

Preference Management in Epistemic Logic L-DINF*

Stefania Costantini¹, Andrea Formisano² and Valentina Pitoni^{1,**}

¹*DISIM, Università di L'Aquila, Italy*

²*DMIF, Università di Udine, Italy*

Abstract

The Logic of “Inferable” L-DINF has been recently proposed as a declarative framework to formally model via epistemic logic the group dynamics of cooperative agents. In this paper, we extend the framework by introducing the possibility to have costs for execution of physical action. Such costs may require the consumption of multiple resources of various types, to be drawn from agents’ budgets. Also, we emphasize that all aspects of Multi-Agent Systems specified in L-DINF can be formalized in a modular way. In particular, concerning the execution of physical actions, dedicated modules allow the specification of a notion of equivalence for actions and a notion of agents’ preference, be used to affect action execution.

Keywords

Multi Agent Systems, Modal Logic, Epistemic Logic

1. Introduction

Nowadays, Multi-Agent Systems (MAS) has become a technology with many fields of application. In previous work [1, 2, 3] we introduced a logical framework based on an epistemic logic called *L-DINF*, which allows to represent knowledge and beliefs of agents, to reason about mental and physical actions of agents, to take into account costs of actions and budgets that are available to afford such costs. In particular, we assume that agents perform mental actions to develop reasoning as a premise to the execution of some physical action. The execution of each step in these inferences has a cost, paid by agents by drawing on a budget which is dedicated to execution of mental actions. Intuitively, in a practical scenario, we could consider this cost as an amount of energy that the agent draws from its internal battery to enable its reasoning functions.

The logical framework *L-DINF* allows the modeling of group dynamics of cooperative agents. Consequently, one can model agents that can form groups and support each other in performing collective mental actions [4, 5]. Moreover, agents can take into account preferences about performing one action instead of another [6]. The logical framework also encompasses the possibility for agents to have *roles* within their group of agents. Roles determine which actions each agent is enabled by its group to perform [7]. A mental action is considered executable if at

CILC'23: 38th Italian Conference on Computational Logic, June 21–23, 2023, Udine, Italy

* Research partially supported by Action COST CA17124 “DigForASP”, by Interdepartmental Project on AI (Strategic Plan UniUD–22-25) and by INdAM-GNCS projects CUP E55F22000270001 and CUP E53C22001930001.

** Corresponding author.

✉ stefania.costantini@univaq.it (S. Costantini); andrea.formisano@uniud.it (A. Formisano);

valentina.pitoni@univaq.it (V. Pitoni)

ORCID 0000-0002-5686-6124 (S. Costantini); 0000-0002-6755-9314 (A. Formisano)

© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

least one agent of the group can perform the action, with the approval of the group and on behalf of the group. An agent can join or leave a group whenever it wants (and, consequently, the role of an agent may change as it joins another group).

The agents of a group can share their beliefs, so that any agent can access beliefs of other agents. This ability opens up the possibility of modeling aspects of “*Theory of Mind*” [8]. For instance, an agent can maintain a version (possibly outdated) of the mental state other agents and perform inferences about such knowledge. Then, it can make predictions and formulate interpretations of other agent’s behaviours. For simplicity, in this paper we do not deal with these aspects and refer the interested reader to [9].

In this context, there are applications where agents can profit from the ability to choose the preferred physical action among a set of actions deemed equivalent. Hence, whenever the inference activity of an agent indicates that a physical action ϕ_A has to be executed, the agent can choose to perform another action ϕ'_A equivalent to ϕ_A . To model this feature we enrich the logical framework so that agents can exploit an equivalence relation among physical actions, together with evaluations of costs of actions, available budget and agents’ own preferences, in order to determine which is the best physical action to be performed. Also, we assume that the execution of each physical action may require some amounts of resources to be “used/available”. Differently from what happens for mental actions (where, for simplicity, only one type of resource, i.e., “energy”, is taken into account), we consider that each agent has different amounts of different resources available, and that each physical action may require multiple resources.

Various logics concerning implicit and explicit belief, as well as some aspects of awareness, have been proposed in the literature. We mention, among others, the seminal work [10] by Fagin & Halpern. Nevertheless, to the best of our knowledge, such proposals make no use of concepts such as ‘reasoning’ and ‘inference’. Instead, the logical framework *L-DINF* accounts for the perceptive and inferential steps leading from agent’s knowledge and beliefs to new beliefs. In this sense, it provides a constructive theory of explicit beliefs. In *L-DINF*, aspects such as “executability” of actions (both mental and physical) and costs related to their execution, can be represented and considered in reasoning activity of the agents.

Epistemic attitudes are modeled similarly to other approaches, among which we mention the dynamic theory of evidence-based beliefs [11]—that exploits, similarly to our approach, a neighborhood semantics for the notion of evidence—, the sentential approach to explicit beliefs dynamics [12], the dynamic theory of beliefs described in [13], and the dynamic logic combining explicit beliefs and knowledge proposed in [14].

Concerning logics of inference, relevant proposals are the one in [15] and the logical system $DES4_n$ described in [16]. In particular, we are indebted to [15] concerning the idea of modeling inference steps by means of dynamic operators in the style of dynamic epistemic logic (DEL). We, however, distinguish and emphasize the notions of explicit belief and background knowledge. Also, we consider issues related to executability and costs.

In developing *L-DINF* we are also indebted to [16], concerning the point of view that agents reach certain belief states by performing inferences, and that making inferences takes time (we tackled the issue of time in previous work, discussed in [1, 17, 18]). Differently from this work however, in *L-DINF* inferential actions are represented both at the syntactic level, via dynamic operators, and at a semantic level as neighborhood-updates. Moreover, *L-DINF* enables an agent to reason on executability of inferential actions.

The notion of explicit beliefs constitutes a major difference, besides others, between *L-DINF* and *active logics* [19, 20]. While active logics provide models of reasoning based on long-term memory and short-term 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. Moreover, *L-DINF* accounts for a collection of mental actions that have not been explored in the active logic literature.

In this paper, we extend the framework described in [1, 2, 3] by introducing the possibility to have costs for execution of physical actions. To meet these costs, the agents draw on resource budgets and this may involve the consumption of multiple amounts of resources of various types. Compared to the formalization of *L-DINF* appeared in previous papers, here we propose a re-engineering of the entire logical framework. This is done, both syntactically and semantically, by distinguishing between a *core* part of the framework (concerning the essential part of the logic, i.e., expressing properties of agents’ knowledge and beliefs) and a *collection of packages*. Each package can be thought as modular extensions of the core part, used to introduce a specific feature (such as preferences, costs, executability constraints, etc.) in the framework.

2. Logical Framework

L-DINF is a logic composed of a static component and a dynamic component. The first, called *L-INF*, is a logic of explicit beliefs and background knowledge. The second component extends the static one with dynamic operators that express the consequences of agents’ mental actions.

2.1. Syntax

Let $Atm = \{p, q, \dots\}$ be a countable set of atomic propositions. The set Atm_A is the set of the physical actions that agents can perform, including “active sensing” actions (e.g., “let’s check whether it rains”, “let’s measure the temperature”, etc.). Let Agt be a set of agents and Grp the set of groups of agents. Moreover, let $Res = \{r_1, \dots, r_\ell\}$ be the set of all (names of) resources.

