Multi-agent Knowledge and Belief Change in the Situation Calculus

Liangda Fang and **Yongmei Liu**
Department of Computer Science
Sun Yat-sen University
Guangzhou 510006, China
fangld@mail2.sysu.edu.cn, ymliu@mail.sysu.edu.cn

Abstract

Belief change is an important research topic in AI. It becomes more perplexing in multi-agent settings, since the action of an agent may be partially observable to other agents. In this paper, we present a general approach to reasoning about actions and belief change in multi-agent settings. Our approach is based on a multi-agent extension to the situation calculus, augmented by a plausibility relation over situations and another one over actions, which is used to represent agents' different perspectives on actions. When an action is performed, we update the agents' plausibility order on situations by giving priority to the plausibility order on actions, in line with the AGM approach of giving priority to new information. We show that our notion of belief satisfies KD45 properties. As to the special case of belief change of a single agent, we show that our framework satisfies most of the classical AGM, KM, and DP postulates. We also present properties concerning the change of common knowledge and belief of a group of agents.

1 Introduction

Knowledge and belief change is an important research topic in logic and AI. In general, an agent may have incomplete and inaccurate information about the world. As she performs physical actions to effect change in the world, or sensing actions to learn new information, she has to modify her knowledge and belief about the world. The issue becomes more perplexing in the presence of multiple agents. In such settings, in addition to first-order beliefs, *i.e.*, beliefs about the world, there are also higher-order beliefs, *i.e.*, beliefs about agents' beliefs, and common beliefs of a group of agents. Moreover, when an agent performs an action, other agents may have different perspectives on it: some may fully observe the action, some may mistake the action for another one, and some may be oblivious of the action.

For example, consider the following scenario adapted from (van Ditmarsch, van Der Hoek, and Kooi 2007). Two stockbrokers Ann and Bob are having a break in a bar. Neither Ann nor Bob knows how the United Agents (UA) company is doing, but Bob believes that UA is not doing well. A messenger comes in and gives Ann a letter marked "urgently requested data on UA". Ann reads the letter and knows that UA is doing well. Bob sees Ann reads the letter, and he still does not know how UA is doing. However, Ann and Bob commonly know that Ann knows how UA is doing.

Knowledge and belief change has been studied in the areas of belief revision, reasoning about actions, and dynamic epistemic logics (DELs) (van Ditmarsch, van Der Hoek, and Kooi 2007). The area of belief revision studies how an agent modifies her beliefs on receiving new information. Various guidelines for belief revision have been proposed, and the most popular ones are the AGM postulates for belief revision (Alchourrón, Gärdenfors, and Makinson 1985), the KM postulates for belief update (Katsuno and Mendelzon 1991), and the DP postulates for iterated belief revision (Darwiche and Pearl 1997). The area of reasoning about actions was historically concerned with physical actions, and was later extended to accommodate epistemic actions. The situation calculus (Reiter 2001) is one of the most popular languages for reasoning about actions. Scherl and Levesque (1993; 2003) proposed an epistemic extension to the situation calculus. Shapiro *et al.* (2011) integrated belief revision wrt accurate sensing into the situation calculus. Delgrande and Levesque (2012) furthered this work by considering noisy sensing and fallible actions. However, all the above works are restricted to the single-agent case.

DELs focus on reasoning about epistemic actions in the multi-agent case. An important concept in DELs is that of an action model, which is a Kripke model of actions, representing the agents' uncertainty about the current action. By the product update operation, an action model may be used to update a Kripke model. Van Benthem (2007) integrated belief revision into DELs. He considered belief change under public announcements of hard and soft facts, and for soft facts, proposed two update rules for changing plausibility orders on states. Baltag and Smets (2008) further presented a general framework for integrating belief revision into DELs. In line with the AGM approach of giving priority to new information, they proposed the action priority update operation: when updating a plausibility model by an action plausibility model, give priority to the action plausibility order.

In this paper, by incorporating action priority update into the situation calculus, we present a general framework for reasoning about actions and belief change in multi-agent settings. We extend the situation calculus by a plausibility relation over situations and another one over actions, and propose a successor state axiom for the plausibility relation over situations. We show that our notion of belief satisfies KD45 properties. As to the special case of belief change of a single agent, we show that our framework satisfies most of the AGM, KM, and DP postulates. We also give properties concerning the change of common knowledge and belief of a group of agents. Finally, we present an illustrating example.

2 Background work

The situation calculus

The situation calculus (Reiter 2001) is a many-sorted first-order language suitable for describing dynamic worlds. There are three disjoint sorts: $action$ for actions, $situation$ for situations, and $object$ for everything else. A situation calculus language has the following components: a constant S_0 denoting the initial situation; a binary function $do(a, s)$ denoting the successor situation to s resulting from performing action a; a binary predicate $s \sqsubset s'$ meaning that situation s is a proper subhistory of situation s'; a binary predicate $\mathrm{Poss}(a, s)$ meaning that action a is possible in situation s; action functions; a finite number of relational and functional fluents, *i.e.*, predicates and functions taking a situation term as their last argument, which denote relations and functions whose values vary from situation to situation; and a finite number of situation-independent predicates and functions.

Based on the situation calculus, a logic programming language Golog (Levesque et al. 1997) was designed for high-level robotic control. The formal semantics of Golog is specified by an abbreviation $Do(\delta, s, s')$, which means executing program δ brings us from situation s to s'. It is inductively defined on δ, and here we only present the definitions we need in this paper.

- Primitive actions:
  $Do(\alpha, s, s') \doteq \mathrm{Poss}(\alpha, s) \wedge s' = do(\alpha, s)$;
- Sequence:
  $Do(\delta_1; \delta_2, s, s') \doteq (\exists s'').Do(\delta_1, s, s'') \wedge Do(\delta_2, s'', s')$.
- Nondeterministic choice of two actions:
  $Do(\delta_1 \,|\, \delta_2, s, s') \doteq Do(\delta_1, s, s') \vee Do(\delta_2, s, s')$.
- Nondeterministic choice of action arguments:
  $Do((\pi\, x)\delta(x), s, s') \doteq (\exists x)Do(\delta(x), s, s')$.

Belief change

Belief change studies how an agent modifies her beliefs in the presence of new information. Here we briefly review the most popular accounts of belief revision, belief update, and iterated belief revision.

Belief revision concerns belief change about static environments due to partial and possibly incorrect information. For illustration, we present the AGM postulates (Alchourrón, Gärdenfors, and Makinson 1985) for belief revision. An agent's beliefs are modeled by a deductively closed set of sentences, BS, called a belief set; hence $BS = Cn(BS)$, where $Cn(BS)$ is the deductive closure of BS. We use $BS * \phi$ for the revision of BS by new information ϕ, and $BS + \phi$ the expansion of BS by ϕ, defined as $Cn(BS \cup \{\phi\})$.

