A Non-monotonic Logic for Specifying and Querying Preferences

Guido Boella
Dipartimento di Informatica
Università di Torino
Italy
guido@di.unito.it

Leendert van der Torre
CWI Amsterdam
and Delft University of Technology
The Netherlands
torre@cwi.nl

1 Introduction

Preferences are becoming of greater interest in many areas of artificial intelligence, such as knowledge representation, multiagent systems, constraint satisfaction, decision making, and decision-theoretic planning. In the logic of preference there is a debate when a set of preferences should be consistent. For example, Bacchus and Grove [1996] criticize ceteris paribus preferences, because $\{p > \neg p, \neg p \wedge q > p \wedge q\}$ should be consistent, and they criticize most existing logics of preference, because $\{p > \neg p, q > \neg q, \neg(p \wedge q) > p \wedge q\}$ should be consistent. In order not to restrict the use of the logic of preference, we propose a minimal logic of preference in which *any* set of specified preferences is consistent. To make it useful for practical applications, we extend this logic to specify preferences with a logic to query preferences, and with a non-monotonic reasoning mechanism.

2 New semantics for the logic of preference

We introduce a logic of preference that distinguishes between specifying and querying, inspired by preference-based deontic logic [van der Torre and Tan, 2000]. We define two types of preference statements, where $p >_! q$ can be read as "it is specified that p is preferred to q" and $p >_? q$ can be read as "the result of the query states that p is preferred to q."

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.*

$L_0 \ni p, q$: $a_i \mid \neg p \mid (p \wedge q)$

$L \ni \phi, \psi$: $(p >_! q) \mid (p >_? q) \mid \neg\phi \mid (\phi \wedge \psi)$

Moreover, disjunction $\vee$, material implication $\rightarrow$ and equivalence $\leftrightarrow$ are defined as usual. We abbreviate formulas of the preference logic using the following order on logical connectives: $\neg \mid \vee, \wedge \mid >_!, >_? \mid \rightarrow, \leftrightarrow$. For example, $\neg p >_! q \wedge r$ is short for $(\neg p >_! (q \wedge r))$.

Typically many conflicts arise during the specification of preferences. Our logic of preference represents *any* set of preferences in a consistent way, which is achieved in two steps. First, we do not use a total pre-order on worlds in the semantics, but a partial pre-order on worlds, i.e., a reflexive and transitive relation. Incomparable worlds indicate some kind of conflict among these worlds. Secondly, and more originally, we formalize a preference $p >_! q$ as the absence of a $\neg p \wedge q$ world that is preferred over a $p \wedge \neg q$ world.

Definition 2 (Specifying) *Let W be the set of propositional interpretations of L_0, and $\succeq$ a partial pre-order on W. We write $w \succ w'$ for $w \succeq w'$ without $w' \succeq w$, and $|\alpha|$ for $\{w \in W \mid w \models \alpha\}$.*

- *$\succeq \models p >_! q$ iff $\forall w \in |p \wedge \neg q|$ and $\forall w' \in |\neg p \wedge q|$ we do not have $w' \succ w$.*

Semantic entailment for more complex formulas, logical entailment and other logical notions are defined as usual. For example, $\phi_1, \ldots, \phi_n \models \phi$ if and only if for all $\succeq$ such that $\succeq \models \phi_i$ for $i = 1 \ldots n$, we have $\succeq \models \phi$.

We define a logic for querying preferences by optimization in a specified preference relation. We use the same semantic structures, but to simplify the definitions we assume that there are no infinitely ascending chains (see, e.g., [Boutilier, 1994] for a discussion). Moreover, p is preferred to q if the optimal $p \vee q$ worlds are p worlds.

Definition 3 (Querying) *We assume $\succeq$ does not contain any infinite ascending chains. We write $\max(\succeq, p)$ for $\{w \in W \mid w \models p$ and $\forall w' \in W : w' \succ w \Rightarrow w' \not\models p\}$.*

- *$\succeq \models p >_? q$ iff $\max(\succeq, p \vee q) \subseteq \max(\succeq, p)$.*

We distinguish models called most connected models.

Definition 4 (Distinguished models) *A model $\succeq_1$ is at least as connected as another model $\succeq_2$, written as $\succeq_1 \sqsubseteq \succeq_2$, if $\succeq_2 \subseteq \succeq_1$, that is, if $\forall w_1, w_2 \in W : w_1 \succeq_2 w_2 \Rightarrow w_1 \succeq_1 w_2$. A model $\succeq_1$ is most connected if there is not another model $\succeq_2$ such that $\succeq_2 \sqsubset \succeq_1$, that is, such that $\succeq_2 \sqsubseteq \succeq_1$ without $\succeq_1 \sqsubseteq \succeq_2$.*

Non-monotonic reasoning is based on distinguished models in the usual way. These distinguished models have sometimes been called preferred models, and non-monotonic entailment has been called preferential entailment. However, to avoid confusion with the preferences formalized in our logical language, we do not use this terminology in this paper.

Definition 5 (Non-monotonic entailment) *A preference specification is a set of specification preferences $PS = \{p_1 >_! q_1, \ldots, p_n >_! q_n\}$. A distinguished model of a preference specification PS is a most connected model of PS. A preference specification non-monotonically entails ϕ, written as $PS \models_{\sqsubseteq} \phi$, if for all distinguished models $\succeq$ of PS we have $\succeq \models \phi$.*

3 Logical properties

Specifying and querying preferences are complementary in the sense that the logic of the former satisfies left and right strengthening, and the logic of the latter satisfies transitivity.

3.1 Ordering

We first show that our desideratum holds. A preference specification which only contains preferences which are individually consistent, is itself consistent.

