Toward an AgentSpeak(L) theory of commitment and intentional learning

Alejandro Guerra-Hernández[1], José Martín Castro-Manzano[1], and Amal El-Fallah-Seghrouchni[2]

[1] Departamento de Inteligencia Artificial
Universidad Veracruzana
Facultad de Física e Inteligencia Artificial
Sebastián Camacho No. 5, Xalapa, Ver., México, 91000
aguerra@uv.mx, e_f_s_s@hotmail.com
[2] Laboratoire d'Informatique de Paris 6
Université Pierre et Marie Curie
Avenue du Président Kennedy, Paris, France, 75016
Amal.Elfalah@lip6.fr

Abstract. This work is about the commitment strategies used by rational agents programmed in *AgentSpeak*(L) and the relationship between single-minded commitment and intentional learning. Although agent oriented languages were proposed to reduce the gap between theory and practice of Multi-Agent Systems, it has been difficult to prove BDI properties of the agents programmed in such languages. For this reason, we introduce some ideas to reason temporally about the intentional state of rational agents in order to prove what kind of commitment strategy is used by *AgentSpeak*(L) agents, based on the operational semantics of the programming language. This enables us to prove that any agent programmed in this language follows by default a limited form of single-minded commitment. Then we analyze how intentional learning can enhance this commitment strategy allowing preemptive abandon of intentions.

1 Introduction

The philosophical foundation for intentional agency is provided by the theory of practical reasoning proposed by Bratman [3]. This theory is innovative because it does not reduce intentions to some combination of beliefs and desires, but indeed it assumes that they are composed by hierarchical, partial plans. Such assumption explains better temporal aspects of practical reasoning as future intentions, persistence, reconsideration and commitment. Three kinds of reconsideration (and nonreconsideration) are identified by Bratman: nondeliberative, based on habits; deliberative, based on belief-desire reasons; and policy based. Since reconsidering the intention that α always open the question of whether α, reconsideration is closely related to commitment. We are interested on how would policy based commitment strategies be approached by intentional learning.

Different multi-modal BDI (Belief-Desire-Intention) logics [11, 13, 14] formalise practical reasoning. They are used to reason about rational agents because of their expressiveness, but not to program them. Agent oriented languages, as *AgentSpeak*(L) [10], were proposed to reduce the gap between the theory and practice of Multi-Agent Systems (MAS). They have a well defined operational semantics, but verifying intentional properties of the agents programmed in them is not evident, since they dropped intentional and time modalities for the sake of efficiency.

The main question approached in this work is what kind of commitment is used by the rational agents implemented in *AgentSpeak*(L)? In order to answer that, we develop some ideas to reason temporally about intentional operators based on the operational semantics of this language. Then we prove that these agents follow a limited form of single minded commitment [9]. Finally, we discuss the use of intentional learning [6–8] to approach full single minded commitment in a similar way, Bratman argues, we form policies of reconsideration.

The paper is organized as follows: section 2 introduces briefly the subject of the commitment strategies in the BDI logics. Section 3 presents the agent oriented language *AgentSpeak*(L) and its operational semantics. Section 4 explains the methodology used to prove BDI properties in such language. Section 5 presents the results on the temporal reasoning approach we use to inquire about the commitment strategy followed by *AgentSpeak*(L) agents and the strategy found. The role of intentional learning in commitment is introduced briefly in section 6. Section 7 closes with discussion and future work.

2 Commitment

Different computational theories of practical reasoning have been proposed to capture the main ideas proposed by Bratman. Cohen and Levesque [4] defined intentions as choice with commitment, based on the concept of persistent goals. A critical examination of this theory [12] suggested that it fails to capture important aspects of commitment, as no-infinite deferral. Alternatively, commitment has been approached as a process of maintenance and revision of intentions, relating current and future intentions. Different types of commitment strategies define different types of agents. Three of them have been extensively studied in the context of BDI_{CTL} [11], where CTL [5] is the well known branched temporal logic:

- **Blind commitment.** An agent intending that inevitably (A, for all time branches) eventually ($\Diamond$, in a branch) is the case that ϕ, inevitably maintains his intentions until (U) he actually believes ϕ (his intention is achieved):

$$\mathsf{INTEND}(\mathsf{A}\Diamond\phi) \implies \mathsf{A}(\mathsf{INTEND}(\mathsf{A}\Diamond\phi)\ \mathsf{U}\ \mathsf{BEL}(\phi)) \tag{1}$$

