An *executability condition for actions* is an expression of the form

$$\mathbf{impossible} \text{ occurs}(a) \text{ if } \text{instance}(a, c), \text{cond} \tag{5}$$

where a is a variable or an object constant, c is the sort *actions* or a subsort of it, and cond is a collection of literals and expressions of the form $\text{occurs}(t)$ or $\neg\text{occurs}(t)$ where t is a variable or an object constant of the sort *actions*. The law says that *an occurrence of an action a of the sort c is impossible when condition cond holds.*

Dynamic causal laws and constraints will be sometimes referred to as *causal laws*. We use the term *head* to refer to l in (2) and (3), and to any of the $p(\bar{t}_i)$, $1 \leq i \leq k$, in (4). We call *body* the expression to the right of the keyword **if** in statements (2), (3), (5), or in any of the statements of (4). Statements not containing variables will be referred to as *ground*.

Definition 3 (Basic action theory – BAT)

A *BAT* is a pair consisting of an action signature Σ and a collection T of statements over Σ (called *axioms* of the theory) such that:

- If f is a basic fluent then
 - T contains a state constraint:

$$\text{dom}_f(X_0, \dots, X_n) \text{ if } f(X_0, \dots, X_n) = Y \tag{6}$$

- No dynamic causal law of T contains an atom formed by dom_f in the head.

- If f is a defined fluent, a static, or an attribute then T contains the definition:

$$\text{dom}_f(X_0, \dots, X_n) \text{ if } f(X_0, \dots, X_n) = Y \tag{7}$$

- T contains definitions of special statics of the hierarchy given in terms of functions *is_a* and *link*:

$$\begin{aligned} &\text{instance}(O, C) \text{ if } \text{is_a}(O, C) \\ &\text{instance}(O, C_2) \text{ if } \text{instance}(O, C_1), \text{link}(C_1, C_2) \\ &\text{has_child}(C_2) \text{ if } \text{link}(C_1, C_2) \\ &\text{has_parent}(C_1) \text{ if } \text{link}(C_1, C_2) \\ &\text{source}(C) \text{ if } \neg\text{has_child}(C) \\ &\text{sink}(C) \text{ if } \neg\text{has_parent}(C) \\ &\text{subsort}(C_1, C_2) \text{ if } \text{link}(C_1, C_2) \\ &\text{subsort}(C_1, C_2) \text{ if } \text{link}(C_1, C), \text{subsort}(C, C_2) \end{aligned} \tag{8}$$

To simplify the notation, in what follows we will often identify a theory with the collection of its axioms. Axioms (6)–(8) above are self-explanatory, with the possible exception of the restriction prohibiting the appearance of dom_f in the head of dynamic causal laws. To understand the latter requirement, it is sufficient to notice that it is not enough to include object O in the domain of basic fluent f – it is also necessary to specify the value of $f(O)$. Otherwise the causal law making $\text{dom}_f(O)$

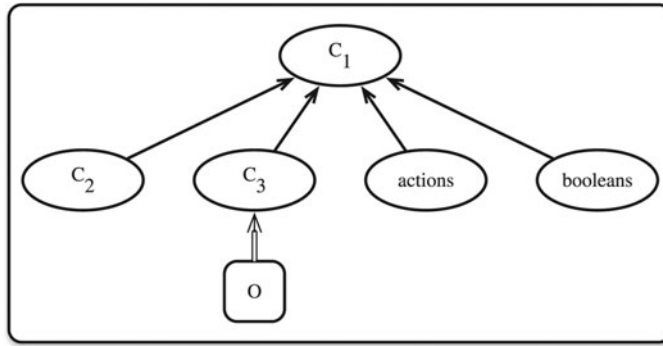


Fig. 1. Hierarchy \mathcal{H}^0 of T^0 .

true would become non-deterministic,² which is not allowed in the current version of \mathcal{ALM} . The presence of a law assigning a value to $f(O)$ makes dynamic causal laws with dom_f in the head unnecessary. It is however useful to allow dynamic causal laws with $\neg dom_f(O)$ in the head as a simple way of removing O from the domain of f .

The following is an example of a BAT .

Example 1 (A basic action theory T^0)

Let us consider an action signature Σ^0 with three sorts, c_1, c_2 , and c_3 , the special sorts *universe* and *actions*, and the pre-defined sort *booleans*, organized in a hierarchy \mathcal{H}^0 in which *universe* is the parent of c_1 , c_1 is the parent of $c_2, c_3, actions$, and *booleans*, and object constant o is of sort c_3 ; attributes $attr_1, attr_2 : actions \rightarrow c_3$; basic fluents $f, g : c_2 \rightarrow c_3$; and special functions like $link, is_a, dom_f, dom_g$. The hierarchy \mathcal{H}^0 can be seen in Figure 1, but we omitted from the picture the sort *universe* whose only child is c_1 .

The BAT T^0 over Σ^0 consists of the causal laws

$$\begin{aligned}
 &occurs(A) \textbf{ causes } f(X) = Y \textbf{ if } \begin{aligned} &instance(A, actions), \\ &attr_1(A) = Y, \\ &g(X) = o \end{aligned} \\
 &occurs(A) \textbf{ causes } \neg dom_f(X) \textbf{ if } \begin{aligned} &instance(A, actions), \\ &attr_2(A) = o \end{aligned} \\
 &false \textbf{ if } \begin{aligned} &\neg dom_g(X), \\ &instance(X, c_2). \end{aligned}
 \end{aligned}$$

The third axiom requires function g to be total.

In addition, T^0 contains standard BAT axioms:

State constraints for the basic fluents:

$$\begin{aligned}
 &dom_f(X) \textbf{ if } f(X) = Y \\
 &dom_g(X) \textbf{ if } g(X) = Y
 \end{aligned}$$

² To see why, consider, for instance, a basic fluent f declared as $f : \{0, 1\} \rightarrow \{0, 1\}$ and a dynamic causal law “ $occurs(a) \textbf{ causes } dom_f(1)$.” Intuitively, the axiom says that after a is executed $f(1)$ must be defined, i.e., $f(1) = 0$ or $f(1) = 1$, which is non-deterministic.

Definitions for the domains of attributes:

$dom_{attr_1}(X)$ if $attr_1(X) = Y$

$dom_{attr_2}(X)$ if $attr_2(X) = Y$

and the collection of axioms from (8).

2.3 Semantics of $\mathcal{BAT}s$

Intuitively, a \mathcal{BAT} T defines the collection of discrete dynamic systems satisfying its axioms. The semantics of T will describe such systems by specifying their transition diagrams, often referred to as *models* of T . Nodes of a transition diagram represent possible states of the dynamic system; arcs of the diagram are labeled by actions. A transition $\langle \sigma_0, a, \sigma_1 \rangle$ says that the execution of action a in state σ_0 may take the system to state σ_1 .

A state of the diagram will be defined by the universe – a collection of objects of the sorts of T , and by a physically possible assignment of values to T 's functions. Moreover, we assume that the sorted universe and the values of statics and attributes are the same in all states, i.e., states only differ by the values of fluents.

To make this precise, it is convenient to partition an interpretation \mathcal{I} of an action signature Σ into two parts: *fluent part* consisting of the universe of \mathcal{I} and the restriction of \mathcal{I} on the sets of fluents, and *static part* consisting of the same universe and the restriction of \mathcal{I} on the remaining elements of the signature. Sometimes, we will refer to the latter as a *static interpretation* of Σ .

We also need the following notation: Given an action signature Σ and a collection U of strings in some fixed alphabet, we denote by Σ_U the signature obtained from Σ by expanding its set of object constants by elements of U , which we assume to be of sort *universe*.

Definition 4 (Pre-model)

Let T be a \mathcal{BAT} with signature Σ and U be a collection of strings in some fixed alphabet. A static interpretation \mathcal{M} of Σ_U is called a *pre-model* of T (with the universe U) if $\mathcal{M}(\text{universe}) = U$ and for every object constant o of Σ_U that is not an object constant of Σ , $\mathcal{M}(o) = o$.

Given a pre-model \mathcal{M} with the universe U , we will often denote signature Σ_U by $\Sigma_{\mathcal{M}}$.

To illustrate this notion, let us consider a pre-model of theory T from Example 1:

Example 2 (A pre-model of basic action theory T^0)

To define a pre-model of \mathcal{BAT} T^0 from Example 1, let us consider a static interpretation \mathcal{M} with the universe $U_{\mathcal{M}} = \{x, y, z, a, b, \text{true}, \text{false}\}$ such that:

$\mathcal{M}(\text{universe}) = \mathcal{M}(c_1) = \{x, y, z, a, b, \text{true}, \text{false}\};$

$\mathcal{M}(c_2) = \{x\};$

$\mathcal{M}(c_3) = \{y, z\},$

$\mathcal{M}(\text{actions}) = \{a, b\};$

$\mathcal{M}(o) = \{y\}$; and
 $\mathcal{M}(\text{attr}_1)(a) = \mathcal{M}(\text{attr}_2)(b) = y$.

In addition: every symbol from $U_{\mathcal{M}}$ is added to Σ_U^0 and mapped into itself; $\text{dom}_{\text{attr}_1} = \{a\}$, $\text{dom}_{\text{attr}_2} = \{b\}$; the interpretation of special function *link* is determined by the hierarchy from Figure 1; the interpretation of *is_a* is extracted from the interpretation of the hierarchy's sorts.

Clearly, \mathcal{M} satisfies the conditions in Definition 4 and hence is a pre-model of T^0 .

A pre-model \mathcal{M} of T uniquely defines a model $T_{\mathcal{M}}$ of T if such a model exists. The definition of $T_{\mathcal{M}}$ will be given in two steps: first, we define $T_{\mathcal{M}}$'s states and then its transitions.

Intuitively, if theory T does not contain definitions, then a *state* of $T_{\mathcal{M}}$ is an interpretation \mathcal{I} with static part \mathcal{M} that satisfies the state constraints of T . The situation is less simple for theories containing definitions (especially recursive ones). Similar to the case of \mathcal{AL} , the definition of a state will be given using logic programs under the answer set semantics; specifically, we will use logic programs with non-Herbrand partial functions in the language $\text{ASP}\{f\}$ (Balduccini 2013).³

Let \mathcal{M} be a pre-model of action theory T .

Program $S_{\mathcal{M}}$:

By $S_{\mathcal{M}}$, we denote a logic program that consists of:

- (a) rules obtained from the state constraints and definitions of T by replacing variables with properly typed object constants of $\Sigma_{\mathcal{M}}$, replacing object constants with their corresponding interpretations in \mathcal{M} , removing the constant *false* from the head of state constraints, and replacing the keyword **if** with \leftarrow ,
- (b) the Closed World Assumption:

$$\neg d(t_0, \dots, t_n) \leftarrow \text{not } d(t_0, \dots, t_n)$$

for every defined function $d : c_0 \times \dots \times c_n \rightarrow \text{booleans}$ and $t_i \in \mathcal{M}(c_i)$, $0 \leq i \leq n$.

end of $S_{\mathcal{M}}$:

Finally, we define a program $S_{\mathcal{I}}$ used in the definition of states of the transition diagram defined by \mathcal{M} .

Program $S_{\mathcal{I}}$:

For every interpretation \mathcal{I} of Σ with static part \mathcal{M} , by $S_{\mathcal{I}}$ we denote the logic program obtained by adding to $S_{\mathcal{M}}$ the set of atoms obtained from \mathcal{I} by removing the defined atoms.

end of $S_{\mathcal{I}}$

Definition 5 (State)

Let \mathcal{M} be a pre-model of a \mathcal{BAT} theory T . An interpretation σ with static part \mathcal{M} is a *state* of the transition diagram $T_{\mathcal{M}}$ defined by \mathcal{M} if σ is *the only* answer set of S_{σ} .

³ Other approaches for introducing non-Herbrand functions in ASP can be seen, for instance, in (Cabalar 2011; Lifschitz 2012; Bartholomew and Lee 2013).

Notice that σ is *not* a state if S_σ has multiple answer sets, a situation that would only occur when the value of some defined function is not completely determined by the values of basic functions. We will return to this issue later, in Section 4.2.

