Dynamic Epistemic Algebra with Post-Conditions to Reason about Robot Navigation

Alexander Horn

Abstract

Dynamic epistemic algebra establishes Galois connections and quantales as a basis for reasoning about knowledge in multi-agent systems. To date, these algebraic-axiomatic methods have been restricted to a positive fragment of dynamic epistemic logic with communication events only. This paper proposes Boolean algebraic extensions to overcome these limitations by generalizing dynamic epistemic algebra to scenarios where events can change facts in form of post-conditions. As an application of post-conditions, we devise and solve a topological map-based robot navigation example for which current axiomatics are insufficient.

1 Introduction

Networks such as the Internet enable autonomous agents to communicate with each other. Such group communication can be understood as information flow. The formal logical analysis of information flow in multi-agent systems is possible with dynamic epistemic logic. Dynamic epistemic logic is sound and complete with respect to relational semantics [4]. These relational semantics trace back to Kripke semantics [14] which shape modal logic (e.g. [5]). Similar to modal logic, traditional dynamic epistemic logic requires fixed valuations where facts never change [24]. To account for factual changes, recent research extended dynamic epistemic logic with post-conditions [25], [11], [21], [23]. The resulting logical expressiveness has found application in the study of market simulations, such as the trade on commodities [22], and the reasoning about information across covert channels [26]. However, relational semantical proofs can be tedious to formalize as agents interact frequently.

In contrast, algebraic semantics of dynamic epistemic logic stylize parts of the traditional relational structures by characterizing the dynamic nature of knowledge in terms of Galois connections and quantales [20], [3]. These mathematical structures tend to make proofs about information flow in multi-agent systems more perspicuous. However, only multi-agent systems in which facts never change are supported [18], [8]. In response, subsequent changes to the algebra aimed at increased flexibility with a converse dynamic modality [17]. However, this converse modality excludes event composition and it formalizes only learning of agents without initial knowledge [13].

To overcome these limitations, this paper proposes a novel algebraic abstraction inspired by the relational semantical extensions of dynamic epistemic logic with post-conditions (e.g. [25], [21], [23]). The novelty is a dynamic epistemic algebra with Boolean algebraic structures and post-conditions for events in a quantale. The former applies to the negation of beliefs (i.e. disbeliefs). The latter enables reasoning about knowledge when events can change facts. As an application of factual changes, I encode the movements of a robot as events with post-conditions. This encoding aligns with topological maps studied in the robotics literature (e.g. [15]). Moreover, it broadens the application of dynamic epistemic logic by yielding a sensor-based robot navigation strategy which demonstrates the formalization and analysis of scenarios in which agents explore their surrounding including explorations from a known initial location.

2 Robot Navigation with Topological Maps

In order to explain the problem, we start with a directed graph $G = (V, E)$. Graph G is called a **topological map** if each vertex corresponds to a distinct **location** in an environment such as an office and edges denote **movements** between locations. Since topological maps are abstractions of the real world, locations could be understood as "distinctive places" which correspond to reference points in the environment [15].

Suppose there is an autonomous robot who knows the topological map of its surrounding. The objective of the robot is to determine its exact location. This objective could require the robot to move about in order to discover more information. During this exploration, the assumption is that the robot can detect all available movements at its current location. When the robot performs one of these movements, we assume that it always reaches the next location according to the topological map.

As an example of robot navigation, consider the topological map in Figure 1 devised by Phillips [18]. Let the robot's initial location be p_2. Since locations p_1 and p_2 are identical in terms of available movements (a and b), the robot considers the possibility of being in either location. Similarly, if it were to start at location p_3, it would not know if it would be in p_3 or p_4. However, as soon as it moves from its unknown initial location p_2 via movement a, it gains knowledge of its position! That is, the robot learns that it has arrived at location p_3. It

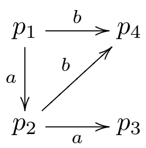

Figure 1: Topological map in which the movement from location p_2 to p_3 reduces the uncertainty of the agent about its position [18], [8].

rules out the possibility of having reached location p_4 because an a movement does not lead to it. Moreover, even though an a movement leads to location p_2, the robot eliminates this possibility as well because it can detect neither a nor b movements once it arrives at location p_3.

Noteworthy, these conclusions cannot be drawn with the original dynamic epistemic algebra [18]. Therefore, our goal is to algebraically formalize and precisely analyze how movements change the robot's uncertainty about locations.

Unlike previous algebraic approaches [18], [8], [7], [17], this paper aims at robot navigation with a more general solution where a robot has a sensor. This sensor-based solution can be described by the following algorithm [13]:

1. The sensor informs the robot of possible movement options (if any).
2. After choosing one of these movements, the robot relocates accordingly.
3. In the meantime, the robot uses its knowledge of the topological map to determine the set of possible destinations.
4. Once the robot arrives at one of these destinations, the sensor probes the environment and broadcasts the set of available movements.
5. By listening to this broadcast, the robot combines the information from the two previous steps to update its uncertainty about its current location.

In fact, this solution appeals to other problems in which agents need to observe changes in their surrounding to learn information. To emphasize this point, we simply represent the fourth step as an honest public announcement by the sensor [13]. In an epistemic setting, the challenge is the relocation of the robot because it causes factual changes.

