Formal Properties of Intentional Reasoning

José Martín Castro-Manzano

Escuela de Filosofía
Universidad Popular Autónoma del Estado de Puebla
21 sur 1103 Barrio de Santiago, Puebla, México 72410
josemartin.castro@upaep.mx

Abstract. Intentional reasoning is a form of defeasible reasoning whose metalogical properties have not been studied completely both logically or philosophically. Our contribution in this paper is the study of such metalogical properties in order to show that intentional reasoning has the right to be called logical reasoning since it behaves as a logic, strictly speaking, as a non-monotonic logic.

Keywords: Defeasible logic, temporal logic, BDI logic.

1 Introduction

Monotony is the property of logical consequence relations that indicates that it is not possible to displace conclusions by adding new premises, in other words, that inferential power does not decrease by adding new information. However, there are situations in which our reasoning does not satisfy this property, given that some inferences are done defeasibly, that is to say, typically but not absolutely. In particular, when we reason using intentions it results quite natural to talk about some intentions that are maintained typically but not absolutely, and so, it is reasonable to conclude that intentions allow some form of defeasible reasoning [10].

A defeasible logic to model this kind of reasoning, which we call intentional, has been proposed previously [5] in terms of the BDI architecture [12], but the metalogical properties of its notion of inference have not been studied completely both logically or philosophically. Our contribution in this paper is the study of such metalogical properties. This study is important by its own nature because defeasible reasoning has certain patterns of inference and therefore the usual challenge is to provide a reasonable description of these patterns. Briefly, the idea is that if monotony is not a property of intentional reasoning and we want to give an adequate description of its notion of inference, then we must study the metalogical properties of intentional inference that occur instead of monotony. Because once monotonicity is given up, a very intuitive question arises: why should we consider intentional reasoning as an instance of a logic *bona fide*?

The paper is organized in the next way. In Section 2 we briefly expose what is understood as intentional reasoning. In Section 3 is our main contribution and finally, in Section 4 we sum up the results obtained.

2 Intentional reasoning

There are two general requirements to be checked out while developing a logical framework for defeasible reasoning: formal and material adequacy [1]. Material adequacy is about capturing an objective phenomenon. Formal adequacy has to do with the metalogical properties that a notion of logical consequence satisfies: it is argued that a well-behaved defeasible logic has to satisfy conditions of Supraclassicality, Reflexivity, Cut and Cautious Monotony [9]. During this study, due to reasons of space, we will focus mainly on the second aspect in order to argue that intentional reasoning can be modeled in a well-behaved defeasible logic.

But just to make some points about material adequacy, let us consider the next example for sake of explanation: assume there is an agent that has an intention of the form $on(X, Y) \leftarrow put(X, Y)$. This means that for such an agent to achieve $on(a, b)$ it typically has to put a on b. If we imagine such an agent is immersed in a dynamic environment, of course the agent will try to put, typically, a on b; nevertheless, a *rational* agent would only do it as long as it is *possible* –and notice how this representation of an intention does not allow us to make such distinctions. Therefore, it is natural to talk about some intentions that are maintained typically but not absolutely if we want to guarantee some level of rationality. And so, it is reasonable to conclude that intentions –in particular policy-based intentions [4]–, allow some form of defeasible reasoning [10] that must comply with some metalogical properties. But before we explore such properties, let us review some previous details.

Considering the problem of material adequacy, let us say, very quickly, that the current logical systems that deal with these topics are built in terms of what we call a bratmanian model. A bratmanian model is a model that *i)* follows the general guidelines of Bratman's theory of practical reasoning [4], *ii)* uses the BDI architecture [12] to represent data structures and *iii)* configures notions of logical consequence based on relations between intentional states. There are several logics based upon a bratmanian model, but we consider there are some problems with these systems [7,13,14]: they fail to recognize that intentional reasoning has not only temporal features but a defeasible nature.

