

Complete Axiomatizations of Finite Syntactic Epistemic States

Thomas Ågotnes and Michal Walicki

Department of Informatics, University of Bergen
PB. 7800, N-5020 Bergen, Norway
{agotnes,walicki}@ii.uib.no

Abstract. An agent who bases his actions upon explicit logical formulae has at any given point time a finite set of formulae he has computed. Closure or consistency conditions on this set cannot in general be assumed – reasoning takes time and real agents frequently have contradictory beliefs. This paper discusses a formal model of knowledge as explicitly computed sets of formulae. It is assumed that agents represent their knowledge syntactically, and that they can only know finitely many formulae at a given time. Existing syntactic characterizations of knowledge seem to be too general to have any interesting properties, but we extend the meta language to include an operator expressing that an agent knows *at most* a particular finite set of formulae. The specific problem we consider is the axiomatization of this logic. A sound system is presented. Strong completeness is impossible, so instead we characterize the theories for which we can get completeness. Proving that a theory actually fits this characterization, including proving weak completeness of the system, turns out to be non-trivial. One of the main results is a collection of algebraic conditions on sets of epistemic states described by a theory, which are sufficient for completeness. The paper is a contribution to a general abstract theory of resource bounded agents. Interesting results, e.g. complex algebraic conditions for completeness, are obtained from very simple assumptions, i.e. epistemic states as arbitrary finite sets and operators for knowing at least and at most.

1 Introduction

Traditional epistemic logics [7, 11], based on modal logic, are logics about knowledge closed under logical consequence – they describe agents who know all the infinitely many consequences of their knowledge. Such logics are very useful for many purposes, including modelling the information implicitly held by the agents or modelling the special case of extremely powerful reasoners. These logics fail, however, to model the explicit knowledge of real reasoners. Models of explicit knowledge are needed e.g. if we want to model agents who base their actions upon their knowledge. An example is when an agent is required to answer questions about whether he knows a certain formula or not. The agent must then decide whether this exact formula is true from his perspective — when he, e.g., is asked whether he knows $q \wedge p$ and he has already computed that $p \wedge q$ is true but not (yet) that $q \wedge p$ is true, then he cannot answer positively before he has performed a (trivial) act of reasoning. Real agents do not have unrestricted memory

or unbounded time available for reasoning. In reality, an agent who bases his actions on explicit logical formulae has at any given time a finite set of formulae he has computed. In the general case, we cannot assume any closure conditions on this set — we cannot assume that the agent has had time to deduce something yet — nor consistency or other connections to reality — real agents often hold contradictory or other false beliefs. The topic of this paper is formal models of knowledge as explicitly computed sets of formulae.

We present an agent simply as a finite set of formulae, called a finite epistemic state. Modal epistemic logics can be seen as describing not only knowledge but also a very particular model *reasoning* which is not valid for resource bounded agents. With a syntactic approach, we can get a theory of knowledge without any unrealistic assumptions about the reasoning abilities of the agents. The logic we present here is a logic about knowledge in a system of resource bounded agents *at a point in time*. We are not concerned with *how* the agents obtain their knowledge, but in reasoning about their static states of knowledge. Properties of reasoning can be modelled in an abstract way by considering only the set of epistemic states which a reasoning mechanism could actually produce. For example, we could choose to consider only epistemic states which does not contain both a formula and its negation. The question is, of course, whether anything interesting can be said about static properties of such general states. That depends on the available language.

Syntactic characterizations of states of knowledge are of course nothing new [5, 12, 8, 7]. The general idea is that the truth value of a formula such as $K_i\phi$, representing the fact that agent i knows the formula ϕ , need not depend on the truth value of any other formula of the form $K_i\psi$. Of course, syntactic characterization is an extremely general approach which can be used for several different models of knowledge – also including closure under logical consequence. It is, however, with the classical epistemic meta language too general to have any interesting logical properties.

The formula $K_i\phi$ denotes that fact that i knows *at least* ϕ – he knows ϕ but he may know more. We can generalize this to finite sets X of formulae:

$$\Delta_i X \equiv \bigwedge \{K_i\phi : \phi \in X\}$$

representing the fact that i knows at least X . In this paper we also use a dual operator, introduced in [2], to denote the fact that i knows *at most* X : $\nabla_i X$. denotes the fact that every formula an agent knows is included in X , but he may not know all the formulae in X . We call the language the agents represent their knowledge in *the object language*. In the case that the object language OL is finite, the operator ∇_i can be defined in terms of K_i (like Δ_i):

$$\nabla_i X = \bigwedge \{\neg K_i\phi : \phi \in OL \setminus X\}$$

But in the general case when OL is infinite, e.g. if OL is closed under propositional connectives, ∇_i is not definable by K_i . We also use a third, derived, epistemic operator: $\diamond_i X \equiv \Delta_i X \wedge \nabla_i X$ meaning that the agent knows exactly X .

The second difference from the traditional syntactic treatments of knowledge, in addition to the new operator, is that we restrict the set of formulae an agent can know at a given time to be finite. The problem we consider in this paper is axiomatizing

the resulting logic. We present a sound axiomatization, and show that it is impossible to obtain strong completeness. The main results are proof-theoretical and semantical characterizations of the sets of premises for which the system *is* complete; these sets include the empty set so the system is weakly complete. Proving completeness turns out to be quite difficult, but this can be seen as a price paid for the treatment of the inherently difficult issue of finiteness.

In the next section, the language and semantics for the logic is presented. In Section 3 it is shown that strong completeness is impossible, and a sound axiomatization presented. The rest of the paper is concerned with finding the sets of premises for which the system *is* complete. Section 5 gives a proof-theoretic account of these premise sets, while a semantic one consisting of complex algebraic conditions on possible epistemic states is given in Section 6. These results build on previous results for a more general logic, presented in Section 4. In Section 7 some actual completeness results, including weak completeness, are shown, and Section 8 concludes.

2 Language and Semantics

The logic is parameterized by an object language OL . The object language is the language in which the agents reason, e.g. propositional logic or first order logic. No assumptions about the structure of OL is made, and the results in this paper are valid for arbitrary object languages, but the interesting case is the usual one where OL is infinite. An example, which is used in this paper, is when OL is closed under the usual propositional connectives. Another possible property of an object language is that it is a subset of the meta language, allowing e.g. the expression of the *knowledge axiom* in the meta language: $\Delta_i\{\alpha\} \rightarrow \alpha$.

