A Logical Toolbox for Knowledge Approximation

(preliminary version)

Frédéric Koriche and Jean Sallantin

LIRMM, UMR 5506, Université Montpellier II CNRS
161, rue Ada 34392 Montpellier Cedex 5, France
koriche@lirmm.fr, js@lirmm.fr

Abstract

It is well-known that the logicist approach to agency is confronted with both epistemological and heuristic problems. On the one hand, the agent's model must be logically adequate: it must provide us a clear picture of what the agent is, and is not, able to deduce from its background knowledge. On the other hand, the agent's program must be adequate in practice: it must generate useful conclusions from input data and given the computational resources that are actually available. In actuality, the agent's need for heuristic adequacy has strong epistemological consequences. Based on this argument, this paper proposes a framework which is based on the paradigm of knowledge approximation and that is flexible enough to incorporate heuristic strategies used in satisfiability algorithms. The framework is used as a "logical toolbox" for modelling resource-bounded agents that have different operational means at their disposal to approximate knowledge. The toolbox consists in a family of *relative relevance logics* which are semantically founded on the notion of *resource* and that include interesting features, such as incremental reasoning and dual approximations.

Keywords: Resource-bounded agents, knowledge approximation, satisfiability, relative relevance logics.

1 Introduction

From early on [23], it has long been recognized that the logicist approach to agency is confronted with two orthogonal issues: *epistemological problems* and *heuristic problems*. On the one hand, the agent's model must be logically adequate: it must represent knowledge and tells us what the agent is able to deduce, or not, from its background knowledge. On the other hand, the agent's program must be adequate in practice: it must determine useful conclusions from input data and given the computational resources that are actually available. One primary lesson of research over the past decade is that epistemological and heuristic problems are *not* separable: each element informs the other [14]. The main reason is that virtually every reasoning problem is very much demanding from a computational point of view. In particular, if the representation formalism is propositional logic, then knowledge deduction is a coNP-complete problem. This barrier of complexity implies that heuristic adequacy has strong epistemological consequences; namely, the strategies involved in solving reasoning problems are central to the formal specifications of reasoning agents as a whole.

Following this line of research, this paper is inspired by two concerns in the setting of propositional reasoning.

Knowledge approximation. This paradigm extends the logicist prescription to reasoning by allowing the agent to return approximate solutions to a given problem. Informally, an approximate solution is a *maybe* answer which provides a middle ground between the exact *yes* and *no* answers. The two major forms of knowledge approximation include *lower approximations* which give *yes* and *maybe no* answers, and *upper approximations* which provide *no* and *maybe yes* answers. The main difficulty in knowledge approximation stems from the fact that, in contrast to numerical optimization problems, there is no explicit metric that defines the accuracy of maybe answers. Therefore, the very notion of approximation has to be grounded on an epistemological basis.

To this end, there have been a number of attempts at devising logics for approximate reasoning, using either proof-theoretic or model-theoretic techniques. On the proof-theoretic side, for example, Konolige in [20] obtains approximate forms of inference by eliminating certain axioms or inference rules from classical propositional logic. Other ways of capturing approximate deduction include Dalal's [5] anytime reasoning systems, Kaplan and Schubert's [19] belief machines and Halpern, Moses and Vardi's [16] logic of algorithmic knowledge. On the model-theoretic side, the work has concentrated on multi-valued logics, especially a fragment of Belnap's relevance logic [2]. The most significant approaches include Levesque's [22] architecture for limited reasoning, Delgrande's [9] framework for explicit belief, Fagin and Halpern and Vardi's [10] nonstandard epistemic logics and Calodi and Schaerf's model of approximate reasoning [24].

Satisfiability. The overall goal of this research field is to find reasoning algorithms with good average performance. In this study, we focus on systematic algorithms that both determine satisfiability and unsatisfiability (and hence deduction), as opposed to local search algorithms, such as hill-climbing techniques, which only determine satisfiability.

