

Representing Multiple Theories

P. Pandurang Nayak

Recom Technologies, NASA Ames Research Center
AI Research Branch, MS 269-2
Moffett Field, CA 94035
nayak@ptolemy.arc.nasa.gov

Abstract

Most Artificial Intelligence programs lack generality because they reason with a single domain theory that is tailored for a specific task and embodies a host of implicit assumptions. *Contexts* have been proposed as an effective solution to this problem by providing a mechanism for explicitly stating the assumptions underlying a domain theory. In addition, contexts can be used to focus reasoning, allow the representation of mutually incoherent domain theories, lift axioms from one context into another, and transcend a context. In this paper we develop a simple propositional logic of context suitable for representing and reasoning with multiple domain theories. We introduce contexts as modal operators, and allow different contexts to have different vocabularies. We analyze the computational properties of the logic, providing the central computational justification for the use of contexts. We show how the logic effectively handles the common uses of contexts. We also discuss the extensions needed to handle first-order logic.

Introduction

Artificial Intelligence is founded upon the idea that any domain of interest can be described within a formal language, and that an artificial agent can draw meaningful conclusions about the domain purely by reasoning within the formal language. The impressive success of various AI programs is testimony to the power of this idea. However, as noted by McCarthy (1987), most of these programs lack generality. A key source of this lack of generality is that most programs reason with a single domain theory. Such a theory is usually tailored for a specific task, embodying a host of implicit assumptions, making it inapplicable for reasoning about other tasks. For example, typical axiomatizations of engineered systems do not account for the extreme environmental conditions encountered in space exploration, e.g., very high or low operating temperatures, gamma radiation, vacuum conditions. This leads to simpler axiomatizations that are adequate for a variety of design tasks, but that are inapplicable for the the design of a new space probe.

The specificity of domain theories is neither avoidable nor undesirable. It is unavoidable because any axiomatization of a real-world domain, however detailed, will invariably miss some subtle nuance. It is desirable because the specificity often buys us computational efficiency in reasoning. The goal, then, is to provide an artificial agent with a variety of domain theories, with varying generality and computational efficiency, and have it automatically select the theories most appropriate to the task at hand.

Contexts

McCarthy (1987; 1993) and Guha (1991) have argued persuasively that the notion of *context* is an effective solution to the problem of representing and reasoning with multiple domain theories. The basic idea is to encapsulate a domain theory within a context, and to explicitly state the assumptions underlying this context.¹ This provides a mechanism for using the simplest, and most efficient, applicable domain theory in every situation.

In addition to providing a mechanism for explicating the assumptions underlying a domain theory, McCarthy and Guha have identified a variety of other uses of contexts. Collecting together a set of related axioms into a context can be used to focus reasoning. Consider SIGMA, a knowledge-base of scientific domain knowledge that supports building consistent, coherent, and executable domain models (Keller, Rimon, & Das 1994). SIGMA contains axioms describing two different application domains: modeling the atmosphere of Titan (a moon of Saturn), and modeling a forest ecosystem. By separating the axioms of these two domains into different contexts, reasoning can be easily focused on just the axioms of the domain of interest.

Encapsulating domain theories within a context supports the representation of mutually incoherent domain theories. Domain theories can be mutually incoherent in a number of ways. First, they can use the same propositions with entirely different meanings.

¹Of course, it is often not possible to state *all* the assumptions underlying a context.

For example, a “low temperature” in the context of Titan is quite different from a “low temperature” in the context of an Earth-based forest ecosystem. Second, different theories can be mutually inconsistent. For example, in modeling the gases in Titan’s atmosphere, two different, and mutually contradictory, axiomatizations are possible: gases can be modeled as being either ideal or non-ideal. Both axiomatizations can co-exist peacefully by placing them in different contexts. Third, different domain theories are possible depending on one’s perspective. Separating these different theories into different contexts simplifies knowledge base construction. For example, different ecosystem processes are best described at different time scales. The growth of trees is best described using a time scale of years, while photosynthesis is best described using a time scale of hours. Describing processes at inappropriate time scales is both difficult and opaque. Separating such descriptions into different contexts solves this problem.

Another important benefit of contexts stems from the ability to *lift* axioms from one context into another context. The simplest form of such lifting involves inheriting all the axioms of one context into another context. As with other forms of inheritance, this supports reuse and facilitates the construction and maintenance of large knowledge bases. For example, both the Titan modeling context and the ecosystem modeling context in SIGMA need axiomatizations of quantities and equations. This axiomatization can be placed in its own context, and reused in the Titan and the ecosystem contexts via inheritance.

