Accumulative Knowledge Under Bounded Resources

Wojciech Jamroga[1,3] and Masoud Tabatabaei[2]

[1] Computer Science and Communication, University of Luxembourg
[2] Interdisc. Centre for Security, Reliability and Trust, University of Luxembourg
[3] Institute of Computer Science, Polish Academy of Sciences
wjamroga@in.tu-clausthal.de masoud.tabatabaei@uni.lu

Abstract. A possible purpose of performing an action is to gather information. Such information-collecting actions are usually resource-consuming. The resources needed for performing them can be for example time or memory, but also money, specialized equipment etc. In this work, we propose a formal framework to study how the ability of an agent to improve its knowledge changes as a result of changing the available resources. We introduce a model for resource-consuming information-collecting actions, and show how the process of accumulating knowledge can be modelled. Based on this model, we propose a modal logic for reasoning about the epistemic abilities of agents. We present some validities of the logic, and show that the model checking problem sits in the first level of polynomial hierarchy. We also discuss the connection between our framework and classic information theory. More specifically, we show that the notion of uncertainty given by Hartley measure can be seen as a special case of an agent's ability to improve its knowledge using information-collecting actions.

Keywords: Reasoning about knowledge, multi-agent systems, belief update, resource-bounded agents, epistemic logic, temporal logic, strategic ability.

1 Introduction

Performing actions is an intrinsic feature of agents. In the real world, execution of an action requires resources. The resources may be time, money, memory, space, etc. Therefore, the abilities ascribed to an agent depend on the amount of available resources. Reasoning about realistic agents should take into account the limitations imposed by resource bounds.

In this work, we are mostly interested in reasoning about the abilities of agents to change their view of the situation. More specifically, we want to capture the way agents with bounded resources modify their knowledge about the environment by performing information-collecting actions, such as sensing and observing. Acquiring knowledge is in many cases essential for an intelligent agent. One example of an agent that performs (resource-consuming) observations in order to refine its knowledge is a robot in a rescue mission that tries to obtain knowledge about the type of danger and the location of people in the danger zone. Another example is a real-time classifier with the task of classifying a given picture within a short time, and with several classification algorithms at hand. We believe that a logic to reason about accumulating knowledge by use of resource-consuming information-collecting actions can help in modelling and analysing the behaviour of agents in many similar scenarios.

We begin by introducing our model of resource-consuming epistemic updates in Section 2. In Section 3, we propose the Logic of Accumulative Knowledge (LAcK). LAcK is a variant of temporal-epistemic logic that can be used to reason about sequences of such updates, and to describe what agents can/must learn if they engage a given amount of resources into their explorative activity. Some validities of LAcK are presented in Section 4; we also show how the well-known Hartley measure of uncertainty can be characterized in our logic. Section 5 investigates the complexity of the verification problem, both in the general case and under some additional assumptions about the model and/or the formula. In Section 6, we propose an extension of LAcK with modalities for strategic exploration. We study the related expressivity of strategic and temporal-epistemic operators for reasoning about accumulative knowledge. We also show that the model checking problem does not become more complex in the enhanced logic, and that Hartley measure gets an even more intuitive characterization. Finally, we discuss related work in Section 7, and present some conclusions in Section 8.

This article is an extended version of the workshop paper [22]. In addition to correcting some minor mistakes and extending the discussion of related work, we have added the whole Section 6 on strategic epistemic exploration.

2 Resource Bounded Model for Accumulative Knowledge

In this section we develop a model that formalizes scenarios in which agents build their knowledge by using resource-consuming actions. We explain the ideas behind our approach with the following motivating example.

Example 1 (Medical agent). Consider a medical assistant agent. The agent is to help diagnosing patients in areas where there are not enough general practitioners. The process of helping a patient starts when the patient informs the agent about his symptoms. The agent then generates a list of all possible diseases consistent with the symptoms. Among the diseases, some are considered as being serious. The agent's duty is finding out whether the patient's disease is serious or not. If it is found out that the disease is not serious, the agent prescribes appropriate medications. Otherwise the agent sends the patient to a medical centre. A set of medical tests is available to determine the seriousness of the disease. Each medical test takes some specific time. Depending on the result of the test, the agent can rule out some of the diseases, and so on.

In principle, the process should continue until the agent finds out if the disease is serious. However, there are some important questions that an intelligent agent might consider before even starting. What are the relevant medical tests for a patient with the given symptoms? If the supply of test kits is limited, is the agent able to find out the seriousness of the disease with the available kits? If, among the possible diseases, there is one that should be diagnosed quickly, is there a sequence of tests that will make the agent certain about this very disease before the condition of the patient gets critical?

2.1 Observation-Based Certainty Model

We will use *possible worlds models* [26] to formalize this and similar scenarios. Each world corresponds to a possible state of affairs. If an agent cannot distinguish between two worlds, this is represented by the corresponding modal accessibility relation. For instance, for the medical agent, the set of possible worlds can consist of all possible diagnoses (i.e., diseases). The agent knows that a given property holds if and only if it holds in all the accessible worlds. For example, if all the possible diseases consistent with the symptoms are caused by an infection, then we say that the agent *knows* that the patient has an infection.

An agent may refine its knowledge by performing information-collecting actions. In this work, we refer to all information-collecting actions as *observations*. The medical agent can, e.g., check the temperature of the patient. Performing an observation may refine the knowledge of an agent by ruling out some of the possible worlds. For example, after learning that the patient does not have high temperature, the medical agent rules out all the diseases that include fever. The agent needs resources (time, memory, space, money, etc.) to perform observations. Thus, in order to analyse the agent's ability to gain the required knowledge, we need to take into account the cost of the observations and the available resources. We assume that agents can proceed with available observations pro-actively, whenever they wish to. In this sense, observations are like tests (or yes/no questions) performed on the environment of the agent.

We formalize the intuitions as follows, drawing inspiration from modal epistemic logic and dynamic epistemic logic.

Definition 1 (Observation-based certainty model). *Having a set of atomic propositions P and a set of agents A, an observation based certainty model is a tuple $M = \langle S, R, V, Obs, obs, cost, cover\rangle$ where:*

- *S is a set of states (possible worlds).*
- *$R \subseteq A \times S \times S$ is the accessibility relation which represents the worlds that are accessible for each agent. We will write $s_1 {\sim_a} s_2$ instead of $(a, s_1, s_2) \in R$. We assume that each binary relation $R(a, \cdot, \cdot)$ is an equivalence relation.*
- *$V : P \to 2^S$ is a valuation of propositions that shows which propositions are true in which worlds.*
- *Obs is a set of labels for binary observations.*
- *$obs : A \to 2^{Obs}$ defines availability of observations to agents.*
- *$cover : Obs \to 2^S$ is the coverage function. It specifies the set of worlds that correspond to the "positive" outcome of an observation. We call $cover(q)$ the* covering set *of the observation q.*

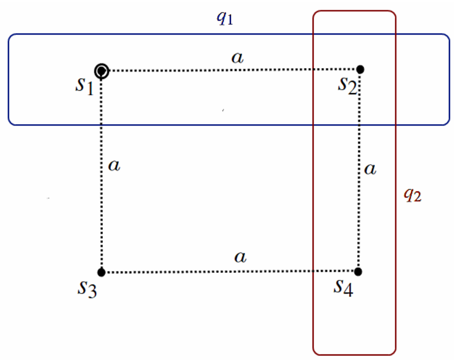

Fig. 1. Model M_1 of simple medical diagnosis. The epistemic accessibility relation for agent a is represented by the dotted lines (modulo transitivity). q_1, q_2 are the available observations; their covering sets are depicted by the rectangles. Moreover, we assume that $cost(q_1) = 1$ and $cost(q_2) = 2$.

– *cost* : *Obs* $\rightarrow \mathcal{C}$ *is the cost function that specifies the amount of resources needed to make the observation. The set of cost values* $\mathcal{C}$ *depends on the context. For example, when the resource in question is time,* $\mathcal{C}$ *can be the set of positive real numbers. For memory usage, costs can be conveniently represented by natural numbers. In case of consumption of multiple resources, the cost can be a vector of numbers, such that each number represents the consumption of a different type of resource. To simplify the presentation, we will assume that* $\mathcal{C} = \mathbb{N} \cup \{0\}$ *throughout the paper.*

A simple model is shown in Figure 1, and discussed in detail in Example 2.

It is fair to mention that actions are not modeled explicitly in our framework. They only appear in the form of observations and the resulting model updates.

2.2 Queries and Updates

Definition 2 (Update by an observation). *Let* $m \subseteq S$ *be a subset of worlds. Typically,* m *will be an abstraction class of* $R(a, \cdot, \cdot)$*, i.e., the set of worlds considered possible by agent* a *at the current moment. Moreover, let* $q \in Obs$ *be an observation, and* $s \in m$ *a state. The update of* m *by observation* q *in state* s *is defined as follows:*

$$m|_q^s = \begin{cases} m \cap cover(q) & \text{if } s \in cover(q) \\ m \setminus cover(q) & \text{if } s \notin cover(q). \end{cases}$$

Definition 3 (Query). *A query is a finite sequence of observations, i.e., a tuple* $l = \langle q_1, \ldots, q_k \rangle$ *where each* q_i *is an observation.*

Definition 4 (Update by a query). *An update of a subset of worlds $m \subseteq S$ by a query $l = \langle q_1, q_2, \ldots q_k \rangle$ in state s is defined recursively as follows:*

$$m|_l^s = m|_{q_1,q_2,\ldots,q_k}^s = \left(m|_{q_1,\ldots,q_{k-1}}^s\right)|_{q_k}^s$$

After updating the initial set m by the first observation in the sequence, the updated set of worlds is the new set of worlds that is used to be updated by the next observation in the sequence. This process continues until updating by the last observation in the sequence is done. Note that, in a way, we treat observations as a source of "negative" information that allows to eliminate some possible worlds from the agent's view of the system. Moreover, we assume that the state of the system does not change. In this sense, our framework can be seen as a special model of belief update, rather than belief revision.