The two major types of procedures include *resolution-based algorithms* and *enumeration-based algorithms*. In resolution, the agent tries to prove unsatisfiability of a clausal theory by deriving the empty-clause from it; if the empty clause cannot be derived, then the theory is satisfiable. A well-known resolution strategy is *directional resolution*, first proposed by Davis and Putnam in [7] and further investigated in [8, 15]. This strategy is based on the observation that a restricted form of resolution performed along a sequence of propositional atoms is sufficient for deciding satisfiability. At any computation step, the algorithm selects an atom, resolves all clauses containing the atom and afterwards deletes the parent clauses, keeping only the newly generated resolvents and those clauses that have not yet been used. In enumeration, the agent tries to prove satisfiability of a clausal theory by generating an interpretation which satisfies the theory; if such an interpretation is not found, then the theory is unsatisfiable. The prototypical strategy here is *backtracking search*, originally introduced by Davis, Logemann, and Loveland in [6], and further studied, notably, in [4, 18, 25]. The essential idea of this strategy is based on the observation that for some atoms, the agent can rapidly decide what values a satisfying interpretation must have. When this decision becomes too complex, the agent simply tests for an atom both possible truth values true and false. The strategy can be represented by a recursive function which performs a depth-first search through the space of possible interpretations.

These strategies have been a focus of extensive research for many years, in both theoretical and empirical sides (e.g. [4, 8, 18, 25]), suggesting perspectives for incorporating them in formal models of agents.

Contributions. We present a framework which is epistemologically based on the paradigm of knowledge approximation and that is flexible enough to model heuristic strategies used in satisfiability algorithms. Our framework can be seen as a "logical toolbox" for specifying resource-bounded agents that have different operational means at their disposal to approximate their knowledge. From a conceptual point of view, the framework is a continuation of the model-theoretic work on approximate reasoning. The toolbox consists in a family of *relative relevance logics* that combines ideas from Belnap's relevance logic [2] with relative modal logics, investigated in rough set theory [1, 11], dynamic logic [17] and boolean modal logic [12, 13]. Based on this formal machinery, the toolbox incorporates several major features. First, the logics are founded on the notion of *resource* which provides a clear picture to the effort required in computing knowledge. Second, the logics use relative accessibility structures that enable us to model *incremental reasoning*: the quality of approximations increases with the accuracy of resources. Third, the modal operators used in these logics allow us to model *dual approximations*: both lower and upper approximations of knowledge can be returned at any step.

This paper is organized as follows. Section 2 presents the logical toolbox. Section 3 concentrates on the formal specifications of approximate elimination and approximate conditioning. Section 4 compares our framework with other models and concludes the paper.

2 The Toolbox

The toolbox consists in three relative relevance logics of increasing expressivity: $\mathbf{RL}_0$, $\mathbf{RL}_1$and $\mathbf{RL}_2$. For sake of simplicity, we consider a single agent and formulas that do have nested modal operators.

2.1 The logic $\mathbf{RL}_0$

Syntax. The linguistic basis consists of a set P of *primitive propositions* and a finite set A of *primitive resources*. The set of *complex propositions* is defined in the standard way from primitive propositions, the connectives $\wedge$ and $\neg$ and the constant $\top$. Other connectives such as $\supset$, $\equiv$ and the constant $\bot$ are defined in the usual way. In this elementary system, the set of *complex resources* is "flat": it is simply given by A. The set of *sentences* is built up from complex propositions and the following rules: if α is a complex resource and ϕ a complex proposition then $[\alpha]\phi$ is a sentence, if ϕ is a sentence then $\neg\phi$ is a sentence, and if ϕ and ψ are sentences then $\phi \wedge \psi$ is a sentence.

The sentence $\langle\alpha\rangle\phi$ is an abbreviation of $\neg[\alpha]\neg\phi$. Intuitively, a sentence such as $[\alpha]\phi$ is read "the agent knows ϕ given the resources α". Dually, $\langle\alpha\rangle\phi$ is read "the agent considers ϕ as possible given the resources α".