More sophisticated forms of lifting allow axioms in one context to be modified before being lifted into another context. Adapting an example from (McCarthy 1993), let context C_1 be the specialization of context C_2 to a specific time t . Since time is fixed in C_1 , it remains implicit. Lifting an axiom from C_1 to C_2 requires the time to be made explicit. For example, the proposition $at(jmc, stanford)$ in C_1 would have to be modified to the proposition $at(jmc, stanford, t)$ before being lifted into C_2 . Similar modifications are needed when a context specializes to a speaker, a hearer, a location, etc. Modifications are also needed when axioms are lifted from one context into another context with fewer underlying assumptions. The implicit assumptions of the former context need to be made explicit.

Finally, the most ambitious use of contexts is to allow an AI system to *transcend* its current context, either by being told how, or by autonomous learning and discovery. For example, a system should be able to transcend the simpler context of Newtonian mechanics and move to the more general context of quantum mechanics.

Logics of context

Recently, there has been much interest in developing logics of context (Guha 1991; Buvač & Mason 1993).

Much of this work has focussed on developing new syntax and semantics for such a logic. However, we feel that, for the purposes of representing and reasoning with multiple domain theories, most aspects of a logic of context are naturally captured by simple extensions to traditional modal logic. In this paper, we focus on propositional modal logics, and introduce a modal operator, C_i , for the i^{th} context and allow different contexts to have different vocabularies. We introduce an axiomatization for the logic, and analyze its computational properties. This analysis provides the central computational justification for the use of contexts. We introduce important relations between contexts, based on the notion of *interpretation functions* (Enderton 1972), and show how these relations can be used to implement important properties of contexts. We conclude with a brief discussion of the extensions needed to handle first-order logic.

The decision to represent contexts as modal operators, rather than as terms, sharply diverges from earlier work on logics of context (Buvač & Mason 1993; Guha 1991). The reasons are two-fold. First, in the propositional case, our $C_i\phi$ is effectively equivalent to $ist(C_i, \phi)$ in (Buvač & Mason 1993) ($ist(C_i, \phi)$ means ϕ is true in context C_i). Second, while introducing contexts as terms in a first-order logic leads to a very expressive logic (Guha 1991), it also leads to a very complex logic. The advantage of contexts as terms is that it allows reasoning about the contexts within the logic. However, as we shall see, much of what we want to say about contexts, and relations between contexts, can be easily stated and used in a meta-theory, and communicated to the logic via axiom schemas. Hence, it is worthwhile to investigate the properties of a simpler logic.

Propositional logic of contexts

The propositional logic of contexts is a propositional modal logic in which we introduce a modal operator, C_i , to represent the i^{th} context. Viewing contexts as modal operators immediately provides us with one of the important properties of contexts: the wffs in different contexts can be mutually incoherent. Different contexts can use the same propositions with different meanings. The wffs in different contexts can be mutually inconsistent (though the contexts themselves remain consistent). Different perspectives can be easily accommodated using different contexts.

Syntax

Given a set $\mathcal{P} = \{P_1, \dots, P_m\}$ of propositions and $\mathcal{C} = \{C_1, \dots, C_n\}$ of context operators, the set of all wffs of the logic is defined by the usual rules of propositional modal logic. However, not all these wffs are *meaningful* because different contexts can use different *vocabularies*. For example, the propositions used to describe a forest ecosystem will be quite different from the propositions used to describe Titan’s atmosphere;

any reference to photosynthesis in the latter context is meaningless. Note that, in general, a context's vocabulary may itself be context dependent. For example, a context describing an Eskimo's beliefs would include multiple words for snow, while the vocabulary of the same context from the point of view of a person living in the tropics would include just one word for snow. Allowing a context's vocabulary to be context dependent leads to a slightly more complex semantics and axiomatization (cf. (Buvač & Mason 1993)). As we will see, our use of contexts does not need context dependent context vocabularies. Hence, we have simplified the presentation by assuming that context vocabularies are context independent.