- **Single-minded commitment.** An agent maintains his intentions as long as he believes they are not achieved or optionally (E) eventually are achievable:

$$\mathsf{INTEND}(\mathsf{A}\Diamond\phi) \implies \mathsf{A}(\mathsf{INTEND}(\mathsf{A}\Diamond\phi)\ \mathsf{U}\ (\mathsf{BEL}(\phi) \vee \neg\mathsf{BEL}(\mathsf{E}\Diamond\phi)) \tag{2}$$

- **Open-minded commitment.** An agent maintains his intentions as long as they are not achieved or they are still desired:

$$\mathsf{INTEND}(\mathsf{A}\Diamond\phi) \implies \mathsf{A}(\mathsf{INTEND}(\mathsf{A}\Diamond\phi)\ \mathsf{U}\ (\mathsf{BEL}(\phi) \vee \neg\mathsf{DES}(\mathsf{A}\Diamond\phi))) \tag{3}$$

BDI_{CTL} is expressive enough to capture these notions of commitment. It is possible to verify if an agent system satisfies them, proving they are valid formulae in such system. Convergence to what is intended, under certain conditions, can also be proved [9]. The problem is that the multimodal logics of rational agency, as BDI_{CTL}, were conceived to reason about agents and not to program them. Agent oriented languages, as $AgentSpeak(L)$, were proposed to fill the gap between the theory and practice of MAS.

3 AgentSpeak(L)

The grammar of $AgentSpeak(L)$ [10], as defined for its interpreter Jason [2], is shown in table 1. As usual an agent ag is formed by a set of plans ps and beliefs bs. Each belief $b_i \in bs$ is a ground first-order term. Each plan $p \in ps$ has the form $trigger\ event : context \leftarrow body$. A trigger event can be any update (addition or deletion) of beliefs (at) or goals (g).The context of a plan is an atom, a negation of an atom or a conjunction of them. A non empty plan body is a sequence of actions (a), goals, or belief updates. $\top$ denotes empty elements, e.g., plan bodies, contexts, intentions. Atoms (at) can be labelled with sources. Two kinds of goals are defined, achieve goals (!) and test goals (?).

$$\begin{array}{lll}
ag & ::= bs\quad ps & \\
bs & ::= b_1 \ldots b_n & (n \geq 0) \\
 & & \\
ps & ::= p_1 \ldots p_n & (n \geq 1) \\
p & ::= te : ct \leftarrow h & \\
te & ::= +at \mid -at \mid +g \mid -g & \\
ct & ::= ct_1 \mid \top & \\
ct_1 & ::= at \mid \neg at \mid ct_1 \wedge ct_1 & \\
h & ::= h_1;\top \mid \top & \\
h_1 & ::= a \mid g \mid u \mid h_1;h_1 &
\end{array}
\qquad
\begin{array}{lll}
at & ::= P(t_1,\ldots,t_n) & (n \geq 0) \\
 & \mid P(t_1,\ldots,t_n)[s_1,\ldots,s_m] & (n \geq 0, \\
 & & m \geq 0) \\
s & ::= \texttt{percept} \mid \texttt{self} \mid id & \\
a & ::= A(t_1,\ldots,t_n) & (n \geq 0) \\
g & ::= !at \mid ?at & \\
u & ::= +b \mid -b &
\end{array}$$

Table 1. Grammar of $AgentSpeak(L)$ [2]

The operational semantics [2] of the language, is given by a set of rules that define a transition system between configurations $\langle ag, C, M, T, s\rangle$, where:

- ag is an agent program formed by a set of beliefs bs and plans ps.
- An agent circumstance C is a tuple $\langle I, E, A\rangle$, where: I is a set of intentions $\{i, i', \dots\}$, each $i \in I$ is a stack of partially instantiated plans $p \in ps$; E is a set of events $\{(te, i), (te', i'), \dots\}$, each te is a trigger event and each i is an intention (internal events) or the empty intention $\top$ (external events); and A is a set of actions to be performed in the environment.
- M is a tuple $\langle In, Out, SI\rangle$ working as a mailbox, where: In is the mailbox of the agent; Out is a list of messages to be delivered by the agent; SI is a register of suspended intentions (intentions that wait for an answer message).
- T is a tuple $\langle R, Ap, \iota, \epsilon, \rho\rangle$ that registers temporary information as follows: R is the set of relevant plans for a given event; Ap is the set of applicable plans (the subset of applicable plans which contexts are believed true); ι, ϵ, and ρ register the current intention, event and applicable plan along one cycle of execution.
- The label s indicates the current step in the reasoning cycle of the agent.