Example 2. Consider the medical agent scenario. In Figure 1, the set of possible worlds $m = \{s_1, s_2, s_3, s_4\}$ represents the diseases consistent with the symptoms of the patient (say, pneumonia, meningitis, leukaemia, and chronic kidney disease). The available medical tests for the medical agent a in this example are the observations q_1 and q_2, which respectively correspond to checking the temperature of the patient and checking her blood pressure. The covering set of the observation q_1 is $\{s_1, s_2\}$, i.e., the diseases with high temperature, and the covering set of q_2 is $\{s_2, s_4\}$, that is, the diseases characterized by high blood pressure. Suppose that the actual disease is s_1 (pneumonia) and the medical agent first checks the temperature and then the blood pressure. It means that we would like to find the update of the set m in state s_1 by observations q_1 and q_2. Checking the temperature tells the agent whether the actual state is in the covering set of q_1 or not. Here the answer is "yes", and thus we have $m|_{q_1}^{s_1} = m \cap \mathit{cover}(q_1) = \{s_1, s_2\}$. Checking the blood pressure after this corresponds to updating the result of the previous update $\{s_1, s_2\}$ by observation q_2. In state s_1, the final result is $\{s_1\}$, so the agent knows precisely that the disease is pneumonia.

Definition 5 (Cost of a query). *Let $\oplus : \mathcal{C} \times \mathcal{C} \rightarrow \mathcal{C}$ be a fixed additive aggregation function [19]. The cost of a query is the aggregation of the costs of its observations: $\mathit{cost}(\langle q_1, \ldots, q_k \rangle) = \mathit{cost}(q_1) \oplus \cdots \oplus \mathit{cost}(q_k)$.*

The aggregation function $\oplus$ is context-dependent, and can be defined in various ways. For example, if the resource is time and observations are made sequentially then the aggregate cost is simply the sum of individual costs. If the observations are applied in parallel, the time needed for the whole query is the maximum of the costs, and so on. In this paper, we assume that $\mathit{cost}(\langle q_1, \ldots, q_k \rangle) = \mathit{cost}(q_1) + \cdots + \mathit{cost}(q_k)$, and leave the general case for future work.

Definition 6 (Relevant observation). *The observation q is called relevant to a set $m \subseteq S$ iff $m \cap \mathit{cover}(q) \neq \emptyset$ and $m \cap \mathit{cover}(q) \neq m$.*

If m is the set of worlds that the agent considers possible, a relevant observation is one that brings new information to the agent. In other words, when an

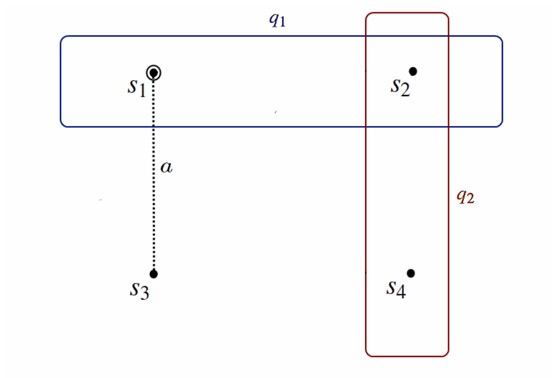

Fig. 2. A model of diagnosis for a more knowledgeable medical agent

observation is not relevant, the agent should know the result of updating even before applying the observation.

Example 3. Consider model M_2 in Figure 2. The set $m = \{s_1, s_3\}$ collects the diseases that the medical agent takes into account. It is easy to see that q_2 is not relevant because the agent already knows that the patient does not have high blood pressure. In other words an update of m by q_2 is equal to m itself. But the agent does not know the result of checking the temperature, therefore q_1 is a relevant observation.

Definition 7 (Relevant query). *Let $a \in A$ and $s \in S$. A query $l = \langle q_1, \ldots, q_k \rangle$ is relevant for agent a in state s iff: (1) $q_i \in obs(a)$ for all i, (2) q_1 is relevant to $\{s' \mid s \sim_a s'\}$, and (3) q_i is relevant to $\{s' \mid s \sim_a s'\}|^s_{q_1,\ldots,q_{i-1}}$ for all $i \geq 2$.*

Note that, while we defined the relevance of an observation with respect to a set of worlds, we use a set *and a state* to define the relevance of a query. This is because in the process of updating a set by a query, in each step, the outcome of the update depends on the actual state. This implies that an agent who does not know what the actual world is, might not know beforehand whether a query is relevant or not. However, the agent knows at each step if the next observation to be applied is relevant or not. Note also that in a state, the same query might be relevant for one agent, and irrelevant for another agent.

Another interesting issue is that, while the order of observations does not change the outcome of updates, it may change the relevance of a query. This is because the relevance of an observation in a query depends only on whether it brings new information *at the moment when it is applied.* In particular, for q_1, q_2 such that $s \in cover(q_1) \subsetneq cover(q_2) \subsetneq m$, we have that query $\langle q_2, q_1 \rangle$ is relevant for m in s, but query $\langle q_1, q_2 \rangle$ is not.

Finally, we remark that modeling of the outcome of observations in terms of global states in a Kripke model can be impractical. One may overcome that by using a higher-level model specification language, for instance one based on interpreted systems [17]. We do not dig deeper into this issue, and discuss only the abstract formulation throughout the paper.

3 A Logic of Accumulative Knowledge

In this section, we introduce a modal language for reasoning about the abilities of agents to refine their knowledge under bounded resources.

3.1 Syntax

The set of formulas of Logic of Accumulative Knowledge (LAcK) is defined by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \psi \mid K_a\varphi \mid \mathfrak{K}_a^l\varphi \mid \Diamond\mathfrak{K}_a^b\varphi \mid \Box\mathfrak{K}_a^b\varphi,$$

where $p \in P$ is an atomic proposition, $a \in A$ is an agent, $l = \langle q_1, ..., q_k \rangle$ is a query such that $q_i \in obs(a)$ for each $i \in [1, ..., k]$, and $b \in \mathcal{B}$ is a resource bound. Unless explicitly stated, we will assume that the set of bounds is $\mathcal{B} = \mathbb{N} \cup \{0, \infty\}$. The other Boolean operators are defined as usual. Additionally, we define $\mathfrak{K}_a\varphi \equiv K_a\varphi \vee K_a\neg\varphi$.

Formula $K_a\varphi$ says that agent a knows that φ. Consequently, $\mathfrak{K}_a\varphi$ expresses that a has no uncertainty about φ, that is, he *knows the truth value* of φ. The formula $\mathfrak{K}_a^l\varphi$ says that a has *observation-based certainty* about φ through query l. Formula $\Diamond\mathfrak{K}_a^b\varphi$ reads as "a can *possibly* (or *potentially*) obtain certainty about φ under resource bound b". Finally, $\Box\mathfrak{K}_a^b\varphi$ expresses that a is *guaranteed* to obtain certainty about φ under bound b.

3.2 Semantics

The semantics of LAcK in observation-based certainty models is defined by the following clauses:

- $M, s \models p$ iff $s \in V(p)$, for any $p \in P$.
- $M, s \models \neg\varphi$ iff $M, s \not\models \varphi$.
- $M, s \models \varphi \vee \psi$ iff $M, s \models \varphi$ or $M, s \models \psi$.
- $M, s \models K_a\varphi$ iff $\forall s' \in m_a(s) :\ M, s' \models \varphi$, where $m_a(s) = \{s' | s \sim_a s'\}$ denotes the set of states indistinguishable from s by agent a.
- $M, s \models \mathfrak{K}_a^l\varphi$ where $l = \langle q_1, \ldots, q_k \rangle$, iff either $\forall s' \in m_a(s)|_l^s :\ M, s' \models \varphi$ or $\forall s' \in m_a(s)|_l^s :\ M, s' \models \neg\varphi$. We call such l an answer query for (a, φ) in s.
- $M, s \models \Diamond\mathfrak{K}_a^b\varphi$ iff for some query l that is relevant for a in s, $M, s \models \mathfrak{K}_a^l\varphi$ and $cost(l) \leq b$.
  Potential certainty expresses that under a given resource bound, the agent has a way to obtain certainty by applying some relevant observations. Note that this does not guarantee that the agent *will* obtain the certainty, since he may not know exactly what query is the right one in the given state s.

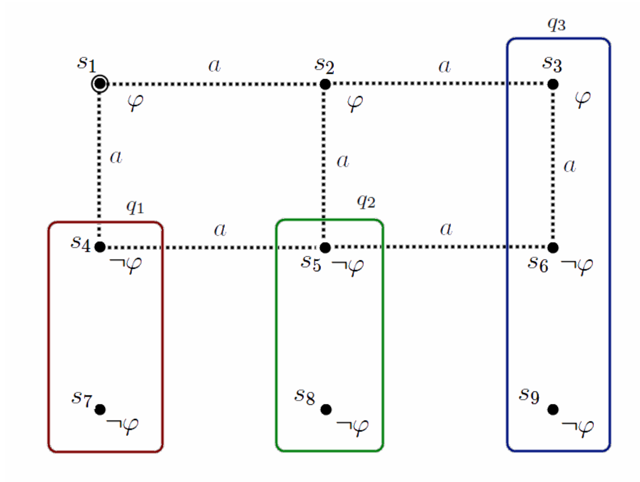

Fig. 3. Observation-based certainty about φ using $\langle q_1, q_2, q_3 \rangle$. The actual state is s_1.

