The Thirty-First International Florida
Artificial Intelligence Research Society Conference (FLAIRS-31)

Reasoning with Doxastic Attitudes in Multi-Agent Domains

Ben Wright, Enrico Pontelli
Department of Computer Science
New Mexico State University
{ bwright, epontell} @cs.nmsu.edu

Abstract

In recent years, we have witnessed a blossoming of research proposals addressing the challenges in reasoning about action and change in domains that include an agent operating in a multi-agent setting. In particular, the recent emphasis has been on dealing with domains that involve agents reasoning not only about the state of the world but also about the knowledge and beliefs of other agents. An open challenge is the management of conflicting and incorrect beliefs. This paper seeks to introduce a solution to this through the use of *doxastic attitudes*. Built on top of the action language $m\mathcal{A}+$, we extend the transition functions of an agent to include this idea of attitudes and showcase how these work in two different examples.

Introduction & Motivation

Reasoning about action and change has long been a field of study; more recently, a push towards reasoning about multi-agent actions has been studied. In particular, researchers have emphasized the importance of reasoning about beliefs and knowledge of agents and the impact that epistemic actions have on them. However, in many of these cases of belief reasoning, either the system is developed from the perspective of a single agent in the system or it is not robust enough to handle beliefs that conflict or might be false. This is illustrated in the following example:

> **Light in the Room Example**: There are two agents A and B. Agents A and B are in Room2. In Room2, there is a light switch. The switch turns the light on or off for that room. The light cannot be seen from a different Room1. An agent can look at the light to determine if it is on or not. Agents can announce to a room if the light is on or off.
>
> Agent A believes the light is on, after seeing Room2. It then moves to Room1. After some time in Room1, Agent B enters Room1 and announces that the light in Room2 is off. Agent A now has an invalid belief (light on in Room2) that will need to be repaired.

To the best of our knowledge, no current epistemic planner has the ability to change the context in which beliefs are reasoned upon during the reasoning process (Baral et al. 2017; Wan et al. 2015). That is, no contextual beliefs would modify an agent's reasoning to change its belief in the moment.

This paper will introduce the concept of *doxastic attitudes* to help represent this idea of changing beliefs over time and to address issues of false or conflicting beliefs. We introduce some relevant background of Kripke Structures and the action language $m\mathcal{A}+$ which we build on, then we define the concept of Doxastic Attitudes. These are then utilized to modify $m\mathcal{A}+$ to handle attitudes, with intuitions of the transition functions given. We introduce both dynamic and static versions of doxastic attitudes. Finally, we work through two different examples to showcase how the attitudes function.

Background

Kripke Structures

A Kripke structure is a formalism commonly used to capture the possible-world semantics for Logic involving epistemic or doxastic operations (Ditmarsch, van der Hoek, and Kooi 2007; Fagin et al. 1995). Given a countable set of propositions P and a finite set of agents A, a Kripke structure is a structure $M = \langle S, R^A, V^P \rangle$, where:

- S is a set of states
- R^A is a function $\forall a \in A$ where $R^A(a) \subseteq S \times S$
- $V^P : P \to 2^S$ is a valuation function, where $\forall p \in P$, $V^P(p) \subseteq S$ is the set of states in which the proposition p is true.

In this structure, each state in S is a possible world of our domain. The function R^A then maps equivalence between possible worlds for each agent. That is, if an agent has a relation between states s and t, then that agent cannot discern between those two possible worlds.

Additionally, a *pointed Kripke Structure* (M, s) is composed of a Kripke structure, M, and a distinguished state $s \in S$—typically representing the "real" state of the world. It is common to express entailment of the truth of an epistemic formula w.r.t. a pointed Kripke structure—where an epistemic formula is built using propositions from P, propositional connectives, and the operator K_a for any $a \in A$. The following rules intuitively capture such notion of entailment, where $p, q \in P, s, t \in S$ and $a \in A$ (Ditmarsch, van der Hoek, and Kooi 2007; Fagin et al. 1995):

- $(M, s) \models p$ iff $s \in V(p)$,
- $(M, s) \models (\varphi \wedge \psi)$ iff $(M, s) \models \varphi$ and $(M, s) \models \psi$,
- $(M, s) \models \neg\varphi$ iff $(M, s) \not\models \varphi$,
- $(M, s) \models K_a(\varphi)$ iff $\forall(s, t) \in R^A(a).(M, t) \models \varphi$.