Figure 1 shows the interpreter for *AgentSpeak*(L) as a transition system. The operational semantics rules [2] define the transitions. Because of space limitations, table 2 shows only the rules that are relevant for the next section.

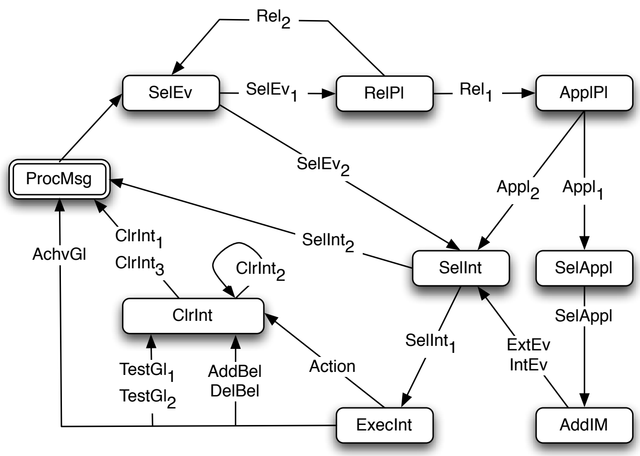

Fig. 1. The interpreter for *AgentSpeak*(L) as a transition system.

Although the operational semantics defines clearly the practical reasoning performed by an agent, it is difficult to prove intentional properties using it. This is due to the abandonment of intentional and temporal modalities. The approach followed in this paper is to define these temporal and intentional operators in terms of the operational semantics, enabling the demonstration of BDI properties.

$$(\mathbf{SelEv_1})\ \frac{\mathcal{S}_{\mathcal{E}}(C_E)=\langle te,i\rangle}{\langle ag,C,M,T,SelEv\rangle \longrightarrow \langle ag,C',M,T',RelPl\rangle} \quad \text{s.t. } \begin{array}{l} C'_E = C_E \setminus \{\langle te,i\rangle\} \\ T'_\epsilon = \langle te,i\rangle \end{array}$$

$$(\mathbf{Rel_1})\ \frac{T_\epsilon=\langle te,i\rangle, RelPlans(ag_{ps},te)\neq\{\}}{\langle ag,C,M,T,RelPl\rangle \longrightarrow \langle ag,C,M,T',AppPl\rangle} \quad \text{s.t. } T'_R = RelPlans(ag_{ps},te)$$

$$(\mathbf{Rel_2})\ \frac{RelPlans(ps,te)=\{\}}{\langle ag,C,M,T,RelPl\rangle \longrightarrow \langle ag,C,M,T,SelEv\rangle}$$

$$(\mathbf{Appl_1})\ \frac{ApplPlans(ag_{bs},T_R)\neq\{\}}{\langle ag,C,M,T,ApplPl\rangle \longrightarrow \langle ag,C,M,T',SelAppl\rangle} \quad \text{s.t. } T'_{Ap} = AppPlans(ag_{bs},T_R)$$

$$(\mathbf{SelAppl})\ \frac{\mathcal{S}_{\mathcal{O}}(T_{Ap})=(p,\theta)}{\langle ag,C,M,T,SelAppl\rangle \longrightarrow \langle ag,C,M,T',AddIM\rangle} \quad \text{s.t. } T'_\rho = (p,\theta)$$

$$(\mathbf{ExtEv})\ \frac{T_\epsilon=\langle te,\top\rangle, T_\rho=(p,\theta)}{\langle ag,C,M,T,AddIM\rangle \longrightarrow \langle ag,C',M,T,SelInt\rangle} \quad \text{s.t. } C'_I = C_I \cup \{[p\theta]\}$$

$$(\mathbf{SelInt_1})\ \frac{C_I\neq\{\}, \mathcal{S}_{\mathcal{I}}(C_I)=i}{\langle ag,C,M,T,SelInt\rangle \longrightarrow \langle ag,C,M,T',ExecInt\rangle} \quad \text{s.t. } T'_\iota = i$$

$$(\mathbf{SelInt_2})\ \frac{C_I=\{\ \}}{\langle ag,C,M,T,SelInt\rangle \longrightarrow \langle ag,C,M,T,ProcMsg\rangle}$$