Even though recent axiomatic extensions of dynamic epistemic logic [23] apply to robot navigation scenarios [13], such Hilbert-style deductions are not necessarily constructive in the sense that they explain the intuition behind machine learning. In Figure 1, for example, the proof that the robot learns its destination when moving from location p_2 to p_3 could consist of tautologies such as $\vdash [(p_3 \vee p_4)!]\Box_\alpha p_3 \leftrightarrow (p_3 \vee p_4) \rightarrow \Box_\alpha(p_4 \rightarrow p_3)$. Intuitively, after the honest public announcement by the sensor, the agent believes that it is in location p_3 if and only if it is neither in p_3 nor in p_4, or it believes that its presence at location p_4 implies that it is at location p_3 [13]. To avoid such indirection, the next section develops a dynamic epistemic algebra with post-conditions which conceptualizes learning as the change of an agent's uncertainty about locations.

3 Dynamic Epistemic Algebra with Post-Conditions

The requirement is to analyze both information flow and factual changes in multi-agent systems. The former has been studied in terms of Galois connections and quantales [2], [20], [3]. The latter is the main contribution of this section. Both features require some level of familiarity with lattice theory (e.g. [9]).

To capture factual changes in an agent's surrounding, we start by defining a countable set of facts as a basis for real world information. Note that the separate treatment of logical propositions deviates from earlier intuitionistic algebraic approaches [20], [3].

Definition 3.1. Define the tuple $\mathcal{P} = \langle P; \vee, \wedge, \neg, \bot, \top \rangle$ to be a Boolean algebra where P is the set of atoms which we call **facts**.

Facts are atomic sentences which are either true or false. By definition (3.1), the Boolean algebra of facts satisfies classical propositional logic properties such as De Morgan's laws. For example, $\neg(p \wedge q) = \neg p \vee \neg q$ for all $p, q \in \mathcal{P}$.

Definition 3.2. A **unital quantale** $\langle E, \bigvee, \bullet, 1 \rangle$ is a sup-lattice E, equipped with the identity element 1 and a monoid $- \bullet -\colon E \to E$ such that

$$\left(\bigvee_j e_j\right) \bullet e = \bigvee_j (e_j \bullet e) \quad \text{(jql)} \qquad e \bullet \left(\bigvee_i e_i\right) = \bigvee_i (e \bullet e_i) \quad \text{(jqr)}$$

In other words, rule (jql) together with rule (jqr) define the monoid to preserve arbitrary joins in both arguments. Intuitively, elements in the quantale are interpreted as events ordered by determinism where $e \vee f$ corresponds to the non-deterministic choice of events [3]. Multiplication in the quantale can be interpreted as sequential event composition [3]. For example, $e \bullet f$ denotes the sequence of events where e happens before event f. Thus, multiplication in the quantale is defined to be noncommutative. Furthermore, quantales are resource-sensitive [20]. For example, $e \not\leq e \bullet e$ and $e \bullet e \not\leq e$ means that the repetition of the same event can be fundamentally different from its single occurrence as the repetitive group interrogation in the Muddy Children Puzzle illustrates (e.g. [10]). Since the quantale is unital, $e \bullet 1 = 1 \bullet e = e$ for all $e \in E$, where the identity element is the **unit event** which does nothing. Apart from the unit event, all other events in the quantale have some sort of effect. This notion of 'change' leads to the definition of a module [20], [3].

Definition 3.3. Let E be a quantale. A **module** over E is a sup-lattice M with a function $- \odot -\colon M \times E \to M$ such that, for all $m \in M$ and $e, f \in E$,

$$m \odot \left(\bigvee_i e_i\right) = \bigvee_i (m \odot e_i) \quad \text{(je)} \qquad \left(\bigvee_j m_j\right) \odot e = \bigvee_j (m_j \odot e) \quad \text{(jm)}$$

$$(m \odot e) \odot f = m \odot (e \bullet f) \quad \text{(co)} \qquad m \odot 1 = m \quad \text{(unit)}$$

Elements in the module embody 'possible worlds' [10], [2]. For this reason, these elements are called **states**. The dynamic nature of multi-agent systems is captured by the module operation $- \odot -$ which reflects changes of beliefs [3]. Since rule (je) and (jm) define this operation to be join-preserving in both arguments, it is monotonic (e.g. [9]) if either argument is fixed. The remaining equalities axiomatize event composition and the unit event [20], [3], [13].

Noteworthy, the pair of quantales and a module without modal operators has been previously studied with denotational models of concurrent processes [1]. Its connection with dynamic epistemic logic has been realized in [20] with a focus on intuitionistic logic where traditional Boolean laws cannot be generally applied. We start to address this limitation by refining the definition of a module to a complete Boolean lattice. Recall that a complete Boolean lattice is the same as a complemented distributive lattice where every subset has a supremum.

Definition 3.4. A **Boolean module** is a module whose lattice is a complete Boolean lattice. In general, it is an infinite lattice [13, pp. 62f.].

Similar to the relational semantics [25], [11], [21], [23], events in the quantale have pre-conditions and post-conditions [13, pp. 51–65]. The former is intrinsic to the update operation: if the pre-condition of an event e is unsatisfied in a state m, then $m \odot e = \bot$. The latter, however, requires the association of facts with elements in the Boolean module [13]. This new requirement leads to the definition of the assignment operator.

Definition 3.5. Let $\mathcal{P}$ be a Boolean algebra of facts. Let M be a Boolean module. Define the **assignment** to be a Boolean homomorphism $-^*\colon \mathcal{P} \to M$.