The last operator $K_a(\varphi)$ indicates that agent a knows φ while in state s of model M. At this point, we use this definition interchangeably with belief (represented as $B_a(\varphi)$) due to the fact that we do not actually check if $(M, s) \models \varphi$. We utilize the $B_A\varphi$ operator further in the paper to represent '*Agent A believes* φ' and use the definition above for its connotation with pointed Kripke Structures. A final note about the B_a operator is that it can be nested to discuss other agents beliefs, such as B_aB_bon which is "*Agent A believes Agent B believes on*".

Action Language $m\mathcal{A}+$

In this section we briefly summarize the structure of the action language $m\mathcal{A}+$ (Baral et al. 2015), used to describe the capabilities of agents operating in a multi-agent setting. $m\mathcal{A}+$ builds on a signature $\langle \mathcal{AG}, \mathcal{F}, \mathcal{A}\rangle$, where $\mathcal{AG}$ is a set of agents, $\mathcal{F}$ is a set of fluents, and $\mathcal{A}$ is a set of actions.

An action theory in $m\mathcal{A}+$ is composed of a set of axioms, describing the capabilities of the agents. The syntax for the basic actions of $m\mathcal{A}+$ is as follows:

- **executable** a **if** ψ
  describes the fact that the property ψ (an epistemic formula) is a pre-condition for the action a.
- a **causes** Φ **if** ψ
  describes the world-changing effects of action a—it causes a set of fluents Φ to become true if ψ currently holds.
- a **determines** p **if** ψ
  describes a sensing action a; it will determine the value of the fluent p if ψ currently holds.
- a **announces** ℓ **if** ψ
  describes an announcement action a; it declares that the fluent literal ℓ is true to the other agents if ψ currently holds.

In addition to these action types, it is important to know the *awareness* of an agent. If an agent is aware of an action being performed, that it will change its knowledge. To keep track of this, $m\mathcal{A}+$ uses the idea of observations:

- X **observes** a **if** φ

where X is a set of agents. An agent in X either observes or is aware of action a when the conditions φ are met. Aside from this state of awareness, there is also the *oblivious* state—which holds when an agent has no awareness or notion that an action or its effects have occurred. Its important to note that these different observation states allow for the delineation between full "announcements" and group or even secret announcements to various agents. If an agent is not observing or aware of an announcement, then it will not receive that "broadcast." Note that there is an intermediate level of observability (see (Baral et al. 2015) for details)—where agents are only partially aware of execution of an action (i.e., they see an action being performed but are unaware of the outcome). For the sake of simplicity, we omit this intermediate level of observability in this work.

The semantics of the action language can be described as the composition of separate transition functions, each describing the behavior of a different type of action. All transition functions are concerned with '*How do we transition from our current state, to this new state given we performed action a?*'. To represent states, $m\mathcal{A}+$ uses pointed Kripke structures (M, s) as defined above.

For $m\mathcal{A}+$ transitions, we provide only an intuition—the reader can find the complete details in (Baral et al. 2015). There are two facets to each action transition, agents who are aware of the action and agents who are *oblivious*. Intuitively, the world-changing actions are the simplest to describe. If the action is performed and observed, then something in the world changes—we go to a new state s'. Likewise, if an observed announcement happens, some relations are modified based on the announced literals. These would be part of the (M', s') structure in Fig. 1 on the right side. Oblivious agents still believe what was before, both about themselves, the state they are in, and *the other agents*. To do this, we create a mirror replica of our pointed Kripke Structure. This is a snapshot the oblivious agent stays in and where its reach lives. This can be seen as the (M, s) structure in Fig. 1 on the right side. The changed state s' of a world changing action or the changed relations of an announced action take place in the non-replica section—then we go through and change all the relations for oblivious agents to point to the mirrored replica (see Fig. 1 for an intuition).

