Algorithms for a Nonmonotonic Logic of Preferences

Souhila Kaci[1] and Leendert van der Torre[2]

[1] Centre de Recherche en Informatique de Lens (C.R.I.L.)–C.N.R.S.
Rue de l'Université SP 16, 62307 Lens Cedex, France
[2] CWI Amsterdam and Delft University of Technology, The Netherlands

Abstract. In this paper we introduce and study a nonmonotonic logic to reason about various kinds of preferences. We introduce preference types to choose among these kinds of preferences, based on an agent interpretation. We study ways to calculate "distinguished" preference orders from preferences, and show when these distinguished preference orders are unique. We define algorithms to calculate the distinguished preference orders.

Keywords: logic of preferences, preference logic

1 Introduction

Preferences guide human decision making from early childhood (e.g., "which ice cream flavor do you prefer?") up to complex professional and organisatioral decisions (e.g., "which investment funds to choose?"). Preferences have traditionally been studied in economics and applied to decision making problems. Moreover, the logic of preference has been studied since the sixties as a branch of philosophical logic. Preferences are inherently a multi-disciplinary topic, of interest to economists, computer scientists, OR researchers, mathematicians, logicians, philosophers, and more.

Preferences are a relatively new topic to artificial intelligence and are becoming of greater interest in many areas such as knowledge representation, multi-agent systems, constraint satisfaction, decision making, and decision-theoretic planning. Recent work in AI and related fields has led to new types of preference models and new problems for applying preference structures [1]. Explicit preference modeling provides a declarative way to choose among alternatives, whether these are solutions of problems to solve, answers of data-base queries, decisions of a computational agent, plans of a robot, and so on. Preference-based systems allow finer-grained control over computation and new ways of interactivity, and therefore provide more satisfactory results and outcomes. Logics of preference are used to compactly represent and reason about preference relations.

A particularly challenging topic in preference logic is concerned with non-monotonic reasoning about preferences. A few constructs have been proposed [6, 14, 11], for example based on mechanisms developed in non-monotonic reasoning such as gravitation towards the ideal, or compactness, but there is no consensus yet in this area. Nevertheless, non-monotonic reasoning about preferences is an important issue, for example when reasoning under uncertainty. When an agent compactly communicates its preferences, another agent has to interpret it and find the most likely interpretation.

A drawback of the present state of the art in the logic of preference is that proposed logics typically formalize only preferences of one kind, formalizing for example strong preferences, defeasible preferences, non-strict preferences, ceteris paribus preferences (interpreted either as "all else being equal" or as "under similar circumstances"), etc. These logics formalize logical relations among one kind of preferences, but relations among distinct kinds of preferences have not been considered. Consequently, when formalizing preferences, one has to choose which kind of preference statements are used for all preferences under consideration. However, often we would like to use several kinds of preference statements at the same time.

We are interested in developing and using a logic with more than one kind of preferences, which we call a logic of preferences – in contrast to the usual reference to the logic of preference. In particular we are interested in nonmonotonic logic of preferences. To interpret the various kinds of preferences we use total pre-orders on worlds, which we call preference orders. We consider the following questions:

1. How to define a logic of preferences to reason about for example strong and weak preferences? How are they related to conditional logics?
2. How to choose among kinds of preferences when formalizing examples?
3. How to calculate "distinguished" preference orders from preferences? Are the distinguished preference orders unique?
4. How can we define algorithms to calculate the distinguished preference orders?

To define our logic of preferences, we define four kinds of strict preferences of p over q as "the best/worst p is preferred over the best/worst q". We define conditionals "if p, then q" as usual as a preference of p and q over p and the absence of q.

To choose among kinds of preferences, we introduce an agent interpretation of the four kinds of preferences studied in this paper. We interpret a preference of p over q as a game between an agent arguing for p and an agent arguing for q. We distinguish locally optimistic, pessimistic, opportunistic and careful preference types.

To calculate a preference order from preferences, we start from a generalization of System Z, which is usually characterized as gravitating towards the ideal for defeasible conditionals, and also known as minimal specificity. We also define the inverse of gravitating towards the worst. In general we need to combine both kinds of mechanisms, for which we study a strict dominance of one of the mechanisms. We provide new algorithms to derive distinguished orders.

The layout of this paper is as follows. We treat each question above mentionned in a subsequent section. Section 2 introduces the logic of preferences we use in this paper. Section 3 introduces the preference types. Section 4 introduces the non-monotonic extensions to define distinguished preference orders. Section 5 introduces algorithms to calculate distinguished preference orders.

2 Logic of preferences

The logical language extends propositional logic with four kinds of preferences. A small m stands for $\min$ and a capital M stands for $\max$, as will be explained in the semantics below.

Definition 1 (Language). *Given a set* $A = \{a_1, \ldots, a_n\}$ *of propositional atoms, we define the set* L_0 *of propositional formulas and the set* L *of preference formulas as follows.*

$L_0 \ni p, q$**:** $a_i \mid (p \wedge q) \mid \neg p$
$L \ni \phi, \psi$**:** $p \;{}^m{>}^m q \mid p \;{}^m{>}^M q \mid p \;{}^M{>}^m q \mid p \;{}^M{>}^M q \mid \neg\phi \mid (\phi \wedge \psi)$

Disjunction $\vee$*, material implication* $\supset$ *and equivalence* $\leftrightarrow$ *are defined as usual. Moreover, we define conditionals in terms of preferences by* $p \;{}^m{\to}^m q =_{def} p \wedge q \;{}^m{>}^m p \wedge \neg q$*, etc. We abbreviate formulas using the following order on logical connectives:* $\neg \mid \vee, \wedge \mid > \mid \supset, \leftrightarrow$*. For example,* $p \vee q > r \supset s$ *is interpreted as* $((p \vee q) > r) \supset s$*.*

In the semantics of the four kinds of preferences, a preference of p over q is interpreted as a preference of $p \wedge \neg q$ over $q \wedge \neg p$. This is standard and known as von Wright's expansion principle [16].

