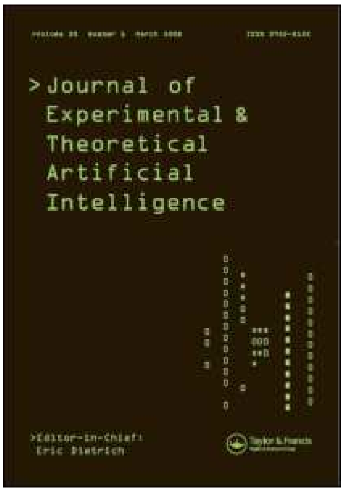

Journal of Experimental & Theoretical Artificial Intelligence
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713723652

From soft constraints to bipolar preferences: modelling framework and solving issues

Stefano Bistarelli [ab]; Maria Silvia Pini [c]; Francesca Rossi [c]; K. Brent Venable [c]
[a] Dipartimento di Scienze, Università 'G. d'Annunzio', Viale Pindaro 87, 65127 Pescara, Italy [b] Istituto di Informatica e Telematica, 56124 Pisa, Italy [c] Dipartimento di Matematica Pura ed Applicata, Università di Padova, 63-35121 Padova, Italy

First Published on: 21 October 2009

Journal of Experimental & Theoretical Artificial Intelligence
2009, 1–24, iFirst

Taylor & Francis
Taylor & Francis Group

From soft constraints to bipolar preferences: modelling framework and solving issues

Stefano Bistarelli[ab], Maria Silvia Pini[c*], Francesca Rossi[c] and K. Brent Venable[c]

[a]*Dipartimento di Scienze, Università 'G. d'Annunzio', Viale Pindaro 87, 65127 Pescara, Italy;* [b]*Istituto di Informatica e Telematica, CNR, Via G.Moruzzi 1, 56124 Pisa, Italy;* [c]*Dipartimento di Matematica Pura ed Applicata, Università di Padova, Via Trieste, 63-35121 Padova, Italy*

(*Received 16 July 2008; final version received 13 January 2009*)

Real-life problems present several kinds of preferences. We focus on problems with both positive and negative preferences, which we call *bipolar preference problems*. Although seemingly specular notions, these two kinds of preferences should be dealt with differently to obtain the desired natural behaviour. We technically address this by generalising the soft constraint formalism, which is able to model problems with one kind of preference. We show that soft constraints model only negative preferences, and we add to them a new mathematical structure which allows to handle positive preferences as well. We also address the issue of the compensation between positive and negative preferences, studying the properties of this operation. Finally, we extend the notion of arc consistency to bipolar problems, and we show how branch and bound (with or without constraint propagation) can be easily adapted to solve such problems.

Keywords: soft constraints; preferences; negative and positive judgements

1. Introduction

Many real-life problems contain statements which can be expressed as preferences. The problem of representing preferences of agents has been deeply investigated in Artificial Intelligence (AI) community in recent past years (Doyle, Shoham and Wellman 1991; Wellman and Doyle 1991; Lang 1996; Bacchus and Grove 1996; Boutilier, Brafman, Hoos and Poole 1999; Benferhat, Dubois, Kaci and Prade 2001; Benferhat, Dubois and Prade 2002). Preference representation is an important issue when we have to represent the desires of users or to reason about them, for example, in recommender systems.

Our long-term goal is to define a general and flexible framework where many kinds of preferences can be naturally modelled and efficiently dealt with. In this article, we focus on problems that present positive and negative preferences, which we call *bipolar preference problems*.

Bipolarity is an important topic in several domains, for example psychology (Osgood and Tannenbaum 1957; Tversky and Kahneman 1992; Cacioppo and Berntson 1997), multi-criteria decision making (Grabisch and Labreuche 2005), and more recently in AI (argumentation (Amgoud and Prade 2005) and qualitative reasoning (Benferhat et al. 2002, 2006; Dubois and Fargier 2005, 2006)). Preferences on a set of possible choices are

*Corresponding author. Email: mpini@math.unipd.it

ISSN 0952–813X print/ISSN 1362–3079 online

DOI: 10.1080/09528130903010212
http://www.informaworld.com

often expressed in two forms: positive and negative statements. In fact, in many real-life situations, agents express what they like and what they dislike, thus often preferences are bipolar.

Let us consider a real-life situation where positive and negative preferences are useful to model the problem (Benferhat et al. 2006). Consider a 3-day summer school, for which each lecturer is asked to express a preferred time slot for scheduling his talk. We assume that talks can be given either on Monday, on Tuesday or on Wednesday, and for each day the talk can be either scheduled in the morning or in the afternoon. A lecturer may provide two kinds of preferences. He may specify negative preferences, which describe unacceptable slots with levels of tolerance. For instance, he may strongly object to lecturing on Monday, while he may weakly refuse to speak on Wednesday. Moreover, a lecturer may also specify positive preferences. For instance, he may state that giving a talk in the morning is preferred to giving it in the afternoon, and scheduling it early in the morning is even better.