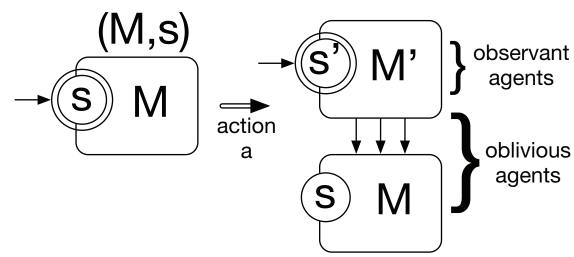

Figure 1: Observability and Action Execution

Doxastic Attitudes

If we are to discuss false beliefs or conflicting beliefs, we will have to ask the question: *"How will the agents decide what to do about their current conflict?"*. One answer to this, is the use of *attitudes* towards beliefs.

Attitudes, or more appropriately propositional attitudes, have been used for sometime (Searle 1983) to convey when a change in belief may occur. When talking about beliefs directly, these can be referred to as *doxastic attitudes*. (Peels 2010) identifies three different results an attitude can produce:

- disbelief: The attitude leads towards the disbelief of whatever was incoming.
- belief: The attitude leads toward believing whatever was incoming.

- suspension of belief: The attitude leads toward removing any belief (for or against) whatever was incoming.

This facet of 'incoming' though, leads towards how attitudes can be categorized or triggered. (Steup 2000) discusses three situations in which an attitude might be attached and have an effect:

- unjustified: there is no previous belief or knowledge or any circumstantial beliefs about this.
- conflicting evidence: there is a belief or knowledge which does not agree with this.
- perceptual evidence: there is knowledge (from sensory information) that is relevant to what is incoming.

These three revolve heavily around the idea of *justification* and/or evidence towards beliefs. This is a large part of philosophical conversations about knowledge and belief (Feldman 2009; Baltag, Fiutek, and Smets 2016).

While these facets of attitudes are useful to this conversation, some parts are less so. There have been many papers linking attitudes with *intention*. Which is to say that many attitudes can be forward-thinking; such as "hoping" or "desiring" something for the future. This focus on future action does not play as large a role in the current pursuit of this paper due to the modelling of the action theory and its possible-world belief space where actions do not inherently exist. Actions are instead what takes us from one pointed Kripke Structure to the next. To this end, some logical models have used Aumann Structures and *event-based modelling* (Cohen and Levesque 1990) to express intentionality or attitudes. However, this loses the granularity of reasoning and possible-world mentality that is so useful through Kripke Structures.

While not looked favorably by some in the past (Ikuenobe 2001), *dogmatic* attitudes can be a source for easy testing in our case. In this sense, dogmatic is the "Always..." or "Never... " type statements where no context is really needed. An example of this can be seen by credulous agents, "I'll believe everything that is told to me."

Some other examples in the context of multi-agent doxastic attitudes could be as follows:

- Agent A will suspend belief on the light if the announcement of the light from agent B contradicts its held beliefs.
- Agent A will believe an announcement on the light from Agent B if and only if it validates their already held beliefs.
- Agent A will always believe they move rooms when they move rooms.
- Agent A will disbelieve agent C's announcement of the coin being heads if they believe Agent B believes it to be tails.

Multi-agent Action Language with Doxastic Attitudes

Syntax for Dynamic Doxastic Attitudes

Using the template of $m\mathcal{A}+$, we can modify the use of observations to the following rules:

$$X \textbf{ observes\_and\_suspends } a \textbf{ if } \varphi$$
$$X \textbf{ observes\_and\_believes } a \textbf{ if } \varphi$$
$$X \textbf{ observes\_and\_disbelieves } a \textbf{ if } \varphi$$

where X is an agent, a is an action, and φ is an executability condition.

Examples:

- A **observes_and_suspends** $announceLightB$ **if** $(RmA \wedge RmB) \vee (\neg RmA \wedge \neg RmB) \wedge B_A \neg on$
  This would read, "Agent A observes and suspends belief of announceLight from agent B if they are in the same room and Agent A believes the light to be off."
- A **observes_and_believes** $announceLightB$ **if** $(RmA \wedge RmB) \vee (\neg RmA \wedge \neg RmB) \wedge B_A on$
  This would read, "Agent A observes and believes announceLight from agent B if they are in the same room and Agent A already believes the light is on."
- A **observes_and_believes** $moveRoomA$ **if** $true$
  This would read, "Agent A observes and believes when it moves Rooms."
- A **observes_and_disbelieves** $announceHeads(C)$ **if** $B_A(B_B tails)$
  This would read, "Agent A observes and disbelieves the announcement of $Heads$ from Agent C if Agent A believes that Agent B believes the coin is showing $Tails$."

