






An Epistemic Logic for Multi-agent Systems with Budget and Costs

Stefania Costantini¹ , Andrea Formisano² , and Valentina Pitoni¹ 

¹ DISIM, Università di L'Aquila, L'Aquila, Italy
stefania.costantini@univaq.it,
valentina.pitoni@graduate.univaq.it

² DMIF, Università di Udine, Udine, Italy
andrea.formisano@uniud.it

Abstract. In Artificial Intelligence, Multi-Agent Systems are able to model many kinds of collective behavior and have a wide range of application. Logic is often used to model aspects of agents' reasoning process. In this paper, we discuss social aspects of such systems. We propose a logical framework (Logic of “Inferable”) which reasons about whether a group of agents can perform an action, highlighting the concepts of action *cost* and *budget* that the group must have available in order to perform actions. The focus is on modeling the group dynamics of cooperative agents: if an agent of a group performs an action, that action to be considered as performed by the whole group, and the group can support a component agent in performing actions not affordable by that agent alone.

Keywords: Multi agents systems · Epistemic logic · Mental actions

1 Introduction

Multi-agent systems are widely employed to model societies whose members are to some extent cooperative towards each other. To achieve better results via cooperation, agents must be able to reason about their own belief states, and those of others. They must also be able to reason about what a group of agents can do, because it is often the case that a group can fulfill objectives that are out of reach for the single agent.

Many kinds of logical frameworks can be found in the literature which try to emulate cognitive aspects of human beings, also from the cooperative point of view. We propose a new logical framework (a new Logic of “Inferable”, called *L-DINF*), that draws inspiration from the concepts of Theory of Mind [20] and of Social Intelligence [21]. We consider the notion of executability of inferential actions, that may require resource consumption (and hence involve a *cost*). So, in order to execute an action the agent must possess the necessary *budget*. In our approach however, when an agent belongs to a group, if that agent does not have enough budget to perform an intended action, it may be supported by its group. So, ‘our’ agents are *aware* of themselves, of the group they belong to, and possibly of other groups. We assume that agents belonging to a group are

Supported by Action COST CA17124 “DigForASP” and by INdAM-GNCS.

cooperative. Hence, an action can be executed by the group if at least one agent therein is able to execute it, and the group can bear (in some way) the cost.

Since the seminal work of Fagin and Halpern [19], logics concerning some aspects of awareness, implicit and explicit belief have been proposed. To the best of our knowledge however, such logics make no use of concepts as ‘reasoning’ or ‘inference’. Instead, *L-DINF* provides a constructive theory of explicit beliefs, so it accounts for the perceptive and inferential steps leading from agent’s knowledge and beliefs to new beliefs, and possibly to perform physical actions. The main point however is that we consider both “executability” of actions and costs related to their execution.

Epistemic attitudes are modeled similarly to other approaches, among which we mention the dynamic theory of evidence-based beliefs [4]—that uses, as we also do, a neighborhood semantics for the notion of evidence—the sentential approach to explicit beliefs and their dynamics [22], the dynamic theory of explicit and implicit beliefs [26], and the dynamic logic of explicit beliefs and knowledge [3].

Concerning logics of inference, the seminal proposals were Velázquez-Quesada [25] and the logical system $DES4_n$ proposed by Duc [16]. We are indebted to Velázquez-Quesada concerning the idea of modeling inference steps by means of dynamic operators in the style of dynamic epistemic logic (DEL). We however emphasize the concepts of explicit belief and of background knowledge, and we introduce issues related to executability and costs. *L-DINF* is also indebted to [16], concerning the point of view that an agent reaches a certain belief state by performing inferences, and that making inferences takes time (we tackled the issue of time in previous work, discussed in [13, 14, 24]). Differently from this work however, in *L-DINF* inferential actions are represented both at the syntactic level, via dynamic operators in the DEL style, and at a semantic level as neighborhood-update operations. Moreover, *L-DINF* enables an agent to reason on executability of inferential actions.

The notion of explicit beliefs constitutes a difference between *L-DINF* and *active logics* [17, 18], besides other important differences. First, while active logics provide models of reasoning based on long-term memory and short-term memory (or working memory) like in our approach, they do not distinguish—as we do—between the notion of explicit belief and the notion of background knowledge, conceived in our case as a radically different kind of epistemic attitude. Second, *L-DINF* accounts for a variety of inferential actions that have not been explored in the active logic literature, whereas they are in our opinion very useful for inferring new beliefs. Note that these actions are *mental operation*, not physical ones. They correspond to basic operations of “mind-reading” in the sense of Theory of Mind [20]. However, the consequence of a mental operation can entail the execution of physical actions, among which “active sensing” actions, where the agent performs to check (aspects of) the state of its environment.

Section 2 introduces syntax and semantics of *L-DINF* and an example of application of our logic. In Sect. 3 we provide an axiomatization of the proposed logical system and state its soundness. The proof of strong completeness of the logic is also shown. In Sect. 4 we briefly discuss complexity and future work, and then conclude.

2 Logical Framework

L - $DINF$ is a logic which consists of a static component and a dynamic one. The static component, called L - INF , is a logic of explicit beliefs and background knowledge. The dynamic component, called L - $DINF$, extends the static one with dynamic operators capturing the consequences of the agents' inferential actions on their explicit beliefs as well as a dynamic operator capturing what an agent can conclude by performing some inferential action in its repertoire.

2.1 Syntax

In this section we provide and illustrate the syntax of the proposed logic. Let $Atm = \{p, q, \dots\}$ be a countable set of atomic propositions. By $Prop$ we denote the set of all propositional formulas, i.e. the set of all Boolean formulas built out of the set of atomic propositions Atm . A subset Atm_A of the atomic propositions represent the physical actions that an agent can perform, including "active sensing" actions (e.g., "let's check whether it rains", "let's measure the temperature"). Moreover, let Agt be a set of agents. The language of L - $DINF$, denoted by \mathcal{L}_{L-DINF} , is defined by the following grammar:

$$\begin{aligned} \varphi, \psi &::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid \mathbf{B}_i \varphi \mid \mathbf{K}_i \varphi \mid do(\phi_A) \mid do^P(\phi_A) \mid exec_G(\alpha) \mid [G : \alpha] \varphi \\ \alpha &::= \uparrow(\varphi, \psi) \mid \cap(\varphi, \psi) \mid \downarrow(\varphi, \psi) \end{aligned}$$

where p ranges over Atm and $i \in Agt$. (Other Boolean operators are defined from \neg and \wedge in the standard manner.) The language of *inferential actions* of type α is denoted by \mathcal{L}_{ACT} . Plainly, the static part L - INF of L - $DINF$, includes only those formulas not having sub-formulas of type α , namely, no inferential operation is admitted.