Example 3 (States of the diagram)

Let \mathcal{M} be the pre-model of theory T^0 from Example 2. The program $S_{\mathcal{M}}$ for this \mathcal{M} looks as follows:

$$\begin{aligned} &\leftarrow \neg dom_g(x), instance(x, c_2) \\ dom_f(x) &\leftarrow f(x) = y \\ dom_f(x) &\leftarrow f(x) = z \\ dom_g(x) &\leftarrow g(x) = y \\ dom_g(x) &\leftarrow g(x) = z \\ \\ dom_{attr_1}(a) &\leftarrow attr_1(a) = y \\ dom_{attr_1}(a) &\leftarrow attr_1(a) = z \\ dom_{attr_2}(a) &\leftarrow attr_2(a) = y \\ dom_{attr_2}(a) &\leftarrow attr_2(a) = z \\ dom_{attr_1}(b) &\leftarrow attr_1(b) = y \\ dom_{attr_1}(b) &\leftarrow attr_1(b) = z \\ dom_{attr_2}(b) &\leftarrow attr_2(b) = y \\ dom_{attr_2}(b) &\leftarrow attr_2(b) = z \end{aligned}$$

and the Closed World Assumptions for the special functions. Recall that according to the definition of an interpretation of a sorted signature, for every $x \in |\mathcal{I}|$, $\mathcal{I}(is_a)(x, c)$ is true iff c is a source node of the sort hierarchy and $\mathcal{I}(x) \in \mathcal{I}(c)$, and for every object o and sort c , $\mathcal{I}(is_a)(\mathcal{I}(o), c)$ is true iff $\langle o, c \rangle$ is a link in our hierarchy. This, together with the condition on the interpretation of *link* guarantees that every state of $T_{\mathcal{M}}$ contains atoms $is_a(x, c_2)$, $is_a(y, c_3)$, and other atoms formed by *is_a* and *link* that define our hierarchy. The collection of these atoms together with the closed world assumptions for *is_a*, *link* and the other defined statics uniquely determine their values. It is easy to check that every state of \mathcal{M} contains literals formed by these special fluents. Every state of $T_{\mathcal{M}}$ also contains $attr_1(a) = y$, $attr_2(b) = y$, and $dom_g(x)$. Overall, $T_{\mathcal{M}}$ has the following six states (for each state, we only show non-special fluents):

$$\begin{aligned} \sigma_1 &= \{f(x) = y, g(x) = y\} & \sigma_2 &= \{f(x) = z, g(x) = y\} \\ \sigma_3 &= \{f(x) = y, g(x) = z\} & \sigma_4 &= \{f(x) = z, g(x) = z\} \\ \sigma_5 &= \{g(x) = y\} & \sigma_6 &= \{g(x) = z\}. \end{aligned}$$

In addition, states $\sigma_1, \sigma_2, \sigma_3$, and σ_4 contain $dom_f(x)$ while states σ_5 and σ_6 , in which f is undefined on x , contain $\neg dom_f(x)$.

To define transitions of the diagram that corresponds to a pre-model \mathcal{M} with the universe U , we construct a logic program $P_{\mathcal{M}}$ whose signature is obtained from the signature of program $S_{\mathcal{M}}$ defined above by

- adding a new sort, *step*, ranging over 0 and 1;
- replacing every fluent $f : c_0 \times \dots \times c_n \rightarrow c$ by function $f : c_0 \times \dots \times c_n \times step \rightarrow c$;

- adding a function symbol $occurs : actions \times step \rightarrow booleans$.

Program $P_{\mathcal{M}}$:

Program $P_{\mathcal{M}}$ is obtained from a theory T and pre-model \mathcal{M} by

- replacing variables by properly typed object constants of $\Sigma_{\mathcal{M}}$;
- replacing object constants by their corresponding interpretations in \mathcal{M} ;
- removing the object constant $false$ from the head of state constraints;
- replacing every occurrence of a fluent term $f(\bar{t})$ in the head of a dynamic causal law by $f(\bar{t}, I + 1)$;
- replacing every other occurrence of a fluent term $f(\bar{t})$ by $f(\bar{t}, I)$;
- removing “ $occurs(a)$ causes” from every dynamic causal law and adding $occurs(a)$ to the body;
- replacing “**impossible** $occurs(a)$ ” in every executability condition by $\neg occurs(a)$;
- replacing $occurs(a)$ by $occurs(a, I)$ and $\neg occurs(a)$ by $\neg occurs(a, I)$;
- replacing the keyword **if** by \leftarrow ;
- adding the Closed World Assumption:

$$\neg d(t_0, \dots, t_n, I) \leftarrow \text{not } d(t_0, \dots, t_n, I)$$

for every defined fluent $d : c_0 \times \dots \times c_n \rightarrow booleans$ and $t_i \in \mathcal{M}(c_i)$, $0 \leq i \leq n$;

- adding the rule:

$$\neg f(t_0, \dots, t_n) \leftarrow \text{not } f(t_0, \dots, t_n)$$

for every defined static of the form $f : c_0 \times \dots \times c_n \rightarrow booleans$ and $t_i \in \mathcal{M}(c_i)$, $0 \leq i \leq n$;

- adding the Inertia Axiom:

$$\begin{aligned} dom_f(t_0, \dots, t_n, I + 1) &\leftarrow dom_f(t_0, \dots, t_n, I), \\ &\text{not } \neg dom_f(t_0, \dots, t_n, I + 1) \\ \neg dom_f(t_0, \dots, t_n, I + 1) &\leftarrow \neg dom_f(t_0, \dots, t_n, I), \\ &\text{not } dom_f(t_0, \dots, t_n, I + 1) \end{aligned}$$

for every basic fluent $dom_f : c_0 \times \dots \times c_n \rightarrow booleans$, and $t_i \in \mathcal{M}(c_i)$, $0 \leq i \leq n$;

- adding the Inertia Axiom:

$$\begin{aligned} f(t_0, \dots, t_n, I + 1) = t &\leftarrow dom_f(t_0, \dots, t_n, I + 1), \\ &f(t_0, \dots, t_n, I) = t, \\ &\text{not } f(t_0, \dots, t_n, I + 1) \neq t \end{aligned}$$

for every basic fluent $f : c_0 \times \dots \times c_n \rightarrow c$ not formed by dom , and $t_i \in \mathcal{M}(c_i)$, $0 \leq i \leq n$, and $t \in \mathcal{M}(c)$.

end of $P_{\mathcal{M}}$

Note that the last axiom is a modification of the standard logic programming version of the Inertia Axiom (see, for instance, (Gelfond and Kahl 2014)), which is stated for total (boolean) functions. The main difference is the addition of the domain statements in the body. The inertia axiom for the function dom_f is of the standard form.

Program $P(\mathcal{M}, \sigma_0, a)$: Let σ_0 be a state of the transition diagram defined by a pre-model \mathcal{M} , and let $a \subseteq \mathcal{M}(\text{actions})$. By $P(\mathcal{M}, \sigma_0, a)$, we denote the logic program formed by adding to $P_{\mathcal{M}}$ the set of atoms obtained from σ_0 by replacing every fluent atom $f(t_0, \dots, t_n) = t$ by $f(t_0, \dots, t_n, 0) = t$ and adding the set of atoms $\{\text{occurs}(x, 0) : x \in a\}$.

end of $P(\mathcal{M}, \sigma_0, a)$

Definition 6 (Transition)

Let σ_0 and σ_1 be states of the transition diagram defined by a pre-model \mathcal{M} and let $a \subseteq \mathcal{M}(\text{actions})$. The triple $\langle \sigma_0, a, \sigma_1 \rangle$ is a *transition* of the transition diagram defined by a pre-model \mathcal{M} of a \mathcal{BAT} theory T if program $P(\mathcal{M}, \sigma_0, a)$ has an answer set A such that $f(t_0, \dots, t_n) = t \in \sigma_1$ iff

- f is an attribute or a static and $f(t_0, \dots, t_n) = t \in A$, or
- f is a fluent and $f(t_0, \dots, t_n, 1) = t \in A$.

Definition 7 (Model)

A transition diagram $T_{\mathcal{M}}$ defined by a pre-model \mathcal{M} of a \mathcal{BAT} T is called a *model* of T if it has a non-empty collection of states.

The following example illustrates the definition.

Example 4 (A model of basic action theory T^0)

To define a model of theory T^0 from Example 1, let us consider the pre-model \mathcal{M} from Example 2. States of the diagram defined by this pre-model were given in Example 3. To define the transitions of the model defined by \mathcal{M} , we use Definition 6. Let us illustrate this by showing that a triple $\langle \sigma_1, b, \sigma_5 \rangle$ is a transition. To do that we need first to construct a program $P(\mathcal{M}, \sigma_1, b)$ (we are only showing rules relevant to our argument):

$$[1] \quad f(x, 1) = y \leftarrow \begin{array}{l} \text{instance}(b, \text{actions}), \\ \text{occurs}(b, 0), \\ \text{attr}_1(b) = y, \\ g(x, 0) = y. \end{array}$$

$$[2] \quad \neg \text{dom}_f(x, 1) \leftarrow \begin{array}{l} \text{instance}(b, \text{actions}), \\ \text{occurs}(b, 0), \\ \text{attr}_2(b) = y. \end{array}$$

$$[3] \quad \begin{array}{l} \text{dom}_f(x, 0) \leftarrow f(x, 0) = y. \\ \text{dom}_f(x, 1) \leftarrow f(x, 1) = y. \\ \text{dom}_g(x, 0) \leftarrow g(x, 0) = y. \\ \text{dom}_g(x, 1) \leftarrow g(x, 1) = y. \end{array}$$

$$[4] \quad f(x, 1) = y \leftarrow \begin{array}{l} \text{dom}_f(x, 1), \\ f(x, 0) = y, \\ \text{not } f(x, 1) \neq y. \end{array}$$

$$g(x, 1) = y \leftarrow \begin{array}{l} \text{dom}_g(x, 1), \\ g(x, 0) = y, \\ \text{not } g(x, 1) \neq y. \end{array}$$

$$[5] \quad \begin{array}{l} \text{dom}_f(x, 1) \leftarrow \text{dom}_f(x, 0), \\ \text{not } \neg\text{dom}_f(x, 1). \end{array}$$

$$\begin{array}{l} \text{dom}_g(x, 1) \leftarrow \text{dom}_g(x, 0), \\ \text{not } \neg\text{dom}_g(x, 1). \end{array}$$

$$[6] \quad \begin{array}{l} f(x, 0) = y. \\ g(x, 0) = y. \\ \text{occurs}(b, 0). \end{array}$$

It is easy to see that the program has a unique answer set, say, S . Since $\sigma_5 = \{g(x) = y\}$, we need to show that the only fluent atom with the step parameter 1 belonging to S is $g(x, 1) = y$. By the second rule from group [5], $\text{dom}_g(x, 1) \in S$. By the second rule of [4], we have that $g(x, 1) = y \in S$. As expected, function g maintains its value by inertia. The situation is different for f . By rule [2], we have that $\neg\text{dom}_f(x, 1) \in S$ and hence neither rule [5] nor [4] for f are applicable. Rule [1] is also not applicable since attr_1 is not defined for b . Therefore, the state defined by S is exactly $\sigma_5 = \{g(x) = y\}$. (Note that the argument would not be possible if we were to use the traditional version of the Inertia Axiom. The modification related to the treatment of dom presented in axioms [4] and [5] is essential.)

Using the same method, one can easily verify that triples $\langle \sigma_2, a, \sigma_1 \rangle$, $\langle \sigma_5, a, \sigma_1 \rangle$, $\langle \sigma_5, b, \sigma_5 \rangle$, etc. are transitions of the transition diagram defined by \mathcal{M} .

2.4 Entailment relation

Let us consider a fixed action theory T with action signature Σ , and define an entailment relation between T and statements of Σ .

Let \mathcal{I} be an interpretation of Σ . A *ground instance* of a statement α of Σ with respect to \mathcal{I} is a statement obtained by replacing variables of α by properly typed object constants in $\Sigma_{\mathcal{I}}$ and replacing object constants of α by their interpretations in \mathcal{I} .

Now, let us consider a model $T_{\mathcal{M}}$ of a *BAT* T defined by a pre-model \mathcal{M} with the universe U and let σ be a state of $T_{\mathcal{M}}$.

Definition 8 (Satisfiability relation for ground statements of a BAT)

- A state σ of $T_{\mathcal{M}}$ satisfies a ground state constraint α if σ contains the head of α whenever it contains its body.
- A state σ of $T_{\mathcal{M}}$ satisfies a ground definition α if σ contains the head of a clause in α iff α contains a clause with the same head and the body belonging to σ .
- A transition $\langle \sigma_0, a, \sigma_1 \rangle$ of $T_{\mathcal{M}}$ satisfies a ground dynamic causal law α that starts with the expression “*occurs*(e) **causes**” if a contains action e and σ_1 contains the head of α whenever σ_0 contains its body.
- A transition $\langle \sigma_0, a, \sigma_1 \rangle$ of $T_{\mathcal{M}}$ satisfies a ground executability condition α that starts with the expression “**impossible** e ” if either (1) a does not contain e or (2) the body of α contains:

- a ground literal l such that $l \notin \sigma_0$, or
- an expression “ $occurs(e_1)$ ” such that $e_1 \notin a$, or
- an expression “ $\neg occurs(e_2)$ ” such that $e_2 \in a$.