The vocabulary of a context C_i is the set of propositions that can be used immediately within the scope of C_i . (A proposition is immediately within the scope of C_i when it is within the scope of C_i but not within the scope of another context operator, C_j , which is within the scope of C_i .) Formally, the vocabulary of the contexts in \mathcal{C} is defined by a function $\text{voc} : \mathcal{C} \rightarrow 2^{\mathcal{P}}$ that maps a context into a subset of \mathcal{P} . The *language* of C_i is just the set of propositional wffs that can be constructed using the propositions in $\text{voc}(C_i)$. A wff ϕ is meaningful with respect to voc if all occurrences of propositions in ϕ are consistent with voc , i.e., if p_i occurs immediately within context C_i , then $p_i \in \text{voc}(C_i)$.

Semantics and axiomatization

The basic semantics of the logic are standard possible worlds semantics. A model M of the logic defined by $(\mathcal{P}, \mathcal{C}, \text{voc})$ is a Kripke structure (W, R, π) , where W is the set of possible worlds, $R = \{R_1, \dots, R_n\}$ is a set of binary accessibility relations on W , and $\pi : W \times \mathcal{P} \rightarrow \{T, F\}$ is the valuation function that assigns truth values to the propositions in \mathcal{P} at each of the worlds in W . Satisfaction at a world $w \in W$ for meaningful wffs is then defined in the standard way with R_i being the accessibility relation for context C_i .

Let \mathcal{M}_n denote the class of all Kripke structures for n contexts. A meaningful wff ϕ is *valid* in \mathcal{M}_n iff it is satisfied at every world of every structure in \mathcal{M}_n . The basic axiomatization, \mathcal{C}_n , of the set of valid meaningful wffs is the usual axiomatization of propositional modal logic (e.g., the axiomatization K_n in (Halpern & Moses 1992)), appropriately restricted to meaningful wffs as follows:

(A1) All meaningful instances of propositional tautologies

(A2) All meaningful wffs of the form

$$(C_i\phi \wedge C_i(\phi \Rightarrow \psi)) \Rightarrow C_i\psi, \quad 1 \leq i \leq n$$

(R1) From $\vdash \phi$ and $\vdash \phi \Rightarrow \psi$ infer $\vdash \psi$

(R2) From $\vdash \phi$ infer $\vdash C_i\phi$ if $C_i\phi$ is meaningful

The usual proof of the soundness and completeness of this axiomatization with respect to \mathcal{M}_n (cf. (Halpern

& Moses 1992)) applies here with slight modifications. Detailed proofs of all results in this paper can be found in (Nayak 1994).

Multiple domain theories

The axiomatization \mathcal{C}_n is weak in the sense that it does not significantly restrict the properties of contexts, or their interrelations. Additional axioms are needed to tailor the logic to particular applications. We now introduce an axiomatization, \mathcal{F}_n , tailored to reasoning about multiple domain theories.

Axiomatization \mathcal{F}_n

The essential element in reasoning about multiple domain theories is that there is no need to reason about nested contexts. Nested contexts are useful when the properties of a context are different depending on the point of view. We have already seen how the vocabulary of a context can itself be context dependent. Another important example is belief contexts, where the beliefs of an agent differ depending on the point of view (my beliefs about your beliefs are almost certainly different from your beliefs). Nested contexts are useful for expressing ignorance about, and for hiding, properties of contexts.

Since our overall goal is to choose amongst multiple domain theories, we see no need to hide any information. On the contrary, we want to be able to use all available information to make the best possible choice. It is due of this that we chose to assume that a context's vocabulary is context independent. To ensure that the properties of contexts are context independent, we introduce the following two axioms in \mathcal{F}_n :

(A3) $C_i\phi \Rightarrow C_jC_i\phi$ for $1 \leq i, j \leq n$

(A4) $\neg C_i\phi \Rightarrow C_j\neg C_i\phi$ for $1 \leq i, j \leq n$

A3 and A4 are generalizations of the *positive introspection* ($C_i\phi \Rightarrow C_iC_i\phi$) and *negative introspection* ($\neg C_i\phi \Rightarrow C_i\neg C_i\phi$) axioms, respectively. They ensure that every context knows about what every other context does and does not know, i.e., the facts true in a context are context independent. To complete the axiomatization \mathcal{F}_n , we require all contexts to be consistent (this axiom is commonly called **D**):

(A5) $C_i\phi \Rightarrow \neg C_i\neg\phi$ for $1 \leq i \leq n$

In summary, the axiomatization \mathcal{F}_n augments \mathcal{C}_n with axioms A3, A4, and A5. Note that \mathcal{F}_n does not include the axiom **T** ($C_i\phi \Rightarrow \phi$) since we do not require a context's facts to be true. This may seem surprising since these contexts are domain theories, and hence ought to include only true domain facts. However, most useful domain theories are only *approximations* of the domain, and hence include facts that are not strictly true. For example, a context may assume that all gases are ideal gases. Since gases are not really ideal, the predictions of such a context are not strictly

true. However, in practice, such contexts are both useful and common.