(AGM1) $BS * \phi$ is deductively closed

(AGM2) $\phi \in BS * \phi$

(AGM3) $BS * \phi \subseteq BS + \phi$

(AGM4) If $\neg\phi \notin BS$, then $BS + \phi \subseteq BS * \phi$

(AGM5) If ϕ is consistent, so is $BS * \phi$

(AGM6) If ϕ and ψ are equivalent, then $BS * \phi = BS * \psi$

(AGM7) $BS * (\phi \wedge \psi) \subseteq (BS * \phi) + \psi$

(AGM8) If $\neg\psi \notin BS * \phi$, then $(BS * \phi) + \psi \subseteq BS * (\phi \wedge \psi)$

Belief update concerns belief change about dynamic environments due to the performance of actions. Katsuno and Mendelzon (1991) presented the following postulates for belief update, where $BS \diamond \phi$ denotes the update of BS by formula ϕ, and $[BS]$ the set of maximum consistent theories containing BS.

(KM1) $BS \diamond \phi$ is deductively closed

(KM2) $\phi \in BS \diamond \phi$

(KM3) If $\phi \in BS$, then $BS \diamond \phi = BS$

(KM4) If both BS and ϕ are consistent, so is $BS \diamond \phi$

(KM5) If ϕ and ψ are equivalent, then $BS \diamond \phi = BS \diamond \psi$

(KM6) $BS \diamond (\phi \wedge \psi) \subseteq (BS \diamond \phi) + \psi$

(KM7) If $\psi \in BS \diamond \phi$ and $\phi \in BS \diamond \psi$, then $BS \diamond \phi = BS \diamond \psi$

(KM8) If BS is complete, then $BS \diamond (\phi \vee \psi) \subseteq Cn((BS \diamond \psi) \cup (BS \diamond \psi))$

(KM9) $BS \diamond \phi = \bigcap_{r \in [BS]} r \diamond \phi$

The above postulates only provide guidelines for one-shot change of beliefs. Darwiche and Pearl (1997) proposed the following postulates regarding iterated belief revision.

(DP1) If $\psi \models \phi$, then $(BS * \phi) * \psi = BS * \psi$

(DP2) If $\psi \models \neg\phi$, then $(BS * \phi) * \psi = BS * \psi$

(DP3) If $\phi \in BS * \psi$, then $\phi \in (BS * \phi) * \psi$

(DP4) If $\neg\phi \notin BS * \psi$, then $\neg\phi \notin (BS * \phi) * \psi$

Action priority update

Baltag and Smets (2008) integrated belief revision into DELs. Their epistemic model is a plausibility model, where for each agent, there is a plausibility order on the set of states. An agent believes ϕ if ϕ holds in the most plausible states. Belief revision is triggered by action plausibility models, which represent each agent's plausibility order on the set of actions. When we update a plausibility model by an action plausibility model, we give priority to the action plausibility order. Here we briefly review the basic concepts in their work.

Definition 2.1 (Locally well-preordered relation) A preorder $\leq$ is a reflexive and transitive binary relation. We use $\sim$ for the associated comparability relation, *i.e.*, $s \sim t$ iff $s \leq t$ or $t \leq s$. The comparability class for an element s, written $[s]$, is the set $\{t \mid s \sim t\}$. We say that $\leq$ is locally well-founded if every non-empty subset of every comparability class has a least element. A relation is locally well-preordered if it is a locally well-founded preorder.

Definition 2.2 (Multi-agent plausibility model) A multi-agent plausibility frame (MPF) is a structure $(S, \leq_i)$, where S is a set of states, and for each agent i, $\leq_i$ is a locally well-preordered relation. A multi-agent plausibility model is a MPF together with a valuation map, which maps each state into a subset of the propositional atoms.

The usual reading of $s \leq t$ is that "s is at least as plausible as t". We use the following notation.

1. $s < t$ iff $s \leq t$ and $t \not\leq s$ (s is more plausible than t);
2. $s \cong t$ iff $s \leq t$ and $t \leq s$ (s and t are equally plausible);
3. $s \to t$ iff t is a least element of $[s]$ (t is a most plausible state comparable to s).

Definition 2.3 (Knowledge and belief) Let M be a multi-agent plausibility model, and s a state of M. We define:

1. $M, s \models K_i\phi$ iff for every t such that $s \sim_i t$, $M, t \models \phi$;
2. $M, s \models B_i\phi$ iff for every t such that $s \to_i t$, $M, t \models \phi$.

Thus, agent i knows ϕ in s if ϕ holds in all states comparable to s, and i believes ϕ in s if ϕ holds in all the most plausible states comparable to s.

Definition 2.4 (Action plausibility model) An action plausibility model is a plausibility frame $(\Gamma, \leq_i)$ together with a map pre, which maps each action point $\gamma \in \Gamma$ into a formula, called the precondition of γ.

Definition 2.5 (Action-Priority Update) Let $M = (S, \leq, V)$ be a plausibility model, and let $N = (\Gamma, \leq, pre)$ be an action plausibility model. The product of M and N, is a new plausibility model $(S', \leq, V')$, defined as follows:

1. $S' = \{(s, \gamma) \mid s \in S, \gamma \in \Gamma, \text{ and } M, s \models pre(\gamma)\}$;
2. $(s, \gamma) \leq_i (s', \gamma')$ iff either $\gamma <_i \gamma'$ and $s \sim_i s'$, or else $\gamma \cong_i \gamma'$ and $s \leq_i s'$;
3. For each $(s, \gamma) \in S'$, $V'((s, \gamma)) = V(s)$.

Thus the updated plausibility order gives priority to the action plausibility relation, and apart from this it keeps as much as possible the old order.

3 Our formal account

In this section, by incorporating action priority update, we develop a formal account of multi-agent knowledge and belief in the situation calculus. We begin with our extension of the language. Then we specify the components of a multi-agent basic action theory. Finally, we define multi-agent knowledge and belief in the extended situation calculus.

An extension of the situation calculus

To model plausibility order, we introduce a special fluent $B(i, s_2, s_1, s)$, which means that in situation s, agent i considers situation s_2 at least as plausible as s_1. Note that unlike other works, our B fluent has 3 situation arguments, this is because an agent's plausibility order at different situations might be different. To model action plausibility order, we introduce a special predicate $A(i, a_2, a_1, a, s)$, meaning that when action a is performed in situation s, agent i considers that a_2 is executed at least as plausible as that a_1 is executed.

We assume that there are two types of primitive actions: ordinary actions which change the world, and epistemic actions which do not change the world but tell the agent that some condition holds in the current situation. We use the action precondition axiom to specify what the epistemic action tells the agent about the world. For example, we may have an epistemic action $ison(i, x)$ which tells agent i that switch x is on. This is axiomatized as[1]: $\text{Poss}(ison(i, x), s) \equiv on(x, s)$. In particular, there is a special epistemic action nil, meaning that nothing happens, with the axiom $\text{Poss}(nil, s) \equiv true$. Note that a sensing action which tells the agent whether ϕ holds can be treated as the nondeterministic choice of two epistemic actions: one is possible iff ϕ holds, and the other is possible iff $\neg\phi$ holds.