Definition 9 (Satisfiability relation for arbitrary statements of a BAT)

Let $T_{\mathcal{M}}$ be a model of a BAT T defined by a pre-model \mathcal{M} with the universe U .

- $T_{\mathcal{M}}$ satisfies a constraint α over signature Σ of T if every state of $T_{\mathcal{M}}$ satisfies all ground instances of α with respect to U . Similarly for definitions.
- $T_{\mathcal{M}}$ satisfies a dynamic causal law α over signature Σ of T if every transition of $T_{\mathcal{M}}$ satisfies all ground instances of α with respect to U . Similarly for executability conditions.

Definition 10 (Entailment)

A statement α is entailed by a theory T ($T \models \alpha$) if α is true in every model of T .

Having the notion of entailment allows us to investigate the relationship between causal laws. For instance, we can show that

$$\{occurs(A) \text{ causes } f \text{ if } p, q; occurs(A) \text{ causes } f \text{ if } \neg p\} \models occurs(A) \text{ causes } f \text{ if } q$$

$$\{occurs(A) \text{ causes } f \text{ if } p, q; q \text{ if } p\} \models occurs(A) \text{ causes } f \text{ if } p$$

etc. Our notion of entailment is somewhat similar to the notion of *subsumption* from (Eiter *et al.* 2010) – a relation between an action description and a query (including queries having the form of causal laws and executability conditions). Our entailment relation can be viewed as a generalization of subsumption from system descriptions to theories. It allows variables and, unlike that of subsumption, is defined in terms of multiple transition diagrams specified by the theory. There are also related formalisms that allow entailment of causal laws and executability conditions (see, for instance (Turner 1999) and (Giunchiglia *et al.* 2004b)). There are many interesting problems related to the \mathcal{ALM} entailment, including that of finding a sound and complete set of inference rules for it. We hope to address these problems in our future work.

3 Language \mathcal{ALM}

In this section, we use examples to introduce the syntax of theories and system descriptions of \mathcal{ALM} and define their semantics. (The full grammar for the language can be seen in the online appendix A, pp. 1–3.).

We begin with describing *unimodule* system descriptions, i.e., system descriptions whose theories consist of exactly one module.

3.1 Unimodule system descriptions

We start with a comparatively simple problem of formalizing the domain described by the following story:

Example 5 (A travel domain)

Consider a travel domain in which there are two *agents*, Bob and John, and three locations, New York, Paris, and Rome. Bob and John can move from one location to another if the locations are connected.

If we were to represent this knowledge in \mathcal{AL} , we would start with identifying objects of the domain including actions such as, say, $go(bob, paris, rome)$ and write \mathcal{AL} axioms describing the relationships between these objects. The use of \mathcal{ALM} suggests a very different methodology.

Methodology of Describing Dynamic Domains in \mathcal{ALM} :

- (1) Determine what *sorts* of objects are relevant to the domain of discourse and how these sorts can be organized into an inheritance hierarchy.
- (2) Use \mathcal{ALM} to describe the \mathcal{BAT} for this type of domains. This should be done in two steps:
 - Describe the action signature of our abstraction by declaring sorts (together with their attributes and the inheritance hierarchy), basic and defined statics and fluents. (Notice that this signature normally will not contain particular objects of our story. It would have no mention of Bob, Paris, etc. However, the signature may include some object constants pertinent to the *general* domain of the story – see for instance the Monkey and Banana Problem in Section 4.1.)
 - Use this action signature to formulate axioms of the theory.
- (3) Populate sorts of your hierarchy with objects relevant to your story and describe these objects and their sort membership in \mathcal{ALM} .

As is the case with other problem-solving methodologies, we begin by choosing a proper level of abstraction for our example. Since the example is used for illustrative purposes, we opted for using the following simple abstraction:

Our domains will contain *things* and *discrete points* in space. Certain things, called *agents*, will be able to move from one point to another if the two points are connected. We are interested in the *relations between points* and the *locations of things*, including changes of these locations caused by a sequence of given *moves*. (Note that our abstraction does not allow a location to be a part of another location, e.g., we will not be able to express that Paris is located in France. It ignores the means of transportation, the possibility that locations may have restrictions on the number of things they can contain, etc.)

Accordingly, our \mathcal{BAT} containing commonsense knowledge about motion formulated in these terms will include sorts *things*, *agents*, *points*, and *move*, together with special sorts *universe* and *actions*, which belong to every action signature.

We call this \mathcal{BAT} T_{bm} . The sorts of T_{bm} will be organized in a hierarchy depicted in Figure 2.

Our next step is to describe T_{bm} in \mathcal{ALM} .




Fig. 2. Sort hierarchy for T_{bm} .

Example 6 (Motion theory in \mathcal{ALM})

The description of a theory in \mathcal{ALM} starts with the keyword **theory** and is followed by a collection of modules. Our theory, called *basic_motion*, consists of only one module *moving*

theory *basic_motion*
module *moving*
 ⟨*module body*⟩

where ⟨*module body*⟩ stands for the declarations of sorts, functions, and axioms of the theory. We assume that *things*, *points*, and *agents* have no attributes, while actions from the sort *move* may come with attribute *actor* indicating the agent involved in the action, and attributes *origin* and *destination* (abbreviated as *dest*) describing the locations of the actor before and after the execution of the action. Syntactically, all this information is specified as:

sort declarations

points, things :: *universe*
agents :: *things*
move :: *actions*

attributes

actor : *agents*
origin : *points*
dest : *points*

The construct `::` is called *specialization* and corresponds to the links of the sort hierarchy; for instance, the link from *agents* to *things* in Figure 2 is recorded by the statement *agents* :: *things*. Multiple links going into the same sort can be recorded by a single statement, as in *points, things* :: *universe*. Note that the special sorts *universe* and *actions* do not have to be declared. In case of a sort hierarchy with multiple links from *c* to *pc₁, ..., pc_k* we will use a specialization statement of the form *c* :: *pc₁, ..., pc_k*. In describing the attributes of actions of the sort *move* we use a shorthand. Attributes of *move* are functions defined on elements of the sort *move*, which means that the definition of, say, attribute *actor* should be written as *actor* : *move* → *agents*. After some deliberation however, we decided to allow to write it simply as *actor* : *agents*. The same agreement holds for attributes

with a larger number of parameters; an attribute of a sort c that has the form $attr_name : c \times c_0 \times \dots \times c_n \rightarrow c_{n+1}$ can be written as $attr_name : c_0 \times \dots \times c_n \rightarrow c_{n+1}$. This completes the description of the syntactic representation of our sort hierarchy in \mathcal{ALM} .

The next step is to syntactically describe functions in the signature. One of the functions mentioned in our informal description specifies whether two points are connected or not. Let us call it *connected*. In general, the value of *connected* can be changed by actions (airports can be closed, roads blocked, etc.) and hence we define *connected* to be a *basic fluent*. In some scenarios, the property *connected* will be a symmetric relation but not in others; similarly, it may be a transitive relation or not. To allow for elaboration tolerance, we introduce two basic static functions, *symmetric_connectivity* and *transitive_connectivity* to characterize the property *connected*. The other function relevant to our domain maps things into points at which they are located. Let us call it *loc_in*. The value of the function can be changed by actions of our domain, hence it is a *fluent*. It is not defined in terms of other functions, thus it is a *basic fluent*. It is also a *total* function, as we assume that the location of every *thing* is defined in every state. In \mathcal{ALM} , these functions are syntactically declared as:

function declarations

statics

basic

symmetric_connectivity : booleans

transitive_connectivity : booleans

fluents

basic

connected : points \times points \rightarrow booleans

total *loc_in* : things \rightarrow points

In this example, the keywords **function declarations** are followed by the lists of statics and fluents. Elements from each list are divided into basic and defined with each total function in the list preceded by the keyword **total**. Naturally, the declaration of a sort, static, or fluent in a module should be unique.

This concludes our description of action signature of T_{bm} ⁴.

Now, we are ready to define the collection of axioms of T_{bm} . In \mathcal{ALM} , we precede this collection by the keyword **axioms**. Each axiom will be ended by a period (.), as in:

axioms

occurs(X) **causes** *loc_in*(A) = D **if** *instance*(X , *move*),
actor(X) = A ,
dest(X) = D .

⁴ The description does not mention object constants, which can be declared in \mathcal{ALM} by statements $o : c$ and $r(c_1, \dots, c_n) : c$. The first statement defines object constant o of sort c ; the second defines the collection of object constants of the form $r(x_1, \dots, x_n)$ where x_1, \dots, x_n are object constants from sorts c_1, \dots, c_n . Example of the latter can be found in module *climbing* of Monkey and Banana representation from Section 4.1.

$connected(X, X).$
 $connected(X, Y) \quad \mathbf{if} \quad connected(Y, X),$
 $\quad \quad \quad symmetric_connectivity.$
 $\neg connected(X, Y) \quad \mathbf{if} \quad \neg connected(Y, X),$
 $\quad \quad \quad symmetric_connectivity.$
 $connected(X, Z) \quad \mathbf{if} \quad connected(X, Y),$
 $\quad \quad \quad connected(Y, Z),$
 $\quad \quad \quad transitive_connectivity.$
 $\mathbf{impossible} \quad occurs(X) \quad \mathbf{if} \quad instance(X, move),$
 $\quad \quad \quad actor(X) = A,$
 $\quad \quad \quad loc_in(A) \neq origin(X).$
 $\mathbf{impossible} \quad occurs(X) \quad \mathbf{if} \quad instance(X, move),$
 $\quad \quad \quad actor(X) = A,$
 $\quad \quad \quad loc_in(A) = dest(X).$
 $\mathbf{impossible} \quad occurs(X) \quad \mathbf{if} \quad instance(X, move),$
 $\quad \quad \quad actor(X) = A,$
 $\quad \quad \quad loc_in(A) = O,$
 $\quad \quad \quad dest(X) = D,$
 $\quad \quad \quad \neg connected(O, D).$

The keyword **total** in the declaration of the basic fluent loc_in stands for the axiom

$false \quad \mathbf{if} \quad \neg dom_{loc_in}(X).$

that would otherwise have to be included among the axioms above. In general, the keyword **total** included in the declaration of a function $f : c_0 \times \dots \times c_n \rightarrow c$ stands for the axiom

$false \quad \mathbf{if} \quad \neg dom_f(X_0, \dots, X_n).$

This completes our description of the $BAT T_{bm}$ in \mathcal{ALM} .

Note that the *semantics of the unimodule \mathcal{ALM} theory basic_motion* is given by the $BAT T_{bm}$ defined by it. In the following sections, we will present other examples of basic action theories and their interpretations represented in \mathcal{ALM} . (Whenever possible we will make no distinction between these theories and their \mathcal{ALM} representations.)

As discussed above, a $BAT T$ is used to define the collection of its models – transition diagrams representing dynamic domains with shared ontology and properties. Usually, a knowledge engineer is interested in one such domain, characterized by particular objects, sorts, and values of statics. If the engineer’s knowledge about this domain is complete, the domain will be represented by a unique model of T . Otherwise, there can be several alternative models.

The syntactic construct of \mathcal{ALM} used to define such knowledge is called a *structure* and has the form

structure *name*
 $\langle structure \ body \rangle$

where $\langle \text{structure body} \rangle$ stands for the definition of objects in the hierarchy of T_{bm} and the values of its statics. Let us illustrate the use of this construct by the following example:

Example 7 (ALM's representation of a specific basic motion domain.)

Let us consider the ALM theory *basic_motion* from Example 6, which encodes the BAT T_{bm} , and use ALM to specify the particular basic motion domain from Example 5.

The ALM definition of the structure used to describe this domain starts with the header:

structure *Bob_and_John*

followed by the definition of *agents* and *points*:

instances

bob, john in agents

new_york, paris, rome in points

To specify particular actions of our domain, we expand our list of instances by

$go(X, P_1, P_2)$ **in** *move* **where** $P_1 \neq P_2$

actor = X

origin = P₁

dest = P₂

Note that the last definition describes several instances simultaneously via the use of variables; we call this type of definition an *instance schema*. The instance schema defining $go(X, P_1, P_2)$ stands for the collection of instance definitions:

$go(\text{bob}, \text{new_york}, \text{paris})$ **in** *move*

actor = bob

origin = new_york

dest = paris

...

