Preferential multi-agent nonmonotonic logics: Preliminary report

Ana Maria Monteiro
Institute of Computing
State University of Campinas, Brazil
anammont@dcc.unicamp.br

Jacques Wainer
Institute of Computing
State University of Campinas, Brazil
wainer@dcc.unicamp.br

Abstract

This paper presents a way of constructing a multi-agent nonmonotonic logic from any propositional preference logic (the base logic). The multi-agent logic will "correctly" extend the base logic for each agent, that is, each agent will reason in the base logic, and the reasoning of each agent is independent of each other. If the base logic is propositional circumscription then the extended logic will also correctly deal with default formulas that mix the knowledge of more than one agent. Some of the properties of the extended logics and an example of nonmonotonic multi-agent reasoning are presented.

1 Introduction

The term agent is frequently used in the area of knowledge representation and reasoning. Several different systems have been proposed to model an agent reasoning about its environment, an environment that in many applications includes other agents. Thus, some logics have been proposed to deal with the knowledge of more than one agent [HM92, FHV91], but these logics are monotonic and hence have a very limited capacity of modeling interesting behaviors and communications among the agents.

There has been some multi-agent nonmonotonic logics developed in the last years based on autoepistemic logics [Mor90, Lak93, HM93, PJ95]. This paper we will develop a preference-based nonmonotonic logic for reasoning about the beliefs of many agents.

More specifically, this work proposes a way of defining a multi-agent preferential logic based on a propositional preference relation in such a way that whatever is captured by the propositional preference relation is carried over "correctly" to the multi-agent logic. Conceptually, if $\mathcal{L}$ is a propositional logic, one can define a preferential logic $\mathcal{L}_{\leq}$ based on a preference relation $\leq$ on the models of $\mathcal{L}$. On the other hand one can define a multi-agent logic $\mathcal{L}^*$ based on $\mathcal{L}$, which includes modal operators to represent the beliefs of n agents. This paper proposes a way of defining a preference relation $\preceq$ among models of $\mathcal{L}^*$ in order to define the multi-agent preference logic $\mathcal{L}^*_{\preceq}$.

The next section describes preference logics and the semantics of the $\mathcal{L}^*$ logic. Section 3 describes how to construct the $\preceq$ preference relation based on the $\leq$ relation. The other section describes some of the properties of the $\mathcal{L}^*_{\preceq}$ logic. Section 5 describe some examples of using the logic.

2 Background

2.1 Preference Logics

This work is based on the model-theoretic approach to constructing nonmonotonic logics proposed by Shoham [Sho87]. The main idea behind preferential logics is that the meaning of a formula is not given by the set of all of its models, as in classical logics, but by some preferred subset of those models, called *minimal models*. Many different preference criteria can be adopted according to one's needs, resulting in different nonmonotonic logic.

Formally, let $\mathcal{L}$ be a standard (monotonic) propositional logic and $\leq$ a pre-order on interpretations for $\mathcal{L}$. If M_1 and M_2 are two interpretations for $\mathcal{L}$, then $M_1 \leq M_2$ means that the interpretation M_1 is as preferred as the interpretation M_2. $\mathcal{L}$ and $\leq$ define a preference logic $\mathcal{L}_{\leq}$. The language of $\mathcal{L}_{\leq}$ is the same of that of $\mathcal{L}$ and its semantics is defined as follows.

Definition 1 *Given an interpretation M of $\mathcal{L}$ and α a sentence of $\mathcal{L}$, then M preferentially satisfies α (written $M \models_{\leq} \alpha$) if $M \models \alpha$, and if there is no other interpretation M' such that $M' < M$ and $M' \models \alpha$. In this case we say that M is a preferred model of α.*

Definition 2 *α is preferentially satisfiable if there exists an M such that $M \models_{\leq} \alpha$.*

Definition 3 *α preferentially entails β (written $\alpha \models_\leq \beta$) if for any M, if $M \models_\leq \alpha$ then $M \models \beta$, or equivalently, if the models of β are a superset of the preferred models of α.*

This framework is general enough to capture the details of many nonmonotonic reasoning systems. Among the minimal model approaches, we can mention circumscription [McC80, McC86], and many proposals for logics that reason about time and change [Sho88, Bak91, Sha95].