Multi-agent basic action theories

We introduce the following abbreviations:

1. $\text{Init}(s) \doteq \neg(\exists a, s').s = do(a, s')$
2. $\text{Exec}(s) \doteq (\forall a, s^*).do(a, s^*) \sqsubseteq s \supset \text{Poss}(a, s^*)$

Intuitively, $\text{Init}(s)$ says s is an initial situation, and $\text{Exec}(s)$ means s is an executable situation, *i.e.*, an action history in which it is possible to perform the actions one after the other.

Using a second-order formula, we define an abbreviation $\text{Lwf}(s)$ saying that $B(i, s_2, s_1, s)$ is locally well-founded:

$$\text{Lwf}(s) \doteq \forall i, s_1, P.$$
$$\forall s_2(P(s_2) \supset B(i, s_2, s_1, s) \vee B(i, s_1, s_2, s)) \wedge \exists s_3 P(s_3)$$
$$\supset \exists s_4(P(s_4) \wedge \forall s_5(P(s_5) \supset B(i, s_4, s_5, s)))$$

Similarly, we can define an abbreviation $\text{Alwf}(a, s)$ which says that $A(i, a_2, a_1, a, s)$ is locally well-founded.

Then the requirement that B is locally well-preordered in initial situations can be axiomatized as follows:

Definition 3.1 $\mathcal{B}_{Init}$ consists of the following axioms:

1. Reflexivity: $\text{Init}(s) \supset B(i, s_1, s_1, s)$;
2. Transitivity: $\text{Init}(s) \wedge B(i, s_2, s_1, s) \wedge B(i, s_3, s_2, s) \supset B(i, s_3, s_1, s)$;
3. Locally well-founded: $\text{Init}(s) \supset \text{Lwf}(s)$.

Similarly, we describe the requirement that A is locally well-preordered in executable situations as below:

Definition 3.2 $\mathcal{A}_{Exec}$ consists of the following axioms:

1. $\text{Exec}(s) \supset A(i, a_1, a_1, a, s)$;
2. $\text{Exec}(s) \wedge A(i, a_2, a_1, a, s) \wedge A(i, a_3, a_2, a, s) \supset A(i, a_3, a_1, a, s)$;
3. $\text{Exec}(s) \supset \text{Alwf}(a, s)$.

To incorporate action priority update, we propose the following successor state axiom for the B fluent:

$$B(i, s_2', s_1', do(a, s)) \equiv \exists s_1, s_2, a_1, a_2.$$
$$s_1' = do(a_1, s_1) \wedge s_2' = do(a_2, s_2) \wedge$$
$$\text{Poss}(a, s) \wedge \text{Poss}(a_1, s_1) \wedge \text{Poss}(a_2, s_2) \wedge$$
$$\{[A(i, a_2, a_1, a, s) \wedge \neg A(i, a_1, a_2, a, s)$$
$$\wedge (B(i, s_2, s_1, s) \vee B(i, s_1, s_2, s))] \vee$$
$$[A(i, a_2, a_1, a, s) \wedge A(i, a_1, a_2, a, s) \wedge B(i, s_2, s_1, s)]\}$$

[1]Throughout this paper, free variables are assumed to be universally quantified from outside.

Intuitively, after action a is performed in situation s, agent i considers situation s'_2 at least as plausible as s'_1 iff s'_i is the result of doing some action a_i in some situation s_i, $i = 1, 2$, a is possible in s, a_i is possible in s_i, $i = 1, 2$, and either agent i considers a_2 more plausible than a_1 and s_2 comparable to s_1, or she thinks that a_2 and a_1 are equally plausible and s_2 is at least as plausible as s_1.

In the multi-agent case, a domain of application is specified by a basic action theory (BAT) of the form:

$$\mathcal{D} = \Sigma \cup \mathcal{B}_{Init} \cup \mathcal{D}_{ap} \cup \mathcal{D}_{ss} \cup \mathcal{D}_{aa} \cup \mathcal{D}_{una} \cup \mathcal{D}_{S_0}, \text{ where}$$

1. Σ are the foundational axioms:
   - $do(a_1, s_1) = do(a_2, s_2) \supset a_1 = a_2 \wedge s_1 = s_2$;
   - $(\neg s \sqsubset S_0) \wedge (s \sqsubset do(a, s') \equiv s \sqsubseteq s')$;
   - $\forall P. \forall s[\text{Init}(s) \supset P(s)] \wedge \forall a, s[P(s) \supset P(do(a, s))] \supset (\forall s) P(s)$;
   - $B(i, s_2, s_1, s) \supset [\text{Init}(s) \equiv \text{Init}(s_2)] \wedge [\text{Init}(s) \equiv \text{Init}(s_1)]$.

   A model of these axioms consists of a forest of isomorphic trees rooted at the initial situations, which can be B-related to only initial situations.
2. $\mathcal{D}_{ap}$ is a set of action precondition axioms, one for each action function C, of the form $\text{Poss}(C(\vec{x}), s) \equiv \Pi_C(\vec{x}, s)$.
3. $\mathcal{D}_{ss}$ is a set of successor state axioms (SSAs), one for each fluent F (including the B fluent), of the form $F(\vec{x}, do(a, s)) \equiv \Phi_F(\vec{x}, a, s)$. These embody a solution to the frame problem (Reiter 1991). The SSA for any ordinary fluent F must satisfy the condition that F is not affected by epistemic actions. This can be easily satisfied if the SSA for F does not mention any epistemic action.
4. $\mathcal{D}_{aa}$ is a set of action plausibility axioms, one for each action function C, of the form
   $$A(i, a_2, a_1, C(\vec{x}), s) \equiv \Psi_C(i, a_2, a_1, \vec{x}, s).$$
5. $\mathcal{D}_{una}$ is the set of unique names axioms for actions.
6. $\mathcal{D}_{S_0}$ is a set of sentences about S_0.
7. $\mathcal{D} \models \mathcal{A}_{Exec}$.

Throughout the paper, we use $\mathcal{D}$ for a BAT of this form.

Then we can show that only executable situations at the same level of the situation forest are B-related, and B is locally well-preordered in all executable situations. We first introduce an abbreviation $\text{Eqlev}(s_1, s)$ which says that s_1 and s are situations at the same level:

$$\text{Eqlev}(s_1, s) \doteq \forall P.$$
$$(\forall s', s'_1)[\text{Init}(s') \wedge \text{Init}(s'_1) \supset P(s'_1, s')] \wedge$$
$$(\forall a, a_1, s', s'_1)[P(s'_1, s') \supset P(do(a_1, s'_1), do(a, s'))] \supset P(s_1, s).$$

Theorem 3.1 *$\mathcal{D}$ entails the following:*

1. $B(i, s_1, s_2, s) \supset Exec(s) \wedge Exec(s_1) \wedge Exec(s_2) \wedge Eqlev(s_1, s) \wedge Eqlev(s_2, s)$;
2. $Exec(s) \wedge Exec(s_1) \wedge Eqlev(s_1, s) \supset B(i, s_1, s_1, s)$;
3. $Exec(s) \supset Lwf(s)$;
4. $B(i, s_2, s_1, s) \wedge B(i, s_3, s_2, s) \supset B(i, s_3, s_1, s)$.