Notice the expression $do(\phi_A)$, where it is required that $\phi_A \in Atm_A$. This expression indicates *actual execution* of action ϕ_A , automatically recorded by the new belief $do^P(\phi_A)$ (postfix "P" standing for "past" action). In fact, do and do^P are not axiomatized, as they are realized by what has been called in [27] a *semantic attachment*, i.e., a procedure which connects an agent with its external environment in a way that is unknown at the logical level. As seen below, in general the execution of actions may have a cost. We impose the meta-constraint that a "physical" action is necessarily determined as a consequence of a mental action, thus it is the latter which bears the cost.

Before introducing a formal semantics, let us provide an intuition about the intended meaning of formulas predicating on beliefs and background knowledge. The formula $\mathbf{B}_i \varphi$ is read "the agent i explicitly believes that φ is true" or, more shortly, "agent i believes φ ". Explicit beliefs are accessible in the working memory and are the basic elements of the agents' reasoning process, according the logic of local reasoning by Fagin and Halpern [19]. In such approach agents cannot distinguish between logically equivalent formulas, i.e., if two facts φ and ψ are logically equivalent and an agent explicitly believes that φ is true, then it believes that ψ is true as well. Unlike explicit beliefs, background knowledge is assumed to satisfy *omniscience* principles, such as closure under conjunction and known implication, closure under logical consequence, and introspection. More specifically, \mathbf{K}_i is nothing but the well-known S5 modal operator often used to model/represent knowledge. The fact that background knowledge is

closed under logical consequence is justified by the fact that we conceive it as a kind of deductively closed *belief base*. We assume the background knowledge to include: facts (formulas) known by the agent from the beginning: plus facts the agent has decided to store in its long-term memory (by means of some decision mechanism not treated here) after having processed them in its working memory, as well their logical consequences. We therefore assume that background knowledge is irrevocable in the sense of being stable over time. A formula of the form $[G : \alpha] \varphi$, with $G \in 2^{Agt}$, states that “ φ holds after the inferential action α has been performed by at least one of the agents in G , and all agents in G have common knowledge about this fact”.

Remark 1. If an action is performed by an agent $i \in G$, the others agents belonging to the same group G have full visibility of this action and, therefore, as we suppose agents to be cooperative, it is as if they had performed the action themselves.

Borrowing from and extending [2], we distinguish three types of inferential actions α which allow us to capture some of the dynamic properties of explicit beliefs and background knowledge: $\vdash(\varphi, \psi)$, $\cap(\varphi, \psi)$ and $\downarrow(\varphi, \psi)$. These actions characterize the basic operations of forming explicit beliefs via inference:

- $\downarrow(\varphi, \psi)$ is the inferential action which consists in inferring ψ from φ in case φ is believed and, according to agent’s background knowledge, ψ is a logical consequence of φ . In other words, by performing this inferential action, an agent tries to retrieve from its background knowledge in long-term memory the information that φ implies ψ and, if it succeeds, it starts believing ψ ;
- $\cap(\varphi, \psi)$ is the inferential action which closes the explicit belief φ and the explicit belief ψ under conjunction. In other words, $\cap(\varphi, \psi)$ characterizes the inferential action of deducing $\varphi \wedge \psi$ from the explicit belief φ and the explicit belief ψ ;
- $\vdash(\varphi, \psi)$ is the inferential action which infers ψ from φ in case φ is believed and, according to agent’s working memory, ψ is logical consequence of φ . This last action operates directly on the working memory without retrieving anything from the background knowledge.

Remark 2. In the mental actions $\vdash(\varphi, \psi)$ and $\downarrow(\varphi, \psi)$, the formula ψ which is inferred and asserted as a new belief can be $do(\phi_A)$, which denotes the actual execution of physical action ϕ_A , where $do^P(\phi_A)$ is the belief to have done such action in the past. In fact, we assume that when inferring $do(\phi_A)$ the action is actually executed, and the corresponding belief $do^P(\phi_A)$ asserted, possibly augmented with a time-stamp. Actions are supposed to succeed by default, in case of failure a corresponding failure event will be perceived by the agent. The do^P beliefs constitute a *history* of the agent’s operation, so they might be useful for the agent to reason about its own past behavior, and/or, importantly, they may be useful to provide *explanations* to human users.

Finally, a formula of the form $exec_G(\alpha)$ expresses executability of inferential actions. It has to be read as: “ α is an inferential action that an agent in G can perform”.

As said in the Introduction, we intend to model agents which, to execute an action, may have to pay a cost, so they must have a consistent budget available. In our approach, agents belong to groups (where the smallest possible group is the single agent), and agents belonging to a group are by definition cooperative. With respect to action

execution, an action can be executed by the group if at least one agent in the group is able to execute it, and the group has the necessary budget available, sharing the cost according to some policy. In order to keep the complexity of the logic reasonable, we have not introduced costs and budget in the language.¹ In fact, by making the assumption that agents are cooperative, we also assume that they are aware of and agree with the cost-sharing policy. So, as seen below, costs and budget are coped with at the semantic level. Variants of the logic can be easily worked out, where the modalities of cost sharing are different from the one shown here, where the group members share an action's cost in equal parts. Below we indicate which are the points that should be modified to change the cost-sharing policy. Moreover, for brevity we introduce a single budget function, and thus, implicitly, a single resource to be spent. Several budget functions, each one concerning a different resource, might be plainly defined.

2.2 Semantics

Definition 1 introduces the notion of *L-INF model*, which is then used to introduce semantics of the static fragment of the logic. As before let Agt be the set of agents.

Definition 1. *A model is a tuple $M = (W, N, \mathcal{R}, E, B, C, V)$ where:*

- W is a set of objects, called worlds (or situations).²
- $\mathcal{R} = \{R_i\}_{i \in Agt}$ is a collection of equivalence relations on W : $R_i \subseteq W \times W$ for each $i \in Agt$;
- $N : Agt \times W \longrightarrow 2^{2^W}$ is a neighborhood function such that for each $i \in Agt$, each $w, v \in W$, and each $X \subseteq W$ these conditions hold:
 - (C1) if $X \in N(i, w)$ then $X \subseteq \{v \in W \mid wR_i v\}$,
 - (C2) if $wR_i v$ then $N(i, w) = N(i, v)$;
- $E : Agt \times W \longrightarrow 2^{\mathcal{L}_{ACT}}$ is an executability function such that for each $i \in Agt$ and $w, v \in W$, it holds that:
 - (E1) if $wR_i v$ then $E(i, w) = E(i, v)$;
- $B : Agt \times W \longrightarrow \mathbb{N}$ is a budget function such that for each $i \in Agt$ and $w, v \in W$, the following holds
 - (E1) if $wR_i v$ then $B(i, w) = B(i, v)$;
- $C : Agt \times \mathcal{L}_{ACT} \times W \longrightarrow \mathbb{N}$ is a cost function such that for each $i \in Agt$, $\alpha \in \mathcal{L}_{ACT}$, and $w, v \in W$, it holds that:
 - (F1) if $wR_i v$ then $C(i, \alpha, w) = C(i, \alpha, v)$;
- $V : W \longrightarrow 2^{Atm}$ is a valuation function.