$$(\mathbf{AchvGl})\ \frac{T_\iota=i[head\leftarrow !at;h]}{\langle ag,C,M,T,ExecInt\rangle \longrightarrow \langle ag,C',M,T,ProcMsg\rangle} \quad \text{s.t. } \begin{array}{l} C'_E = C_E \cup \{\langle +!at, T_\iota\rangle\} \\ C'_I = C_I \setminus \{T_\iota\} \end{array}$$

$$(\mathbf{ClrInt_1})\ \frac{T_\iota=[head\leftarrow \top]}{\langle ag,C,M,T,ClrInt\rangle \longrightarrow \langle ag,C',M,T,ProcMsg\rangle} \quad \text{s.t. } C'_I = C_I \setminus \{T_\iota\}$$

$$(\mathbf{ClrInt_2})\ \frac{T_\iota=i[head\leftarrow \top]}{\langle ag,C,M,T,ClrInt\rangle \longrightarrow \langle ag,C',M,T,ClrInt\rangle} \quad \text{s.t. } \begin{array}{l} C'_I = (C_I \setminus \{T_\iota\}) \cup \\ \{k[(head' \leftarrow h)\theta]\} \\ \text{if } i = k[head' \leftarrow g;h] \\ \text{and } g\theta = TrEv(head) \end{array}$$

$$(\mathbf{ClrInt_3})\ \frac{T_\iota\neq[head\leftarrow\top] \wedge T_\iota\neq i[head\leftarrow\top]}{\langle ag,C,M,T,ClrInt\rangle \longrightarrow \langle ag,C,M,T,ProcMsg\rangle}$$

Table 2. Some rules of the operational semantics of *AgentSpeak*(L), relevant for the definition of intentional and temporal operators.

4 Methodology

Following Bordini [1] we define the intentional modalities of BDI_{CTL} in terms of *AgentSpeak*(L) operational semantics. First the auxiliary function achievement goals is defined:

$$\begin{aligned} agoals(\top) &= \{\}, \\ agoals(i[p]) &= \begin{cases} \{at\} \cup agoals(i) & \text{if } p = +!at : ct \leftarrow h, \\ agoals(i) & \text{otherwise} \end{cases} \end{aligned}$$

which returns the set of atomic formulae (at) subject to an addition of achievement goal ($+!$) in the trigger events of the plans composing a given intention ($i[p]$ denotes an intention i which top plan is p).

Then the operators for BEL, DES, and INTEND are defined in terms of an agent ag and its circumstance C, given in a configuration:

$$\mathsf{BEL}_{\langle ag,C\rangle}(\phi) \equiv bs \models \phi. \tag{4}$$

$$\mathsf{INTEND}_{\langle ag,C\rangle}(\phi) \equiv \phi \in \bigcup_{i \in C_I} agoals(i) \vee \phi \in \bigcup_{\langle te,i\rangle \in C_E} agoals(i). \tag{5}$$

$$\mathsf{DES}_{\langle ag,C\rangle}(\phi) \equiv \langle +!\phi, i\rangle \in C_E \vee \mathsf{INTEND}_{\langle ag,C\rangle}(\phi). \tag{6}$$

These definitions are used to prove the assymetry thesis [3] (See table 3). The thesis expresses that intention-belief inconsistency is closer to irrationality than intention-belief incompletness (AT1-AT3), and the same for intention-desire (AT4-AT6) and desire-belief (AT7-AT8). Bordini and Moreira [1] prove that, under close world assumption, all *AgentSpeak*(L) agents do not satisfy the asymmetry thesis AT1, AT5, and AT7, but they satisfy the rest of them. This means that *AgentSpeak*(L) is not equivalent to any of the BDI modal systems studied previously by Rao and Georgeff [11].

Label	**Theorem**
AT1	$\models \mathsf{INTEND}_{\langle ag,C\rangle}(\phi) \implies \mathsf{BEL}_{\langle ag,C\rangle}(\phi)$
AT2	$\not\models \mathsf{INTEND}_{\langle ag,C\rangle}(\phi) \implies \mathsf{BEL}_{\langle ag,C\rangle}(\phi)$
AT3	$\not\models \mathsf{BEL}_{\langle ag,C\rangle}(\phi) \implies \mathsf{INTEND}_{\langle ag,C\rangle}(\phi)$
AT4	$\models \mathsf{INTEND}_{\langle ag,C\rangle}(\phi) \implies \mathsf{DES}_{\langle ag,C\rangle}(\phi)$
AT5	$\not\models \mathsf{INTEND}_{\langle ag,C\rangle}(\phi) \implies \mathsf{DES}_{\langle ag,C\rangle}(\phi)$
AT6	$\not\models \mathsf{DES}_{\langle ag,C\rangle}(\phi) \implies \mathsf{INTEND}_{\langle ag,C\rangle}(\phi)$
AT7	$\models \mathsf{DES}_{\langle ag,C\rangle}(\phi) \implies \mathsf{BEL}_{\langle ag,C\rangle}(\phi)$
AT8	$\not\models \mathsf{DES}_{\langle ag,C\rangle}(\phi) \implies \mathsf{BEL}_{\langle ag,C\rangle}(\phi)$
AT9	$\not\models \mathsf{BEL}_{\langle ag,C\rangle}(\phi) \implies \mathsf{DES}_{\langle ag,C\rangle}(\phi)$