The bratmanian model we use tries to respect this double nature and by following the general guidelines of Bratman's theory of practical reasoning distinguishes that intentions have properties that configure a notion of inference [4,5]: functional (pro-activity, inertia, admissibility), descriptive (partiality, dynamism, hierarchy) and normative (internal, external consistency and coherence). However, the metalogical properties of such notion have been barely studied. To capture this notion of inference in a formal fashion the next framework has been proposed in terms of $AgentSpeak(L)$[3] (see Appendix):

Definition 1 *(Non-monotonic intentional framework) A non-monotonic intentional framework is a tuple* $\langle B, I, F_B, F_I, \vdash, \mid\!\sim, \dashv, \sim\!\mid, \succ \rangle$ *where:*

- *B denotes the belief base.*

- *I denotes the set of intentions.*
- *$F_B \subseteq B$ denotes the basic beliefs.*
- *$F_I \subseteq I$ denotes the basic intentions.*
- *$\vdash$ and $\dashv$ are strong consequence relations.*
- *$\mid\!\sim$ and $\sim\!\mid$ are weak consequence relations.*
- *$\succ \subseteq I^2$ s.t. $\succ$ is acyclic.*

The item B denotes the beliefs, which are literals. F_B stands for the beliefs that are considered as basic; and similarly F_I stands for intentions considered as basic. Each intention $\phi \in I$ is a structure $te : ctx \leftarrow body$ where te represents the goal of the intention –so we preserve *proactivity*–, ctx a context and the rest denotes the body. When ctx or $body$ are empty we write $te : \top \leftarrow \top$ or just te. Also it is assumed that plans are *partially* instantiated.

Internal consistency is preserved by allowing the context of an intention denoted by $ctx(\phi)$, $ctx(\phi) \in B$ and by letting te be the head of the intention. So, *strong consistency* is implied by internal consistency (given that strong consistency is $ctx(\phi) \in B$). *Means-end coherence* will be implied by *admissibility* –the constraint that an agent will not consider contradictory options– and the *hierarchy* of intentions is represented by the order relation, which we require to be acyclic in order to solve conflicts between intentions. And with this framework we can arrange a notion of inference where we will say that ϕ is strongly (weakly) derivable from a sequence Δ if and only if there is a proof of $\Delta \vdash \phi$ ($\Delta \mid\!\sim \phi$). And also, that ϕ is not strongly (weakly) provable if and only if there is a proof of $\Delta \dashv \phi$ ($\Delta \sim\!\mid \phi$), where $\Delta = \langle B, I \rangle$.

2.1 The system $NBDI^{CTL}_{AS(L)}$

We start with $CTL_{AgentSpeak(L)}$ [11] as a logical tool for the formal specification. Of course, initially, the approach is similar to a BDI^{CTL} system defined after $B^{KD45}D^{KD}I^{KD}$ with the temporal operators: *next* ($\bigcirc$), *eventually* ($\Diamond$), *always* ($\Box$), *until* (U), *optional* (E), *inevitable* (A), and so on, defined after $CTL*$ [6,8].

Syntax of $BDI^{CTL}_{AS(L)}$ $CTL_{AgentSpeak(L)}$ may be seen as an instance of BDI^{CTL}. The idea is to define some BDI^{CTL} semantics in terms of $AgentSpeak(L)$ structures. So, we need a language able to express temporal and intentional states. Thus, we require in first place some way to express these features.

Definition 2 *(Syntax of $BDI^{CTL}_{AS(L)}$) If ϕ is an $AgentSpeak(L)$ atomic formula, then $\mathsf{BEL}(\phi)$, $\mathsf{DES}(\phi)$ and $\mathsf{INT}(\phi)$ are well formed formulas of $BDI^{CTL}_{AS(L)}$.*

To specify the temporal behavior we use $CTL*$ in the next way.

Definition 3 *($BDI^{CTL}_{AS(L)}$ temporal syntax) Every $BDI^{CTL}_{AS(L)}$ formula is a state formula s:*

- $s ::= \phi | s \wedge s | \neg s$
- $p ::= s | \neg p | p \wedge p | \mathsf{E}p | \mathsf{A}p | \bigcirc p | \Diamond p | \Box p | p \ \mathsf{U}\ p$

Semantics of $BDI^{CTL}_{AS(L)}$ Initially the semantics of BEL, DES and INT is adopted from [2]. So, we assume the next function:

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

which gives us the set of atomic formulas (at) attached to an achievement goal ($+!$) and $i[p]$ denotes the stack of intentions with p at the top.