Definition 2 (Semantics). *Let* A *be a finite set of propositional atoms,* L *a propositional logic based on* A*,* W *the set of propositional interpretations of* L*, and* $\succeq$ *a total pre-order on* W*. We write* $w \succ w'$ *for* $w \succeq w'$ *without* $w' \succeq w$*, we write* $\max(p, \succeq)$ *for* $\{w \in W \mid w \models p, \forall w' \in W : w' \models p \Rightarrow w \succeq w'\}$*, and we write* $\min(p, \succeq)$ *for* $\{w \in W \mid w \models p, \forall w' \in W : w' \models p \Rightarrow w' \succeq w\}$*.*

$\succeq \models p \;{}^m{>}^m q$ *iff* $\forall w \in \min(p \wedge \neg q, \succeq)$ *and* $\forall w' \in \min(\neg p \wedge q, \succeq)$ *we have* $w \succ w'$
$\succeq \models p \;{}^m{>}^M q$ *iff* $\forall w \in \min(p \wedge \neg q, \succeq)$ *and* $\forall w' \in \max(\neg p \wedge q, \succeq)$ *we have* $w \succ w'$
$\succeq \models p \;{}^M{>}^m q$ *iff* $\forall w \in \max(p \wedge \neg q, \succeq)$ *and* $\forall w' \in \min(\neg p \wedge q, \succeq)$ *we have* $w \succ w'$
$\succeq \models p \;{}^M{>}^M q$ *iff* $\forall w \in \max(p \wedge \neg q, \succeq)$ *and* $\forall w' \in \max(\neg p \wedge q, \succeq)$ *we have* $w \succ w'$

Moreover, logical notions are defined as usual, in particular:

– $\succeq \models \{\phi_1, \ldots, \phi_n\}$ *iff* $\succeq \models \phi_i$ *for* $1 \leq i \leq n$*,*
– $\models \phi$ *iff for all* $\succeq$*, we have* $\succeq \models \phi$*,*
– $S \models \phi$ *iff for all* $\succeq$ *such that* $\succeq \models S$*, we have* $S \models \phi$*.*

The ${}^m{>}^M$'s preference is the strongest one while ${}^M{>}^m$'s preference is the weakest one [15]. The following example illustrates the logic of preferences.

Example 1. We have $\models p \;{}^M{>}^M q \leftrightarrow (p \wedge \neg q) \vee (\neg p \wedge q) \;{}^M{\to}^M p$, which expresses a well-known relation between a defeasible conditional ${}^M{\to}^M$ and preferences ${}^M{>}^M$. Moreover, we have $\models p \;{}^m{>}^M q \supset p \;{}^M{>}^M q$, which expresses that strong preferences ${}^m{>}^M$ imply defeasible preferences ${}^M{>}^M$.

The following definition illustrates how a preference order – represented in a *qualitative* form by a total pre-order $\succeq$ on worlds – can also be represented by a well ordered partition of W. This is an equivalent representation, in the sense that each preference order corresponds to one ordered partition and vice versa. This equivalent representation as an ordered partition makes some definitions easier to read.

Definition 3 (Ordered partition). *A sequence of sets of worlds of the form* $(E_1, \cdots, E_n)$ *is an ordered partition of* W *iff* $\forall i$*,* E_i *is nonempty,* $E_1 \cup \cdots \cup E_n = W$ *and* $\forall i, j$*,* $E_i \cap E_j = \emptyset$ *for* $i \neq j$*. An ordered partition of* W *is associated with pre-order* $\succeq$ *on* W *iff* $\forall \omega, \omega' \in W$ *with* $\omega \in E_i, \omega' \in E_j$ *we have* $i \leq j$ *iff* $\omega \succeq \omega'$*.*

3 Preference types as agent types

The logic of preferences now forces us to choose among the four kinds of preferences when we formalize an example in the logic. From the literature it is only known how to choose among monopolar preferences such as "I prefer p", or more involved "Ideally p", "p is my goal", "I desire p", "I intend p", etc. In such cases we can distinguish two notions of lifting worlds to sets of worlds.

Definition 4 (Agent types for the lifting problem). *Let S be a set ordered by a total pre-order $\succeq$. The lifting problem is the selection of an element of S. We define the following agent types for the lifting problem:*

- **Optimistic agent:** *The agent selects the elements of S which are maximal w.r.t. $\succeq$.*
- **Pessimistic agent:** *The agent selects the elements of S which are minimal w.r.t. $\succeq$.*

However, this cannot directly be used for our four kinds of preferences, due to the bipolar representation of preferences. To choose among these kinds of preferences, we introduce an agent interpretation of preferences. We interpret a preference of p over q as a game between an agent arguing for p and an agent arguing for q. Thus, the agent argues that p is better than q against a (possibly hypothetical) opponent.

Example 2. Assume an agent is looking for a flight ticket on the web, and it prefers web-service FastTicket to web-service TicketNow. If the agent is opportunistic, it is optimistic about FastTicket and pessimistic about TicketNow, but when it is careful, it is pessimistic about FastTicket, and optimistic about TicketNow. Clearly, an opportunistic agent has many preferences, whereas a careful agent has only a few preferences.

Preference types can now be defined in terms of agent types.