$go(\text{john}, \text{paris}, \text{rome})$ **in** *move*

actor = john

origin = paris

dest = rome

The condition **where** $P_1 \neq P_2$ ensures that Bob and John do not move to a destination identical to the origin.

The following would also be a valid instance schema:

$go(X, P)$ **in** *move*

actor = X

dest = P

if we were interested only in the destinations of Bob and John's movements, but not in their origins.

In our example, connectivity between points is both symmetric and transitive: This is captured syntactically by the following:⁵

values of statics

symmetric_connectivity.

transitive_connectivity.

This concludes our definition of *Bob_and_John* structure.

To syntactically relate a theory with its structure, we use the construct of \mathcal{ALM} called *system description*. In our case, it will look as follows:

system description *travel*

theory *basic_motion*

module *moving*

<module body>

structure *Bob_and_John*

<structure body>

where *<module body>* and *<structure body>* are defined in Examples 6 and 7. The system description *travel* contains all the information we considered relevant to our particular travel domain. It is not difficult to see that this knowledge is complete and therefore describes exactly one model (i.e., one transition diagram) of *basic_motion*. This is exactly the model we *intended* for our domain. A part of this model can be seen in Figure 3. We only show fluent *loc_in* and assume that in every state of the part of the diagram shown in the picture Paris and Rome are connected to each other, but neither of them is connected to New York; we use shorthands *b*, *j*, *ny*, *p*, and *r* for *bob*, *john*, *new_york*, *paris*, and *rome* respectively; and we only show arcs that are labeled by a single action.

The model is unique because we specified the membership of our objects in the source nodes of the hierarchy. This information is sufficient to uniquely define the universe and the interpretations of all the sorts.

The next example illustrates how incomplete information about a domain can lead to multiple models of the system description of this domain:

Example 8 (System description with multiple models)

Consider a system description *underspecified_hierarchy* consisting of a theory *professors* and a structure *alice*:

system.description *underspecified_hierarchy*

theory *professors*

module *professors*

sort declarations

professor :: *person*

assistant, associate, full :: *professor*

⁵ If a theory contains an object constant *o* then its value, say *y*, can be declared as:

object constants

o = y

If the structure contains no assignment of value to constant *o*, we assume that *o* belongs to the structure's universe and is mapped into itself.




Fig. 3. (Partial) transition diagram for system description *travel*.




Fig. 4. Underspecified hierarchy.

axioms

false if $instance(X, C_1), instance(X, C_2),$
 $link(C_1, professor), link(C_2, professor),$
 $C_1 \neq C_2.$

structure *alice*

instances

alice in professor

The theory describes a simple hierarchy. The structure populates the hierarchy with one member, *Alice* (see Figure 4). Unfortunately, all we know about *Alice* is that she is a professor. It is not difficult to check that this system description has three models. In the first one *Alice* is an assistant professor, in the second she is an associate professor, and in the third one—a full professor.

We hope that these examples gave the reader a sufficient insight in the meaning of unimodule \mathcal{ALM} theories and system descriptions. In general, the semantics of a syntactically correct unimodule theory T of \mathcal{ALM} is given by the unique *BAT* defined by T . Similarly, the semantics of a system description D of \mathcal{ALM} is given

by models of the BAT theory defined by \mathcal{T} and by the set of interpretation defined by the structure of D .

3.2 Organizing knowledge into modules

So far, we only considered very simple \mathcal{ALM} theories consisting of one module. To create theories containing a larger body of knowledge, we need multiple modules organized into a module hierarchy. To illustrate this concept, let us consider an extension of BAT T_{bm} of motion by an additional sort of things called *carriables*, which can be *carried* between connected points by *agents* that are holding them. Recall from Example 6, that we represented the original T_{bm} as an \mathcal{ALM} theory called *basic_motion*, with a unique module *moving*. We will use the name *motion* for the \mathcal{ALM} theory that will specify the extension of T_{bm} . The new theory will contain the *moving* module developed above as well as a new module called *carrying_things*:

theory *motion*

module *moving*

⟨*module body*⟩

module *carrying_things*

⟨*module body*⟩

In addition to sorts, fluents, and axioms from module *moving*, the signature of the new module *carrying_things* will contain two new sorts, *carriables* and *carry*; a new inertial fluent, *holding*; and a defined fluent, *is_held*. Informally, *holding* will be understood as *having in one's hands* and *carry* as *moving while holding*, which will allow us to define *carry* as a special case of *move*.

The dependency of *carrying_things* on *moving* is expressed in \mathcal{ALM} by the syntactic construct **depends on** called *module dependency* as follows:

module *carrying_things*

depends on *moving*

This says that the sorts and functions explicitly declared in *carrying_things* depend on sorts and functions declared in the module *moving*. We say that the declarations of *moving* are *implicit* in module *carrying_things*. We *require all sorts and functions appearing in a module to be either explicitly or implicitly declared in that module*. By means of the module dependency construct, a theory of \mathcal{ALM} can be structured into a *hierarchy of modules*. *The dependency relation of this hierarchy should form a DAG*. Now, we define the body of the new module:

sort declarations

carriables :: *things*

carry :: *move*

attributes

carried_object : *carriables*

Note that since *carry* is defined as a special case of *move*, it automatically inherits the attributes of *move*; hence those attributes do not have to be repeated in the declaration of *carry*. Next, the module contains the declarations of functions:

function declarations**fluents****basic****total** $holding : agents \times things \rightarrow booleans$ **defined** $is_held : things \rightarrow booleans$

and the new axioms:

axioms $loc_in(C) = P$ **if** $holding(T, C),$
 $loc_in(T) = P.$ $loc_in(T) = P$ **if** $holding(T, C),$
 $loc_in(C) = P.$ $is_held(X)$ **if** $holding(T, X).$ **impossible** $occurs(X)$ **if** $instance(X, move),$
 $actor(X) = A,$
 $is_held(A).$ **impossible** $occurs(X)$ **if** $instance(X, carry),$
 $actor(X) = A,$
 $carried_object(X) = C,$
 $\neg holding(A, C).$

The first two axioms say that an agent and an object he is holding have the same location. The next defines fluent $is_held(X)$ – object X is held by someone or something. The first executability condition states that to move an actor should be free (i.e., not held). The second states that it is impossible to carry a thing without holding it.

Structuring a theory of \mathcal{ALM} into a hierarchy of modules has several advantages. First, this supports the step-wise development of a knowledge base by allowing parts of its theory to be developed and tested independently from other parts. Second, it increases the readability of \mathcal{ALM} theories, due to the more manageable size of their modules.⁶ And finally, this approach facilitates the creation of knowledge libraries. Theories containing very general information can be stored in a library and imported from there when constructing system descriptions. For instance, imagine that our *motion* theory is stored in a library called *commonsense_library*. The system description *travel* could then be re-written by importing this theory as follows:

system description *travel***import theory** *motion* **from** *commonsense_library***structure** *Bob_and_John* $\langle structure\ body \rangle$

We hope that these examples gave the reader some insight into the meaning of theories of \mathcal{ALM} that have more than one module. The accurate semantics for such

⁶ For greater readability, we recommend maintaining a balance between a manageable module size and a relatively shallow module dependence hierarchy.

a theory T is given by its *flattening*, i.e., by translating T into the unimodular theory with the same intuitive meaning.

First, we will give the semantics of theories satisfying the semantic conditions given in the following definition, theories that we call *semantically coherent*.

Definition 11 (Semantically coherent theory)

A theory of \mathcal{ALM} is *semantically coherent* if it satisfies the following conditions:

- All sorts and functions appearing in a module of T are (explicitly or implicitly) declared in that module.
- The module hierarchy of T defined by relation “*depends on*” forms a DAG, G . (The nodes of G correspond to modules of T . An arc $\langle M_2, M_1 \rangle$ is in G if and only if module M_2 contains the statement “*depends on* M_1 ”.)
- No two modules of a theory contain different declarations of the same sort or the same function name.

The last condition in Definition 11 can be weakened to allow the use of the same name for a function and its restriction on a smaller sort. This and other similar features however can somewhat distract from the main ideas of \mathcal{ALM} and will not be included in the original version of \mathcal{ALM} .

The flattening $f(T)$ of an \mathcal{ALM} theory T is constructed by the following algorithm:

- (1) Select modules M_1 and M_2 of T such that M_1 contains the statement “*depends on* M_2 ”.
- (2) Replace M_1 and M_2 by the new module M obtained by uniting *depends on* statements, sort declarations, object constant declarations, function declarations, and axioms of M_2 with those of M_1 .
- (3) Remove the statement “*depends on* M_2 ” from M .
- (4) Replace M_1 and M_2 in all the statements of T of the form “*depends on* M_1 ” and “*depends on* M_2 ” by M .
- (5) Repeat until no dependent modules exist.
- (6) Construct a new module with declarations and axioms defined as unions of the corresponding declarations and axioms of the remaining modules.
- (7) Return the resulting unimodule theory $f(T)$.

The second condition in Definition 11 guarantees that the algorithm will terminate. The first and second conditions ensure that the result of the algorithm does not contain the *depends on* statement and that all sorts and functions within module M of step (2) have unique (explicit or implicit) declarations. Thanks to condition three, this property is preserved by step (6) of the algorithm and hence $f(T)$ is indeed a unimodule theory.

As expected, *the semantics of an \mathcal{ALM} theory T with more than one module is given by the semantics of the unimodule theory $f(T)$.*

For illustrative purposes, we give the result of applying the flattening algorithm to the *motion* theory given above:

```
theory flat_motion
module flat_motion
```

sort declarations

points, things :: universe
agents, carriables :: things
move :: actions

attributes

actor : agents
origin : points
dest : points

carry :: move

attributes

carried_object : carriables

function declarations**statics****basic**

symmetric_connectivity : booleans
transitive_connectivity : booleans

fluents**basic**

total *loc_in* : things \rightarrow points
total *holding* : agents \times things \rightarrow booleans
connected : points \times points \rightarrow booleans

defined

is_held : things \rightarrow booleans

axioms

occurs(*X*) **causes** *loc_in*(*A*) = *D* **if** *instance*(*X*, *move*),
actor(*X*) = *A*, *dest*(*X*) = *D*.

connected(*X*, *X*).

connected(*X*, *Y*) **if** *connected*(*Y*, *X*), *symmetric_connectivity*.

\neg *connected*(*X*, *Y*) **if** \neg *connected*(*Y*, *X*), *symmetric_connectivity*.

connected(*X*, *Z*) **if** *connected*(*X*, *Y*), *connected*(*Y*, *Z*),
transitive_connectivity.

loc_in(*C*) = *P* **if** *holding*(*T*, *C*), *loc_in*(*T*) = *P*.

loc_in(*T*) = *P* **if** *holding*(*T*, *C*), *loc_in*(*C*) = *P*.

is_held(*C*) **if** *holding*(*T*, *C*).

impossible *occurs*(*X*) **if** *instance*(*X*, *move*), *actor*(*X*) = *A*,
origin(*X*) \neq *loc_in*(*A*).

impossible *occurs*(*X*) **if** *instance*(*X*, *move*), *actor*(*X*) = *A*,
dest(*X*) = *loc_in*(*A*).

impossible *occurs*(*X*) **if** *instance*(*X*, *move*), *actor*(*X*) = *A*,
loc_in(*A*) = *O*, *dest*(*X*) = *D*,
 \neg *connected*(*O*, *D*).

impossible *occurs*(*X*) **if** *instance*(*X*, *move*),
actor(*X*) = *A*, *is_held*(*A*).

impossible $occurs(X)$ **if** $instance(X, carry), actor(X) = A,$
 $carried_object(X) = C, \neg holding(A, C).$

For readability, we selected the same names for the theory and its module. This theory will be used in the online appendix C (pp. 11–17), for the purpose of comparing \mathcal{ALM} and MAD .

Finally, *the semantics of a system description with a theory T consisting of multiple modules is given by the collection of models of the BAT defined by $f(T)$ and the collection of interpretations defined by the system's structure.*

This concludes our introduction to the syntax and semantics of \mathcal{ALM} .

4 Methodology of language use

In this section, we further illustrate the methodology of using \mathcal{ALM} for knowledge representation and for solving various computational tasks.