Semantics. The basic building blocks of the semantics are states, which determine the interpretation of propositions, and relative accessibility relations which capture the epistemic nature of resources. An *atom* is a primitive proposition or the constant $\top$ and a *literal* is an atom or its negation. Given a literal l, the *adjunct* of l, denoted l^* is defined as follows: for any atom p, $p^* = \neg p$ and $(\neg p)^* = p$. A *state* s is a set of literals such that $\top \in s$ and $\perp \notin s$. Given a state s, the *adjunct* of s, denoted s^*, is defined as follows: for any literal l, $l \in s^*$ iff $l^* \notin s$. In the following, the set of all states generated from P is denoted S. A *filter* in S is a subset T of S such that for any states s and t, if $s \in T$ and $s \subseteq t$ then $t \in T$. A *prime filter* in S is a filter T in S such that for any states s and t, if $s, t \in T$ then $s \cap t \in T$. Interestingly, the notion of prime filter captures an important structural property, namely, a proposition is true in a primer filter T iff it is satisfied by its "meet" element $\cap T$. A *relative accessibility relation* for $\mathbf{RL}_0$ is a map R from complex resources into binary relations of S such that for every resource α and every state s, $R(\alpha)(s)$ is a prime filter. We are now ready to assign truth values to sentences:

$R, s \models \top$,
$R, s \models p$ iff $p \in s$,
$R, s \models \neg\phi$ iff $R, s^* \not\models \phi$,
$R, s \models \phi \wedge \psi$ iff $R, s \models \phi$ and $R, s \models \psi$,
$R, s \models [\alpha]\phi$ iff $R, t \models \phi$, for all $t \in S$ such that $t \in R(\alpha)(s)$.

A *world* is a state w such that $w^* = w$. A *model* of $\mathbf{RL}_0$ is a pair (R, w) such that R is a relative accessibility relation and w is a world. A sentence ϕ is called *valid* is $\mathbf{RL}_0$(written $\models_{RL_0} \phi$) iff $R, w \models \phi$ for every model (R, w). A sentence ϕ is called *satisfiable* in $\mathbf{RL}_0$ iff $R, w \models \phi$ for some model (R, w). Notice that a formula ϕ is valid in $\mathbf{RL}_0$ iff $\neg\phi$ is not satisfiable in $\mathbf{RL}_0$.

Axiomatization. We now present a sound and complete Hilbert-style system for our logic. In the following, we use the notation ϕ_{CNF} (ϕ_{DNF}) for the conjunctive (disjunctive) normal form of the proposition ϕ. The logic $\mathbf{RL}_0$ is the smallest set of sentences containing classical propositional logic and the following axiom schemata, and that is closed under modus ponens:

(A1) $[\alpha]\phi \equiv [\alpha]\phi_{CNF} \equiv [\alpha]\phi_{DNF}$
(A2) $[\alpha]\phi \wedge [\alpha]\psi \equiv [\alpha](\phi \wedge \psi)$
(A3) $[\alpha]\phi \vee [\alpha]\psi \equiv [\alpha](\phi \vee \psi)$

A sentence ϕ is *provable* in $\mathbf{RL}_0$(written $\vdash_{RL_0} \phi$) iff $\phi \in \mathbf{RL}_0$.

The axiom schema (A1) states that relative modalities respect the standard properties such as absorption, idempotence, commutativity, associativity, and De Morgan's laws. Axiom schemata (A2) and (A3) are the properties of conjunctive and disjunctive knowledge. We remark that, in contrast to standard logics of explicit beliefs [9, 22], all instances of the schema $[\alpha](\phi \vee \psi) \supset [\alpha]\phi \vee [\alpha]\psi$ are theorems of $\mathbf{RL}_0$. This schema captures the prime filter condition of relative accessibility relations.

Theorem 2.1 (completeness). *$\models_{RL_0} \phi$ iff $\vdash_{RL_0} \phi$.*

The soundness proof of the above theorem is obtained by a straightforward inductive argument. The completeness proof is an extension of a standard construction in modal logics [3] of the canonical model. The key idea is to show that any consistent sentence is satisfiable. We begin by extending a given consistent sentence ϕ to a maximally consistent set T by Lindenbaum's lemma. Then we build the canonical model (R_T, w_T) of T. The world w_T is given as follows: for every literal l, $l \in w_T$ iff $l \in T$. The mapping R_T is defined as follows: for every resource α, $s \in R_T(\alpha)(w_T)$ iff $s \models \phi$ for every sentence ϕ such that $[\alpha]\phi \in T$. The remaining proof is built upon two central lemmas. The *construction lemma* claims that the canonical model is a model of $\mathbf{RL}_0$. The *truth lemma* shows that membership in T and truth on the canonical model amount to the same thing. Soundness and completeness results of the remaining logics discussed in this paper are obtained by suitably modifying this proof.