The axiomatization \mathcal{F}_n achieves our goal of not hiding any information in the following sense: any meaningful wff is equivalent to a wff with no nested contexts. Let a *propositional literal* be either a proposition in \mathcal{P} or its negation, and a *propositional clause* be a disjunction of propositional literals. Let a *contextual literal* be a context operator or its negation applied to a propositional clause. Let a *contextual clause* be a disjunction of contextual literals, and a *general clause* be a disjunction of propositional and contextual literals. Finally, let a wff be in *conjunctive normal form* (CNF) iff it is a conjunction of general clauses.

Lemma 1 *Every meaningful wff, ϕ , of \mathcal{F}_n is equivalent to a meaningful wff, ϕ' , in CNF, i.e., $\mathcal{F}_n \vdash \phi \Leftrightarrow \phi'$.*

Since a wff in CNF has no nested contexts, it follows that every meaningful wff is equivalent to a meaningful wff without nested contexts.

Correspondences

To assist in the complexity analysis of \mathcal{F}_n , we investigate the correspondences between its axioms and the accessibility relations of its models. It is well known that the axiom A5 corresponds to the accessibility relations being *serial*. (An accessibility relation R_i is serial on the set W if for all $w \in W$ there is a $w' \in W$ such that $(w, w') \in R_i$.) The correspondences of A3 and A4 are generalizations of transitive and euclidean accessibility relations, respectively. Let us say that the set $R = \{R_1, \dots, R_n\}$ of accessibility relations on W is *hyper-transitive* iff for all $x, y, z \in W$ and $1 \leq i, j \leq n$ we have:

$$(x, y) \in R_j \wedge (y, z) \in R_i \Rightarrow (x, z) \in R_i \quad (1)$$

R is said to be *hyper-euclidean* iff for all $x, y, z \in W$ and $1 \leq i, j \leq n$ we have:

$$(x, y) \in R_j \wedge (x, z) \in R_i \Rightarrow (y, z) \in R_i \quad (2)$$

One can show that axiom A3 corresponds to a hyper-transitive accessibility relation, while A4 corresponds to a hyper-euclidean accessibility relation.

The above correspondence results can be used to show that, in the case of \mathcal{F}_n , we can further restrict our attention to particularly simple Kripke structures with one distinguished world, intuitively describing the “real” world, and a set of worlds for each context, which are the worlds considered possible by that context in every world. More formally, say that ϕ is \mathcal{F}_n -consistent iff $\mathcal{F}_n \not\vdash \neg\phi$. We have:

Lemma 2 *If a meaningful wff ϕ is \mathcal{F}_n -consistent, then ϕ is satisfiable in a structure $M = (W, R, V)$ such that $W = \{w_0\} \cup_{1 \leq i \leq n} W_i$ and $R_i = W \times W_i$, $1 \leq i \leq n$.*

The world w_0 is the “real” world, and the worlds W_i are the worlds that context C_i considers possible in every world.

Complexity

Let us now consider the complexity of the satisfiability problem of \mathcal{F}_n . Since \mathcal{F}_n contains propositional logic, satisfiability is certainly NP-hard. Furthermore, it is easy to see that satisfiability is NP-hard both in the number of propositions and the number of context operators. To show that satisfiability is NP-complete, we can use Lemma 2 to show that satisfiable \mathcal{F}_n wffs are satisfiable in Kripke structures with “very few” states. In the following, let $|\phi|$ denote the length of ϕ .

Lemma 3 *A meaningful wff, ϕ , of \mathcal{F}_n is satisfiable iff it is satisfiable in a Kripke structure with at most $|\phi|$ states.*

The following theorem is an immediate consequence:

Theorem 1 *The satisfiability problem for \mathcal{F}_n is NP-complete.*

The above theorem is interesting because the satisfiability problems of other common modal logics with $n > 1$ modal operators (K_n , T_n , $S4_n$, $S5_n$, $KD45_n$) are all PSPACE-complete (Halpern & Moses 1992).