Table 3. Asymmetry thesis expressed in *AgentSpeak*(L)

5 Results

The main contributions of this paper are the preliminary definition of temporal operators to reason about *AgentSpeak*(L) programs; and a demonstration that *AgentSpeak*(L) agents are not blind commited, but perform a limited single-minded commitment.

5.1 Time

Since *AgentSpeak*(L) abandons time modalities for the sake of efficiency, it is necessary to redefine them. As it is well known, temporal modalities are defined after a Kripke Structure $\langle S, R, L \rangle$ where S is a set of states, L the labeling for each state in S and R is a total relation on $S \times S$. Roughly, the states in *AgentSpeak*(L) correspond to agent configurations $\langle ag, C, M, T, s \rangle$; R is defined by the operational semantics of the system, being certainly total ($\forall k\ \exists t\ (k, t) \in R$ s.t. $k, t \in S$), as shown in figure 1. L is the label of primitive formulae valid at a given state. Validity for intentional operators is defined as in the previous section. Paths, as usual, are sequences of states (configurations) $c_0, \dots, c_n$.

Then, the definition of next is:

$$\models_{c_0} \bigcirc \alpha \ \equiv\ T_\iota = i[head \leftarrow \alpha; h] \tag{7}$$

for $\alpha \in \{a, g, u\}$ (action, goal, or belief update). c_0 is the current configuration, where formulae are evaluated. Observe that $T_\iota = \_$ always, except when an intention has been successfully selected to be executed at time t in $s =$ `SelInt`, so that at time $t+1$ the system will be in $s =$ `ExecInt` since the selection was successful (otherwise the system goes to `ProcMsg`), then α will occur in the next state of the system. It is evident that we can now define the semantics for expressions like $\bigcirc \mathsf{BEL}(\phi)$, at least for intentional updates, e.g., changes that result from the execution of intentions.

As part of our current work, we are exploring formal definitions for until:

$$\models_{c_0} \phi \ \mathsf{U}\ \psi \ \equiv\ \exists k > 0\ \models_{c_k} \psi \wedge \forall 0 < j \leq k \models_{c_j} \phi \tag{8}$$

and eventually:

$$\models_{c_0} \Diamond \phi \ \equiv\ \exists k > 0 \models_{c_k} \phi \tag{9}$$

With these definitions we already can prove some properties about the commitment strategies of *AgentSpeak*(L) agents.

5.2 Commitment strategies in *AgentSpeak*(L)

The main question at this stage of our research on commitment and intentional learning was what kind of commitment strategy is used by *AgentSpeak*(L) agents? Knowing that is important, because intentional learning seems to be irrelevant for a blindly committed agent, while it can be really useful and explanatory if agents are single or open minded. So our first step was to prove that *AgentSpeak*(L) agents do not satisfy the blind commitment axiom under no-infinite deferral.

Proposition 1. *AgentSpeak*(L) *agents satisfy the no-infinite deferral axiom* $\mathsf{INTEND}(\phi) \Rightarrow \mathsf{A}\Diamond(\neg \mathsf{INTEND}(\phi))$.