4.1 Representing knowledge in \mathcal{ALM}

We exemplify the methodology of representing knowledge in \mathcal{ALM} by considering a benchmark commonsense example from the field of reasoning about action and change – the Monkey and Banana Problem (McCarthy 1963; McCarthy 1968). (Another, more realistic, example of the use of \mathcal{ALM} can be found in the online appendix B, pp. 4–10.)

Problem 1 (Monkey and banana)

A monkey is in a room. Suspended from the ceiling is a bunch of bananas, beyond the monkey's reach. In the room, there is also a box. The ceiling is just the right height so that a monkey standing on the box under the bananas can reach the bananas. The monkey can move around, carry other things around, climb on the box, and grasp the bananas. What is the best sequence of actions for the monkey to get the bananas?

In accordance with the basic methodology of declarative programming, we will first *represent knowledge about the problem domain and then reduce the problem's solution to reasoning with this knowledge*. Based on our current experience, we recommend to divide the process of representation into the following steps:

Methodology of Creating Modular Representations in \mathcal{ALM} :

- Build a hierarchy of *actions* pertinent to the domain.
- Starting from the top of the hierarchy *gradually* build and *test* modules capturing properties of its actions. If necessary, add *general* non-action modules (e.g., a module defining a sequence of actions). Whenever feasible, use existing library modules.
- Build a module *main* containing *specific* information needed for the problem solution.
- Populate the hierarchy with the domain's objects.

Here are a few comments about the second step listed above: When deciding how many actions to describe in one module, consider balancing the size of the module

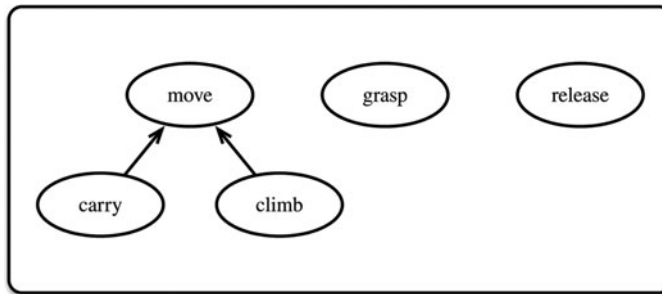


Fig. 5. Action hierarchy for the monkey and banana problem.

with the depth of the (part of the) hierarchy that it captures; also consider the resulting depth of the module dependency hierarchy. For instance, an action and its opposite are normally included in the same module. So are actions that usually occur together and share common fluents and sorts. To facilitate the discovery of relevant library modules, we assume that a dictionary indexed by action classes will be available to knowledge engineers. Action classes will be associated with the library modules in which they are described. The signature and axioms of library modules will be viewable by the knowledge engineer.

Let us illustrate the methodology by solving the Monkey and Banana problem. The story is clearly about an agent moving around, and grasping and carrying things between various points. The hierarchy of actions pertinent to the story is illustrated in Figure 5.

Note that unlike other actions, action *release* does not explicitly appear in the story. However, it is often advisable to consider actions together with their opposites, so our hierarchy contains *release* together with *grasp*.

To gradually build a theory *monkey_and_banana* containing the knowledge needed to solve the Monkey and Banana problem, we start with selecting a root of the action hierarchy – in our case action *move*. The inheritance hierarchy pertinent to *move* appears in Figure 2. We already discussed the module *moving* describing the properties of *move*. The theory consisting of this module can be tested on a number of specific domains using ASP-based methods discussed in the next section. Next, we select three actions *carry*, *grasp*, and *release* understood as *move while holding*, *take and hold*, and *stop holding* respectively. Since these actions share a fluent *holding*⁷ and sorts *things* and *agents*, and since a things-carrying agent usually also executes actions *grasp* and *release*, knowledge about these actions can be put in the same module. To do that we extend the inheritance hierarchy by a subclass *carriables* of *things* and expand module *carrying_things* from Section 3.2 by information about another two actions. Sort declarations of *carrying_things* from 3.2 will now also include

⁷ For simplicity, we assume that an agent can only hold one thing at a time. A more general module may allow to grasp a collection of things up to a certain capacity.

grasp :: actions
attributes
grasper : agents
grasped_thing : things

and

release :: actions
attributes
releaser : agents
released_thing : things

The section **function declarations** of the new module will contain the additional function *can_reach* needed as a precondition for the executability of *grasp*. The function will be defined in terms of locations of things.

defined

can_reach : agents \times things \rightarrow booleans

The set of **axioms** will be expanded as follows. The first two axioms below describe the direct effects of our new actions: action *grasp* results in the grasper holding the thing he grasped; this is no longer true after the thing is released.

occurs(*A*) **causes** *holding*(*X*, *Y*) **if** *instance*(*A*, *grasp*),
grasper(*A*) = *X*,
grasped_thing(*A*) = *Y*.
occurs(*A*) **causes** \neg *holding*(*X*, *Y*) **if** *instance*(*A*, *release*),
releaser(*A*) = *X*,
released_thing(*A*) = *Y*.

The constraint

\neg *holding*(*X*, *Y*₂) **if** *holding*(*X*, *Y*₁), *Y*₁ \neq *Y*₂

ensures that only one thing can be held at a time (and hence to grasp a thing an agent must have his hands free).

This is followed by the executability conditions: one cannot grasp a thing he is already holding or a thing that is out of his reach; one cannot release a thing unless he is holding it.

impossible *occurs*(*A*) **if** *instance*(*A*, *grasp*),
grasper(*A*) = *X*,
grasped_thing(*A*) = *Y*,
holding(*X*, *Y*).
impossible *occurs*(*A*) **if** *instance*(*A*, *grasp*),
grasper(*A*) = *X*,
grasped_thing(*A*) = *Y*,
 \neg *can_reach*(*X*, *Y*).
impossible *occurs*(*A*) **if** *instance*(*A*, *release*),
releaser(*A*) = *X*,
released_thing(*A*) = *Y*,
 \neg *holding*(*X*, *Y*).

We also need a simple definition of *can_reach* – an agent can always reach an object he shares a location with.

$$\text{can_reach}(M, O) \text{ if } \text{loc_in}(M) = \text{loc_in}(O).$$

This definition will later be expanded to describe the specific geometry of our domain.

This completes our construction of the new module *carrying_things*.

After testing the theory consisting of *moving* and *carrying_things* we proceed to constructing a new module, *climbing*, which axiomatizes action *climb* understood as *moving from the bottom of a thing to its top*. We assume that one can climb only on tops of a special type of things called *elevations*, which will be added to our hierarchy as a subset of *things*. The corresponding declarations look as follows:

```

module climbing
  depends on moving
  sort declarations
    elevations :: things
    climb :: move
  attributes
    elevation : elevations

```

Now, we introduce notation for *points associated with the tops of elevations*. The points are represented by *object constants* of the form $\text{top}(E)$ where E is an elevation. In \mathcal{ALM} , this is expressed by the following:

```

object constants
  top(elevations) : points

```

(Notice that *top* here is not a function symbol; if e is an elevation, then $\text{top}(e)$ is simply a point.)

The module contains axioms saying that $\text{top}(E)$ is the destination of climbing an elevation E :

$$\text{dest}(A) = \text{top}(E) \text{ if } \text{elevation}(A) = E.$$

and that a thing cannot be located on its own top:

$$\text{false if } \text{loc_in}(E) = \text{top}(E).$$

The last axiom prohibits an attempt by an agent to climb an elevation from a distance:

$$\begin{aligned} \text{impossible } \text{occurs}(X) \text{ if } & \text{instance}(X, \text{climb}), \\ & \text{actor}(X) = A, \\ & \text{elevation}(X) = O, \\ & \text{loc_in}(O) \neq \text{loc_in}(A). \end{aligned}$$

After testing the existing modules, we concentrate on the specific information needed for the problem solution. It will be presented in a module called *main*.

module main

depends on *carrying_things, climbing*

The main goal of the module is to define when the monkey can reach the banana. We start by dividing our sort *points* into three parts: *floor_points*, *ceiling_points*, and *movable_points*:

sort declarations

floor_points, ceiling_points, movable_points :: *points*

where the latter correspond to tops of movable objects. We will see the use of these sorts a little later. Now, we move to function declarations. The story is about three particular entities: the monkey, the banana, and the box. They will be defined as constants of our module.

object constants

monkey : *agents*

box : *carriables, elevations*

banana : *carriables*

We will also need a function *under*, such that $under(P, T)$ is true when point P is located under the thing T . Note that if we consider this function to be defined for arbitrary points, it will be dynamic – $under(top(box), banana)$ can be true in one state and false in another. This will force us to declare this function as a fluent, causing an unnecessary complication. Instead, we define *under* for floor points only, which is sufficient for our purpose and is substantially simpler.

function declarations**statics**

basic $under : floor_points \times things \rightarrow booleans$

To define our function *can_reach* we need the following axiom:

axioms

$can_reach(monkey, banana)$ **if** $loc_in(box) = P,$
 $under(P, banana),$
 $loc_in(monkey) = top(box).$

Finally, we need the following axioms for the basic fluent *connected*:

$connected(top(box), P)$ **if** $loc_in(box) = P,$
 $instance(P, floor_points).$
 $\neg connected(top(box), P)$ **if** $loc_in(box) \neq P,$
 $instance(P, floor_points).$
 $connected(P_1, P_2)$ **if** $instance(P_1, floor_points),$
 $instance(P_2, floor_points).$
 $\neg connected(P_1, P_2)$ **if** $instance(P_1, ceiling_points),$
 $instance(P_2, points),$
 $P_1 \neq P_2.$

This completes the construction of module *main* as well as theory *monkey_and_banana* that we will use to solve the Monkey and Banana problem. It is easy to see that

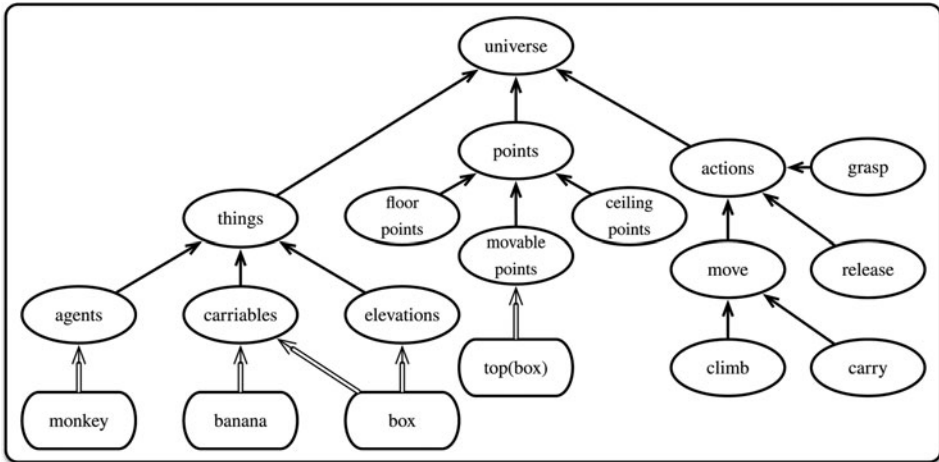


Fig. 6. Sort hierarchy for the monkey and banana problem.




Fig. 7. Module hierarchy for the monkey and banana problem.

the theory is semantically coherent, as it satisfies the conditions in Definition 11. Figures 6 and 7 represent the sort hierarchy and module hierarchy of this theory, respectively.

To complete the description of our domain, we introduce the structure containing three points located on the floor of the room and one point located on the ceiling, as well as movable points and particular actions mentioned in the story:

structure *monkey_and_banana*

instances

under_banana, initial_monkey, initial_box **in** *floor_points*

initial_banana **in** *ceiling_points*

top(box) **in** *movable_points*

move(P) **in** *move* **where** *instance(P, points)*

actor = monkey

dest = P

```

carry(box, P) in carry where instance(P, floor_points)
  actor = monkey
  carried_object = box
  dest = P
grasp(C) in grasp where instance(C, carriables)
  grasper = monkey
  grasped_thing = C
release(C) in release where instance(C, carriables)
  releaser = monkey
  released_thing = C
climb(box) in climb
  actor = monkey
  elevation = box

```

values of statics

```

under(under_banana, banana).
symmetric_connectivity.
¬transitive_connectivity.

```

The structure specifies that the relation *connected* is symmetric, but not transitive. The latter prevents the monkey from moving from its initial location directly on top of the box.

The theory and structure described above can be combined into a system description *monkey_and_banana* as follows:

```

system description monkey_and_banana_problem
  theory monkey_and_banana
    import motion from commonsense_library
    module main
      ⟨module body⟩
    structure monkey_and_banana
      ⟨structure body⟩

```

Note that the import statement above is a directive to import all of the modules of the library theory *motion* into the theory *monkey_and_banana*.

