

Expressive power and succinctness of propositional languages for preference representation

Sylvie Coste-Marquis

CRIL - Univ. d'Artois
Lens, France

coste@cril.univ-artois.fr

Jérôme Lang

IRIT - Univ. Paul Sabatier
Toulouse, France

lang@irit.fr

Paolo Liberatore

DIS - Univ. La Sapienza
Roma, Italy

paolo@liberatore.org

Pierre Marquis

CRIL - Univ. d'Artois
Lens, France

marquis@cril.univ-artois.fr

Abstract

Several logical languages have been considered in AI for encoding compactly preference relations over a set of alternatives. In this paper, we analyze both the expressiveness and the spatial efficiency (succinctness) of such preference representation languages. The first issue is concerned with the nature of the preorders that can be encoded (for instance, all preorders, all complete preorders). The second issue is about how succinctly a preference relation can be expressed in those languages. We give polynomial-size translations in some cases, and prove the impossibility of such translations in other cases.

Introduction

As soon as decision making is concerned, one has to represent the *preferences* of the agent over the set of feasible alternatives. This applies to a large variety of domains, such as decision-theoretic planning, combinatorial auctions, game playing, group decision making (including resource sharing and vote), and agent coordination (including cooperation, negotiation, and communication).

In the following, we assume that the set \mathcal{A} of *feasible alternatives* is fixed (and finite). There are several possible ways for modeling preference over \mathcal{A} ; however, most models fall in one of these two classes:

- *utilitarian preferences* consist of a utility function u from \mathcal{A} to a numerical scale (generally \mathbb{R});
- *ordinal preferences* consist of a preorder \succeq on \mathcal{A} .

In many real-world domains, the set of alternatives is the set of assignments of a value to each of a given set of variables. In such cases, the alternatives are exponentially many. It is not reasonable to ask agents to report their preference in an explicit way when the set of alternatives is exponentially large, as this amounts to listing the exponentially many alternatives together with their utility assessment or their ranking.

For this reason, AI researchers have been more and more concerned with *languages for preference representation* aiming at enabling a succinct representation of the description of the problem, without having to enumerate a

prohibitive number of alternatives. Such preference representation languages are often built up on propositional logic, and allow for a much more concise representation of the preference structure than an explicit enumeration, while preserving a good readability and hence a similarity with the way agents express their preferences in natural language. The latter point is of the utmost value in the perspective of preference elicitation, which is typically a difficult problem when the set of alternatives has exponentially many elements, due to its combinatorial nature. In this paper we choose to consider only languages based on full propositional logic, therefore we leave graphical languages for utility representation (Bacchus & Grove 1995; Boutilier, Bacchus, & Brafman 2001; Gonzales & Perny 2004), CP-nets¹ (Boutilier *et al.* 1999; Domshlak & Brafman 2002), and valued constraint satisfaction (Bistarelli *et al.* 1999) to a further study.

A fundamental issue is how to choose among the many preference representation languages. Many parameters are relevant to such a choice. Some of them are domain-dependent (e.g., preference may be easier to express in one language than in another one for a specific application). Among the domain-independent ones are expressiveness and succinctness. On the one hand, assumptions are usually made about the nature of preferences (e.g., preferences are often considered completely ordered) so it is essential to choose a representation language suited to such structural assumptions, i.e., allowing for the representation of the expected kind of preorder. On the other hand, the spatial efficiency of the languages, which is their ability to represent information in little space (Cadoli *et al.* 1996), must also be considered. Indeed, a language in which a given ordering can be expressed using only exponentially long expressions is not only problematic from a computational point of view (a large amount of memory is required to represent orderings) but also indicates that the language is intuitively not suited for representing such an ordering. On the contrary, the motivation of using preference languages for representing orderings \succeq is to avoid explicitly storing the fact that $M \succeq M'$ for each ordered pair of alternatives

¹CP-nets are nevertheless close to the language of *ceteris paribus* statements considered in this paper. We will return on it later.

(M, M') in \mathcal{A} for which the relation holds. If representing the ordering requires exponential space in a language, however, the gain is lost and the explicit representation may be even better.

In this paper, we investigate both the expressiveness and the relative space efficiency of preference languages. While it is known that some orderings require exponential space in the explicit form but are more compact in preference representation languages, no result is known to the authors about comparing the ability of the languages to represent orderings compactly. The results of this paper are of two kinds: the first results show that all orderings expressed in a language L_1 can be translated into language L_2 with a polynomially large increase of size; the second ones prove that such a translation is not always possible. In the first case, the language L_2 is at least as good as language L_1 , as it can polynomially express the same orderings L_1 can polynomially express. In the second case, this is not true; as a result, L_2 is not as good as L_1 , as some orderings are polynomial in L_1 but exponential in L_2 . Some pairs of preference languages L_1 and L_2 can also be incomparable, in the sense that some orderings are exponential in the first one but polynomial in the second one, while the inverse holds for other preference relations. Given a set of computational tasks (e.g., determining whether a world is undominated) which can be solved with the same computation costs when preferences are encoded in two different languages, it is natural for efficiency reasons to prefer the language enabling the more compact encoding.

Encoding preference relations (for various uses) must not be confused with taking advantage of preferential information for defining inference (which is another issue, even if it is related to preference representation). If some problems in the two fields are similar (e.g., determining whether a decision is undominated is just the model checking problem for nonmonotonic KR formalisms), this is not the case for all of them. For instance, it can make sense to determine the diameter of a preference ordering, or the number of its connected components (since this reflects in some sense how refined the preferences are) while there is no significant counterpart for it when inference is concerned.

The difference between the two objectives is very salient in light of the translatability functions used for comparing the spatial efficiency of the logic-based languages we consider. When inference is concerned, it is sufficient to preserve the set of preferred models (or the set of consequences of preferred models over the original language); this is reflected in the few notions of polyspace translatability which can be found in the literature (Cadoli *et al.* 1996; Gogic *et al.* 1995; Gottlob 1995; Janhunen 1998). Relaxing the polyspace requirement, every inference relation from a propositional base which has a preferred model semantics can be turned into an equivalent base classically interpreted (any base whose classical models are the preferred models of the original base does the job). When the purpose is preference representation, such a translation is impossible in the general case (it is only possible when the set of alternatives quotiented by the equivalence relation induced by the preorder has two elements). Indeed, we need a much stronger

notion of translatability when preference representation is concerned (the whole preorder must be preserved, not only the set of undominated elements). As a corollary, existing translatability results for inference cannot be used as such for our purpose.