Since the new assignment operator maps logical propositions to states, it shares much in common with valuations such as in modal logic, for example. The difference is that the algebraic association of states with facts is accomplished by the partial order relation on the complete Boolean lattice of the module. Similar to valuations, for more complex expressions such as $(p \wedge \neg q)^*$, we can conclude that it is equal to $p^* \wedge \neg q^*$ because the assignment operator is defined to be a Boolean homomorphism. This Boolean homomorphism builds the basis for the algebraic specification of post-conditions which we define next.

Definition 3.6. Let $\mathcal{P}$ be a Boolean algebra of facts. Let E be a quantale and M be a Boolean module over E. Let $p \in \mathcal{P}$, $m \in M$ and $e \in E$. We say m **satisfies** p if $m \leq p^*$. We call proposition p the **post-condition** of event e in state m if $m \odot e$ satisfies p.

Intuitively, $m \leq p^*$ means that in state m proposition p is true. We refer to such an inequality as **satisfaction relation**. By definition (3.5), satisfaction relations such as $m \leq (p \wedge \neg q)^*$ and $m \leq p^* \wedge \neg q^*$ are equivalent. As special cases of satisfaction relations, post-condition specifications are inequalities of the form $m \odot e \leq p^*$. Intuitively, such inequalities assert that proposition p is true after the update of a state m with an event e. Noteworthy, these updates can capture factual changes. Of course, it could also be the case that state m does not even satisfy the pre-condition of event e in which case the inequality $m \odot e = \bot \leq p^*$ is vacuously true. Such vacuous conditions are in accordance with the meta-logical treatment of pre-conditions in the relational approach where the outcome of an event is undefined if its pre-condition has not been satisfied [24]. Thus, post-condition specifications tend to be much more concise in both models without non-trivial cases. The introduction of post-conditions, however, warrants the formal differentiation between epistemic events and the new class of non-epistemic events.

Definition 3.7. Let $\mathcal{P}$ be a Boolean algebra of facts. Let E be a quantale and M be a Boolean module over E. Let $e \in E$. Then, e is called an **epistemic event** if $p^* \odot e \leq p^*$ for all $p \in \mathcal{P}$. Otherwise, we call the event **non-epistemic**.

In other words, if there exists $p \in \mathcal{P}$ such that $p^* \odot e \not\leq p^*$, then e is called a non-epistemic event. Informally, non-epistemic events can change the propositional assignment of the state being updated, whereas epistemic events preserve all propositional assignments. The latter corresponds to 'atomic permanence' in proof systems for dynamic epistemic logic without factual changes [4]. In fact, the definition of an epistemic event is also similar to the definition of **stable facts** in the intuitionistic approach [20], [3]. However, our approach achieves a clearer separation between stable facts and epistemic propositions by virtue of the additional assignment operator introduced as part of definition (3.5).

Definition 3.8. A **Boolean system** is a quadruple consisting of a Boolean algebra with set of facts P together with a quantale E, a Boolean module M over E and an assignment. We write $\langle P, M, E, \odot, ^* \rangle$ for a Boolean system.

Finally, we augment the Boolean system with the lax sup-endomorphisms defined for static intuitionistic approaches to dynamic epistemic logic [2], [20], [3].

Definition 3.9. Let $\mathcal{A}$ be a finite set of agents and $\alpha \in \mathcal{A}$. Let $u^E_\alpha \colon E \to E$ and $u^M_\alpha \colon M \to M$ be sup-endomorphisms. Define the pair $\left(u^E_\alpha, u^M_\alpha\right)$ to be a **lax sup-endomorphism of a Boolean system** $\langle P, M, E, \odot, ^* \rangle$ such that, for all $m \in M$ and $e, f \in E$,

$$u^M_\alpha(m \odot e) \leq u^M_\alpha(m) \odot u^E_\alpha(e) \tag{uui}$$

$$u^E_\alpha(e \bullet f) \leq u^E(e) \bullet u^E_\alpha(f) \tag{uci}$$

$$1 \leq u^E_\alpha(1) \tag{usi}$$

Semantically, the function u^M_α encodes the uncertainty of agent α about states [3]. Similarly, u^E_α is an agent's α uncertainty about events [3]. Examples for both sup-endomorphisms and their interpretations appear in [20], [3], [13].

Clearly, uncertainties could change as agents interact with their environment. The epistemic effects of such interactions are modeled by the uncertainty update inequality (uci) which captures the notion of learning [20], [3], [13]. Secondly, (uci) relates the uncertainty about an event composition to the individual sequential events [20], [3]. Lastly, (usi) asserts that when no event occurs an agent must consider the possibility that, in fact, nothing has happened [20].

Notice that the assignment operator integrates the Boolean algebra $\mathcal{P}$ with these sup-endomorphisms. However, we treat expressions such as $u^M_\alpha(p^*)$ as proof-theoretic constants which cannot be simplified for propositions $p \in \mathcal{P}$. More accurately, the general assumption is that the exact value of $u^M_\alpha(p^*)$ is unknown. The rationale is similar to the argument that we would not generally compute $u_\alpha^M(\top)$ because it requires an explicit construction of the entire Boolean module. However, it turns out that we can gain much ground without knowing the exact value of $u_\alpha^M(p^*)$. For simplicity, we *could* assume that agents have total uncertainty about states which satisfy the proposition p. Of course, the claim is *not* that $u_\alpha^M(p^*) = \top$. A simple counterexample is $u_\alpha^M(\bot^*) = u_\alpha^M(\bot) = \bot$ by definition (3.5) and remark (3.10).