Definition 5 (Preference types). *Consider an agent expressing its preference of p over q. We define the following preference types:*

- **Locally optimistic:** *the agent is optimistic about p and optimistic about q.*
- **Locally pessimistic:** *the agent is pessimistic about p and pessimistic about q.*
- **Opportunistic:** *the agent is optimistic about p and pessimistic about q.*
- **Careful:** *the agent is pessimistic about p and optimistic about q.*

The following example illustrates that the preference types are a useful metaphor to distinguish among the kinds of preferences, but that their use should not be taken too far.

Example 3 (Continued). The agent types are very strong, which makes them useful in practice but which also has the consequence that one has to be careful when using them, for example when formalizing examples. This is illustrated by several properties about preference types in the logic. For example, when a careful agent prefers FastTicket to TicketNow, an opportunistic agent with the same preference order holds the same preference. Moreover, if a careful agent prefers FastTicket to TicketNow, then it follows that it cannot hold the inverse preference of TicketNow over FastTicket at the same time. An opportunistic agent, however, can hold both inverse preferences at the same time.

It seems that careful preference type is too weak. However it may be useful when all other preference types give an empty set of models [15]:

Example 4. Let j and f be two propositional variables which stand for marriage with John and Fred, respectively. Let $\mathcal{P}_{xy} = \{\top \;{}^{x}\!\!\rightarrow^{y} j, \top \;{}^{x}\!\!\rightarrow^{y} f, \top \;{}^{x}\!\!\rightarrow^{y} \neg(j \wedge f)\}$ be a set of Sue's preferences about its marriage with John or Fred. $\mathcal{P}_{xy}$ induces the following set of constraints: $\{j \;{}^{x}\!\!>^{y} \neg j, f \;{}^{x}\!\!>^{y} \neg f, \neg(j \wedge f) \;{}^{x}\!\!>^{y} (j \wedge f)\}$. The first constraint means that Sue prefers to be married to John over not being married to him. The second constraint means that Sue prefers to be married to Fred over not being married to him and the last constraint means that Sue prefers not to be married to both. There is no pre-order satisfying any of the sets $\mathcal{P}_{MM}$, $\mathcal{P}_{mM}$ and $\mathcal{P}_{mm}$ while the following pre-order $(\{j\neg f, \neg j f\}, \{jf, \neg j \neg f\})$ satisfies $\mathcal{P}_{Mm}$.

4 Nonmonotonic logic of preferences

We study fragments of the logic that consist of sets of preferences only. We call such sets of preferences a preference specification.

Definition 6 (Preference Specification). *A preference specification is a tuple $\langle \mathcal{P}_{MM}, \mathcal{P}_{Mm}, \mathcal{P}_{mM}, \mathcal{P}_{mm} \rangle$ where $\mathcal{P}_{xy}$ ($xy \in \{MM, Mm, mM, mm\}$) is a set of preferences of the form $\{p_i \;{}^{x}\!\!>^{y} q_i : i = 1, \cdots, n\}$.*

In this section we consider the problem of finding pre-orders $\succeq$ that satisfy each desire of a single set $\mathcal{P}_{xy}$ – i.e., *models* of $\mathcal{P}_{xy}$. In the following section, we consider models of two or more sets of preferences.

Definition 7 (Model of a set of preferences). *Let $\mathcal{P}_{xy}$ be a set of preferences and $\succeq$ be a total pre-order. $\succeq$ is a model of $\mathcal{P}_{xy}$ iff $\succeq$ satisfies each preference $p_i \;{}^{x}\!\!>^{y} q_i$ in $\mathcal{P}_{xy}$.*

Shoham [13] characterizes nonmonotonic reasoning as a mechanism that selects a subset of the models of a set of formulas, which we call distinguished models in this paper. Shoham calls these models "preferred models", but we do not use this terminology as this meta-logical terminology may be confused with preferences in logical language and preference orders in semantics.

In this paper we compare total pre-orders based on the so-called specificity principle. The minimal specificity principle is gravitating towards the least specific pre-order, while the maximal specificity principle is gravitating towards the most specific pre-order. These have been used in non-monotonic logic to define the distinguished model of a set of conditionals of the kind ${}^{M}\!\!\rightarrow^{M}$, sometimes called defeasible conditionals.

Definition 8 (Minimal/Maximal specificity principle). *Let $\succeq$ and $\succeq'$ be two total pre-orders on a set of worlds W represented by ordered partitions $(E_1, \cdots, E_n)$ and $(E'_1, \cdots, E'_n)$ respectively. We say that $\succeq$ is at least as specific as $\succeq'$, written as $\succeq \sqsubseteq \succeq'$, iff $\forall \omega \in W$, if $\omega \in E_i$ and $\omega \in E'_j$ then $i \leq j$. $\succeq$ is said to be the least (resp. most) specific pre-order among a set of pre-orders $\mathcal{O}$ if there is no $\succeq'$ in $\mathcal{O}$ such that $\succeq' \sqsubset \succeq$, i.e., $\succeq' \sqsubseteq \succeq$ without $\succeq \sqsubseteq \succeq'$ (resp. $\succeq \sqsubset \succeq'$).*

The following example illustrates minimal and maximal specificity.

Example 5. Consider the rule $p \ {}^{x}\!\!\rightarrow^{y} q$. Applying the minimal specificity principle on $p \ {}^{M}\!\!\rightarrow^{M} q$ or $p \ {}^{m}\!\!\rightarrow^{M} q$ gives the following model $\succeq = (\{pq, \neg pq, \neg p\neg q\}, \{p\neg q\})$. The preferred worlds in this model are those which do not violate the rule. More precisely pq belongs to the set of preferred worlds since it satisfies the rule but $\neg pq$ and $\neg p\neg q$ are preferred too since they do not violate the rule even if they do not satisfy it. Now applying the maximal specificity principle on $p \ {}^{m}\!\!\rightarrow^{m} q$ gives the following model $\succeq' = (\{pq\}, \{\neg pq, p\neg q, \neg p\neg q\})$. We can see that the preferred worlds are only those which satisfy the rule.