Proof. Given the definition for intend (eq. 5), the no-infinite deferral axiom expresses that if a plan p with context $+!\phi$ is adopted to form an intention, this plan is eventually retired from C_I (active intentions) and C_E (suspended intentions). Given the finite nature of the plans and providing that intentions and events are always possible to be selected in *SelInt* and *SelEv* steps, there are different paths satisfying $\mathsf{A}\Diamond\neg\mathsf{INTEND}_{\langle ag,C\rangle}(\phi)$, all of them via $\bigcirc\neg\mathsf{INTEND}_{\langle ag,C\rangle}(\phi) \equiv \bigcirc-!\phi$ (eq. 7) or $s = ClrInt$: i) Succesful execution of a plan: when the plan body becomes empty, the plan is removed from $i[p] \in C_I$ by $ClrInt_2$; ii) Succesful execution of intention: when an intention becomes empty, the full intention is removed from C_I by $ClrInt_1$; iii) Keep going execution of the intention: when the body of the plan in the top of an intention is not empty, the cycle continues and the intention will be eventually selected again by $SelInt$ arriving, if everything goes right, to one of the previous situations. If something goes wrong, a failure mechanism is activated by an event of the form $\langle -!\phi, i[p]\rangle$. Although Bordini et al. [2] only formalizes failures in finding relevant plans (Rel_2) which discards suspended intentions in C_E, other forms of failure detection have been considered in the context of intentional learning [6, 8]. By successful or failed execution of the plans every adopted intention is inevitable eventually dropped. □

Proposition 2. *AgentSpeak(L) agents do not satisfy the blind commitment axiom* $\mathsf{INTEND}(\mathsf{A}\Diamond\phi) \implies \mathsf{A}(\mathsf{INTEND}(\mathsf{A}\Diamond\phi)\ \mathsf{U}\ \mathsf{BEL}(\phi))$.

Proof. Given that the no-infinite-deferral axiom is safisfied by $AgentSpeak(L)$ agents, the blind axiom can be reduced to $\mathsf{INTEND}(\mathsf{A}\Diamond(\phi)) \implies (\mathsf{A}\Diamond(\mathsf{BEL}(\phi))$, given the blind commitment axiom (eq. 1) and assuming weak until [9]. In order to prove the proposition we will define an agent that does not satisfy the reduced blind commitment axiom. Consider an agent s.t. $ag = \langle bs, ps\rangle$ where $bs = \{\}$ and $ps = \{+b(t_1) : \top \leftarrow p(t_2).\quad +!p(t_2) : \top \leftarrow +b(t_3).\}$ at its initial configuration. Suppose that from perception of the environment a belief $b(t_1)$ is added to the $ag_{bs} = \{b(t_1)\}$. An event is generated by this belief update, so that $C_E = \{\langle +b(t_1), \top\rangle\}$. Then following the state transitions defined by the semantic rules $SelEv_1$, Rel_1, $ApplPl_1$, we obtain a configuration where $C_I = \{[+b(t_1) : \top \leftarrow !p(t_2).]\}$ and $C_E = \{\}$, Then proceeding with the reasoning steps $SelAppl$, $ExtEv$, $SelInt_1$, $AchvGl$ we obtain a configuration where $C_E\langle +!p(t_2), +b(t_1) : \top \leftarrow \top\rangle$, $C_I = \{\}$. At this moment, the agent $\mathsf{DES}_{\langle ag,C\rangle}(p(t_2))$ (eq. 6). If we apply then $SelEv_1$, Rel_1, $AppPl_1$, $SelAppl$ then we obtain a configuration where $C_I = \{[+!p(t_2) : \top \leftarrow\ +b(t_3).]\}$ and $C_E = \{\}$, where the agent $\mathsf{INTEND}_{\langle ag,C\rangle}(p(t_2))$ (eq. 5). Then proceeding with $IntEv$, $SelInt_1$, $AddBel$ then $C_E = \langle +b(t_3)[self], \top\rangle$, $ag_{bs} = \{b(t_1)\}$ and $C_I = \{[+b(t_1) : \top \leftarrow \top \ddagger +!p(t_2) : \top \leftarrow \top]$ and $bs = \{b(t_1), b(t_3)\}$. The intention about $p(t_2)$ is maintained. Observe that the plan bodies in the intention are empty, so the $ClrInt$ rules will discard the whole intention, so that $\neg\mathsf{INTEND}_{\langle ag,C\rangle}(p(t_2))$ and $\neg\mathsf{BEL}_{\langle ag,C\rangle}(p(t_2))$. $\mathsf{INTEND}(\mathsf{A}\Diamond(\phi)) \implies (\mathsf{A}\Diamond(\mathsf{BEL}(\phi))$ is not satisfied for this agent. □

In fact, our agent does not safisfy the extended blind commitment axiom (eq. 1), since the agent did not keep its intention about $p(t_2)$ until she believed it. This reasoning is similar to the demonstration of intention-belief incompleteness (AT2) for $AgentSpeak(L)$ [1].