Transition Functions for Observational Doxastic Attitudes

The idea for the transition functions for Doxastic Attitudes is very similar to how observations were originally handled in $m\mathcal{A}+$. However, there are more cases to check now. Instead of just "Observe" we have three varieties, which we can separate out as:

$$S_D(a, M, s) = \{X \in AG | \, [X \textbf{ observes\_and\_suspends } a \textbf{ if } \varphi] \in D \textit{ s.t. } (M, s) \models \varphi\}$$
$$B_D(a, M, s) = \{X \in AG | \, [X \textbf{ observes\_and\_believes } a \textbf{ if } \varphi] \in D \textit{ s.t. } (M, s) \models \varphi\}$$
$$D_D(a, M, s) = \{X \in AG | \, [X \textbf{ observes\_and\_disbelieves } a \textbf{ if } \varphi] \in D \textit{ s.t. } (M, s) \models \varphi\}$$

Intuitively, we can base our ideas off of how observation worked originally except for one large crucial step. These transitions need to add relations back into the Pointed Kripke Structure. Previously, knowledge only grew in $m\mathcal{A}+$. However, with the introduction of "suspending" or "disbelieving" attitudes, we need to remove beliefs. This is done by adding back in the proper relations in the R^A function for that agent. Put another way, adding in relations increases the unknown to the Structure, while removing relations increases the amount of known. With this caveat in mind, the difference between each type of attitude is essentially which relations are added and which are removed. A break down for each type of action follows.

World Changing Action Transitions using Observable Doxastic Attitudes We can start with the *observe and believe* attitude, which would work just like the previous *observe* from $m\mathcal{A}+$ rules.

Intuitively, given an action of a **causes** Φ **if** ψ, we know that we will have a state change for the pointed Kripke structure from (M, s) to (M, s') such that $\Phi \in s'$.

The attitudes can change relations however. If we *observe and believe*, then we can remove any relations that go to a state where $\neg\Phi$ exists, and add in relations where Φ exists. Likewise, *observe and disbelieve* would be the opposite. We can add in any relations that go to a state were $\neg\Phi$ exists and remove any relations going to a state where Φ holds.

This just leaves *observe and suspend belief*. For this, we want neither Φ nor $\neg\Phi$ to hold. Therefore, we add in relations that go to both Φ and $\neg\Phi$, removing no relations. This ensures we cannot find either belief to hold.

Sensing Transitions using Observable Doxastic Attitudes Sensing actions, a **determines** φ **if** ψ, would work very similarly to the world changing transitions. Except that no state changes at the beginning. Like the world changing transition, the *observe and believe* attitude is the same as the traditional *observe* from $m\mathcal{A}+$. Here we remove any relations that end in a state where $\neg\varphi$ hold and we add in the relations where they end in a state where φ holds. Then, *observe and disbelieve* would behave in the opposite manner, we would add in relations that end in a state where $\neg\varphi$ holds and remove any relations that end in a state where φ holds. Finally, just like in the world-changing action, *observe and suspend* would add in relations for both φ and $\neg\varphi$ and not remove any relations.

Announcement Transitions using Observable Doxastic Attitudes Announcement actions, a **announces** φ **if** ψ, are practically identical to Sensing Actions. The difference comes from generally how many agents are affected. Sensing is usually internal, while Announcing is usually external.

Just like the two previous examples, the *observe and believe* attitude is the same as *observe* from $m\mathcal{A}+$. We remove any relations that end in a state where $\neg\varphi$ holds and we add in the relations where they end in a state where φ holds. *Observe and disbelieve* would function in the opposite manner, adding in relations that end in a state where $\neg\varphi$ holds and removing any relations that end in a state where φ holds. Lastly, *observe and suspend* would add in relations for both φ and $\neg\varphi$ and not remove any relations.