$\wp^{fin}(OL)$ is the set of all finite epistemic states, and a state $T \in \wp^{fin}(OL)$ is used as a term in an expression such as $\Delta_i T$. In addition, we allow set-building operators \sqcup, \sqcap on terms in order to be able to express things like $(\Delta_i T \wedge \Delta_i U) \rightarrow \Delta_i T \sqcup U$ in the meta language. TL is the language of all terms:

Definition 1 ($TL(OL)$) $TL(OL)$, or just TL , is the least set such that

- $\wp^{fin}(OL) \subseteq TL$
- If $T, U \in TL(OL)$ then $(T \sqcup U), (T \sqcap U) \in TL(OL)$

The *interpretation* $[T] \in \wp^{fin}(OL)$ of a term $T \in TL$ is defined as expected: $[X] = X$ when $X \in \wp^{fin}(OL)$, $[T \sqcup U] = [T] \cup [U]$, $[T \sqcap U] = [T] \cap [U]$. \square

An expression like $\Delta_i T$ relates the current epistemic state of an agent to the state described by the term T . In addition, we allow reasoning about the relationship between the two states denoted by terms T and U in the meta language by introducing formulae of the form $T \doteq U$, meaning that $[T] = [U]$.

The meta language EL , and the semantical structures, are parameterized by the number of agents n and a set of primitive propositions Θ , in addition to the object language. The primitive propositions Θ play a very minor role in the rest of this paper; they are only used to model an arbitrary propositional language which is then extended with epistemic (and term) formulae. Particularly, no relation between OL and Θ is assumed.

Definition 2 ($EL(n, \Theta, OL)$) Given a number of agents n , a set of primitive formulae Θ , and an object language OL , the epistemic language $EL(n, \Theta, OL)$, or just EL , is the least set such that:

- $\Theta \subseteq EL$
- If $T \in TL(OL)$ and $i \in [1, n]$ then $\Delta_i T, \nabla_i T \in EL$
- If $T, U \in TL(OL)$ then $(T \doteq U) \in EL$
- If $\phi, \psi \in EL$ then $\neg\phi, (\phi \wedge \psi) \in EL$ □

The usual derived propositional connectives are used, in addition to $T \preceq U$ for $T \sqcup U \doteq U$ and $\diamond_i \phi$ for $(\Delta_i \phi \wedge \nabla_i \phi)$. The operators Δ_i, ∇_i and \diamond_i are called epistemic operators. A boolean combination of formulae of the form $T \doteq U$ is called a *term formula*. Members of OL will be denoted α, β, \dots , of EL ϕ, ψ, \dots , and of TL T, U, \dots

The semantics of EL is defined as follows. Again, Θ and its interpretation does not play an important role in this paper.

Definition 3 (Knowledge Set Structure) A Knowledge Set Structure (KSS) for n agents, primitive propositions Θ and object language OL is an $n + 1$ -tuple

$$M = (s_1, \dots, s_n, \pi) \text{ where } s_i \in \wp^{fin}(OL)$$

and $\pi : \Theta \rightarrow \{\mathbf{true}, \mathbf{false}\}$ is a truth assignment. s_i is the epistemic state of agent i , and the set of all epistemic states is $\mathcal{S}^f = \wp^{fin}(OL)$. The set of all KSSs is denoted \mathcal{M}_{fin} . The set of all truth assignments is denoted Π . □

Truth of an EL formula ϕ in a KSS M , written $M \models_f \phi$, is defined as follows (the subscript f means “finite” and the reason for it will become clear later).

Definition 4 (Satisfaction) Satisfaction of a EL -formula ϕ in a KSS $M = (s_1, \dots, s_n, \pi) \in \mathcal{M}_{fin}$, written $M \models_f \phi$ (M is a model of ϕ), is defined as follows:

$$\begin{array}{lll} M \models_f p & \Leftrightarrow & \pi(p) = \mathbf{true} \\ M \models_f \neg\phi & \Leftrightarrow & M \not\models_f \phi \\ M \models_f (\phi \wedge \psi) & \Leftrightarrow & M \models_f \phi \text{ and } M \models_f \psi \\ M \models_f \Delta_i T & \Leftrightarrow & [T] \subseteq s_i \\ M \models_f \nabla_i T & \Leftrightarrow & s_i \subseteq [T] \\ M \models_f T \doteq U & \Leftrightarrow & [T] = [U] \end{array} \quad \square$$

As usual, if Γ is a set of formulae then we write $M \models_f \Gamma$ iff M is a model of all formulae in Γ and $\Gamma \models_f \phi$ (ϕ is a logical consequence of Γ) iff every model of Γ is also a model of ϕ . If $\emptyset \models_f \phi$, written $\models_f \phi$, then ϕ is valid. The set of all models of Γ is denoted $mod^f(\Gamma)$.

The logic consisting of the language EL , the set of structures \mathcal{M}_{fin} and the relation \models_f describes the current epistemic states of agents and how epistemic states are related to each other — without any restrictions on the possible epistemic states. For example, the epistemic states are neither required to be consistent – an agent can know both a formula and its negation – nor closed under any form of logical consequence – an agent

can know $\alpha \wedge \beta$ without knowing $\beta \wedge \alpha$. Both consequence conditions and closure conditions can be modelled by a set of structures $\mathcal{M}' \subseteq \mathcal{M}_{fin}$ where only epistemic states not violating the conditions are allowed. For example, we can construct a set of structures allowing only epistemic states not including both a formula α and $\neg\alpha$ at the same time, or including $\beta \wedge \alpha$ whenever $\alpha \wedge \beta$ is included. If we restrict the class of models considered under logical consequence to \mathcal{M}' , we get a new variant of the logic. We say that “ $\Gamma \models_f \phi$ with respect to \mathcal{M}' ” if every model of Γ in \mathcal{M}' is a model of ϕ .

The question of how to completely axiomatize these logics, the general logic described by \mathcal{M}_{fin} and the more special logics described by removing “illegal” epistemic states, is the main problem considered in this paper and is introduced in the next section.