Remark 3.10. *Recall that $\bot = \bigvee \emptyset$. By join-preservation of $- \odot -$, we conclude $m \odot \bot = \bot \odot e = \bot$ for all $m \in M$ and $e \in E$. Similarly, $u_\alpha^M(\bot) = \bot$.*

Most notably, since u_α^M preserves arbitrary joins, we conclude that it has a Galois connection $u_\alpha^M(-) \dashv \Box_\alpha -$ (e.g. [9]). In fact, the right adjoint $\Box_\alpha -$ encodes the belief modality for agent α [20], [3]. Let $m \in M$ be a state, $e \in E$ be an event and $p \in \mathcal{P}$ be a proposition. By definition of Galois connection, $u_\alpha^M(m) \leq p^*$ if and only if $m \leq \Box_\alpha p^*$. In other words, if the uncertainty of the agent about state m includes only states which satisfy proposition p, then $\Box_\alpha p^*$ can be read as "agent α believes proposition p is true". Likewise, since $- \odot e$ preserves arbitrary joins, it has an adjoint denoted by $[e]-$. Recall that if the inequality $m \odot e \leq p^*$ is satisfied, then the post-condition of event e in state m is proposition p by definition (3.6). By adjunction $- \odot e \dashv [e]-$, this inequality holds if and only if $m \leq [e]p^*$. Intuitively, the expression $[e]p^*$ means that "after event e, proposition p holds". Of course, unless the pre-condition is actually satisfied, the inequality $m \leq [e]p^*$ is vacuously true because $m \odot e = \bot \leq p^*$.

4 Robot Movements as Post-Conditions

This section exemplifies the algebraic treatment of post-conditions. For this purpose, we devise a robot navigation example based on the topological map shown in Figure 2. Notice that the robot can uniquely identify location p_1 because no other location has c movements. Similarly, location p_3 is unique. Therefore, the robot knows its destination when it moves from either of these known initial locations to any other. Recall that such valid conclusions cannot be generally drawn from algebraic semantics which require the converse dynamic modality [13]. In contrast, the extended dynamic epistemic algebra can be used to prove these and other scenarios [13].

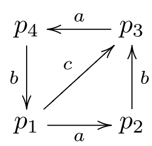

Figure 2: Robot navigation example which demonstrates the algebraic treatment of post-conditions by proving that the robot knows its destination when it moves from a known initial location such as p_1 or p_3.

Before we can illustrate this point, we must agree on a precise problem specification. For this purpose, we use the extended dynamic epistemic algebra defined in the previous section. Let $\mathcal{P}$ be a Boolean algebra where each fact in the set P represents a location of the robot according to the topological map. For example, the fact $p_1 \in P$ means that the robot is in location p_1. Since the robot can never be at two locations at the same time, we define $p_i \wedge p_j = \bot$ for all $p_i, p_j \in P$ such that $p_i \neq p_j$. Let M be a Boolean module such that for all facts $p_i \in P$ there exists an atom $m_i \in M$ where $m_i \leq p_i^*$. By definition (3.6), $m_i \leq p_i^*$ means that state m_i satisfies the fact that the robot is at location p_i. Since $m_i \neq \bot$, $p_i^* \neq \bot$ for all facts $p_i \in P$. Let movements between locations be events in the quantale E. Since one of the assumptions is that the robot has absolute certainty about its movements, $u_\alpha^E(e) := e$ for all $e \in E$. Define post-conditions by $p_i^* \odot e \leq p_j^*$, for all $p_i, p_j \in P$ and $e \in E$, such that p_j is the fact for the location which can be reached from location p_i with movement e. If such a movement is impossible, then $p_i^* \odot e = \bot$. Prior to any such movement, the uncertainty about initial locations is solely determined by available movements. Formally, we define $m_i \in M$ to be an **initial state** if m_i is an atom and $m_i \neq m_j \odot e$ for all $m_j \in M$ and $e \in E$ such that $e \neq 1$. The set of all such initial states is denoted by $I(M)$. Finally, define the uncertainty about an initial state m_i by $u_\alpha^M(m_i) := \bigvee\{m_j \in I(M) \mid events(m_i) = events(m_j)\}$ where $events(m_i)$ is the set of available movements at location p_i.

To develop an intuition for these definitions, reconsider the topological map in Figure 2. Since the robot can move from location p_1 to p_2 via movement a, the inequality $p_1^* \odot a \leq p_2^*$ is true. Furthermore, $u_\alpha^E(a) = a$ because the assumption is that the robot has absolute certainty about its movements. Finally, note that the robot has no uncertainty about the initial location p_1 because movement c is unique to it. Therefore, $u_\alpha^M(m_1) = m_1$ where $m_1 \in I(M)$ and $m_1 \leq p_1^*$.

Next, we use these definitions to prove that the robot has no uncertainty about its destination when moving from a known initial location such as p_1. Formally, $m_1 \leq [a]\Box_\alpha p_2^*$ where $m_1 \in I(M)$ and $m_1 \leq p_1^*$. The left side of the inequality corresponds to the robot's known initial location p_1. On the right side, the dynamic modality describes the robot's movement from location p_1 to p_2. After this movement, $\Box_\alpha p_2^*$ means that the agent believes that it is in location p_2. Before establishing this dynamic knowledge property, the next lemma proves that the robot, in fact, reaches location p_2 after starting from p_1.