Static Causal Doxastic Attitudes

Static Causal Doxastic Attitude Syntax and Intent

We can have static causal doxastic attitude rules as follows:

X **suspends_belief_on** f **if** φ
X **believes** l **if** φ
X **disbelieves** l **if** φ

where X is an agent, f is a fluent, l is a literal, and φ is an executability condition. Some examples would look like:

- A **suspends_belief_on** on **if** $\neg RmA$
- B **believes** on **if** $\neg RmA$
- A **disbelieves** on **if** $\neg RmA$

This represents the idea that Agent A will suspend all beliefs on the light being on or off while not in the room the light is in. At first glance, this seems to work fairly well for our representation. However, what happens when agent B comes along and announces that the light is off? Currently, due to the static law - Agent A *cannot* do anything with the value of the light besides suspend belief. This seems a bit overbearing as it seems there might be some mixture of attitudes and announcements in which there is a case in which Agent A would have an attitude to suspend belief, someone announces something about that belief, and then Agent A modifies their beliefs based on that announcement.

At this point, we then need to think about the ordering or priority of attitudes. Let's say that agent A has the attitude,"If I hear an announcement from agent B, I will believe it". Now we can have a situation in which Agent A is not in the room and also hears an announcement from B about the light! Which attitude do we follow in this case?

Based on that example, it seems that the more action-focused attitudes should take precedent over the default/global attitudes. So it should be expected that Agent A would then have a belief about the light based on agent B's announcement - even though its still not in the room. However, this is not the case. Below we show how the transitions for static causal doxastic attitudes work - based on similar transition rules from Static Causal Laws in action language $\mathcal{B}$.

Transitions for Static Causal Doxastic Attitudes

Previously, static casual laws followed a transition function similar to:

$$s' = Cn_Z(E(A, s) \cup (s \cap s'))$$

where Cn is a closure function, Z is the set of all static laws of D, and $E(A, s)$ is the set of effects from dynamic actions A that are possible in state s. This is taken from Action Language $\mathcal{B}$ (Gelfond and Lifschitz 1998):

So, our version would be akin to:

$$\Phi_D(a, (M, s)) = \{Cn_Z((M', s'))\}$$

Where (M', s') is defined by the process in Φ_D^a,Φ_D^s, and Φ_D^w.

Instead of the one version though, we have three forms of static laws - just like the observational doxastic attitudes above. The process for this is quite similar to the above transitions.

For X **suspends_belief_on** f **if** φ, we want $\neg B_X f \wedge \neg B_X \neg f$ to hold in (M, s) if φ is true. For this to happen, we add in any missing relations that go to a state in which f or $\neg f$ hold.

For X **believes** l **if** φ, we want $B_X l$ to hold in (M, s) if φ is true. For this to happen, we add in any missing relations that go to a state in which l holds and we remove any relations that go to a state in which $\neg l$ holds.

For X **disbelieves** l **if** φ, we want $\neg B_X l$ to hold in (M, s) if φ is true. For this to happen, we add in any missing relations that go to a state in which $\neg l$ and remove any relations that go to a state in which l holds.

By adding these to Cn_Z, this has the added consequence of having static attitudes be of a higher priority than dynamic attitudes. For instance, lets assume we have the following two attitudes in our $Domain$:

A **suspends_belief_on** on **if** $\neg RmA$
A **observes_and_believes** $announceLightB$ **if** $(\neg RmA \wedge \neg RmB)$

Here, if Agent A is not in the Room along w/ Agent B and B announces $announceLightB$ then no matter what B announced, A will continue to $suspendbelief$ – due to Closure happening after the generation of $(M", s")$.

Example using Doxastic Attitudes

With these new attitudes, we can look at the example given in the Introduction again. The interesting scenario in this example is when Agent A is in the room without the light and has a belief about the light when Agent B enters and announces the light value.

From this example, let's identify any attitudes that might stand out. If we were using a skeptical approach for these agents, it would follow that if an agent A cannot sense the light, and agent B announces the opposite of what it believes - agent A wouldn't necessarily believe agent B:

A **observes_and_believes** $announceLightA$ **if** $true$
A **observes_and_believes** $announceLightB$ **if** $(RmA \wedge RmB) \vee (\neg RmA \wedge \neg RmB)$
A **observes_and_suspends** $announceLightB$ **if** $(RmA \wedge RmB) \vee (\neg RmA \wedge \neg RmB) \wedge B_A \neg on$

These say that agent A will always believe itself when it announces the light. However if agent B announces the light to A and agent A believes the light to be off, it will suspend belief instead.

Transition Examples

Using these attitudes, we can see what happens when Agent A is in Room1.