Definition 4 *($BDI^{CTL}_{AS(L)}$ semantics) The operators* BEL*,* DES *and* INT *are defined in terms of an agent ag and its configuration* $\langle ag, C, M, T, s\rangle$*:*

$$\mathsf{BEL}_{\langle ag,C,M,T,s\rangle}(\phi) \equiv \phi \in bs$$

$$\mathsf{INT}_{\langle ag,C,M,T,s\rangle}(\phi) \equiv \phi \in \bigcup_{i \in C_I} agoals(i) \vee \bigcup_{\langle te,i\rangle \in C_E} agoals(i)$$

$$\mathsf{DES}_{\langle ag,C,M,T,s\rangle}(\phi) \equiv \langle +!\phi, i\rangle \in C_E \vee \mathsf{INT}(\phi)$$

where C_I denotes current intentions and C_E suspended intentions.

And now some notation: we will denote an intention ϕ with head g by $\phi[g]$. Also, a negative intention is denoted by $\phi[g^c]$, i.e., the intention ϕ with $\neg g$ as the head. The semantics of this theory will require a Kripke structure $K = \langle S, R, V\rangle$ where S is the set of agent configurations, R is an access relation defined after the transition system Γ and V is a valuation function that goes from agent configurations to true propositions in those states.

Definition 5 *Let $K = \langle S, \Gamma, V\rangle$, then:*

- *S is a set of agent configurations $c = \langle ag, C, M, T, s\rangle$.*
- *$\Gamma \subseteq S^2$ is a total relation such that for all $c \in \Gamma$ there is a $c' \in \Gamma$ s.t. $(c, c') \in \Gamma$.*
- *V is valuation s.t.:*
  - *$V_{\mathsf{BEL}}(c, \phi) = \mathsf{BEL}_c(\phi)$ where $c = \langle ag, C, M, T, s\rangle$.*
  - *$V_{\mathsf{DES}}(c, \phi) = \mathsf{DES}_c(\phi)$ where $c = \langle ag, C, M, T, s\rangle$.*
  - *$V_{\mathsf{INT}}(c, \phi) = \mathsf{INT}_c(\phi)$ where $c = \langle ag, C, M, T, s\rangle$.*
- *Paths are sequences of configurations $c_0, \ldots, c_n$ s.t. $\forall i(c_i, c_{i+1}) \in R$. We use x^i to indicate the i-th state of path x. Then:*

S1 $K, c \models \mathsf{BEL}(\phi) \Leftrightarrow \phi \in V_{\mathsf{BEL}}(c)$
S2 $K, c \models \mathsf{DES}(\phi) \Leftrightarrow \phi \in V_{\mathsf{DES}}(c)$
S3 $K, c \models \mathsf{INT}(\phi) \Leftrightarrow \phi \in V_{\mathsf{INT}}(c)$
S4 $K, c \models \mathsf{E}\phi \Leftrightarrow \exists x = c_1, \ldots \in K | K, x \models \phi$
S5 $K, c \models \mathsf{A}\phi \Leftrightarrow \forall x = c_1, \ldots \in K | K, x \models \phi$
P1 $K, c \models \phi \Leftrightarrow K, x^0 \models \phi$ *where ϕ is a state formula.*
P2 $K, c \models \bigcirc\phi \Leftrightarrow K, x^1 \models \phi$.

P3 $K, c \models \Diamond\phi \Leftrightarrow K, x^n \models \phi$ *for* $n \geq 0$
P4 $K, c \models \Box\phi \Leftrightarrow K, x^n \models \phi$ *for all* n
P5 $K, c \models \phi \mathsf{U} \psi \Leftrightarrow \exists k \geq 0$ *s.t.* $K, x^k \models \psi$ *and for all* $j, k, 0 \leq j < k | K, c^j \models \phi$ *or* $\forall j \geq 0 : K, x^j \models \phi$

A notion of inference comes in four cases: if the sequence is $\Delta \vdash \phi$, we say ϕ is strongly provable; if it is $\Delta \dashv \phi$ we say ϕ is not strongly provable. If is $\Delta \mid\!\sim \phi$ we say ϕ is weakly provable and if it is $\Delta \sim\!\mid \phi$, then ϕ is not weakly provable.