- $M, s \models \Box\mathfrak{K}_a^b\varphi$ iff, for all queries l which are relevant for a in s and $cost(l) \leq b$, we have either $M, s \models \mathfrak{K}_a^l\varphi$, or there exists a query l' so that $M, s \models \mathfrak{K}_a^{l \cdot l'}\varphi$ and $cost(l \cdot l') \leq b$.
  Guaranteed certainty expresses that the agent, by applying relevant and possible observations in any order, obtains certainty without running out of resources.

We can equivalently define guaranteed and potential certainty as follows:

- $M, s \models \Box\mathfrak{K}_a^b\varphi$ iff, for all relevant queries l for agent a in s which are maximal under bound b (meaning that adding any observation to the query makes its cost more than b or makes the query not relevant), l resolves the uncertainty of a about φ in s.
- $M, s \models \Diamond\mathfrak{K}_a^b\varphi$ iff for some query l that is relevant for a in s and maximal under bound b, we have $M, s \models \mathfrak{K}_a^l\varphi$ and $cost(l) \leq b$.

Proposition 1. *Certainty about φ is equivalent to certainty about $\neg\varphi$. Thus, the following equivalences are trivially valid in LAcK:*

$$\mathfrak{K}_a^l\varphi \leftrightarrow \mathfrak{K}_a^l\neg\varphi, \qquad \Diamond\mathfrak{K}_a^b\varphi \leftrightarrow \Diamond\mathfrak{K}_a^b\neg\varphi, \qquad \Box\mathfrak{K}_a^b\varphi \leftrightarrow \Box\mathfrak{K}_a^b\neg\varphi.$$

On the other hand, potential and guaranteed certainty are not *dual, i.e., the following formulae are* not *valid in LAcK:*

$$\neg\Diamond\mathfrak{K}_a^b\varphi \leftrightarrow \Box\mathfrak{K}_a^b\neg\varphi, \qquad \neg\Box\mathfrak{K}_a^b\varphi \leftrightarrow \Diamond\mathfrak{K}_a^b\neg\varphi.$$

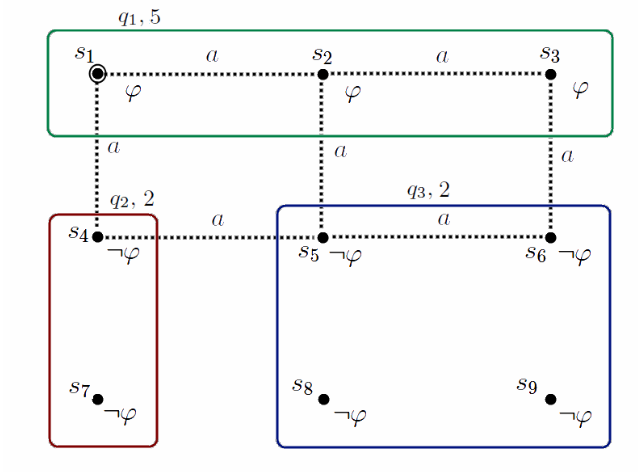

Fig. 4. Guaranteed certainty about φ under bound 5

3.3 Examples

Example 4. Consider model M_3 in Figure 3. The initial set of possible diseases for the medical agent a is $\{s_1, s_2, \ldots, s_6\}$. Moreover, proposition φ is true in a state if the corresponding disease is dangerous, otherwise it is false. In this example, in some of the possible worlds for the agent the formula φ is true, while in some it is not. Therefore initially the agent is not certain about seriousness of the disease. There are three types of available medical tests, corresponding to observations q_1, q_2 and q_3. After updating its initial set of possible worlds with $l_1 = \langle q_1, q_2, q_3 \rangle$, the agent gets $\{s_1, s_2\}$. Since φ is true in both s_1 and s_2, we have that the agent can use l_1 to become certain about the truth value of φ. We denote this by $M_3, s_1 \models \mathfrak{K}_a^{l_1} \varphi$. If the agent prescribes only the medical tests q_2 and q_3, then the updated set is $\{s_1, s_2, s_4\}$. As φ is neither true in all elements of this set, nor false in all of them, we have $M_3, s_1 \not\models \mathfrak{K}_a^{l_2} \varphi$, $(l_2 = \langle q_2, q_3 \rangle)$.

Example 5. Now consider model M_4 in Figure 4. As before, the available medical tests are represented with observations q_1, q_2, and q_3, and the actual disease is s_1. Next to each observation there is a number that shows the cost of applying that medical test. In our example the cost of a test is the time needed to execute it. So, the time needed for test q_1 is 5 hours, and for tests q_2, q_3 it is 2 hours each. The medical agent is not certain about the seriousness of the disease, but it is guaranteed to be certain about it if it has at least 5 hours for doing the tests. To see why, first note that initially all the three observations are relevant and their costs are all lower than the bound. Thus, the agent has three choices. If it chooses q_1 as the first test, the updated set is $\{s_1, s_2, s_3\}$. In all the worlds in this set φ is true, so the agent has obtained certainty (in this case, it knows that the disease is dangerous). If the agent chooses q_2 first, the updated set is $\{s_1, s_2, s_3, s_5, s_6\}$.

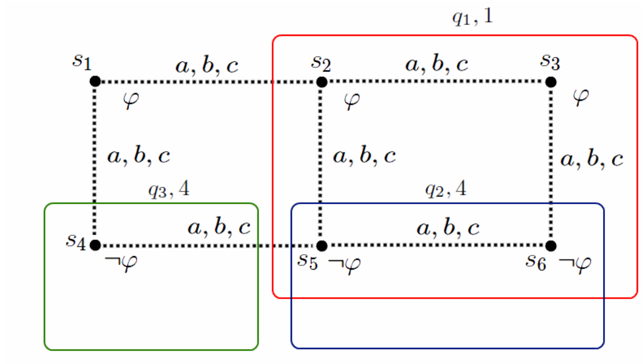

Fig. 5. Model M_5. If $obs(a) = \{q_1\}$, $obs(b) = \{q_1, q_2\}$, and $obs(c) = \{q_1, q_3\}$, then agent a has observation-based certainty about certainty of agents b and c.

The agent is not certain yet, and has to continue applying observations. The time needed for applying q_2 is 2, so after applying q_2 the agent has $5 - 2 = 3$ hours left, during which it only can apply q_3. After updating by q_3, the updated set is $\{s_1, s_2, s_3\}$, and the certainty is gained. The result of applying q_3 first is similar to the previous case, except that this time the second observation is inevitably q_2. So, if the agent at each step chooses only relevant and feasible observations, it will attain certainty about φ under bound 5. Therefore in this example $M_4, s_1 \models \Box\mathfrak{K}_a^5\varphi$.

If the agent had only 3 hours for tests, the only possible choices would be q_2 and q_3. But after updating the set of its possible worlds with any of these observations the agent would still be uncertain about the seriousness of the disease and the remaining time would not suffice for applying any more observations. Therefore $M_4, s_1 \not\models \Diamond\mathfrak{K}_a^3\varphi$.

Note, finally, that if the actual disease is s_4, the agent might obtain certainty within 3 hours. This is possible by choosing observation q_2. On the other hand, if a chooses q_3 first, it will not obtain certainty within the same time. Thus, in state s_4, the agent has potential but not guaranteed certainty about φ, i.e., $M_4, s_4 \models \Diamond\mathfrak{K}_a^3\varphi \wedge \neg\Box\mathfrak{K}_a^3\varphi$.

Example 6. Nested formulas refer to an agent's certainty about its own, or another agent's certainty. Consider medical agent a that is not sufficiently equipped to become certain about the seriousness of the disease (Figure 5). Instead, a has to decide to which specialized medical center the patient should be sent. The medical centre b specializes in brain diseases, and the medical center c specializes in heart diseases. Test q_1 is a general test and is available for all the agents a, b and c. Test q_2 is only available at the brain centre, and test q_3 is only available at the heart center. So in this model $obs(a) = \{q_1\}$, $obs(b) = \{q_1, q_2\}$ and $obs(c) = \{q_1, q_3\}$. By applying the medical test q_1 the medical agent a is able to determine which specialized center are competent to do the diagnosis in a given time (say, up to 5 hours), and which are not: $M_5, s_1 \models \mathfrak{K}_a^{\langle q_1\rangle}\Diamond\mathfrak{K}_b^5\varphi \wedge \mathfrak{K}_a^{\langle q_1\rangle}\Diamond\mathfrak{K}_c^5\varphi$.

In consequence, we also have that $M_5, s_1 \models \Diamond\mathfrak{K}^1_a\Diamond\mathfrak{K}^5_b\varphi \wedge \Diamond\mathfrak{K}^1_a\Diamond\mathfrak{K}^5_c\varphi$ and $M_5, s_1 \models \Box\mathfrak{K}^1_a\Diamond\mathfrak{K}^5_b\varphi \wedge \Box\mathfrak{K}^1_a\Diamond\mathfrak{K}^5_c\varphi$: the agent has both potential and guaranteed certainty to learn about b and c's epistemic abilities within 1 hour.