While the above theorem applies to satisfiability of arbitrary wffs, in practice we are often interested in wffs of a restricted syntactic form. Say that a clause is *weakly Horn* if it has at most one positive propositional or contextual literal. For example, $\neg p_1 \vee \neg C_1(p_2 \vee p_3) \vee C_2 p_4$ is a weakly Horn clause. Say that a clause is *strongly Horn* if it is weakly Horn, and each contextual literal in the clause is the result of applying a context operator (or its negation) to a propositional Horn clause. For example, $\neg p_1 \vee \neg C_1(\neg p_2 \vee p_3) \vee C_2(\neg p_1 \vee p_4)$ is a strongly Horn clause. Weak and strong Horn clauses allow us to infer facts from one context to another, while restricting the forms of disjunction that can be stated. These restrictions buy us computational efficiency as follows:

Theorem 2 *Let Σ be a set of clauses of \mathcal{F}_n , $|\mathcal{C}|$ be the number of contexts in \mathcal{C} , $|\mathcal{P}|$ the number of propositions in \mathcal{P} , and P_{max} be the number of propositions in the vocabulary of the context with the largest vocabulary.*

1. *if the clauses in Σ are weakly Horn then deciding the satisfiability of Σ is polynomial in $|\mathcal{C}|$, but remains NP-hard in $|\mathcal{P}|$;*
2. *in part 1, if the clauses are contextual clauses, then satisfiability is NP-hard in P_{max} (rather than in $|\mathcal{P}|$); and*
3. *if the clauses in Σ are strongly Horn then the satisfiability of Σ can be decided in time polynomial in both $|\mathcal{C}|$ and $|\mathcal{P}|$.*

A procedure analogous to bottom-up evaluation in deductive databases (Ullman 1988) allows us to prove the above theorem. This theorem is the central computational justification for the use of contexts in partitioning large KBs to focus reasoning. In particular, item 2 in the theorem shows the benefits of breaking

up a large KB with a single large vocabulary, into a number of smaller contexts with smaller vocabularies.

Relations between contexts

To effectively use multiple contexts, it is important to represent and reason about relations between contexts and the assumptions underlying contexts, and to lift axioms between contexts. Since our logic does not treat contexts as terms, the representation and use of context properties must be done in a meta-theory, and communicated to the logic via axiom schemas. (In the following, a context's theory is just the set of propositional wffs true in the context.)

The simplest relation between contexts is the *weaker than* relation. A context C_i is weaker than a context C_j if every fact true in C_i is also true in C_j , i.e., every wff of C_i is lifted into C_j . Naturally, C_i can be weaker than C_j only if C_i 's vocabulary is a subset of C_j 's vocabulary. We capture this axiomatically using the following axiom schema:

(A6) $C_i\phi \Rightarrow C_j\phi$ for all ϕ such that $C_i\phi$ is meaningful

We have already seen how lifting all axioms from one context into another facilitates the construction and maintenance of a large KB. Furthermore, meta-theoretic knowledge of the weaker than relation is useful for focusing reasoning. For example, to prove $C_j\phi$, it suffices to prove $C_i\phi$. This is particular useful if the theory of C_i is Horn. Similarly, to prove $\neg C_i\phi$, it suffices to prove $\neg C_j\phi$.

Compositional modeling (Falkenhainer & Forbus 1991; Nayak, Joskowicz, & Addanki 1992) provides a variant of the weaker than relation which involves explicitly representing the assumptions underlying a context. In compositional modeling, a model fragment is a piece of knowledge that is applied if the assumptions underlying it are true. We can represent this by introducing a context C_i for each model fragment m_i . Let A_i denote the conjunction of the assumptions underlying m_i ,² and let C be a problem solving context, representing the domain description composed out of the applicable model fragments. The relationship between A_i , C_i , and C is represented by the following axiom schema:

$A_i \Rightarrow (C_i\phi \Rightarrow C\phi)$ for all ϕ s.t. $C_i\phi$ is meaningful

i.e., the axioms of C_i are lifted into the problem solving context if the assumptions underlying C_i hold. Additional relations between assumptions are expressed using appropriate wffs. To answer a user query in C , Falkenhainer and Forbus provide a meta-theoretic algorithm to compose a "simplest" theory for C by selecting an appropriate set of consistent assumptions.