3 Axiomatizations

The usual terminology and notation for Hilbert-style proof systems are used. A proof system is *sound* with respect to $\mathcal{M}' \subseteq \mathcal{M}_{fin}$ iff $\Gamma \vdash \phi$ implies that $\Gamma \models_f \phi$ wrt. \mathcal{M}' , *weakly complete* wrt. \mathcal{M}' iff $\models_f \phi$ wrt. \mathcal{M}' implies that $\vdash \phi$, and *strongly complete* wrt. \mathcal{M}' iff $\Gamma \models_f \phi$ wrt. \mathcal{M}' implies that $\Gamma \vdash \phi$.

When it comes to completeness, it is easy to see that it is impossible to achieve *full* completeness with respect to \mathcal{M}_{fin} with an axiomatization without rules with infinitely many antecedents. Let Γ_1 be the following theory:

$$\Gamma_1 = \{\Delta_1\{p\}, \Delta_1\{\Delta_1\{p\}\}, \Delta_1\{\Delta_1\{\Delta_1\{p\}\}\}, \dots\}$$

Clearly, this theory is not satisfiable, intuitively since it describes an agent with an infinite epistemic state. However, a proof of its inconsistency would necessarily include every single formula in the theory and be of infinite length. Another illustrating example is the following theory:

$$\Gamma_2 = \{\Delta_1\{\alpha, \beta\} \rightarrow \Delta_1\{\alpha \wedge \beta\} : \alpha, \beta \in OL\}$$

Unlike Γ_1 , Γ_2 is satisfiable, but only in a structure in which agent 1’s epistemic state is the empty set. Thus, $\Gamma_2 \models_f \nabla_1\emptyset$. But again, a proof of $\nabla_1\emptyset$ from Γ_2 would be infinitely long (because it would necessarily use infinitely many instances of the schema Γ_2), and an axiomatization without an infinite deduction rule would thus be (strongly) incomplete since then $\Gamma_2 \not\vdash \nabla_1\emptyset$.

3.1 The Basic System

Since we cannot get strong completeness, the natural question is whether we can construct a weakly complete system for the logic described by \mathcal{M}_{fin} . The answer is positive. The following system *EC* is sound and weakly complete with respect to \mathcal{M}_{fin} . The completeness proof is, however, not trivial, and is given in the rest of this paper.

Definition 5 (EC) The epistemic calculus EC is the logical system for the epistemic language EL consisting of the following axiom schemata:

All substitution instances of tautologies of propositional calculus	Prop
A sound and complete axiomatization of term formulae	TC
$\Delta_i \emptyset$	E1
$(\Delta_i T \wedge \Delta_i U) \rightarrow \Delta_i(T \sqcup U)$	E2
$(\nabla_i T \wedge \nabla_i U) \rightarrow \nabla_i(T \sqcap U)$	E3
$(\Delta_i T \wedge \nabla_i U) \rightarrow T \preceq U$	E4
$(\nabla_i(U \sqcup \{\alpha\}) \wedge \neg \Delta_i \{\alpha\}) \rightarrow \nabla_i U$	E5
$\Delta_i T \wedge U \preceq T \rightarrow \Delta_i U$	KS
$\nabla_i T \wedge T \preceq U \rightarrow \nabla_i U$	KG

and the following transformation rule

$$\frac{\Gamma \vdash \phi, \Gamma \vdash \phi \rightarrow \psi}{\Gamma \vdash \psi} \quad \text{MP} \quad \square$$

A sound and complete term calculus can be found in [2]. The main axioms of EC are self-explaining. **KS** and **KG** stand for “knowledge specialization” and “knowledge generalization”. It is easy to see that the deduction theorem (DT) holds for EC .

Theorem 1 (Soundness) If $\Gamma \vdash \phi$ then $\Gamma \models_f \phi$ □

3.2 Extensions

In Section 2 we mentioned that a logic with closure conditions or consistency conditions on the epistemic states can be modelled by a class $\mathcal{M}' \subseteq \mathcal{M}_{fin}$ by restricting the set of possible epistemic states. Such subclasses can often be described by axioms. For example, the soon following axiom **D** describes agents who never will believe both a formula and its negation.

The next question is whether if we add an axiom to EC the resulting system will be complete with respect to the class of models of the axiom; e.g. if EC extended with **D** will be complete with respect to the class of all models with epistemic states without both a formula and its negation.

Weak completeness of EC does, of course, entail (weak) completeness of EC extended with a *finite* set of axioms (DT). An axiom schema such as **D**, however, represents an *infinite* set of axioms, so completeness EC extended with such an axiom schema (with respect to the models of the schema) does not necessarily follow. The completeness proof to be presented is actually more than a proof of weak completeness of EC : it is a characterization of those sets of premises for which EC is complete, called *finitary theories*, and a method for deciding whether a given theory is finitary. Thus, if we extend EC with a finitary theory, the resulting logic is weakly complete with respect to the corresponding models.

Examples If we assume that OL is closed under the usual propositional connectives, some common axioms can be written in EL as follows:

$\Delta_i \{(\alpha \rightarrow \beta)\} \rightarrow (\Delta_i \{\alpha\} \rightarrow \Delta_i \{\beta\})$	Distribution	K
$\Delta_i \{\alpha\} \rightarrow \neg \Delta_i \{\neg \alpha\}$	Consistency	D
$\Delta_i \{\alpha\} \rightarrow \Delta_i \{\Delta_i \{\alpha\}\}$	Positive Introspection	4
$\neg \Delta_i \{\alpha\} \rightarrow \Delta_i \{\neg \Delta_i \{\alpha\}\}$	Negative Introspection	5

The system EC extended with axiom Φ will be denoted $EC\Phi$; e.g. the axioms above give the systems ECK , ECD , $EC4$, $EC5$.

4 More General Epistemic States

The results in the two next sections, builds upon an existing completeness result for a related logic. In this section we briefly describe the logic and quote the result. Details can be found in [1]¹.