The language of *L-DINF*, denoted by \mathcal{L}_{L-DINF} , is defined by the following grammar:

$$\begin{aligned} \varphi, \psi & ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid \mathbf{B}_i\varphi \mid \mathbf{K}_i\varphi \mid \\ & \quad do_G^P(\phi_A) \mid do_G(\phi_A) \mid can_do_G(\phi_A) \mid \\ & \quad intend_G(\phi_A) \mid exec_G(\alpha) \mid pref_do_i(\phi_A, d) \mid pref_do_G(i, \phi_A) \mid \\ & \quad [G : \alpha]\varphi \mid Cl(\phi_A, \phi'_A) \mid fCl_i(\phi_A) \\ \alpha & ::= +\varphi \mid \vdash(\varphi, \psi) \mid \cap(\varphi, \psi) \mid \downarrow(\varphi, \psi) \mid \neg(\varphi, \psi) \end{aligned}$$

where p ranges over Atm , $\phi_A, \phi'_A \in Atm_A$, $i \in Agt$, $d \in \mathbb{N}$, and $G \in Grp$. Other Boolean operators are defined from \neg and \wedge in the standard manner.¹ The language of *mental actions* of type α is denoted by \mathcal{L}_{ACT} . The static part *L-INF* of *L-DINF*, includes only those formulas not having sub-formulas of the form $[G : \alpha]\varphi$.

Before introducing the formal semantics let us briefly describe the intended informal meaning of basic formulas of *L-INF*. As mentioned, we are interested in modelling the reasoning of agents

¹For simplicity, whenever $G = \{i\}$ we will write i as subscript in place of $\{i\}$. So, for instance, we often write $exec_i(\phi_A)$ instead of $exec_{\{i\}}(\phi_A)$ and similarly for other constructs.

acting cooperatively. We consider the set of agents as partitioned in groups: each agent $i \in \text{Agt}$ always belongs to a unique group in Grp . We assume that all agents initially belong to an initial group. Any agent i , at any time, can perform a (physical) action $\text{join}_A(i, j)$, for $j \in \text{Agt}$, in order to change her group and join j 's group. The special case in which $i = j$ denotes the action that allows agent i to leave her current group and form the new singleton group $\{i\}$.

The formula $\text{intend}_i(\phi_A)$ indicates the intention of agent i to perform the physical action ϕ_A , in the sense of the BDI agent model [21]. Formulas of this form can be part of agent's knowledge base from the beginning or it can be derived later. In this paper we do not cope with the formalization of BDI, for which the reader may refer, e.g., to [22]. Hence, we will deal with intentions rather informally, also assuming that $\text{intend}_G(\phi_A)$ holds whenever all agents of group G intend to perform ϕ_A .

The formula $\text{do}_i(\phi_A)$ indicates the *actual execution* of action ϕ_A by agent, automatically recorded by the new belief $\text{do}_i^P(\phi_A)$ (postfix “P” standing for “past” action). Note that, we do not provide an axiomatization for do (and similarly for do_G , that indicates the actual execution of ϕ_A by the group of agents G). In fact, we assume that in any concrete implementation of the logical framework, do_i and do_G are realized by means of a *semantic attachment* [23], that is, a procedure which connects an agent with its external environment in a way that is unknown at the logical level. The axiomatization only concerns the relationship between doing and being enabled to do.

The expressions $\text{can_do}_i(\phi_A)$ and $\text{pref_do}_i(\phi_A, d)$ are closely related to $\text{do}_i(\phi_A)$. In particular, $\text{can_do}_i(\phi_A)$ must be seen as an enabling condition, indicating that the agent i is enabled to perform the action ϕ_A , while $\text{pref_do}_i(\phi_A, d)$ indicates the level d of preference/willingness of agent i to perform ϕ_A .

The formula $\text{pref_do}_G(i, \phi_A)$ indicates that agent i exhibits the *maximum level* of preference on performing action ϕ_A within all group members. Notice that, if a group of agents intends to perform an action ϕ_A , this will entail that the entire group intends to do ϕ_A , that will be enabled to be actually executed only if at least one agent $i \in G$ can do it, i.e., it can derive $\text{can_do}_i(\phi_A)$.

The formula $\text{Cl}(\phi_A, \phi'_A)$ denoted the equivalence of the two physical actions ϕ_A and ϕ'_A . Intuitively, this means that in the specific practical context at hand, the two actions have “something in common”, i.e., for instance, they use similar resources, perform in a similar way, can be used by an agent to obtain equivalent results, etc. Notice that the predicate Cl induces a partition of Atm_A in a collection of equivalence classes.

Agents modeled through *L-DINF* deal with two kind of memories, namely, a *working memory* used to represent beliefs, i.e., facts and formulas acquired via perceptions during an agent's operation, and a *long-term memory* used to model agent's background knowledge. Such knowledge is assumed to satisfy *omniscience* principles, such as: closure under conjunction and known implication, closure under logical consequence, and introspection.

Background knowledge of an agent i is specified by means of the modal operator \mathbf{K}_i , which is actually the usual S5 modal operator often used to model knowledge. The fact that background knowledge is closed under logical consequence is justified because we conceive it as a kind of stable and reliable *knowledge base*. The modal operator \mathbf{B}_i , instead, is used to represent the beliefs of agents i kept in i 's working memory. The contents of the working memory is determined by the mental actions i has executed (cf., Section 2.2.2). We assume the background knowledge to include: facts/formulas known by the agent from the beginning, and facts the agent

subsequently decided to store in its long-term memory (via a decision-making mechanism not covered here) after processing them in its working memory. We therefore assume that background knowledge is irrevocable, in the sense of being stable over time.

Whenever an agent wants to perform a physical action ϕ'_A , it can exploit the equivalence described by the facts of the form $Cl(\phi_A, \phi'_A)$ to execute a most convenient action ϕ_A (in terms of resources requires, preferences, etc.) drawn from the equivalence class of ϕ'_A . The formula $fCl_i(\phi_A)$ indicates that ϕ_A is the more convenient action among those in the set $\{\phi'_A \mid Cl(\phi_A, \phi'_A)\}$.

The formulas $exec_G(\alpha)$ express executability of mental actions by a group G (which is a consequence of the fact that any member of the group is able to perform the action). They have to be read as: “ α is a mental action that an agent in G can perform”.

A formula of the form $[G:\alpha]\varphi$, where α must be a mental action, states that “ φ holds after 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”.