The aforementioned difference between encoding a preference relation over a set of possible alternatives or possible worlds, and using preferential information for defining or controlling inference, is the reason why we do not consider in this paper “preference-based” approaches to logic programming or default logic (see for instance (Brewka & Eiter 1999; Delgrande, Schaub, & Tompits 2003)). In these approaches, preference bears on syntactical items (rules) and is used for discriminating among answer sets (or extensions) so as to produce “preferred” ones, but not to produce a full ordering on the set of all possible interpretations² (note, by the way, that the word “preference” in these approaches does not have exactly the same meaning as the decision-theoretic one).

The paper is organized as follows. In the next section, we state the notations used in this paper. The following section contains a panorama of the propositional preference representation languages so far studied in the literature. The results of this paper are summarized in another section; due to lack of space, full proofs are omitted and we only describe succinctly the proof methodology and give a few proof sketches. In the last section of the paper, we point out some further issues.

Preliminaries

In this paper, \mathcal{L} is a propositional language built upon a finite set of propositional variables $VAR = \{x_1, \dots, x_n\}$, the usual connectives, and the symbols \top (tautology) and \perp (contradiction). A literal is a propositional variable or its negation. $W = 2^{VAR}$ is the set of all interpretations (worlds) over VAR . Interpretations are denoted by M, M' etc. They are represented as tuples of literals over VAR ; for instance, if $VAR = \{a, b, c, d\}$, then $M = [a, \neg b, \neg c, d]$ is the interpretation mapping a and d to true, and b and c to false. If a world M satisfies a formula $G \in \mathcal{L}$, we write $M \models G$. $Mod(G)$ and $Var(G)$ denote the set of models of G and the set of variables mentioned in G , respectively. If $X \subseteq W$ then $Form(X)$ denotes the formula – unique up to logical equivalence – such that $Mod(Form(X)) = X$.

Propositional Languages for Preference Representation

The relative preference of alternatives can be expressed in several ways. We will not discuss here the pros and cons of each – this important question has been discussed for long by researchers in decision theory and cognitive psychology. Rather, we investigate how good they are in expressing a given preference ordering in little space. Most preference relation can be formalized either as an utilitarian preference

²This does not truly apply to the recent work on logic programming with ordered disjunction (Brewka 2002); more on ordered disjunction will be given later in the paper.

(a function giving the goodness of the alternatives) or as an ordinal preference (a binary relation over the pairs of alternatives).

In this paper, we mainly focus on this second kind of preferences. A *preference relation* \succeq is a preorder, i.e., a reflexive and transitive binary relation on \mathcal{A} . $M \succeq M'$ means that alternative M is at least as good (to the agent) as alternative M' . Such a relation \succeq is not necessarily complete, that is, it may be that neither $M \succeq M'$ nor $M' \succeq M$ holds for a pair of alternatives M and M' in \mathcal{A} . We note $M \succ M'$ for $M \succeq M'$ and not $(M' \succeq M)$ (strict preference of M over M'), and $M \sim M'$ for $M \succeq M'$ and $M' \succeq M$ (indifference). It is important to note that $M \sim M'$ means that the agent takes M and M' to be equally preferred, while the incomparability between M and M' ($M \not\succeq M'$ and $M' \not\succeq M$) simply means that no preference between them is expressed.

These definitions are about preferences over an arbitrary set of alternatives \mathcal{A} . In this paper, we consider propositional languages expressing preferences: such languages express preferences over the set of possible interpretations W over a given alphabet VAR . A refinement of this definition is that of assuming that the set of possible alternatives excludes some interpretations of W . In this case, we assume that a formula K is given: this formula represents “integrity constraints” on the set of *feasible* alternatives, i.e., the only interpretations we accept as possible alternatives are those of $Mod(K)$, i.e., $\mathcal{A} = Mod(K)$. For instance, in a decision making problem consisting of recruiting at least one and at most two of three candidates a, b and c , the feasible alternatives are the models of $K = (a \vee b \vee c) \wedge (\neg a \vee \neg b \vee \neg c)$.

Even if we focus on ordinal preference in this paper, we sometimes nevertheless make use of utility functions u , but only because utility functions are a way among others of encoding complete preference relations: namely, each u induces \succeq_u defined by $M \succeq_u M'$ iff $u(M) \geq u(M')$.

We now briefly recall the propositional languages for preference representation we study. In the following, the formulas G_i are propositional formulas representing elementary *goals*. The input of a logically-represented preference relation is a pair $\Delta = \langle K, GB \rangle$ where K is the propositional formula restricting the possible alternatives (the integrity constraints) and GB (the *goal base*) is a set of elementary goals, generally associated with extra data such as weights, priorities, contexts or distances. $\succeq_{K,GB}$ (or simply \succeq_{GB} when there is no risk of ambiguity) denotes the preference relation induced by GB over $Mod(K)$.

Penalties ($R_{penalties}$)

In this natural and frequently used preference representation language, the agent expresses her preferences in terms of a set of propositional formulas that she wants to be satisfied. In order to compare alternatives (models), formulas are associated with weights (usually, numbers), which tell how important the satisfaction of the formula is considered.

Formally, the preferences of an agent are expressed as a finite set of goals, where each goal is a propositional formula with an associated weight. The complete preference is given by a set of these goals: $GB = \{\langle \alpha_1, G_1 \rangle, \dots, \langle \alpha_n, G_n \rangle\}$,

where each α_i is an integer and each G_i is a propositional formula.

The degree of preference of a model is measured as follows: for any $M \in Mod(K)$, we define $p_{GB}(M) = \sum \{\alpha_i | M \not\models G_i\}$ to be the penalty of M . The preference relation \succeq_{GB}^{pen} is defined by $M \succeq_{GB}^{pen} M'$ if and only if $p_{GB}(M) \leq p_{GB}(M')$ (with the convention $\sum(\emptyset) = 0$).

Example 1 Let

- $K = \top$
- $GB = \{\langle 2, a \wedge b \rangle, \langle 2, \neg a \wedge c \rangle, \langle 3, c \rightarrow d \rangle\}$.

We have

$$\begin{aligned} p_{GB}([a, b, c, d]) &= 0 + 2 + 0 = 2; \\ p_{GB}([\neg a, \neg b, c, d]) &= 2 + 0 + 0 = 2; \\ p_{GB}([\neg a, \neg b, \neg c, \neg d]) &= 2 + 2 + 0 = 4 \\ \text{and } p_{GB}([\neg a, b, c, \neg d]) &= 2 + 0 + 3 = 5, \end{aligned}$$

therefore

$$\begin{aligned} & [a, b, c, d] \sim_{GB}^{pen} [\neg a, \neg b, c, d] \\ \succ_{GB}^{pen} & [\neg a, \neg b, \neg c, \neg d] \\ \succ_{GB}^{pen} & [\neg a, b, c, \neg d] \end{aligned}$$

Many references can be given regarding the use of penalties in a logical framework. Here are a few of them: (Haddawy & Hanks 1992) for a utility representation language with penalties and rewards, (Pinkas 1995; Dupin de Saint-Cyr, Lang, & Schiex 1994) for a more systematic study.