¹ We intend to use this logic in practice, to formalize memory in DALI agents, where DALI is a logic-based agent-oriented programming language [5, 6, 15]. So, computational effectiveness was crucial. Assuming that agents share the cost is reasonable when agents share resources, or cooperate to a common goal, as discussed, e.g., in [7, 8].

² Concerning Definition 1, a world is just an arbitrary object. No “internal structure” is required. In Sect. 3 we will take advantage of this and define worlds as set of formulas.

To simplify the notation, let $R_i(w)$ denote the set $\{v \in W \mid wR_iv\}$, for $w \in W$. The set $R_i(w)$ identifies the situations that agent i considers possible at world w . It is the *epistemic state* of agent i at w . In cognitive terms, $R_i(w)$ can be conceived as the set of all situations that agent i can retrieve from its long-term memory and reason about.

While $R_i(w)$ concerns background knowledge, $N(i, w)$ is the set of all facts that agent i explicitly believes at world w , a fact being identified with a set of worlds. Hence, if $X \in N(i, w)$ then, the agent i has the fact X under the focus of its attention and believes it. We say that $N(i, w)$ is the explicit *belief set* of agent i at world w .

The executability of actions is determined by the function E . For an agent i , $E(i, w)$ is the set of inferential actions that agent i can execute at world w . The value $B(i, w)$ is the budget the agent has available to perform actions. Similarly, the value $C(i, \alpha, w)$ is the cost to be paid by agent i to execute the action α in the world w .

Constraint **(C1)** imposes that agent i can have explicit in its mind only facts which are compatible with its current epistemic state. Moreover, according to constraint **(C2)**, if a world v is compatible with the epistemic state of agent i at world w , then agent i should have the same explicit beliefs at w and v . In other words, if two situations are equivalent as concerns background knowledge, then they cannot be distinguished through the explicit belief set. Analogous properties are imposed by constraints **(D1)**, **(E1)**, and **(F1)**. Namely, **(D1)** imposes that agent i always knows which actions it can perform and those it cannot. **(E1)** states that agent i always knows the available budget in a world (potentially needed to perform actions). Finally, **(F1)** determines that agent i always knows how much it costs to perform an inferential action.

Truth values for formulas of $L\text{-DINF}$ are inductively defined. Given a model $M = (W, N, \mathcal{R}, E, B, C, V)$, $i \in \text{Agt}$, $G \subseteq \text{Agt}$, $w \in W$, and a formula $\varphi \in \mathcal{L}_{L\text{-INF}}$, we introduce a shorthand notation for the set of all words R_i -related to w that satisfy φ :

$$\|\varphi\|_{i,w}^M = \{v \in W : wR_iv \text{ and } M, v \models \varphi\}$$

whenever $M, v \models \varphi$ is well-defined (see below). Then, we set:

- $M, w \models p$ iff $p \in V(w)$
- $M, w \models \text{exec}_G(\alpha)$ iff there exists $i \in G$ with $\alpha \in E(i, w)$
- $M, w \models \neg\varphi$ iff $M, w \not\models \varphi$
- $M, w \models \varphi \wedge \psi$ iff $M, w \models \varphi$ and $M, w \models \psi$
- $M, w \models \mathbf{B}_i \varphi$ iff $\|\varphi\|_{i,w}^M \in N(i, w)$
- $M, w \models \mathbf{K}_i \varphi$ iff $M, v \models \varphi$ for all $v \in R_i(w)$

For any inferential action α performed by any agent i , we set:

- $M, w \models [G : \alpha]\varphi$ iff $M^{[G:\alpha]}, w \models \varphi$

where we put $M^{[G:\alpha]} = \langle W; N^{[G:\alpha]}, \mathcal{R}, E, B^{[G:\alpha]}, C, V \rangle$, representing the fact that the execution of an inferential action α affects the sets of beliefs of agent i and modifies the available budget. Such operation can add new beliefs by direct perception, by means of one inference step, or as a conjunction of previous beliefs. Hence, when introducing new beliefs (i.e., performing mental actions), the neighborhood must be extended accordingly.

A key aspect in the definition of the logic is the following, which states under which conditions, and by which agent(s), an action may be performed:

$$enabled_w(G, \alpha) \equiv_{Def} \exists j \in G (\alpha \in E(j, w) \wedge \frac{C(j, \alpha, w)}{|G|} \leq \min_{h \in G} B(h, w)).$$

This condition as defined above expresses the fact that an action is enabled when: at least an agent can perform it; and the “payment” due by each agent, obtained by dividing the action’s cost equally among all agents of the group, is within each agent’s available budget. In case more than one agent in G can execute an action, we implicitly assume the agent j performing the action is the one corresponding to the lowest possible cost. Namely, j is such that $C(j, \alpha, w) = \min_{h \in G} C(h, \alpha, w)$. This definition reflects a parsimony criterion reasonably adoptable by cooperative agents sharing a crucial resource such as, e.g., energy or money.

Remark 3. Notice that the policy we have specified to enable the action, share the costs, and select the executor of the action is just one among many possible options. Other choices might be viable, for example, depending on the specific implementation choices of an agent system or on the characteristics of the concrete real-world application at hand. So variations of this logic can be easily defined by devising some other enabling condition and policy for cost sharing, or even by introducing differences in neighborhood update. The semantics is, in a sense, parametric w.r.t. such choice. Notice, moreover, that the definition of the enabling function basically specifies the “**role**” that agents take while concurring with their own resources to actions’ execution. Also, in case of specification of different resources, different corresponding enabling functions should be defined.

The updated neighborhood $N^{[G:\alpha]}$ is as follows.

$$N^{[G:\downarrow(\psi, \chi)]}(i, w) = \begin{cases} N(i, w) \cup \{|\chi|_{i,w}^M\} & \text{if } i \in G \text{ and } enabled_w(G, \downarrow(\psi, \chi)) \text{ and} \\ & M, w \models \mathbf{B}_i\psi \wedge \mathbf{K}_i(\psi \rightarrow \chi) \\ N(i, w) & \text{otherwise} \end{cases}$$

$$N^{[G:\cap(\psi, \chi)]}(i, w) = \begin{cases} N(i, w) \cup \{|\psi \wedge \chi|_{i,w}^M\} & \text{if } i \in G \text{ and } enabled_w(G, \cap(\psi, \chi)) \text{ and} \\ & M, w \models \mathbf{B}_i\psi \wedge \mathbf{B}_i\chi \\ N(i, w) & \text{otherwise} \end{cases}$$

$$N^{[G:\uparrow(\psi, \chi)]}(i, w) = \begin{cases} N(i, w) \cup \{|\chi|_{i,w}^M\} & \text{if } i \in G \text{ and } enabled_w(G, \uparrow(\psi, \chi)) \text{ and} \\ & M, w \models \mathbf{B}_i\psi \wedge \mathbf{B}_i(\psi \rightarrow \chi) \\ N(i, w) & \text{otherwise} \end{cases}$$

Notice that after an action α has been performed by an agent $j \in G$, all agents $i \in G$ see the same update in the neighborhood. Conversely, for any agent $h \notin G$ the neighborhood remains unchanged (i.e., $N^{[G:\alpha]}(h, w) = N(h, w)$). However, even for agents in G , the neighborhood remains unchanged if the required preconditions, on explicit beliefs, knowledge, and budget, do not hold (and hence the action is not executed). Notice also that we might devise variations of the logic by making different decisions about neighborhood update to implement, for instance, partial visibility within a group.