If Agent A beliefs the light to be on, then Agent A has no relations that lead to a state where the light is off this is by definition of $(M, s) \models B_A on$. If Agent B then announces $\neg on$ to Agent A, the $observe\_and\_suspend$ attitude would trigger for Agent A. This would add in relations for agent A in which lead to states where the light is off, making agent A then have the following belief about on: $\neg B_A on \wedge \neg B_A \neg on$.

If, at this point, agent B were to announce again $\neg on$, the observe and believe attitude would trigger! Since agent A no longer believes the opposite of what agent B announces, the $observe\_and\_believe$ attitude would trigger having agent A remove any relations that end with a state that holds on. Thus causing agent A's belief about on to change to $B_A \neg on$.

So a series of actions that this example could use to arrive here could be the following: $\{checkLight(A), checkLight(B), moveRoom(A), flipLight(B), checkLight(B), moveRoom(B), announceLight(B), announceLight(B)\}$.

Another Example

While the above shows some perspective for attitudes in a multi-agent domain, a more complicated example is below.

Example Description

> There are 4 agents, A, B, C, and D. Agents A and B read the paper, while Agents C and D watch the news on TV. Agent A believes what they see from the newspaper, Agent B disbelieves it. Likewise, Agent C believes what they see on the TV and Agent D disbelieves it. No agents will change their beliefs based on the same source they can access, however if they hear someone mention from a different source, they will suspend their current belief. Additionally, if an agent doesn't believe or disbelieve a fact and hears a belief for it the 2nd time, they will believe it. Given this setup, is it possible for all 4 agents to belief the same fact?

This example is interesting for a few reasons. First, it utilizes only the sensing and announcement type actions. Second, it shows much more variety in attitude changes. Third, each agent *tells* the other agents what they announce is the same time. That is, all the tells are broadcasts. So every agent's beliefs will change differently based on which agent did the announcement.

Example Domain

Here is a basic domain D setup of this example using $m\mathcal{A}+$ and the new attitude rules.
$\mathcal{AG} = \{A, B, C, D\}$
$\mathcal{F} = \{factoid\}$
$\mathcal{A} = \{tell(\varphi, source), readPaper, watchNews\}$

$readPaper$ **determines** $factoid$
$watchNews$ **determines** $factoid$
$tell(\varphi, source)$ **announces** φ

A **observes_and_believes** $readPaper$
B **observes_and_disbelieves** $readPaper$
C **observes_and_believes** $watchNews$
D **observes_and_disbelieves** $watchNews$

$\{A, B\}$ **observe_and_suspend** $tell(\varphi, news)$ **if** $\neg\varphi$
$\{C, D\}$ **observe_and_suspend** $tell(\varphi, paper)$ **if** $\neg\varphi$

$\{A, B\}$ **observe_and_believe** $tell(\varphi, paper)$ **if** $B_i\varphi$
$\{C, D\}$ **observe_and_believe** $tell(\varphi, news)$ **if** $B_i\varphi$

$\{A, B\}$ **observe_and_disbelieve** $tell(\varphi, paper)$ **if** $\neg B_i\varphi$
$\{C, D\}$ **observe_and_disbelieve** $tell(\varphi, news)$ **if** $\neg B_i\varphi$

$\{A, B, C, D\}$ **observe_and_believe** $tell(\varphi, source)$ **if** $\neg(B_i\varphi \vee \neg B_i\varphi)$
where, $source = \{A, B, C, D, paper, news\}$ and $i \in Agents$ for that rule. So if $\{A, B\}$ is one rule, the i for that would either be A or B respectively.

Goal, Initial Domain

With this setup, we can say that initially $factoid$ is $true$ and our goal for reasoning is to determine if we can reach a state in which $(B_A factoid \wedge B_B factoid \wedge B_C factoid \wedge B_D factoid) \vee (\neg B_A factoid \wedge \neg B_B factoid \wedge$

Initially	A tells(F,paper)	C tells(F,paper)	D tells(F,C)
$B_A(F)$	$B_A(F)$	$B_A(F)$	$B_A(F)$
$\neg B_B(F)$	$\neg B_B(F)$	$\neg B_B(F) \wedge B_B(F)$	$B_B(F)$
$B_C(F)$	$B_C(F)$	$B_C(F)$	$B_C(F)$
$\neg B_D(F)$	$\neg B_D(F) \wedge B_D(F)$	$B_D(F)$	$B_D(F)$

Table 1: Belief Progression in Facts Example

$\neg B_C factoid \wedge \neg B_D factoid)$. That is all 4 agents belief the same thing, either $factoid$ or $\neg factoid$.