In this paper we will be restricted to propositional logics, and we will develop examples of default reasoning (as opposed to, say, temporal reasoning). There is a trivial simplification of circumscription that allows one to represent and reasoning about defaults in a propositional language. We call it propositional circumscription. Default formulas of the kind α then normally β are represented as:

$$\alpha \wedge \neg abn_1 \rightarrow \beta$$

The preference criteria is the one that tries to falsify abn_1. That is: $M_1 < M_2$ iff

- for all propositional symbols p in the language other than abn_1 $M_1 \models p$ iff $M_2 \models p$.
- $M_1 \models \neg abn_1$ and $M_2 \models abn_1$

All the standard circumscription variations: priority among abnormal predicates, formula circumscription and allowing predicates to vary, can be implemented in propositional circumscription.

2.2 The $\mathcal{L}^*$ logic

In this section we describe how to extend a monotonic base logic $\mathcal{L}$ to a monotonic multi-agent logic $\mathcal{L}^*$. The language of $\mathcal{L}^*$ is the same of $\mathcal{L}$, augmented with a set of new modal operators $\{B_1, \ldots, B_n\}$ to represent the beliefs of each of the agents 1 …, n. That is, if α is a well-formed formula of the language of $\mathcal{L}$, then it is also a well-formed formula of $\mathcal{L}^*$, and if α is a well-formed formula of $\mathcal{L}^*$, so is $B_i\alpha$ for $i \in \{1, \ldots, n\}$.

To model the states of belief of the agents we use the belief structures introduced by Fagin, Halpern and Vardi [FHV91]. The idea of possible worlds is employed here in a lightly different way of that of traditional Kripke structures. The worlds are defined inductively:

Definition 4 [1] *A 0th-order belief assignment, f_0 is a truth assignment to the primitive propositions. We call $\langle f_0 \rangle$ a 1-ary world, (since its "length" is 1).*

Assume inductively that k-ary worlds (or k-worlds, for short) have been defined. Let W_k be the set of all k-worlds. A **k-th-order belief assignment** *is a function $f_k : A \rightarrow 2^{W_k}$, where A is the set of agents. A* **(k+1)-sequence** *of belief assignment is a sequence $\langle f_0, \ldots, f_k \rangle$, where f_i is a i-th-order belief assignment.*

A **(k+1)-world** *is a (k+1)-sequence of belief assignments that satisfy the following restrictions for each agent i:*

1. *$f_k(i)$ is nonempty if $k \geq 1$.*
2. *If $\langle g_0, \ldots, g_{k-1} \rangle \in f_k(i)$ and $k > 1$, then $g_{k-1}(i) = f_{k-1}(i)$.*
3. *$\langle g_0, \ldots, g_{k-2} \rangle \in f_{k-1}(i)$ iff there is a (k-1)st-order belief assignment g_{k-1} such that $\langle g_0, \ldots, g_{k-1} \rangle \in f_k(i)$, if $k > 1$.*

Intuitively, a 1-ary world is a description of reality and a belief assignment f_k associates with each agent a set of "possible k-worlds"; that is, the worlds in $f_k(i)$ are the k-worlds that agent i thinks are possible descriptions of the reality (as described by k-worlds).

The belief structures are defined based on these k-worlds.

Definition 5 *An infinite sequence $\langle f_0, f_1, \ldots \rangle$ is called a belief structure if each prefix $\langle f_0, \ldots, f_k \rangle$ is a k-world, for each k.*

Thus a k-world describes beliefs of depth $k - 1$, where the depth of a formula is roughly the number of nested belief operators in the formula. A belief structure describes beliefs of arbitrary depth.

For a formula α, the worlds that have the information needed to evaluate it are the k-worlds such that $k \geq$ depth(α).