The logic is actually a generalization of the logic in this paper, in which more epistemic states, henceforth called *general epistemic states*, are allowed:

$$\mathcal{S} = \wp(OL) \cup \wp^{fin}(OL \cup \{*\})$$

$*$ is a fixed formula, which is not a member of OL . In addition to the finite epistemic states \mathcal{S}^f , general epistemic states include states s where:

1. s is an infinite subset of OL : the agent knows infinitely many formulae
2. $s = s' \cup \{*\}$, where $s' \in \wp^{fin}(OL)$: the agent knows finitely many formulae but one of them is the special formula $*$

There is not space here to discuss what a state containing $*$ really represents, and it is not necessary since we will only use general epistemic states for technical intermediate results in this paper. Meta language and satisfiability is as in Section 2. It is assumed that an epistemic state can contain a the special formula $*$ $\notin OL$, and since EL is defined over OL , e.g. $\Delta_i \{*\}$ is not a well formed formula. It turns out that this crucial point makes our logical system EC (Def. 5) strongly complete with respect to this semantics. *General Knowledge Set Structures (GKSSs)* are like KSSs, but with general epistemic states instead of just finite epistemic states. \mathcal{M} is set of all GKSSs. To discern between the two logics we use the symbol \models for GKSSs and \models_f for KSSs. The set of all GKSS models of Γ is denoted $mod(\Gamma)$.

Theorem 2 (Completeness) For every $\Gamma \subseteq EL, \phi \in EL: \Gamma \models \phi \Leftrightarrow \Gamma \vdash \phi$ □

5 Finitary Theories and Completeness

Since EC is not strongly complete with respect to \mathcal{M}_{fin} , it is of interest to characterize exactly the theories for which EC is complete, i.e. for which Γ 's $\Gamma \models_f \phi \Rightarrow \Gamma \vdash \phi$ for

¹ And in a forthcoming paper (manuscript under preparation).

every $\phi \in EL$. In this section we provide a characterization of such theories. We define the concept of a *finitary theory*, and show that the set of finitary theories is exactly the set of theories for which *EC* is complete. The proof builds upon the completeness result for the more general logic described in the previous section.

Definition 6 (Finitary Theory) A theory Γ is *finitary* iff it is consistent and for all ϕ ,

$$\begin{aligned} \Gamma \vdash (\nabla_1 T_1 \wedge \cdots \wedge \nabla_n T_n) \rightarrow \phi \text{ for all terms } T_1, \dots, T_n \in TL \\ \Downarrow \\ \Gamma \vdash \phi \quad \square \end{aligned}$$

Informally speaking, a theory is finitary if provability of a formula under arbitrary upper bounds on epistemic states implies provability of the formula itself.

We use the intermediate definition of a finitarily open theory, and its relation to that of a finitary theory, in order to prove completeness.

Definition 7 (Finitarily Open Theory) A theory Γ is *finitarily open* iff there exist terms T_1, \dots, T_n such that

$$\Gamma \not\vdash \neg(\nabla_1 T_1 \wedge \cdots \wedge \nabla_n T_n) \quad \square$$

Informally speaking, a theory is finitarily open if it can be consistently extended with some upper bound on the epistemic state of each agent.

Lemma 1

1. A finitary theory is finitarily open.
2. If Γ is a finitary theory and $\Gamma \not\vdash \phi$, then $\Gamma \cup \{\neg\phi\}$ is finitarily open. \square

PROOF

1. Let Γ be a finitary theory. If Γ is not finitarily open, $\Gamma \vdash \neg(\nabla_1 T_1 \wedge \cdots \wedge \nabla_n T_n)$ for all terms T_1, \dots, T_n . Then, for an arbitrary ϕ , $\Gamma \vdash (\nabla_1 T_1 \wedge \cdots \wedge \nabla_i T_n) \rightarrow \phi$ for all T_1, \dots, T_n and thus $\Gamma \vdash \phi$ since Γ is finitary. By the same argument $\Gamma \vdash \neg\phi$, contradicting the fact that Γ is consistent.
2. Let Γ be a finitary theory, and let $\Gamma \not\vdash \phi$. Then there must exist terms $T_1^\phi, \dots, T_n^\phi$ such that $\Gamma \not\vdash (\nabla_1 T_1^\phi \wedge \cdots \wedge \nabla_n T_n^\phi) \rightarrow \phi$. By **Prop** we must have that $\Gamma \not\vdash \neg\phi \rightarrow \neg(\nabla_1 T_1^\phi \wedge \cdots \wedge \nabla_n T_n^\phi)$ and thus that $\Gamma \cup \{\neg\phi\} \not\vdash \neg(\nabla_1 T_1^\phi \wedge \cdots \wedge \nabla_n T_n^\phi)$, which shows that $\Gamma \cup \{\neg\phi\}$ is finitarily open. \blacksquare

It is difficult in practice to show whether a given theory satisfies a proof theoretic condition such as those for finitary or finitarily open theories, but we have a tool to convert the problem to a semantic one: the completeness result for GKSSs in the previous section (Theorem 2). For example, to show that $\Gamma \vdash \phi$, it suffices to show that $\Gamma \models \phi$ (with respect to GKSSs). This result can be used to see that the claims of *non-finitaryness* in the following example hold.

Example 1 The following are examples of non-finitary theories (let $n = 2$ and $p \in \Theta$):

1. $\Gamma_1 = \{\Delta_1\{p\}, \Delta_1\{\Delta_1\{p\}\}, \Delta_1\{\Delta_1\{\Delta_1\{p\}\}\}, \dots\}$. Γ_1 is not finitarily open, and describes an agent with an infinite epistemic state.
2. $\Gamma_2 = \{\neg \nabla_1 T : T \in TL\}$. Γ_2 is not finitarily open, and describes an agent which cannot be at any finite point.
3. $\Gamma_3 = \{\nabla_1 T \rightarrow \neg \nabla_2 T' : T, T' \in TL\}$. Γ_3 is not finitarily open, and describes a situation where agents 1 and 2 cannot *simultaneously* be at finite points.
4. $\Gamma_4 = \{\nabla_1 T \rightarrow p : T \in TL\}$. Γ_4 is finitarily open, but not finitary. To see the former, observe that there is no T such that $\Gamma_4 \vdash \neg \nabla_1 T$ or $\Gamma_4 \vdash \neg \nabla_2 T$. To see the latter, observe that $\Gamma_4 \not\vdash p$ but $\Gamma_4 \vdash (\nabla_1 T_1 \wedge \nabla_2 T_2) \rightarrow p$ for all T_1, T_2 . \square