Multi-agent knowledge and belief

We are now ready to define mental attitudes knowledge, belief, and conditional belief in the situation calculus. We begin with several relations derived from the B fluent:

Definition 3.3 (K, MPB and ConMPB relations)

1. $K(i, s_2, s_1, s) \doteq B(i, s_2, s_1, s) \vee B(i, s_1, s_2, s)$;
2. $\text{MPB}(i, s_2, s_1, s) \doteq \forall s_3. K(i, s_3, s_1, s) \supset B(i, s_2, s_3, s)$;
3. $\text{ConMPB}(i, \psi(now), s_2, s_1, s) \doteq K(i, s_2, s_1, s) \wedge \psi(s_2) \wedge \forall s_3. K(i, s_3, s_1, s) \wedge \psi(s_3) \supset B(i, s_2, s_3, s)$,

where $\psi(s)$ is a formula with a single free variable s, "now" is a placeholder for a situation argument.

Intuitively, $K(i, s_2, s_1, s)$ means that in situation s, agent i considers s_2 comparable to s_1. $\text{MPB}(i, s_2, s_1, s)$ means that according to agent i, s_2 is a most plausible situation in the comparability class of s_1. $\text{ConMPB}(i, \psi, s_2, s_1, s)$ states that according to agent i, s_2 is a most plausible one among situations which are comparable to s_1 and where ψ holds.

We first define mental attitudes in a pair of situations (s_1, s) where s_1 is the actual state, and s determines the plausibility order. Then we define mental attitudes in a situation s as mental attitudes in (s, s).

Definition 3.4 (Knowledge, belief and conditional belief) Let $\phi(s)$ be a formula with a single free variable s.

1. Agent i knows ϕ in situation pair (s_1, s):
   $\text{Know}(i, \phi(now), s_1, s) \doteq \forall s_2. K(i, s_2, s_1, s) \supset \phi(s_2)$;
2. Agent i believes ϕ in situation pair (s_1, s):
   $\text{Bel}(i, \phi(now), s_1, s) \doteq \forall s_2. \text{MPB}(i, s_2, s_1, s) \supset \phi(s_2)$;
3. Agent i believes ϕ after learning ψ in (s_1, s):
   $\text{ConBel}(i, \phi(now_1), \psi(now), s_1, s) \doteq \forall s_2. \text{ConMPB}(i, \psi(now), s_2, s_1, s) \supset \phi(s_2)$.

Then agent i knows ϕ in situation s is represented as $\text{Know}(i, \phi(now), s, s)$. Belief and conditional belief in a situation are similarly defined.

The "now" arguments are placeholders for situation arguments, and are often omitted when no confusion is incurred.

When we write a formula $\text{Know}(i, \phi(now), s, s)$, to improve readability, when s is the second situation argument of a mental attitude sub-formula, we omit this argument. For example, let $p(s)$ be a unary fluent. The statement that in situation s, agent i knows that agent j knows p is written as

$$\text{Know}(i, \text{Know}(j, p(now), now, s), s, s).$$

Instead, we write $\text{Know}(i, \text{Know}(j, p(now), now), s)$. Then we omit "now", and write $\text{Know}(i, \text{Know}(j, p), s)$.

Definition 3.5 (Objective formulas) We say that a situation calculus formula ϕ is *objective*, written $\phi \in \mathcal{L}_{obj}$, if it does not use the B fluent or the A predicate.

We use an abbreviation: Agent i knows whether ϕ holds:
$\text{KW}(i, \phi, s) \doteq \text{Know}(i, \phi, s) \vee \text{Know}(i, \neg\phi, s)$.

Example 1 We now formalize the letter example from the introduction. Let fluent $p(s)$ mean that "UA is doing well in situation s". We introduce an epistemic action $read(i, e)$, which tells agent i that p holds if $e = 1$ and that $\neg p$ holds if $e = 0$. The axioms are:

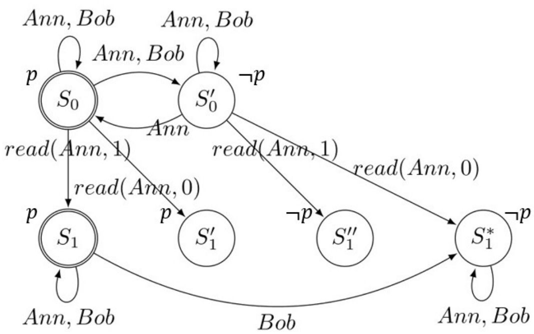

Figure 1: The Letter Example

- $\text{Poss}(read(i, e), s) \equiv (e = 1 \equiv p(s))$;
- $A(j, a_2, a_1, read(i, e), s) \equiv a_1 = a_2 \vee$
  $j \neq i \wedge \exists e_1, e_2(a_1 = read(i, e_1) \wedge a_2 = read(i, e_2))$;
- $p(S_0) \wedge \neg\text{KW}(Ann, p, S_0) \wedge \neg\text{KW}(Bob, p, S_0) \wedge$
  $\text{Bel}(Bob, \neg p, S_0)$.

The A axiom specifies the action plausibility order when action $read(i, e)$ is executed. Agent i knows the action she is performing, so the two actions $read(i, 1)$ and $read(i, 0)$ are not comparable to i. Agent $j \neq i$ cannot distinguish between the two actions, so she considers them equally plausible.

Let $S_1 = do(read(Ann, 1), S_0)$. Then we have that $\mathcal{D} \models \text{Know}(Ann, p, S_1) \wedge \neg\text{KW}(Bob, p, S_1) \wedge \text{Bel}(Bob, \neg p, S_1)$. We illustrate this with a model M of $\mathcal{D}$, and the first two levels of the situation forest of M is shown in Figure 1.

There are two initial situations S_0 where p holds and S'_0 where $\neg p$ holds. Note that our B fluent is of the form $B(i, s_2, s_1, s)$. The interpretation of the B fluent at S_0 (*i.e.*, when the last situation argument is S_0) is indicated by the arrows labeled with the agents. To Ann, both S_0 and S'_0 are comparable to S_0, so $\neg\text{KW}(Ann, p, S_0)$ holds; similarly, $\neg\text{KW}(Bob, p, S_0)$ holds. To Bob, S'_0 is the single most plausible situation comparable to S_0, so $\text{Bel}(Bob, \neg p, S_0)$ holds.

There are four situations at level 2: S_1, S'_1, S''_1, and S^*_1, resulting from performing the respective actions, as shown by the figure. Among them, only S_1 and S^*_1 are executable situations. By the SSA for the B fluent, we get its interpretation at S_1 from that at S_0. To Ann, S_1 and S^*_1 are not comparable since the two actions $read(Ann, 1)$ and $read(Ann, 0)$ are not comparable. To Bob, S^*_1 is more plausible than S_1 since the two actions are equally plausible and S'_0 is more plausible than S_0. It is easy to check that $\text{Know}(Ann, p, S_1) \wedge \neg\text{KW}(Bob, p, S_1) \wedge \text{Bel}(Bob, \neg p, S_1)$ holds.