Definition 6 *A $(k+1)$-world $\langle f_0, \ldots, f_k \rangle$ satisfies a formula α, written $\langle f_0, \ldots, f_k \rangle \models \alpha$ if $k \geq depth(\alpha)$ and:*

- *$\langle f_0, \ldots, f_k \rangle \models \alpha$ iff $f_0 \models \alpha$ for α a propositional formula.*
- *$\langle f_0, \ldots, f_k \rangle \models \neg\alpha$ iff $\langle f_0, \ldots, f_k \rangle \not\models \alpha$;*
- *$\langle f_0, \ldots, f_k \rangle \models \alpha \wedge \beta$ iff $\langle f_0, \ldots, f_k \rangle \models \alpha$ and $\langle f_0, \ldots, f_k \rangle \models \beta$;*
- *$\langle f_0, \ldots, f_k \rangle \models B_i\alpha$ iff for each $\langle g_0, \ldots, g_{k-1} \rangle \in f_k(i)$, $\langle g_0, \ldots, g_{k-1} \rangle \models \alpha$*

But it suffices to consider the $depth(\alpha)$-worlds.

Proposition 1 *Assume that $depth(\alpha) = r$ and $k \geq r$. Then, $\langle f_0, \ldots, f_k \rangle \models \alpha$ iff $\langle f_0, \ldots, f_r \rangle \models \alpha$.*

In view of the last result the notion of satisfiability of a formula is captured by the following definition:

[1] The definition of belief-world is similar to the one presented in [FHV91] with the exception of restriction 1. This change yields a logic in which each B_i follows the axioms of KD45 (for belief) instead of [FHV91]'s S5 axioms (for knowledge).

Definition 7 *We say that the belief structure* $f = \langle f_0, f_1, \ldots \rangle$ *satisfies* α, *written* $f \models \alpha$, *if* $\langle f_0, \ldots, f_r \rangle \models \alpha$, *where* $r = depth(\alpha)$.

Definition 8 *We say that a formula* α *is valid, written* $\models \alpha$ *if* $f \models \alpha$ *for each belief structure* f.

The logic $\mathcal{L}^*$ just defined has the following properties.

Proposition 2

- *All substitution instances of the tautologies are valid.*
- $\models B_i\alpha \rightarrow \neg B_i \neg \alpha$.
- $\models B_i\alpha \rightarrow B_i B_i \alpha$.
- $\models \neg B_i\alpha \rightarrow B_i \neg B_i \alpha$.
- $\models B_i\alpha \wedge B_i(\alpha \rightarrow \beta) \rightarrow B_i\beta$.

Thus, each of B_i satisfy the requirements of the logic KD45, which are the most accepted logic for belief.

3 The logic $\mathcal{L}^*_{\preceq}$

3.1 The $\preceq$ preference relation

Given the monotonic multi-agent logic $\mathcal{L}^*$, if we define a preference relation $\preceq$ on the set of belief structures of $\mathcal{L}^*$, then we would have a multi-agent preference logic. This section describes how to define a preference relation $\preceq$ based on a preference relation $\leq$ defined on the set of propositional interpretation of $\mathcal{L}$.

We would like to use the inductive definition of k-worlds to define a preference relation among them. Clearly, $\leq$ is exactly the preference relation defined for 1-worlds (which are just a propositional assignment). The difficulty is that a 2-world attributes to each agent, a *set of* 1-worlds. Thus we have to define a preference relation among sets of 1-worlds, given that there is a preference $\leq$ among 1-worlds. Or in general, given a pre-order $\leq$ on a set A, we must define a pre-order $\sqsubseteq$ on the set 2^A.

We will follow idea of elementary improvement discussed in [Wai93]. The set A_1 is an elementary improvement ($\sqsubset_e$) over the set A_2, if they agree in all worlds, except for a worlds in A_1 which is better (under the $\leq$ order) than its correspondent in A_2. More formally, $A_1 \sqsubset_e A_2$, if $A_1 = A \cup \{a_1\}$ and $A_2 = A \cup \{a_2\}$ and $a_1 \leq a_2$; or $A_2 = A_1 \cup \{a_2\}$ and there is a a_1 in A_1 such that $a_1 \leq a_2$[2].

We are interested in the transitive, reflexive closure of the $\sqsubset_e$ relation, which has a rather simple formulation:

[2] This last case corresponds to the situation where the world that "got better" in the passage from A_2 to A_1 "landed" on a world that was "already there".

Definition 9 *Let* $\leq$ *a pre-order defined on a set* A. *We define the relation* $\sqsubseteq$ *on the set* 2^A *as:* $A_1 \sqsubseteq A_2$ *iff for all* $a \in A_1$ *there is* $b \in A_2$ *such that* $a \leq b$ *and for all* $b \in A_2$ *there is* $a \in A_1$ *such that* $a \leq b$.