Let us now introduce the dynamic component of the framework. Borrowing from [6, 24], we distinguish five types of mental actions α that capture some of the dynamic properties of explicit beliefs and background knowledge. $+\varphi$, $\downarrow(\varphi, \psi)$, $\cap(\varphi, \psi)$, $\neg(\varphi, \psi)$, and $\vdash(\varphi, \psi)$. These actions characterize the basic operations of belief formation through inference:

- $+\varphi$: learning perceived belief: the mental operation that serves to form a new belief from a perception φ . A perception may become a belief whenever an agent becomes “aware” of the perception and takes it into explicit consideration.
- $\downarrow(\varphi, \psi)$ is the mental action which consists in inferring ψ from φ , where ψ is an atom: an agent, believing that φ is true and having in its long-term memory that φ implies ψ , starts believing that ψ is true.
- $\cap(\varphi, \psi)$ is the mental action which closes the beliefs φ and belief ψ under conjunction. Namely, $\cap(\varphi, \psi)$ characterizes the mental action of deducing $\varphi \wedge \psi$ from φ and ψ .
- $\neg(\varphi, \psi)$, where φ and ψ are atoms, is the mental action that performs a simple form of “belief revision”, i.e., it removes ψ from the belief set, in case φ is believed and, according to the background knowledge, $\neg\psi$ is logical consequence of φ .
- $\vdash(\varphi, \psi)$, where ψ is an atom; by means of this mental action, an agent believing that φ is true (i.e., it is in the working memory) and that φ implies ψ , starts believing that ψ is true. This last action operates exclusively on the working memory without recovering anything from the background knowledge.

We conclude this section by introducing some pieces of notation that will be useful in the following. Recall $Res = \{r_1, \dots, r_\ell\}$ is the set of all (names of) resources. Each resource is available in a certain amount. For simplicity, let us assume that all amounts are natural numbers. We will use the writing $r_j:n_j$ to denote an amount of n_j units of resource r_j . Hence, the writing $(r_1:n_1, \dots, r_\ell:n_\ell)$ describes the amounts of all existing resources. More in general, let $Amounts = \{(r_1:n_1, \dots, r_\ell:n_\ell) \mid n_1, \dots, n_\ell \in \mathbb{N}\}$ be the set of all possible descriptions of amounts of all resources. Moreover, given $\bar{r} = (r_1:n_1, \dots, r_\ell:n_\ell)$ and $\bar{t} = (r_1:m_1, \dots, r_\ell:m_\ell)$ in $Amounts$, we write $\bar{r} \leq \bar{t}$ iff $\forall i n_i \leq m_i$ and denote by $sum(\bar{r})$ the value $\sum_{i=1}^{\ell} n_i$. Also, in case $\bar{t} \leq \bar{r}$, we denote by $\bar{r} - \bar{t}$ the tuple $\bar{r} = (r_1:(n_1 - m_1), \dots, r_\ell:(n_\ell - m_\ell))$.

2.2. Semantics

Many relevant aspects of an agent's behaviour are specified in the definition of *L-INF model*, including what mental and physical actions an agent can perform, what is the cost of an action and what is the budget that the agent has at its disposal, what is the degree of preference of the agent to perform each action, what is the degree of preference of the agent to use a particular resource. This choice has the advantage of keeping the complexity of the logic under control and making these aspects modular. Definitions 2.1 and 2.2 introduce the notion of *L-INF model*, which is then used to introduce semantics of the static fragment *L-INF*. A model M is composed of two parts. A *core part* \mathcal{C}_M and a collection of *packages* \mathcal{P}_M . More specifically:

Definition 2.1. *The core part \mathcal{C}_M of a model M is a tuple $(W, N, \mathcal{R}, V, S)$, where*

- W is a set of worlds (or situations);
- $\mathcal{R} = \{R_i\}_{i \in \text{Agt}}$ is a collection of equivalence relations on W : $R_i \subseteq W \times W$;
- $N : \text{Agt} \times W \rightarrow 2^{2^W}$ is a neighborhood function such that, for each $i \in \text{Agt}$, each $w, v \in W$, and each $X \subseteq W$ these conditions hold:
 - (C1) if $X \in N(i, w)$ then $X \subseteq \{v \in W \mid wR_iv\}$,
 - (C2) if wR_iv then $N(i, w) = N(i, v)$;
- $V : W \rightarrow 2^{\text{Atm}}$ is a valuation function;
- $S : W \rightarrow 2^{\{do_G(\phi_A), do_i^P(\phi_A) \mid \phi_A \in \text{Atm}_A, i \in \text{Agt}, G \in \text{Grp}\}}$ is a valuation function for formulas of the forms $do_G(\phi_A)$ and $do_i^P(\phi_A)$.

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 . 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. This aspect of the semantics can be extended in future work to allow agents make plausible assumptions.

The packages of a model can be thought as modular extensions of the core part. Each package is used to specify a specific feature, such as preferences, costs, executability, etc. Ideally, each package, (may) correspond to some syntactic element of the syntax of *L-INF*. The connection between the syntactic elements and the corresponding package will be established by a suitable component of the semantics (so be seen). The following are some possible packages. Note that we are focusing on those of interests for the purposes of this paper. Plainly, the designer of a particular MAS may decide to include only part of the following packages or even to add/model other features (also providing a suitable adaptation of the notion of truth).

Definition 2.2. Given a core model $\mathcal{C}_M = (W, N, \mathcal{R}, V, S)$, the packages \mathcal{P}_M are:

EXECUTABILITY FOR MENTAL ACTIONS

- $E : \text{Agt} \times W \rightarrow 2^{\mathcal{L}_{\text{ACT}}}$ is an executability function of mental actions such that, for each $i \in \text{Agt}$ and $w, v \in W$, it holds that:

$$(D1) \text{ if } wR_iv \text{ then } E(i, w) = E(i, v);$$

BUDGET AND COSTS FOR MENTAL ACTIONS

- $B_1 : \text{Agt} \times W \rightarrow \mathbb{N}$ is a budget function such that, for each $i \in \text{Agt}$ and $w, v \in W$, the following holds

$$(E1) \text{ if } wR_iv \text{ then } B_1(i, w) = B_1(i, v);$$

- $C_1 : \text{Agt} \times \mathcal{L}_{\text{ACT}} \times W \rightarrow \mathbb{N}$ is a cost function such that, for each $i \in \text{Agt}$, $\alpha \in \mathcal{L}_{\text{ACT}}$, and $w, v \in W$, it holds that:

$$(F1) \text{ if } wR_iv \text{ then } C_1(i, \alpha, w) = C_1(i, \alpha, v);$$

EXECUTABILITY FOR PHYSICAL ACTIONS

- $A : \text{Agt} \times W \rightarrow 2^{\text{Atm}_A}$ is an executability function for physical actions such that, for each $i \in \text{Agt}$ and $w, v \in W$, it holds that:

$$(G1) \text{ if } wR_iv \text{ then } A(i, w) = A(i, v);$$

BUDGET AND COSTS FOR PHYSICAL ACTIONS

- $B_2 : \text{Agt} \times W \rightarrow \text{Amounts}$ is a budget function for physical action, such that, for each $i \in \text{Agt}$, and $w, v \in W$, it holds that:

$$(E2) \text{ if } wR_iv \text{ then } B_2(i, w) = B_2(i, v);$$

- $C_2 : \text{Agt} \times \text{Atm}_A \times W \rightarrow \text{Amounts}$ is a cost function for physical action, such that, for each $i \in \text{Agt}$, $\phi_A \in \text{Atm}_A$, and $w, v \in W$, it holds that:

$$(F2) \text{ if } wR_iv \text{ then } C_2(i, \phi_A, w) = C_2(i, \phi_A, v);$$

AGENTS' ROLES

- $H : \text{Agt} \times W \rightarrow 2^{\text{Atm}_A}$ is an enabling function for physical actions such that, for each $i \in \text{Agt}$ and $w, v \in W$, it holds that:

$$(G2) \text{ if } wR_iv \text{ then } H(i, w) = H(i, v);$$

PREFERENCES ON PHYSICAL ACTIONS

- $P : \text{Agt} \times W \times \text{Atm}_A \rightarrow \mathbb{N}$ is a preference function for physical actions ϕ_A such that, for each $i \in \text{Agt}$ and $w, v \in W$, it holds that:

$$(H1) \text{ if } wR_iv \text{ then } P(i, w, \phi_A) = P(i, v, \phi_A);$$

For each i and w , the function P induces a preference order $\preceq_{i,w}$ on Atm_A , such that $\phi_A \preceq_{i,w} \phi'_A$ iff $P(i, w, \phi_A) \leq P(i, w, \phi'_A)$.

EQUIVALENCE OF PHYSICAL ACTIONS

- $Q : \text{Atm}_A \times W \rightarrow 2^{\text{Atm}_A}$ is a function describing a partition of Atm_A in equivalence classes (i.e., Q associates each physical action with its equivalence class), such that for each $i \in \text{Agt}$ and $w, v \in W$, it holds that:

$$(I1) \text{ if } wR_iv \text{ then } Q(\phi_A, w) = Q(\phi_A, v);$$

- $F : \text{Agt} \times W \times \mathcal{L}_{\text{ACT}} \rightarrow \mathcal{L}_{\text{ACT}}$ is a selector function for physical actions that, given i and w , selects one physical action $F(i, w, \phi_A)$ from the equivalence class of ϕ_A . Namely, it holds that $F(i, w, \phi_A) \in Q(\phi_A, w) \wedge \forall \phi'_A \in Q(\phi_A, w) \phi'_A \preceq_{i,w} \phi_A$. For each $i \in \text{Agt}$ and $w, v \in W$, it holds that:

$$\text{(I2)} \text{ if } wR_iv \text{ then } F(i, w, \phi_A) = F(i, v, \phi_A).$$

Let us briefly describe the intended features shaped by the packages introduced by Def. 2.2. Notice that the concrete implementation, in a real MAS, the specification of some packages might depend on other packages (for example, in what follows we will describe a possible implementation of F that relies on the function P).

For an agent i , $E(i, w)$ is the set of mental actions that i can execute at world w . To execute a mental action, i has to pay the cost $C_1(i, \alpha, w)$. $B_1(i, w)$ is the budget that i has (in w) to perform mental actions. As mentioned, concerning physical actions, we are interested in modeling situations where performing an action may require multiple resources. Hence, the cost $C_2(i, \phi_A, w)$ of an action ϕ_A (for agent i in world w) is a tuple in *Amounts*, while the available budget is described by $B_2(i, w)$. For an agent i , the set of physical actions it can execute at w is $A(i, w)$. Equivalence between physical actions is determined by function Q . That is, $Q(\phi_A, w)$ is the set of physical actions that are equivalent to ϕ_A in w . Roles of agents (that, as we will see, affects the capability of agents in a group to execute actions) is described through H . Namely, $H(i, w)$ is the set of physical actions that agent i is enabled by its group to perform (recall that, at each time instant, an agent belongs to a single group). Agent's preference on execution of physical actions is determined by the function P . For an agent i , and a physical action ϕ_A , the value of $P(i, w, \phi_A)$ should be intended as a *degree of willingness* of agent i to execute ϕ_A at world w . Analogously to property **(C2)** imposed in Def. 2.1, the constrain **(D1)** imposes that agent i always knows which mental actions it can perform and those it cannot, but if two situations/worlds are equivalent as concerns background knowledge, then they cannot be distinguished through the executability of actions. Similar “indistinguishability” requirements are imposed for each package by conditions **(E1)**, **(F1)**, **(F2)**, **(G1)**, **(H1)**, **(G2)**, **(H2)**, **(I1)**, and **(I2)**.

Let us give some hints on how the functions P and F might be actually implemented in a concrete MAS. As concerns P , we assume defined (e.g., by the MAS designer) a preference relation among (equivalent) actions, for any agent i . In practice, this relation might be obtained by exploiting some specific reasoning module. Some possibilities in this sense are described in [25, 26]. Similarly, as for all packages, a specific module in the MAS implementation may be devoted to realize the selector function F . Here we outline a simple option in defining F , relying on the availability of functions B_2 , C_2 , and Q . Given an agent i , a world w , and an action ϕ_A , let $\mathcal{A} = \{\phi'_A \mid \phi'_A \in Q(\phi_A, w) \wedge C_2(i, \phi'_A, w) \leq B_2(i, w)\}$ and $\mathcal{A}' \subseteq \mathcal{A}$ such that for each $\phi'_A \in \mathcal{A}'$ the value $\text{sum}(C_2(i, \phi'_A, w))$ is minimal among the elements of \mathcal{A} . Finally, select the preferred element to be the $\preceq_{i,w}$ -maximal element of \mathcal{A}' (i.e., the action ϕ''_A with the larger value of $P(i, w, \phi''_A)$). In case of multiple options, any deterministic criterion can be applied).

2.2.1. Truth Conditions

Truth values of L -*DINF* formulas are inductively defined as follows. Given a model M (cf., Definitions 2.1 and 2.2), $i \in \text{Agt}$, $G \in \text{Grp}$, $w \in W$, and a formula $\varphi \in \mathcal{L}_{L\text{-INF}}$, we introduce

this shorthand notation: $\|\varphi\|_{i,w}^M = \{v \in W : wR_iv \text{ and } M, v \models \varphi\}$, whenever $M, v \models \varphi$ is well-defined (see below). Then, we set:

1. $M, w \models p$ iff $p \in V(w)$
2. $M, w \models Cl(\phi_A, \phi'_A)$ iff $\phi'_A \in Q(\phi_A, w)$
3. $M, w \models exec_G(\alpha)$ iff $\exists i \in G$ with $\alpha \in E(i, w)$
4. $M, w \models pref_do_i(\phi_A, d)$ iff $\phi_A \in A(i, w_I)$ and $P(i, w_I, \phi_A) = d$
5. $M, w \models pref_do_G(i, \phi_A)$ iff $M, w \models pref_do_i(\phi_A, d)$ for $d \in \mathbb{N}$ such that $d = \max\{P(j, w, \phi_A) \mid j \in G \wedge \phi_A \in A(j, w) \cap H(j, w)\}$
6. $M, w \models can_do_G(\phi_A)$ iff $\exists i \in G$ with $\phi_A \in A(i, w) \cap H(i, w)$ and $\phi_A = F(i, w, \phi_A)$
7. $M, w \models fCl_i(\phi_A)$ iff $\phi_A = F(i, w, \phi_A)$
8. $M, w \models \neg\varphi$ iff $M, w \not\models \varphi$
9. $M, w \models \varphi \wedge \psi$ iff $M, w \models \varphi$ and $M, w \models \psi$
10. $M, w \models \mathbf{B}_i\varphi$ iff $\|\varphi\|_{i,w}^M \in N(i, w)$
11. $M, w \models \mathbf{K}_i\varphi$ iff $M, v \models \varphi$ for all $v \in R_i(w)$
12. $M, w \models \varphi$, for φ of the forms $do_G(\phi_A)$ and $do_i^P(\phi_A)$, iff $\varphi \in S(w)$

As mentioned, a physical action can be performed by a group of agents if at least one agent of the group can do it. In this case, the level of preference for performing this action is set to the maximum among those of the agents enabled to execute the action. In addition, the agent selects (using F) among the enabled equivalent actions the one it prefers the most.