By our definition, C_i can be weaker than C_j only if C_i 's vocabulary is a subset of C_j 's vocabulary. However, this is not always the case, and yet one can often

²In (Falkenhainer & Forbus 1991) an assumption is just a proposition. In general, it can be any meaningful wff.

say that everything expressible in one context is also expressible in another, e.g., earlier we saw that the proposition $at(jmc, stanford)$ in C_1 is expressible as the proposition $at(jmc, stanford, t)$ in C_2 . We capture this using an *interpretation function*, based upon the interpretation between theories discussed in (Enderton 1972) and following the semantic theory of abstractions presented in (Nayak & Levy 1994). We say that the vocabulary, $voc(C_i)$, of context C_i can be interpreted in C_j iff there is an interpretation function f that assigns to each $p \in voc(C_i)$ a wff $f(p)$ in the language of C_j . The interpretation function f is intended to state that p in C_i "expresses the same thing" as $f(p)$ in C_j , e.g., $f(at(jmc, stanford)) = at(jmc, stanford, t)$. It can be extended to wffs in the language of C_i in the natural way: $f(\phi \wedge \psi) = f(\phi) \wedge f(\psi)$, $f(\neg \phi) = \neg f(\phi)$.

Using interpretation functions, and following the terminology in (Giunchiglia & Walsh 1992), we define theorem decreasing (TD), theorem conserving (TC), and theorem increasing (TI) abstractions with the following axiom schemas (in all cases C_i is the more abstract context and ϕ is any wff in the language of C_i):

(A7) $C_i\phi \Rightarrow C_j f(\phi)$ (TD-abstraction)

(A8) $C_i\phi \Leftrightarrow C_j f(\phi)$ (TC-abstraction)

(A9) $C_i\phi \Leftarrow C_j f(\phi)$ (TI-abstraction)

TD-abstractions are like weaker contexts except that they allow vocabulary changes. TC-abstractions are the strongest possible TD-abstractions. In these cases, f specifies how a wff ϕ is to be lifted from C_i to C_j . One can see that specializing a context to a particular time, location, speaker, or hearer will correspond to a TD or a TC abstraction. Other common relations between contexts, such as structural or behavioral abstractions are also TD/TC-abstractions.

Meta-theoretic knowledge of the TD/TC-abstraction relation can be used to control diagnostic reasoning, where the goal is to find a theory (context) that is consistent with the observations. The diagnostic strategy is based on the observation that axioms A7 and A8 ensure that if C_i is inconsistent with the observations, then so is C_j (see (Struss 1992)).

TI-abstractions are best viewed as TC-abstractions under certain simplifying assumptions: adding the simplifying assumptions to the context C_j increases the theorems of C_j and makes C_i a TC-abstraction of C_j . In particular, if ψ is the simplifying assumption underlying the TI-abstraction, then we can incorporate ψ using the following axiom schema:

(A10) $C_i\phi \Leftrightarrow C_j(\psi \Rightarrow f(\phi))$

i.e., ϕ is lifted from C_i into C_j as $\psi \Rightarrow f(\phi)$. For example, the ideal electrical conductor context is a TI-abstraction of the electrical resistor context under the assumption that the resistance is zero. As with TD/TC-abstractions, TI-abstractions can also be used for meta-theoretic control of diagnosis. The difference is that if C_i is inconsistent with the observations, either

C_j is inconsistent or the simplifying assumption ψ does not hold (see (Struss 1992)). (Giunchiglia & Walsh 1992) contains a survey of the uses of TI-abstractions in diverse areas including planning, theorem proving, and common sense reasoning.

The logic itself does not provide a means for transcending its current context (nor, for that matter, can any other logic). However, the process of transcending a context C_j to a context C_i is equivalent to constructing the interpretation function f and the simplifying assumption ψ in axiom A10. For example, in transcending the context of Newtonian mechanics, we need to specify how Newtonian mechanics is interpreted in quantum mechanics (f), and under what simplifying assumptions it holds (ψ).