Many other operators can be used, in place of the sum, for aggregating weights of violated (or symmetrically, satisfied) formulas (see (Lafage & Lang 2000) for a general discussion). The other usual choice, namely the maximum – leading to *possibilistic logic* – will be covered by the best-out preference relation in the next paragraph. This principle can be extended so as to introduce polarities between *goals* (inducing positive rewards when satisfied) and *constraints* (inducing negative rewards when violated); see for instance (van der Torre & Weydert 2001; Benferhat *et al.* 2002; Kaci 2002).

Hamming Distance to Goals (R_H)

The preference relation based on penalties only makes a distinction between models satisfying a formula and models violating it. On the other hand, if an agent prefers a formula G_i to be satisfied, we could infer that she also prefers models “close” to this formula than models “far”. The measure of closeness between models mostly used is the Hamming distance, so that the distance from a model to a goal is the number of variables that must be flipped in the model in order to make it satisfy the formula.

Let d be the Hamming distance between models (i.e., $d(M, M')$ is the number of variables that are assigned different values in M and M' .) The “distance” between a model M and a formula G is defined as follows:

$$d(M, G) = \min_{M' \models G} d(M, M').$$

A goal base is a finite set of pairs $\langle \alpha_i, G_i \rangle$, as for penalties languages. In this case, however, we make each formula

contribute to the total weight in a measure that is proportional to the distance between the formula and the model:

$$d(M, GB) = \sum_i \{\alpha_i \cdot d(M, G_i)\}.$$

\succeq_{GB}^H is defined (see for instance (Lafage & Lang 2001)) by

$$M \succeq_{GB}^H M' \text{ if and only if } d(M, GB) \leq d(M', GB)$$

Example 2 *Let (again)*

- $K = \top$
- $GB = \{\langle 2, a \wedge b \rangle, \langle 2, \neg a \wedge c \rangle, \langle 3, c \rightarrow d \rangle\}$.

Let $M = [\neg a, \neg b, c, \neg d]$. We have

$$d(M, a \wedge b) = 2, d(M, \neg a \wedge c) = 0 \text{ and } d(M, c \rightarrow d) = 1$$

hence $d(M, GB) = 2 \times 2 + 2 \times 0 + 3 \times 1 = 7$.

Let $M' = [\neg a, \neg b, \neg c, \neg d]$. We have

$$d(M', a \wedge b) = 2, d(M', \neg a \wedge c) = 1 \text{ and } d(M', c \rightarrow d) = 0$$

hence $d(M', GB) = 2 \times 2 + 2 \times 1 + 3 \times 0 = 6$ and therefore $M' \succ_{GB}^H M$.

Distances between logical worlds have been used many times in knowledge representation, especially in belief revision (Dalal 1988; Magidor & Schlechta 2001), belief update (Katsuno & Mendelzon 1991) and belief merging (Liberatore & Schaerf 1998; Lin & Mendelzon 1999; Revesz 1997; Konieczny & Pino-Pérez 1998; Konieczny, Lang, & Marquis 2002). Their use in logical preference representation is more recent (Lafage & Lang 2000; 2001). While we use here the Hamming distance d and the weighted sum, an arbitrary distance function and an arbitrary aggregation function can be used instead – note however that succinctness results depend on the chosen distance.

Prioritized Goals

The languages defined above allow for compensations among goals (the violation of a goal may be compensated by the satisfaction of a sufficient number of goals of lower importance). Prioritization is used when such a compensation should not be possible, and does not need any numerical data.

While a few approaches make use of a partial priority preorders, most of them make the assumption that the priority relation is complete (or, if it is defined as a strict order, that it is ranked, i.e., its complement is transitive), and for the sake of simplicity, we make this assumption too. When the priority relation is complete, priorities of formulas can be equivalently expressed (and in a simpler way) by a function r from integers into integers. A goal base is then a finite set of formulas with an associated function: $GB = \{\{G_1, \dots, G_n\}, r\}$. If $r(i) = j$, then j is called the rank of the formula G_i . By convention, a lower rank means a higher priority. In the following we use the convention $\min \emptyset = +\infty$

The question is now how to extend the priority on goals to a preference relation on alternatives. The following three choices are the most frequent ones:

best-out ordering ($R_{prio}^{bestout}$) (Benferhat *et al.* 1993)

Let us define

$$r_{GB}(M) = \min\{r(i) \mid M \not\models G_i\}$$

Then we define

$$\mid M \succeq_{GB}^{bo} M' \text{ if and only if } r_{GB}(M) \geq r_{GB}(M')$$

discrimin ordering ($R_{prio}^{discrimin}$) (Brewka 1989; Geffner 1992; Benferhat *et al.* 1993)

Let

$$discr_{GB}^+(M, M') = \{i \mid M \models G_i \text{ and } M' \not\models G_i\}$$

and

$$discr_{GB}(M, M') = discr_{GB}^+(M, M') \cup discr_{GB}^+(M', M)$$

Then:

$$\begin{array}{l} M \succ_{GB}^{discrimin} M' \\ \text{if and only if} \\ \min_{i \in discr_{GB}^+(M, M')} r(i) < \min_{j \in discr_{GB}^+(M', M)} r(j) \\ \\ M \succeq_{GB}^{discrimin} M' \\ \text{if and only if} \\ M \succ_{GB}^{discrimin} M' \text{ or } discr_{GB}(M, M') = \emptyset. \end{array}$$

leximin ordering ($R_{prio}^{leximin}$) (Benferhat *et al.* 1993; Lehmann 1995)

Let $d_k(M)$ be the cardinal of $\{i \mid M \models G_i \text{ and } r(i) = k\}$.

$$\begin{array}{l} M \succ_{GB}^{leximin} M' \text{ if and only if there is a } k \text{ such that} \\ d_k(M) > d_k(M') \text{ and } \forall j < k, d_j(M) = d_j(M'); \\ \\ M \succeq_{GB}^{leximin} M' \text{ if and only if} \\ M \succ_{GB}^{leximin} M' \text{ or } d_i(M) = d_i(M') \text{ for any } i. \end{array}$$

Note that $\succeq_{GB}^{leximin}$ and \succeq_{GB}^{bo} are complete preference relations while $\succeq_{GB}^{discrimin}$ is generally not. We moreover have the following chain of implications: $M \succ_{GB}^{bo} M' \Rightarrow M \succ_{GB}^{discrimin} M' \Rightarrow M \succ_{GB}^{leximin} M'$.