2.2 The logic $\mathbf{RL}_1$

Syntax. The language of $\mathbf{RL}_1$ is obtained from the language of the previous logic by adjoining the conjunction operation to the class of operations acting on resource expressions. The set of complex resources of $\mathbf{RL}_1$ is the smallest set that contains A and that is closed under the connective $\wedge$. Under some reasonable assumptions, the connective $\wedge$ captures the notion of "parallel composition" in dynamic logics [17]. A sentence such as $[\alpha \wedge \beta]\phi$ is read "the agent knows ϕ given the composition of resources α and β".

Semantics. A *relative accessibility relation* for $\mathbf{RL}_1$ is a map R from complex resources into binary relations on S that satisfies the prime filter condition and such that $R(\alpha \wedge \beta) = R(\alpha) \cap R(\beta)$. Interestingly, we remark that for any state s, the prime filter $R(\alpha \wedge \beta)(s)$ is the greatest lower bound of the prime filters $R(\alpha)(s)$ and $R(\beta)(s)$.

Axiomatization. Given a complex resource α, we denote α_{SET} the subset of A defined by induction as follows: $a_{\text{SET}} = \{a\}$, $(\alpha \wedge \beta)_{\text{SET}} = \alpha_{\text{SET}} \cup \beta_{\text{SET}}$. We define the *normal form* of a complex resource α, by the (finite) formula $\bigwedge \alpha_{\text{SET}}$. The logic $\mathbf{RL}_1$ is the smallest extension of $\mathbf{RL}_0$ containing the following axiom schemata:

(A4) $[\alpha]\phi \equiv [\alpha_{NF}]\phi$
(A5) $[\alpha]l \vee [\beta]l \equiv [\alpha \wedge \beta]l$

The axiom schema (A4) summarizes the algebraic properties of the conjunctive connective, such as idempotence, commutativity and associativity. The axiom schema (A5) reflects the following fact: the system knows a literal l under the composite resource $\alpha \wedge \beta$ iff it knows l under the resources α or β, considered separately. It can be easily observed that the axiom of parallel composition $[\alpha]\phi \vee [\beta]\phi \supset [\alpha \wedge \beta]\phi$ is a theorem of the logic $\mathbf{RL}_1$.

Theorem 2.2 (completeness). $\models_{RL_1} \phi$ *iff* $\vdash_{RL_1} \phi$.

2.3 The logic $\mathbf{RL}_2$

Syntax. The linguistic basis of $\mathbf{RL}_2$ is the language obtained from $\mathbf{RL}_1$ enriched by adjoining the disjunction operation to resource expressions. The set of *complex resources* of $\mathbf{RL}_2$ is the smallest set that contains A and that is closed under the connectives $\wedge$ and $\vee$. Intuitively, the connective $\vee$ captures the idea of "nondeterministic choice" advocated in dynamic logics. A sentence such as $[\alpha \vee \beta]\phi$ can be read "the agent knows ϕ given either the resources α or the resources β".

Semantics. A *relative accessibility relation* for $\mathbf{RL}_2$ is a map R from complex resources into binary relations on S such that R is defined according to the conditions of relative accessibility relations for $\mathbf{RL}_1$ and the additional rule: $R(\alpha \vee \beta) = R(\alpha) \cup R(\beta)$. In this setting, we remark that for any state s, the prime filter $R(\alpha \vee \beta)(s)$ is the least upper bound of the prime filters $R(\alpha)(s)$ and $R(\beta)(s)$.