Shoham defines non-monotonic consequences of a logical theory as all formulas which are true in the distinguished models of the theory. An attractive property occurs when there is only one distinguished model, because in that case it can be decided whether a formula non-monotonically follows from a logical theory by calculating the unique distinguished model, and testing whether the formula is satisfied by the distinguished model. Likewise, all non-monotonic consequences can be found by calculating the unique distinguished model and characterizing all formulas satisfied by this model.

Theorem 1. *The following table summarizes uniqueness of distinguished models.*

$\mathcal{P}_{mm}$		$\mathcal{P}_{mM}$		$\mathcal{P}_{Mm}$		$\mathcal{P}_{MM}$	
least	*most*	*least*	*most*	*least*	*most*	*least*	*most*
no	*yes [9]*	*yes [5]*	*yes*	*no*	*no*	*yes [12, 3]*	*no*

Proof. Most of the uniqueness proofs have been given in the literature, as indicated in the table. The only exception is the uniqueness of most specific model of $\mathcal{P}_{mM}$, which can be derived from the uniqueness of the least specific model of $\mathcal{P}_{mM}$. We do not give the details here – it follows from the more general Theorem 3 below. Here we give counterexamples for the uniqueness in the other cases. Let $A = \{p, q\}$ such that we have four distinct worlds.

Non-uniqueness of most specific models of ${}^{M}\!\!>^{M}$:
$\mathcal{P}_{MM}\{p \ {}^{M}\!\!>^{M} \neg p\}$, $\succeq = (\{pq\}, \{p\neg q, \neg pq, \neg p\neg q\})$, $\succeq' = (\{p\neg q\}, \{\neg pq, \neg p\neg q, pq\})$.

Non-uniqueness of least specific models of ${}^{m}\!\!>^{m}$:
$\mathcal{P}_{mm}\{p \ {}^{m}\!\!>^{m} \neg p\}$, $\succeq = (\{pq, p\neg q, \neg pq\}, \{\neg p\neg q\})$, $\succeq' = (\{pq, p\neg q, \neg p\neg q\}, \{\neg pq\})$.

Non-uniqueness of least specific models of ${}^{M}\!\!>^{m}$:
$\mathcal{P}_{Mm}\{p \ {}^{M}\!\!>^{m} \neg p\}$, $\succeq = (\{pq, p\neg q, \neg pq\}, \{\neg p\neg q\})$, $\succeq' = (\{pq, p\neg q, \neg p\neg q\}, \{\neg pq\})$.

Non-uniqueness of most specific models of ${}^{M}\!\!>^{m}$:
$\mathcal{P}_{Mm}\{p \ {}^{M}\!\!>^{m} \neg p\}$, $\succeq = (\{pq\}, \{p\neg q, \neg pq, \neg p\neg q\})$, $\succeq' = (\{p\neg q\}, \{pq, \neg pq, \neg p\neg q\})$

There are two consequences of Theorem 1 which are relevant for us now. First, as we are interested in developing algorithms for unique distinguished models, in the remainder of this paper we only focus on ${}^{M}\!\!>^{M}$, ${}^{m}\!\!>^{M}$ and ${}^{m}\!\!>^{m}$ preference types. Secondly, constraints of the form ${}^{m}\!\!>^{M}$ are in between ${}^{M}\!\!>^{M}$ and ${}^{m}\!\!>^{m}$, in the sense that there is a unique least specific model for ${}^{m}\!\!>^{M}$ and ${}^{M}\!\!>^{M}$, and there is a unique most specific model for ${}^{m}\!\!>^{M}$ and ${}^{m}\!\!>^{m}$.

5 Algorithms for nonmonotonic logic of preferences

We now consider distinguished models of sets of preferences of distinct types. It directly follows from Theorem 1 that our only hope to find a unique least or most specific model of a set of preferences is that we may find a unique least specific model for preferences for constraints of both ${}^{m}\!\!>^{M}$ and ${}^{M}\!\!>^{M}$, and a unique most specific model for ${}^{m}\!\!>^{M}$ and ${}^{m}\!\!>^{m}$. In all other cases we already do not have a unique distinguished model for one of the preferences. However, it does not follow from Theorem 1 that a least specific model of a set of ${}^{m}\!\!>^{M}$ and ${}^{M}\!\!>^{M}$ *together* is unique, and it does not follow from the theorem that a most specific model for ${}^{m}\!\!>^{M}$ and ${}^{m}\!\!>^{m}$ *together* is unique! We therefore consider the two following questions in this section:

1. Is a least specific model of a set of ${}^{m}\!\!>^{M}$ and ${}^{M}\!\!>^{M}$ *together* unique? Is a most specific model for ${}^{m}\!\!>^{M}$ and ${}^{m}\!\!>^{m}$ *together* unique? If so, how can we find these unique models?
2. How can we define distinguished models that consists of all three kinds of preferences?

5.1 $\mathcal{P}_{MM}$ and $\mathcal{P}_{mM}$

The following definition derives a unique distinguished model from $\mathcal{P}_{MM}$ and $\mathcal{P}_{mM}$ together. This algorithm generalizes the algorithms given in [3, 5], in the sense that when one of the sets is empty, we get one of the original algorithms.