Definition 6 *(Proof) A proof of* ϕ *from* Δ *is a finite sequence of beliefs and intentions satisfying:*

1. $\Delta \vdash \phi$ *iff*
   1.1. $\Box\mathsf{A}(\mathsf{INT}(\phi))$ *or*
   1.2. $\Box\mathsf{A}(\exists\phi[g] \in F_I : \mathsf{BEL}(ctx(\phi)) \wedge \forall\psi[g'] \in body(\phi) \vdash \psi[g'])$
2. $\Delta \mid\!\sim \phi$ *iff*
   2.1. $\Delta \vdash \phi$ *or*
   2.2. $\Delta \dashv \neg\phi$ *and*
      2.2.1. $\Diamond\mathsf{E}(\mathsf{INT}(\phi) \mathsf{U} \neg\mathsf{BEL}(ctx(\phi)))$ *or*
      2.2.2. $\Diamond\mathsf{E}(\exists\phi[g] \in I : \mathsf{BEL}(ctx(\phi)) \wedge \forall\psi[g'] \in body(\phi) \mid\!\sim \psi[g'])$ *and*
         2.2.2.1. $\forall\gamma[g^c] \in I, \gamma[g^c]$ *fails at* Δ *or*
         2.2.2.2. $\psi[g'] \succ \gamma[g^c]$
3. $\Delta \dashv \phi$ *iff*
   3.1. $\Diamond\mathsf{E}(\mathsf{INT}(\neg\phi))$ *and*
   3.2. $\Diamond\mathsf{E}(\forall\phi[g] \in F_I : \neg\mathsf{BEL}(ctx(\phi)) \vee \exists\psi[g'] \in body(\phi) \dashv \psi[g'])$
4. $\Delta \sim\!\mid \phi$ *iff*
   4.1. $\Delta \dashv \phi$ *and*
   4.2. $\Delta \vdash \neg\phi$ *or*
      4.2.1. $\Box\mathsf{A}\neg(\mathsf{INT}(\phi) \mathsf{U} \neg\mathsf{BEL}(ctx(\phi)))$ *and*
      4.2.2. $\Box\mathsf{A}(\forall\phi[g] \in I : \neg\mathsf{BEL}(ctx(\phi)) \vee \exists\psi[g'] \in body(\phi) \sim\!\mid \psi[g'])$ *or*
         4.2.2.1. $\exists\gamma[g^c] \in I$ *s.t.* $\gamma[g^c]$ *succeds at* Δ *and*
         4.2.2.2. $\psi[g'] \not\succ \gamma[g^c]$

3 Formal adequacy

So once monotonicity is given up, a very intuitive question arises: why should we consider intentional reasoning as an instance of a logic *bona fide*? That is to say, why should we treat intentional reasoning as an authentic form of reasoning? We indirectly answer this question by arguing that intentional reasoning under this bratmanian model is well-behaved [9].

3.1 Consistency

A square of opposition has been suggested to depict logical relationships of consistency and coherence [5].

Proposition 1 *(Subalterns$_1$) If $\vdash \phi$ then $\mathrel{|\!\sim} \phi$.*

Proposition 2 *(Subalterns$_2$) If $\mathrel{\sim\!|} \phi$ then $\dashv \phi$.*

Proposition 3 *(Contradictories$_1$) There is no ϕ s.t. $\vdash \phi$ and $\dashv \phi$.*

Proposition 4 *(Contradictories$_2$) There is no ϕ s.t. $\mathrel{|\!\sim} \phi$ and $\mathrel{\sim\!|} \phi$.*

Proposition 5 *(Contraries) There is no ϕ s.t. $\vdash \phi$ and $\mathrel{\sim\!|} \phi$.*

Proposition 6 *(Subcontraries) For all ϕ either $\mathrel{|\!\sim} \phi$ or $\dashv \phi$.*

These propositions form the next square of opposition where c denotes contradictories, s subalterns, k contraries and r subcontraries.