It can be shown that $\sqsubseteq$ is a pre-order in 2^A.

Now we can inductively define the pre-order $\preceq_k$ among k-worlds.

Definition 10 *Let* $\leq$ *be a pre-order defined on the set of propositional interpretations, then one can inductively define* $\preceq_k$ *as a pre-order on the set of k-worlds as:*

- $\preceq_1 = \leq$.
- *if* $\preceq_k$ *has been defined on the set of k-worlds, then* $\preceq_{k+1}$ *is defined on the set of* $k+1$*-worlds as:*

  $\langle f_0, \ldots, f_k \rangle \preceq_{k+1} \langle f'_0, \ldots, f'_k \rangle$ *iff* $\langle f_0, \ldots, f_{k-1} \rangle \preceq_k \langle f'_0, \ldots, f'_{k-1} \rangle$ *and for each* i *such that* $1 \leq i \leq n$, $f_k(i) \sqsubseteq_k f'_k(i)$, *where* $\sqsubseteq_k$ *is defined from* $\preceq_k$, *as in definition 9.*

And finally we can define the preference relation $\preceq$ among belief structures:

Definition 11 *Let* $\leq$ *a pre-order on a set of propositional interpretations. If* f *and* f' *are belief structures, then* $f \preceq f'$ *iff for each prefix* $\langle f_0, \ldots, f_{k-1} \rangle$ *of* f *and for each prefix* $\langle f'_0, \ldots, f'_{k-1} \rangle$ *of* f', $\langle f_0, \ldots, f_{k-1} \rangle \preceq_k \langle f'_0, \ldots, f'_{k-1} \rangle$.

$\mathcal{L}^*$ together with the pre-order $\preceq$ defines a multi-agent preferential logic $\mathcal{L}^*_{\preceq}$ that allows one to model the beliefs of agents that can reason non-monotonically.

4 Properties of the logic $\mathcal{L}^*_{\preceq}$

The logic $\mathcal{L}^*_{\preceq}$ defined above have some interesting properties. The proofs of the theorems in this section and the next can be found at [Mon96].

The first important and potentially controversial property of the logic is that all agents have the same reasoning power. In particular, when reasoning about formulas that do not contain modal operators, all agents use the preference logic defined by $\leq$. Section 4.2 discusses the implications of this theorem, and shows that the assumption that all agents reason with the same preference does not limit the applicability of the logic $\mathcal{L}^*_{\preceq}$ in a multi-agent situation.

Theorem 1 *If* α *and* β *are propositional formulas then* $\alpha \models_{\leq} \beta$ *iff* $B_i\alpha \models_{\preceq} B_i\beta$.

The next theorem states that it is common knowledge that all agents reason (about the world) based on the same propositional preference:

Theorem 2 *If* $\alpha \models_{\leq} \beta$ *then* $B_{i_1} \ldots B_{i_r}\alpha \models_{\preceq} B_{i_1} \ldots B_{i_r}\beta$*, whenever* $r > 0$*.*

Another result is related to the independence of the reasoning performed within different belief contexts.

Theorem 3 *If* $\alpha \models_{\leq} \beta$ *and* $\gamma \models_{\leq} \delta$ *then* $B_i\alpha \wedge B_j\gamma \models_{\preceq} B_i\beta \wedge B_j\delta$ *if* $i \neq j$*.*

This states that inferences that could be performed independently in $\mathcal{L}_{\leq}$ can still be performed if they are in different belief contexts. In particular, it could be the case that in $\mathcal{L}_{\leq}$, α and γ could interfere with each other, for example $\alpha \wedge \gamma \models_{\leq} \neg\beta$, but that interference is prevented by the different belief contexts.

Theorem 3 can be extended to arbitrary belief contexts. To state the extended theorem we need to define a syntactic operation on sequences of B_i operators: $\overline{Q}$ is a operation that substitutes in Q all subsequences of adjacent B_i by a single B_i. For example:

$$\overline{B_3B_3B_7B_7B_7B_2B_2B_7} = B_3B_7B_2B_7$$

Theorem 4 *If* $\alpha \models_{\leq} \beta$ *and* $\gamma \models_{\leq} \delta$ *then* $Q_a\alpha \wedge Q_b\gamma \models_{\preceq} Q_a\beta \wedge Q_b\delta$ *where* Q_a *and* Q_b *are any sequence of* B_i *operators, provided that* $\overline{Q_a} \neq \overline{Q_b}$*.*

The intuition for the restriction on Q_a and Q_b is that if $\overline{Q_a} = \overline{Q_b}$ then $Q_a\alpha$ is equivalent to $Q_b\alpha$.