Example 3

- $K = \top$;
- $GB = (\{G_1, G_2, G_3\}, r)$

where

- $G_1 = a \wedge b$; $G_2 = \neg a \wedge c$; $G_3 = c \rightarrow d$;
- $r(3) = 1, r(1) = r(2) = 2$.

Let $M_1 = [a, b, c, \neg d]$, $M_2 = [a, b, c, d]$, $M_3 = [a, \neg b, c, d]$ and $M_4 = [\neg a, b, c, d]$.

best-out

We have $r_{GB}(M_1) = 1, r_{GB}(M_2) = r_{GB}(M_3) = r_{GB}(M_4) = 2$, hence

$$M_1 \prec_{GB}^{bo} M_2 \sim_{GB}^{bo} M_3 \sim_{GB}^{bo} M_4.$$

discrimin

We have

$$\text{discr}_{GB}^+(M_1, M_2) = \emptyset; \text{discr}_{GB}^+(M_2, M_1) = \{3\};$$

$$\min_{i \in \text{discr}_{GB}^+(M_1, M_2)} r(i) = +\infty;$$

$$\min_{i \in \text{discr}_{GB}^+(M_2, M_1)} r(i) = 1$$

therefore $M_2 \succ_{GB}^{\text{discrimin}} M_1$. As for M_2 and M_4 , we have $\text{discr}_{GB}^+(M_2, M_4) = \{1\}$, $\text{discr}_{GB}^+(M_4, M_2) = \{2\}$, hence M_2 and M_4 are incomparable w.r.t. $\succeq_{GB}^{\text{discrimin}}$. We have the following:

$$M_1 \prec_{GB}^{\text{discrimin}} M_3$$

$$M_3 \prec_{GB}^{\text{discrimin}} M_2$$

$$M_3 \prec_{GB}^{\text{discrimin}} M_4$$

M_2, M_4 incomparable.

leximin

We have the following:

| | M_1 | M_2 | M_3 | M_4 |
|-------|-------|-------|-------|-------|
| d_1 | 0 | 1 | 1 | 1 |
| d_2 | 1 | 1 | 0 | 1 |

Therefore, we have

$$M_1 \prec_{GB}^{\text{leximin}} M_3 \prec_{GB}^{\text{leximin}} M_2 \sim_{GB}^{\text{leximin}} M_4$$

Let us note at this point that *qualitative choice logic* (QCL) (Brewka, Benferhat, & Berre 2002) falls into the class of priority languages. Indeed, (Brewka, Benferhat, & Berre 2002) gives a polynomial-size translation from QCL formulas to stratified knowledge bases (interpreted with the leximin ordering) preserving the preference relation induced; now, it can be shown that there exists as well a polynomial size translation in the other direction, which implies that both languages have the same succinctness power and that results obtained for $\succeq_{GB}^{\text{leximin}}$ carry over to QCL. However, they do not carry over to logic programming with ordered disjunction (Brewka 2002), the succinctness of which is left for further research.

Conditional Logics

Each goal G_i is attached to a *context* C_i , so that $GB = \{C_1 : G_1, \dots, C_n : G_n\}$. Each conditional desire $C_i : G_i$ is interpreted as “ideally C_i if G_i ” (Boutilier 1994), or: “in all of the most preferred alternatives in which C_i holds, G_i holds as well.” Formally, a complete preorder \succeq satisfies $C_i : G_i$ if and only if every model of $\max(\text{Mod}(C_i), \leq)$ is a model of $\text{Mod}(G_i)$. On the other hand, more than one complete preorder may satisfy this condition and there are several ways for deriving a preference relation from this set of complete preorders. The standard preference relation R_{cond}^S consists in considering that $M \succeq_{GB} M'$ if and only if this holds for all relations \succeq that satisfy the condition above for all goals $C_i : G_i$, while the preference R_{cond}^Z derived from the Z -completion (Pearl 1990) selects *one* specific preorder relation, by making worlds gravitate towards preference (Boutilier 1994) and allows much more consequences to be derived and therefore makes more sense. To be more precise:

Standard preference relation (R_{cond}^S). R_{cond}^S consists in considering that an alternative is at least as good as another one if and only if this holds in *all* models of GB .

Formally: $M \succeq_{GB}^{cond,S} M'$ if and only if for any \succeq satisfying GB we have $M \succeq M'$. Note that $\succeq_{GB}^{cond,S}$ is only a partial preorder which is generally very weak, often much too weak (it does not enable enough comparisons) to be a good alternative for preference representation.

Preference relation based on Z -ranking (R_{cond}^Z). While R_{cond}^S considered *all* models satisfying a set of conditionals, the approach based on the Z -completion of GB , at work in System- Z (Pearl 1990) and similar approaches, selects *one* model and allows much more consequences to be derived. Given a theory $\Delta = \langle K, GB \rangle$ where K is a set of hard constraints (propositional formulas) and $GB = \{\delta_1, \dots, \delta_n\}$ a set of conditional rules, where $\delta_i = C_i : G_i$ (with the condition that $C_i \wedge G_i$ is consistent), System- Z proceeds by determining a partition $\{GB_0, \dots, GB_q\}$ of GB by the following procedure, for which we need first this definition: a conditional rule $C : G$ is *tolerated* by a set of conditional rules $\{C_1 : G_1, \dots, C_m : G_m\}$ (w.r.t. K) if and only if $C \wedge G \wedge (C_1 \rightarrow G_1) \wedge \dots \wedge (C_m \rightarrow G_m) \wedge K$ is consistent.

$k := 0; R := GB;$

repeat

$k := k + 1$

$R_k := \emptyset;$

for each $\delta \in R$

if δ is tolerated by $R \setminus \{\delta\}$

then add δ to R_k and remove it from R

endif

endfor

until $R = \emptyset; \text{maxrank} := k$

This procedure stops because each individual rule in Δ is consistent (that is, $C_i \wedge G_i$ is consistent). If a conditional rule δ is in R_k then we let $\text{rank}(\delta) = k$. Ranks respect specificity relations between rules, i.e., more specific rules are assigned higher ranks. Now, let $\text{MatImp}(\delta)$ be the material implication $C \rightarrow G$ associated with $\delta = C : G$. Then, $\succeq_{GB}^{cond,Z}$ is defined by the *best-out* preference relation generated by $GB' = \langle \{\text{MatImp}(\delta_1), \dots, \text{MatImp}(\delta_n)\}, r \rangle$ where $r(i) = \text{maxrank} - \text{rank}(\delta_i) + 1$.