Remark 1. The semantics of nested modalities in LAcK is somewhat counterintuitive when the nested and the outer modalities refer to the same agent, like in formula $\Diamond\mathfrak{K}^b_a\Diamond\mathfrak{K}^{b'}_a\mathsf{p}$. In that case, one would expect agent a in subformula $\Diamond\mathfrak{K}^{b'}_a\mathsf{p}$ to make use of the observations already performed while evaluating the outer modality. Instead, our semantics makes the agent collect knowledge anew.

One possible way out could be to change our update operator so that it updates directly the epistemic relation of agent a, rather than the additional set m representing the temporary view of the agent during the query execution. That would work similar to *public announcement* in dynamic epistemic logic [34]. Another possibility is to evaluate formulae in a model, state, *and a set of states* m representing the current epistemic scope of evaluation. Then, m can be "passed" to the subformula, thus implementing "knowledge transfer" between the agents referred to by the outer and the nested modalities. We leave a more thorough semantic investigation of nested observation-based certainty for future work.

4 Some Properties

In this section, we present some interesting properties that can be expressed in LAcK. We begin by listing some validities that capture interesting general properties of accumulative knowledge. Then, in Section 4.2, we show how the basic information-theoretic notion of *Hartley measure* can be characterized in our framework.

4.1 Interesting Validities

Below we list some interesting validities of LAcK. We give only some of the proofs; the others are either straightforward or analogous.

Theorem 1. *The following formulas are valid in LAcK:*

1. $K_a\varphi \to \mathfrak{K}^l_a\varphi$.
   Certainty cannot be destroyed by observations.
2. $\mathfrak{K}^l_a\varphi \to \mathfrak{K}^{l\cdot l'}_a\varphi$.
   A more general variant of 1.
3. $\mathfrak{K}^l_a\varphi \wedge \mathfrak{K}^l_a\psi \to \mathfrak{K}^l_a(\varphi \wedge \psi)$.
   Outcomes of a query combine.
4. $\mathfrak{K}^l_a\varphi \wedge \mathfrak{K}^{l'}_a\psi \to \mathfrak{K}^{l\cdot l'}_a(\varphi \wedge \psi)$.
   Combining queries yields combined outcomes.
5. $K_a\varphi \to \Diamond\mathfrak{K}^b_a\varphi$ *and* $K_a\varphi \to \Box\mathfrak{K}^b_a\varphi$.
   A variant of 1 for potential and guaranteed observation-based certainty.
6. $\Diamond\mathfrak{K}^b_a\varphi \to \Diamond\mathfrak{K}^{b+b'}_a\varphi$ *and* $\Box\mathfrak{K}^b_a\varphi \to \Box\mathfrak{K}^{b+b'}_a\varphi$.
   Monotonicity of observation-based certainty wrt resource bounds.

7. $\Diamond\mathfrak{K}_a^b\varphi \wedge \Diamond\mathfrak{K}_a^{b'}\psi \rightarrow \Diamond\mathfrak{K}_a^{b+b'}(\varphi \wedge \psi)$.
   Variant of 4 for potential certainty.
8. $\Box\mathfrak{K}_a^b\varphi \wedge \Box\mathfrak{K}_a^{b'}\psi \rightarrow \Diamond\mathfrak{K}_a^{max(b,b')}(\varphi \wedge \psi)$.
   Variant of 4 for guaranteed certainty.
9. $\Box\mathfrak{K}_a^b\varphi \rightarrow \Diamond\mathfrak{K}_a^b\varphi$.
   Guaranteed certainty implies potential certainty.
10. $\Box\mathfrak{K}_a^\infty\varphi \leftrightarrow \Diamond\mathfrak{K}_a^\infty\varphi$.
   For unlimited resources, potential and guaranteed certainty coincide. Note that this validity can be also given in a stronger form: for $b \geq \sum_{o\in obs(a)} cost(o)$, $\Box\mathfrak{K}_a^b\varphi \leftrightarrow \Diamond\mathfrak{K}_a^b\varphi$ *is also valid.*

Proof.
Ad. 7: From the antecedent, we know that there is an answer query l for (a,φ) such that $cost(l) \leq b$, and there is an answer query l' for (a,ψ) such that $cost(l') \leq b'$. Therefore $l \cdot l'$ is answer query for $(a, \varphi \wedge \psi)$, and $cost(l \cdot l') = cost(l) + cost(l') \leq b + b'$.

Ad. 8: Assume that $max(b,b') = b$. Then by definition, any relevant maximal query l under bound b is an answer query for (a,φ). Then there is a relevant maximal query l under b such that $l = l_1 \cdot l_2$ for some l_1 and l_2, with l_1 being a relevant maximal query under bound b'. From $\Box\mathfrak{K}_a^{b'}\psi$ we know that l_1 is an answer query for (a,ψ). Therefore $l = l_1 \cdot l_2$ is also an answer query for (a,ψ). As l is an answer query both for (a,φ) and for (a,ψ), it is an answer query for $(a, \varphi \wedge \psi)$. The proof is similar in the case that $max(b,b') = b'$.

Ad. 10: Inferring $\Diamond\mathfrak{K}_a^\infty\varphi$ from $\Box\mathfrak{K}_a^\infty\varphi$ is a direct result of the previous property. For proving the other direction, first note that changing the order of the observations in a query does not change the updated set of worlds, and adding some observations to a query cannot make an answer query a non-answer query. Now if we have $\Diamond\mathfrak{K}_a^\infty\varphi$, then there is an answer query l for (a,φ). Therefore any query l' which consists of all the available observations is also an answer query for (a,φ). As the upper limit for the resource is infinity, the agent can choose the observations in any order and it is guaranteed to be certain about φ without running out of resources, hence $\Box\mathfrak{K}_a^\infty\varphi$. □

4.2 Relation to Information Theory

In the previous sections we have defined a framework for reasoning about agents that collect information in order to become certain about a given property. In other words, the agents reduce their uncertainty about the property by accumulating observations. There seems to be an intuitive connection to the classic definition of uncertainty, and in particular Hartley measure of uncertainty. In this section, we look at the relationship.

Two most established measures of uncertainty are *Hartley measure* and *Shannon entropy*. Hartley measure is based on possibility theory, whereas Shannon entropy is based on probability theory. Hartley measure quantifies uncertainty in terms of a finite set of possible outcomes. Let X be the set of all alternatives under consideration, out of which only one is considered the *correct one*. Note that this can be seen as corresponding to the set of possible worlds and the actual world, respectively. It was shown by Hartley [21] that the only sensible way to measure the uncertainty about the correct alternative in a set of alternatives X is to use the function:

$$H(X) = \lceil \log_2 |X| \rceil.$$

The unit of uncertainty measured by $H(X)$ is *bit*. The intuition behind Hartley measure is that $\log_2 |X|$ is the minimal number of binary questions that guarantees identifying the correct alternative, provided that the set of questions is rich enough. We will now use the intuition to characterize Hartley measure in LAcK.

Definition 8 (Bisective Observations). *Let $n[i]$ denote the ith bit in the binary unfolding of n. A set of observations O is* bisective *for states S iff there is a bijective ordering of states $ord : S \to \{1, \ldots, |S|\}$ and a bijective mapping $bitno : O \to \{1, \ldots, \lceil \log |S| \rceil\}$ such that $cover(q) = \{s \in S \mid \big(ord(s)\big)[bitno(q)] = 1\}$ for every $q \in Q$. Thus, each observation q separates states with different values of $bitno(q)$ in their binary unfolding. Putting it differently, we see S as a k-dimensional binary cube, with each $q \in Q$ "cutting across" a different dimension.*

Definition 9 (Distinguishing model). *A possible worlds model M is* distinguishing by formulas $\psi_1, \ldots, \psi_k$ *iff for every state s_i in M there exists ψ_i which holds exactly in s_i.*

Definition 10 (Hartley model, Hartley formula). *We say that an observation-based certainty model $M = \langle S, R, V, Obs, obs, cost, cover\rangle$ is a* Hartley model *iff:*

1. *M consists of a single agent a (the "observer"),*
2. *$R(a, \cdot, \cdot) = S \times S$,*
3. *M is distinguishing by some formulas $\psi_1, \ldots, \psi_k$,*
4. *$obs(a)$ is bisective for S, and*
5. *The cost of every observation is* 1.

The Hartley formula of M under bound b *is defined as:* $\chi(M, b) \equiv \bigwedge_{i\in|S|} \Diamond \mathfrak{K}_a^b \psi_i$.

Intuitively, Hartley formula in a Hartley model expresses that the observer will identify the actual world in at most b steps if he iteratively bisects the set of possibilities.

Theorem 2. *Let M be a Hartley model with state space S. Then, for all $s \in S$, we have $M, s \models_{\mathsf{LAcK}} \chi(M, H(S))$.*

Proof. Take a query consisting of all the (bisective) observations in M. Clearly, the query updates any set of indistinguishable states yielding the singleton set containing only the actual state. Moreover, it consists of at most $H(S)$ steps, which concludes the proof. □

5 Model Checking

In this section, we look at the complexity of verification for accumulative knowledge. Similarly to many problems where agents' uncertainty is involved, it turns out to be **NP**-hard. We also show that the hardness of the problem is due to bounded resources. Finally, we prove that verification becomes tractable in many realistic scenarios where resource bounds are relatively tight.

5.1 General Result

The (local) model checking problem for LAcK is formally defined as follows.