Axiomatization. As previously, we first introduce a canonical presentation for resources. Given a complex resource α, we denote α_{SET} the subset of the powerset of A defined by induction as follows: $a_{\text{SET}} = \{\{a\}\}$, $(\alpha \wedge \beta)_{\text{SET}} = \{\text{A} \cup \text{B} : \text{A} \in \alpha_{\text{SET}} \text{ and } \text{B} \in \beta_{\text{SET}}\}$, and $(\alpha \vee \beta)_{\text{SET}} = \alpha_{\text{SET}} \cup \beta_{\text{SET}}$. The *normal form* of a complex resource α, denoted α_{NF}, is given by the (finite) formula $\bigvee\{\bigwedge \text{A} : \text{A} \in \alpha_{\text{SET}}\}$. The logic $\mathbf{RL}_2$ is the smallest extension of $\mathbf{RL}_1$ containing the following axiom schema:

(A6) $[\alpha]\phi \wedge [\beta]\phi \equiv [\alpha \vee \beta]\phi$

In the logic $\mathbf{RL}_2$ the axiom schema (A3) summarizes the algebraic properties of both conjunctive and disjunctive operators, namely idempotence, associativity, commutativity and distributivity. The schema (A6) is the axiom of nondeterministic choice [17].

Theorem 2.3 (completeness). $\models_{RL_2} \phi$ *iff* $\vdash_{RL_2} \phi$.

3 Knowledge Approximation

To this point we have investigated the "backbone" of propositional knowledge approximation by examining a hierarchy of logics for knowledge and resources. In this section, we illustrate that the components of our toolbox can be instantiated to provide semantics for approximate resolution and approximate enumeration.

3.1 Approximate Resolution

An important heuristic issue of approximate reasoning is to use efficient resolution strategies. We focus here on directional resolution [7, 8]. In this setting, the agent iteratively checks for satisfiability of a clausal theory ϕ by means of an increasing sequence of variables $(p_1, \cdots, p_n)$. At any step i, the agent resolves all clauses in ϕ over the variable p_i and then eliminates remaining clauses containing p_i in the resulting theory. The agent answers "yes" if the resulting theory ϕ is empty, "no" if it contains the empty clause, and "maybe" otherwise.

The logic ARL. Approximate resolution is semantically founded on the logic **ARL**, an instance of the system $\mathbf{RL}_1$. In this logic, the set of atomic resources A is the set of *atoms* generated from P. A *relative accessibility relation* for **ARL** is a r.a.r. R for $\mathbf{RL}_1$ such that for any atom p and any states s and t, we have $t \in R(p)(s)$ iff either $p \in t \cap s$ or $\neg p \in t \cap s$. It is easy to check that, for every atom p and every state s, $R(p)(s)$ is a prime filter in S. The logic **ARL** is the smallest extension of $\mathbf{RL}_1$ containing the additional axiom schemata:

(AR1) $[\alpha]\phi \supset \phi$
(AR2) $[p](q \vee \neg q)$, where $p = q$
(AR3) $\langle p\rangle(q \wedge \neg q)$, where $p \neq q$

The axiom schema (AR1), often called T in the logic literature [3], demonstrates that reasoning under the scope of the modal operator $[\alpha]$ is sound. Axioms (AR2) and (AR3) play a key role in approximate resolution.

More precisely, if the available resources contain the atom p, then the agent necessarily infers the tautology $p \vee \neg p$. Dually, if the available resources do not contain any occurrence of p, then the agent can infer the antilogy $p \wedge \neg p$.

Theorem 3.1 (completeness). *$\models_{ARL} \phi$ iff $\vdash_{ARL} \phi$.*

Properties. Given two resources α and β, α is called *less accurate than β* iff $\alpha_{\text{SET}} \subseteq \beta_{\text{SET}}$. In the following, ϕ is a CNF proposition (clausal theory) and α and β are resources such that α is less accurate that β.