Example 4 Let

• $K = \top;$

• $GB = \{\top : a, a : \neg b, \neg a : b, \neg a \wedge \neg b : c\}$

The first iteration gives $R_1 = \{\top : a, a : \neg b\}$, the second one gives $R_2 = \{\neg a : b\}$ and the last one $R_3 = \{\neg a \wedge \neg b : c\}$. We have $\text{maxrank} = 3$. Thus, $GB' = \langle \{G_1, G_2, G_3, G_4\}, r \rangle$ with $G_1 = a, G_2 = a \wedge \neg b, G_3 = \neg a \wedge b, G_4 = \neg a \wedge \neg b \wedge c$, and $r(4) = 1, r(3) = 2, r(1) = r(2) = 3$. Let $M_1 = [\neg a, \neg b, \neg c], M_2 = [\neg a, \neg b, c], M_3 = [a, b, c], M_4 = [a, \neg b, c]$. We have

$$M_1 \prec_{GB}^{cond,Z} M_2 \prec_{GB}^{cond,Z} M_3 \prec_{GB}^{cond,Z} M_4$$

Intuitively speaking, $\succeq_{GB}^{cond,Z}$ is the preference relation, among those satisfying GB , maximizing preference world

by world ((Boutilier 1994), page 79). The obtained relation $\succeq_{GB}^{cond,Z}$ is much more discriminant (hence much better) than $\succeq_{GB}^{cond,S}$. Nevertheless, one drawback of both $R^{cond,Z}$ and $R^{cond,S}$ is the so-called “drowning effect” (some goals are ignored while they should not); this can be remedied for instance by adding extra constraints expressing that violating a conditional desire induces an explicit utility loss (Lang 1996; Lang, van der Torre, & Weydert 2002).

Ceteris Paribus Preferences (R_{cp})

In this language, preferences are expressed in terms of statements like: “all other things being equal, I prefer these alternatives over these other ones.” Formally, let C , G , and G' be three propositional formulas and V being a subset of VAR such that $Var(G) \cup Var(G') \subseteq V$. The *ceteris paribus* desire $C : G > G'[V]$ means: “all irrelevant things being equal, I prefer $G \wedge \neg G'$ to $\neg G \wedge G'$ ”, where the “irrelevant things” are the variables that are not in V . The definitions proposed in various places (Doyle & Wellman 1991; Doyle, Wellman, & Shoham 1991; Tan & Pearl 1994; Boutilier *et al.* 1999) differ somehow. We take as a basis the definition by (Doyle & Wellman 1991), slightly generalized by introducing a variable set V that explicitly tells the variables that are referred to in the clause “all other things being equal” (namely, the other things are the variables not in V .) For natural reasons, and to remain consistent with the original definitions, we impose that $Var(G) \cup Var(G') \subseteq V$. Furthermore, we add to the original definition the ability to express *indifference statements* – without them, $M \sim M'$ could not be expressed.

Let $GB = \mathcal{D}_P \cup \mathcal{D}_I$, where \mathcal{D}_P and \mathcal{D}_I are defined as follows.

$$\begin{aligned} \mathcal{D}_P &= \{C_1 : G_1 > G'_1[V_1], \dots, C_m : G_m > G'_m[V_m]\} \\ \mathcal{D}_I &= \{C_n : G_n \sim G'_n[V_n], \dots, C_p : G_p \sim G'_p[V_p]\} \end{aligned}$$

We call the elements of \mathcal{D}_P as “preference desires” while elements of \mathcal{D}_I are “indifference desires”. For all i , C_i , G_i and G'_i are propositional formulas and $Var(G_i) \cup Var(G'_i) \subseteq V_i \subseteq VAR$. We define the preference induced by a single desire $D_i = C_i : G_i > G'_i[V_i]$, denoted by $M >_{D_i} M'$, by the following three conditions:

1. $M \models C_i \wedge G_i \wedge \neg G'_i$;
2. $M' \models C_i \wedge \neg G_i \wedge G'_i$;
3. M and M' coincide on all variables in $VAR \setminus V_i$.

If the above conditions 1-3 are satisfied for an indifference desire $D_i = C_i : G_i \sim G'_i[V_i]$ in \mathcal{D}_I , then we say that M and M' are *indifferent* with respect to D_i , denoted by $M \sim_{D_i} M'$. Now, the preference order \succeq_{GB}^{cp} is defined from the above dominance relations by transitive closure of their union: $M \succeq_{GB}^{cp} M'$ holds if and only if there exists a finite chain $M_0 = M, M_1, \dots, M_{q-1}, M_q = M'$ of alternatives such that for all $j \in \{0, \dots, q-1\}$ there is a $D_i \in GB$ such that $M_j >_{D_i} M_{j+1}$ or such that $M_j \sim_{D_i} M_{j+1}$.

Example 5 Let $K = \top$ and $GB = \{\top : a > \neg a [\{a\}], a : b > c [\{b, c\}], \neg a : c > b [\{b, c\}], \top : d > \neg d [\{c, d\}], \top :$

$e \sim \neg e [\{e\}]\}$. Then we have the following:

$$\begin{aligned} & \gamma_{GB}^{cp} [a, b, \neg c, d, e] \\ & \gamma_{GB}^{cp} [a, \neg b, c, d, e] \text{ (using } a : b > c [\{b, c\}]) \\ & \gamma_{GB}^{cp} [\neg a, \neg b, c, d, e] \text{ (using } \top : a > \neg a [\{a\}]) \\ & \gamma_{GB}^{cp} [\neg a, b, \neg c, d, e] \text{ (using } \neg a : c > b [\{b, c\}]) \\ & \gamma_{GB}^{cp} [\neg a, b, c, \neg d, e] \text{ (using } \top : d > \neg d [\{c, d\}]) \\ & \sim_{GB}^{cp} [\neg a, b, c, \neg d, \neg e] \text{ (using } \top : e \sim \neg e [\{e\}]) \end{aligned}$$

Synthesis of Results

We now present the results we have obtained when considering both the expressiveness and the spatial efficiency dimensions to compare the preference representation languages described before.

Expressiveness

The first results of this paper tell which orders can be expressed in the languages we consider. Namely, we show that some languages are only able to express complete preorders, while other ones are able to express all preorders.

Theorem 1

1. $R_{penalties}$, R_H , $R_{prio}^{bestout}$, $R_{prio}^{leximin}$ and R_{cond}^Z can express all complete preorders (and nothing more);
2. $R_{prio}^{discrimin}$ and R_{cp} can express all preorders;
3. R_{cond}^S cannot even express all complete preorders.