The aim of this article is to propose a tool to represent these two types of preferences in a single framework and to provide algorithms that, taken in input a problem with these two kinds of preferences, return its best solutions.

Positive and negative preferences can be thought as two symmetric concepts, and thus one can think that they can be dealt with via the same operators. However, this would not model what one usually expects in real scenarios. In fact, usually combination of positive preferences should produce a higher (positive) preference, while combination of negative preferences should produce a lower (negative) preference.

For example, if Paul likes (resp., dislikes) meat and if he likes (resp., dislikes) tomatoes, then, if he does not say explicitly that he dislikes (resp., likes) a meal with meat and tomatoes, then certainly he will like (resp., dislikes) more to eat both meat and tomatoes. As another example, consider a decision d that is obtained by combining two decisions d_1 and d_2, that both produce advantages, i.e. they are both associated to positive preferences. Then, the global decision d will produce a higher advantage than d_1 and d_2 and thus, as expected, a higher positive preference.

On the other hand, if d_1 and d_2 produce only disadvantages, i.e. if they are both associated with negative preferences, it is clear that decision d will produce a higher disadvantage; thus it will be associated with a more negative preference than those of d_1 and d_2.

Note that, even if positive and negative preferences could be seen as symmetric concepts, we will show that it is not sufficient to use a single structure to model, for example, the negative preferences and to transform every positive preference into a negative one. Moreover, the representation of positive and negative preferences in a separate way is motivated by recent studies in psychology showing that the distinction between positive and negative preferences make sense. In fact, they are processed separately in the brain and they are felt as different dimensions by people (Cacioppo and Berntson 1997, 1999).

When dealing with both kinds of preferences, it is also natural to express indifference, which means that we express neither a positive nor a negative preference over an object. Then, a desired behaviour of indifference is that, when combined with any preference (either positive or negative it should not influence the overall preference.

Finally, besides combining preferences of the same type, we also want to combine positive with negative preferences. We strongly believe that the most natural and intuitive way to do so is to allow for compensation. Comparing positive against negative aspects and compensating them w.r.t. their strength is one of the core features of decision-making processes, and it is, undoubtedly, a tactic universally applied to solve many real-life problems.

This approach follows the same idea proposed in chapter IV of Bentham (1789) for evaluating the tendency of an act. Such an idea consists of summing up all the values of all the pleasures produced by the considered act on one side, and those of all the pains produced by it on the other, and then balancing these two resulting values in a value which can be on the side of pleasure or on the side of pain. If this value is on the side of pleasure, then the tendency of the act is good, otherwise the tendency is bad.

For example, assume to have to choose among two decisions, say d_1 and d_2. Assume also to know that decision d_1 has 10 argumentations in favour and 5 argumentations against, and that decision d_2 has 15 argumentations in favour and 6 argumentations against. If all the argumentations have the same importance, as in this case, in order to choose the best decision, a classical approach (Bentham 1789) requires to compensate (via the subtraction operator) the positive and the negative argumentations for every decision and then to choose the one with the highest compensation value. In this example, the best solution is therefore decision d_2, since globally it produces more advantages than d_1. In fact, globally d_1 has $10-5=5$ argumentation in favour, while d_2 has $15-6=9$ argumentations in favour. In general, compensation may be less simple than this, since, for example, we may have qualitative considerations over the argumentations, rather than just quantitative ones.

Positive and negative preferences might seem as just two different criteria to reason with, and thus techniques such as those usually adopted by multi-criteria optimisation (Ehrgott and Gandibleux 2002) could appear suitable for dealing with them. However, this interpretation would hide the fundamental nature of bipolar preferences, that is, that positive preferences are naturally opposite of negative preferences.

For example, pros and cons of a decision, that can be seen as positive and negative preferences, are opposite elements, thus it would not be reasonable to consider them separately as done usually in a multi-criteria approach. Moreover, in multi-criteria optimisation it is often reasonable to use a Pareto-like approach, thus associating tuples of values to each solution, and comparing solutions according to tuple dominance. Instead, in bipolar problems, it would be very unnatural to force such an approach in all contexts, or to associate to a solution a preference which can be neither a positive nor a negative one as done in Fargier and Wilson (2007).

Soft constraints (Bistarelli, Montanari and Rossi 1997; Bistarelli 2004) are a useful formalism to model problems with quantitative preferences. However, they can only model negative preferences, since in this framework preference combination returns lower preferences. In this article we adopt the soft constraint formalism based on semirings to model negative preferences. We then define a new algebraic structure to model positive preferences. To model bipolar problems, we link these two structures and we set the highest negative preference to coincide with the lowest positive preference to model indifference. We then define a combination operator between positive and negative preferences to model preference compensation, and we study its properties.