Lemma 4.1. *Let $m_1 \in M$ and $p_1, p_2 \in P$. Assume $m_1 \leq p_1^*$ and $p_1^* \odot a \leq p_2^*$. Then, $m_1 \odot a \leq p_2^*$.*

Proof. By assumption, $m_1 \leq p_1^*$ and $p_1^* \odot a \leq p_2^*$. Since $- \odot a$ is monotonic, $m_1 \odot a \leq p_1^* \odot a$. By transitivity, $m_1 \odot a \leq p_2^*$. ■

Lemma (4.1) is a post-condition specification according to definition (3.6). It is independent from the robot's knowledge of the topological map. Since the next proposition, however, proves that the robot knows its destination, it also requires the additional assumption about the agent's uncertainty about its initial location p_1, i.e. $u_\alpha^M(m_1) = m_1$ where $m_1 \in I(M)$ and $m_1 \leq p_1^*$.

Proposition 4.2. *Assume that agent α knows the topological map in Figure 2. Let $m_1 \in I(M)$ where $m_1 \leq p_1^*$ and $p_1, p_2 \in P$. Then, $m_1 \leq [a]\Box_\alpha p_2^*$.*

Proof. By the adjunction $- \odot a \dashv [a]-$, the claim is equivalent to $m_1 \odot a \leq \Box_\alpha p_2^*$. By the adjunction $u_\alpha^M \dashv \Box_\alpha$, this inequality is equivalent to $u_\alpha^M(m_1 \odot a) \leq p_2^*$. By (uui), it suffices to show $u_\alpha^M(m_1) \odot u_\alpha^E(a) \leq p_2^*$. By the topological map definition, $m_1 \odot a \leq p_2^*$. By lemma (4.1), we conclude that the claim is true. ∎

The dynamic epistemic algebra with post-conditions also supports more sophisticated situations when the agent does not know its initial location [13]. Furthermore, the algebra integrates with communication events such as honest public announcements [13]. The combination of these features are sufficiently expressive, for example, for the sensor-based robot navigation strategy (p. 3) and the topological map in Figure 1 [13]:

Proposition 4.3. *Assume agent α knows the topological map in Figure 1 (p. 2). Let $m_2 \in I(M)$ such that $m_2 \leq p_2^*$. Then, $m_2 \leq [a][(p_3^* \vee p_4^*)!]\Box_\alpha p_3^*$ where $(p_3^* \vee p_4^*)!$ denotes the honest public announcement by the sensor.*

Proof. The proof appears in Appendix A. ∎

5 Disbelief

Since previous algebraic semantics were built for intuitionistic logic [20], it has been difficult to express the negation of beliefs such as $\neg\Box_\alpha\phi$ because it required another Galois connection [20] without direct proof-theoretic applications [13]. In this section, we identify another Boolean algebraic characterization of the negation of agent's beliefs. Before we do this, the next definition suggests a more natural reading of an agent's negated beliefs.

Definition 5.1. Let M be a Boolean module. Let $\phi \in M$. Define **disbelief** of an agent α about ϕ by $\neg\Box_\alpha\phi$.

The next proposition contributes to the algebraic reasoning about disbeliefs.

Proposition 5.2. *Let M be a Boolean module. Let $m, \phi \in M$ be states. Then, $m \leq \neg\Box_\alpha\phi$ if and only if, for all $x \in M$, $u_\alpha^M(x) \leq \phi$ implies $m \wedge x = \bot$.*

Proof. The proof appears in Appendix B. ∎

Proposition (5.2) relates disbeliefs to the greatest lower bound (infimum) of elements in the Boolean module. The next theorem proves a satisfying condition of disbelief for the special case in which one of these elements is an atom.

Theorem 5.3. *Let M be a Boolean module. Let $m, \phi \in M$. If m is an atom and $u_\alpha^M(m) \not\leq \phi$, then $m \leq \neg\Box_\alpha\phi$.*

Proof. The proof appears in Appendix B. ∎

The next example demonstrates a proof about disbelief as a consequence of theorem (5.3).

Example 5.4. Consider a robot who knows the topological map in Figure 2. Let $m_2, m_4 \in I(M)$ be initial states and $p_2, p_4 \in P$ be facts. Assume $m_2 \leq p_2^*$ and $m_4 \leq p_4^*$. To show that the robot cannot clearly identify its initial location p_2, we must prove $m_2 \leq \neg\Box_\alpha p_2^*$. By definition of $I(M)$, both m_2 and m_4 are atoms. By theorem (5.3), it suffices to show that $u_\alpha^M(m_2) \not\leq p_2^*$. Since locations p_2 and p_4 have both only a movements, the agent considers the possibility of being in either initial location. Formally, $u_\alpha^M(m_2) = m_2 \vee m_4$. Since $m_4 \not\leq p_2^*$, we conclude that $u_\alpha^M(m_2) \not\leq p_2^*$. Therefore, $m_2 \leq \neg\Box_\alpha p_2^*$.