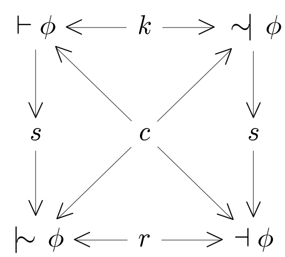

Propositions 1 and 2 represent Supraclassicality; Propositions 3 and 4 stand for Consistency while the remaining statements specify the coherence of the square, and thus, the overall coherence of the system.

3.2 Soundness

It has also been suggested that the framework is Sound with respect to its semantics [5] in such a way that:

Proposition 7 *The following relations hold:*

$$a)\ If \vdash \phi\ then \models \phi \qquad b)\ If \mathrel{|\!\sim} \phi\ then \mathrel{|\!\approx} \phi$$

$$c)\ If \dashv \phi\ then \mathrel{=\!|} \phi \qquad d)\ If \mathrel{\sim\!|} \phi\ then \mathrel{\approx\!|} \phi$$

3.3 Other formal properties

But there are other formal properties that may be used to explore and define the rationality of intentional reasoning, i.e., its good behavior. In first place, it results quite reasonable to impose Reflexivity on the consequence relation so that if $\phi \in \Delta$, then $\Delta \mid\sim \phi$.

Further, Supraclassicality requires that if ϕ follows from Δ in a monotonic way, then it must also follow according to a non-monotonic approach. Thus, in second place, we need the reasonable requirement that intentions strongly mantained have to be also weakly mantained, but no the other way around:

Proposition 8 *(Supraclassicality) If $\Delta \vdash \phi$, then $\Delta \mid\sim \phi$.*

Proof. See Proposition 1. ■

Another property, a very strong one, is Consistency Preservation. This property tells us that if some intentional set is classically consistent, then so is the set of defeasible consequences of it.

Proposition 9 *(Consistency preservation) If $\Delta \mid\sim \bot$, then $\Delta \vdash \bot$.*

Proof. Let us consider the form of the intention $\bot$. Such intention is the intention of the form $\phi \wedge \neg\phi$, which is, therefore, impossible to achieve, that is to say, for all agent runs, $\mid\sim \bot$ is never achieved. Thus $\Delta \mid\sim \bot$ is false, which makes the whole implication true. ■

A more interesting property is a form of Cautious Cut which is a form cautious transitivity. It dictates that if ϕ is a consequence of Δ, then ψ is a consequence of Δ and ϕ only if it is already a consequence of Δ alone. In other words, that adding to Δ some intentions that are already a consequence of Δ does not lead to any *increase* of information. It can be interpreted in terms of the size of a proof [1]: the idea is that the size of the proof does not affect the degree to which the assumptions support the conclusion.

Proposition 10 *(Cautious cut) If $\Delta \mid\sim \phi$ and $\Delta, \phi \mid\sim \psi$ then $\Delta \mid\sim \psi$.*

Proof. Let us start by transforming the original proposition into the next one: if $\Delta \sim\!\dashv \psi$ then it is not the case that $\Delta \mid\sim \phi$ and $\Delta, \phi \mid\sim \psi$. Further, this proposition can be transformed again: if $\Delta \sim\!\dashv \psi$ then either $\Delta \sim\!\dashv \phi$ or $\Delta, \phi \sim\!\dashv \psi$ from which, using Proposition 2, we can infer: if $\Delta \dashv \psi$ then either $\Delta \dashv \phi$ or $\Delta, \phi \dashv \psi$. Now, let us assume that $\Delta \dashv \psi$ but it is not the case that either $\Delta \dashv \phi$ or $\Delta, \phi \dashv \psi$, i.e., that $\Delta \dashv \psi$ but $\Delta \vdash \phi$ and $\Delta, \phi \vdash \psi$. Considering the expression $\Delta, \phi \vdash \psi$ we have two alternatives: either $\psi \in body(\phi)$ or $\psi \notin body(\phi)$. In the first case, given that $\Delta \vdash \phi$ then, since $\psi \in body(\phi)$ it follows that $\vdash \psi$, but that contradicts the assumption that $\Delta \dashv \psi$. In the remaining case, if $\Delta, \phi \vdash \psi$ but $\psi \notin body(\phi)$, then $\Delta \vdash \psi$, which contradicts the assumption that $\Delta \dashv \psi$. ■