Definition 11 (Model checking for LAcK).
Input: *Observation-based certainty model M, state s in M, LAcK formula φ;*
Output: *yes iff $M, s \models_{\mathsf{LAcK}} \varphi$.*

We will show that the problem sits in the first level of polynomial hierarchy, more precisely between $\mathbf{NP} \cup \mathbf{coNP}$ and $\mathbf{\Delta_2^P}$ (where $\mathbf{\Delta_2^P} = \mathbf{P^{NP}}$ is the class of problems that can be solved in polynomial by a deterministic Turing machine asking adaptive queries to an **NP** oracle). We start by showing the upper bound.

Proposition 2. *Model checking LAcK is in $\mathbf{\Delta_2^P}$.*

Proof. We demonstrate the upper bound by the following algorithm.

$\underline{mcheck(M, s, \varphi)}$:

Case $\varphi \equiv p$ **:** return($s \in V(p)$);
Cases $\varphi \equiv \neg\psi, \psi_1 \wedge \psi_2, K_a\psi$ **:** standard;
Case $\varphi \equiv \mathfrak{K}_a^l\psi$ **:** $X := \{s' \in S \mid mcheck(M, s', \psi)\}$;
return($m_a(s)|_l^s \subseteq X$ or $m_a(s)|_l^s \subseteq S \setminus X$);
Case $\varphi \equiv \Diamond\mathfrak{K}_a^b\psi$ **:** return($oracle_1(M, s, a, \psi)$);
Case $\varphi \equiv \Box\mathfrak{K}_a^b\psi$ **:** return(not $oracle_2(M, s, a, \psi)$);

$\underline{oracle_1(M, s, a, \psi)}$:

$X := \{s' \in S \mid mcheck(M, s', \psi)\}$;
guess a query l with no repeated observations;
return$\big(cost(l) \leq b$ and $(m_a(s)|_l^s \subseteq X$ or $m_a(s)|_l^s \subseteq S \setminus X)\big)$;

$\underline{oracle_2(M, s, a, \psi)}$:

$X := \{s' \in S \mid mcheck(M, s', \psi)\}$;
guess a relevant query l;
$maximal := (cost(l) \leq b$ and for all $q \notin l$: $cost(lq) > b$ or lq is not relevant);
return$\big(maximal$ and $m_a(s)|_l^s \not\subseteq X$ and $m_a(s)|_l^s \not\subseteq S \setminus X\big)$; □

To prove the lower bound, we will use an old result by Karp [24].

Definition 12 ([24]). SetCovering *is the following decision problem.*
Input: *Domain of elements D, a finite family of finite sets $\mathcal{S} = \{S_1, \ldots, S_n\} \subseteq 2^D$, and a number $k \in \mathbb{N}$;*
Output: *yes iff there exists a family of k sets $\mathcal{T} = \{T_1, T_2, \ldots, T_k\} \subseteq \mathcal{S}$ such that $\bigcup_j T_j = \bigcup_i S_i$.*

Proposition 3 ([24]). SetCovering *is* **NP***-complete.*

Lemma 1. *Model checking of the LAcK formula $\Diamond\mathfrak{K}_a^b p$ is* **NP***-complete.*

Proof. Inclusion in **NP** follows from the algorithm in the proof of Proposition 2. The lower bound is obtained by a reduction of SetCovering. Let M include:

- $S = D \cup \{s_0\}$ for some $s_0 \notin D$;
- $A = \{a\}$, and $\sim_a = S \times S$;
- $Obs = \{q_1, \ldots, q_n\}$, $obs(a) = Obs$, and $cover(q_i) = S_i$;
- $cost(q_i) = 1$ for every i;
- single atomic proposition p_0 with $V(p_0) = \{s_0\}$.

Now, SetCovering$(D, \{S_1, \ldots, S_n\}, k)$ iff $M, s_0 \models_{\mathsf{LAcK}} \Diamond\mathfrak{K}_a^k p_0$. □

The following is a straightforward consequence (note that we can use negation to obtain the complement of a problem expressible in LAcK).

Proposition 4. *Model checking LAcK is* **NP***-hard and* **coNP***-hard.*

Thus, finally, we obtain the following result.

Theorem 3. *Model checking LAcK is between* $(\mathbf{NP} \cup \mathbf{coNP})$ *and* $\mathbf{\Delta_2^P}$*.*

5.2 Closer Look

What is the hard part of the verification problem for LAcK? The next result shows that the hardness is due to bounded resources, since with unlimited resources the problem becomes easy.

Proposition 5. *If $\mathcal{B} = \{\infty\}$ then model checking LAcK is in* **P***.*

Proof. First, observe that $M, s \models \Diamond\mathfrak{K}_a^\infty\varphi$ iff $M, s \models \mathfrak{K}_a^l\varphi$ for l being the "grand query" collecting all the observations available for a in M. Moreover, $M, s \models \Box\mathfrak{K}_a^\infty\varphi$ iff $M, s \models \Diamond\mathfrak{K}_a^\infty\varphi$ by Theorem 1, point 10. For the other cases, we proceed according to the algorithm in the proof of Proposition 2. It is easy to see that the new algorithm terminates in time $O(|S| \cdot |Obs| \cdot |\varphi|)$. □

Finally, we want to suggest that the pessimistic view of Theorem 3 is not always justified. True, verification is **NP**-hard in general. However, we argue that it only makes sense to engage in checking $M, s \models \Diamond\mathfrak{K}_a^b\varphi$ or $M, s \models \Box\mathfrak{K}_a^b\varphi$ if a's observations are relatively expensive compared to the available resources b. After all, if observations were cheap, a might as well skip deliberation and start observing right away. The following result shows that when the relation between costs and bounds is tight, the model checking problem becomes easy again.

Proposition 6. *Let $\alpha > 1$ be given and fixed. Model checking $\Diamond\mathfrak{K}_a^b p$ and $\Box\mathfrak{K}_a^b p$ in a model such that* $\min\{cost(q) \mid q \in Obs\} \geq \frac{b}{\alpha(\log|S|+\log|Obs|+\log b)}$ *is in* $\mathbf{P}$.

Proof. If $\min\{cost(q) \mid q \in Obs\} \geq \frac{b}{\alpha(\log|S|+\log|Obs|+\log b)}$ then every query that consists of more than $\alpha(\log|S| + \log|Obs| + \log b)$ observations will cost more than b. Thus, it suffices to check the outcome of at most $2^\alpha \cdot b \cdot |S| \cdot |Obs|$ queries, which is polynomial in the size of the model. Note that, for this result, it is essential that α is not a parameter of the problem. □

6 Reasoning about Strategic Exploration

LAcK provides modalities for reasoning about what an agent *might* discover if he asks the right questions, and what he *must* discover if he queries the environment exhaustively. In this section, we propose a variant that allows to reason about what agents can unveil by systematic *strategic* exploration of the environment.

We begin by formally extending LAcK to handle strategic exploration. Then, we show that the extension is substantial in the sense that it allows to express properties that cannot be captured by the temporal-epistemic modalities of LAcK alone. As it turns out, the temporal-epistemic modalities of LAcK also cannot be simulated by the new strategic modality. In consequence, both kinds of modalities offer complementary views of explorative abilities. Finally, we show that the increased expressivity comes with no computational price: model checking of the broader language is still in the first level of the polynomial hierarchy.

6.1 Strategic Extension of LAcK

We extend LAcK with formulas $\blacklozenge\mathfrak{K}_a^b\varphi$ reading: "agent a knows how to explore the environment in order to obtain certainty about φ using at most b units of resources." Formally, the syntax of *Logic of Strategic Accumulative Knowledge* (SLAcK) is defined by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \psi \mid K_a\varphi \mid \mathfrak{K}_a^l\varphi \mid \Diamond\mathfrak{K}_a^b\varphi \mid \Box\mathfrak{K}_a^b\varphi \mid \blacklozenge\mathfrak{K}_a^b\varphi,$$

We refer to the fragment of SLAcK using only modalities $K_a, \Diamond\mathfrak{K}_a^b, \Box\mathfrak{K}_a^b$ as the *temporal-epistemic part of SLAcK (SLAcK^{te})*. The fragment of SLAcK using only the strategic modality $\blacklozenge\mathfrak{K}_a^b$ is called the *strategic part of SLAcK (SLAcK^{s})*.

In order to propose semantics of the new modality, we first define a proper notion of strategic behavior.

Definition 13 (Exploration strategy). *Let A be a set of agents, $a \in A$ an agent, and $M = \langle S, R, V, Obs, obs, cost, cover\rangle$ an observation-based certainty model. An* exploration strategy for player a *is any function of type $f_a : 2^S \to obs(a)$. That is, an exploration strategy assigns exploration-related choices (questions, observations) to any potential epistemic position of agent a.*

The outcome of strategy f_a is understood as the sequence of epistemic positions of agent a that are effected by subsequent updates prescribed by f_a.

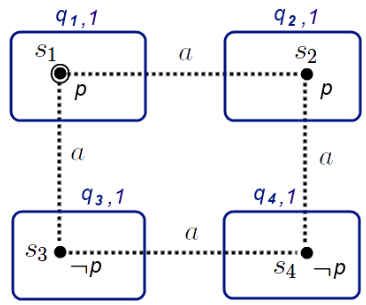

Fig. 6. Model M_6: difference between possible certainty, guaranteed certainty, and strategic certainty

Definition 14 (Outcome of exploration). *Given an observation-based certainty model M, a state $s \in S$ in the model, and an exploration strategy f_a for player a, the function $out(s, f_a)$ returns the (infinite) sequence $\langle X_0, X_1, X_2, ...\rangle$, such that $X_0 = \{s' \mid s{\sim_a}s'\}$ and $X_{k+1} = X_k|^s_{f_a(X_k)}$.*

We have defined the outcome as an infinite sequence for mathematical convenience. Note, however, that if the set of states *or* the set of observations is finite then the sequence will stabilize after a finite number of steps – in fact, no more than $\min(|obs(a)|, 2^{|S|})$. In that case, only a finite initial segment of the sequence matters in the context of strategic exploration.