Notice that we generalise soft constraint formalism to model bipolar preferences to be able to exploit the generality of the semiring-based soft constraint framework, that allows us to represent several kinds of preferences and also partially ordered preferences, and to exploit and adapt the existing solving machinery for soft constraint problems to obtain optimal solutions for bipolar preference problems.

Non-associativity of preference compensation occurs in many contexts, thus we think it is too restrictive to focus just on associative environments. For example, non-associativity of compensation arises when either positive or negative preferences are aggregated with an idempotent operator (such as min or max), while compensation is instead non-idempotent (such as sum). Our framework allows for non-associativity, since we want to give complete freedom to choose the positive and negative algebraic structures. However, we also describe a technique that, given a negative structure, builds a corresponding positive structure and an associative compensation operator.

Finally, we consider the problem of finding optimal solutions of bipolar problems, by suggesting a possible adaptation of constraint propagation and branch-and-bound (BB) to the generalised scenario.

Summarising, the main results of this article are:

- a formal definition of an algebraic structure to model bipolar preferences;
- the study of the notion of compensation and of its properties (such as associativity);
- a technique to build a bipolar preference structure with an associative compensation operator;
- the adaptation of BB to solve bipolar problems;
- the definition of bipolar preference propagation and its use within a BB solver.

This article is organised as follows. Section 2 recalls the main notions of semiring-based soft constraints. Section 3 describes how to model negative preferences using usual soft constraints and Section 4 shows how to model positive preferences. Section 5 shows how to model both positive and negative preferences and Section 6 defines constraint problems with both positive and negative preferences. Section 7 shows how to modify the bipolar framework to handle also negative preferences that cannot be compensated. Section 8 shows that it is important to have a bipolar structure for expressing both positive and negative preferences, since expressing all the problems' requirements in a positive (or negative) form might lead to different optimal solutions. Section 9 shows that very often the compensation operator is not associative and it describes a technique to build a bipolar preference structure with an associative compensation operator. Section 10 shows how to adapt BB to solve bipolar problems, how to define bipolar propagation and its use within a BB solver. Finally, Section 11 describes the existing related work and gives some hints for future work.

Earlier versions of parts of this article have appeared in Bistarelli, Pini, Rossi and Venable (2006, 2007a).

2. Semiring-based soft constraints

A soft constraint (Bistarelli et al. 1997; Bistarelli 2004) is just a classical constraint (Dechter 2003) where each instantiation of its variables has an associated value from a (totally or partially ordered) set. Combining constraints will then have to take into account such additional values, and thus the formalism has also to provide suitable operations for combination ($\times$) and comparison ($+$) of tuples of values and constraints. This is why this formalisation is based on the concept of c-semiring, which is just a set plus two operations. More precisely, a c-semiring is a tuple $\langle A, +, \times, \mathbf{0}, \mathbf{1}\rangle$ (when $+$ (respectively $\times$) is applied to a two-element set we will use symbol + (respectively $\times$) in infix notation, while in general we will use the symbol $\sum$ (respectively $\prod$) in prefix notation.) such that:

- A is a set and $\mathbf{0}, \mathbf{1} \in A$;
- $+$ is commutative, associative, idempotent, $\mathbf{0}$ is its unit element and $\mathbf{1}$ is its absorbing element;
- $\times$ is associative, commutative, distributes over $+$, $\mathbf{1}$ is its unit element and $\mathbf{0}$ is its absorbing element.

Given the relation $\leq_S$ over A such that $a \leq_S b$ iff $a + b = b$, it is possible to prove that:

- $\leq_S$ is a partial order;
- $+$ and $\times$ are monotone on $\leq_S$;
- $\mathbf{0}$ is its minimum and $\mathbf{1}$ is its maximum;
- $\langle A, \leq_S \rangle$ is a lattice and, for all $a, b \in A$, $a + b = lub(a, b)$.

Moreover, if $\times$ is idempotent, then $\langle A, \leq_S \rangle$ is a distributive lattice and $\times$ is its greatest lower bound. Informally, the relation $\leq_S$ gives us a way to compare (some of the) tuples of values and constraints. In fact, when we have $a \leq_S b$, we will say that *b is better than a*.

Given a c-semiring $S = (A, +, \times, \mathbf{0}, \mathbf{1})$, a finite set D (the domain of the variables), and an ordered set of variables V, a constraint is a pair $\langle def, con \rangle$ where $con \subseteq V$ and $def : D^{|con|} \rightarrow A$. Therefore, a constraint specifies a set of variables (the ones in *con*), and assigns to each tuple of values of D of these variables an element of the semiring set A.