Definition 9. *Given two sets of preferences* $\mathcal{P}_{MM} = \{C_i = p_i \ {}^{M}\!\!>^{M} q_i : i = 1, \ldots, n\}$ *and* $\mathcal{P}_{mM} = \{C'_j = p'_j \ {}^{m}\!\!>^{M} q'_j : j = 1, \ldots, n'\}$*, let associated constraints be sets of pairs* $\mathcal{C} = \{(L(C_i), R(C_i))\} \cup \{(L(C'_j), R(C'_j))\}$*, where* $L(C_i) = |p_i \wedge \neg q_i|$, $R(C_i) = |\neg p_i \wedge q_i|$, $L(C'_j) = |p'_j \wedge \neg q'_j|$ *and* $R(C'_j) = |\neg p'_j \wedge q'_j|$ *(where* $|\alpha|$ *is* $\{s \in W \mid w \models \alpha\}$*). Algorithm 1.1 computes a unique distinguished model of* $\mathcal{P}_{MM} \cup \mathcal{P}_{mM}$*.*

Algorithm 1.1: Handling mixed preferences ${}^{M}\!\!>^{M}$ and ${}^{m}\!\!>^{M}$.

```
begin
  l ← 0 ;
  while W ≠ ∅ do
    – l ← l + 1 ;
    – E_l = {ω : ∀(L(C_i), R(C_i)), (L(C'_j), R(C'_j)) ∈ C, ω ∉ R(C_i), ω ∉ R(C'_j)} ;
    if E_l = ∅ then
      Stop (inconsistent constraints)
    – W = W − E_l ;
    – remove from C each (L(C_i), R(C_i)) such that L(C_i) ∩ E_l ≠ ∅ ;
    – replace each (L(C'_j), R(C'_j)) in C by (L(C'_j) − E_l, R(C'_j));
    – remove from C each (L(C'_j), R(C'_j)) such that L(C'_j) is empty;
  return (E_1, · · · , E_l)
end
```

We first explain the algorithm, then we illustrate it by an example, and finally we show that the distinguished model computed is the unique least specific one. At each step of the algorithm, we look for worlds which can have the actual highest ranking in the preference order. This corresponds to the actual minimal value l. These worlds are those which do not appear in any right part of the actual set of constraints $\mathcal{C}$ i.e., they do not falsify any constraint. Once these worlds are selected, the two types of constraints have different treatments:

1. We remove constraints $(L(C_i), R(C_i))$ such that $L(C_i) \cap E_l \neq \emptyset$, because such constraints are satisfied. Worlds in $R(C_i)$ will necessarily belong to E_j with $j > l$, i.e., they are less preferred than worlds in the actual set E_l.
2. Concerning the constraints $(L(C'_j), R(C'_j))$, we reduce their left part by removing the elements of the actual set E_l. While $L(C'_j) \neq \emptyset$, such a constraint is not yet satisfied since the constraint $p'_j \ {}^{m}\!\!>^{M} q'_j$ induces a constraint stating that each $p'_j \wedge \neg q'_j$ world should be preferred to all $\neg p'_j \wedge q'_j$ worlds. A pair $(L(C'_j), R(C'_j))$ is then removed only when $L(C'_j) \subseteq E_l$.

The least specific criterion can be checked by construction. At each step l we put in E_l all worlds which do not appear in any $R(C_i)$ or $R(C'_j)$ and which are not yet put in some E_j with $j < l$. If $\omega \in E_l$, then it necessarily falsifies some constraints which are not falsified by worlds of E_j for $j < l$. If we would put some ω of E_l in E_j with $j < l$, then we get a contradiction.

Example 6. Let r, j and w be three propositional variables which stand respectively for "it rains", "to do jogging" and "put a sport wear". Let $\{\omega_0 : \neg r \neg j \neg w, \omega_1 : \neg r \neg j w, \omega_2 : \neg r j \neg w, \omega_3 : \neg r j w, \omega_4 : r \neg j \neg w, \omega_5 : r \neg j w, \omega_6 : r j \neg w, \omega_7 : r j w\}$. Let $\mathcal{P} = \{C_1 : r \wedge \neg j \ {}^{M}\!\!>^{M} r \wedge j, C_2 : (j \vee r) \wedge w \ {}^{M}\!\!>^{M} (j \vee r) \wedge \neg w, C_3 : \neg j \wedge \neg w \ {}^{m}\!\!>^{M} \neg j \wedge w\}$. The first constraint means that if it rains then the agent prefers to do jogging. The second constraint means that if the agent does jogging or it rains then it prefers to put a sport wear and the third constraint means that if the agent will not do jogging then it prefers to not put a sport wear.

We have $\mathcal{C} = \{(L(C_1), R(C_1)), (L(C_2), R(C_2)), (L(C_3), R(C_3))\}$, i.e., $\{(\{\omega_4, \omega_5\}, \{\omega_6, \omega_7\}), (\{\omega_3, \omega_5, \omega_7\}, \{\omega_2, \omega_4, \omega_6\}), (\{\omega_0, \omega_4\}, \{\omega_1, \omega_5\})\}$. We put in E_1 worlds which do not appear in any $R(C_i)$. Then $E_1 = \{\omega_0, \omega_3\}$. We remove $(L(C_2), R(C_2))$ and replace $(L(C_3), R(C_3))$ by $(L(C_3) - E_1, R(C_3)) = (\{\omega_4\}, \{\omega_1, \omega_5\})$. Then $\mathcal{C} = \{(\{\omega_4, \omega_5\}, \{\omega_6, \omega_7\}), (\{\omega_4\}, \{\omega_1, \omega_5\})$. Now $E_2 = \{\omega_2, \omega_4\}$ so both constraints in $\mathcal{C}$ are removed. Lastly $E_3 = \{\omega_1, \omega_5, \omega_6, \omega_7\}$. Finally, the computed distinguished model of $\mathcal{P}$ is $\succeq = (\{\omega_0, \omega_3\}, \{\omega_2, \omega_4\}, \{\omega_1, \omega_5, \omega_6, \omega_7\})$.