Notice that, in the above described semantics, a specific evaluation function S deals with formulas of the forms $do_G(\phi_A)$ and $do_i^P(\phi_A)$. These kind of formulas are, nevertheless, left *not axiomatized*. This because $do_G(\phi_A)$ refers to the practical execution of an action by some kind of actuator, where in a robotic application this action can have physical effects. To find a way for accounting for such expression, we choose to resort to a concept that has been called by Weyhrauch in the seminal work [23] a *semantic attachment*, i.e., it is assumed that some device exists, which connects an agent with its external environment in a way that is unknown at the logical level. The aim of [23] was exactly to explain how formal systems could be used in AI by being “mechanized” in a practical way, by providing ideas about a principled though potentially running implementation of these systems. In our setting, an action is meant to be executed by means of such a device, and, whenever successfully completed, it will be then recorded by means of atoms of the form $do_i^P(\phi_A)$; such records can greatly aid the agent’s subsequent reasoning process and support the ability to provide explanations. Hence, we assume that the function S reflects at the semantic level the presence of such semantic attachment mechanism. So that, the semantics is concerned only with the relationship between doing and being enabled to do. A similar treatment is done for join actions. Performing $join_A(i, j)$ imply that agents i, j are now in the same group. We assume that the execution of $join_A(i, j)$ affects the contents of the working memories of the agents i and j (and consequently of the other members of their groups).

As mentioned, formulas of the form $intend_G(\phi_A)$ express agents’ intentions of performing physical actions. In this paper we do not cope with the formalization of BDI (for which the reader may refer, e.g., to [22]). So, we do not provide a specific semantics for intentions and treat them rather informally, assuming also that $intend_G(\phi_A)$ represents the fact that all agents in G intend to perform ϕ_A .

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

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

where the model $M^{[G:\alpha]}$ is an updated version of the model M , that takes into account the effect that the execution of the mental action α has on the sets of beliefs of i and on the available budget. Hence, $M^{[G:\alpha]}$ is obtained from M by replacing the function N and B_1 , with the function $N^{[G:\alpha]}$ and $B_1^{[G_I:\alpha]}$, resp., defined as described below.

The action α may 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. The following condition characterizes the circumstances in which a mental action may be performed, and by which agent(s):

$$enabled_w(G, \alpha) : \exists j \in G (\alpha \in E(j, w) \wedge \frac{C_1(j, \alpha, w)}{|G|} \leq \min_{h \in G} B_1(h, w))$$

To handle the case of multiple enabled agents, we assume defined a predicate $doer_w(i, G, \alpha)$ to univocally select one among the enabled agents. Its definition might rely on any criteria, even involving background knowledge and belief sets. For simplicity, let us define such predicate as:

$$doer_w(i, G, \alpha) \leftrightarrow i = \min\{j \mid enabled_w(j, G, \alpha)\}.$$

This condition, as defined above, expresses the fact that a mental 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 (this choice is inherited from L-DINF). 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 $C_1(j, \alpha, w) = \min_{h \in G} C_1(h, \alpha, w)$. Other choices might be viable, so variations of this logic can be easily defined simply by devising some other enabling condition and, possibly, introducing differences in neighborhood update.

2.2.2. Belief Update

Updating an agent’s beliefs accounts to modify the neighborhood of the present world. The updated neighborhood $N^{[G:\alpha]}$ resulting from execution of a mental action α by a group G of agents is as follows.

- If α is $+\varphi$, then, for each $i \in Agt$ and $w \in W$,

$$N^{[G:+\varphi]}(i, w) = N(i, w) \cup \{ \|\varphi\|_{i,w}^M \}$$
 if $i \in G$ and $doer_w(i, G, +\varphi)$ holds. Otherwise, the neighborhood function does not change (i.e., $N^{[G:+\varphi]}(i, w) = N(i, w)$).
- If α is $\downarrow(\psi, \chi)$, then, for each $i \in Agt$ and $w \in W$,

$$N^{[G:\downarrow(\psi, \chi)]}(i, w) = N(i, w) \cup \{ \|\chi\|_{i,w}^M \}$$
 if $i \in G$ and $doer_w(i, G, \downarrow(\psi, \chi))$ and $M, w \models \mathbf{B}_i\psi \wedge \mathbf{K}_i(\psi \rightarrow \chi)$ hold. Otherwise, the neighborhood function does not change (i.e., $N^{[G:\downarrow(\psi, \chi)]}(i, w) = N(i, w)$).
- If α is $\cap(\psi, \chi)$, then, for each $i \in Agt$ and $w \in W$,

$$N^{[G:\cap(\psi, \chi)]}(i, w) = N(i, w) \cup \{ \|\psi \wedge \chi\|_{i,w}^M \}$$
 if $i \in G$ and $doer_w(i, G, \cap(\psi, \chi))$ and $M, w \models \mathbf{B}_i\psi \wedge \mathbf{B}_i\chi$. Otherwise, the neighborhood function remains unchanged (i.e., $N^{[G:\cap(\psi, \chi)]}(i, w) = N(i, w)$).

- If α is $\vdash(\psi, \chi)$, then, for each $i \in \text{Agt}$ and $w \in W$,

$$N^{[G:\vdash(\psi, \chi)]}(i, w) = N(i, w) \cup \{ \|\chi\|_{i,w}^M \}$$
if $i \in G$ and $\text{doer}_w(i, G, \vdash(\psi, \chi))$ and $M, w \models \mathbf{B}_i\psi \wedge \mathbf{B}_i(\psi \rightarrow \chi)$. Otherwise, the neighborhood function remains unchanged: $N^{[G:\vdash(\psi, \chi)]}(i, w) = N(i, w)$.
- If α is $\neg(\psi, \chi)$, then, for each $i \in \text{Agt}$ and $w \in W$,

$$N^{[G:\neg(\psi, \chi)]}(i, w) = N(i, w) \setminus \{ \|\chi\|_{i,w}^M \}$$
if $i \in G$ and $\text{doer}_w(i, G, \neg(\psi, \chi))$ and $M, w \models \mathbf{B}_i\psi \wedge \mathbf{K}_i(\psi \rightarrow \neg\chi)$. Otherwise, the neighborhood does not change (i.e., $N^{[G:\neg(\psi, \chi)]}(i, w) = N(i, w)$).