The system describes a unique hierarchy and a unique transition diagram, τ . Note that the hierarchy contains properly typed constants *monkey*, *box*, and *banana* declared in our module *main*; and that some of our functions, e.g., *under*, are partial.

It is not difficult to check that there is a path in τ that starts with the initial state of our problem and is generated by actions *move(initial_box)*, *grasp(box)*, *carry(box, under_banana)*, *release(box)*, *climb(box)*, *grasp(banana)*. The final state of this path will contain a fluent *holding(monkey, banana)*. In the next section, we discuss how ASP based reasoning can be used to automatically find such sequences.

4.2 \mathcal{ALM} 's use in solving computational tasks

A system description of \mathcal{ALM} describes a collection of transition diagrams that specifies some dynamic system. System descriptions can be used to solve

computational tasks such as temporal projection or planning, using a methodology similar to that developed for non-modular action languages like \mathcal{AL} (see, for instance, (Gelfond and Kahl 2014)).

4.2.1 Temporal projection

Normally, system descriptions of \mathcal{ALM} are used in conjunction with the description of the *system's recorded history* – a collection of facts about the values of fluents and the occurrences of actions at different time steps in a trajectory. (Since we are only dealing with discrete systems, such steps are represented by non-negative integers). Together, the system description and the history define the collection of possible trajectories of the system up to the current step. In our methodology of solving temporal projection tasks, possible trajectories are obtained by computing the answer sets of a logic program. To formally describe this methodology, we need the following definitions.

Definition 12 (History – adapted from (Balduccini and Gelfond 2003a))

By the *recorded history* Γ_n of a system description \mathcal{D} up to time step n we mean a collection of observations, i.e., facts of the form:

- (1) $observed(f(\bar{t}), v, i)$ – fluent $f(\bar{t})$ was observed to have value v at time step i , where $0 \leq i \leq n$.
- (2) $happened(a, i)$ – action a was observed to happen at time step i , where $0 \leq i < n$.

(There are two small differences between this and the definition of a history by Balduccini and Gelfond (2003a): the latter only allows boolean fluents and observations that have the form $observed(l, i)$ where l is a fluent or its negation. Similarly, for the next definitions in this subsection.)

We say that the *initial situation* of Γ_n is *complete* if, for every user-defined basic fluent f and any sequence of ground terms \bar{t} such that $observed(dom_f(\bar{t}), true, 0) \in \Gamma_n$, Γ_n also contains a fact of the form $observed(f(\bar{t}), v, 0)$.

Example 9 (History)

A possible recorded history for the system description *monkey_and_banana_problem* in Section 4.1 may look as follows:

$$\Gamma_1 =_{def} \{observed(loc_in(monkey), initial_monkey, 0), \\ observed(loc_in(box), initial_box, 0), \\ happened(move(initial_box), 0)\}$$

which says that initially, the monkey was at point *initial_monkey* and the box was at *initial_box*; the monkey went to the initial location of the box.

The semantics of a history Γ_n is given by the following definition:

Definition 13 (Model of a history – adapted from (Balduccini and Gelfond 2003a))

Let Γ_n be a history of a system description \mathcal{D} up to time step n .

- (a) A trajectory $\langle \sigma_0, a_0, \sigma_1, \dots, a_{n-1}, \sigma_n \rangle$ is a *model* of Γ_n if:

- (1) $a_i = \{a : \text{happened}(a, i) \in \Gamma_n\}$, for every $0 \leq i < n$.
- (2) if $\text{observed}(f(\bar{t}), v, i) \in \Gamma_n$ then $f(\bar{t}) = v \in \sigma_i$, for every $0 \leq i \leq n$.
- (b) Γ_n is *consistent* if it has a model.
- (c) An atom $f(\bar{t}) = v$ *holds* in a model M of Γ_n at time $0 \leq i \leq n$ if $f(\bar{t}) = v \in \sigma_i$;
 A literal $f(\bar{t}) \neq v$ *holds* in a model M of Γ_n at time $0 \leq i \leq n$ if $\text{dom}_f(\bar{t}) = \text{true} \in \sigma_i$ and $f(\bar{t}) = v \notin \sigma_i$;
 Γ_n *entails* a literal l at time step $0 \leq i \leq n$ if, for every model M of Γ_n , l holds in M .

Example 10 (Model of a history)

History Γ_1 from Example 9 is consistent. Its model is the trajectory:

$$M = \langle \{ \text{loc_in(monkey)} = \text{initial_monkey}, \text{loc_in(box)} = \text{initial_box}, \dots \}, \\ \text{move(initial_box)}, \\ \{ \text{loc_in(monkey)} = \text{initial_box}, \text{loc_in(box)} = \text{initial_box}, \dots \} \rangle.$$

(We do not show the values of *connected* since they are unchanged by our actions). Γ_1 entails, for example, $\text{loc_in(monkey)} = \text{initial_box}$ at time step 1.

Note that a consistent history may have more than one model if non-deterministic actions are involved or the initial situation is not complete.

Next, we define some useful vocabulary.

Definition 14 (Set of literals defining a sequence – adapted from (Balduccini and Gelfond 2003a)) Let Γ_n be a history of \mathcal{D} and A be a set of literals over signature Σ . We say that A *defines* the sequence

$$\langle \sigma_0, a_0, \sigma_1, \dots, a_{n-1}, \sigma_n \rangle$$

if:

- (a) $\sigma_i = \{f(t_0, \dots, t_n) = t : f(t_0, \dots, t_n) = t \in A \text{ and } f \text{ is a static or attribute}\} \cup \{f(t_0, \dots, t_n) = t : f(t_0, \dots, t_n, i) = t \in A \text{ and } f \text{ is a fluent}\}$
 for any $0 \leq i \leq n$, and
- (b) $a_k = \{a : \text{occurs}(a, k) \in A\}$ for any $0 \leq k < n$.

Definition 15 (Program Ω_{tp} – adapted from (Balduccini and Gelfond 2003a))

If Γ_n is a history of system description \mathcal{D} up to time step n , then by Ω_{tp} we denote the ASP{f} program constructed as follows:

- (1) For every action a such that $\text{happened}(a, i) \in \Gamma_n$, Ω_{tp} contains:
 $\text{occurs}(a, i) \leftarrow \text{happened}(a, i)$.
- (2) For every expression $\text{observed}(f(\bar{t}), v, 0) \in \Gamma_n$, Ω_{tp} contains:
 $f(\bar{t}, 0) = v \leftarrow \text{observed}(f(\bar{t}), v, 0)$.
- (3) For every expression $\text{observed}(f(\bar{t}), v, i) \in \Gamma_n, i > 0$, Ω_{tp} contains the *reality check axiom*:
 $\leftarrow \text{observed}(f(\bar{t}), v, i),$
 $\text{dom}_f(\bar{t}, i),$
 $f(\bar{t}, i) \neq v.$

Our methodology of finding trajectories by computing answer sets of a logic program is designed for system descriptions that match the intuition that defined functions are only shorthands, and their values are fully determined by those of basic statics and fluents. We call such system descriptions well-founded and define them formally as follows.

Definition 16 (Well-founded system description – adapted from (Gelfond and Incelezan 2013))

Let \mathcal{D} be a system description whose theory encodes the \mathcal{BAT} theory T , and whose structure defines a collection \mathcal{S} of models of T . \mathcal{D} is *well-founded* if, for every model \mathcal{M} in \mathcal{S} , and every interpretation \mathcal{I} with static part \mathcal{M} , the program $S_{\mathcal{I}}$ (defined as in Section 2.3) has at most one answer set.

The system description *monkey_and_banana_problem* from Section 4.1 is well-founded. An example of a system description that is *not well-founded* is *n_w_f* shown below and adapted from (Gelfond and Incelezan 2013). The two defined fluents of *n_w_f* are not defined in terms of basic statics or fluents but rather in terms of one another by mutually recursive axioms.

```

system description n_w_f
  theory n_w_f
    module main
      sort declarations
        c :: universe
      fluent declarations
        defined
          f : c → booleans
          g : c → booleans
      axioms
        f(X) if ¬g(X).
        g(X) if ¬f(X).
    structure n_w_f
      instances
        x in c

```

In the case of the non-modular action language \mathcal{AL} , there is a known syntactic condition that guarantees that a system description is well-founded (Gelfond and Incelezan 2013). This condition can be easily expanded to \mathcal{ALM} due to close connections between \mathcal{ALM} and \mathcal{AL} .

Trajectories of a dynamic system specified by a well-founded system description are computed using a logic program Π that consists of the ASP{f} encoding of the system description, the system's recorded history, and the program Ω_{tp} connecting the recorded history with the system description.

To simplify the presentation, in what follows we limit ourselves to well-founded system descriptions that describe domains in which there is *complete* information

about the sort memberships of objects of the domain.⁸ Let us consider system description \mathcal{D} that meets this requirement, and let \mathcal{M} be a model of \mathcal{D} 's theory. Then, the program $P_{\mathcal{M}}$ obtained from the theory of \mathcal{D} and \mathcal{M} as described in Section 2.3 will be used as the $\text{ASP}\{f\}$ encoding of \mathcal{D} .

Definition 17 (Program $\Pi_{tp}(\mathcal{D})$)

If Γ_n is a history of \mathcal{D} up to step n , then $\Pi_{tp}(\mathcal{D})$ is the logic program defined as

$$\Pi_{tp}(\mathcal{D}) =_{\text{def}} P_{\mathcal{M}} \cup \Gamma_n \cup \Omega_{tp}$$

such that the sort *step* in the signature of $\Pi_{tp}(\mathcal{D})$ ranges over the set $\{0, \dots, n\}$.

Proposition 1

If Γ_n is a *consistent* history of \mathcal{D} such that the initial situation of Γ_n is complete, then M is a model of Γ_n iff M is defined by some answer set of program $\Pi_{tp}(\mathcal{D})$.

This proposition can be proven using techniques similar to the ones employed in Lemma 5 in Balduccini and Gelfond (2003a).⁹

We used the above methodology of solving temporal projection tasks to create a question answering system in the context of the Digital Aristotle project (Incezan and Gelfond 2011). Our system was capable of answering complex end-of-the-chapter questions on cell division, extracted from a well-known biology textbook.

4.2.2 Planning

In planning problems, in addition to the history of the dynamic system up to the current time point, information about the goal to be achieved is also provided. Given a system description of \mathcal{ALM} whose theory describes a \mathcal{BAT} T , a *goal* is a collection G of ground user-defined fluent literals over the signature of T . For instance, for the Monkey and Banana problem in Section 4.1, the goal is $G_{mb} = \{\text{holding}(\text{monkey}, \text{banana})\}$. Goals can be encoded as logic programming rules, as described in the following definition:

Definition 18 (Goal encoding)

Given a goal G , we call *encoding of G* , denoted by $lp(G)$ the rule

$$\text{goal}(I) \leftarrow \text{body}$$

where *body* is defined as follows:

$$\text{body} =_{\text{def}} \{f(\bar{t}, I) = v : f(\bar{t}) = v \in G\} \cup \{f(\bar{t}, I) \neq v : f(\bar{t}) \neq v \in G\}.$$

⁸ This is not a serious restriction; it can be easily lifted by adding to the ASP encoding of the \mathcal{ALM} system description rules of the type

$$\text{is}_a(x, c) \text{ or } \neg \text{is}_a(x, c)$$

for every object x and every source node c in the hierarchy of sorts.

⁹ The proof and text of Lemma 5 appear on page 29 of the version of Balduccini and Gelfond (2003a) available at <http://arxiv.org/pdf/cs/0312040v1.pdf>. Retrieved on August 3, 2014.

In order to solve planning problems, a slightly different logic programming module will be needed than for solving temporal projection tasks. This module is defined in CR-Prolog (Balduccini and Gelfond 2003b), an extension of ASP designed to handle, among other things, rare events. In addition to regular ASP rules, programs in CR-Prolog may contain *consistency restoring* rules that have the following syntax:

$$h_1 \text{ or } \dots \text{ or } h_k \stackrel{+}{\leftarrow} l_1, \dots, l_m, \text{not } l_{m+1}, \dots, \text{not } l_n.$$

Informally, this statement says that an intelligent agent who believes l_1, \dots, l_m and has no reason to believe l_{m+1}, \dots, l_n may believe one of h_i 's, $1 \leq i \leq k$, but only if no consistent set of beliefs can be formed otherwise. For the formal semantics of CR-Prolog, we refer the reader to Balduccini and Gelfond (2003b). An extension of ASP{f} by consistency restoring rules is defined in Balduccini and Gelfond (2012). Solvers for CR-Prolog are described in Balduccini (2007) and Balai *et al.* (2012).