These results are not surprising. They are rather positive: all languages considered (except R_{cond}^S) are fully expressive for what they are designed to. Notice also that extending R_{cond}^S with *negated conditionals* ($\neg(C : G)$) enables expressing all preorders.

This result has two obvious consequences:

- let R_1 a language expressing all preorders and R_2 a language expressing only complete preorders. Then there cannot be any translation from R_1 to R_2 (because it is not possible to find a translation in R_2 of a preference item of R_1 inducing an incomplete preorder).
- let R_1 and R_2 be two languages of the same expressive power (that is, both expressing all preorders, or both expressing all complete preorders). Then R_1 and R_2 can be translated to each other. *Note however that the result of such a translation can be exponentially large.*

Succinctness

The second kind of results are about the existence of a polynomial translation of any preorder from a language into another one in another language. This is simply proved by exhibiting the translation.

The largest part of the work, however, is composed of results of impossibility of always translating preferences from a language to another one while preserving the size of the original goal base. Three methods are used to prove these results:

1. show that a specific preorder can be expressed in a language by a goal base whose size is polynomial in the number of variables, while it cannot in the other language. We consider three classes of such specific preorders:

- *exponentially long chains.* An exponentially long chain is a preference relation of the form $M_1 \succ M_2 \succ \dots \succ M_{2^n}$ where n is the number of propositional variables. Some of the languages (namely $R_{penalties}$, $R_{distances}$, $R_{prio}^{leximin}$, $R_{prio}^{discrimin}$, R_{cp}) can express such chains in polynomial space while some others (namely $R_{prio}^{bestout}$, R_{cond}^S , R_{cond}^Z) cannot. Therefore, there is no polynomial size translation from one language of the former class into a language of the latter class.
 - *exponentially many equivalence classes.* We consider the (partial) preference relation such that $M \succeq M'$ if and only if M and M' are “opposite” models (this preference relation has 2^{n-1} equivalence classes, two models of different equivalence classes being incompatible). Among the languages expressing all orderings, $R_{prio}^{discrimin}$ can express this preference relation in polynomial size while R_{cp} cannot.
 - *exponentially large equivalence classes.* We consider the preference relation where $M \succeq M'$ holds for all M, M' . R_{cond}^S cannot express it in polynomial space while all other languages can.
2. show that the existence of such a translation contradicts some results on circuit complexity, such as the impossibility of expressing the majority function with a CNF circuit of size polynomial in the number of variables. Technically, this is done by showing that the preference relation such that $M \succeq M'$ if and only if M has at least as many positive literals as M' can be expressed in polynomial space in some languages but not in other ones. In particular, the languages $R_{distances}$, $R_{penalties}$, $R_{prio}^{leximin}$, and R_{cp} can express this preference relation in polynomial size. The languages $R_{prio}^{discrimin}$, $R_{prio}^{bestout}$, R_{cond}^S , and R_{cond}^Z can express this ordering only using formulae that, combined in some way, express exactly the majority function. The majority function can be represented in polynomial space in general, but do not if we restrict to specific syntactic forms such as CNF. As a result, the languages $R_{prio}^{discrimin}$, $R_{prio}^{bestout}$, R_{cond}^S , and R_{cond}^Z cannot express the preference relation based on the number of positive literals only if we restrict to CNF. We remark that this result only holds under this language restriction, but is not conditioned to the collapse of the polynomial hierarchy.
 3. prove a result based on complexity classes; such results are conditioned to the non-collapse of the polynomial hierarchy. The impossibility of translating a preference language from a language to another one in polynomial time can be proved as follows: prove that checking whether $M \succeq M'$ is hard in one language and easy in another one. If a translation from the first language from the second were possible in polynomial time, then some parts of the polynomial hierarchy would collapse, contradicting a widely-accepted conjecture. Such a result, however, only indicates that the translation would require super-polynomial time; it may very well be that the result of the translation is polynomial in size. Since the size of the result is what matters in this paper, we cannot directly use the conventional complexity classes. We instead use the

compilability classes, which characterize the complexity of problems when preprocessing is allowed (Cadoli *et al.* 2002; Liberatore 2001). Intuitively, we can view the translation of the problem $M \succeq M'$ from a language to another one as a problem of finding a new goal base in the second language; we however do not require this translation to take only a polynomial amount of time, but only that the result is polynomially large. This is exactly the kind of problems the classes of compilability characterize. Somehow, proving impossibility in this way gives a stronger result, as a polysize translation does not exist even if we are allowed to translate the models M and M' ; on the other hand, these results are conditioned to the non-collapse of the polynomial hierarchy.

Theorem 2 *The results about translations and impossibility of translations are summarized in Table 1.*

This table deserves a few comments.

- **no** means a provable impossibility of a polysize translation. We indicate the type of proof used for showing this impossibility:
 - ¹ for an impossibility due to the fact that the first language can express all preorders while the second does not (typically, the second language expresses only complete preorders);
 - ² for an impossibility due to the fact that there are specific classes of preorders that the first language can express in polynomial space while the second one cannot.
- **no-c** (“conditional no”) means that there cannot be any polysize translation unless the polynomial hierarchy collapses.
- **w-no** (“weak no”) means that there cannot be any polysize translation if we require the formulae of the resulting language to be in CNF.

One of the cells of the table contains both a “w-no” and a “no-c”. These two results do not imply each other: “w-no” means that no polysize translation is possible if we restrict formulae to a specific syntactic form, but this result is not conditioned (*i.e.*, is proved for sure); “no-c” means that the impossibility of translation holds in general (*i.e.*, even if we do not restrict the syntactic form of formulae) but only if the polynomial hierarchy does not collapse.

One more class of results has not been reported in the table in order to make it more readable. These results are about translations that require the introduction of new variables. Indeed, all other results (both positive and negative) about translations have been proved by assuming that the goal base and its translation contain exactly the same variables. This restriction is not irrelevant: the addition of new variables increases the number of possible formulae that represent the same function, and some of these formulae may be very short. For example, as it has been proved for a number of nonmonotonic formalisms (Cadoli *et al.* 1999), the introduction of new variables may superpolynomially reduce the minimal size of formulae needed for representing a piece

Table 1: Existence of Polynomial-Size Translations.