Given a subset of variables $I \subseteq V$, and a soft constraint $c = \langle def, con \rangle$, the projection of c over I, written $c \Downarrow_I$, is a new soft constraint $\langle def', con' \rangle$, where $con' = con \cap I$ and $def(t') = \sum_{\{t | t\downarrow_{con'} = t'\}} def(t)$. The scope, con' of the projection constraint contains the variables that *con* and I have in common, and thus $con' \subseteq con$. Moreover, the preference associated to each assignment to the variables in con' denoted with t' is the highest ($\sum$ is the additive operator of the c-semiring) among the preferences associated by *def* to any completion of t' t, to an assignment to *con*.

A soft constraint satisfaction problem (SCSP) is a set of soft constraints over a set of variables.

Choosing a specific semiring means selecting a class of preferences. For example,

- the semiring $S_{FCSP} = ([0, 1], \max, \min, 0, 1)$ allows one to model fuzzy CSPs (Schiex 1992; Ruttkay 1994), which associate to each element allowed by a constraint a preference between 0 and 1 (with 0 being the worst and 1 being the best preference), and gives to each complete assignment a preference that is the minimal among all preferences selected in the constraints. The optimal solutions are then those solutions with the maximal preference;
- the semiring $S_{CSP} = (\{false, true\}, \vee, \wedge, false, true)$ allows one to model classical CSPs (Rossi et al. 2006), without preferences. The only two preferences that can be given are *true*, indicating that the tuple is allowed and *false*, indicating that the tuple is forbidden. The optimal solutions are those solutions with preference *true* on all constraints;
- the semiring $S_{WCSP} = (\Re^+, \min, +, +\infty, 0)$, allows one to model weighted SCSPs (Schiex, Fargier and Verfaillie 1995). Preferences are interpreted as costs from 0 to $+\infty$, and each complete assignment is associated to a cost that is obtained by summing all costs selected in the constraints. The optimal solutions are those solutions with the minimal cost and

- the semiring $S_{PCSP} = ([0, 1], \max, \times, 0, 1)$, allows one to model probabilistic SCSPs (Fargier and Lang 1993). Preferences are interpreted as probabilities ranging from 0 to 1. Every solution is associated to a probability, that is, the joint probability, which is obtained by multiplying the probabilities selected in the constraints.

Given an assignment s to all the variables of an SCSP, we can compute its preference value $pref(s)$ by combining the preferences associated by each constraint to the subtuples of the assignments referring to the variables of the constraint. More precisely, given an SCSP P defined by a set of constraints C over a c-semiring S and by a set of variables V, the preference of an assignment s to all the variables of P is $pref(s) = \prod_{\langle def, con \rangle \in C, con \subseteq V} def(s \downarrow con)$, where $\prod$ is the combination operator of the c-semiring S. An optimal solution of an SCSP is then a complete assignment t such that there is no other complete assignment t' with $pref(t) <_S pref(t')$.

3. Negative preferences

The structure we use to model negative preferences is exactly a c-semiring, as defined in Section 2. In fact, in a c-semiring the element which acts as indifference is $\mathbf{1}$, since $\forall a \in A$, $a \times \mathbf{1} = a$. This element is the best in the ordering, which is consistent with the fact that indifference is the best preference when using only negative preferences. Moreover, in a c-semiring, combination goes down in the ordering, since $a \times b \leq a, b$. This can be naturally interpreted as the fact that combining negative preferences worsens the overall preference.

Example 1: The above interpretation is very natural when considering, for example, the weighted semiring $(R^+, \min, +, +\infty, 0)$. In fact, in this case the real numbers are costs and thus negative preferences. The sum of two costs is never better than the two costs w.r.t. the ordering induced by the additive operator (that is, min) of the semiring.

Example 2: If, instead, we consider the fuzzy semiring, that is, $([0, 1], \max, \min, 0, 1)$, according to this interpretation, giving preference 1 to a tuple means that there is nothing negative about such a tuple. Instead, giving a preference strictly less than 1 (e.g. 0.6) means that there is at least a constraint which such tuple does not satisfy at the best. Moreover, combining two fuzzy preferences means taking the minimum and thus the worst among them.

Example 3: When considering classical constraints via the c-semiring $S_{\text{CSP}} = (\{false, true\}, \vee, \wedge, false, true)$, we just have two elements to model preferences: *true* and *false*. *True* is here the indifference, while *false* means that we do not like the object. This interpretation is consistent with the fact that, when we do not want to say anything about the relation between two variables, we just omit the constraint, which is equivalent to having a constraint where all instantiations are allowed (thus they are given value *true*).

From now on, we use $(N, +_n, \times_n, \perp_n, \top_n)$ and call it a negative preference structure, the c-semiring to model negative preferences.

4. Positive preferences

When dealing with positive preferences, we want two main properties to hold: combination should bring to better preferences, and indifference should be lower than all the other positive preferences. These properties can be found in the following structure.