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

2.2.3. Budget Update for Mental Actions

If a mental action α is executed by a group G , agent in G has to contribute to cover the cost of execution by consuming part of its budget. Hence, for each $i \in \text{Agt}$ and each $w \in W$, we set

$$B_1^{[G:\alpha]}(i, w) = B_1(i, w) - C_1(i, \alpha, w)/|G|,$$

if $i \in G$, $\text{doer}_w(i, G, \alpha)$ holds, and, depending on α , the same conditions described before to enable neighborhood updates are satisfied. Otherwise, the budget is preserved, i.e., $B_1^{[G:\alpha]}(i, w) = B_1(i, w)$. Clearly, the budget is preserved for those agents that are not in G .

2.2.4. Budget Update for Physical Action

Also execution of physical actions involves consumption of some amounts of resources. Hence, the budget available for physical actions has to be updated accordingly. If an action ϕ_A is performed by an agent i at a world w , this involves a transition from w to another world w' . Moreover, if an action is executed, it means that enough resources are available and are consumed to complete the action. The budget function satisfies this condition:

$$B_2(i, w') = B_2(i, w) - C_2(i, \phi_A, w)$$

Remark 2.1. *A comment is due concerning the action $\text{join}_A(i, j)$. We assume that whenever an agent $i \in G$ joins the group of another agent j (by executing $\text{join}_A(i, j)$), the neighborhood function $N(i, w)$ becomes equal to $N(j, w)$, for each $h \in \text{Agt}$. In case $i \in G$ executes $\text{join}_A(i, i)$ (i.e., it leaves G and forms a new singleton group) then it maintains its current neighborhood function, but without any binding with the belief set of the remaining agents in G .*

Remark 2.2. *In the actions $\vdash(\varphi, \psi)$ and $\downarrow(\varphi, \psi)$, the formula ψ which is inferred and asserted as a new belief can be of the forms $\text{can_do}_G(\phi_A)$ or $\text{do}_i(\phi_A)$. Conclusion $\text{do}_i(\phi_A)$ (from $\text{can_do}_i(\phi_A)$ and possibly other conditions) implies that physical action ϕ_A is actually performed by i . Actions are supposed to succeed by default, in case of failure a corresponding failure event will be perceived by the agent (again, we rely on semantic attachment). The do_i^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.*

2.3. Axiomatization

In previous works [1, 4, 3], a sound and complete axiomatic system has been proposed for *L-INF* and *L-DINF*. For simplicity, below we recall the axiomatization of the core part of the logic only (corresponding to the notion of core model introduced by Def. 2.1). The *L-INF* and *L-DINF* core axioms and inference rules are (together with the usual axioms of propositional logic):

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. $[G:\alpha](\varphi \wedge \psi) \leftrightarrow [G:\alpha]\varphi \wedge [G:\alpha]\psi$;
12. $[G:\alpha]\mathbf{K}_i\varphi \leftrightarrow \mathbf{K}_i([G:\alpha]\varphi)$;
13. $[G:+\varphi]\mathbf{B}_i\psi \leftrightarrow (\mathbf{B}_i([G:+\varphi]\psi) \vee (\text{doer}_w(i, G, +\varphi) \wedge \mathbf{K}_i([G:+(\varphi)]\psi \leftrightarrow \varphi)))$;
14. $[G:\downarrow(\varphi, \psi)]\mathbf{B}_i\chi \leftrightarrow (\mathbf{B}_i([G:\downarrow(\varphi, \psi)]\chi) \vee (\text{doer}_w(i, G, \downarrow(\varphi, \psi)) \wedge \mathbf{B}_i\varphi \wedge \mathbf{K}_i(\varphi \rightarrow \psi) \wedge \mathbf{K}_i([G:\downarrow(\varphi, \psi)]\chi \leftrightarrow \psi)))$;
15. $[G:\cap(\varphi, \psi)]\mathbf{B}_i\chi \leftrightarrow (\mathbf{B}_i([G:\cap(\varphi, \psi)]\chi) \vee (\text{doer}_w(i, G, \cap(\varphi, \psi)) \wedge \mathbf{B}_i\varphi \wedge \mathbf{B}_i\psi \wedge \mathbf{K}_i([G:\cap(\varphi, \psi)]\chi \leftrightarrow (\varphi \wedge \psi))))$;
16. $[G:\vdash(\varphi, \psi)]\mathbf{B}_i\chi \leftrightarrow (\mathbf{B}_i([G:\vdash(\varphi, \psi)]\chi) \vee (\text{doer}_w(i, G, \vdash(\varphi, \psi)) \wedge \mathbf{B}_i\varphi \wedge \mathbf{B}_i(\varphi \rightarrow \psi) \wedge \mathbf{K}_i([G:\vdash(\varphi, \psi)]\chi \leftrightarrow \psi)))$;
17. $[G:\dashv(\varphi, \psi)]\mathbf{B}_i\chi \leftrightarrow (\mathbf{B}_i([G:\dashv(\varphi, \psi)]\chi) \vee (\text{doer}_w(i, G, \dashv(\varphi, \psi)) \wedge \mathbf{B}_i\varphi \wedge \mathbf{K}_i(\varphi \rightarrow \neg\psi) \wedge \mathbf{K}_i([G:\dashv(\varphi, \psi)]\chi \leftrightarrow \neg\psi)))$;
18. $\text{intend}_G(\phi_A) \leftrightarrow \forall i \in G \text{intend}_i(\phi_A)$;
19. $\frac{\psi \leftrightarrow \chi}{\varphi \leftrightarrow \varphi[\psi/\chi]}$.

We write $L-DINF \vdash \varphi$ to denote that φ is a theorem of *L-DINF*.

The above axiomatization is sound for the class of *L-INF* models. Namely, all axioms are valid and the two inference rules (8) and (19) preserve validity. In particular, soundness of axioms (13)–(17) follows from the semantics of $[G:\alpha]\varphi$, for each mental action α , as previously defined. Recall that, as mentioned earlier in the paper, the axiomatization does not deal with formulas of the forms $\text{do}_G(\phi_A)$, as they are intended to be realized by a *semantic attachment*, that connects an agent with its external environment.