Proposition 3.2 (Monotonicity). *If $[\alpha]\phi$ is satisfiable, then $[\beta]\phi$ is satisfiable, and if $\langle\alpha\rangle\phi$ is unsatisfiable, then $\langle\beta\rangle\phi$ is unsatisfiable.*

Proof. The result is based on the observation that if α is less accurate than β, then $R(\beta) \subseteq R(\alpha)$ for every relative accessibility relation R. Suppose we are given two states s and t such that $t \in R(\beta)(s)$. For every atom $p \in \beta_{\text{SET}}$, either $p \in t \cap s$ or $\neg p \in t \cap s$. Since $\alpha_{\text{SET}} \subseteq \beta_{\text{SET}}$, it follows that for every $p \in \alpha_{\text{SET}}$, either $p \in t \cap s$ or $\neg p \in t \cap s$. Thus, $t \in R(\alpha)(s)$. Now suppose that $[\alpha]\phi$ is satisfiable and $[\beta]\phi$ is unsatisfiable. We show that this leads to a contradiction. Clearly, there is a model (R, w) such that $R, w \models [\alpha]\phi$ and $R, w \not\models [\beta]\phi$. Thus, $R, s \not\models \phi$ for some state $s \in R(\beta)(w)$. Since $R(\beta) \subseteq R(\alpha)$, it follows that $s \in R(\alpha)(w)$. Therefore, $R, w \not\models [\alpha]\phi$, hence contradiction. A dual argument applies to the second part.

Proposition 3.3 (Duality). *If $[\alpha]\phi$ is satisfiable then ϕ is satisfiable, and if $\langle\alpha\rangle\phi$ is unsatisfiable then ϕ is unsatisfiable.*

Proof. By application of the axiom schema (AR1).

Proposition 3.4 (Complexity). *There exists an algorithm for deciding whether $[\alpha]\phi$ and $\langle\alpha\rangle\phi$ is satisfiable, which runs in $O(|\phi| \cdot 2^{|\alpha_{\text{SET}}|})$ time.*

Proof. We examine $[\alpha]\phi$. This sentence is satisfiable iff $R, w \models [\alpha]\phi$ for some model (R, w). The complexity result is based on three observations. First, we remark that there exists exactly one accessibility relation $R(\alpha)$ that satisfies the semantical conditions of **ARL**. Second, based on the prime filter condition of $R(\alpha)$, we observe that $R, w \models [\alpha]\phi$ holds iff $\cap R(\alpha)(w) \models \phi$ holds. As a direct consequence, checking whether a model (R, w) supports the truth of the sentence $[\alpha]\phi$ can be done in $O(|\phi|)$ time. Third and finally, we observe that the number of distinct prime filters of $R(\alpha)$, defined over W is given by $2^{|\alpha_{\text{SET}}|}$. Hence, checking whether $[\alpha]\phi$ is satisfiable can be done in $O(|\phi| \cdot 2^{|\alpha_{\text{SET}}|})$ time. A dual argument applies to $\langle\alpha\rangle\phi$.

3.2 Approximate Enumeration

Another important research avenue in knowledge approximation is to incorporate efficient enumeration techniques. We use a resource-bounded version of backtracking enumeration [6] which can be shown incremental using an iterative deepening search technique. The agent checks for satisfiability of a clausal theory ϕ by means of an increasing sequence of "search trees" $(\alpha_1, \cdots, \alpha_n)$. For any i, the agent answers "yes" if one of the leaves of the tree α_i satisfies ϕ, "no" if all the leaves falsify ϕ, and "maybe" otherwise.

The logic TRL. In order to model tree-like resources, we advocate an instance of the logic $\mathbf{RL}_2$. The set of atomic resources is the collection of literals generated from P. The set of complex resources is defined using the so-called Shannon rules: $\top, \bot$ are complex resources, and if l is a literal and α and β are complex resources, then $(l \wedge \alpha) \vee (l^* \wedge \beta)$ is a complex resource. A *relative accessibility relation* for **TRL** is a r.a.r. for $\mathbf{RL}_2$ such that for any literal l and any states s, t, we have $t \in R(l)(s)$ iff $l \in t \cap s$. It can be easily shown that for every literal l and every state s, $R(l)(s)$ is a prime filter in S. The logic **TRL** is the smallest extension of $\mathbf{RL}_2$ containing the additional axiom schemata:

(TR1) $\langle\alpha\rangle\top \supset ([\alpha]\phi \supset \phi)$
(TR2) $[l]l'$, where $l = l'$
(TR3) $l \supset \langle l \rangle l'$, where $l \neq l'$

The axiom schema (TR1) states that reasoning under the scope of $[\alpha]$ is sound, provided that α is consistent. Axioms (TR2) and (TR3) are essentially (AR2) and (AR3) adapted for tree-like resources.