The potential benefits of using the above relations, and thereby having to introduce the above axiom schemas, does not increase the worst-case complexity of satisfiability. This is a consequence of correspondences between the above axiom schemas and a model's accessibility relations. Say that a world w_2 is an f -abstraction of a world w_1 (denoted by $f-abs(w_1, w_2)$) iff for every proposition $p \in voc(C_i)$, p is true at w_2 iff $f(p)$ is true at w_1 . We have the following correspondences (w, w_1, w_2 are worlds; R_i and R_j are accessibility relations corresponding to C_i and C_j , respectively):

Theorem 3 *A6 is a sound and complete axiomatization of the class of Kripke structures in which $R_j \subseteq R_i$. A7 is a sound and complete axiomatization of the class of Kripke structures in which for all w, w_1*

$$(w, w_1) \in R_j \Rightarrow (\exists w_2 (w, w_2) \in R_i \wedge f-abs(w_1, w_2))$$

A9 is a sound and complete axiomatization of the class of Kripke structures in which for all w, w_2

$$(w, w_2) \in R_i \Rightarrow (\exists w_1 (w, w_1) \in R_j \wedge f-abs(w_1, w_2))$$

The correspondence for A8 is the combination of the ones for A7 and A9. The correspondence for A10 is like the one for A8, except that only those R_j accessible worlds in which ψ is satisfied are considered. Since the above properties can be checked in polynomial time, satisfiability remains in NP.

First-order logic of context

In this section we briefly discuss the extensions needed for a first-order logic of contexts. The syntax is similar to traditional first-order modal logics, appropriately restricted to handle different context vocabularies. The vocabulary of a context now includes constants, functions, and relations. Functions and relations can have different arities in different contexts; occurrences of the same relation (function) with different arities just correspond to different relations (functions). The semantics of first-order modal logics are essentially possible worlds semantics, where each world is a first-order model. The primary difficulty with choosing a satisfactory semantics is related to deciding how the first-order

models at each world relate to each other. In particular, two decisions need to be made: (a) how are the domains at the worlds related; and (b) how are the term interpretations at the worlds related.

The simplest first-order modal logic is obtained by assuming that (a) domains at all worlds are identical; and (b) terms are *rigid*, i.e., term denotations are the same at every world. This semantics is easily axiomatized by adding the principles of first-order logic to the principles of a propositional modal logic (e.g., C_n or \mathcal{F}_n) and the *Barcan formula* ($\forall x C_i \phi \Rightarrow C_i \forall x \phi$) (Garson 1984). However, this is inadequate as a semantics for contexts. First, this semantics requires that all contexts have the same domain, which is clearly inadequate, e.g., the domain of forest ecosystems is clearly different from the domain of Titan's atmosphere. Second, since this semantics allows only rigid terms, it precludes the satisfactory treatment of indexicals, i.e., terms like "I," and "now" whose denotations are clearly context dependent.

We address these shortcomings as follows. First, we let different worlds have different domains, so that different contexts can have different domains. Second, we allow the denotation of terms to be context dependent. Note that this differs from many traditional first-order modal logics where the denotations of terms are *world dependent* (Garson 1984). There is no need to make term denotations world dependent since it is perfectly natural to have rigid term denotations *within* each context (i.e., in all worlds considered possible by the context); it is just that we don't want rigid term denotations *across* contexts.

Space restrictions preclude a detailed description of the resulting logic (see (Nayak 1994)). However, it is worth noting that any sentence, ϕ , of this logic is equivalent to a sentence, ϕ' , of the simpler logic with fixed domains and rigid terms, where ϕ' is the result of (a) replacing every occurrence of $\forall x \dots$ in ϕ with $\forall x E(x) \Rightarrow \dots$, where E is a special predicate denoting the existents at a world; and (b) every constant d and function f occurring immediately within context C is replaced by a constant d^C and function f^C , respectively, in ϕ' . This equivalence is useful because when we extend \mathcal{F}_n to a first-order logic, ϕ' can be converted into a canonical form similar to the one in Lemma 1. Furthermore, we can define and use abstractions as discussed in (Nayak & Levy 1994).

Related work

McCarthy first noted the importance of contexts to common sense reasoning (McCarthy 1987; 1993). He argues that any axiom has limited generality since it is true only in a certain context, but that overly general axiomatizations are often inconvenient. He suggests that formalizing the notion of context is a way out of this dilemma. Guha built upon McCarthy's ideas and developed a first-order logic of contexts in which contexts are incorporated as terms (Guha 1991). Im-

portant contributions of his work include the development of default lifting axioms and the presentation of a large array of examples of the use of contexts in CYC (Lenat & Guha 1990). The primary difference between this work and ours is that we represent contexts as modal operators rather than as terms, leading to a simpler, albeit less expressive, logic. However, as we have shown, important relations between contexts can still be expressed using axiom schemas.