4.1 Mixed defaults

The theorems above state generic properties of the $\mathcal{L}^*_{\preceq}$ logic, independently of the $\leq$ base preference relation. We will now show that if the base logic is propositional circumscription then default rules that combine subformulas with and without the modal operators are correctly dealt with. We call such default rules as **mixed defaults**.

If $\leq_a$ is the preference that minimizes the propositional symbol abn_1, and if $\preceq_a$ is the corresponding preference for the $\mathcal{L}^*$ models, then the following results are true:

- $B_i\alpha \wedge (B_i\alpha \wedge \neg abn_1 \rightarrow \beta) \models_{\preceq_a} \beta$.
- $\alpha \wedge (\alpha \wedge \neg abn_1 \rightarrow B_i\beta) \models_{\preceq_a} B_i\beta$.
- $\alpha \wedge (\alpha \wedge \neg abn_1 \rightarrow B_i\beta) \wedge B_i(\neg\beta) \not\models_{\preceq_a} B_i\beta$

In fact, results above hold when B_i is substituted by any sequence $B_{i_1}B_{i_2} \ldots B_{i_r}$ for $r > 0$. The first two examples show that if the modal sub-formula is either the antecedent or the consequent of a default rule, and the antecedent is true, then the default consequent would be inferred. The last example shows that the logic when determining if a default consequence can be consistently asserted, will take into consideration the logic of the B_i operators.

4.2 Discussion

We must address now the question brought up by theorem 1 which states that all agents reason propositionally using the same preference $\leq$. This may seems too strict, and one could think that this assumption would limit the usefulness of the logic for multi-agent applications. The underlying question is: if all the agents reason with the same preference then it is not the case that they are all the "same" agent.

First we would like to point out that the use of the term preference in this paper refers to a pre-order relation among models, and not a relation among propositions. Thus the logic presented here does not model statements like "Agent 1 prefers to have a toothache than to go to the dentist," which one would certainly not want to be the same for all agents. The preference relation in this paper is the implementation of a form of nonmonotonic reasoning. It is not a form of representing this preference aspect of the mental model of agents ([WD91, Wai94] are attempts to represent this mental aspect).

As for the fact that all the agents reason based on the same preference, one should notice that either there is a single preference relation for all the agents, or there is an infinite number of them. It is not just that agent 1 could use a different preference than agent 2, but also that agent 1 could attribute to agent 2 a different preference, which could be also different than the preference agent 1 thinks agent 2 thinks agent 1 uses, and so on. In other words, one would need a preference to reason within the context of a single B_1, possibly another for reasoning within the context of B_1B_2, possibly another to reason within the context of $B_1B_2B_1$, and so on.

There are applications of preference logic, specially temporal reasoning [Sho88, Bak91, Sha95], for which it would be acceptable to propose that a single general preference relation that apply to all agents. For temporal reasoning in particular, these logics usually propose a particular preference relation that minimizes "meaningful" predicates or formulas in such a way that reasoning about time and change can be correctly performed. If that particular preference relation is the essence of common-sense reasoning about time and change it is very acceptable to assume that all agents use that same preference to reason about time and change.

Other applications of preference logic, such as taxonomic reasoning and default reasoning, which is the main interest of this work, could possibly benefit from multiple preferences. The crucial difference is that preference logics are used as an *implementation* of default or taxonomic reasoning. One minimizes "meaningless" or arbitrary propositional symbols and thus there is no claim that the preference is unique. For example $\alpha \wedge \neg abn_1 \rightarrow \beta$ implements the default rule that if α then normally β, if the preference relation minimizes abn_1. But so does $\alpha \wedge \neg abn_2 \rightarrow \beta$, if the preference relation minimizes abn_2. Since the preference relation is in some way arbitrary, it is unreasonable to expect that all agents use the same one.