Theorem 3 A theory Γ is finitarily open if and only if it is satisfiable in \mathcal{M}_{fin} . \square

PROOF Γ is finitarily open iff there exist T_i ($1 \leq i \leq n$) such that $\Gamma \not\vdash \neg(\nabla_i T_1 \wedge \dots \wedge \nabla_n T_n)$; iff, by Theorem 2, there exist T_i such that $\Gamma \not\vdash \neg(\nabla_i T_1 \wedge \dots \wedge \nabla_n T_n)$; iff there exist T_i and a GKSS $M \in \mathcal{M}$ such that $M \models \Gamma$ and $M \models \nabla_1 T_1 \wedge \dots \wedge \nabla_n T_n$; iff there exist T_i and $M = (s_1, \dots, s_n, \pi) \in \mathcal{M}$ such that $s_i \subseteq [T_i]$ ($1 \leq i \leq n$) and $M \models \Gamma$; iff there exist $s_i \in \wp^{fin}(OL)$ ($1 \leq i \leq n$) such that $(s_1, \dots, s_n, \pi) \models \Gamma$; iff Γ is satisfiable in \mathcal{M}_{fin} . \blacksquare

Theorem 4 Let $\Gamma \subseteq EL$. $\Gamma \models_f \phi \Rightarrow \Gamma \vdash \phi$ for all ϕ iff Γ is finitary. \square

PROOF Let Γ be a finitary theory and let $\Gamma \models_f \phi$. By Lemma 1.1 Γ is finitarily open and thus satisfiable by Theorem 3. $\Gamma \cup \{\neg\phi\}$ is unsatisfiable in \mathcal{M}_{fin} , and thus not finitarily open, and it follows from Lemma 1.2 that $\Gamma \vdash \phi$.

For the other direction, let $\Gamma \models_f \phi \Rightarrow \Gamma \vdash \phi$ for all ϕ , and assume that $\Gamma \not\vdash \phi$. Then, $\Gamma \not\models_f \phi$, that is, there is a $M = (s_1, \dots, s_n, \pi) \in \text{mod}^f(\Gamma)$ such that $M \not\models_f \phi$. Let T_i ($1 \leq i \leq n$) be terms such that $[T_i] = s_i$. $M \models_f \nabla_i T_1 \wedge \dots \wedge \nabla_n T_n$, and thus $M \not\models_f (\nabla_1 T_1 \wedge \dots \wedge \nabla_n T_n) \rightarrow \phi$. By soundness (Theorem 1) $\Gamma \not\vdash (\nabla_1 T_1 \wedge \dots \wedge \nabla_n T_n) \rightarrow \phi$, showing that Γ is finitary. \blacksquare

Lemma 2 Let $\Gamma \subseteq EL$. The following statements are equivalent:

1. Γ is finitary.
2. $\Gamma \models_f \phi \Rightarrow \Gamma \vdash \phi$, for any ϕ
3. $\Gamma \models_f \phi \Rightarrow \Gamma \models \phi$, for any ϕ
4. $(\exists M \in \text{mod}(\Gamma) M \models \phi) \Rightarrow (\exists M \in \text{mod}^f(\Gamma) M \models_f \phi)$, for any ϕ
5. $\Gamma \not\vdash \phi \Rightarrow \Gamma \cup \{\neg\phi\}$ is finitarily open, for any ϕ . \square

Lemma 2.4 is a finite model property, with respect to the models of Γ .

We have now given a proof-theoretic definition of all theories for which EC is complete: the finitary theories. We have also shown some examples of non-finitary theories. We have not, however, given any examples of *finitary* theories. Although the problem of proving that EC is complete for a theory Γ have been reduced to proving that the theory is finitary according to Definition 6, the next problem is how to show that a given theory in fact is finitary. For example, is the empty theory finitary? If it is, then EC is weakly complete. We have not been able to find a trivial or easy way to prove finitariness. In the next section, we present results which can be used to prove finitariness. The results are semantic conditions for finitariness, but can only be used for theories of a certain class and we are only able to show that they are *sufficient*.

6 Semantic Finitaryness Conditions

Epistemic axioms are axioms which describe legal epistemic states, like “an agent cannot know both a formula and its negation”. In Section 4 we presented the notion of a *general* epistemic state, and epistemic axioms can be seen as describing sets of legal general epistemic states as well as sets of legal finite epistemic states. Although we are ultimately interested in the latter, in this section we will be mainly interested in the former – we will present conditions on the algebraic structure of sets of general epistemic states in $mod(\Phi)$ which are sufficient for the axioms Φ to be finitary.

First, epistemic axioms and their correspondence with sets of legal general epistemic states are defined. Then, conditions on these sets are defined, and it is shown that the GKSSs of a given set of epistemic axioms – being (essentially) the Cartesian product of the corresponding sets of legal general states – exhibit the finite model property if the sets of legal general states fulfil the conditions. The set of axioms is then finitary by Lemma 2.4.

6.1 Epistemic Axioms

Not all formulae in EL should be considered as candidates for describing epistemic properties. One example is $p \rightarrow \Delta_i\{p\}$. This formula does not solely describe the *agent* – it describes a relationship between the agent and the world. Another example is $\Diamond_i\{p\} \rightarrow \Diamond_j\{q\}$, which describes a constraint on one agent’s belief set contingent on another agent’s belief set. Neither of these two formulae describe purely *epistemic* properties of an agent. In the following definition, EF is the set of epistemic formulae and Ax is the set of candidate epistemic axioms.