Theorem 3.5 (completeness). $\models_{TRL} \phi$ *iff* $\vdash_{TRL} \phi$.

Properties. We say that α is *less accurate* than β if α is a subtree of β, i.e. $\forall A \in \alpha_{\mathrm{SET}}$, $\exists B \in \beta_{\mathrm{SET}}$ such that $A \subseteq B$. In the following, ϕ is a CNF proposition and α and β are resources such that α is less accurate that β.

Proposition 3.6 (Monotonicity). *If $\alpha \wedge [\alpha]\phi$ is satisfiable, then $[\beta]\phi$ is satisfiable. If α is valid and $\langle\alpha\rangle\phi$ is unsatisfiable, then $\langle\beta\rangle\phi$ is unsatisfiable.*

Proof. Suppose that $\alpha \wedge [\alpha]\phi$ is satisfiable. It follows that $R, w \models \alpha \wedge [\alpha]\phi$ for some model (R, w). In particular, there exists a term $A \in \alpha_{\mathrm{SET}}$ such that $w \models A$. Since α is less accurate than β, there exists a term $B \in \beta_{\mathrm{SET}}$ such that $A \subseteq B$. As a direct consequence, $R(\beta)(w) \subseteq R(\alpha)(w)$. Now assume that $[\beta]\phi$ is unsatisfiable. Thus, $R, s \not\models \phi$ for some state $s \in R(\beta)(w)$. Since

$R(\beta)(w) \subseteq R(\alpha)(w)$, it follows that $s \in R(\alpha)(w)$. Therefore, $R, w \not\models [\alpha]\phi$, hence contradiction. A dual argument applies to the second part.

Proposition 3.7 (Duality). *If $\alpha \wedge [\alpha]\phi$ is satisfiable then ϕ is satisfiable. If α is valid and $\langle\alpha\rangle\phi$ is unsatisfiable then ϕ is unsatisfiable.*

Proof. Suppose that $\alpha \wedge [\alpha]\phi$ is satisfiable. Thus $R, w \models \alpha \wedge [\alpha]\phi$ for some model (R, w). Clearly, $R(\alpha)(w)$ is not empty and thus $R, w \models \langle\alpha\rangle\top$. By application of axiom schema (TR1), it follows that $R, w \models \phi$. Hence ϕ is satisfiable, as desired. A dual strategy applies to the second part.

Proposition 3.8 (Complexity). *There exists an algorithm for deciding whether $[\alpha]\phi$ and $\langle\alpha\rangle\phi$ is satisfiable, which runs in $O(|\phi| \cdot |\alpha_{\text{SET}}|)$ time.*

Proof. The proof is very similar to that of proposition 3.4. The only difference stems from the fact that the number of distinct prime filters of $R(\alpha)$ defined over W is given by $|\alpha_{\text{SET}}|$.

4 Conclusions

We have introduced a "logical toolbox" for knowledge approximation based on a family of relative relevance logics. We have illustrated that the logics can be instantiated to model several strategies investigated in propositional satisfiability. The toolbox is founded on the notion of resource and includes interesting features such as incremental reasoning and dual approximations.

The overall motivation behind this study is similar to the one presented in [19, 16]. However, these approaches do not focus on the heuristic issues involved in propositional reasoning. In contrast, our toolbox proposes an "algebra" of resource expressions that can be instantiated to specific strategies and tells us what sort of semantics underlies the strategies. Our contribution is also compatible with Dalal's [5] architecture based on unit resolution. However, Dalal's semantics is "endogenous", i.e. founded on an implicit notion of computation. In contrast, our logical paradigm is "exogenous" and explicitly represents the effort involved in computation.