6 Conclusions

The main contribution of this paper is an algebraic framework for the formalization and analysis of agents which can observe and reason about changes in their environment. This achievement has been possible with the extension of dynamic epistemic algebra with post-conditions. For this purpose, the extended dynamic epistemic algebra integrates Galois connections, quantales and modules [20], [3] with an additional Boolean algebraic assignment operator. The significance of the assignment operator is threefold. Firstly, it achieves a clearer separation between stable facts and epistemic propositions. Secondly, it enables the algebraic abstraction of factual changes in form of events with post-conditions. Lastly, it unifies the algebraic treatment of both factual and epistemic changes under previously developed lax sup-endomorphisms. This unification eliminates the need for converse dynamic modalities (e.g. [18], [8], [7], [17]). Unfortunately, this elimination also reduces the expressiveness about temporal properties. In turn, however, the simplicity of dynamic epistemic algebra has been restored even for events which change facts. The algebraic specification of these events was shown to serve as a proof-theoretic basis for the encoding of robot movements.

7 Further Research

Future research could aim at the extension with fuzzy Galois connections [6]. Another research direction is the integration with Kleene algebra to simplify the algebraic characterization of common knowledge in resemblance to [16]. The resulting algebraic semantics could be a candidate for automated theorem provers (e.g. [12]). Such an implementation strategy could be compared and contrasted to epistemic model checking.

8 Acknowledgements

I am grateful to M. Sadrzadeh whose research questions sparked this paper. Her generous feedback improved the presentation of these results.

References

[1] S. Abramsky and S. Vickers. Quantales, observational logic and process semantics. *Mathematical Structures in Computer Science*, 3(02):161–227, 1993.

[2] A. Baltag, B. Coecke, and M. Sadrzadeh. Algebra and sequent calculus for epistemic actions. *Electronic Notes in Theoretical Computer Science*, 126:27–52, 2005. Proceedings of the 2nd International Workshop on Logic and Communication in Multi-Agent Systems (2004).

[3] A. Baltag, B. Coecke, and M. Sadrzadeh. Epistemic actions as resources. *Journal of Logic and Computation*, 17:555–585, 2007.

[4] A. Baltag, L. S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. In *TARK '98: Proceedings of the 7th conference on Theoretical aspects of rationality and knowledge*, pages 43–56, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc.

[5] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2001.

[6] R. Bělohlávek. Fuzzy galois connections. *Math. Log. Q.*, 45:497–504, 1999.

[7] Philips C., Precup D. Panangaden P., and Sadrzadeh M. Reasoning about factual games using information updates. In *9th Conference on Logic and the Foundations of Game and Decision Theory*, Toulouse, July 2010.

[8] Philips C., Precup D. Panangaden P., and Sadrzadeh M. An algebraic approach to dynamic epistemic logic. In *23rd International Workshop on Description Logics*, Waterloo, May 2010.

[9] B. A. Davey and H. A. Priestley. *Introduction to lattices and order*. Cambridge University Press, Cambridge, second edition, 2002.

[10] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.

[11] A. Herzig and T. De Lima. Epistemic actions and ontic actions: A unified logical framework. In *IBERAMIA-SBIA 2006, LNAI 4140*, pages 409–418. Springer, 2006.

[12] P. Höfner and G. Struth. Automated reasoning in kleene algebra. In *CADE-21: Proceedings of the 21st international conference on Automated Deduction*, pages 279–294, Berlin, Heidelberg, 2007. Springer-Verlag.

[13] A. Horn. Reasoning about learning in robot navigation via algebraic dynamic epistemic logic. Master's thesis, University of Oxford, 2010.

[14] S. A. Kripke. A semantical analysis of modal logic I: Normal modal propositional calculi. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 9:67–96, 1963.

[15] B. Kuipers and Yung-Tai Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Journal of Robotics and Autonomous Systems*, 8:47–63, 1991.

[16] B. Möller. Knowledge and games in modal semirings. In *RelMiCS'08/AKA'08: Proceedings of the 10th international conference on Relational and kleene algebra methods in computer science, and 5th international conference on Applications of kleene algebra*, pages 320–336, Berlin, Heidelberg, 2008. Springer-Verlag.

[17] P. Panangaden and M. Sadrzadeh. Learning in a changing world via algebraic modal logic. Technical Report RR-10-10, OUCL, June 2010.

[18] C. Phillips. An algebraic approach to dynamic epistemic logic. Master's thesis, McGill University, 2009.

[19] H. Rasiowa and R. Sikorski. *The mathematics of metamathematics.* Pańistwowe Wydawn Naukowe, 1968.

[20] M. Sadrzadeh. *Actions and Resources in Epistemic Logic.* PhD thesis, Université du Québec à Montréal, 2006.

[21] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Inf. Comput.*, 204(11):1620–1662, 2006.

[22] H. van Ditmarsch. The logic of pit. *Synthese*, 149(2), 2006.

[23] H. van Ditmarsch and B. Kooi. Semantic results for ontic and epistemic change. In G Bonanno, W van der Hoek, and M Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory*, Texts in Logic and Games, pages 87–117. Amsterdam University Press, 2008.

[24] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Springer, 2007.

[25] H. van Ditmarsch, W. van der Hoek, and B. P. Kooi. Dynamic epistemic logic with assignment. In *AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 141–148, New York, NY, USA, 2005. ACM.

[26] J. van Ditmarsch, H. van Eijck and W. Wu. One hundred prisoners and a lightbulb - logic and computation. In *F. Lin and U. Sattler (editors).* AAAI Press, 2010.

Appendix A: Robot Navigation

This appendix proves proposition (4.3) which exemplifies how a robot learns its position as it moves from an *unknown* initial location.