The next form of monotony dictates that adding a consequence ϕ back into Δ does not lead to any *decrease* of information, that is to say, that adding implicit information is a monotonic task. Cautious Monotony then tells us that intentional reasoning is a cumulative process whenever we use implicit information.

Proposition 11 *(Cautious monotony) If $\Delta \mathrel{|\!\sim} \psi$ and $\Delta \mathrel{|\!\sim} \gamma$ then $\Delta, \psi \mathrel{|\!\sim} \gamma$.*

Proof. Let us transform the original proposition: if $\Delta, \psi \mathrel{\sim\!|} \gamma$ then it is not the case that $\Delta \mathrel{|\!\sim} \psi$ and $\Delta \mathrel{|\!\sim} \gamma$. Thus, if $\Delta, \psi \mathrel{\sim\!|} \gamma$ then either $\Delta \mathrel{\sim\!|} \psi$ or $\Delta \mathrel{\sim\!|} \gamma$, and by Proposition 2, if $\Delta, \psi \dashv \gamma$ then either $\Delta \dashv \psi$ or $\Delta \dashv \gamma$. Now, let us suppose that $\Delta, \psi \dashv \gamma$ but it is false that either $\Delta \dashv \psi$ or $\Delta \dashv \gamma$, this is to say, that $\Delta, \psi \dashv \gamma$ and $\Delta \vdash \psi$ and $\Delta \vdash \gamma$. Regarding the expression $\Delta, \psi \dashv \gamma$ we have two alternatives: either $\gamma \in \mathit{body}(\psi)$ or $\gamma \notin \mathit{body}(\psi)$. In the first case, since $\gamma \in \mathit{body}(\psi)$ and $\Delta \dashv \psi$, then $\dashv \gamma$, which contradicts the assumption that $\Delta \vdash \gamma$. On the other hand, if we consider the second alternative, $\Delta \dashv \gamma$, but that contradicts the assumption that $\Delta \vdash \gamma$. ■

This is important because when tied, Cut and Cautious Monotony tell us that if ϕ is a consequence of Δ then for any ψ, ψ is a consequence of Δ if and only if it is a consequence of Δ together with ϕ [1].

4 Conclusion

It seems reasonable to conclude that this bratmanian model of intentional reasoning can be modeled in a well-behaved defeasible logic that satisfies conditions of Consistency, Soundness, Supraclassicality, Reflexivity, Consistency Preservation, Cautious Cut and Cautious Monotony. In other words, it is plausible to conclude that intentional reasoning has the right to be called *logical reasoning* since it behaves, *mutatis mutandis*, as a logic, strictly speaking, as a non-monotonic logic.

Acknowledgements. The author would like to thank the anonymous reviewers for their helpful comments and precise corrections; and the School of Philosophy at UPAEP for all the assistance. This work has also been supported by the CONACyT scholarship 214783.

References

1. Antonelli, A.: Grounded Consequence for Defeasible Logic. Cambridge: Cambridge University Press (2005)
2. Bordini, R.H., Moreira, Á.F.: Proving BDI properties of agent-oriented programming languages. Annals of Mathematics and Artificial Intelligence 42, 197–226 (2004)
3. Bordini, R.H., Hübner, J.F., Wooldridge, M.: Programming Multi-Agent Systems in AgentSpeak using Jason. Wiley, England (2007)
4. Bratman, M.: Intention, Plans, and Practical Reason. Harvard University Press, Cambridge (1987)
5. Castro-Manzano, J.M., Barceló-Aspeitia. A.A. and Guerra-Hernández, A.: Consistency and soundness for a defeasible logic of intention. Advances in soft computing algorithms, Research in Computing Science vol. 54, (2011)
6. Clarke, E. M. Jr., Grumberg, O. and Peled. D. A.: Model Checking. MIT Press, Boston, MA., USA, (1999)