The aim of a strategy is to obtain certainty before running out of resources. This is formalized by the following definition.

Definition 15 (Achieving certainty). *We say that strategy f_a* achieves certainty about φ from state s under bound b *iff, for $out(s, f_a) = \langle X_0, X_1, X_2, ...\rangle$, there exists k such that: (i) $\sum_{j=0}^{k-1} cost(f_a(X_j)) \leq b$, and (ii) either $M, s' \models \varphi$ for all $s' \in X_k$ or $M, s' \models \neg\varphi$ for all $s' \in X_k$.*

The semantics of SLAcK extends the one for LAcK by the following clause:

– $M, s \models \blacklozenge\mathfrak{K}^b_a\varphi$ iff there is an exploration strategy f_a for agent a such that, for every $s \sim_a s'$, we have that f_a achieves certainty about φ from s' under bound b.

Note that $\blacklozenge\mathfrak{K}^b_a\varphi$ requires that a *single exploration strategy* is successful from all the states indistinguishable from the current state. In this sense, it is closely related to the notion of "knowing how to play" (a.k.a. subjective ability) formalized in modal logics of strategic ability [31, 23, 13].

Example 7. Consider model M_6 in Figure 6 that includes one agent (a) and 4 states, each of them initially indistinguishable to the agent. The agent can test if the system is in a particular state – hence 4 observations: q_1, q_2, q_3, q_4. The cost of each observation is 1.

Let us consider an arbitrary state s_i. It is easy to see that there is a query obtaining certainty about property p in one step (namely, by testing q_i). Thus, $M_6, s_i \models \Diamond\mathfrak{K}_a^1\mathsf{p}$. Since this holds for every state in the indistinguishability class, we even have that $M_6, s_i \models K_a\Diamond\mathfrak{K}_a^1\mathsf{p}$. On the other hand, no exploration strategy yields certainty about p in one step from all the indistinguishable states. Hence, $M_6, s_i \models \neg\blacklozenge\mathfrak{K}_a^1\mathsf{p}$.

Strategic exploration can succeed from all the states only in two steps (example strategy: first test q_1, then test q_2). Thus, the agent knows how to obtain certainty under bound 2: $M_6, s_i \models \blacklozenge\mathfrak{K}_a^2\mathsf{p}$. On the other hand, $M_6, s_i \models \neg\Box\mathfrak{K}_a^2\mathsf{p}$. This is because, if the agent does not plan strategically, he can for instance execute the relevant sequence $\langle q_{(i \, mod\, 4)+1}, q_{((i+2)\, mod\, 4)+1}\rangle$. Consider now the relevant sequence consisting of all the observations except q_i. Since this is the only relevant query that is maximal under bound 3, we do have that $M_6, s_i \models \Box\mathfrak{K}_a^3\mathsf{p}$, but the certainty cannot be guaranteed under bound 2.

The following proposition is straightforward.

Proposition 7. *For any observation-based certainty model M and a state s in M, we have that:*

$$M, s \models \Box\mathfrak{K}_a^b\varphi \quad \Rightarrow \quad M, s \models \blacklozenge\mathfrak{K}_a^b\varphi \quad \Rightarrow \quad M, s \models \Diamond\mathfrak{K}_a^b\varphi.$$

Thus, $\Box\mathfrak{K}_a^b$ and $\Diamond\mathfrak{K}_a^b$ can be seen as upper/lower approximations of explorative ability. We have also seen in Example 7 that these are indeed approximations, i.e., the ability to explore is equivalent to neither possible nor guaranteed certainty. Can we capture strategic exploration by a more complicated combination of possible/guaranteed certainty (perhaps together with epistemic operators)? In the next subsection we show that it is not possible.

6.2 Expressive Power of Strategic Exploration

We argue that the new modality $\blacklozenge\mathfrak{K}_a^b$ offers significant "added value" when reasoning about epistemic exploration. To this end, we will show that properties of strategic exploration cannot be expressed by a combination of knowledge operators and possible/guaranteed certainty modalities. In fact, the converse also applies. In consequence, the temporal-epistemic and strategic fragments of SLAcK offer different views of accumulative knowledge and dynamics of certainty.

We begin by a brief introduction of the notions of *distinguishing power* and *expressive power* (cf. e.g. [35]).

Definition 16 (Distinguishing and expressive power). *Let $L_1 = (\mathcal{L}_1, \models_1)$ and $L_2 = (\mathcal{L}_2, \models_2)$ be two logical systems over the same class of Kripke models $\mathcal{M}$. By $[\![\varphi]\!]_{\models} = \{(M, s) \mid M, s \models \varphi\}$, we denote the class of pointed models that satisfy φ in the semantics given by $\models$. Likewise, $[\![\varphi, M]\!]_{\models} = \{s \mid M, s \models \varphi\}$ is the set of states (or, esuivalently, pointed models) that satisfy φ in a given structure M.*

L_2 is at least as expressive *as L_1 (written: $L_1 \preceq_e L_2$) iff for every formula $\varphi_1 \in \mathcal{L}_1$ there exists $\varphi_2 \in \mathcal{L}_2$ such that $[\![\varphi_1]\!]_{\models_1} = [\![\varphi_2]\!]_{\models_2}$.*

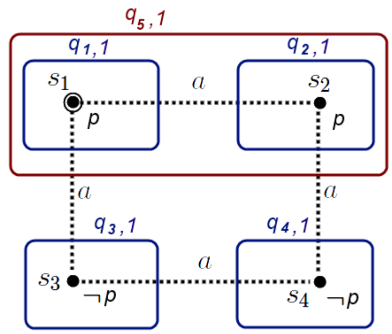

Fig. 7. Model M_7

L_2 *is* at least as distinguishing *as* L_1 *(written:* $L_1 \preceq_d L_2$*) iff for every model* M *and formula* $\varphi_1 \in \mathcal{L}_1$ *there exists* $\varphi_2 \in \mathcal{L}_2$ *such that* $[\![\varphi_1, M]\!]_{\models_1} = [\![\varphi_2, M]\!]_{\models_2}$.[4]

It is easy to see that $L_1 \preceq_e L_2$ implies $L_1 \preceq_d L_2$. By transposition, we also have that $L_1 \not\preceq_d L_2$ implies $L_1 \not\preceq_e L_2$.

Theorem 4. *$\mathsf{SLAcK}^s \not\preceq_d \mathsf{SLAcK}^{te}$ (and hence also $\mathsf{SLAcK}^s \not\preceq_e \mathsf{SLAcK}^{te}$).*

We will first prove the following lemma:

Lemma 2. *Consider models M_6 from Figure 6 and M_7 from Figure 7. We claim that every formula φ of SLAcK^{te} is satisfied in exactly the same states in both models. Moreover, the set of states satisfying φ can be only $\emptyset$, $\{s_1, s_2\}$, $\{s_3, s_4\}$, or $\{s_1, s_2, s_3, s_4\}$.*

Proof (of Lemma 2). We prove the lemma by induction on the structure of formula φ:

Case $\varphi \equiv \mathsf{p}$: obvious.
Cases $\varphi \equiv \neg\psi, \psi_1 \vee \psi_2$: straightforward induction.
Case $\varphi \equiv K_a\psi$: $M_6, s_i \models K_a\psi$ iff ψ holds in all states of M_6. Thus, $K_a\psi$ holds either in all or no state of M_6. The same applies to $M_7, s_i \models K_a\psi$. Moreover, by induction, ψ hold in all states of M_6 iff it holds in all states of M_7, so $K_a\psi$ must be satisfied in the same states in M_6 and M_7.
Case $\varphi \equiv \Diamond\mathfrak{K}^b_a\psi$: By induction, ψ holds in states $X = \emptyset$, $\{s_1, s_2\}$, $\{s_3, s_4\}$, or $\{s_1, s_2, s_3, s_4\}$, and in the same set of states in M_6, M_7. For $X = \emptyset, \{s_1, s_2, s_3, s_4\}$, we have that $M_6/M_7, s_i \models \Diamond\mathfrak{K}^b_a\psi$ for all $b \geq 0$, hence $\Diamond\mathfrak{K}^b_a\psi$ holds in $\{s_1, s_2, s_3, s_4\}$ in both M_6, M_7. For $X = \{s_1, s_2\}, \{s_3, s_4\}$, achieving certainty about ψ is possible in one step regardless of the state: $M_6/M_7, s_i \models \Diamond\mathfrak{K}^b_a\psi$ iff $b \geq 1$. Thus, in both M_6 and M_7, $\Diamond\mathfrak{K}^b_a\psi$ is satisfied in $\emptyset$ for $b < 1$ and in $\{s_1, s_2, s_3, s_4\}$ otherwise.

[4] Equivalently: for every pair of pointed models that can be distinguished by some $\varphi_1 \in \mathcal{L}_1$ there exists $\varphi_2 \in \mathcal{L}_2$ that distinguishes these models.