Lemma 8.1. *Let $m_1, m_2 \in I(M)$ such that $m_1 \leq p_1^*$ and $m_2 \leq p_2^*$. Assume $p_1^* \odot a \leq p_2^*$ and $p_2^* \odot a \leq p_3^*$. Then, $m_1 \odot a \leq p_2^*$ and $m_2 \odot a \leq p_3^*$*

Proof. By assumption, monotonicity of $- \odot a$ and transitivity. ■

Lemma (8.1) is a post-condition specification according to definition (3.6). The first inequality specifies the post-condition of event a in state m_1 to be the fact that the robot reaches location p_2. Similarly, the post-condition of $m_2 \odot a$ is the fact for location p_3.

The next two lemmas aim at the formalization of the sensor's honest public announcement which is essential for the robot to learn its position.

Lemma 8.2. *Let $\mathcal{P}$ be a Boolean algebra of facts P. Let $p_i, p_j \in P$ be distinct facts (i.e. $i \neq j$) such that $p_i \wedge p_j = \bot$. If $p_i^* \neq \bot$, then $p_i^* \not\leq p_j^*$.*

Proof. Assume $p_i \wedge p_j = \bot$ and $p_i^* \neq \bot$. Assume $p_i^* \leq p_j^*$. Then, $p_i^* \wedge p_j^* = p_i^*$. Since $-^*$ is a Boolean homomorphism, $p_i^* \wedge p_j^* = (p_i \wedge p_j)^* = \bot^* = \bot = p_i^*$ contradicting the assumption. Hence, $p_i^* \not\leq p_j^*$. ■

Lemma 8.3. *Let $\mathcal{P}$ be a Boolean algebra of facts P. Let $p_2, p_3, p_4 \in P$ be distinct facts. Then, $p_2^* \odot (p_3^* \vee p_4^*)! \leq \bot$ where $(p_3^* \vee p_4^*)!$ denotes the honest public announcement of the proposition $p_3 \vee p_4$.*

Proof. By lemma (8.2), $p_2^* \not\leq p_3^*$ and $p_2^* \not\leq p_4^*$. Therefore, $p_2^* \not\leq p_3^* \vee p_4^*$. Hence, p_2^* does not satisfy the pre-condition of the honest public announcement. ■

The next lemma eliminates the possibility of the robot reaching location p_3 or p_4 from the initial location p_1 via movement a due to the truthful information sharing by the sensor.

Lemma 8.4. *Let $m_1 \in I(M)$ such that $m_1 \leq p_1^*$. Then, $m_1 \odot a \odot (p_3^* \vee p_4^*)! = \bot$.*

Proof.

$$
\begin{aligned}
m_1 \odot a &\leq p_2^* && \{\text{lemma (8.1)}\} \\
m_1 \odot a \odot (p_3^* \vee p_4^*)! &\leq p_2^* \odot (p_3^* \vee p_4^*)! && \{\text{monotonicity of } - \odot -\} \\
p_2^* \odot (p_3^* \vee p_4^*)! &\leq \bot && \{\text{lemma (8.3)}\} \\
m_1 \odot a \odot (p_3^* \vee p_4^*)! &\leq \bot && \{\text{transitivity}\}
\end{aligned}
$$

■

Lemma 8.5. *Let $m_2 \in I(M)$ such that $m_2 \leq p_2^*$. Then, $m_2 \odot a \odot (p_3^* \vee p_4^*)! \leq p_3^*$.*

Proof. Similar to lemma (8.4) except that $m_2 \odot a$ satisfies the pre-condition of the honest public announcement $(p_3^* \vee p_4^*)!$ because $m_2 \odot a \leq p_3^*$. ■

Finally, we prove the claim that the robot learns its destination after moving from location p_2 to p_3 via movement a and then listening to the honest public announcement by the sensor. Since this proof requires the assumption that the robot knows the topological map in Figure 1, we can gain additional clarity by stating the robot's uncertainty about its initial location p_2.

Lemma 8.6. *Assume that agent α knows the topological map in Figure 1 (p. 2). Let $m_1, m_2 \in I(M)$ such that $m_1 \leq p_1^*$ and $m_2 \leq p_2^*$. Then, $u_\alpha^M(m_2) = m_1 \vee m_2$.*

Proof. Since there are the same available movements at location p_1 and p_2, $events(m_1) = events(m_2)$. By definition of the uncertainty about initial states, we conclude that $u_\alpha^M(m_2) = m_1 \vee m_2$. ∎

Proposition 8.7. *Assume agent α knows the topological map in Figure 1 (p. 2). Let $m_2 \in I(M)$ such that $m_2 \leq p_2^*$. Then, $m_2 \leq [a][(p_3^* \vee p_4^*)!]\Box_\alpha p_3^*$.*

Proof.