Definition 1: A positive preference structure is a tuple $(P, +_p, \times_p, \perp_p, \top_p)$ such that

- P is a set and $\top_p, \perp_p \in P$;
- $+_p$, the additive operator, is commutative, associative, idempotent, with $\perp_p$ as its unit element ($\forall a \in P, a +_p \perp_p = a$) and $\top_p$ as its absorbing element ($\forall a \in P, a +_p \top_p = \top_p$);
- $\times_p$, the multiplicative operator, is associative, commutative and distributes over $+_p$ $(a \times_p (b +_p c) = (a \times_p b) +_p (a \times_p c))$, with $\perp_p$ as its unit element and $\top_p$ as its absorbing element. (The absorbing nature of $\top_p$ can be derived from the other properties.)

Notice that the additive operator of this structure has the same properties as the corresponding one in c-semirings, and thus it induces a partial order over P in the usual way: $a \leq_p b$ iff $a +_p b = b$. This allows to prove that $+_p$ is monotone over $\leq_p$ and that it computes the least upper bound in the lattice $(P, \leq_p)$. On the other hand, the multiplicative operator has different properties. More precisely, the best element in the ordering ($\top_p$) is now its absorbing element, while the worst element ($\perp_p$) is its unit element. This reflects the desired behaviour of the combination of positive preferences.

Theorem 1: *Given a positive preference structure* $(P, +_p, \times_p, \perp_p, \top_p)$, *consider the relation* $\leq_p$ *over* P. *Then*:

- $\times_p$ *is monotone over* $\leq_p$ (*i.e. for any* $a, b \in P$ *such that* $a \leq_p b$, *then* $a \times_p d \leq_p b \times_p d$, $\forall d \in P$) *and*
- $\forall a, b \in P$, $a \times_p b \geq_p a +_p b \geq_p a, b$.

Proof: Since $a \leq_p b$ iff $a +_p b = b$, then $b \times_p d = (a +_p b) \times_p d = (a \times_p d) +_p (b \times_p d)$. Thus $a \times_p d \leq_p b \times_p d$. Also, $a \times_p b = a \times_p (b +_p \perp_p) = (a \times_p b) +_p (a \times_p \perp_p) = (a \times_p b) +_p a$. Thus $a \times_p b \geq_p a$ (the same for b). Finally: $a \times_p b \geq a, b$. Thus $a \times_p b \geq lub(a, b) = a +_p b$.

In a positive preference structure, $\perp_p$ is the element modelling indifference. In fact, it is the worst one in the ordering and it is the unit element for the combination operator $\times_p$. These are exactly the desired properties for indifference w.r.t. positive preferences. The role of $\top_p$ is to model a very high preference, much higher than all the others. In fact, since it is the absorbing element of the combination operator, when we combine any positive preference a with $\top_p$, we get $\top_p$.

Example 4: An example of a positive preference structure is $P_1 = (R^+, \max, +, 0, +\infty)$, where preferences are positive reals. The smallest preference that can be assigned is 0. It represents the lack of any positive aspect and can thus be regarded as indifference. Preferences are aggregated taking the sum and are compared taking the max.

Example 5: Another example is $P_2 = ([0, 1], \max, \max, 0, 1)$. In this case preferences are reals between 0 and 1, as in the fuzzy semiring for negative preferences. However, the combination operator is max, which gives, as a resulting preference, the highest one among all those combined.

Example 6: As an example of a partially ordered positive preference structure, consider the Cartesian product of the two structures described above: $(R^+ \times [0, 1], (\max, \max), (+, \max), (0, 0), (+\infty, 1))$. Positive preferences, here, are ordered pairs where the first element is a positive preference of type P_1 and the second one is a positive preference of type P_2. Consider, for example, the (incomparable) pairs $(8, 0.1)$ and $(3, 0.8)$. Applying the aggregation operator $(+, \max)$ gives the pair $(11, 0.8)$ which, as expected, is better than both pairs, since $\max(8, 3, 11) = 11$ and $\max(0.1, 0.8, 0.8) = 0.8$.

5. Bipolar preference structures

Once we are given a positive and a negative preference structure, a naive way to combine them would be in performing the Cartesian product of the two structures. For example, if we have a positive structure P and a negative structure N, taking their Cartesian product would mean that, given a solution, it will be associated with a pair (p, n), where $p \in P$ is the overall positive preference and $n \in N$ is the overall negative preference. Such a pair is in general neither an element of P nor of N, so it is neither positive nor negative, unless one or both of p and n are the indifference element. Moreover, the ordering induced over these pairs is the well-known Pareto ordering, which may induce a lot of incomparability among the solutions. These two features imply that compensation is not allowed at all. Instead, we believe that it should be allowed, if desired. We will therefore now describe a way to combine these two preference structures that allows for it.