Definition 19 (Planning module (Balduccini 2004; Gelfond and Kahl 2014))

Given a goal G , the planning module Ω_{pl} extends module Ω_{tp} from Section 4.2.1 by the following rules:

$$\begin{array}{ll} \text{success} & \leftarrow \text{goal}(I), I \leq n \\ & \leftarrow \text{not success} \\ r_1(A, I) : \text{occurs}(A, I) & \stackrel{+}{\leftarrow} \text{instance}(A, \text{actions}) \\ \text{smtg_happened}(I) & \leftarrow \text{occurs}(A, I) \\ & \leftarrow \text{not smtg_happened}(I), \\ & \text{smtg_happened}(I + 1). \end{array}$$

Ω_{pl} computes minimal plans of maximum length n by the use of the consistency restoring rule r_1 and the two regular rules that follow it.

The actual program for computing plans is constructed similarly as before.

Definition 20 (Program $\Pi_{pl}(\mathcal{D})$)

If Γ_n is a history of \mathcal{D} up to step n and G is a goal over \mathcal{D} , then $\Pi_{pl}(\mathcal{D}, G)$ is the logic program defined as

$$\Pi_{pl}(\mathcal{D}, G) =_{def} P_{\mathcal{M}} \cup \Gamma_n \cup \Omega_{pl} \cup lp(G)$$

such that the sort *step* in the signature of $\Pi_{pl}(\mathcal{D}, G)$ ranges over the set $\{0, \dots, n\}$.

The following proposition specifies how answer sets of the logic program defined above can be mapped into plans for achieving given goals.

Proposition 2

If Γ_n is a *consistent* history of \mathcal{D} such that the initial situation of Γ_n is complete and G is a goal over \mathcal{D} , then the collection of atoms of the form $\text{occurs}(a, i)$ from an answer set of $\Pi_{pl}(\mathcal{D}, G)$ defines a minimal plan for achieving goal G , and every such plan is represented by the *occurs* atoms of some answer set of $\Pi_{pl}(\mathcal{D}, G)$.

Example 11 (Planning in the monkey and banana problem)

If we consider the Monkey and Banana problem with the initial situation

$$\Gamma_{mb} = \{ \text{observed}(\text{loc_in}(\text{monkey}), \text{initial_monkey}, 0), \\ \text{observed}(\text{loc_in}(\text{box}), \text{initial_box}, 0) \}$$

and the goal

$$G_{mb} = \{ \text{holding}(\text{monkey}, \text{banana}) \}$$

defined earlier, then an answer set of program $\Pi_{pl}(\text{monkey_and_banana_problem}, G_{mb})$ will contain the following *occurs* atoms:

$$\{ \text{occurs}(\text{move}(\text{initial_box}), 0), \quad \text{occurs}(\text{grasp}(\text{box}), 1), \\ \text{occurs}(\text{carry}(\text{box}, \text{under_banana}), 2), \quad \text{occurs}(\text{release}(\text{box}), 3), \\ \text{occurs}(\text{climb}(\text{box}), 4), \quad \text{occurs}(\text{grasp}(\text{banana}), 5) \}$$

defining a minimal plan $\langle \text{move}(\text{initial_box}), \text{grasp}(\text{box}), \text{carry}(\text{box}, \text{under_banana}), \text{release}(\text{box}), \text{climb}(\text{box}), \text{grasp}(\text{banana}) \rangle$ resulting in the monkey holding the banana at time step 6. The program will also find the second minimal plan in which $\text{carry}(\text{box}, \text{under_banana})$ at step 2 is replaced by $\text{move}(\text{under_banana})$. Since the first action is more specific than the second one, the first plan seems to be preferable. This can easily be expressed by a slightly modified planning module allowing only most specific actions.

5 Related work

Many ideas of \mathcal{ALM} , such as the notions of action language, module, sort hierarchy, attribute defined as a partial function, etc., are well known from the literature on programming languages and knowledge representation. Some of the basic references to these notions were given in the text. In this section, we briefly comment on the relationship between \mathcal{ALM} and the previously existing modular action languages *MAD* (Erdoğan and Lifschitz 2006; Lifschitz and Ren 2006; Desai and Singh 2007), *TAL-C* (Gustafsson and Kvarnström 2004), and the earlier version of \mathcal{ALM} (Gelfond and Incelezan 2009).

We start with summarizing the differences between the two versions of \mathcal{ALM} . There are a number of changes in the syntax of the language. For instance, theories of the new version of \mathcal{ALM} may contain non-boolean fluents¹⁰ and constants that substantially simplify \mathcal{ALM} 's use for knowledge representation. Axioms of a theory, which in the old version were included in the theory's declarations, are now put in a separate section of the theory. This removed the problem of deciding which fluent or action declaration should contain an axiom, and improved the readability of the language. There are also substantial improvements in the syntax of axioms, etc. Another collection of changes is related to the semantics of the language. First, the new semantics, based on the notions of *BAT* and its models, clarified and generalized the old definition and allowed the introduction of the entailment relation.

¹⁰ In the field of logic programming, an early discussion on the introduction of functions appears in (Hanus 1994).

Second, the semantics is now defined for structures with possibly underspecified membership relations of its objects in the sort hierarchy, which simplifies reasoning with incomplete information. Third, the semantics was initially given in terms of action language \mathcal{AL} (Turner 1997; Baral and Gelfond 2000), where the \mathcal{AL} semantics is defined by a translation into ASP; now, we give the semantics of our language directly in ASP – in fact, in an extension of ASP with non-Herbrand functions, $\text{ASP}\{\text{f}\}$ (Balduccini 2013). We believe that decoupling \mathcal{ALM} from \mathcal{AL} will allow us to combine \mathcal{ALM} with action languages that correspond to other intuitions.

Another modular language is TAL-C (Gustafsson and Kvarnström 2004), which allows definitions of classes of objects that are somewhat similar to those in \mathcal{ALM} . TAL-C, however, seems to have more ambitious goals: the language is used to describe and reason about various dynamic scenarios, whereas in \mathcal{ALM} the description of a scenario and that of reasoning tasks are not viewed as part of the language. The more rigid structure of \mathcal{ALM} supports the *separation of concerns* design principle and makes it easier to give a formal semantics of the language. These differences led to vastly distinct knowledge representation styles reflected in these languages.

There are smaller, but still very substantial, differences between \mathcal{ALM} and *MAD*. The two languages are based on non-modular action languages with substantially different semantics and underlying assumptions, use very different constructs for creating modules and for defining actions as special cases, etc. A more detailed comparison between the two approaches can be found in the online appendix C, pp. 11–17.

6 Conclusions and future work

In this paper, we have presented a methodology of representing and reasoning about dynamic systems. A knowledge engineer following this methodology starts with finding a proper generalization of a particular dynamic system D , finds the sorts of objects pertinent to this generalization, organizes these sorts into an inheritance hierarchy and uses causal laws, definitions, and executability conditions to specify relevant properties of the sorts elements. The resulting *BAT*, say T , gives the first mathematical model of the system. In the next step of the development, a knowledge engineer refines this model by providing its description in the high-level action language \mathcal{ALM} . The language has means for precisely representing the signature of T including its sort hierarchy. It is characterized by a modular structure, which improves readability and supports the step-wise development of a knowledge base, reuse of knowledge, and creation of knowledge libraries. \mathcal{ALM} 's description of T can be used to specify multiple dynamic systems with different collections of objects and statics. A particular system D can be specified by populating sorts of T by objects of D and defining values of D 's statics. This step is also supported by \mathcal{ALM} , which clearly separates the definition of *sorts of objects* of the domain (given in T) from the definition of *instances* of these sorts (given by an \mathcal{ALM} structure). This, together with the means for defining objects of the domain as special cases

of previously defined ones, facilitates the step-wise development and testing of the knowledge base and improves its elaboration tolerance.

A close relationship between \mathcal{ALM} and ASP allows the use of \mathcal{ALM} system descriptions for non-trivial reasoning problems including temporal projection, planning, and diagnosis. This is done by an automatic translation of an \mathcal{ALM} system description into logic programs whose answer sets correspond to solutions of the corresponding problems. The existence of efficient answer set solvers that allow to compute these answer sets substantially increases the practical value of this approach.

The above methodology has been illustrated by two examples: the well-known benchmark Monkey and Banana problem and a more practical problem of formalization of knowledge and answering questions about biological processes such as the cell division (see online appendix B, pp. 4–10). It is possible (and even likely) that further experience with \mathcal{ALM} will suggest some useful extensions of the language but the authors believe that the version presented in this paper will remain relatively stable and provide a good basis for such extensions.

We conclude by briefly outlining a number of questions about \mathcal{ALM} that we believe deserve further investigation:

- Investigating mathematical properties of \mathcal{ALM} and its entailment relation. This includes but is not limited to studying compositional properties of \mathcal{ALM} modules, axiomatizing its entailment relation, and establishing a closer relationship between \mathcal{ALM} and modular logic programming.
- Developing more efficient reasoning algorithms exploiting the modular structure of \mathcal{ALM} 's theories and the available information about the sorts of objects in \mathcal{ALM} 's system descriptions. Among other things, it is worth investigating the possible use of modular logic programming as well as the methods from Gebser *et al.* (2011), Gebser *et al.* (2011), and Balai *et al.* (2012). It may also be interesting to see if the implementation could benefit from hybrid approaches combining description logics with ASP (e.g., (Eiter *et al.* 2008)) or from typed logic programming (e.g., (Pfenning 1992)).
- Designing and implementing a development environment to facilitate the use of \mathcal{ALM} in applications, the creation and storage of libraries, and the testing and debugging of theories and modules.
- Extending \mathcal{ALM} with the capability of representing knowledge about *hybrid* domains, i.e., domains that allow both discrete and continuous change. In particular, it may be a good idea to combine \mathcal{ALM} with action language \mathcal{H} (Chintabathina *et al.* 2005; Chintabathina 2012).
- Developing the core of an \mathcal{ALM} library of commonsense knowledge. (In particular, we would like to create an \mathcal{ALM} library module containing a theory of intentions in the style of Blount *et al.* (2014).) This work would allow us to extend our study on the capabilities of our language, while simultaneously providing a tool for members of our community to use when building their reasoning systems.

Acknowledgments

We are grateful to Evgenii Balai, Justin Blount, Vinay Chaudhri, Vladimir Lifschitz, Yana Todorova, and the anonymous reviewers for useful comments and discussions. This work was partially supported by NSF grant IIS-1018031.

References

- AKMAN, V., ERDOĞAN, S. T., LEE, J., LIFSCHITZ, V. AND TURNER, H. 2004. Representing the zoo world and the traffic world in the language of the causal calculator. *Artificial Intelligence* 153, 1–2 (March), 105–140.
- BALAI, E., GELFOND, M. AND ZHANG, Y. 2012. SPARC – sorted ASP with consistency restoring rules. In *Proc. of Answer Set Programming and Other Computing Paradigms (ASPOCP 2012)*, M. Fink and Y. Lierler, Eds. Proceedings published online in arXiv at <http://arxiv.org/html/1301.2215v1>, 19–33.
- BALDUCCINI, M. 2004. USA-smart: Improving the quality of plans in answer set planning. In *Proc. PADL'04*, B. Jayaraman, Ed. Lecture Notes in Artificial Intelligence (LNCS). Springer, Berlin, 135–147.
- BALDUCCINI, M. 2007. CR-MODELS: An inference engine for CR-prolog. In *Proc. of LPNMR-07*, C. Baral, G. Brewka and J. S. Schlipf, Eds., Springer, Berlin, 18–30.
- BALDUCCINI, M. 2013. ASP with non-Herbrand partial functions: A language and system for practical use. *Theory and Practice of Logic Programming* 13, 4–5, 547–561.
- BALDUCCINI, M. AND GELFOND, M. 2003a. Diagnostic reasoning with A-prolog. *Theory and Practice of Logic Programming* 3, 425–461.
- BALDUCCINI, M. AND GELFOND, M. 2003b. Logic programs with consistency-restoring rules. In *International Symposium on Logical Formalization of Commonsense Reasoning*, P. Doherty, J. McCarthy and M.-A. Williams, Eds., AAAI 2003 Spring Symposium Series, Palo Alto, CA, 9–18.
- BALDUCCINI, M. AND GELFOND, M. 2012. Language ASP{f} with arithmetic expressions and consistency-restoring rules. In *Proc. of Answer Set Programming and Other Computing Paradigms (ASPOCP 2012)*, M. Fink and Y. Lierler, Eds. Proceedings published online in arXiv at <http://arxiv.org/pdf/1301.1387v1.pdf>, 35–49.
- BARAL, C. 2003. *Knowledge Representation, Reasoning, and Declarative Problem Solving*. Cambridge University Press.
- BARAL, C., DZIFCAK, J. AND TAKAHASHI, H. 2006. Macros, macro calls and use of ensembles in modular answer set programming. In *Logic Programming*, S. Etalle and M. Truszczyski, Eds. Lecture Notes in Computer Science, vol. 4079. Springer, Berlin Heidelberg, 376–390.
- BARAL, C. AND GELFOND, M. 2000. *Reasoning Agents in Dynamic Domains*. Kluwer Academic Publishers, Norwell, MA, 257–279.
- BARTHOLOMEW, M. AND LEE, J. 2013. On the stable model semantics for intensional functions. *Journal of Theory and Practice of Logic Programming (TPLP)* 13, 4–5, 863–876.
- BLOUNT, J., GELFOND, M. AND BALDUCCINI, M. 2014. Towards a theory of intentional agents. In *AAAI 2014 Spring Symposium Series*, M. Sridharan, F. Yang, S. Ramamoorthy, V. Patoglu and E. Erdem, Eds. AAAI Press, Palo Alto, CA.
- CABALAR, P. 2011. Functional answer set programming. *Journal of Theory and Practice of Logic Programming (TPLP)* 11, 2–3, 203–233.
- CALIMERI, F. AND IANNI, G. 2006. Template programs for disjunctive logic programming: An operational semantics. *AI Communications* 19, 3, 193–206.