As concerns completeness of the (core) axiomatization, a standard notion of *canonical L-INF model* can be introduced. Then, strong completeness of the axiomatization can be proved by applying a standard canonical-model argument (cf., [3, 6, 9]) and this leads to the following result:

Theorem 2.1. *L-DINF is strongly complete for the class of L-INF models.*

3. Problem Specification and Inference: An Example

In this section, we propose an example of problem specification and inference in *L-DINF*. Consider a group of n agents, e.g., four, who are colleagues, work together over the weekend to prepare banquets. They know that three of them are both able to prepare dishes, while the other one is both able to prepare the *mise en place*. Below we show how our logic is able to represent the situation, and the proceedings of this work. Each agent will initially have in its knowledge base the fact $\mathbf{K}_i(\textit{intend}_G(\textit{prepare-banquet}))$. The physical actions are the following:

$$\textit{prepare-mise-en-place}, \quad \textit{prepare-starters}, \quad \textit{prepare-main-dish}, \quad \textit{prepare-dessert}. \quad (1)$$

Assume that the knowledge base of each agent i contains the following rule, that specifies how to reach the intended goal in terms of actions to perform:

$$\mathbf{K}_i(\textit{intend}_G(\textit{prepare-banquet}) \rightarrow \textit{intend}_G(\textit{prepare-mise-en-place}) \wedge \textit{intend}_G(\textit{prepare-starters}) \wedge \textit{intend}_G(\textit{prepare-main-dish}) \wedge \textit{intend}_G(\textit{prepare-dessert})).$$

By axiom18, every agent will also have the following:

$$\mathbf{K}_i(\textit{intend}_i(\textit{prepare-banquet}) \rightarrow \textit{intend}_i(\textit{prepare-mise-en-place}) \wedge \textit{intend}_i(\textit{prepare-starters}) \wedge \textit{intend}_i(\textit{prepare-main-dish}) \wedge \textit{intend}_i(\textit{prepare-dessert})).$$

Therefore, the following is entailed for each of the agents ($1 \leq i \leq 4$):

$$\begin{aligned} \mathbf{K}_i(\textit{intend}_i(\textit{prepare-banquet}) &\rightarrow \textit{intend}_i(\textit{prepare-mise-en-place})) \\ \mathbf{K}_i(\textit{intend}_i(\textit{prepare-banquet}) &\rightarrow \textit{intend}_i(\textit{prepare-starters})) \\ \mathbf{K}_i(\textit{intend}_i(\textit{prepare-banquet}) &\rightarrow \textit{intend}_i(\textit{prepare-main-dish})) \\ \mathbf{K}_i(\textit{intend}_i(\textit{prepare-banquet}) &\rightarrow \textit{intend}_i(\textit{prepare-dessert})). \end{aligned} \quad (2)$$

Assume now that the knowledge base of each agent i contains also the following rule, for $\phi_A = \textit{prepare-mise-en-place}, \textit{prepare-starters}, \textit{prepare-main-dish}, \textit{prepare-dessert}$:

$$\mathbf{K}_i(\textit{intend}_i(\phi_A) \wedge \textit{can_do}_i(\phi_A) \wedge \textit{fCl}_i(\phi_A) \rightarrow \textit{do}_i(\phi_A))$$

As previously stated, whenever an agent derives $\textit{do}_i(\phi_A)$ for any physical action ϕ_A , the action is supposed to have been performed via some kind of *semantic attachment* which links the agent to the external environment. However, $\textit{do}_i(\phi_A)$ will be derived by means of some mental action based upon the available rules. Such mental action can have a cost, that can be paid either by the agent itself or by the group (according to the adopted policy of cost-sharing for this group). According to the above rules, an agent can execute an action ϕ_A if it is able to derive $\textit{can_do}_i(\phi_A)$, and ϕ_A is the selected one among the viable equivalent actions (i.e., $\textit{fCl}_i(\phi_A)$ has also been derived). Such conclusion will be drawn on the basis of the assessment performed in external modules that concretely implement the model packages. These modules provide the decision according to some kind of reasoning process in some formalism, with respect to which the logic *L-DINF* is completely agnostic: modules will add the corresponding facts to each agent's knowledge base.

To get the agents do the actions listed in (1) four sequences of mental actions have to be executed, yielding, respectively, conclusions of the forms $do_G(\text{prepare-mise-en-place})$, $do_G(\text{prepare-starters})$, $do_G(\text{prepare-main-dish})$, and $do_G(\text{prepare-dessert})$, and causing their addition to agents' working memory. Such reasoning would consist in mental actions of kind \cap to form conjunctions from single facts, and mental actions of kind \downarrow to apply knowledge rule, i.e., given their preconditions, draw the conclusions. In particular, given the initial general intention by the group, it will be possible to derive the practical goal, in terms of the conjunction of actions to be performed by the group. From its own specialized rules and the available facts about enabling and willingness, the execution of each action by some agent i will be then derived. Note that, there can be unlucky situations where no agent is enabled to perform some action, or that the one allowed is not willing, or that there is not enough budget. In this case, the goal fails.

Let $\alpha_1-\alpha_4$ be the last mental actions performed at the end of the mentioned four sequences of mental inferences (that lead to derive the $do_i(\phi_A)$, for ϕ_A among the actions in (1)), respectively. For how such mental actions are treated we can refer to [6]. Let us focus on physical actions and their equivalence classes that we assume to be specified by function Q so that:

$$Q(\text{prepare-mise-en-place}) = \{\text{prepare-mise-en-place}\}$$

$$Q(\text{prepare-starters}) = \{\text{prepare-starters}, \text{prepare-main-dish}, \text{prepare-dessert}\}$$

Moreover let the cost of physical actions be (where we do not list resources with null amounts):

$$C_2(i, \text{prepare-mise-en-place}, w_1) = (\text{table} : 5, \text{cutlery} : 10, \text{table-center} : 5),$$

$$C_2(i, \text{prepare-dessert}, w_1) = (\text{sugar} : 4, \text{flour} : 5, \text{egg} : 2),$$

$$C_2(i, \text{prepare-starters}, w_1) = (\text{appetizers} : 5, \text{drinks} : 5),$$

$$C_2(i, \text{prepare-main-dish}, w_1) = (\text{sauce} : 2, \text{pasta} : 10)$$

and the their budget are:

$$B_2(1, w_1) = (\text{table} : 6, \text{cutlery} : 10, \text{table-center} : 7),$$

$$B_2(2, w_1) = (\text{sugar} : 5, \text{flour} : 7, \text{egg} : 4, \text{appetizers} : 4, \text{drinks} : 4, \text{sauce} : 2, \text{pasta} : 9)$$

$$B_2(3, w_1) = (\text{sugar} : 4, \text{flour} : 5, \text{egg} : 1, \text{appetizers} : 5, \text{drinks} : 5)$$

$$B_2(4, w_1) = (\text{sugar} : 3, \text{flour} : 5, \text{egg} : 2, \text{appetizers} : 1, \text{drinks} : 1, \text{sauce} : 5, \text{pasta} : 20)$$

Considering the available resources, only agent 1 can perform *prepare-mise-en-place*, for the other three action we choose the right one using the function F . By assuming the definition of F described in Section 2.2, we have that agent 2 can perform *prepare-dessert*, agent 3 can perform *prepare-starters* and agent 4 can perform *prepare-main-dish*. After the execution of these actions the new budget becomes (cf. Section 2.2.4):

$$B_2(1, w_2) = (\text{table} : 5, \text{table-center} : 2),$$

$$B_2(2, w_2) = (\text{sugar} : 1, \text{flour} : 2, \text{egg} : 2, \text{appetizers} : 4, \text{drinks} : 4, \text{sauce} : 2, \text{pasta} : 9)$$

$$B_2(3, w_2) = (\text{sugar} : 4, \text{flour} : 5, \text{egg} : 1)$$

$$B_2(4, w_2) = (\text{sugar} : 3, \text{flour} : 5, \text{egg} : 2, \text{appetizers} : 1, \text{drinks} : 1, \text{sauce} : 3, \text{pasta} : 10)$$

It is relevant to comment about the role of past events. If the set of past events, which is a part of an agent's short-term memory, is made available to the external modules defining actions enabling and degree of willingness, such recordings might be used, for instance, to define constraints concerning actions execution.