Definition 8 (EF, EF^i, Ax)

- $EF \subseteq EL$ is the least set such that

$$\text{If } T \in TL \text{ then } \Delta_i T, \nabla_i T \in EF \quad (1 \leq i \leq n)$$

$$\text{If } \phi, \psi \in EF \text{ then } \neg\phi, (\phi \wedge \psi) \in EF$$

- $EF^i = \{\phi \in EF : \text{Every epistemic operator in } \phi \text{ is an } i\text{-op.}\} \quad (1 \leq i \leq n)$
- $Ax = \bigcup_{1 \leq i \leq n} EF^i$ □

An example of an epistemic axiom schema is, if we assume that OL is closed under conjunction,

$$\Delta_i\{\alpha \wedge \beta\} \rightarrow \Delta_i\{\alpha\} \wedge \Delta_i\{\beta\} \quad (1)$$

Definition 9 ($\mathcal{M}^\phi, \mathcal{S}_i^\phi, \mathcal{M}^\Phi, \mathcal{S}_i^\Phi$) For each epistemic formula $\phi \in EF^i$,

$$\mathcal{M}^\phi = \mathcal{S}_1^\phi \times \dots \times \mathcal{S}_n^\phi \times \Pi$$

where $\mathcal{S}_j^\phi = \mathcal{S}$ for $j \neq i$ and \mathcal{S}_i^ϕ is constructed by structural induction over ϕ as follows:

$$\begin{aligned} \mathcal{S}_i^{\Delta_i T} &= \{X \in \mathcal{S} : [T] \subseteq X\} & \mathcal{S}_i^{\nabla_i T} &= \{X \in \mathcal{S} : X \subseteq [T]\} \\ \mathcal{S}_i^{\neg\psi} &= \mathcal{S} \setminus \mathcal{S}_i^\psi & \mathcal{S}_i^{\psi_1 \wedge \psi_2} &= \mathcal{S}_i^{\psi_1} \cap \mathcal{S}_i^{\psi_2} \end{aligned}$$

If $\Phi \subseteq Ax$ then

$$S_i^\Phi = \left(\bigcap_{\phi \in \Phi \cap EF^i} S_i^\phi \right) \cap \mathcal{S} \quad \mathcal{M}^\Phi = S_1^\Phi \times \cdots \times S_n^\Phi \times \Pi \quad \square$$

In the construction of \mathcal{M}^Φ we remove the impossible (general) epistemic states by restricting the set of epistemic states to S_i^ϕ . The epistemic states which are not removed are the possible states — an agent can be placed in any of these states and will satisfy the epistemic axiom ϕ . That \mathcal{M}^Φ indeed is the class of models of Φ can easily be shown.

Lemma 3 If $\Phi \subseteq Ax$, $\mathcal{M}^\Phi = \text{mod}(\Phi)$ □

Note that \emptyset is trivially a set of epistemic axioms, and that $S_i^\emptyset = \mathcal{S}$ and $\mathcal{M}^\emptyset = \mathcal{M}$.

Thus, the model class for epistemic axioms is constructed by removing certain states from the set of legal epistemic states. For example, (1) corresponds to removing epistemic states where the agent knows a conjunction without knowing the conjuncts.

6.2 Finitaryness of Epistemic Axioms

Lemmas 2.1 and 2.4 say that Γ is finitary iff $\text{mod}(\Gamma)$ has the finite model property. We make the following intermediate definition, and the following Lemma is an immediate consequence.

Definition 10 (Finitary set of GKSSs) A class of GKSSs $\mathcal{M}' \subseteq \mathcal{M}$ is finitary iff, for all ϕ :

$$\exists M \in \mathcal{M}' M \models \phi \Rightarrow \exists M^f \in \mathcal{M}'^f M^f \models \phi$$

where $\mathcal{M}'^f = \mathcal{M}' \cap \mathcal{M}_{fin}$. □

Lemma 4 Let $\Gamma \subseteq EL$. Γ is finitary iff $\text{mod}(\Gamma)$ is finitary. □

In the definition of the conditions on sets of general epistemic states, the following two general algebraic conditions are used.

Directed Set A set A with a reflexive and transitive relation \leq is *directed* iff for every finite subset B of A , there is an element $a \in A$ such that $b \leq a$ for every $b \in B$. In the following directedness of a set of sets is implicitly taken to be with respect to subset inclusion.

Cover A family of subsets of a set A whose union includes A is a *cover* of A .

The main result is that the following conditions on sets of general epistemic states are sufficient for the corresponding GKSSs to be finitary (Def. 10), and furthermore, if the sets are induced by epistemic axioms, that the axioms are finitary. The conditions are quite complicated, but simpler ones are given below.

Definition 11 (Finitary Set of Epistemic States) If $S \subseteq \mathcal{S}$ is a set of epistemic states and $s \in \wp(OL)$, then the set of finite subsets of s included in S is denoted

$$S|_s^f = S \cap \wp^{fin}(s)$$

S is *finitary* iff both:

1. For every infinite $s \in S$:
 - (a) $S|_s^f$ is directed
 - (b) $S|_s^f$ is a cover of s
2. $\forall_{s \cup \{*\} \in S} \forall_{s' \in \wp^{fin}(OL)} \exists_{\alpha \notin s'}$:
 - (a) $\exists_{s^f \in S \cap \wp(s \cup \{\alpha\})} s' \cap s \subseteq s^f$
 - (b) $\exists_{s^f \in S \cap \wp(s \cup \{\alpha\})} s^f \not\subseteq s'$
 - (c) $S \cap \wp(s \cup \{\alpha\})$ is directed \square

The following Lemma is the main technical result in this section. The proof is quite long and complicated, and must be left out due to space restrictions. It can be found in [1]².

Lemma 5 If S_1, \dots, S_n are finitary sets of epistemic states (Def. 11), then

$$S_1 \times \dots \times S_n \times II$$

is a finitary set of GKSSs (Def. 10). \square

Recall that a set Φ of epistemic axioms induces sets of legal epistemic states S_i^Φ (Def. 9).

Theorem 5 If Φ is a set of epistemic axioms such that $S_1^\Phi, \dots, S_n^\Phi$ are finitary sets of epistemic states, then Φ is finitary. \square

PROOF Since Φ are epistemic axioms, $\mathcal{M}^\Phi = S_1^\Phi \times \dots \times S_n^\Phi \times II$. Since all S_i^Φ are finitary, by Lemma 5 \mathcal{M}^Φ is a finitary set of GKSSs. Since $\mathcal{M}^\Phi = \text{mod}(\Phi)$ (Lemma 3), Φ is finitary by Lemma 4. \blacksquare

Theorem 5 shows that the conditions in Def. 11 on the set of legal epistemic states induced by epistemic axioms are sufficient to conclude that the axioms are finitary. In the following Corollary, we present several alternative sufficient conditions which are stronger. It can easily be shown that these conditions imply Def. 11.