The above algorithm computes the least specific model of $\mathcal{P}_{MM} \cup \mathcal{P}_{mM}$ which is unique. To show the uniqueness property, we follow the line of the proofs given in [4, 5]. We first define the maximum of two preference orders.

Definition 10. *Let $\succeq$ and $\succeq'$ be two preference orders represented by their well ordered partitions $(E_1, \cdots, E_n)$ and $(E'_1, \cdots, E'_{n'})$ respectively. We define the $\mathcal{MAX}$ operator by $\mathcal{MAX}(\succeq, \succeq') = (E''_1, \cdots, E''_{min(n,n')})$, such that $E''_1 = E_1 \cup E'_1$ and $E''_k = (E_k \cup E'_k) - (\bigcup_{i=1,\cdots,k-1} E''_i)$ for $k = 2, \cdots, min(n, n')$, and the empty sets E''_k are eliminated by renumbering the non-empty ones in sequence.*

We put $\mathcal{P} = \mathcal{P}_{MM} \cup \mathcal{P}_{mM}$. Let $\mathcal{M}(\mathcal{P})$ be the set of models of $\mathcal{P}$ in the sense of Definition 7. Given Definition 10, the following lemma shows that the $\mathcal{MAX}$ operator is internal to $\mathcal{M}(\mathcal{P})$.

Lemma 1. *Let $\succeq$ and $\succeq'$ be two elements of $\mathcal{M}(\mathcal{P})$. Then,*

1. *$\mathcal{MAX}(\succeq, \succeq') \in \mathcal{M}(\mathcal{P})$,*
2. *$\mathcal{MAX}(\succeq, \succeq')$ is less specific than $\succeq$ and $\succeq'$,*
3. *If $\succeq^*$ is less specific than both $\succeq$ and $\succeq'$ then it is less specific than $\mathcal{MAX}(\succeq, \succeq')$.*

Proof. The proof of item 1 is given in the appendix. The proofs of item 2 and 3 can be found in [4].

We also have the following Lemma:

Lemma 2. *There exists a unique preference order in $\mathcal{M}(\mathcal{P})$ which is the least specific one, denoted by $\succeq_{spec}$, and defined by: $\succeq_{spec} = \mathcal{MAX}\{\succeq : \succeq \in \mathcal{M}(\mathcal{P})\}$.*

Proof. From point 1 of Lemma 1, $\succeq_{spec}$ belongs to $\mathcal{M}(\mathcal{P})$. Suppose now that $\succeq_{spec}$ is not unique. This means that there exists another preference order $\succeq^*$ which also belongs to $\mathcal{M}(\mathcal{P})$ and $\succeq_{spec}$ is not less specific than $\succeq^*$. Note that $\succeq_{spec}$ is the result of combining elements of $\mathcal{M}(\mathcal{P})$ using the $\mathcal{MAX}$ operator. Now supposing that $\succeq_{spec}$ is not less specific than $\succeq^*$ contradicts point 2 of Lemma 1.

We can now conclude:

Theorem 2. *Algorithm 1.1 computes the least specific model of $\mathcal{M}(\mathcal{P})$.*

Proof. Following Lemma 1 it computes a preference order which belongs to the set of the least specific models and following Lemma 2, this preference order is unique.

5.2 $\mathcal{P}_{mm}$ and $\mathcal{P}_{mM}$

Algorithm 1.2. computes a distinguished model of $\mathcal{P}_{mM} \cup \mathcal{P}_{mm}$. This algorithm is structurally similar to Algorithm 1.1., and the proof that this algorithm produces the most specific model of these preferences is analogous to the proof of Theorem 2.

Let $\mathcal{P}_{mm} = \{C_i = p_i \;{}^{m}\!\!>^{m} q_i : i = 1, \cdots, n\}$ and $\mathcal{P}_{mM} = \{C'_j = p'_j \;{}^{m}\!\!>^{M} q'_j : j = 1, \cdots, n'\}$. Let $\mathcal{C} = \{(L(C_i), R(C_i))\} \cup \{(L(C'_j), R(C'_j))\}$, where $L(C_i) = \mid p_i \wedge \neg q_i \mid$, $R(C_i) = \mid \neg p_i \wedge q_i \mid$, $L(C'_j) = \mid p'_j \wedge \neg q'_j \mid$ and $R(C'_j) = \mid \neg p'_j \wedge q'_j \mid$.

Example 7 (Continued). Let $\mathcal{P}_{mM} = \{\neg j \wedge \neg w \;{}^{m}\!\!>^{M} \neg j \wedge w\}$ and $\mathcal{P}_{mm} = \{\neg j \wedge w \wedge r \;{}^{m}\!\!>^{m} \neg j \wedge w \wedge \neg r\}$.

Following Algorithm 1.2, we have $\succeq_{mM,mm} = (\{\omega_0, \omega_4\}, \{\omega_5\}, \{\omega_1, \omega_2, \omega_3, \omega_6, \omega_7\})$.

Theorem 3. *Let $\mathcal{P} = \mathcal{P}_{mM} \cup \mathcal{P}_{mm}$. Then Algorithm 1.2 computes the most specific model of $\mathcal{P}$ which is unique.*

Proof (sketch). Follows the same lines as the proof of Theorem 2. It can also be derived from Theorem 2 using symmetry of the two algorithms.