Definition 2: A bipolar preference structure is a tuple $(N, P, +, \times, \bot, \Box, \top)$, where

- $(N, +_{|_N}, \times_{|_N}, \bot, \Box)$ is a negative preference structure;
- $(P, +_{|_P}, \times_{|_P}, \Box, \top)$ is a positive preference structure;
- $+ : (N \cup P)^2 \rightarrow (N \cup P)$ is such that $a_n + a_p = a_p$ for any $a_n \in N$ and $a_p \in P$; this operator induces as partial ordering on $N \cup P$: $\forall a, b \in P \cup N$, $a \leq b$ iff $a + b = b$ and
- $\times : (N \cup P)^2 \rightarrow (N \cup P)$ is an operator (called the *compensation operator*) that, for all a, b, $c \in N \cup P$, satisfies the following properties:
  - commutativity: $a \times b = b \times a$;
  - monotonicity: if $a \leq b$, then $a \times c \leq b \times c$.

In the following, we will write $+_n$ instead of $+_{|_N}$ and $+_p$ instead of $+_{|_P}$. Similarly for $\times_n$ and $\times_p$. Moreover, we will sometimes write $\times_{np}$ when operator $\times$ will be applied to a pair in $(N \times P)$.

Notice that bipolar preference structures generalise both negative preference structures and positive preference structures via a bipolar structure. In fact, a negative preference structure is just a bipolar preference structure with a single positive preference: the indifference element, which, in such a case, is also the top element of the structure. Similarly, positive preference structures are just bipolar preference structures with a single negative preference: the indifference element. By symmetry, in such cases the indifference element coincides with the bottom element of the structure.

Given the ordering induced by $+$ on $N \cup P$, we have $\bot \leq \Box \leq \top$. Thus, there is a unique maximum element ($\top$) and a unique minimum element (i.e. $\bot$); the element $\Box$ is smaller than any positive preference and greater than any negative preference, and it is used to model indifference. The shape of a bipolar preference structure is shown in Figure 1. $\times_p$ and $+_p$ are the combination and the additive operators of the positive preference structure, while $\times_n$ and $+_n$ are the combination and the additive operators of the negative preference structure.

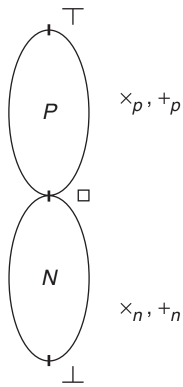

Figure 1. The shape of a bipolar preference structure.

Notice that, although all positive preferences are strictly above negative preferences, our framework does not prevent from using the same scale, or partially overlapping scales, to represent positive and negative preferences. In fact, if we have the same scale L for both the positive preferences and the negative ones, then we can map, via an isomorphism, the positive preference scale into another scale L' such that the bottom element of L' coincides with the top element of L. A similar procedure can be used also when one wishes to use partially overlapping scales. So our framework is not restrictive and it can also handle these cases.

A bipolar preference structure allows us to have different ways to model and reason about positive and negative preferences. In fact, we can have different lattices $(P, \leq_p)$ and $(N, \leq_n)$. This is common in real-life problems, where negative and positive statements are not necessarily expressed using the same granularity. For example, we could be satisfied with just two levels of negative preferences, while requiring several levels of positive preferences. Our framework allows us also to model cases in which the two structures are isomorphic (Section 9).

Notice that the combination of a positive and a negative preference is a preference which is higher than, or equal to, the negative one and lower than, or equal to, the positive one.

Theorem 2: *Given a bipolar preference structure* $(N, P, +, \times, \bot, \Box, \top)$, *for all* $p \in P$ *and* $n \in N$, $n \leq p \times n \leq p$.

Proof: For any $n \in N$ and $p \in P$, $\Box \leq p$ and $n \leq \Box$. By monotonicity of $\times$, we have: $n \times \Box \leq n \times p$ and $n \times p \leq \Box \times p$. Hence: $n = n \times \Box \leq n \times p \leq \Box \times p = p$.

This means that the compensation of positive and negative preferences lies in one of the chains between the two combined preferences. Notice that all such chains pass through the indifference element $\Box$. Possible choices for combining strictly positive with strictly negative preferences are thus the average or the median operator, or also the minimum or the maximum.

Moreover, by monotonicity, we can show that if $\top \times \bot = \bot$, then the result of the compensation between any positive preference and the bottom element is the bottom element. Also, if $\top \times \bot = \top$, then the compensation between any negative preference and the top element is the top element.

Theorem 3: *Given a bipolar preference structure* $(N, P, +, \times, \bot, \Box, \top)$, *we have that*:

- *if* $\top \times \bot = \bot$, *then* $\forall p \in P$, $p \times \bot = \bot$,
- *if* $\top \times \bot = \top$, *then* $\forall n \in N$, $n \times \top = \top$.

Proof: Assume $\top \times \bot = \bot$. Since for all $p \in P$, $p \leq \top$, then, by monotonicity of $\times$, $p \times \bot \leq \top \times \bot = \bot$, hence $p \times \bot = \bot$.