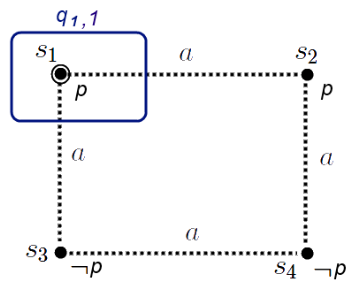

Fig. 8. Model M_8

Case $\varphi \equiv \Box\mathfrak{K}_a^b\psi$**:** Again, by induction, ψ must hold in M_6, M_7 in the same set $X = \emptyset$, $\{s_1, s_2\}$, $\{s_3, s_4\}$, or $\{s_1, s_2, s_3, s_4\}$. For $X = \emptyset, \{s_1, s_2, s_3, s_4\}$, we have $M_6/M_7, s_i \models \Box\mathfrak{K}_a^b\psi$ for all $b \geq 0$, hence $\Box\mathfrak{K}_a^b\psi$ holds in $\{s_1, s_2, s_3, s_4\}$ in both models. For $X = \{s_1, s_2\}, \{s_3, s_4\}$, achieving certainty about ψ is guaranteed in three steps (and no less) regardless of the state: $M_6/M_7, s_i \models \Box\mathfrak{K}_a^b\psi$ iff $b \geq 3$. Thus, in both models $\Box\mathfrak{K}_a^b\psi$ is satisfied in $\emptyset$ for $b < 3$ and in $\{s_1, s_2, s_3, s_4\}$ otherwise. $\Box$

Proof (of Theorem 4). Let us assume by convention that M' denotes model M in which labels of all the states have been "primed". Take model $M_{67} = M_6 \cup (M_7)'$, and consider the pointed models: (M_{67}, s_i) and (M_{67}, s_i'). By Lemma 2, they satisfy exactly the same formulae of $\mathsf{SLAcK^{te}}$. On the other hand, we have that $M_{67}, s_i \not\models \blacklozenge\mathfrak{K}_a^1\mathsf{p}$ and $M_{67}, s_i' \models \blacklozenge\mathfrak{K}_a^1\mathsf{p}$. $\Box$

Theorem 5. *$\mathsf{SLAcK^{te}} \not\preceq_d \mathsf{SLAcK^{s}}$ (and hence also $\mathsf{SLAcK^{te}} \not\preceq_e \mathsf{SLAcK^{s}}$).*

Proof. Consider the pointed models (M_8, s_1) and (M_8, s_2) from Figure 8. We claim that they satisfy exactly the same formulae of $\mathsf{SLAcK^s}$. The proof is analogous to Lemma 2, based on the fact that (i) (M_8, s_1) and (M_8, s_2) agree on atomic propositions, and (ii) a has only the trivial explorative abilities (formally: for every $s_i \in S$, we have $M_8, s_i \models \blacklozenge\mathfrak{K}_a^b\varphi$ iff φ holds in all states).

On the other hand, $M_8, s_1 \models \lozenge\mathfrak{K}_a^1\mathsf{p}$ and $M_8, s_2 \not\models \lozenge\mathfrak{K}_a^1\mathsf{p}$. $\Box$

Corollary 1. *The temporal-epistemic part of SLAcK and the strategic part of SLAcK have incomparable expressive and distinguishing powers.*

6.3 Model Checking SLAcK

We will now prove that extending LAcK with strategic modalities $\blacklozenge\mathfrak{K}_a^b$ does not increase the complexity of verification significantly.

Theorem 6. *Model checking SLAcK is in $\mathbf{\Delta_2^P}$ (hence, between $(\mathbf{NP} \cup \mathbf{coNP})$ and $\mathbf{\Delta_2^P}$).*

Proof. We first observe that verification of $M, s \models \blacklozenge\mathfrak{K}_a^b\varphi$ can be done in nondeterministic polynomial time by the following algorithm:

1. Let $\mathit{Indist} = \{s' \mid s \sim_a s'\}$, and $f_a = \emptyset$;
2. For every $s_i \in \mathit{Indist}$, guess a sequence O_i of at most $|obs(a)|$ observations;
3. For $i = 1$ to $|\mathit{Indist}|$ do
   - $X := \mathit{Indist}$;
   - For $j = 1$ to $|O_i|$ do
     - If $\forall o.(X, o) \notin f_a$ then add $(X, O_i[j])$ to f_a;
     - $X := X|^{s_i}_{f_a(X)}$, where $f_a(X)$ denotes the unique o st. $(X, o) \in f_a$;
   - If there is no certainty about φ in X or the costs of O_i sum up to more than b then return(false);
4. return(true).

By combining the above with Proposition 2, we get that model checking of an arbitrary SLAcK formula is in $P^{NP} = \mathbf{\Delta_2^P}$. □

6.4 Relation to Information Theory Revisited

We wrote in Section 4.2 that the Hartley measure $H(X)$ computes, intuitively, the minimal number of binary questions that guarantees identifying the correct alternative in X, provided that the set of questions is rich enough. However, our characterization of $H(X)$ in LAcK was somewhat forced in the sense that we "hardwired" the optimal exploration strategy (i.e., the bisection strategy) in the model. With the new modality of strategic exploration, we can rephrase the result from Section 4.2 in a more natural way.

Definition 17 (Strategic Hartley model, strategic Hartley formula). *We say that an observation-based certainty model* $M = \langle S, R, V, Obs, obs, cost, cover\rangle$ *is a* strategic Hartley model *iff:*

1. *M consists of a single agent a (the "observer"),*
2. *$R(a, \cdot, \cdot) = S \times S$,*
3. *M is distinguishing by some formulas $\psi_1, \ldots, \psi_k$,*
4. *$obs(a)$* includes *a set of bisective observations for S,*[5] *and*
5. *The cost of every observation is* 1*.*

The strategic Hartley formula of M under bound b *is defined as:* $\sigma(M, b) \equiv \bigwedge_{s_i \in S} \blacklozenge\mathfrak{K}_a^b\psi_i$.

Intuitively, the strategic Hartley formula in a strategic Hartley model expresses that the observer knows how to systematically identify the actual world in at most b steps. Consequently, we get the following reformulation of Theorem 2.

[5] For example, the "fully observational" model where a can test for any subset of S satisfies this requirement.

Theorem 7. *Let M be a strategic Hartley model with state space S. Then, for all $s \in S$, we have $M, s \models_{\mathsf{SLAcK}} \sigma(M, H(S))$.*

Proof. Straightforward, from Theorem 2 and the fact that the bisection strategy guarantees obtaining certainty about the actual world in a minimal number of steps. □

It would be interesting to consider a more general setting where different alternatives have different likelihood or different importance. Some preliminary research in that direction has been reported in [33]. We leave a proper treatment of the topic for future work.

7 Related Work

The inspiration for this work can be traced back to Herbert A. Simon who introduced the term of bounded rationality [32]. More recently, several approaches by e.g. Halpern et al. [20, 16], Rantala [29] and Konolige [25] have been proposed to deal with the logical omniscience problem, and enable reasoning about agents who are not necessarily perfect reasoners. Reasoning about agents' abilities under bounded resources has also become a major topic in the community of temporal and strategic logic, cf. the works by Alechina, Logan et al. [3, 4, 8, 5–7] and Bulling and Farwer [12, 11]. We review the most relevant approaches below.

In the syntactic approach to knowledge [15, 28, 25], the known sentences are explicitly listed for each possible world. This enables to capture an agent's limited ability to reason about knowledge. However, such models cannot capture changes in epistemic abilities of an agent when the available resources change. For each new amount of available resource, a new model must be built to reason about the new situation.

The awareness approach [16] has many similarities with the syntactic approach. It assumes that an agent cannot know something unless the agent is aware of it. In this approach, like the syntactic approach, for an agent to know that φ holds in a given state, φ must be in the set of "explicit awareness formulae" of the agent in that state. But unlike the syntactic approach, this is not enough. The agent must also know that φ holds by the definition of the classical modal epistemic logic, i.e., φ should be true in all the indistinguishable worlds.

The impossible worlds approach [29] avoids the assumption of perfect reasoning by limiting the immediate knowledge of tautologies and logical consequences. To this end some worlds are considered possible where those tautologies or logical consequences are false. Then, an agent knows φ if and only if, apart from knowing it in the classical way of modal epistemic logic, φ is also true in all the "impossible" worlds accessible via the agent's epistemic relation.

The algorithmic knowledge [20] distinguishes between implicit knowledge and explicit knowledge. Implicit knowledge is what an agent *may know* by the structure of the model, similar to the definition of knowledge in modal epistemic logic. But an agent who, by the model, has implicit knowledge about some fact, might not be able to conclude that fact due to different reasons. The algorithmic knowledge approach models this inability by assuming that, to check the truth value of φ in a given state, an agent uses an algorithm represented as a part of the state. An agent has explicit knowledge about φ if and only if its current algorithm enables the agent to conclude that φ holds.

The main difference between our work and the approaches of impossible worlds, algorithmic knowledge, and awareness is that, like for syntactic knowledge, neither of the approaches has an explicit concept of resources, and none of them allows for expressing properties of resource-bounded reasoning.

In [3, 4], a notion of delayed belief was introduced. The approach assumes that an agent is a perfect reasoner in an arbitrary decidable logic, but only derives the consequences of its beliefs after some delay. An advantage of delayed belief is that we can represent situations where an agent does not yet know a statement, but it can learn the statement by doing some action(s). Still, there is no notion of quantitative resource in this approach. So, e.g., we cannot reason about the effect of changing the available resources on the epistemic state of an agent. In [8], Timed Reasoning Logics (TRL) were introduced for describing resource-bounded reasoners that use time to derive consequences of their knowledge. TRL use a discrete model of time to capture the dynamics of agents' knowledge. However, it lacks the explicit notion of resources (although it is possible to reason about the time consumed when performing actions).