4 Properties of knowledge and belief

In this section, we analyze properties of knowledge and belief in our framework. We first show that our notion of knowledge is S5 knowledge (*i.e.*, knowledge is truthful and both positively and negatively introspective), our notion of belief is KD45 belief (*i.e.*, belief is consistent and introspective), and our knowledge entails belief.

Theorem 4.1 *$\mathcal{D}$ entails the following:*

1. $Exec(s) \wedge Know(i, \phi, s) \supset \phi(s)$;
2. $Exec(s) \wedge Know(i, \phi, s) \supset Know(i, Know(i, \phi), s)$;
3. $Exec(s) \wedge \neg Know(i, \phi, s) \supset Know(i, \neg Know(i, \phi), s)$.

Theorem 4.2 *$\mathcal{D}$ entails the following:*

1. $Exec(s) \supset \neg Bel(i, false, s)$;
2. $Exec(s) \wedge Bel(i, \phi, s) \supset Bel(i, Bel(i, \phi), s)$;
3. $Exec(s) \wedge \neg Bel(i, \phi, s) \supset Bel(i, \neg Bel(i, \phi), s)$;
4. $Exec(s) \wedge Know(i, \phi, s) \supset Bel(i, \phi, s)$.

We now show properties reducing knowledge and belief after an action is performed to the current knowledge and belief. We need some relations derived from the A predicate.

Definition 4.1 (KA relation)
$KA(i, a_2, a_1, a, s) \doteq A(i, a_2, a_1, a, s) \vee A(i, a_1, a_2, a, s)$.

Theorem 4.3 (Reduction law for knowledge) *$\mathcal{D}$ entails:*
$Exec(s) \wedge Poss(a, s) \supset$
$\{Know(i, \phi, do(a, s)) \equiv \forall a_1[KA(i, a_1, a, a, s) \supset$
$Know(i, Poss(a_1, now) \supset \phi(do(a_1, now)), s)]\}$.

Intuitively, this says: after action a is done in situation s, agent i knows ϕ iff in s, for any action a_1 comparable to a, she knows that if a_1 is possible, ϕ holds after a_1 is executed. We illustrate this law with the following example.

Example 2 Delgrande and Levesque (2012) considered fallible actions: an agent wants to push button m, but she may actually push button n such that $|m - n| \leq 1$. We introduce an action $push(i, m, n)$, meaning that agent i wants to push button m but ends up pushing button n. The axioms are:

- $Poss(push(i, m, n), s) \equiv |m - n| \leq 1$;
- $on(n, do(a, s)) \equiv \exists m.a = push(i, m, n)$;
- $A(i, a_2, a_1, push(j, m, n), s) \equiv a_1 = a_2 \vee$
  $\{i = j \supset \exists n_1, n_2[|m - n_1| \leq 1 \wedge |m - n_2| \leq 1 \wedge$
  $a_1 = push(j, m, n_1) \wedge a_2 = push(j, m, n_2)]\} \wedge$
  $\{i \neq j \supset \exists n_1, n_2, m_1, m_2$
  $[a_1 = push(j, m_1, n_1) \wedge a_2 = push(j, m_2, n_2)]\}$.

The A axiom specifies the action plausibility order when action $push(j, m, n)$ is executed. Agent j cannot distinguish between any two actions $push(j, m, n_1)$ and $push(j, m, n_2)$ such that $|m - n_1| \leq 1$ and $|m - n_2| \leq 1$, so she considers them equally plausible. Agent $i \neq j$ can only observe that agent j pushes a button, so she considers any two pushing actions of j equally plausible.

We would like to show that $\mathcal{D}$ entails:
$\text{Exec}(s) \wedge \text{Poss}(push(i, m, n), s) \supset \text{Know}(i, on(m - 1) \vee on(m) \vee on(m + 1), do(push(i, m, n), s))$.

By the reduction law for knowledge, it suffices to show that $\mathcal{D}$ entails the following, which obviously holds:
$\text{Exec}(s) \wedge \text{Poss}(push(i, m, n), s) \supset$
$\text{Know}(i, \text{Poss}(push(i, m, m - 1), now) \supset$
$on(m - 1, do(push(i, m, m - 1), now)), s) \wedge$
$\text{Know}(i, \text{Poss}(push(i, m, m), now) \supset$
$on(m, do(push(i, m, m), now)), s) \wedge$
$\text{Know}(i, \text{Poss}(push(i, m, m + 1), now) \supset$
$on(m + 1, do(push(i, m, m + 1), now)), s)$.

Definition 4.2 (KAP and MPAP relations)

1. $KAP(i, a_2, a_1, a, s) \doteq$
   $\exists s_1.K(i, s_1, s, s) \wedge Poss(a_2, s_1) \wedge KA(i, a_2, a_1, a, s)$;

2. $MPAP(i, a_2, a_1, a, s) \doteq KAP(i, a_2, a_1, a, s) \wedge$
$\forall a_3. KAP(i, a_3, a_1, a, s) \supset A(i, a_2, a_3, a, s)$.

Intuitively, MPAP(i, a_2, a_1, a, s) holds iff a_2 is a most plausible one among those actions which are comparable to a_1 and possible in some situation comparable to s. When this is the case, we say that agent i considers a_2 MPAP to a_1 when action a is performed in situation s.

Theorem 4.4 (Reduction law for belief) *$\mathcal{D}$ entails:*
$Exec(s) \wedge Poss(a, s) \supset$
$\{Bel(i, \phi, do(a, s)) \equiv \forall a_1[MPAP(i, a_1, a, a, s) \supset$
$ConBel(i, Poss(a_1, now_1) \supset \phi(do(a_1, now_1)),$
$\exists a_2(MPAP(i, a_2, a, a, s) \wedge Poss(a_2, now)), s)]\}$.

This says: after action a is performed in situation s, agent i believes ϕ iff in s, for any action a_1 MPAP to a, she believes that if a_1 is possible then ϕ holds after a_1 is done, after learning that there is an action MPAP to a and possible now.

5 Postulate soundness

In this section, we examine the extent to which our framework satisfies the AGM, DP and KM postulates. Since these postulates only concern objective formulas, this section focuses on knowledge and belief of objective formulas.

We begin with belief revision, and consider noisy sensing actions which may return false information. We first formalize noisy sensing actions in our framework. We introduce an epistemic action $ns_\phi(i, e)$, which tells agent i that ϕ holds in the current situation, and if e is 1, this information is true, and false otherwise. This is axiomatized as follows:

- Poss$(ns_\phi(i, e), s) \equiv (e = 1 \equiv \phi(s))$;
- $A(j, a_2, a_1, ns_\phi(i, e), s) \equiv a_1 = a_2 \vee \exists e_1, e_2.$
  $a_1 = ns_\phi(i, e_1) \wedge a_2 = ns_\phi(i, e_2) \wedge (j = i \supset e_2 = 1)$.

The A axiom says that when action $ns_\phi(i, e)$ is executed, agent i considers truth of ϕ more plausible than falsity of ϕ, and agent $j \neq i$ considers them equally plausible.