Proposition 3. *AgentSpeak(L) agents satisfy a limited single-minded commitment* $\mathsf{INTEND}(\mathsf{A}\Diamond\phi) \implies \mathsf{A}(\mathsf{INTEND}(\mathsf{A}\Diamond\phi)\ \mathsf{U}\ \neg\mathsf{BEL}(\mathsf{E}\Diamond\phi))$.

Proof. Following the no-infinite-deferral demonstration, in the failure cases the agent will eventually satisfy $\bigcirc\neg\mathsf{INTEND}_{\langle ag,C\rangle}(\phi)$ because Rel_2 which means that for an event $\langle te, i[+!\phi : c \leftarrow h.]\rangle$ there were not relevant plans and the associated intention will be discarded, e.g., there is not a path of configurations where eventually ϕ, so that it is rational to drop $\mathsf{INTEND}_{\langle ag,C\rangle}(\phi)$. □

This is a limited form of single-mind commitment because $\neg\mathsf{BEL}(\mathsf{E}\Diamond\phi)$ is not represented explicitly by the agent. In fact, she only can not continue intending ϕ because there are no plans to do it and the full intention fails. It is also limited because of intention-belief incompleteness that can be avoided dropping the close world assumption [1]; or using the intentional and temporal definitions for studying the necessary conditions in the operational semantics and definition of the agents to warrant the expected properties of intentions, e.g., equivalence to a KD modal system.

Intentional learning provides a third alternative approach to achieve a full single-minded strategy, enabling the explicit representation of the reasons to abandon and intention in a less reactive basis, i.e., preemptive abandon of intentions.

6 Intentional learning and commitment

An agent can learn about his practical reasons. Learning examples are composed by the beliefs that supported the adoption of a plan as an intention. Because of the no-infinite deferral axiom, all examples become eventually labelled as success or failure cases. Then, first-order induction of logical decision trees is used to learn hypothesis about the reasons for successful adoption of intentions, in order to update the context of plans that have failed. If the examples of an agent do not offer enough evidence to learn, the agent can ask other agents sharing the plan for more examples. We have called this intentional learning [6–8].

For example, suppose we have an agent situated in the blocks world who perceives the state shown at figure 2, State-a; and desires to achieve the state shown at State-b. So, he forms an intention after the event $\langle +!on(b,c), \top\rangle$ using a relevant applicable plan. Suppose such plan has the form $[p1]\ +!on(X,Y) : \top \leftarrow .take(X);\ .put(X,Y)$. It means our agent is bold (or naive) about stacking objects, he believes X can be stacked on Y in any circumstance using plan $p1$. Now, suppose that after forming an intention with this plan, the agent perceives the State-c. The internal action $.put(X/a, Y/b)$ will fail, the intention will eventually fail, and the goal $on(b,c)$ will not be achieved.

An intentional learning agent can find out that the right context for the failed plan is $clear(Y) \wedge handfree(Ag)$ and modify his plan definition accordingly. Then the next time he processes the event $\langle +!on(b,c), \top\rangle$ at a state similar to State-c, it will be the case that $bs \not\models clear(c)$, and the plan will not be applicable anymore. This is better than the original non-learning approach, but does not avoid the original problem. Considering also the failure branches of the induced logic tree, in order to add beliefs as $abandon(p1) \leftarrow not(clear(Y))$, enables preemptive abandon of intentions, via a cleaning mechanism to deal events of

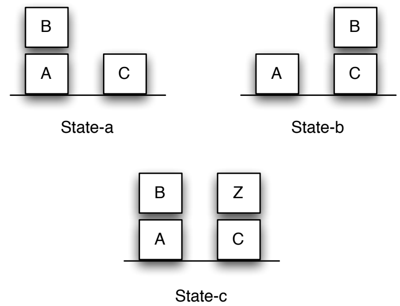

Fig. 2. Blocks world as perceived (State-a), desired (State-b), and perceived after forming an intention to put B on C, but before execute it (State-c).

the form $\langle +abandon(X), i\rangle$, as in Rel_2. We are currently experimenting with different forms of cleaning, implementing this scenario in Jason [2].

7 Discussion and future work

We have extended the methodology proposed by Bordini and Moreira [1] to reason about $AgentSpeak(L)$ agents. The extension consists in defining temporal operators based on the operational semantics of this agent oriented programming language. Then we proved that any $AgentSpeak(L)$ agent is not blindly committed, but follows a limited form of single-mind commitment. The main limitations for these agents are intention-belief incompleteness and the lack of an explicit representation for abandoning reasons. We have argued that intentional learning provides a solution for the second problem. Interestingly, modifying the context of the plans in ag_{ps} involves changes in the accessibility while forming intentions. We plan to study if intentional learning can approach intention-belief completeness in this way.