Another group of approaches was proposed in the agent logics community. RTL [11, 12] is a resource-bounded extension of the Computation Tree Logic CTL. In RTL, each transition between states can consume some resources and produce other resources. Thus, RTL has an explicit notion of resources, and enables reasoning about changes of abilities of agents due to changes in available resources. However, it has no semantic representation of knowledge. In order to reason about the evolution of knowledge, one would have to define new propositions that describe the knowledge of agents, and find out (by using other methods) what an agent knows in each state in order to determine the valuation of these propositions. The same applies to several other logics for resource-bounded agents, such as Coalition Logic for Resource Games [5], Resource Bounded Alternating-time Temporal Logic [6], Priced Resource Bounded ATL [27], and Resource Bounded Coalition Logic [7]. In these approaches the main focus is reasoning about how agents use actions to achieve their goals under resource bounds, with no specific machinery to capture information flow.

A framework where accumulation of knowledge plays an important role is offered by the Dynamic Epistemic Logic [34] and its extensions [1]. However, DEL takes into account neither limited observational resources nor imperfections of reasoning by real agents.

Reasoning about the outcome of accumulated observations has been also studied in belief revision and AI planning. Classical belief revision is syntactic in nature [2], though there are also formalizations based on possible worlds [18]. Still, both strands focus on inference by perfect reasoners using cost-free information-collecting actions. Within the belief revision literature, our work comes especially close to the agenda of [9]. However, in contrast to [9] where the stream of observations is determined externally by the environment, we assume that agents are free to choose the observations they are to make. This corresponds better to our notion of pro-active exploration of the environment, rather than passive sensing. Also, [9] uses no notion of resources, focusing on properties that may be learned with unbounded effort.

AI planning approaches that take into account epistemic actions are also mostly based on the syntactic approach to knowledge. For instance, the C-BURIDAN planner [14] represents each state by the set of propositions that hold in it. An information-collecting action does not change the propositions, but adds some labels to the state. These labels represent the observations that the agent has collected, and can later be used to block application of other actions. Although C-BURIDAN and similar planners capture information-collecting actions, they do not include a notion of resource. One can set preconditions for actions, and by this restrict the availability of an action in a state, but there is no way of representing the amount of resources needed to perform the action. Also, the main concern in AI planning is to find a sequence of actions that transforms the initial state of the world to an "objective" goal state, and not to a given *epistemic* state. Outcomes of observations can be used to find out the needed sequence of actions, but cannot define the goal itself.

Variants of AI planning which come much closer to our approach are belief planning [30] and dynamic epistemic planning [10]. They are both based on Kripke semantics, and focus on goals formulated in terms of epistemic states. Still, they lack the notion of quantitative resource, and do not address the impact of available resources on the outcome of plans (nor on the outcome of planning).

In summary, some existing frameworks allow for reasoning about information flow and epistemic change, albeit in a purely qualitative way. Other approaches allow us to model agents with limited epistemic abilities (for example because of bounded resources), but do not include the concept of resource, and do not facilitate reasoning about the relationship between agents' ability to gain knowledge, and changes in the resources. Yet another group includes the notion of resource and supports reasoning about resource-dependent abilities of agents, but lacks a semantic representation of knowledge and its dynamics. In this paper, we propose a logic-based approach that on the one hand relates epistemic abilities to resources, and on the other hand represents the process of refining knowledge in a semantically sensible way.

8 Conclusions

Intelligent agents usually choose their actions based on their knowledge about the environment. In order to gain or refine this knowledge, agents may perform information-collecting actions. Information-collecting actions like all other actions require resources. Therefore, the abilities of agents to improve their knowledge are limited by the resources available to them. In this work, we propose a modal approach to modeling, analyzing, and reasoning about agents that build their knowledge by using resource consuming information-collecting actions.

Our approach is based on several simplifying assumptions, which might not hold in realistic scenarios. Nevertheless, we believe the approach to be useful, especially with respect to simple scenarios. In more complex contexts, refinements of the framework could be needed.

Acknowledgements. Wojciech Jamroga acknowledges the support of the National Research Fund (FNR) Luxembourg under the project GALOT (INTER/DFG/12/06), as well as the support of the 7th Framework Programme of the European Union under the Marie Curie IEF project ReVINK (PIEF-GA-2012-626398). Masoud Tabatabaei acknowledges the support of the National Research Fund Luxembourg (FNR) under project GAIVS (AFR Code:5884506).

The authors thank the anonymous reviewers of CLIMA and JLC for their extremely useful comments.

References

1. T. Ågotnes and H. van Ditmarsch. Coalitions and announcements. In *Proceedings of AAMAS*, pages 673–680, 2008.
2. C.E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
3. N. Alechina and B. Logan. Logical omniscience and the cost of deliberation. In *Logic for Programming, Artificial Intelligence, and Reasoning*, volume 2250 of *LNCS*, pages 100–109. Springer, 2001.
4. N. Alechina and B. Logan. Ascribing beliefs to resource bounded agents. In *Proceedings of AAMAS*, pages 881–888. ACM, 2002.
5. N. Alechina, B. Logan, N.H. Nga, and A. Rakib. Verifying properties of coalitional ability under resource bounds. In *Proceedings of the Logics for Agents and Mobility (LAM)*, 2009.
6. N. Alechina, B. Logan, H.N. Nguyen, and A. Rakib. Resource-bounded alternating-time temporal logic. In *Proceedings of AAMAS*, pages 481–488, 2010.
7. N. Alechina, B. Logan, H.N. Nguyen, and A. Rakib. Logic for coalitions with bounded resources. *Journal of Logic and Computation*, 21(6):907–937, 2011.
8. N. Alechina, B. Logan, and M. Whitsey. A complete and decidable logic for resource-bounded agents. In *Proceedings of AAMAS*, pages 606–613. IEEE Computer Society, 2004.
9. A. Baltag, N. Gierasimczuk, and S. Smets. Belief revision as a truth-tracking process. In *Proceedings of TARK XIII*, pages 187–190. ACM, 2011.
10. T. Bolander and M. Birkegaard Andersen. Epistemic planning for single- and multi-agent systems. *Journal of Applied Non-Classical Logics*, 21(1):9–34, 2011.
11. N. Bulling and B. Farwer. Expressing properties of resource-bounded systems: The logics RTL* and RTL. In *Proceedings of CLIMA*, volume 6214 of *Lecture Notes in Computer Science*, pages 22–45, 2010.
12. N. Bulling and B. Farwer. On the (un-)decidability of model checking resource-bounded agents. In *Proceedings of ECAI*, volume 215 of *Frontiers in Artificial Intelligence and Applications*, pages 567–572. IOS Press, 2010.

13. N. Bulling and W. Jamroga. Comparing variants of strategic ability: How uncertainty and memory influence general properties of games. *Journal of Autonomous Agents and Multi-Agent Systems*, 28(3):474–518, 2014.
14. D. Draper, S. Hanks, and D. Weld. A probabilistic model of action for least-commitment planning with information gathering. In *Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence*, pages 178–186, 1994.
15. Rolf A. Eberle. A logic of believing, knowing, and inferring. *Synthese*, 26:356–382, 1974.
16. R. Fagin and J. Y. Halpern. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34(1):39–76, 1987.
17. R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge.* MIT Press, 1995.
18. N. Friedman and J.Y. Halpern. A knowledge-based framework for belief change, Part II: Revision and update. In *Proceedings of KR'94*, pages 190–200, 1994.
19. M. Grabisch, J-L. Marichal, R. Mesiar, and E. Pap. *Aggregation Functions.* Cambridge University Press, 2009.
20. J.Y. Halpern, Y. Moses, and M.Y. Vardi. Algorithmic knowledge. In *Proceedings of TARK'94*, pages 255–266, 1994.
21. R.V.L. Hartley. Transmission of information. *The Bell System Technical Journal*, 7(3):535–563, 1928.
22. W. Jamroga and M. Tabatabaei. Accumulative knowledge under bounded resources. In *Proceedings of CLIMA XIV*, volume 8143 of *LNCS*, pages 206–222. Springer, 2013.
23. W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2–3):185–219, 2004.
24. R.M. Karp. Reducibility among combinatorial problems. In R.E. Miller and J.W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. 1972.
25. Kurt Konolige. *A Deduction Model of Belief.* Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1986.
26. S. Kripke. Semantic analysis of modal logic. *Zeitschrift fur Mathematische Logik und Grund lagen der Mathematik*, 9:67–96, 1963.
27. D. Della Monica, M. Napoli, and M. Parente. On a logic for coalitional games with priced-resource agents. *Electronic Notes in Theoretical Computer Science*, 278:215–228, 2011.
28. R.C. Moore and Hendrix G. *Computational models of beliefs and the semantics of belief sentences.* SRI International, 1979.
29. V. Rantala. A modal logic for coalitional power in games. *Acta Philosophica Fennica*, 35:18–24, 1982.
30. S. Sardiña, G. De Giacomo, Y. Lespérance, and H.J. Levesque. On the limits of planning over belief states under strict uncertainty. In *Proceedings of KR*, pages 463–471, 2006.
31. P. Y. Schobbens. Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science*, 85(2):82–93, 2004.
32. H.A. Simon. Theories of bounded rationality. In C.B McGuire and R. Radner, editors, *In Decision and Organization*, pages 161–176. Amsterdam: North-Holland, 1972.
33. M. Tabatabaei. Accumulative knowledge under bounded resources: A logical account. Master thesis, University of Luxembourg, 2012.
34. H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic.* Springer, 2007.

35. Y. Wang and F. Dechesne. On expressive power and class invariance. *CoRR*, abs/0905.4332, 2009.