Since each agent in G has to contribute to cover the costs of execution by consuming part of its available budget, an update of the budget function is needed. As before, for an action α , we require $enabled_w(G, \alpha)$ to hold and assume that $j \in G$ executes α . Then, depending on α , we have:

$$B^{[G:\downarrow(\psi,\chi)]}(i,w) = \begin{cases} B(i,w) - \frac{C(j,\downarrow(\psi,\chi),w)}{|G|} & \text{if } i \in G \text{ and } enabled_w(G, \downarrow(\psi,\chi)) \text{ and} \\ & M, w \models \mathbf{B}_i\psi \wedge \mathbf{K}_i(\psi \rightarrow \chi) \\ B(i,w) & \text{otherwise} \end{cases}$$

$$B^{[G:\cap(\psi,\chi)]}(i,w) = \begin{cases} B(i,w) - \frac{C(j,\cap(\psi,\chi),w)}{|G|} & \text{if } i \in G \text{ and } enabled_w(G, \cap(\psi,\chi)) \text{ and} \\ & M, w \models \mathbf{B}_i\psi \wedge \mathbf{B}_i\chi \\ B(i,w) & \text{otherwise} \end{cases}$$

$$B^{[G:\vdash(\psi,\chi)]}(i,w) = \begin{cases} B(i,w) - \frac{C(j,\vdash(\psi,\chi),w)}{|G|} & \text{if } i \in G \text{ and } enabled_w(G, \vdash(\psi,\chi)) \text{ and} \\ & M, w \models \mathbf{B}_i\psi \wedge \mathbf{B}_i(\psi \rightarrow \chi) \\ B(i,w) & \text{otherwise} \end{cases}$$

We write $\models_{L-DINF} \varphi$ to denote that $M, w \models \varphi$ holds for all worlds w of every model M .

Property 1. As consequence of previous definitions, for any set of agents G and each $i \in G$, we have the following:

- $\models_{L-INF} (\mathbf{K}_i(\varphi \rightarrow \psi)) \wedge \mathbf{B}_i\varphi \rightarrow [G : \downarrow(\varphi, \psi)] \mathbf{B}_i\psi$.
Namely, if an agent has φ among beliefs and $\mathbf{K}_i(\varphi \rightarrow \psi)$ in its background knowledge, then as a consequence of the action $\downarrow(\varphi, \psi)$ the agent starts believing ψ .
- $\models_{L-INF} (\mathbf{B}_i\varphi \wedge \mathbf{B}_i\psi) \rightarrow [G : \cap(\varphi, \psi)] \mathbf{B}_i(\varphi \wedge \psi)$.
Namely, if an agent has φ and ψ as beliefs, then as a consequence of the action $\cap(\varphi, \psi)$ the agent starts believing $\varphi \wedge \psi$.
- $\models_{L-INF} (\mathbf{B}_i(\varphi \rightarrow \psi)) \wedge \mathbf{B}_i\varphi \rightarrow [G : \vdash(\varphi, \psi)] \mathbf{B}_i\psi$.
Namely, if an agent has φ among its beliefs and $\mathbf{B}_i(\varphi \rightarrow \psi)$ in its working memory, then as a consequence of the action $\vdash(\varphi, \psi)$ the agent starts believing ψ .

Proof. Let $i \in G$, $M = \langle W, N, \mathcal{R}, E, B, C, V \rangle$, and $w \in W$.

- Let $M, w \models \mathbf{K}_i(\varphi \rightarrow \psi) \wedge \mathbf{B}_i\varphi$. We have to show that $M, w \models [G:\downarrow(\varphi, \psi)] \mathbf{B}_i\psi$ holds. This holds iff $M^{[G:\downarrow(\varphi, \psi)]}, w \models \mathbf{B}_i\psi$, with $M^{[G:\downarrow(\varphi, \psi)]} = \langle W, N^{[G:\downarrow(\varphi, \psi)]}, \mathcal{R}, E, B^{[G:\downarrow(\varphi, \psi)]}, C, V \rangle$, where $N^{[G:\downarrow(\varphi, \psi)]}(i, w) = N(i, w) \cup \{ \parallel \psi \parallel_{i,w}^M \}$, because $M, w \models (\mathbf{K}_i(\varphi \rightarrow \psi) \wedge \mathbf{B}_i\varphi)$ and $i \in G$, by hypothesis. $M^{[G:\downarrow(\varphi, \psi)]}, w \models \mathbf{B}_i\psi$ holds because $\parallel \psi \parallel_{i,w}^{M^{[G:\downarrow(\varphi, \psi)]}}$ is member of $N^{[G:\downarrow(\varphi, \psi)]}(i, w)$.
- Let $M, w \models \mathbf{B}_i\varphi \wedge \mathbf{B}_i\psi$. We have to show that $M, w \models [G:\cap(\varphi, \psi)] \mathbf{B}_i(\varphi \wedge \psi)$. This holds iff $M^{[G:\cap(\varphi, \psi)]}, w \models \mathbf{B}_i(\varphi \wedge \psi)$, with $M^{[G:\cap(\varphi, \psi)]} = \langle W, N^{[G:\cap(\varphi, \psi)]}, \mathcal{R}, E, B^{[G:\cap(\varphi, \psi)]}, C, V \rangle$ and $N^{[G:\cap(\varphi, \psi)]} = N(i, w) \cup \{ \parallel \varphi \wedge \psi \parallel_{i,w}^M \}$, because $M, w \models \mathbf{B}_i\varphi \wedge \mathbf{B}_i\psi$, by hypothesis. Then, $M^{[G:\cap(\varphi, \psi)]}, w \models \mathbf{B}_i(\varphi \wedge \psi)$ holds.
- Let $M, w \models (\mathbf{B}_i(\varphi \rightarrow \psi) \wedge \mathbf{B}_i\varphi)$. The proof that $M, w \models [G:\vdash(\varphi, \psi)] \mathbf{B}_i\psi$ follows the same line of the proof developed for the case of action $\downarrow(\varphi, \psi)$.

2.3 Problem Specification and Inference: An Example

In this section we propose an example of problem specification and inference in *L-DINF*. Note that an agent performs physical actions to interact with other agents or with the surrounding environment in consequence to some internal inference. Therefore, we consider inferential actions as a prerequisite for physical ones, and so it is inferential actions which bear costs.