$$
\begin{aligned}
m_2 &\leq [a][(p_3^* \vee p_4^*)!]\Box_\alpha p_3^* && \{\text{claim}\}\\
m_2 \odot a &\leq [(p_3^* \vee p_4^*)!]\Box_\alpha p_3^* && \{\text{gal}\}\\
(m_2 \odot a) \odot (p_3^* \vee p_4^*)! &\leq \Box_\alpha p_3^* && \{\text{gal}\}\\
u_\alpha^M((m_2 \odot a) \odot (p_3^* \vee p_4^*)!) &\leq p_3^* && \{\text{gal}\}\\
u_\alpha^M(m_2 \odot a) \odot u_\alpha^E((p_3^* \vee p_4^*)!) &\leq p_3^* && \{\text{uui}\}\\
\big(u_\alpha^M(m_2) \odot u_\alpha^E(a)\big) \odot u_\alpha^E((p_3^* \vee p_4^*)!) &\leq p_3^* && \{\text{uui}\}\\
((m_1 \vee m_2) \odot a) \odot (p_3^* \vee p_4^*)! &\leq p_3^* && \{\text{lemma (8.6)}\}\\
(m_1 \odot a \vee m_2 \odot a) \odot (p_3^* \vee p_4^*)! &\leq p_3^* && \{\text{jm}\}\\
m_1 \odot a \odot (p_3^* \vee p_4^*)! \vee m_2 \odot a \odot (p_3^* \vee p_4^*)! &\leq p_3^* && \{\text{jm}\}\\
m_2 \odot a \odot (p_3^* \vee p_4^*)! &\leq p_3^* && \{\text{lemma (8.4)}\}\\
p_3^* &\leq p_3^* && \{\text{lemma (8.5)}\}
\end{aligned}
$$

∎

The second, third and fourth proof lines appeal to the Galois connections for the a movement, the honest public announcement by the sensor and the belief modality respectively. The uncertainty update inequality is applied twice to reduce a more complex expression to a simpler one by separating the operands of the module operation. The resulting expression is in terms of the uncertainty about events, $u_\alpha^E(a)$ and $u_\alpha^E((p_3^* \vee p_4^*)!)$, in addition to the uncertainty about the initial state. By definition of movements and honest public announcements, the uncertainty of both events is the identity function. By lemma (8.6) and lemma (8.1), the robot's uncertainty $u_\alpha^M(m_2)$ about its initial location p_2 lets it conclude that the subsequent a movement must lead either to location p_2 or p_3. After the honest public announcement by the sensor, the robot can eliminate the possibility of being in location p_2 by lemma (8.4). By lemma (8.5), it remains only the possibility of being in location p_3 proving the claim.

Appendix B: Disbelief

For the proof of proposition (5.2), we use the next lemma which appears in meta-mathematical discussions about Boolean entailment[1] [19].

Lemma 8.8. *Let M be a Boolean module. Let $x, y \in M$. Then,*

$$x \wedge y = \bot \quad \textit{iff} \quad x \leq \neg y$$

The next lemma states the well-known result that every sup-homomorphism, which preserves arbitrary joins, has a right adjoint (e.g. [9]).

Lemma 8.9. *Let P and Q be ordered sets and $f \colon P \to Q$ be a function that preserves arbitrary joins. Then, there exists a function $g \colon Q \to P$ such that $f \dashv g$. In fact, $g(y) = \bigvee \{x \in P \mid f(x) \leq y\}$.*

The last lemma concerns the join of a set of elements in a complete lattice.

Lemma 8.10. *Let P be a complete lattice. Let $p \in P$ and $Q \subseteq P$. Then,*

$$\bigvee Q \leq p \quad \textit{iff} \quad x \leq p \textit{ for all } x \in Q$$

Proof. The proof is trivial by definition of least upper bound of the set Q. ■

The combination of these lemmas together with full distributivity of the Boolean module proves the algebraic characterization of an agent's disbeliefs.

Proposition 8.11. *Let M be a Boolean module. Let $m, \phi \in M$ be states. Then, $m \leq \neg \Box_\alpha \phi$ if and only if, for all $x \in M$, $u^M_\alpha(x) \leq \phi$ implies $m \wedge x = \bot$.*

Proof.

$$\begin{aligned} m \leq \neg\Box_\alpha\phi \quad &\text{iff} \quad m \wedge \Box_\alpha \phi = \bot && \{\text{lemma (8.8)}\} \\ &\text{iff} \quad m \wedge \bigvee \left\{x \in M \mid u^M_\alpha(x) \leq \phi\right\} = \bot && \{\text{lemma (8.9)}\} \\ &\text{iff} \quad \bigvee \left\{m \wedge x \in M \mid u^M_\alpha(x) \leq \phi\right\} = \bot && \{\text{distributivity}\} \\ &\text{iff} \quad (\forall x \in M)\, u^M_\alpha(x) \leq \phi \text{ implies } m \wedge x = \bot && \{\text{lemma (8.10)}\} \end{aligned}$$

■

What remains to show is theorem (5.3) as a special case of proposition (5.2).

Theorem 8.12. *Let M be a Boolean module. Let $m, \phi \in M$. If m is an atom and $u^M_\alpha(m) \not\leq \phi$, then $m \leq \neg\Box_\alpha\phi$.*

Proof. Assume m is an atom and $u^M_\alpha(m) \not\leq \phi$. Let $x \in M$. By proposition (5.2), the consequent is equivalent to the implication if $u^M_\alpha(x) \leq \phi$, then $m \wedge x = \bot$. Assume $u^M_\alpha(x) \leq \phi$. Since m is an atom, the infimum of m and x is the least element provided that $m \not\leq x$. For the purpose of reaching a contradiction, assume that $m \leq x$. Since $u^M_\alpha(-)$ preserves arbitrary joins, it is monotonic. Therefore, $u^M_\alpha(m) \leq u^M_\alpha(x)$. By assumption that $u^M_\alpha(m) \not\leq \phi$, we conclude that $u^M_\alpha(x) \not\leq \phi$. However, we assumed that $u^M_\alpha(x) \leq \phi$. We reached a contradiction. Therefore, $m \not\leq x$ proving the claim. ■

[1] $\neg x \vee y = \top$ if and only if $x \leq y$ for all elements x, y in a Boolean algebra.