Assume $\top \times \bot = \top$. Since for all $n \in N$, $\bot \leq n$, then, by monotonicity of $\times$, $\top = \top \times \bot \leq \top \times n$, hence $\top \times n = \top$.

Example 7: In Table 1 we give three examples of bipolar preference structures. The structure described in the first row uses real numbers as positive and negative preferences. Compensation is obtained by summing the preferences, while the ordering is given by the max operator. In the second structure we have positive preferences between 0 and 1 and negative preferences between -1 and 0. Aggregation of positive preferences is max and of negative preferences is min, while compensation between positive and negative preferences is sum, and the order is given by max. In the third structure we use positive preferences between 1 and $+\infty$ and negative preferences between 0 and 1. Compensation is obtained by multiplying the preferences and ordering is again obtained via max. Notice that, if $\top \times \bot \in \{\top, \bot\}$, then compensation in the first and in the third structure is associative.

6. Bipolar preference problems

A bipolar constraint is just a constraint where each assignment of values to its variables is associated to one of the elements in a bipolar preference structure. A bipolar CSP (V, C) is then just a set of variables V and a set of bipolar constraints C over V. There could be many ways of defining the optimal solutions of a bipolar CSP. Here we propose one which avoids problems due to the possible non-associativity of the compensation operator, since compensation never involves more than two preference values.

Definition 3: Given a bipolar preference structure $(N, P, +, \times, \bot, \Box, \top)$, a solution of a bipolar CSP (V, C) over this structure is a complete assignment to all variables in V, say s, and an associated preference which is computed as follows: $\text{pref}(s) = (p_1 \times_p \cdots \times_p p_k) \times (n_1 \times_n \cdots \times_n n_l)$, where $k + l = |C|$, $p_i \in P$ for $i = 1, \ldots, k$ and $n_j \in N$ for $j = 1, \ldots, l$, and where $\exists \langle def_i, con_i \rangle \in C$ such that $p_i = def_i(s \downarrow_{con})$, $\exists \langle def_j, con_j \rangle \in C$ such that $n_j = def_j(s \downarrow_{con})$. *A* solution s is an optimal solution if there is no other solution s' with $pref(s') > pref(s)$.

In this definition, the preference of a solution s is obtained by combining all the positive preferences associated to its projections over the constraints via operator $\times_p$, by combining

Table 1. Examples of bipolar preference structures.

N, P	$+$	$\times$	$\bot, \Box, \top$
R^-, R^+	$+_n = +_p = +_{np} =$ max	$\times_n = \times_p = \times_{np} =$ sum	$-\infty, 0, +\infty$
$[-1, 0], [0, 1]$	$+_n = +_p = +_{np} =$ max	$\times_n =$ max, $\times_p =$ min, $\times_{np} =$ sum	$-1, 0, 1$
$[0, 1], [1, +\infty]$	$+_n = +_p = +_{np} =$ max	$\times_n = \times_p = \times_{np} =$ prod	$0, 1, +\infty$

all the negative preferences associated to its projections over the constraints via operator $\times_n$ and then by combining the two preferences obtained so far (one positive and one negative) via operator $\times_{np}$. Such a definition follows the same idea proposed in chapter IV of Bentham (1789), however many others can be defined according to several principles.

For example, we may avoid aggregation of positive and negative preferences via the two operators $\times_p$ and $\times_n$, since it may lead to poor discrimination among solutions. In fact, if we have a finite preference scale with few elements, then aggregating means obtaining one of the preferences in the scale to associate to a solution. Thus, if there are k elements in the scale, we can have at most k different evaluations for a solution. Thus, if the number of solutions is much higher than k, many solutions will end up with the same evaluation and will thus result to be indistinguishable. A possible solution to this problem is thus to avoid aggregation, and rather to maintain, for each solution, the tuple of all preferences given by the single constraints to the solution. In this way, if we have c constraints, then the number of different evaluations passes from k to k^c. Thus a greater discriminating power is achieved. Assuming no aggregation, a solution is associated to a tuple of positive preferences and a tuple of negative preferences. Different solutions can then be compared by ordering the elements of the two tuples (according to $+_p$ and $+_n$) for each solution, and then by comparing the ordered tuples by a lexicographic order.

Notice that in this article we have not followed this last approach, but we have applied preference aggregation in order to be able to exploit the machinery to solve soft constraint problems (Bistarelli et al. 1997; Bistarelli 2004; Rossi, Van Beek and Walsh 2006) (i.e. bipolar preference problems where only negative preferences are present), that just relays on preference aggregation.

Example 8: Consider the scenario in which we want to buy a car and we have preferences over some features. In terms of colour, we like red, we are indifferent to white and we hate black. Also, we like convertible cars a lot and we do not care much for big cars (for example, SUVs). In terms of engines, we like diesel. However, we do not want a diesel convertible.