The degree of boldness and cautiousness for a given agent is something hard to define. It is well known that in dynamic environments a very cautious agent performs better than a bold one; and inversely, in static environments boldness pays better. The relevance of learning intentionally is that the right degree of cautioness is learned by the agents, instead of being established once and forever by the programmers.

Immediate future work includes to propose a $CTL_{AgentSpeak(L)}$ logic to reason about these agents. The preliminary results reported here are very encouraging in this sense. The main difficulty here is that CTL is propositional, while the content of $AgentSpeak(L)$ intentional operators is first-order, complicating the definition of L or a valuating function in the Kripke structure supporting temporal semantics. An extended $AgentSpeak(L)$ operational semantics that deals with intentional learning, for both incremental and batch inductive methods, has been proposed [8]. So, it is possible to arrive to a full theory of commitment and intentional learning using the techniques presented here.

This computational theory should consider the concept of policy based reconsideration and commitment in practical reasoning. This is relevant because it brings $AgentSpeak(L)$ closer to its philosophical foundation. But also, because policy based (non)reconsideration seems to be the more interesting of the three cases considered by Bratman. It is not so hard wired as non-deliberative cases, nor is so costly as deliberative ones.

Acknowledgments. The first and third authors are supported by Conacyt CB-2007-1 (project 78910) funding for this research. The second author is supported by Conacyt scholarship 214783.

References

1. Bordini, R.H., Moreira, Á.F.: Proving BDI properties of agent-oriented programming languages. Annals of Mathematics and Artificial Intelligence 42, 197–226 (2004)
2. Bordini, R.H., Hübner, J.F., Wooldridge, M.: Programming Multi-Agent Systems in AgentSpeak using Jason. Wiley, England (2007)
3. Bratman, M.: Intention, Plans, and Practical Reason. Harvard University Press, Cambridge (1987)
4. Cohen, P., Levesque, H.: Intention is choice with commitment. Artificial Intelligence 42(3), 213–261 (1990)
5. Emerson, A.: Temporal and modal logic. In: van Leeuwen, J., (ed.) Handbook of Theoretical Computer Science, pp. 996–1072. Elsevier, Amsterdam (1990)
6. Guerra-Hernández, A., El-Fallah-Seghrouchni, A., Soldano, H.: Learning in BDI Multi-agent Systems. In: Dix, J., Leite, J. (eds.) CLIMA IV. LNCS, vol. 3259, pp. 218–233. Springer, Heidelberg (2004)
7. Guerra-Hernández, A., Ortíz-Hernández, G.: Toward BDI sapient agents: Learning intentionally. In: Mayorga, R.V., Perlovsky, L.I. (eds.) Toward Artificial Sapience: Principles and Methods for Wise Systems, pp. 77–91. Springer, London (2008)
8. Guerra-Hernández, A., Ortíz-Hernández, G., Luna-Ramírez, W.A.: Jason smiles: Incremental BDI MAS learning. In: MICAI 2007 Special Session, IEEE, (In press).
9. Rao, A.S., Georgeff, M.P.: Modelling Rational Agents within a BDI-Architecture. In: Huhns, M.N., Singh, M.P., (eds.) Readings in Agents, pp. 317–328. Morgan Kaufmann (1998)
10. Rao, A.S.: AgentSpeak(L): BDI agents speak out in a logical computable language. In: de Velde, W.V., Perram, J.W. (eds.) MAAMAW. LNCS, vol. 1038, pp. 42–55. Springer, Heidelberg (1996)
11. Rao, A.S., Georgeff, M.P.: Decision procedures for BDI logics. Journal of Logic and Computation 8(3), pp. 293–342 (1998)
12. Singh, M.P.: A critical examination of the Cohen-Levesque Theory of Intentions. In: Proceedings of the European Conference on Artificial Intelligence (1992).
13. Singh, M.P., Rao, A.S., Georgeff, M.P.: Formal Methods in DAI: Logic-Based Representation and Reasoning. In: Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence, pp. 331–376. MIT Press, Cambridge (1999)
14. Wooldridge, M.: Reasoning about Rational Agents. MIT Press, Cambridge (2000)