Action Transitions in Example

Due to the announcements effecting all the different agents every time, this example's actions and agent beliefs can be viewed in the given Table 1.

The basic follow through of this example is as follows, Agent A tells everyone $factoid$ based on the $newspaper$. This causes agents B and C to keep their beliefs as is because C already had that and B disbelieves what the $newspaper$ says. Agent D however changes its belief to $suspension$ due to hearing from a different source than the TV and it being not what it currently believes.

The second tell from Agent C also is of $factoid$ now based on the TV. This causes Agent D to believe $factoid$ due to it having a suspended belief previously, and causes agent B to start having a suspended belief. Agent A continues to believe $factoid$ as that is what it already believed.

Finally, Agent D tells $factoid$ based on its source of agent C who was the agent to convince it of $factoid$. This final tell makes Agent B also believe $factoid$ based on it previously suspending its belief.

Discussion & Conclusion

The use of attitudes to represent and reason about the changing beliefs of agents seems to be a viable and worthwhile endeavor to pursue. It addresses the idea of conflicting and false belief, and allows for agents to change behavior in the middle of a reasoning scenario. This can have implications to defining agent roles or representing such things as dishonest or lying actions. Additionally, it is not known to the authors if any other research has pursued the concept of static belief rules for multi-agent domains.

While the examples given are fairly rudimentary, it is the intent of the authors to have these systems be utilized in future research for more specific domains, though there are none currently.

Some limitations do exist however. This action language does not fix the complexity issues of $m\mathcal{A}+$ when multiple oblivious agents are continually encountered. Each time an agent is oblivious, the world-space doubles. That still holds in these transitions. Another issue to note, discussed in the Static Causal Doxastic Attitudes section, is the order in which dynamic and static attitudes are evaluated. Currently dynamic is evaluated before static, so a static attitude can overwrite a dynamic attitude (as seen in the Light Room example). This may not be what some systems want, having rather dynamic attitudes supersede any statically defined attitudes.

It is our intent to develop an implementation of these new attitude transitions in the near future. This implementation will be used to pursue epistemic reasoning in the context of Dishonest Agents. It is the authors' intent that these doxastic attitudes can play a pivotal role in the reasoning of agents when beliefs are not guaranteed to be true.

References

Baltag, A.; Fiutek, V.; and Smets, S. 2016. Beliefs and evidence in justification models. In *Advances in Modal Logic*, 156–176.

Baral, C.; Gelfond, G.; Pontelli, E.; and Son, T. C. 2015. An action language for multi-agent domains: Foundations. *CoRR* abs/1511.01960.

Baral, C.; Bolander, T.; van Ditmarsch, H.; and McIlrath, S. 2017. Epistemic planning (dagstuhl seminar 17231). In *Dagstuhl Reports*, volume 7. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

Cohen, P. R., and Levesque, H. J. 1990. Intention is choice with commitment. *Artificial intelligence* 42(2-3):213–261.

Ditmarsch, H. v.; van der Hoek, W.; and Kooi, B. 2007. *Dynamic Epistemic Logic*. Springer Publishing Company, Incorporated, 1st edition.

Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1995. *Reasoning About Knowledge*. MIT Press.

Feldman, R. 2009. Evidentialism, higher-order evidence, and disagreement. *Episteme* 6(3):294–312.

Gelfond, M., and Lifschitz, V. 1998. Action languages. *Electronic Transactions on AI* 3.

Ikuenobe, P. 2001. Questioning as an epistemic process of critical thinking. *Educational Philosophy and Theory* 33(3-4):325–341.

Peels, R. 2010. What is ignorance? *Philosophia* 38(1):57–67.

Searle, J. R. 1983. *Intentionality: An essay in the philosophy of mind*. Cambridge university press.

Steup, M. 2000. Doxastic voluntarism and epistemic deontology. *Acta Analytica* 24:25–56.

Wan, H.; Yang, R.; Fang, L.; Liu, Y.; and Xu, H. 2015. A complete epistemic planner without the epistemic closed world assumption. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.