Surprisingly, the logic $\mathcal{L}^*_{\preceq}$ can deal with a finite number of multiple preference relations. This is done by bringing this finite set of these preferences into the base-logic preference relation $\leq$, provided that the preferences are not contradictory. For example, let us suppose that agent 1 reasons by minimizing the propositional symbol abn_1, and he believes agent 2 reasons by minimizing the preference abn_2 and that he thinks agent 2 thinks he reasons by minimizing abn_3. One can construct a propositional preference relation $\leq$ that simultaneously minimizes abn_1, abn_2, and abn_3. If all formulas within the scope of a single B_1 do not contain either abn_2 or abn_3, then reasoning based on the $\leq$ relation would derive the correct conclusions for the beliefs of agent 1. Similarly if all formulas within the scope of B_1B_2 do not contain either abn_1 or abn_3 then again all reasoning performed in that belief context would be the same as having a particular preference of minimizing abn_2 applied for only that context.

Finally, an important part of default reasoning using preference logic, in particular circumscription, is defining the circumscription policy, that is, which predicates are allow to vary in the minimization process and so on. In a propositional preference logic, circumscription policy is restricted to setting up the priority among the propositional symbols that will be minimized. By "bringing" all these preferences into the $\leq$ relation, one is able to set different priorities for each "particular preference" as to allow the correct chaining of inferences across beliefs contexts, as the example 2 discussed in the section below illustrates, which would not be possible if all preferences were isolated.

5 Examples

Example 1. In order to illustrate an example of multi-agent preferential reasoning, let us assume the following situation. Agent 1 are sitting with a friend (agent 2) in a cafe and waiting for third friend (agent 3). If agent 3 leaves work at 5 (which we will abbreviate as $l5$) she will normally arrive at the cafe at 6 ($a6$), but not so if there is a traffic jam (tj). Furthermore, let assume that all agents know these facts.

The first problem the logic developed herein solves is a representational problem. One can represent, using a preference logic, the default rule that leaving at 5 will normally entail arriving at 6 as:

$$l5 \wedge \neg abn_1 \rightarrow a6 \tag{1}$$

if the preferential logic is based on the preference relation $\leq_x$ that circumscribes the propositional symbol abn_1. But it is not clear how to represent that agent 1, or agent 2 believes (1) in such a way that the default would work correctly. The main result of this work is that one can develop a multi-agent preference logic based on an already understood propositional preference logic, and one of the results is that each agent in the society will reason (about the real world) using the propositional preference logic. In the logic $\mathcal{L}^*_{\preceq_x}$, where the preference relation $\preceq_x$ is derived from $\leq_x$, one can represent the fact that agent 2 believes (1), as one would expect:

$$B_2(l5 \wedge \neg abn_1 \rightarrow a6)$$

and this representation works correctly, that is

$$B_2(l5) \;\wedge\; B_2(l5 \wedge \neg abn_1 \rightarrow a6) \;\models_{\preceq_x}\; B_2(a6)$$

Example 1 (cont). Now let us suppose that agent 3 left at 5 and all the agents know it. Agent 2 believes that there is a traffic jam, and he believes that 3 does not know it. On the other hand, agent 1 believes that there is no traffic jam. From this situation one can conclude that

- agent 1 believes agent 3 will arrive at six,
- agent 2 believes agent 3 will not
- agent 2 also believes that agent 3 thinks she will arrive at six .

This reasoning can be captured by the logic $\mathcal{L}^*_{\preceq_x}$.

Let us make the following abbreviation where α is the knowledge that if agent 3 leaves at 5, she would usually arrive by 6, unless there is a traffic jam.

$$\alpha \quad \equiv \quad (l5 \wedge \neg abn_1 \rightarrow a6) \wedge (l5 \wedge tj \rightarrow \neg a6)$$

Then, the situation above is captured by:

$$B_1 \begin{pmatrix} l5 \wedge \\ \neg tj \wedge \\ \alpha \end{pmatrix} \wedge B_2 \begin{pmatrix} l5 \wedge \\ tj \wedge \\ \alpha \end{pmatrix} \wedge B_2B_3 \begin{pmatrix} l5 \wedge \\ \alpha \end{pmatrix}$$
$$\models_{\preceq_x} \; B_1(a6) \wedge B_2(\neg a6) \wedge B_2B_3(a6)$$

The example above illustrate that defaults that are contained within a B_i operator work as expected.