Algorithm 1.2: Handling mixed preferences ${}^{m}\!\!>^{M}$ and ${}^{m}\!\!>^{m}$.

begin
$l \leftarrow 0$;
while *($W \neq \emptyset$)* **do**
$l \leftarrow l + 1$;
$E_l = \{\omega : \forall (L(C_i), R(C_i)), \forall (L(C'_j), R(C'_j)) \in \mathcal{C}, \omega \notin L(C_i), \omega \notin L(C'_j)\}$;
if $E_l = \emptyset$ **then**
Stop (inconsistent constraints)
- Remove from W elements of E_l;
- Remove from $\mathcal{C}$ constraints s.t. $R(C_i) \cap E_l \neq \emptyset$;
- Replace each $(L(C'_j), R(C'_j))$ in $\mathcal{C}$ by $(L(C'_j), R(C'_j) - E_l)$;
- Remove from $\mathcal{C}$ constraints with empty $R(C'_j)$
return $(E'_1, \cdots, E'_l)$ *s.t.* $\forall 1 \leq j \leq l, E'_j = E_{l-j+1}$
end

5.3 $\mathcal{P}_{MM}$, $\mathcal{P}_{mm}$ and $\mathcal{P}_{mM}$

To find a distinguished model of three kinds of preferences, we want to combine the two algorithms. It has been argued in [2, 8] that, in the context of preference modeling, the minimal specificity principle models constraints which should not be violated while the maximal specificity principle models what is really desired by the agent. In our setting, this combination of the least specific and the most specific models leads to a refinement of the first one by the latter.

Definition 11. *Let $\succeq''$ be the result of combining $\succeq$ and $\succeq'$ corresponding to the least specific and the most specific models respectively. Then,*

- *if $\omega \succ \omega$ then $\omega \succ'' \omega'$,*
- *if $\omega \simeq \omega'$ then ($\omega \succeq'' \omega'$ iff $\omega \succeq' \omega'$).*

Example 8 (Continued from Examples 6 and 7). We have a unique least specific pre-order $\succeq_{MM,mM} = (\{\omega_0, \omega_3\}, \{\omega_2, \omega_4\}, \{\omega_1, \omega_5, \omega_6\omega_7\})$, and a unique most specific pre-order $\succeq_{mM,mm} = (\{\omega_0, \omega_4\}, \{\omega_5\}, \{\omega_1, \omega_2, \omega_3, \omega_6, \omega_7\})$. Following the combination method of Definition 11, we get the following unique distinguished model:
$(\{\omega_0\}, \{\omega_3\}, \{\omega_4\}, \{\omega_2\}, \{\omega_5\}, \{\omega_1, \omega_6, \omega_7\})$.

6 Summary

In this paper we introduce and study a logic of preferences, which we understand as a logic that formalizes reasoning about various kinds of preferences.

To define mixed logics of preference, we use total orders on worlds called the preference order. We define four kinds of strict preferences of p over q as ”the best/worst p is preferred over the best/worst q”.

To choose among types of preferences, we introduce an agent interpretation of preferences. We interpret a preference of p over q as a game between an agent arguing for p and an agent arguing for q. For an ordered set S an optimistic agent selects the maximal element of S, and a pessimistic agent selects the minimal element of S. For a preference of p over q, a locally optimistic agent is optimistic about p and optimistic about q, a locally pessimistic agent is pessimistic about p and pessimistic about q, an opportunistic agent is optimistic about p and pessimistic about q, and a careful agent is pessimistic about p and optimistic about q.

To calculate a preference order from preferences, we start from a generalization of System Z, which is usually characterized as gravitating towards the ideal. $\max$ is gravitating towards the ideal or minimal specificity, $\min$ is gravitating towards the worst or maximal specific for ${}^{M}{>}^{M}$ and ${}^{m}{>}^{M}$, and most specific for ${}^{m}{>}^{m}$ and ${}^{m}{>}^{M}$. We show that also for ${}^{M}{>}^{M}$ and ${}^{m}{>}^{M}$ preferences together the least specific model is unique, and we show that for ${}^{m}{>}^{m}$ and ${}^{m}{>}^{M}$ preferences together the most specific preference order is unique. For these cases, we have provided algorithms to compute the unique models. We also propose a way to compute a distinguished model of ${}^{M}{>}^{M}$, ${}^{m}{>}^{M}$ and ${}^{m}{>}^{m}$ preferences toegther, combining the developed algorithms.

The results in this paper can be generalized to ceteris paribus preferences using frames [7] or Hansson functions [10]. This is subject of future research. We will also consider consequences of our framework for the discussion on bipolarity [2, 8], distinguishing between bipolarity in logic (left hand side and right hand side of constraint) and in nonmonotonic reasoning (least or most specific).

References

1. Special issue on preferences of computational intelligence. *Computational intelligence*, 20(2), 2004.
2. S. Benferhat, D. Dubois, S. Kaci, and H. Prade. Bipolar representation and fusion of preferences in the possibilistic logic framework. In *8th International Confenrence on Principle of Knowledge Representation and Reasoning (KR'02)*, pages 421–432, 2002.
3. S. Benferhat, D. Dubois, and H. Prade. Representing default rules in possibilistic logic. In *Proceedings of 3rd International Conference of Principles of Knowledge Representation and Reasoning (KR'92)*, pages 673–684, 1992.
4. S. Benferhat, D. Dubois, and H. Prade. Possibilistic and standard probabilistic semantics of conditional knowledge bases. *Logic and Computation*, 9:6:873–895, 1999.
5. S. Benferhat and S. Kaci. A possibilistic logic handling of strong preferences. In *International Fuzzy Systems Association (IFSA'01)*, pages 962–967, 2001.
6. C. Boutilier. Toward a logic for qualitative decision theory. In *Proceedings of the 4th International Conference on Principles of Knowledge Representation, (KR'94)*, pages 75–86, 1994.
7. J. Doyle and M. P. Wellman. Preferential semantics for goals. In *National Conference on Artif. Intellig. AAA'91*, pages 698–703, 1991.
8. D. Dubois, S. Kaci, and H. Prade. Bipolarity in reasoning and decision – an introduction. the case of the possibility theory framework. In *Proceedings of Information Processing and Management of Uncertainty in Knowledge-Based Systems Conference, IPMU'04*, pages 959–966, 2004.
9. D. Dubois, S. Kaci, and H. Prade. Ordinal and absolute representations of positive information in possibilistic logic. In *Proceedings of the International Workshop on Nonmonotonic Reasoning (NMR' 2004), Whistler, June*, pages 140–146, 2004.
10. S.O. Hansson. What is ceteris paribus preference? *Journal of Philosophical Logic*, 25:307–332, 1996.