| from \ to | R_{pen} | R_H | R_{prio}^{bo} | R_{prio}^{lexi} | R_{prio}^{disc} | R_{cond}^S | R_{cond}^Z | R_{cp} |
|-------------------|-----------------|-----------------|-----------------|-------------------|-------------------|-----------------|-----------------|-----------------|
| R_{pen} | yes | yes | no ² | ? | w-no | no ² | no ² | no ² |
| R_H | no-c | yes | no ² | no-c | w-no/no-c | no ² | no ² | ? |
| R_{prio}^{bo} | yes | yes | yes | yes | yes | no ² | yes | yes |
| R_{prio}^{lexi} | yes | yes | no ² | yes | w-no | no ² | no ² | no ² |
| R_{prio}^{disc} | no ¹ | no ¹ | no ¹ | no ¹ | yes | no ² | no ¹ | no ² |
| R_{cond}^S | no ¹ | no ¹ | no ¹ | no ¹ | ? | yes | no ¹ | ? |
| R_{cond}^Z | yes | yes | yes | yes | yes | no ² | yes | yes |
| R_{cp} | no ¹ | no ¹ | no ¹ | no ¹ | w-no | no ² | no ¹ | yes |

of information. In the setting of representing preference information, we can introduce new variables by assuming that a goal base GB that contains only the propositional variables X can be translated into a goal base GB' (in another language) that contains the variables $X \cup Y$. Intuitively, the new variables Y are used as names for “propositional macros”, *i.e.*, the goal base contains formulae like $y \equiv F$, where F is a formula that occurs often; this way, we can replace each occurrence of F with y (this is how size can be reduced.) Technically, we can no more define a translation from a language L_1 to L_2 by enforcing $M \succeq_{GB}^{L_1} M'$ to be equivalent to $M \succeq_{GB'}^{L_2} M'$ because M and M' are models over X , while GB' is a goal base over variables $X \cup Y$. The problem is that $M \succeq_{GB'}^{L_2} M'$ only makes sense if M and M' are models over the variables of GB' . This problem can be overcome by extending the definition to partial models: $M \succeq_{GB'}^{L_2} M'$ holds if and only if there exists two models M_t and M'_t over $X \cup Y$ that extend M and M' , respectively (*i.e.*, M_t gives the same evaluation of M to the variables in X), and $M_t \succeq_{GB'}^{L_2} M'_t$. Using this definitions, the constraint that $M \succeq_{GB}^{L_1} M'$ is equivalent to $M \succeq_{GB'}^{L_2} M'$ now makes sense, as the latter formula is well-defined. Translations in this weak sense exist from $R_{prio}^{leximin}$ to $R_{prio}^{discrimin}$, from $R_{prio}^{discrimin}$ to R_{cp} and from $R_{prio}^{leximin}$ to R_{cp} . A formalization of this kind of translations where new variables can be introduced, but their values must be determinable from the value of the old ones in polynomial time, are the model-preserving translation (Cadoli *et al.* 2000).

Conclusion

In this paper, we have shown that some preference representation languages are more succinct than others. However, no language is the most succinct one: for any language, there is another one that cannot be translated into it in polynomial space. On the other hand, $R_{prio}^{bestout}$ and R_{cond}^Z can be translated into any other language except R_{cond}^S (which is however incomparable to everything else) and can therefore be considered as the least succinct languages of our list. While spatial efficiency results were known for knowledge representation formalisms such as nonmonotonic logics and revision (Cadoli *et al.* 1996) and action representation for planning (Nebel 2000), nothing so far was known about the spatial efficiency of preference representation languages.

Even if some of our impossibility results make use of complexity results for some decision problems, it is important to note that our results *cannot be directly* inferred from the complexity of the decision problems associated with these languages³. Indeed, even if a problem is hard in a language and easy in another one, the existence of a translation is not guaranteed: see for example R_{cp} , whose comparison problem is PSPACE-hard, but cannot be translated into the simpler $R_{prio}^{discrimin}$, where comparison is polynomial-time. Our results that are conditioned to the non-collapse of the polynomial hierarchy are not obtained using “standard” complexity classes.

We do not claim that the list of languages considered here is exhaustive. It is however representative of the languages developed by researchers in AI and especially in qualitative decision theory for representing *preference relations* compactly *using propositional logic*: this means that we consciously omitted from the study “numerical” languages for describing compactly utility functions, such as (Boutilier, Bacchus, & Brafman 2001) as well as CP-nets and constraints and let their succinctness as an open problem. Note however that these languages are close to propositional preference languages, up to syntactical restrictions – the fact that variables are not binary having no significant influence on succinctness. Furthermore, CP-nets with binary variables can be encoded in a straightforward way (in linear time) as *ceteris paribus* statements. However, CP-nets cannot encode all preorders, in contrast to the language R_{cp} ⁴.

Expressivity and succinctness of languages are particularly significant to *preference elicitation* (which, roughly speaking, consists in obtaining preference items directly from the agent until the preference relation is “sufficiently” described): indeed, the interactive elicitation process using some preference representation language may require exponential time, while polynomial time would be sufficient if a more succinct language were considered. Another impact of our results concerns the design of “mixed” languages based on two (or more) types of representations (which is particularly relevant when each language “cognitively fits” some

³Complexity results for some of the languages considered here can be found, or easily derived from results in (Friedman & Halpern 1994; Eiter & Lukasiewicz 2000; Nebel 1998; Lang 2002).

⁴Thanks to Christophe Gonzalès and Patrice Perny for pointing this out to us.

particular types of preferences); the existence or the impossibility of polysize translation help us identifying the languages that may or that may not be used together in such a mixed language.

Acknowledgements

Sylvie Coste-Marquis and Pierre Marquis would like to thank the IUT de Lens, the IRCICA, and the European Community FEDER program for their support.