7. Cohen, P., Levesque, H.: Intention is choice with commitment. Artificial Intelligence 42(3), 213-261 (1990)
8. Emerson, A.: Temporal and modal logic. In: Handbook of Theoretical Computer Science, Elsevier Science Publishers B.V., Amsterdam, (1990)
9. Gabbay, D. M.: Theoretical foundations for nonmonotonic reasoning in expert systems. in K. Apt (ed.), Logics and Models of Concurrent Systems, Berlin and New York: Springer Verlag, pp. 439459 (1985).
10. Governatori, G., Padmanabhan, V. and Sattar, A.: A Defeasible Logic of Policy-based Intentions. In AI 2002: Advances in Artificial Intelligence, LNAI-2557. Springer Verlag (2002)
11. A. Guerra-Hernández, J. M. Castro-Manzano, A. El-Fallah-Seghrouchni.: CTLAgentSpeak(L): a Specification Language for Agent Programs. Journal of Algorithms in Cognition, Informatics and Logic, (2009)
12. 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).
13. 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).
14. Wooldridge, M.: Reasoning about Rational Agents. MIT Press, Cambridge (2000).

Appendix

AgentSpeak(L) syntax An agent *ag* is formed by a set of plans *ps* and beliefs *bs* (grounded literals). Each plan has the form $te : ctx \leftarrow h$. The context *ctx* of a plan is a literal or a conjunction of them. A non empty plan body *h* is a finite sequence of actions $A(t_1, \ldots, t_n)$, goals g (achieve ! or test ? an atomic formula $P(t_1, \ldots, t_n)$), or beliefs updates u (addition + or deletion −). $\top$ denotes empty elements, e.g., plan bodies, contexts, intentions. The trigger events *te* are updates (addition or deletion) of beliefs or goals. The syntax is shown in Table 1.

$ag ::= bs \quad ps$	$h ::= h_1; \top \mid \top$
$bs ::= b_1 \ldots b_n \ (n \geq 0)$	$h_1 ::= a \mid g \mid u \mid h_1; h_1$
$ps ::= p_1 \ldots p_n \ (n \geq 1)$	$at ::= P(t_1, \ldots, t_n) \ (n \geq 0)$
$p ::= te : ctx \leftarrow h$	$a ::= A(t_1, \ldots, t_n) \ (n \geq 0)$
$te ::= +at \mid -at \mid +g \mid -g$	$g ::= !at \mid ?at$
$ctx ::= ctx_1 \mid \top$	$u ::= +b \mid -b$
$ctx_1 ::= at \mid \neg at \mid ctx_1 \wedge ctx_1$	

Table 1. Sintax of *AgentSpeak*(L).

AgentSpeak(L) semantics The operational semantics of *AgentSpeak*(L) are defined by a transition system, as showed in Figure 1, between configurations $\langle ag, C, M, T, s \rangle$:

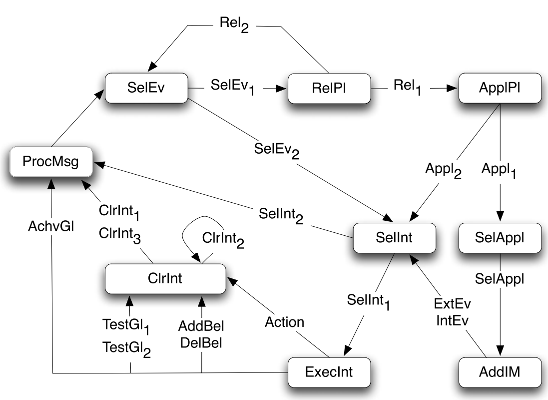

Fig. 1. The interpreter for *AgentSpeak*(L) as a transition system.

Under such semantics a run is a set $Run = \{(\sigma_i, \sigma_j) | \Gamma \vdash \sigma_i \rightarrow \sigma_j\}$ where Γ is the transition system defined by the *AgentSpeak*(L) operational semantics and σ_i, σ_j are agent configurations.