Consider a group of n agents, where each agent manages a smart home, which is a prosumer (producer+consumer) of energy. The electricity is produced by solar panels during the day. The budget available for the night is the difference between energy produced and energy consumed. More energy can be bought at high cost from the outside, so agents try to avoid this extra cost. Assume that the agents are available to lend energy to others. Now, assume that an agent i would like to use some appliance (e.g., air conditioning system, washing machine, etc.) during the night, but its own budget is insufficient. Nevertheless, agent i could use the needed appliance if the group as a whole has sufficient budget. To consider a more concrete situation, let $n = 4$ and assume that in world w_1 these four agents have the following budgets to perform actions: $B(1, w_1) = 10$, $B(2, w_1) = 7$, $B(3, w_1) = 8$, and $B(4, w_1) = 20$. The physical actions any agent can perform are, e.g.,: *switch-on-airconditioning_A*, *switch-on-washing-machine_A*, *close-electric-shutter_A*.

Among the various possible inferential actions that agents might be able to do, let us, for simplicity, consider only the following ones:

$$\begin{aligned}\alpha_1 &: \downarrow(\textit{temperature-high}, \textit{do}(\textit{switch-on-airconditioning}_A)) \\ \alpha_2 &: \downarrow(\textit{dirty-clothes}, \textit{do}(\textit{switch-on-washing-machine}_A)) \\ \alpha_3 &: \downarrow(\textit{night} \wedge \textit{thieves-fear}, \textit{do}(\textit{close-electric-shutter}_A)) \\ \alpha_4 &: \cap(\textit{night}, \textit{thieves-fear})\end{aligned}$$

Assume that their costs are $C(i, \alpha_1, w) = 20$, $C(i, \alpha_2, w) = 12$, $C(i, \alpha_3, w) = 8$, $C(i, \alpha_4, w) = 1$; that $\alpha_j \in E(i, w)$ holds for each world w , each agent i , and each action α_j ; and that the knowledge base of each agent i contains the following rules:

1. $\mathbf{K}_i(\textit{temperature-high} \rightarrow \textit{do}(\textit{switch-on-airconditioning}_A))$

This rule indicates that an agent knows that if the temperature inside the house is high, it can switch on the air conditioner;

2. $\mathbf{K}_i(\textit{do}^P(\textit{switch-on-airconditioning}_A) \rightarrow \textit{do}(\textit{close-electric-shutter}_A))$

This rule indicates that if an agent knows that someone has switched on the air conditioning (past action, postfix “ P ”), it can close the electric shutter so as not to let the heat in from the outside;

3. $\mathbf{K}_i(\textit{dirty-clothes} \rightarrow \textit{do}(\textit{switch-on-washing-machine}_A))$

This rule indicates that if an agent knows that there are dirty clothes inside the washing machine, it can switch it on;

4. $\mathbf{K}_i(\textit{night} \wedge \textit{thieves-fear} \rightarrow \textit{do}(\textit{close-electric-shutter}_A))$

This rule indicates that if an agent knows that it is night and someone has the fear of thieves, it can close the electric shutter.

Assume also that the agents have the following beliefs:

$$\begin{aligned}\mathbf{B}_1(\textit{temperature-high}) \quad \mathbf{B}_2(\textit{dirty-clothes}) \quad \mathbf{B}_3(\textit{thieves-fear}) \quad \mathbf{B}_3(\textit{night}) \\ \mathbf{B}_4(\textit{temperature-high} \rightarrow \textit{do}(\textit{switch-on-airconditioning}_A))\end{aligned}$$

The latter formula—which states that if the temperature in the house is high, then agent 4 can switch on the air conditioner—, represents an inference that agent 4 may perform by exploiting its working memory (i.e., its own present beliefs). This implication allows agent 4 to infer $\mathbf{B}_4(do(\text{switch-on-airconditioning}_A))$ depending on the contents of its own working memory. In particular such inference requires the presence of the belief $\mathbf{B}_4(\text{temperature-high})$. Compare this formula with rule (1) shown earlier, as part of the knowledge base of the agent. There, the implication concerns the agent’s long-term memory and the inference would thus exploit background knowledge.

Suppose agent 1 wants to perform α_1 . It alone cannot perform α_1 , because it does not have enough budget. But, using the inferential action

$$[G : \downarrow(\text{temperature-high}, do(\text{switch-on-airconditioning}_A))],$$

with $G = \{1, 2, 3, 4\}$, the other agents can lend its part of their budgets to share the cost, so the group can perform α_1 , because $\frac{C(1, \alpha_1, w_1)}{|G|} \leq \min_{h \in G} B(h, w_1)$. Hence, $\mathbf{B}_1(do(\text{switch-on-airconditioning}_A))$ can be inferred by agent 1 and this determines the execution of the concrete physical action. Note that each agent $i \in G$ adds $\mathbf{B}_i(do(\text{switch-on-airconditioning}_A))$ to its beliefs. Indeed, the inferential action is considered as performed by the whole group and each agent of G updates its neighborhood. After the execution of the action the budget of each agent is updated (cf., Sect. 2.2) as follows: $B(1, w_2) = 5$, $B(2, w_2) = 2$, $B(3, w_2) = 3$, and $B(4, w_2) = 15$, where, for simplicity, we name w_2 the situation reached after executing the action.

Let us now consider the case in which, in such situation, agent 2 wants to perform $do(\text{switch-on-washing-machine}_A)$, enabled by the inferential action

$$\downarrow(\text{dirty-clothes}, do(\text{switch-on-washing-machine}_A)).$$

In this case, the right precondition $\mathbf{B}_2(\text{dirty-clothes})$ holds, but, even considering the entire group G , the available budgets do not satisfy the constraint $\frac{C(2, \alpha_2, w_2)}{|G|} = 3 \leq \min_{h \in G} B(h, w_2)$ (in particular, because the available budget of agent 2 is 2).

Let us, instead, assume that agent 3 wants to perform α_3 (in w_2), to enable the physical action $close-electric-shutter_A$. This cannot be done directly, because before executing the inferential action $\downarrow(\text{night} \wedge \text{thieves-fear}, do(\text{close-electric-shutter}_A))$, it has to perform the inferential action $\cap(\text{night}, \text{thieves-fear})$ in order to infer the belief $\mathbf{B}_3(\text{night} \wedge \text{thieves-fear})$. Considering its current budget, the execution of $[\{3\} : \cap(\text{night}, \text{thieves-fear})]$ can be completed (and, after that, the budget for agent 3 becomes 2). So, agent 3 obtains the belief needed as precondition to the execution of $\downarrow(\text{night} \wedge \text{thieves-fear}, do(\text{close-electric-shutter}_A))$. Nonetheless, in order to execute such action it needs the help of other agents (because its budget does not suffice), and the new belief $\mathbf{B}_3(do(\text{close-electric-shutter}_A))$ will be inferred through $[G : \downarrow(\text{night} \wedge \text{thieves-fear}, do(\text{close-electric-shutter}_A))]$. Again, all agents in G acquire the belief inferred by agent 3 and extend their belief sets, The condition on cost sharing is also satisfied for action α_3 , and the budgets after the execution become 3, 0, 0, 13, for the agents 1, 2, 3, 4, respectively. At this point, since agents 2 and 3 have exhausted their budgets, they cannot perform any other action.