Example 1 (another solution). The solution above uses the same abnormal symbol for all belief context, or in other words, all agents use the same preference, the one that minimizes abn_1. Let us develop another solution, where each agent uses its own abnormal propositional symbol and its "own" preference. Let us use the following abbreviations:

$$\begin{aligned} \alpha &\equiv (l5 \wedge \neg abn_1 \rightarrow a6) \wedge (l5 \wedge tj \rightarrow \neg a6) \\ \alpha' &\equiv (l5 \wedge \neg abn_2 \rightarrow a6) \wedge (l5 \wedge tj \rightarrow \neg a6) \\ \alpha'' &\equiv (l5 \wedge \neg abn_3 \rightarrow a6) \wedge (l5 \wedge tj \rightarrow \neg a6) \end{aligned}$$

Now the situation described in the second part of Example 1 is represented as:

$$B_1 \begin{pmatrix} l5 \wedge \\ \neg tj \wedge \\ \alpha \end{pmatrix} \wedge B_2 \begin{pmatrix} l5 \wedge \\ tj \wedge \\ \alpha' \end{pmatrix} \wedge B_2B_3 \begin{pmatrix} l5 \wedge \\ \alpha'' \end{pmatrix} \tag{2}$$

In order to define a multi-agent preference logic, one need to bring those three preferences (minimizing abn_1 in the belief context B_1, minimizing abn_2 in the context B_2 and abn_3 in the context B_2B_3) into a single propositional preference. This is done by defining the propositional preference $\leq_y$ which minimizes in parallel (that is with the same priority) abn_1, abn_2, and abn_3. That is, $M_1 \leq_y M_2$ iff:

- for all propositional symbols p in the language other than abn_1, abn_2, abn_3,
  mbox $M_1 \models p$ iff $M_2 \models p$.
- there is no i such that $M_1 \models abn_i$ and $M_2 \models \neg abn_i$

Once $\leq_y$ is defined, the logic $\mathcal{L}^*_{\preceq_y}$ would allow one to derive the correct conclusions from (2), that is:

$$(2) \quad \models_{\preceq_y} \quad B_1(a6) \wedge B_2(\neg a6) \wedge B_2B_3(a6)$$

Example 2. Let us see now that mixed defaults also work as expected. Let us add to the common knowledge the fact that if it is raining (ra) then usually there is a traffic jam, and the fact that if it is raining then agent 2 knows it (because he hears the noise).

We will use the same scheme as the first solution to example 1, that is to use the same abnormal propositional function in all contexts where it is appropriate. Let us make the following abbreviations:

$$\begin{aligned} \alpha &\equiv (l5 \wedge \neg abn_1 \rightarrow a6) \wedge (l5 \wedge tj \rightarrow \neg a6) \\ \beta &\equiv (ra \wedge \neg abn_2 \rightarrow tj) \\ \gamma &\equiv (ra \wedge \neg abn_3 \rightarrow B_2 ra) \end{aligned}$$

α is the statement that leaving at 5 normally results in arriving at 6, unless there is a traffic jam, β states that if it is raining then normally there is a traffic jam, and γ states that it it is raining then normally agent 2 will know it.

Let us now suppose that agent 1 believes (or sees) that it is raining, then

- he concludes that agent 2 knows it (by γ)
- he concludes that agent 2 believes there is a traffic jam (by β)
- he concludes that agent 2 believes that agent 3 will not arrive at six.

For the reasoning above to work one needs that defaults represented by β and γ to be stronger than the competing default α. But it is important also to notice that these defaults will be operating in different belief contexts: γ is used inside the B_1 operator and β and α are used inside the B_1B_2 sequence of operators.