Buvač and Mason develop a propositional logic of context, and provide soundness and completeness results (Buvač & Mason 1993). Following Guha, they introduce contexts as arguments to *ist*: $ist(C_i, \phi)$ says that ϕ is true in context C_i . As mentioned earlier, in the propositional case, this is effectively equivalent to introducing C_i as a modal operator. They allow context vocabularies to be context dependent, and hence need to define satisfaction and provability with respect to context sequences. While we could have done the same, our main interest (the logic \mathcal{F}_n) does not require context dependent context vocabularies, and so we chose a simpler formalism.

Shoham discusses the ubiquity of contexts with a series of examples drawn from a variety of domains (Shoham 1991). He discusses various relations between contexts, and introduces a set of 14 benchmark sentences against which one may evaluate different semantics for contexts. He does not try and pin down “the right semantics” for contexts, but does suggest introducing contexts as propositions, and representing “ p is true in context q ” with the (material, intuitionistic, or relevant) implication $q \rightarrow p$.

Conclusions

In this paper we have developed a simple logic of context that is well suited for representing and reasoning with multiple domain theories. A key feature of this logic is that we introduce contexts as modal operators, rather than as terms. Our analysis of the computational properties of the resulting logic provided us with the central computational justification for the use of contexts. We showed that, for the purposes of representing and reasoning with multiple theories, this logic is able to effectively handle common uses of contexts. This seems to suggest that the difficulties in understanding contexts lie not so much with their logical properties, but with their heuristic properties. What are the useful and common types of contexts? How does one decide which context to use in a particular situation? How does one detect that the current context is inadequate? How does one transcend the current context? To investigate these heuristic questions, we have augmented SIGMA’s frame representation language with contexts, and are currently evaluating the utility of using contexts in SIGMA. We are also evaluating the utility of contexts in developing practical diagnostic engines for complex physical systems.

Acknowledgements

I would specially like to thank Surajit Chaudhuri and Alon Levy for helpful discussions and useful comments on earlier drafts of this paper. Thanks also to Saša Buvač, Adam Grove, R. V. Guha, Mike Lowry, David Thompson, and the anonymous reviewers of this paper.

References

- Buvač, S., and Mason, I. A. 1993. Propositional logic of context. In *Proceedings of AAAI-93*.
- Enderton, H. B. 1972. *A Mathematical Introduction to Logic*. Academic Press, Inc.
- Falkenhainer, B., and Forbus, K. D. 1991. Compositional modeling: Finding the right model for the job. *Artificial Intelligence* 51.
- Garson, J. W. 1984. Quantification in modal logic. In Gabbay, D., and Guenther, F., eds., *Handbook of Philosophical Logic, Volume II*. D. Reidel Publishing Company.
- Giunchiglia, F., and Walsh, T. 1992. A theory of abstraction. *Artificial Intelligence* 57(2-3).
- Guha, R. V. 1991. *Contexts: A Formalization and Some Applications*. Ph.D. Dissertation, Stanford University. Technical Report No. STAN-CS-91-1399.
- Halpern, J. Y., and Moses, Y. 1992. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence* 54(3).
- Keller, R. M.; Rimon, M.; and Das, A. 1994. A knowledge based prototyping environment for construction of scientific modeling software. *Automated Software Engineering* 1(1).
- Lenat, D. B., and Guha, R. V. 1990. *Building Large Knowledge-Based Systems*. Addison Wesley.
- McCarthy, J. 1987. Generality in artificial intelligence. *CACM* 30(12).
- McCarthy, J. 1993. Notes on formalizing context. In *Proceedings of IJCAI-93*.
- Nayak, P. P., and Levy, A. Y. 1994. A semantic theory of abstractions: Preliminary report. In *Proceedings of the Workshop on Theory Reformulation and Abstraction*, 1994.
- Nayak, P. P.; Joskowicz, L.; and Addanki, S. 1992. Automated model selection using context-dependent behaviors. In *Proceedings of AAAI-92*.
- Nayak, P. P. 1994. Representing multiple theories. Technical report, NASA Ames Research Center.
- Shoham, Y. 1991. Varieties of context. In Lifschitz, V., ed., *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*. Academic Press.
- Struss, P. 1992. What’s in SD? Towards a theory of modeling for diagnosis. In Hamscher, W.; Console, L.; and de Kleer, J., eds., *Readings in Model-Based Diagnosis*. Morgan Kaufmann.
- Ullman, J. D. 1988. *Principles of Database and Knowledge-Base Systems, Vol. I*. Computer Science Press.