The non-executability depends on the policy adopted to share action cost among agents. For instance, a policy requiring proportional sharing of costs with respect to agents' budgets, could be adopted. By applying this criterion, the execution of action α_1 in world w_1 , by agent 1 as part of G , would have generated the following budgets 6, 4, 4, 11 for the agents 1, 2, 3, 4, respectively, because agents would have contributed paying 4, 3, 4, 9, respectively (where we rounded values to the closest integer). Similarly, with a proportional sharing of costs even in the the last situation of the example, agents of G would collectively have the budget to perform more actions.

3 Axiomatization and Strong Completeness

In this section we present an axiomatization of our logic and discuss the proof of its strong completeness w.r.t. the proposed class of models.

The *L-INF* and *L-DINF* axioms and inference rules are the following:

1. $(\mathbf{K}_i \varphi \wedge \mathbf{K}_i(\varphi \rightarrow \psi)) \rightarrow \mathbf{K}_i \psi$;
2. $\mathbf{K}_i \varphi \rightarrow \varphi$;
3. $\neg \mathbf{K}_i(\varphi \wedge \neg \varphi)$;
4. $\mathbf{K}_i \varphi \rightarrow \mathbf{K}_i \mathbf{K}_i \varphi$;
5. $\neg \mathbf{K}_i \varphi \rightarrow \mathbf{K}_i \neg \mathbf{K}_i \varphi$;
6. $\mathbf{B}_i \varphi \wedge \mathbf{K}_i(\varphi \leftrightarrow \psi) \rightarrow \mathbf{B}_i \psi$;
7. $\mathbf{B}_i \varphi \rightarrow \mathbf{K}_i \mathbf{B}_i \varphi$;
8. $\frac{\varphi}{\mathbf{K}_i \varphi}$;
9. $[G : \alpha]p \leftrightarrow p$;
10. $[G : \alpha]\neg\varphi \leftrightarrow \neg[G : \alpha]\varphi$;
11. $exec_G(\alpha) \rightarrow \mathbf{K}_i(exec_G(\alpha))$;
12. $[G : \alpha](\varphi \wedge \psi) \leftrightarrow [G : \alpha]\varphi \wedge [G : \alpha]\psi$;
13. $[G : \alpha]\mathbf{K}_i \varphi \leftrightarrow \mathbf{K}_i([G : \alpha]\varphi)$;
14. $[G : \downarrow(\varphi, \psi)]\mathbf{B}_i \chi \leftrightarrow \mathbf{B}_i([G : \downarrow(\varphi, \psi)]\chi) \vee ((\mathbf{B}_i \varphi \wedge \mathbf{K}_i(\varphi \rightarrow \psi)) \wedge \mathbf{K}_i([G : \downarrow(\varphi, \psi)]\chi \leftrightarrow \psi))$;
15. $[G : \cap(\varphi, \psi)]\mathbf{B}_i \chi \leftrightarrow \mathbf{B}_i([G : \cap(\varphi, \psi)]\chi) \vee ((\mathbf{B}_i \varphi \wedge \mathbf{B}_i \psi) \wedge \mathbf{K}_i([G : \cap(\varphi, \psi)]\chi \leftrightarrow (\varphi \wedge \psi)))$;
16. $[G : \vdash(\varphi, \psi)]\mathbf{B}_i \chi \leftrightarrow \mathbf{B}_i([G : \vdash(\varphi, \psi)]\chi) \vee ((\mathbf{B}_i \varphi \wedge \mathbf{B}_i(\varphi \rightarrow \psi)) \wedge \mathbf{B}_i([G : \vdash(\varphi, \psi)]\chi \leftrightarrow \psi))$;
17. $\frac{\psi \leftrightarrow \chi}{\varphi \leftrightarrow \varphi[\psi/\chi]}$;

We write $L-DINF \vdash \varphi$ to denote that φ is a theorem of *L-DINF*. It is easy to verify that the above axiomatization is sound for the class of *L-INF* models, namely, all axioms are valid and inference rules preserve validity. In particular, soundness of axioms (14)–(16) immediately follows from the semantics of $[G : \alpha]\varphi$, for each inferential action α , as defined in Sect. 2.2. As before let Ag_t be a set of agents. For the proof that *L-INF* is strongly complete we use a standard canonical-model argument.

Definition 2. *The canonical L-INF model is a tuple $M_c = \langle W_c, N_c, \mathcal{R}_c, E_c, B_c, C_c, V_c \rangle$ where:*

- W_c is the set of all maximal consistent subsets of \mathcal{L}_{L-INF} ;

- $\mathcal{R}_c = \{R_{c,i}\}_{i \in \text{Agt}}$ is a collection of equivalence relations on W_c such that, for every $i \in \text{Agt}$ and $w, v \in W_c$, $wR_{c,i}v$ if and only if (for all φ , $\mathbf{K}_i \varphi \in w$ implies $\varphi \in v$)
- For $w \in W_c$, $\varphi \in \mathcal{L}_{L\text{-INF}}$ let $A_\varphi(i, w) = \{v \in R_{c,i}(w) \mid \varphi \in v\}$. Then, we put $N_c(i, w) = \{A_\varphi(i, w) \mid \mathbf{B}_i \varphi \in w\}$.
- $E_c : \text{Agt} \times W_c \longrightarrow 2^{\mathcal{L}_{\text{ACT}}}$ is such that for each $i \in \text{Agt}$ and $w, v \in W_c$, if $wR_{c,i}v$ then $E_c(i, w) = E_c(i, v)$;
- $B_c : \text{Agt} \times W_c \longrightarrow \mathbb{N}$ is such that for each $i \in \text{Agt}$ and $w, v \in W_c$, if $wR_{c,i}v$ then $B_c(i, w) = B_c(i, v)$;
- $C_c : \text{Agt} \times \mathcal{L}_{\text{ACT}} \times W_c \longrightarrow \mathbb{N}$ is such that for each $i \in \text{Agt}$, $\alpha \in \mathcal{L}_{\text{ACT}}$, and $w, v \in W_c$, if $wR_{c,i}v$ then $C_c(i, \alpha, w) = C_c(i, \alpha, v)$;
- $V_c : W_c \longrightarrow 2^{\text{Atm}}$ is such that $V_c(w) = \text{Atm} \cap w$.

Note that, analogously to what done before, $R_{c,i}(w)$ denotes the set $\{v \in W_c \mid wR_{c,i}v\}$, for each $i \in \text{Agt}$.