Corollary 1 A set of epistemic states $S \subseteq \mathcal{S}$ is finitary if either one of the following three conditions hold:

1. For every $s \subseteq OL$:
 - (a) $S|_s^f$ is directed
 - (b) $S|_s^f$ is a cover of s
2. (a) $S|_s^f$ is directed for every $s \subseteq OL$
 (b) $\{\alpha\} \in S$ for every $\alpha \in OL$
3. (a) $S|_s^f$ is directed for every infinite $s \in S$
 (b) $\{\alpha\} \in S$ for every $\alpha \in OL$
 (c) $\forall_{s \cup \{*\} \in S} \forall_{s' \in \wp^{fin}(OL)} \exists_{\alpha \notin s'} s \cup \{\alpha\} \in S$ \square

7 Some Completeness Results

For a given axiom schema Φ , the results from Sections 4, 5 and 6 can be used to test whether the system $EC\Phi$ is weakly complete, henceforth in this section called only “complete”, with respect to $\text{mod}^f(\Phi) \subseteq \mathcal{M}_{fin}$. First, check that Φ is an epistemic axiom (Def. 8). Second, construct the GKSS (see Sec. 4) models of Φ , $\mathcal{M}^\Phi = S_1^\Phi \times \dots \times S_n^\Phi \times$

² Or by contacting the authors.

$\Pi = \text{mod}(\Phi)$ (Def. 9, Lemma 3). Third, check that each S_i^Φ is finitary (Def. 11) – it suffices that they each satisfy one of the simpler conditions in Corollary 1. If these tests are positive, $EC\Phi$ is complete with respect to $\mathcal{M}_{fin}^\Phi = \text{mod}^f(\Phi)$, the KSSs included in \mathcal{M}^Φ , by Theorems 5 and 3. The converse does not hold; \mathcal{M}_{fin}^Φ is not necessarily *incomplete* with respect to the corresponding models if the tests are negative. Many of the properties discussed in Section 5 can, however, be used to show incompleteness.

These techniques are used in Theorem 6 below to prove the assertion from Section 3 about weak completeness of EC , in addition to results about completeness of the systems ECK , ECD , $EC4$ and $EC5$ from Section 3.2. For the latter results it is assumed that OL is closed under the usual propositional connectives.

Theorem 6 (Completeness Results)

1. EC is sound and complete with respect to \mathcal{M}_{fin}
2. ECK is sound and complete with respect to $\mathcal{M}_{fin}^{\mathbf{K}}$
3. ECD is sound and complete with respect to $\mathcal{M}_{fin}^{\mathbf{D}}$
4. $EC4$ is not complete with respect to \mathcal{M}_{fin}^4
5. $EC5$ is not complete with respect to \mathcal{M}_{fin}^5 □

PROOF Soundness, in the first three parts of the theorem, follows immediately from Theorem 1 and the fact that \mathbf{K} and \mathbf{D} are valid in $\mathcal{M}_{fin}^{\mathbf{K}}$ and $\mathcal{M}_{fin}^{\mathbf{D}}$, respectively. The strategy for the completeness proofs, for the first three parts of the theorem, is as outlined above. (Weak) completeness of EC can be considered by “extending” EC by the empty set, and show that the empty set is a finitary theory. The empty set is trivially a set of epistemic axioms, and the axiom schemas \mathbf{K} and \mathbf{D} also both represent sets of epistemic axioms, with GKSS models constructed from the following sets of general epistemic states respectively:

$$\begin{aligned} S_i^\emptyset &= \mathcal{S} \\ S_i^{\mathbf{K}} &= \mathcal{S} \setminus \{X \in \mathcal{S} : \exists \alpha, \beta \in OL \alpha \rightarrow \beta, \alpha \in X; \beta \notin X\} \\ S_i^{\mathbf{D}} &= \mathcal{S} \setminus \{X \in \mathcal{S} : \exists \alpha \in OL \alpha, \neg \alpha \in X\} \end{aligned}$$

We show that these sets all are finitary sets of epistemic states by using Corollary 1. It follows by Theorem 5 that the theories \emptyset , \mathbf{K} and \mathbf{D} are finitary theories, and thus that EC , ECK and ECD are (weakly) complete by Theorem 4. For the two last parts of the theorem, we show that 4 and 5 are not finitary theories; it follows by Theorem 4 that $EC4$ and $EC5$ are incomplete.

1. Corollary 1.1 holds for $S_i^\emptyset = \mathcal{S}$: Let $s \subseteq OL$. $S|_s^f = \mathcal{S} \cap \wp^{fin}(s) = \wp^{fin}(s)$. $\wp^{fin}(s)$ is directed, because for every finite subset $B \subset \wp^{fin}(s)$, $\cup_{s' \in B} s' \in \wp^{fin}(s)$. $\wp^{fin}(s)$ is a cover of s , because $s \subseteq \cup \wp^{fin}(s)$.
2. Corollary 1.3 holds for $S_i^{\mathbf{K}}$:
Corollary 1.3.(a): Let $s', s'' \in S_i^{\mathbf{K}} \cap \wp^{fin}(s)$, and let:

$$\begin{aligned} s_0 &= s' \cup s'' \\ s_j &= s_{j-1} \cup \{\beta : \alpha \rightarrow \beta, \alpha \in s_{j-1}\} \quad 0 < j \\ s^f &= \bigcup_j s_j \end{aligned}$$

It is easy to show that $s^f \in S_i^{\mathbf{K}}$, each s_j is a finite subset of s , and s^f is finite.

Corollary 1.3.(b): Clearly, $\{\alpha\} \in S_i^K$ for every $\alpha \in OL$.

Corollary 1.3.(c): Let $s \cup \{*\} \in S_i^K$ and $s' \in \wp^{fin}(OL)$. Let $\alpha \in OL$ be s. t.:

- $\alpha \rightarrow \beta \notin s$ for any $\beta \in OL$
- $\alpha \notin s'$
- The main connective in α is not implication

It is easy to see that there exist infinitely many α satisfying these three conditions and it can easily be shown that $s \cup \{\alpha\} \in S_i^K$.