4. Conclusions

In this paper we extended an epistemic logical framework previously introduced in [4, 3] and originally designed to enable modeling and reasoning on group dynamics of cooperative agents. In such framework, agents can perform actions and inferences on the base of their knowledge and their beliefs. Agents can reason about mental and physical actions. The extension presented in this paper enriches such a framework with the possibility of modeling costs involving multiple resources consumption and specifying through particular components of the concrete MAS (modules/packages) which physical actions to choose according to agent's preferences, actions' costs, and available resources.

References

- [1] S. Costantini, A. Formisano, V. Pitoni, Timed memory in resource-bounded agents, in: C. Ghidini, B. Magnini, A. Passerini, P. Traverso (Eds.), *AI*IA 2018 - Advances in Artificial Intelligence - XVIIth International Conference of the Italian Association for Artificial Intelligence*, 2018, Proceedings, volume 11298 of *LNCS*, Springer, 2018, pp. 15–29.
- [2] V. Pitoni, Memory management with explicit time in resource-bounded agents, in: S. A. McIlraith, K. Q. Weinberger (Eds.), *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18)*, AAAI Press, 2018, pp. 8133–8134.
- [3] S. Costantini, A. Formisano, V. Pitoni, An epistemic logic for multi-agent systems with budget and costs, in: W. Faber, G. Friedrich, M. Gebser, M. Morak (Eds.), *Logics in Artificial Intelligence - 17th European Conference, JELIA 2021, Proceedings*, volume 12678 of *LNCS*, Springer, 2021, pp. 101–115.
- [4] S. Costantini, V. Pitoni, Towards a logic of “inferable” for self-aware transparent logical agents, in: C. Musto, D. Magazzeni, S. Ruggieri, G. Semeraro (Eds.), *Proc. of the Italian Workshop on Explainable AI co-located with 19th International Conference of AI*IA*, 2020, volume 2742 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2020, pp. 68–79.
- [5] S. Costantini, A. Formisano, V. Pitoni, Cooperation among groups of agents in the epistemic logic L-DINF, in: G. Governatori, A. Turhan (Eds.), *Proc. of RuleML+RR'22*, volume 13752 of *LNCS*, Springer, 2022, pp. 280–295.
- [6] S. Costantini, A. Formisano, V. Pitoni, An epistemic logic for modular development of multi-agent systems, in: N. Alechina, M. Baldoni, B. Logan (Eds.), *Proc. of EMAS'21, Revised Selected papers*, volume 13190 of *LNCS*, Springer, 2021, pp. 72–91.
- [7] S. Costantini, A. Formisano, V. Pitoni, Modelling agents roles in the epistemic logic L-DINF, in: O. Arieli, G. Casini, L. Giordano (Eds.), *NMR'22*, volume 3197 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022, pp. 70–79.
- [8] A. I. Goldman, Theory of mind, in: E. Margolis, R. Samuels, S. P. Stich (Eds.), *The Oxford Handbook of Philosophy of Cognitive Science*, volume 1, Oxford University Press, 2012, pp. 402–424.

- [9] S. Costantini, A. Formisano, V. Pitoni, An epistemic logic for formalizing group dynamics of agents, *Interaction Studies* 23 (2023) 391–426.
- [10] R. Fagin, J. Y. Halpern, Belief, awareness, and limited reasoning., *Artif. Intell.* 34 (1987) 39–76.
- [11] J. van Benthem, E. Pacuit, Dynamic logics of evidence-based beliefs, *Studia Logica* 99 (2011) 61–92.
- [12] M. Jago, Epistemic logic for rule-based agents, *Journal of Logic, Language and Information* 18 (2009) 131–158.
- [13] F. R. Velázquez-Quesada, Dynamic epistemic logic for implicit and explicit beliefs, *Journal of Logic, Language and Information* 23 (2014) 107–140.
- [14] P. Balbiani, D. Fernández-Duque, E. Lorini, The dynamics of epistemic attitudes in resource-bounded agents, *Studia Logica* 107 (2019) 457–488.
- [15] F. R. Velázquez-Quesada, Explicit and implicit knowledge in neighbourhood models, in: D. Grossi, O. Roy, H. Huang (Eds.), *Logic, Rationality, and Interaction - 4th International Workshop, LORI 2013*, volume 8196 of *LNCS*, Springer, 2013, pp. 239–252.
- [16] H. N. Duc, Reasoning about rational, but not logically omniscient, agents, *J. Log. Comput.* 7 (1997) 633–648.
- [17] V. Pitoni, S. Costantini, A temporal module for logical frameworks, in: B. Bogaerts, E. Erdem, P. Fodor, A. Formisano, G. Ianni, D. Inlezan, G. Vidal, A. Villanueva, M. De Vos, F. Yang (Eds.), *Proc. of ICLP 2019 (TC)*, volume 306 of *EPTCS*, 2019, pp. 340–346.
- [18] S. Costantini, V. Pitoni, Memory management in resource-bounded agents, in: M. Alviano, G. Greco, F. Scarcello (Eds.), *AI*IA 2019 - Advances in Artificial Intelligence - XVIIIth International Conference of the Italian Association for Artificial Intelligence*, 2019, Proceedings, volume 11946 of *LNCS*, Springer, 2019, pp. 46–58.
- [19] J. J. Elgot-Drapkin, S. Kraus, M. Miller, M. Nirkhe, D. Perlis, *Active Logics: A Unified Formal Approach to Episodic Reasoning*, Technical Report, UMIACS—University of Maryland, 1999. CS-TR-4072.
- [20] J. J. Elgot-Drapkin, M. I. Miller, D. Perlis, Life on a desert island: ongoing work on real-time reasoning, in: F. M. Brown (Ed.), *The Frame Problem in Artificial Intelligence*, Morgan Kaufmann, 1987, pp. 349–357.
- [21] A. S. Rao, M. Georgeff, Modeling rational agents within a BDI architecture, in: *Proc. of the Second Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'91)*, Morgan Kaufmann, 1991, pp. 473–484.
- [22] H. V. Ditmarsch, J. Y. Halpern, W. V. D. Hoek, B. Kooi, *Handbook of Epistemic Logic*, College Publications, 2015. Editors.
- [23] R. W. Weyhrauch, Prolegomena to a theory of mechanized formal reasoning, *Artif. Intell.* 13 (1980) 133–170.
- [24] P. Balbiani, D. F. Duque, E. Lorini, A logical theory of belief dynamics for resource-bounded agents, in: *Proc. of AAMAS'16*, ACM, 2016, pp. 644–652.
- [25] S. Costantini, A. Formisano, Modeling preferences and conditional preferences on resource consumption and production in ASP, *J. Algorithms* 64 (2009) 3–15.
- [26] S. Costantini, A. Formisano, Augmenting weight constraints with complex preferences, in: *Logical Formalizations of Commonsense Reasoning, Papers from the 2011 AAAI Spring Symposium*, AAAI Press, USA, 2011.