If one had a preference for each belief context, it is not clear how to set up the priorities among the defaults. But in our approach these preferences are brought together at the $\leq$ preference relation. If one wants γ and β to be stronger than α, one defines a preference relation $\leq_z$ where the propositional symbols abn_2 and abn_3 are minimized with higher priority than abn_1. The logic that extends the relation $\leq_z$ will yield the correct deductions:

$$\begin{gathered} B_1 ra \wedge B_1(ra \wedge \neg abn_3 \rightarrow B_2 ra) \wedge \\ B_1B_2(l5) \wedge B_1B_2(\beta) \\ \models_{\preceq_z} \; B_1B_2(\neg a6) \end{gathered}$$

This example illustrates a form of chaining across belief contexts: defaults in the B_1's belief space ($ra \wedge (ra \wedge \neg abn_3 \rightarrow B_2 ra)$) trigger defaults in agent 2 belief space ($B_2(l5 \wedge \beta)$). In agent 2 belief space the default that $l5$ normally entail $a6$ is disabled by the higher priority of the default that if ra then $\neg a6$.

6 Conclusions and Future work

As far as the authors know this is the first preference based nonmonotonic logic for multi-agents. The other examples of non-monotonic multi-agent logics [Mor90, Lak93, HM93, PJ95] follow the autoepistemic tradition. One important difference between the preferential and autoepistemic logics is that the former allows one to easily introduce priorities among the defaults, as we did in example 2. It will be an interesting extension of this work to compare a flat version of the preference multi-agent logic (that is, where all minimizations are done in parallel) with some of these autoepistemic logics.

The work presented here is still preliminary. First, we need to discover other properties of the logic $\mathcal{L}^*_{\preceq}$. For example, we cannot yet present any result on how lack of knowledge ($\neg B_i$) interacts with the preference ordering.

It is also important to extend the results for a quantified logic since preference logics have naturally been based on quantified logics.

Acknowledgments

The first author was supported by Universidad Nacional de La Plata and by a scholarship from CNPq, and the second author was partially supported by CNPq grant 521464/93-5. This work is part of the CNPq PROTEM III project LOGIA.

References

[Bak91] A. B. Baker. Nonmonotonic reasoning in the framework of situation calculus. *Artificial Intelligence*, 49:5–23, 1991.

[FHV91] R. Fagin, J. Y. Halpern, and M. Y. Vardi. A model-theoretic analysis of knowledge. *Journal of the A.C.M.*, 32(2):382–428, 1991.

[HM92] J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54(3):319–379, 1992.

[HM93] J. Y. Halpern and Y. O. Moses. Reasoning about only knowing with many agents. In *Proc. of the 11th National Conference on Artificial Intelligence (AAAI)*, pages 655–661, 1993.

[Lak93] G. Lakemeyer. All they know: a study in multi-agent autoepestemic reasoning. In *Proc. Thirteenth International Joint Conference on Artificial Intelligence (IJCAI '93)*, 1993. Unpublished manuscript.

[McC80] J. McCarthy. Circumscription—a form of non-monotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980.

[McC86] J. McCarthy. Applications of circumscription to formalizing common-sense reasoning. *Artificial Intelligence*, 28:89–116, 1986.

[Mon96] A. M. Monteiro. Estensões epsitemicas de duas familias de logicas não monotonicas Master thesis *in portuguese.* Instituto de Computação, UNICAMP, Campinas 13083-970, Brazil, 1996

[Mor90] L. Morgenstern. A theory of multiple agent nonmonotonic reasoning. In *Proc of the AAAI-90*, pages 538–544, 1990.

[PJ95] Y. Permpoontanalarp and J. Y. Jiang. Generalized proof-theory for multi-agent autoepistemic reasoning. In *Proceedings of the ICMAS*, pages 304–311, 1995.

[Sha95] M. Shanahan. A circumscriptive calculus of events. *Artificial Intelligence*, 77(2):249–284, 1995.

[Sho87] Y. Shoham. A semantical approach to nonmonotonic logics. In *Proc. of the Tenth International Joint Conference on Artificial Intelligence*, pages 388–392, 1987.

[Sho88] Y. Shoham. *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence.* MIT Press, Cambridge, MA, 1988.

[Wai93] J. Wainer. Epistemic extension of preference logics. In *Proc. of the 13th International Joint Conference on Artificial Intelligence*, pages 382–387, 1993.

[Wai94] J. Wainer Yet another semantics of goals and goal priorities. In *Proc. of the 11th European Conference on Artificial Intelligence*, pages 269-273, 1994.

[WD91] M. P. Wellman and J. Doyle. Preferential semantics for goals. In *Proc. of the 9th National Conference on Artificial Intelligence*, pages 698-703, 1991.