Then a noisy sensing action which tells agent i that ϕ holds is defined as the nondeterministic choice of two actions $ns_\phi(i, 0)$ and $ns_\phi(i, 1)$, as follows:

Definition 5.1 $nns_\phi(i) \doteq ns_\phi(i, 0) \mid ns_\phi(i, 1)$.

Proposition 5.1 *Let ϕ and ψ be objective. $\mathcal{D}$ entails:*

1. $Exec(s) \wedge Do(nns_\phi(i), s, s') \wedge \neg Know(i, \neg\phi, s) \supset$
   $Bel(i, \phi, s')$;
2. $Exec(s) \wedge Do(nns_\phi(i), s, s') \wedge Know(i, \neg\phi, s) \supset$
   $[Bel(i, \psi, s) \equiv Bel(i, \psi, s')]$.

This says: when agent i is told ϕ, if she does not know $\neg\phi$, she believes ϕ, otherwise there won't be any change in her belief about objective formulas.

Shapiro *et al.* (2011) integrated belief revision into the situation calculus. Here we follow their approach in defining belief sets and belief operators. First, given a model M of $\mathcal{D}$, we define the knowledge set (resp. belief set) of agent i in situation s as the set of objective formulas agent i knows (resp. believes) in situation s:

Definition 5.2 (Knowledge and belief set)

1. $KSet(i, s) \doteq \{\psi \in \mathcal{L}_{obj} \mid M \models Know(i, \psi, s)\}$;
2. $BSet(i, s) \doteq \{\psi \in \mathcal{L}_{obj} \mid M \models Bel(i, \psi, s)\}$.

We define belief revision and belief expansion as follows:

Definition 5.3 (Belief revision) $BSet(i, s * \phi) \doteq$
$\{\psi \in \mathcal{L}_{obj} \mid M \models \forall s'. Do(nns_\phi(i), s, s') \supset Bel(i, \psi, s')\}$.

Intuitively, BSet$(i, s * \phi)$ is the belief set of agent i after being told ϕ in situation s.

Definition 5.4 (Belief expansion)
$BSet(i, s + \phi) \doteq \{\psi \in \mathcal{L}_{obj} \mid M \models Bel(i, (\phi \supset \psi), s)\}$.

So BSet$(i, s+\phi)$ is the set of formulas believed to be implied by ϕ in situation s.

Recall that (AGM2) states that the new information ϕ should always be included in the new belief set. However, by Proposition 5.1, when an agent knows $\neg\phi$, she wouldn't believe ϕ. So (AGM2) is not satisfied in our framework. Similarly, (AGM8), (DP1) and (DP2) are not satisfied, either. Here we present weak and more reasonable versions of these postulates. The weak postulates are presented in our notation, and other postulates can be similarly translated.

Definition 5.5 (Weak AGM and DP postulates)

(AGM2′) *If* $\neg\phi \notin KSet(i, s)$, $\phi \in BSet(i, s * \phi)$

(AGM8′) *If* $\neg(\phi \wedge \psi) \notin KSet(i, s)$ *and* $\neg\psi \notin BSet(i, s * \phi)$, *then* $BSet(i, (s * \phi) + \psi) \subseteq BSet(i, s * (\phi \wedge \psi))$

(DP1′) *If* $\neg\psi \notin KSet(i, s)$ *and* $\psi \models \phi$, *then* $BSet(i, (s * \phi) * \psi) = BSet(i, s * \psi)$

(DP2′) *If* $\neg\psi \notin KSet(i, s)$ *and* $\psi \models \neg\phi$, *then* $BSet(i, (s * \phi) * \psi) = BSet(i, s * \psi)$

Existing postulates only concern beliefs, while our versions deal with interaction of knowledge and belief. For example, (AGM2) says that an agent believes ϕ when she is told ϕ, while (AGM2′) says that when an agent does not know $\neg\phi$ and is told ϕ, she believes ϕ. In the presence of both knowledge and belief, our version is more reasonable than the original one. We consider the support of representation of more reasonable postulates an advantage of our framework.

Theorem 5.1 *When $*$ is defined as in Definition 5.3, for any objective formulas ϕ and ψ and executable situation s, (AGM1), (AGM2′), (AGM3)-(AGM7), (AGM8′), (DP1′), (DP2′), (DP3) and (DP4) postulates are satisfied.*

Finally, we consider belief update:

Definition 5.6 (Belief update) Let α be an update action for ϕ, that is, an action such that $M \models$ Bel$(i, \phi, do(\alpha, s))$.
BSet$(i, s \diamond_\alpha \phi) \doteq \{\psi \in \mathcal{L}_{obj} \mid M \models$ Bel$(i, \psi, do(\alpha, s))\}$.

So when α is an action whose execution makes agent i believe ϕ, BSet$(i, s \diamond_\alpha \phi)$ is the belief set of i after α is done.

Theorem 5.2 *When $\diamond_\alpha$ is defined as in Definition 5.6, for any objective formulas ϕ and ψ and executable situation s, (KM1), (KM2), (KM4)-(KM8) postulates are satisfied.*

Note that (KM3) is not satisfied because an update action for ϕ may have other effects.

6 Common knowledge and belief

In this section, we study change of common knowledge and belief in 3 multi-agent scenarios, *i.e.*, public announcement of soft facts, public sensing, and secret communication.

We begin with the definition of common knowledge and belief. In the following, we let $\mathcal{G}$ be the set of agents, and ε a subset of $\mathcal{G}$. We let CK$(\varepsilon, s_2, s_1, s)$ denote the transitive closure of $\exists i \in \varepsilon.K(i, s_2, s_1, s)$, which can be defined with a second-order formula:

Definition 6.1 (CK relation) $CK(\varepsilon, s_2, s_1, s) \doteq \forall P.$
$\forall i \in \varepsilon, u, v, w\{[K(i, u, w, s) \supset P(u, w, s)] \wedge$
$[P(u, v, s) \wedge K(i, v, w, s) \supset P(u, w, s)]\} \supset P(s_2, s_1, s)$

Similarly, we can define CMPB$(\varepsilon, s_2, s_1, s)$ as the transitive closure of $\exists i \in \varepsilon.MPB(i, s_2, s_1, s)$.

Definition 6.2 (Common knowledge and belief)

1. *Group ε commonly knows ϕ in situation pair (s_1, s):* $CKnow(\varepsilon, \phi, s_1, s) \doteq \forall s_2.CK(\varepsilon, s_2, s_1, s) \supset \phi(s_2)$;
2. *Group ε commonly believes ϕ in situation pair (s_1, s):* $CBel(\varepsilon, \phi, s_1, s) \doteq \forall s_2.CMPB(\varepsilon, s_2, s_1, s) \supset \phi(s_2)$.

Example 3 Continuing the letter example, we have that $\mathcal{D} \models$ CKnow$(\{Ann, Bob\}, \text{KW}(Ann, p), S_1)$. Again, we illustrate with the model M of $\mathcal{D}$ from Figure 1. We see that S_1 and S_1^* are all the situations reachable from S_1 by a path of K_{Ann} and K_{Bob} edges. S_1 is the only situation K_{Ann}-related to itself, so Know(Ann, p, S_1, S_1) holds. Similarly, Know$(Ann, \neg p, S_1^*, S_1)$ holds. Thus CKnow$(\{Ann, Bob\}, \text{KW}(Ann, p), S_1)$ holds.