It is easy to verify that M_c is an $L\text{-INF}$ model as defined in Definition 1, since, it satisfies conditions **(C1)**,**(C2)**,**(D1)**,**(E1)**,**(F1)**. Hence, it models the axioms and the inference rules (1)–(17) introduced before. Consequently, the following properties hold too. Let $w \in W_c$, then

- given $\varphi \in \mathcal{L}_{L\text{-INF}}$, it holds that $\mathbf{K}_i \varphi \in w$ if and only if $\forall v \in W_c$ such that $wR_{c,i}v$ we have $\varphi \in v$;
- for $\varphi \in \mathcal{L}_{L\text{-INF}}$, if $\mathbf{B}_i \varphi \in w$ and $wR_{c,i}v$ then $\mathbf{B}_i \varphi \in v$;

Thus, $R_{c,i}$ -related worlds have the same knowledge and N_c -related worlds have the same beliefs. By proceeding similarly to what is done in [2] we obtain the proof of strong completeness. Let us start with some preliminary results:

Lemma 1. For all $w \in W_c$ and $\mathbf{B}_i \varphi, \mathbf{B}_i \psi \in \mathcal{L}_{L\text{-INF}}$, if $\mathbf{B}_i \varphi \in w$ but $\mathbf{B}_i \psi \notin w$, it follows that there exists $v \in R_{c,i}(w)$ such that $\varphi \in v \leftrightarrow \psi \notin v$.

Proof. Let $w \in W_c$ and φ, ψ be such that $\mathbf{B}_i \varphi \in w$ and $\mathbf{B}_i \psi \notin w$. Assume now that for every $v \in R_{c,i}(w)$ we have $\varphi \in v \wedge \psi \in v$ or $\varphi \notin v \wedge \psi \notin v$; then, from previous statements it follows that $\mathbf{K}_i(\varphi \leftrightarrow \psi) \in w$ so that by axiom (6), $\mathbf{B}_i \psi \in w$ which is a contradiction.

Lemma 2. For all $\varphi \in \mathcal{L}_{L\text{-INF}}$ and $w \in W_c$ it holds that $\varphi \in w$ iff $M_c, w \models \varphi$.

Proof. We have to prove the statement for all $\varphi \in \mathcal{L}_{L\text{-INF}}$. The proof is by induction on the structure of formulas. For instance, if $\varphi = p$ and $w \in W_c$, then $p \in w$ iff $p \in V_c(w)$ and this means that $M_c, w \models p$ by the semantics defined in Sect. 2.2. The case of formulas of the form $\mathbf{B}_i \varphi$ is the most involved: assume $\mathbf{B}_i \varphi \in w$ for $w \in W_c$. We have that $A_\varphi(i, w) = \{v \in R_{c,i}(w) \mid \varphi \in v\}$. By the definition of W_c and of $\|\cdot\|_{i,w}^M$ in Sect. 2.2, we have $A_\varphi(i, w) = \|\varphi\|_{i,w}^{M_c} \cap R_{c,i}(w)$. Hence, by the definition of $N_c(i, w)$ it follows that $\mathbf{B}_i \varphi \in w$ and then, $M_c, w \models \mathbf{B}_i \varphi$.

Suppose $\mathbf{B}_i \varphi \notin w$, so $\neg \mathbf{B}_i \varphi \in w$ and we have to prove $\|\varphi\|_{i,w}^{M_c} \cap R_{c,i}(w) \notin N_c(i, w)$. Choose $A \in N_c(i, w)$: by definition we know that $A = A_\psi(i, w)$ for some ψ with $\mathbf{B}_i \psi \in w$. By Lemma 1 there is some $v \in R_{c,i}(w)$ such that $\varphi \in v \leftrightarrow \psi \notin v$. By induction hypothesis, we obtain that either $v \in (\|\varphi\|_{i,w}^{M_c} \cap R_{c,i}(w)) \setminus A_\psi(i, w)$

or $v_I \in A_\psi(i, w) \setminus (\|\varphi\|_{i,w}^{M_c} \cap R_{c,i}(w))$ holds. Consequently, in both cases, $A_\psi(i, w) \neq \|\varphi\|_{i,w}^{M_c} \cap R_{c,i}(w)$. Thanks to the arbitrariness in the choice of A in $N_c(i, w)$ we conclude that $\|\varphi\|_{i,w}^{M_c} \cap R_{c,i}(w) \notin N_c(i, w)$. Hence $M_c, w \not\models \mathbf{B}_i \varphi$.

A crucial result states that each L - $DINF$ formula has an equivalent L - INF formula:

Lemma 3. *For all $\varphi \in \mathcal{L}_{L-DINF}$ there exists $\tilde{\varphi} \in \mathcal{L}_{L-INF}$ such that $L-DINF \vdash \varphi \leftrightarrow \tilde{\varphi}$.*

Proof. We have to prove the statement for all $\varphi \in \mathcal{L}_{L-DINF}$ but we show the proof only for $\varphi = p$, because the others are proved analogously. By the axiom (9) we have $[G : \alpha]p \leftrightarrow p$, and by rule (3) we have $\frac{[G:\alpha]p \leftrightarrow p}{\varphi \leftrightarrow \varphi[[G:\alpha]p/p]}$ which means that we can obtain $\tilde{\varphi}$ by replacing $[G : \alpha]p$ with p in φ .

The previous lemmas allow us to prove the following theorems.

Theorem 1. *L - INF is strongly complete for the class of L - INF models.*

Proof. Any consistent set φ may be extended to a maximal consistent set of formulas $w^* \in W_c$ and $M_c, w^* \models \varphi$ by Lemma 2. Then, L - INF is strongly complete for the class of L - INF models.

Theorem 2. *L - $DINF$ is strongly complete for the class of L - INF models.*

Proof. If K is a consistent set of \mathcal{L}_{L-DINF} formulas then, by Lemma 3, we can obtain the set $\tilde{K} = \{\tilde{\varphi} \mid \varphi \in K\}$, which is a consistent set of \mathcal{L}_{L-INF} formulas. By Theorem 1 $M_c, w \models \tilde{K}$. Since L - $DINF$ is sound and for each $\varphi \in K$, L - $DINF \vdash \varphi \leftrightarrow \tilde{\varphi}$, and it follows $M_c, w \models K$ then L - $DINF$ is strongly complete for the class of L - INF models.

4 Discussion and Future Work

In this paper we discussed some cognitive aspects of autonomous systems, concerning executability of actions in a group of agents, depending upon the available budget. To model these aspects we have proposed the new epistemic logic L - $DINF$, that we have shown “at work” via an example, and of which we have proved some useful properties among which strong completeness. The logic is easily extensible to accommodate kinds of resources, and kinds of agents’ “roles”, meaning capabilities of executing actions, and amounts they are required to spend according to their role.

The complexity of other logics which are based on the same principles as ours (Kripke semantics, canonical models, update of the neighborhood upon performing mental actions, proof of strong completeness via a standard canonical-model argument) has been thoroughly studied, thus, ‘mutatis mutandis’, we can borrow from there. After re-perusing those proofs we can in fact safely claim that, like in the analogous cases, the satisfiability problem is NP-complete in the single-agent case and it is, instead, PSPACE-complete in the multi-agent case.