- CHINTABATHINA, S. 2012. Planning and scheduling in hybrid domains. *Frontiers in Artificial Intelligence and Applications* 241, 59–70.
- CHINTABATHINA, S., GELFOND, M. AND WATSON, R. 2005. Modeling hybrid domains using process description language. In *Proc. of ASP '05 Answer Set Programming: Advances in Theory and Implementation*, M. De Vos and A. Provetti, Eds., CEUR-WS, CEUR-WS.org/Aachen, 303–317.
- DESAI, N. AND SINGH, M. P. 2007. A modular action description language for protocol composition. In *Proc. of the 22nd AAAI Conference on Artificial Intelligence, July 22–26, 2007, Vancouver, British Columbia, Canada*, AAAI Press, Palo Alto, CA, 962–967.
- DOVIER, A., FORMISANO, A. AND PONTELLI, E. 2007. Multivalued action languages with constraints in CLP(FD). *Logic Programming: Lecture Notes in Computer Science* 4670, 255–270.
- EITER, T., ERDEM, E., FINK, M. AND SENKO, J. 2010. Updating action domain descriptions. *Artif. Intell.* 174, 15 (October), 1172–1221.
- EITER, T., FABER, W., LEONE, N., PFEIFER, G. AND POLLERES, A. 2004. Approach to knowledge-state planning: Semantics and complexity. *ACM Transactions on Computational Logic* 5, 206–263.
- EITER, T., IANNI, G., LUKASIEWICZ, T., SCHINDLAUER, R. AND TOMPITS, H. 2008. Combining answer set programming with description logics for the semantic web. *Artificial Intelligence* 172, 12–13 (August), 1495–1539.
- ERDOĞAN, S. AND LIFSCHITZ, V. 2006. Actions as special cases. In *Principles of Knowledge Representation and Reasoning: Proceedings of the International Conference*, P. Doherty, J. Mylopoulos and C. A. Welty, Eds., AAAI Press, Palo Alto, CA, 377–387.
- ERDOĞAN, S. T. 2008. *A Library of General-Purpose Action Descriptions*. Ph.D. thesis, University of Texas at Austin, Austin, TX, USA.
- FODOR, P. AND KIFER, M. 2011. Modeling hybrid domains using process description language. In *Proc. of the 27th International Conference on Logic Programming (ICLP)*, J. P. Gallagher and M. Gelfond, Eds., Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Wadern, 162–174.
- GEBSER, M., GROTE, T., KAMINSKI, R. AND SCHAUB, T. 2011. Reactive answer set programming. In *LPNMR*, J. P. Delgrande and W. Faber, Eds. Lecture Notes in Computer Science, vol. 6645. Springer, Berlin, 54–66.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B. AND SCHAUB, T. 2012. *Answer Set Solving in Practice*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan and Claypool Publishers.
- GEBSER, M., SABUNCU, O. AND SCHAUB, T. 2011. An incremental answer set programming based system for finite model computation. *AI Commun.* 24, 2, 195–212.
- GELFOND, M. AND INCLEZAN, D. 2009. Yet another modular action language. In *Proc. of SEA-09*. University of Bath Opus: Online Publications Store, 64–78.
- GELFOND, M. AND INCLEZAN, D. 2013. Some properties of system descriptions in ALd. *Journal of Applied Non-Classical Logics* 23, 105–120.
- GELFOND, M. AND KAHL, Y. 2014. *Knowledge Representation, Reasoning, and the Design of Intelligent Agents*. Cambridge University Press.
- GELFOND, M. AND LIFSCHITZ, V. 1988. The stable model semantics for logic programming. In *Proc. of ICLP-88*, R. A. Kowalski and K. A. Bowen, Eds., MIT Press, Cambridge, MA, 1070–1080.
- GELFOND, M. AND LIFSCHITZ, V. 1991. Classical negation in logic programs and disjunctive databases. *New Generation Computing* 9, 3–4, 365–386.
- GELFOND, M. AND LIFSCHITZ, V. 1993. Representing action and change by logic programs. *Journal of Logic Programming* 17, 2–4, 301–321.

- GELFOND, M. AND LIFSCHITZ, V. 1998. Action languages. *Electronic Transactions on AI* 3, 16, 193–210.
- GIUNCHIGLIA, E., LEE, J., LIFSCHITZ, V., MCCAIN, N. AND TURNER, H. 2004a. Nonmonotonic causal theories. *Artificial Intelligence* 153, 1–2, 105–140.
- GIUNCHIGLIA, E., LEE, J., LIFSCHITZ, V., MCCAIN, N. AND TURNER, H. 2004b. Nonmonotonic causal theories. *Artificial Intelligence* 153, 105–140.
- GIUNCHIGLIA, E. AND LIFSCHITZ, V. 1998. An action language based on causal explanation: Preliminary report. In *Proc. of National Conference on Artificial Intelligence (AAAI)*, J. Mostow and C. Rich, Eds., AAAI Press, Palo Alto, CA, 623–630.
- GROSOFF, B., DEAN, M. AND KIFER, M. 2009. The SILK system: Scalable higher-order defeasible rules. In *International RuleML Symposium on Rule Interchange and Applications*. Available online at: <http://silk.semwebcentral.org/talk-ruleml2009-silk-demo.pdf>
- GUNNING, D., CHAUDHRI, V. K., CLARK, P., BARKER, K., CHAW, S.-Y., GREAVES, M., GROSOFF, B., LEUNG, A., McDONALD, D., MISHRA, S., PACHECO, J., PORTER, B., SPAULDING, A., TECUCI, D. AND TIEN, J. 2010. Project Halo—progress toward digital aristotle. *AI Magazine* 31, 3, 33–58.
- GUSTAFSSON, J. AND KVARNSTRÖM, J. 2004. Elaboration tolerance through object-orientation. *Artificial Intelligence* 153, 239–285.
- HANUS, M. 1994. The integration of functions into logic programming: From theory to practice. *Journal of Logic Programming* 19–20, Supplement 1, 583–628.
- HENSCHER, A. AND THIELSCHER, M. 2000. The LMW Traffic world in the fluent calculus. *Linköping Electronic Articles in Computer and Information Science* 5 (14). Available at <http://www.ep.liu.se/ea/cis/2000/014/>.
- INCLEZAN, D. 2010. Computing trajectories of dynamic systems using ASP and Flora-2. Paper presented at NonMon@30: Thirty Years of Nonmonotonic Reasoning Conference, Lexington, Kentucky, 22–25 October. G. Brewka, V. Marek and M. Truszczyński, Eds., Available at <http://www.depts.ttu.edu/cs/research/kr/lab/pdfs/papers/di10.pdf>
- INCLEZAN, D. AND GELFOND, M. 2011. Representing biological processes in modular action language ALM. In *Proc. of the 2011 AAAI Spring Symposium on Formalizing Commonsense*. AAAI Press, Palo Alto, CA, 49–55.
- KAKAS, A. AND MILLER, R. 1997. A simple declarative language for describing narratives with actions. *Journal of Logic Programming* 31, 1–3, 157–200.
- LEONE, N., PFEIFER, G., FABER, W., EITER, T., GOTTLÖB, G., PERRI, S. AND SCARCELLO, F. 2006. The DLV system for knowledge representation and reasoning. *ACM Transactions on Computational Logic* 7, 3, 499–562.
- LIERLER, Y. AND TRUSZCZYŃSKI, M. 2013. Modular answer set solving. In *Proc. of the 27th AAAI Conference on Artificial Intelligence (AAAI-13)*, AAAI Press, Palo Alto, CA, 68–70.
- LIFSCHITZ, V. 2012. Logic programs with intensional functions. In *Proc. of International Conference on Principles of Knowledge Representation and Reasoning (KR)*, G. Brewka, T. Eiter and S. A. McIlraith, Eds., AAAI Press, Palo Alto, CA, 24–31.
- LIFSCHITZ, V. AND REN, W. 2006. A modular action description language. In *Proc. of the 21st National Conference on Artificial Intelligence (AAAI)*, AAAI Press, Palo Alto, CA, 853–859.
- MAREK, V. W. AND TRUSZCZYŃSKI, M. 1999. Stable models and an alternative logic programming paradigm. In *The Logic Programming Paradigm: A 25-Year Perspective*, Springer, Verlag, Berlin, 375–398.
- MCCARTHY, J. 1963. Situations, actions, and causal laws. Tech. Rep. Memo 2, Stanford University.
- MCCARTHY, J. 1968. Programs with common sense. In *Semantic Information Processing*, MIT Press, 403–418.

- MCCARTHY, J. 1998. Elaboration tolerance. In *Proceedings of Commonsense Reasoning*. Available online at: <http://www-formal.stanford.edu/jmc/elaboration/elaboration.html>
- NIEMELÄ, I. 1998. Logic programs with stable model semantics as a constraint programming paradigm. *Annals of Mathematics and Artificial Intelligence* 25, 72–79.
- NIEMELÄ, I. AND SIMONS, P. 1997. Smodels - an implementation of the stable model and well-founded semantics for normal logic programs. In *Proc. of the 4th International Conference on Logic Programming and Non-Monotonic Reasoning (LPNMR-97)*. Lecture Notes in Artificial Intelligence (LNCS), J. Dix, U. Furbach and A. Nerode, Eds., vol. 1265. Springer, Berlin, 420–429.
- OIKARINEN, E. AND JANHUNEN, T. 2006. Modular equivalence for normal logic programs. In *Proceedings of 17th European Conference on Artificial Intelligence (ECAI)*, G. Brewka, S. Coradeschi, A. Perini and P. Traverso, Eds., IOS Press, Amsterdam, 412–416.
- PFENNING, F., Ed. 1992. *Types in Logic Programming*. MIT Press.
- SANDEWALL, E. 1999. Logic modelling workshop: Communicating axiomatizations of actions and change. URL: <http://www.ida.liu.se/ext/etai/lmw>.
- STRASS, H. AND THIELSCHER, M. 2012. A language for default reasoning about actions. In *Correct Reasoning: Essays in Honor of Vladimir Lifschitz*, E. Erdem, J. Lee, Y. Lierler, and D. Pearce, Eds. Lecture Notes in Computer Science, vol. 7265. Springer, 527–542.
- TURNER, H. 1997. Representing actions in logic programs and default theories: A situation calculus approach. *Journal of Logic Programming* 31, 1-3 (June), 245–298.
- TURNER, H. 1999. A logic of universal causation. *Artificial Intelligence* 113, 87–123.
- VAN GELDER, A., ROSS, K. A. AND SCHLIPF, J. S. 1991. The well-founded semantics for general logic programs. *Journal of the ACM* 38, 619–649.
- WIRTH, N. 1971. Program development by stepwise refinement. *Communications of the ACM* 14, 4 (April), 221–227.