Public announcement of soft facts

Consider the scenario of publicly announcing ϕ, which might be false, but it is more plausible that ϕ is true. We introduce an epistemic action $sf_\phi(e)$, which tells the agents that ϕ holds, and if e is 1, this information is true, and false otherwise. The axioms are:

- Poss$(sf_\phi(e), s) \equiv (e = 1 \equiv \phi(s))$;
- $A(i, a_2, a_1, sf_\phi(e), s) \equiv$
  $a_1 = a_2 \vee \exists e_1.a_1 = sf_\phi(e_1) \wedge a_2 = sf_\phi(1)$.

Then publicly announcing ϕ is defined as the nondeterministic choice of two actions:

Definition 6.3 $nsf_\phi \doteq sf_\phi(0) \mid sf_\phi(1)$.

Proposition 6.1 *Let ϕ be an objective formula. $\mathcal{D}$ entails* $Exec(s) \wedge Do(nsf_\phi, s, s') \supset$
$CKnow(\mathcal{G}, \forall i(Bel(i, \phi) \vee Know(i, \neg\phi)), s')$.

So after public announcement of soft fact ϕ, the agents commonly know that each agent believes ϕ or knows $\neg\phi$.

Public sensing

We now consider the setting that agent i accurately senses if ϕ holds; the sensing action is observable to the other agents, but the sensing result is not. We introduce an epistemic action $pn_\phi(i, e)$, which tells agent i that ϕ holds if $e = 1$ and that $\neg\phi$ holds if $e = 0$. The axioms are:

- Poss$(pn_\phi(i, e), s) \equiv (e = 1 \equiv \phi(s))$;
- $A(j, a_2, a_1, pn_\phi(i, e), s) \equiv a_1 = a_2 \vee$
  $j \neq i \wedge \exists e_1, e_2(a_1 = pn_\phi(i, e_1) \wedge a_2 = pn_\phi(i, e_2))$.

The A axiom says that $pn_\phi(i, 1)$ and $pn_\phi(i, 0)$ are not comparable to agent i but equally plausible to agent $j \neq i$.

Proposition 6.2 *Let ϕ be an objective formula. $\mathcal{D}$ entails* $Exec(s) \wedge \phi(s) \wedge s' = do(pn_\phi(i, 1), s) \supset$
$Know(i, \phi, s') \wedge CKnow(\mathcal{G}, Know(i, \phi) \vee Know(i, \neg\phi), s')$

So after agent i publicly senses ϕ with result 1, she knows ϕ, and the agents commonly know she knows whether ϕ holds.

Secret communication

Suppose that a set ε of agents are truthfully told that ϕ holds; other agents are oblivious and think that nothing happens. We use an epistemic action $sc_\phi(\varepsilon)$, with the axioms:

- Poss$(sc_\phi(\varepsilon), s) \equiv \phi(s)$.
- $A(i, a_2, a_1, sc_\phi(\varepsilon), s) \equiv a_1 = a_2 \vee i \notin \varepsilon \wedge a_1 \neq nil$

The A axiom says that agents not in ε think that nil is more plausible than any other actions, which are equally plausible. We use $\overline{\varepsilon}$ for the complement of ε.

Proposition 6.3 *Let ϕ and ψ be objective. $\mathcal{D}$ entails:* $Exec(s) \wedge \phi(s) \wedge s' = do(sc_\phi(\varepsilon), s) \supset$
$CKnow(\varepsilon, \phi, s') \wedge [CBel(\overline{\varepsilon}, \psi, s) \equiv CBel(\overline{\varepsilon}, \psi, s')]$

So after secret communication of ϕ, the agents in ε commonly know ϕ, and the common beliefs of agents not in ε about objective formulas remain unchanged.

7 An illustrating example

In this section, we use a simplified and adapted version of Levesque's Squirrels World to illustrate multi-agent knowledge and belief change. Squirrels and acorns live in a 5×5 grid. Each acorn and squirrel is located at some point on the grid, and each point can contain any number of squirrels and acorns. Squirrels can do the following actions:

1. $left(i)$: Squirrel i moves left a unit. Similarly, there are actions $right(i)$, $up(i)$ and $down(i)$.
2. $pick(i)$: Squirrel i picks up an acorn, which is possible when she is not holding an acorn and there is at least one acorn at her location.
3. $drop(i)$: Squirrel i drops the acorn she is holding.
4. $learn(i, m, n)$: Squirrel i learns that there are m acorns at her location when the actual number is n. This is a noisy sensing action, and it is possible when $m \geq 0$ and $|n - m| \leq 1$. We let $smell(i, m)$ denote $(\pi n)learn(i, m, n)$.
5. $notice(i, j, x, y)$: Squirrel i notices that squirrel j is at location (x, y), which is possible when their distance is at most one. This action is used to simulate passive sensor.

A squirrel can observe the action of another one within a distance of one, but if the action is a sensing action, the result is not observable. There are four squirrels: Nutty, Skwirl, Edgy, and Wally. Initially, there are two acorns at each point, and the squirrels' locations are as shown by the top left grid of Figure 2. Initially, the squirrels all hold no acorns, and have no belief about the number of acorns at each point, and the above is common knowledge. We use 3 ordinary fluents:

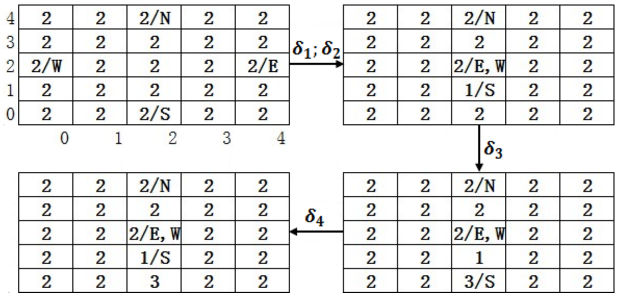

Figure 2: Evolution of Squirrels World

1. $hold(i, s)$: Squirrel i is holding an acorn in situation s.
2. $loc(i, x, y, s)$: Squirrel i is at location (x, y).
3. $acorn(x, y, n, s)$: There are n acorns at location (x, y).