From a technical point of view, our toolbox combines ideas from four-valued logic with relative accessibility relations. To this end, the toolbox extends in several directions previous model-theoretical approaches to limited reasoning [9, 10, 22, 24]. In particular, it can be observed that the logic **ARL**, presented in an earlier version of this work [21], is a direct generalization of Schaerf and Cadoli's [24] model of knowledge approximation.

We believe that our toolbox will be a useful architecture for finer analysis of the connections between knowledge and computation.

References

[1] P. Balbiani and E. Orlowska. A hierarchy of modal logics with relative accessibility relations. *Journal of Applied Non-Classical Logics*, 9:303–328, 1998.

[2] N. D. Belnap. A useful four-valued logic. In *Modern Uses of Multiple-Valued Logic*, pages 8–37. Reidel, 1977.

[3] B. F. Chellas. *Modal Logic, an Introduction.* Cambridge University Press, 1980.

[4] J. M. Crawford and L. D. Auton. Experimental results on the crossover point in random 3-SAT. *Artificial Intelligence*, 81:13–59, 1996.

[5] M. Dalal. Anytime clausal reasoning. *Annals of Mathematics and Artificial Intelligence*, 22(3-4):297–318, 1998.

[6] M. Davis, G. Logemann, and D. Loveland. A machine program for theorem proving. *Communications of the ACM*, 5:394–397, 1962.

[7] M. Davis and H. Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7:201–215, 1960.

[8] R. Dechter and I. Rish. Directional resolution: The Davis-Putnam procedure, revisited. In *Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning*, pages 134–145, 1994.

[9] J. P. Delgrande. A framework for logics of explicit belief. *Computational Intelligence*, 11(1):47–88, 1995.

[10] R. Fagin, J. Y. Halpern, and M. Y. Vardi. A nonstandard approach to the logical omniscience problem. *Artificial Intelligence*, 79(2):203–240, 1995.

[11] L. Farinas del Cerro and E. Orlowska. DAL: a logic for data analysis. *Theorical Computer Science*, 36:251–264, 1985.

[12] G. Gargov and S. Passy. A note on boolean modal logic. In *Mathematical Logic*, pages 299–309. Plenum Press, 1990.

[13] G. Gargov, S. Passy, and T. Tinchev. Modal environment for boolean speculations. In *Mathematical Logic and its Applications*, pages 253–263. Plenum Press, 1987.

[14] M. L. Ginsberg. Epistemological and heuristic adequacy revisited. *ACM Computing Surveys*, 27(3):331–333, 1995.

[15] A. Goert. Davis-Putnam resolution versus unrestricted resolution. *Annals of Mathematics and Artificial Intelligence*, 6:169–184, 1992.

[16] J. Y. Halpern, Y. Moses, and M. Y. Vardi. Algorithmic knowledge. In *Proceedings of the 5th Conference on Theorical Aspects of Rationality and Knowledge*, pages 255–266. Morgan Kaufmann, 1994.

[17] D. Harel. Dynamic logic. In *Handbook of Philosophical Logic*, volume 3, pages 497–604. Reidel, 1984.

[18] R. Jesroslow and J. Wang. Solving propositional satisfiability problems. *Annals of Mathematics and Artificial Intelligence*, 1:167–187, 1990.

[19] A. N. Kaplan and L. K. Schubert. A computational model of belief. *Artificial Intelligence*, 120:119–160, 2000.

[20] K. Konolige. *A Deduction Model of Belief.* Morgan Kauffman, 1986.

[21] F. Koriche. A logic for anytime deduction and anytime compilation. In *Logics in Artificial Intelligence*, volume 1489, pages 324–342. Springer Verlag, 1998.

[22] H. J. Levesque. A logic of implicit and explicit belief. In *Proceedings of the 6th National Conference on Artificial Intelligence*, pages 198–202, 1984.

[23] J. McCarthy and P.J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. *Machine Intelligence*, 4:463–502, 1969.

[24] M. Schaerf and M. Cadoli. Tractable reasoning via approximation. *Artificial Intelligence*, 74:249–310, 1995.

[25] H. Zhang and M. E. Stickel. Implementing the Davis-Putnam method. *Journal of Automated Reasoning*, 24(1/2):277–296, 2000.