11. J. Lang, L. Van Der Torre, and E. Weydert. Utilitarian desires. *Autonomous Agents and Multi-Agent Systems*, 5:329–363, 2002.
12. J. Pearl. System z: A natural ordering of defaults with tractable applications to default reasoning. In R. Parikh. Eds, editor, *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge (TARK'90)*, pages 121–135. Morgan Kaufmann, 1990.
13. Y. Shoham. Nonmonotonic logics: Meaning and utility. In *Procs of IJCAI 1987*, pages 388–393, 1987.
14. S. Tan and J. Pearl. Qualitative decision theory. In *Proceedings of the National Conference on Artificial Intelligence (AAAI'94)*, pages 928–933, 1994.
15. L. van der Torre and E. Weydert. Parameters for utilitarian desires in a qualitative decision theory. *Applied Intelligence*, 14:285–301, 2001.
16. G. H. von Wright. *The Logic of Preference*. University of Edinburgh Press, 1963.

Appendix

Proposition 1
Let $\succeq$ and $\succeq'$ be two elements of $\mathcal{M}(\mathcal{P})$. Then,

1. $\mathcal{MAX}(\succeq, \succeq') \in \mathcal{M}(\mathcal{P})$.

Proof
Let $\mathcal{P} = \mathcal{P}_{MM} \cup \mathcal{P}_{mM}$. Let $\succeq$ and $\succeq'$ be two elements of $\mathcal{M}(\mathcal{P})$.
Suppose that $\succeq$ and $\succeq'$ are represented by $(E_1, \cdots, E_n)$ and $(E'_1, \cdots, E'_h)$ respectively. Let $\succeq'' = \mathcal{MAX}(\succeq, \succeq')$. To show that $\succeq'' \in \mathcal{M}(\mathcal{P})$, we show that $\succeq''$ satisfies all constraints $p \; {}^{M}\!\!>^{M} q$ and $p' \; {}^{m}\!\!>^{M} q'$ in $\mathcal{P}$.
Let $(E''_1, \cdots, E''_{min(n,m)})$ be the well ordered partition associated to $\succeq''$. Recall that the best models of $p \wedge q$ w.r.t. $\succeq$ are defined by $max(p \wedge q, \succeq) = \{\omega : \omega \models p \wedge q$ s.t. $\nexists \omega', \omega' \models p \wedge q$ with $\omega \in E_i, \omega' \in E_j$ and $j < i\}$.
Similarly the worst models of $p \wedge q$ w.r.t. $\succeq$ are defined by $min(p \wedge q, \succeq) = \{\omega : \omega \models p \wedge q$ s.t. $\nexists \omega', \omega' \models p \wedge q$ with $\omega \in E_i, \omega' \in E_j$ and $j > i\}$.
Let $p \; {}^{M}\!\!>^{M} q$ be a constraint in $\mathcal{P}$.
Following Definition 7, $\succeq$ belongs to $\mathcal{M}(\mathcal{P})$ means that $max(p \wedge \neg q, \succeq) \subseteq E_i$ and $max(\neg p \wedge q, \succeq) \subseteq E_j$ with $i < j$. Also $\succeq'$ belongs to $\mathcal{M}(\mathcal{P})$ means that $max(p \wedge \neg q, \succeq') \subseteq E'_k$ and $max(\neg p \wedge q, \succeq') \subseteq E'_m$ with $k < m$.
Following Definition 10, $max(p \wedge \neg q, \succeq'') \subseteq E''_{min(i,k)}$ and $max(\neg p \wedge q, \succeq'') \subseteq E''_{min(j,m)}$. Now since $i < j$ and $k < m$, we have $min(i,k) < min(j,m)$. Hence $\succeq''$ satisfies $p \; {}^{M}\!\!>^{M} q$.
Similarly we show that each constraint $p' \; {}^{m}\!\!>^{M} q'$ in $\mathcal{P}$ is satisfied by $\succeq''$.
$\succeq$ (resp. $\succeq'$) satisfies $p' \; {}^{m}\!\!>^{M} q'$ means that $min(p' \wedge \neg q', \succeq) \subseteq E_i$ (resp. $min(p' \wedge \neg q', \succeq') \subseteq E'_k$) and $max(\neg p' \wedge q', \succeq) \subseteq E_j$ (resp. $max(\neg p' \wedge q', \succeq') \subseteq E'_m$) s.t. $i < j$ (resp. $k < m$). Following Definition 10, $min(p' \wedge \neg q', \succeq'') \subseteq E''_{min(i,k)}$ and $max(\neg p' \wedge q', \succeq'') \subseteq E''_{min(j,m)}$. Again since $i < j$ and $k < m$ then $min(i,k) < min(j,m)$. Hence $\succeq''$ satisfies $p' \; {}^{m}\!\!>^{M} q'$.