We use an abbreviation $dle1(i, j, s)$, omitting the definition, which means that in situation s, the distance of squirrels i and j is at most one. For illustration, we only present some axioms of $\mathcal{D}$:

1. Poss$(learn(i, m, n), s) \equiv |m - n| \leq 1 \wedge m \geq 0 \wedge$
   $\exists x, y.loc(i, x, y, s) \wedge acorn(x, y, n, s)$
2. $hold(i, do(a, s)) \equiv a = pick(i) \vee$
   $a \neq drop(i) \wedge hold(i, s)$
3. $A(i, a_2, a_1, notice(j, k, x, y), s) \equiv a_1 = a_2 \vee$
   $(\neg dle1(i, k, s) \vee \neg dle1(i, j, s)) \wedge a_1 \neq nil$
4. $A(i, a_2, a_1, learn(j, m, n), s) \equiv a_1 = a_2 \vee$
   $\{i = j \supset \exists m_1, m_2[|m - m_1| = 1 \wedge |m - m_2| = 1 \wedge$
   $a_1 = learn(j, m, m_1) \wedge$
   $(a_2 = learn(j, m, m) \vee a_2 = learn(j, m, m_2))]\} \wedge$
   $\{dle1(i, j, s) \wedge i \neq j \supset \exists n_1, n_2, m_1, m_2$
   $[a_1 = learn(j, m_1, n_1) \wedge a_2 = learn(j, m_2, n_2)]\} \wedge$
   $\{\neg dle1(i, j, s) \supset a_1 \neq nil\}$
5. CKnow$(\{E, S, W, N\}, \forall i \neg hold(i) \wedge$
   $(\forall i, x, y, n)\neg$Bel$(i, \neg acorn(x, y, n)), S_0)$

Item 4 specifies the action plausibility order when action $learn(j, m, n)$ is executed. There are three cases. Squirrel j herself thinks that among the three alternatives $n = m$, $n = m + 1$, and $n = m - 1$, the first is more plausible than the other two, which are equally plausible. When $i \neq j$ and the distance of squirrels i and j is at most one, i considers any two learning actions of j equally plausible. When the distance of i and j is more than one, i thinks that nil is the unique most plausible action.

We now consider 4 complex actions δ_1-δ_4. Figure 2 shows the evolution of the world state as δ_1-δ_4 are performed. Then $\mathcal{D}$ entails the following formulas, where the second and the last concern higher-order knowledge and beliefs.

1. $Do(\delta_1, S_0, s) \supset$ CKnow$(\{E, S, W\}$,
   $loc(E, 2, 2) \wedge loc(W, 2, 2) \wedge loc(S, 2, 1), s)$.
   Here δ_1 is the complex action that E moves left two units, W moves right two units, S moves up, together with the associated noticing actions. For example, the action sequence $left(E)$; $left(E)$; $right(W)$ is followed by the actions $notice(W, E, 2, 2)$; $notice(E, W, 1, 2)$. The formula says that after δ_1 is performed, E, S and W commonly know the locations of each other. This is because they can see each other.
2. $Do(\delta_1; \delta_2, S_0, s) \supset$ Bel$(S, acorn(2, 1, 2), s) \wedge$
   CKnow$(\{E, S, W\}, hold(S) \wedge$
   $\exists n$Bel$(S, acorn(2, 1, n)), s)$.
   Here δ_2 is $smell(S, 3); pick(S)$. Note that the smelling action incorrectly tells S there are three acorns at her location. The formula says that after $\delta_1; \delta_2$ is performed, S believes that there are two acorns at $(2, 1)$, and E, S and W commonly know that S is holding an acorn and S has belief about the number of acorns at $(2, 1)$. This is because E, S, and W can see each other.
3. $Do(\delta_1; \delta_2; \delta_3, S_0, s) \supset$
   CBel$(\{E, W\}, loc(S, 2, 0) \wedge hold(S), s)$.
   Here δ_3 is $down(S); drop(S)$. Note that $down(S)$ is observable to E and W, but $drop(S)$ is not. So the common belief that S is at $(2, 0)$ is true, but the common belief that S is holding an acorn is false.
4. $Do(\delta_1; \delta_2; \delta_3; \delta_4, S_0, s) \supset$ Bel$(S, acorn(2, 1, 1), s) \wedge$
   Bel$(N, \forall i \neg hold(i) \wedge$
   $\forall i, x, y, n \neg$Bel$(i, \neg acorn(x, y, n)), s)$.
   Here δ_4 consists of $up(S)$, $smell(S, 1)$, together with the associated noticing actions. Since the execution of δ_1-δ_4 is not observable to N, she maintains her initial belief.

8 Conclusions

In this paper, by incorporating action priority update from DELs into the situation calculus, we have developed a general framework for reasoning about actions, both physical and epistemic, and change, not only world change but also knowledge and belief change, in multi-agent scenarios. This is simply achieved by extending the situation calculus with two plausibility relations and encoding action priority update with the successor state axiom for the B fluent. To the best of our knowledge, our work is the most general one based on the situation calculus that handles a wide range of multi-agent scenarios, knowledge and belief, noisy sensing and belief revision. Since DELs are propositional and the situation calculus is first-order, an advantage of our work is the gain of more expressiveness and succinctness in representation. As shown by the examples, in the situation calculus, the action plausibility order can be easily axiomatized, and domains like Squirrels World (and even versions with infinite grids) can be succinctly specified. Despite its prohibitive complexity, our framework can serve as the semantic foundation for implementations and applications of multi-agent knowledge and belief change. In the future, we would like to pursue a computational investigation of multi-agent knowledge and belief change via the techniques of language restriction and limited reasoning. Based on this, we would like to explore multi-agent high-level program execution and develop interesting applications of it.

Acknowledgments

We thank the anonymous reviewers for helpful comments. This work was supported by the Natural Science Foundation of China under Grant No. 61073053.

References

Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic* 50(2):510–530.

Baltag, A., and Smets, S. 2008. A qualitative theory of dynamic interactive belief revision. In *Texts in Logic and Games, Vol 3*. Amsterdam University Press. 9–58.

Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artificial Intelligence* 89:1–29.

Delgrande, J. P., and Levesque, H. J. 2012. Belief revision with sensing and fallible actions. In *Proc. KR-12*, 148–157.

Katsuno, H., and Mendelzon, A. O. 1991. On the difference between updating a knowledge base and revising it. In *Proc. KR-91*, 387–394.

Levesque, H. J.; Reiter, R.; Lespérance, Y.; Lin, F.; and Scherl, R. B. 1997. GOLOG: A logic programming language for dynamic domains. *J. Logic Programming* 31(1-3).

McCarthy, J., and Hayes, P. J. 1969. Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*, volume 4, 463–502. Edinburgh University Press.

Reiter, R. 1991. The frame problem in the situation calculus: a simple solution (sometimes) and a completeness result for goal regression. In Lifschitz, V., ed., *Aritificial Intelligence and Mathematical Thoery of Computation: Papers in Honor of John McCarthy*. Academic Press. 359–380.

Reiter, R. 2001. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. The MIT Press.

Scherl, R. B., and Levesque, H. J. 1993. The frame problem and knowledge-producing actions. In *Proc. AAAI-93*.

Scherl, R. B., and Levesque, H. J. 2003. Knowledge, action, and the frame problem. *Artificial Intelligence* 144(1-2):1–39.

Shapiro, S.; Pagnucco, M.; Lespérance, Y.; and Levesque, H. J. 2011. Iterated belief change in the situation calculus. *Artificial Intelligence* 175:165–192.

van Benthem, J. 2007. Dynamic logic for belief revision. *Journal of Applied Non-classical Logics* 17(2):129–155.

van Ditmarsch, H.; van Der Hoek, W.; and Kooi, B. P. 2007. *Dynamic epistemic logic*. Springer.