References

- Bacchus, F., and Grove, A. 1995. Graphical models for preference and utility. In *Proc. of UAI'95*, 3–10.
- Benferhat, S.; Cayrol, C.; Dubois, D.; Lang, J.; and Prade, H. 1993. Inconsistency management and prioritized syntax-based entailment. In *Proc. of IJCAI'93*, 640–645.
- Benferhat, S.; Dubois, D.; Kaci, S.; and Prade, H. 2002. Bipolar representation and fusion of preferences in the possibilistic logic framework. In *Proc. of KR'02*, 421–432. Morgan Kaufmann.
- Bistarelli, S.; Fargier, H.; Montanari, U.; Rossi, F.; Schiex, T.; and Verfaillie, G. 1999. Semiring-based CSPs and valued CSPs: Frameworks, properties, and comparison. *CONSTRAINTS: An International Journal* 4(3):199–240.
- Boutilier, C.; Bacchus, F.; and Brafman, R. 2001. UCP-networks: a directed graphical representation of conditional utilities. In *Proc. of UAI'01*, 56–64.
- Boutilier, C.; Brafman, R.; Hoos, H.; and Poole, D. 1999. Reasoning with conditional *ceteris paribus* statements. In *Proc. of UAI'99*, 71–80.
- Boutilier, C. 1994. Toward a logic for qualitative decision theory. In *Proc. of KR'94*, 75–86.
- Brewka, G., and Eiter, T. 1999. Preferred answer sets for extended logic programs. *Artificial Intelligence* 109:297–356.
- Brewka, G.; Benferhat, S.; and Berre, D. L. 2002. Qualitative choice logic. In *Proc. of KR'02*, 158–169.
- Brewka, G. 1989. Preferred subtheories: an extended logical framework for default reasoning. In *Proc. of IJCAI'89*, 1043–1048.
- Brewka, G. 2002. Logic programming with ordered disjunction. In *Proc. of AAAI'02*, 100–105.
- Cadoli, M.; Donini, F.; Liberatore, P.; and Schaerf, M. 1996. Comparing space efficiency of propositional knowledge representation formalisms. In *Proc. of KR'96*, 364–373.
- Cadoli, M.; Donini, F. M.; Liberatore, P.; and Schaerf, M. 1999. The size of a revised knowledge base. *Artificial Intelligence* 115(1):25–64.
- Cadoli, M.; Donini, F.; Liberatore, P.; and Schaerf, M. 2000. Space efficiency of propositional knowledge representation formalisms. *Journal of Artificial Intelligence Research* 13:1–31.
- Cadoli, M.; Donini, F.; Liberatore, P.; and Schaerf, M. 2002. Preprocessing of intractable problems. *Information and Computation* 176(2):89–120.
- Dalal, M. 1988. Investigations into a theory of knowledge base revision: preliminary report. In *Proc. of AAAI'88*, 475–479.
- Delgrande, J.; Schaub, T.; and Tompits, H. 2003. A framework for compiling preferences in logic programs. *Theory and Practice of Logic Programming* 3(2):129–187.
- Domshlak, C., and Brafman, R. 2002. CP-nets: reasoning and consistency testing. In *Proc. of KR'02*, 121–132.
- Doyle, J., and Wellman, M. P. 1991. Preferential semantics for goals. In *Proc. of AAAI'91*, 698–703.
- Doyle, J.; Wellman, M. P.; and Shoham, Y. 1991. A logic of relative desire (preliminary report). In *Proc. of ISMIS'91*, 16–31.
- Dupin de Saint-Cyr, F.; Lang, J.; and Schiex, T. 1994. Penalty logic and its link with Dempster-Shafer theory. In *Proc. of UAI'94*, 204–211. Morgan Kaufmann.
- Eiter, T., and Lukasiewicz, T. 2000. Complexity results for default reasoning from conditional knowledge bases. In *Proc. of KR'00*, 62–73.
- Friedman, N., and Halpern, J. 1994. On the complexity of conditional logics. In *Proc. of KR'94*, 202–213.
- Geffner, H. 1992. Default reasoning: causal and conditional theories. *MIT Press*.
- Gogic, G.; Kautz, H.; Papadimitriou, C.; and Selman, B. 1995. The comparative linguistics of knowledge representation. In *Proc. of IJCAI'95*, 862–869.
- Gonzales, C., and Perny, P. 2004. Gai networks for utility representation. In *Proceedings of KR-2004 (this Volume)*.
- Gottlob, G. 1995. Translating default logic into standard autoepistemic logic. *Journal of the ACM* 42:711–740.
- Haddawy, P., and Hanks, S. 1992. Representations for decision theoretic planning: utility functions for deadline goals. In *Proc. of KR'92*, 71–82.
- Janhunen, T. 1998. On the intertranslatability of autoepistemic, default and priority logics, and parallel circumscription. In *Proc. of JELIA'98*, volume 1489 of *Lectures Notes in Artificial Intelligence*, 216–232. Springer-Verlag.
- Kaci, S. 2002. *Fusion de connaissances et de préférences en logique possibiliste*. Ph.D. Dissertation, Université Paul Sabatier.
- Katsuno, H., and Mendelzon, A. 1991. On the difference between updating a knowledge base and revising it. In *Proceedings of KR'91*, 387–394.
- Konieczny, S., and Pino-Pérez, R. 1998. On the logic of merging. In *Proc. of KR'98*, 488–498.
- Konieczny, S.; Lang, J.; and Marquis, P. 2002. Distance-based merging: a general framework and some complexity results. In *Proc. of KR'02*, 97–108.
- Lafage, C., and Lang, J. 2000. Logical representation of preferences for group decision making. In *Proc. of KR'00*, 457–468.

- Lafage, C., and Lang, J. 2001. Propositional distances and preference representation. In *Proc. of ECSQARU 01*, 48–59.
- Lang, J.; van der Torre, L.; and Weydert, E. 2002. Utilitarian desires. *International Journal on Autonomous Agents and Multi-Agent Systems* 5:329–363.
- Lang, J. 1996. Conditional desires and utilities - an alternative logical approach to qualitative decision theory. In *Proc. of ECAI'96*, 318–322.
- Lang, J. 2002. From preference representation to combinatorial vote. In *Proc. of KR'02*, 277–288.
- Lehmann, D. 1995. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence* 15(1):61–82.
- Liberatore, P., and Schaerf, M. 1998. Arbitration (or how to merge knowledge bases). *IEEE Transactions on Knowledge and Data Engineering* 10(1):76–90.
- Liberatore, P. 2001. Monotonic reductions, representative equivalence, and compilation of intractable problems. *Journal of the ACM* 48(6):1091–1125.
- Lin, J., and Mendelzon, A. O. 1999. Knowledge base merging by majority. In *Dynamic Worlds: From the Frame Problem to Knowledge Management*. Kluwer.
- Magidor, M., and Schlechta, K. 2001. Distance semantics for belief revision. *Journal of Symbolic Logic* 66(1):295–317.
- Nebel, B. 1998. How hard is it to revise a knowledge base? In *Handbook of Defeasible Reasoning and Uncertainty Management Systems*. Kluwer Academics. 77–145.
- Nebel, B. 2000. On the compilability and expressive power of propositional planning formalisms. *Journal of Artificial Intelligence Research* 12:271–315.
- Pearl, J. 1990. System Z : a natural ordering of defaults with tractable applications for default reasoning. In *Proc. of TARK'90*, 121–135.
- Pinkas, G. 1995. Reasoning, nonmonotonicity and learning in connectionist networks that capture propositional knowledge. *Artificial Intelligence* 77:203–247.
- Revesz, P. 1997. On the semantics of arbitration. *Int. Journal of Algebra and Computation* 133–160.
- Tan, S., and Pearl, J. 1994. Specification and evaluation of preferences for planning under uncertainty. In *Proc. of KR'94*, 530–539.
- van der Torre, L., and Weydert, E. 2001. Parameters for utilitarian desires in a qualitative decision theory. *Applied Intelligence* 14(3):285–302.