Concerning related work, in alternating time temporal logics [23] costs appears explicitly in the language, and it is even possible to ask, e.g., what is the minimal amount of a resource that makes a given goal achievable; but, decision problems are

strictly more complex. However, in the present work we did not intend to design a logic to reason about coalitions and strategies like done, e.g., in [23], rather we meant to model the internal mental processes of an agent which is a member of a group, with a certain “role”. In this sense the two approaches are orthogonal rather than in competition. There has been a considerable amount of work on logics concerning coalitions’ strategic abilities where agents’ actions consume resources, or both produce and consume resources. For a review of this work and a discussion of the complexity of this kind of logics, the reader may refer to [1]. We have done ourselves some work on resource consumption/production, with preferences concerning which resources to spend or to save [9–12], for the single-agent case; the add-on is that we have devised a prototypical (freely available) implementation (see <http://users.dimi.uniud.it/~andrea.formisano/raspberry/>).

In future work, we mean to extend our logic so as to integrate temporal aspects, i.e., in which instant or time interval an action has been or should be performed, and how this may affect resource usage, and agent’s and group’s functioning.

References

1. Alechina, N., Demri, S., Logan, B.: Parameterised resource-bounded ATL. In: The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, 7–12 February 2020, pp. 7040–7046. AAAI Press (2020)
2. Balbiani, P., Duque, D.F., Lorini, E.: A logical theory of belief dynamics for resource-bounded agents. In: Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, AAMAS 2016, pp. 644–652. ACM (2016)
3. Balbiani, P., Fernández-Duque, D., Lorini, E.: The dynamics of epistemic attitudes in resource-bounded agents. *Studia Logica* **107**(3), 457–488 (2019)
4. van Benthem, J., Pacuit, E.: Dynamic logics of evidence-based beliefs. *Studia Logica* **99**(1–3), 61–92 (2011)
5. Costantini, S., Tocchio, A.: A logic programming language for multi-agent systems. In: Flesca, S., Greco, S., Ianni, G., Leone, N. (eds.) JELIA 2002. LNCS (LNAI), vol. 2424, pp. 1–13. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45757-7_1
6. Costantini, S., Tocchio, A.: The DALI logic programming agent-oriented language. In: Alferes, J.J., Leite, J. (eds.) JELIA 2004. LNCS (LNAI), vol. 3229, pp. 685–688. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30227-8_57
7. Costantini, S., De Gasperis, G.: Flexible goal-directed agents’ behavior via DALI mass and ASP modules. In: 2018 AAAI Spring Symposia, Stanford University, Palo Alto, California, USA, 26–28 March 2018. AAAI Press (2018)
8. Costantini, S., De Gasperis, G., Nazzicone, G.: DALI for cognitive robotics: principles and prototype implementation. In: Lierler, Y., Taha, W. (eds.) PADL 2017. LNCS, vol. 10137, pp. 152–162. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-51676-9_10
9. Costantini, S., Formisano, A.: Modeling preferences and conditional preferences on resource consumption and production in ASP. *J. Algorithms* **64**(1), 3–15 (2009). <https://doi.org/10.1016/j.jalgor.2009.02.002>
10. Costantini, S., Formisano, A.: Answer set programming with resources. *J. Log. Comput.* **20**(2), 533–571 (2010). <https://doi.org/10.1093/logcom/exp071>

11. Costantini, S., Formisano, A.: Weight constraints with preferences in ASP. In: Delgrande, J.P., Faber, W. (eds.) LPNMR 2011. LNCS (LNAI), vol. 6645, pp. 229–235. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-20895-9_24
12. Costantini, S., Formisano, A., Petturiti, D.: Extending and implementing RASP. *Fundam. Inform.* **105**(1–2), 1–33 (2010)
13. Costantini, S., Formisano, A., Pitoni, V.: Timed memory in resource-bounded agents. In: Ghidini, C., Magnini, B., Passerini, A., Traverso, P. (eds.) AI*IA 2018. LNCS (LNAI), vol. 11298, pp. 15–29. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-03840-3_2
14. Costantini, S., Pitoni, V.: Memory management in resource-bounded agents. In: Alviano, M., Greco, G., Scarcello, F. (eds.) AI*IA 2019. LNCS (LNAI), vol. 11946, pp. 46–58. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-35166-3_4
15. De Gasperis, G., Costantini, S., Nazzicone, G.: Dali multi agent systems framework. DALI GitHub Software Repository, July 2014, DALI. <http://github.com/AAAI-DISIM-UnivAQ/DALI>. <https://doi.org/10.5281/zenodo.11042>
16. Duc, H.N.: Reasoning about rational, but not logically omniscient, agents. *J. Log. Comput.* **7**(5), 633–648 (1997)
17. Elgot-Drapkin, J., Kraus, S., Miller, M., Nirkhe, M., Perlis, D.: Active logics: a unified formal approach to episodic reasoning. Technical report, UMIACS–University of Maryland (1999). cS-TR-4072
18. Elgot-Drapkin, J.J., Miller, M.I., Perlis, D.: Life on a desert island: ongoing work on real-time reasoning. In: Brown, F.M. (ed.) *The Frame Problem in Artificial Intelligence*, pp. 349–357. Morgan Kaufmann (1987)
19. Fagin, R., Halpern, J.Y.: Belief, awareness, and limited reasoning. *Artif. Intell.* **34**(1), 39–76 (1987)
20. Goldman, A., et al.: Theory of mind. In: *The Oxford Handbook of Philosophy of Cognitive Science*, vol. 1. Oxford University Press (2018)
21. Herzig, A., Lorini, E., Pearce, D.: Social intelligence. *AI Soc.* **34**(4), 689 (2019)
22. Jago, M.: Epistemic logic for rule-based agents. *Journal Logic Lang. Inf.* **18**(1), 131–158 (2009)
23. Nguyen, H.N., Alechina, N., Logan, B., Rakib, A.: Alternating-time temporal logic with resource bounds. *J. Log. Comput.* **28**(4), 631–663 (2018)
24. Pitoni, V., Costantini, S.: A temporal module for logical frameworks. In: Bogaerts, B., et al. (eds.) *Proceedings of ICLP 2019 (Technical communications)*. EPTCS, vol. 306, pp. 340–346 (2019)
25. Velázquez-Quesada, F.R.: Explicit and implicit knowledge in neighbourhood models. In: Grossi, D., Roy, O., Huang, H. (eds.) LORI 2013. LNCS, vol. 8196, pp. 239–252. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40948-6_19
26. Velázquez-Quesada, F.R.: Dynamic epistemic logic for implicit and explicit beliefs. *Journal Logic Lang. Inf.* **23**(2), 107–140 (2014)
27. Weyhrauch, R.W.: Prolegomena to a theory of mechanized formal reasoning. *Artif. Intell.* **13**(1–2), 133–170 (1980)