3. Corollary 1.3 holds for S_i^D :

Corollary 1.3.(a): Let $s', s'' \in S_i^D \cap \wp^{fin}(s)$, and let $s^f = s' \cup s''$. It can easily be shown that $s^f \in S_i^D$, and $s^f \in \wp^{fin}(s)$ trivially.

Corollary 1.3.(b): Clearly, $\{\alpha\} \in S_i^D$ for every $\alpha \in OL$.

Corollary 1.3.(c): Let $s \cup \{*\} \in S_i^D$ and $s' \in \wp^{fin}(OL)$. Let $\alpha \in OL$ be s. t.:

- $\neg\alpha \notin s$
- $\alpha \notin s'$
- α does not start with negation

It is easy to see that there exist infinitely many α satisfying these three conditions and it can easily be shown that $s \cup \{\alpha\} \in S_i^D$.

4. Let $1 \leq i \leq n$, and let $M = (s_1, \dots, s_n, \pi) \in \mathcal{M}_{fin}$ such that $M \models_f \mathbf{4}$. Clearly, s_i must be the empty set – otherwise it would not be finite. Thus, $\mathbf{4} \models_f \nabla_i \emptyset$. $\mathbf{4}$ does, however, have *infinite* models, so $\mathbf{4} \not\models \nabla_i \emptyset$. Lemma 2 gives that $\mathbf{4}$ is not finitary.

5. It is easy to see that $\mathbf{5}$ is not satisfiable in \mathcal{M}_{fin} (i.e. that a model for $\mathbf{5}$ must be infinite). By Theorem 3 and Lemma 1, $\mathbf{5}$ is not finitary. ■

Although the results in Theorem 6 are hardly surprising, they are surprisingly hard to prove.

8 Discussion and Conclusions

This paper presents a general and very abstract theory of resource bounded agents. We assumed that agents' epistemic states are arbitrary finite sets of formulae. The addition of the “knowing at most” operator ∇_i gives a more expressive language for a theory of knowledge without any unrealistic assumptions about the reasoning abilities of the agents. Properties of reasoning can be modelled in an abstract way by considering only the set of epistemic states which a reasoning mechanism could actually produce. If a more detailed model of reasoning is needed, the framework can be extended with a model describing transitions between finite epistemic states. This is exactly what is done in [3].

The main results in this paper are an axiomatization of the logic, and two characterizations of the theories for which the logic is complete. The first, the notion of finitary theories, is a proof-theoretic account of all such theories. The second, algebraic conditions on certain sets of epistemic states, is a semantic one, but is only a sufficient condition for finitariness. The latter was used to show finitariness of the empty theory and thus weak completeness of the system. It follows from these results that the logic *EC* is decidable. Interesting results have been obtained from very weak assumptions: finite memory and a “knowing at most” operator in the meta language give complex algebraic conditions for axiomatizability.

Related works include the classical syntactic treatment of knowledge as mentioned in Section 1. [7] presents this approach in the form of *standard syntactic assignments*. It is easy to see that KSSs are equivalent to standard syntactic assignments restricted to assigning finite knowledge to each agent. The ∇_i operator is new in the context of syntactic models. It is, however, similar to Levesque's *only knowing* operator \mathbf{O} [10]. $\mathbf{O}\alpha$ means that the agent does not know more than α , but knowledge in this context means knowledge closed under logical consequence and "only knowing α " is thus quite different from "knowing at most" a finite set of formulae syntactically. Models of reasoning as transition between syntactic states, as mentioned above, include *Konolige's deduction model* [9], *active logics* [6] and *timed reasoning logics (TRL)* [4].

Possibilities for future work include further development of the identification of finitary theories. For the case of epistemic axioms, the presented algebraic conditions are sufficient but not necessary and tighter conditions would be interesting. Deciding finitariness of general, not necessarily epistemic, axioms should also be investigated, and particularly interesting is the knowledge axiom (mentioned on p. 3).

References

1. Thomas Ågotnes. *A Logic of Finite Syntactic Epistemic States*. PhD thesis, Department of Informatics, University of Bergen, 2004.
2. Thomas Ågotnes and Michal Walicki. A logic for reasoning about agents with finite explicit knowledge. In Bjørnar Tessem, Pekka Ala-Siuru, Patrick Doherty, and Brian Mayoh, editors, *Proceedings of the 8th Scandinavian Conference on Artificial Intelligence (SCAI'03)*, Frontiers in Artificial Intelligence and Applications, pages 163–174, Bergen, Norway, Nov 2003. IOS Press.
3. Thomas Ågotnes and Michal Walicki. Syntactic knowledge: A logic of reasoning, communication and cooperation. In *Proceedings of the Second European Workshop on Multi-Agent Systems (EUMAS)*, Barcelona, Spain, December 2004.
4. Natasha Alechina, Brian Logan, and Mark Whitsey. A complete and decidable logic for resource-bounded agents. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2004)*, pages 606–613. ACM Press, Jul 2004.
5. R. A. Eberle. A logic of believing, knowing and inferring. *Synthese*, 26:356–382, 1974.
6. Jennifer Elgot-Drapkin, Sarit Kraus, Michael Miller, Madhura Nirkhe, and Donald Perlis. Active logics: A unified formal approach to episodic reasoning. Technical Report CS-TR-4072, 1999.
7. Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge*. The MIT Press, Cambridge, Massachusetts, 1995.
8. Joseph Y. Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, July 1990.
9. Kurt Konolige. *A Deduction Model of Belief and its Logics*. PhD thesis, Stanford University, 1984.
10. H. J. Levesque. All I know: a study in autoepistemic logic. *Artificial Intelligence*, 42:263–309, 1990.
11. J.-J. Ch. Meyer and W. van der Hoek. *Epistemic Logic for AI and Computer Science*. Cambridge University Press, Cambridge, England, 1995.
12. R. C. Moore and G. Hendrix. Computational models of beliefs and the semantics of belief sentences. Technical Note 187, SRI International, Menlo Park, CA, 1979.