Proposition 1 *If there is a model for each preference in PS, then there is a model of PS.*
Proof. The identity relation $\{w \succeq w \mid w \in W\}$ is such a model of PS.

The logical relations among preference specifications are characterized by left and right additivity.

LA $\models (p \vee q >_! r) \leftrightarrow (p >_! r) \wedge (q >_! r)$

RA $\models (p >_! q \vee r) \leftrightarrow (p >_! q) \wedge (p >_! r)$

Consequently, reading the additivity formulas from left to right, specification preferences are strong preferences, in the sense that they satisfy the following properties of left and right strengthening.

LS $\models (p >_! q) \rightarrow (p \wedge r >_! q)$

RS $\models (p >_! q) \rightarrow (p >_! q \wedge r)$.

Moreover, reading the additivity formulas from right to left, specification preferences satisfy left and right disjunction.

LOR $\models (p >_! r) \wedge (q >_! r) \rightarrow (p \vee q >_! r)$

ROR $\models (p >_! q) \wedge (p >_! r) \rightarrow (p >_! q \vee r)$

In case of a finite set of atoms, the combination of LS, RS, LOR and ROR give a simple way to derive all implied preferences from a preference specification. Call a formula a complete sentence when it implies p or $\neg p$ for each proposition $p \in A$. Given a preference specification PS, first use LS and RS to derive preferences among complete sentences. Secondly, use LOR and ROR to derive all preferences among formulas.

Finally we consider some borderline cases. As usual, we cannot have that a proposition is preferred to itself. This is due to the fact that $\succeq$ is reflexive.

Id $\models \neg(p >_! p)$

Moreover, an important distinction among preference logics is whether they satisfy transitivity. Our logic to specify preferences does not, because we may have r worlds preferred to p worlds, and q worlds incomparable to both p and r worlds.

T $\not\models (p >_! q) \wedge (q >_! r) \rightarrow (p >_! r)$

3.2 Optimizing

Optimizing preferences have been studied in the context of total pre-orders. For example, Weydert [1991] defines elementary qualitative magnitude logic as follows, in a propositional preference logic extended with actual worlds:

QM1 $(p > q) \wedge (q > r) \rightarrow (p > r)$ (transitivity)

QM2 $(p > q) \wedge (p > r) \leftrightarrow (p > q \vee r)$ (additivity)

QM3 $\neg(q > p) \wedge (q > r) \rightarrow (p > r)$ (maximality)

QM4 $\bot > \bot \wedge ((\bot > p) \leftrightarrow (p > p))$ (quasi-reflexivity)

QM5 $\bot > p \rightarrow \neg p$ (correctness)

QM6 $(p > q) \leftrightarrow (r > s)$ if $\vdash (p \leftrightarrow r) \wedge (q \leftrightarrow s)$ (extensionality)

Of these formulas QM5 cannot be expressed in our logic, and the relevant properties to check are QM1, QM2 and QM3. It can be verified that QM1 and QM2 also hold for our logic based on partial pre-orders, but QM3 does not.

In some cases, it may be useful to define a stronger notion of specification preferences, as a combination of $>_!$ and $>_?$.

Definition 6 $p >_0 q = p >_! q \wedge p >_? q$

Consequently, if we have $\succeq \models p >_0 q$ then we also have $\succeq \models p >_? q$, a property that does not hold for $>_!$. That is, if we have $\succeq \models p >_! q$, then we do not necessarily have $\succeq \models p >_? q$.

3.3 Modal characterization

It is well known that the combination of two logics in a single logical system may be problematic, and that the axiomatization of such a combined logic is a non-trivial task. To obtain an axiomatization of our logic, we follow [Boutilier, 1994] by providing a modal characterization of the logic. We define (or *simulate*) the preference logic in a normal bimodal system, where $\Box$ is a universal modal operator and $\Box'$ is a normal modal operator. See, e.g., [Boutilier, 1994] for an axiomatization of the bimodal logic.

Definition 7 *Let L' be a modal logic with two modal operators $\Box$ and $\Box'$. Moreover, let a model of L' be a tuple $M = \langle W, \succeq, V \rangle$, with W a set of worlds, $\succeq$ a partial preorder on W, and V a valuation function of propositions at the worlds. The satisfiability relation for modal formulas is defined as follows.*

- *$M, w \models \Box p$ when $\forall w' \in W$ we have $M, w' \models p$.*
- *$M, w \models \Box' p$ when $\forall w' \in W : w' \succeq w \Rightarrow M, w' \models p$.*

Moreover, we add the following definitions to L':

- $\Diamond p = \neg \Box \neg p$, $\Diamond' p = \neg \Box' \neg p$
- $p >_! q = \Box(p \rightarrow \Box' \neg q)$
- $p >_? q = \Box(q \rightarrow \Diamond'(p \wedge \neg q \wedge \Box'(p \vee q \rightarrow p \wedge \neg q)))$

References

[Bacchus and Grove, 1996] F. Bacchus and A. Grove. Utility independence in a qualitative decision theory. In *Proceedings of the KR-96*, pages 542–552, 1996.

[Boutilier, 1994] C. Boutilier. Conditional logics of normality: a modal approach. *Artificial Intelligence*, 68:87–154, 1994.

[van der Torre and Tan, 2000] L. van der Torre and Y. Tan. Two-phase deontic logic. *Logique et Analyse*, 171-172:411–456, 2000.

[Weydert, 1991] E. Weydert. Qualitative magnitude reasoning. In *Nonmonotonic and Inductive Logic*, volume 543 of *Lecture Notes in Computer Science*, pages 138–160. Springer, 1991.