We represent positive preferences via positive integers, negative preferences via negative integers and we maximise the sum of all kinds of preferences. This can be modelled by a bipolar preference structure where $N=[-\infty, 0]$, $P=[0, +\infty]$, $+=\max$, $\times=\text{sum}$, $\bot=-\infty$, $\Box=0$, $\top=+\infty$. Figure 2 shows the structure (variables, domains, constraints and preferences) of such a bipolar CSP, where preferences have been chosen consistently with the above informal specification, and 0 is used to model indifference (also when tuples are not shown). Consider solution $s_1=(C=\text{red}, T=\text{convertible}, E=\text{diesel})$. Its preference is $pref(s_1)=(def_1(\text{red})\times def_2(\text{convertible})\times def_3(\text{diesel}))\times def_4(\text{convertible}, \text{diesel})=(10+20+10)+(-20)=20$. We can see that the optimal solution is $(C=\text{red}, T=\text{convertible}, E=\text{gasoline})$ with global preference of 30.

Consider now a different bipolar preference structure, which differs from the previous one only for $\times_p$, which is now max. We want to see if, using a bipolar preference structure where the compensation operator is not associative as in the previous case and where the positive combination operator is idempotent, the optimal solution remains the same or it changes. Now solution s_1 has preference $pref(s_1)=(def_1(\text{red})\times def_2(\text{convertible})\times def_3(\text{diesel}))\times def_4(\text{convertible}, \text{diesel})=\max(10, 20, 10)+(-20)=0$. It is easy to see that now an optimal solution has preference 20. There are two of such solutions: one is the same as the above optimal solution and the other one is $(C=\text{white}, T=\text{convertible}, E=\text{gasoline})$.

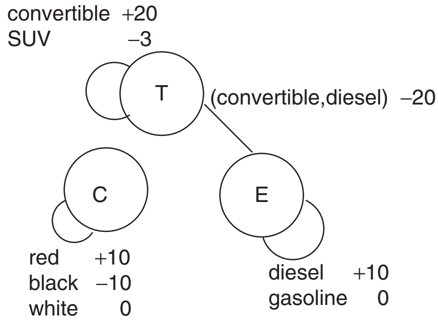

Figure 2. The bipolar CSP of the car's preferences.

The two cars have the same features except for the colour. A white convertible is just as good as a red convertible because we decided to aggregate positive preference by taking the maximum element rather than by summing them.

7. Strong bipolar preference structure

In some real-life problems there are situations where some strong negative statements are so negative that we would not like them to be compensated even by the best positive statements. For example, if we are allergic to the ingredients of a medication, then, even if it would solve some other health problem, we do not want to use it. Moreover, there are also statements that need to be expressed as hard constraints, which have to be satisfied for a scenario to be feasible. For example, if a classroom cannot fit more than 100 students, then, no matter the other features of the room, we cannot choose it for a class of 150 students.

In these situations we need to have not only positive and negative preferences, possibly compensating each other, but also negative preferences that cannot be compensated.

To provide a framework where such situations can be expressed, it is useful to consider an extension of the bipolar preference structure defined in Section 5, where it presents an additional set of preferences, which represents negative statements that cannot be compensated by any positive preference. Since such preferences are negative, it is reasonable that they are aggregated as it is usual for negative preferences, that is, like in a negative preference structure. However, the compensation operator is not defined on pairs including one of such strong negative preferences. We call this new structure a *strong bipolar preference structure*.

We build the strong bipolar structure by linking a bipolar preference structure with a negative preference structure, which models strong negative preferences, and by setting the highest strong negative preference to coincide with the lowest negative preference. In this structure the compensation operator is not defined on pairs including one of such strong negative preferences.

Definition 4 (strong bipolar preference structure): A strong bipolar preference structure is a tuple $\langle R, N, P, +, \times, \bot_r, \bot, \Box, \top \rangle$ where

- $\langle R, +_{|_R}, \times_{|_R}, \bot_r, \bot \rangle$ is a negative preference structure;

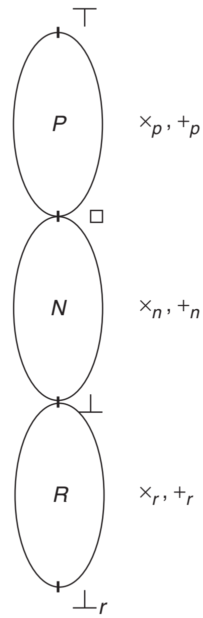

Figure 3. The shape of a strong bipolar preference structure.

- $\langle N, P, +_{|_{(P\cup N)}}, \times_{|_{(P\cup N)}}, \bot, \Box, \top\rangle$ is a bipolar preference structure;
- $+ : (N \cup P \cup R)^2 \to (N \cup P \cup R)$ is such that $a_r + a_{pn} = a_{pn}$ for any $a_{pn} \in P \cup N$ and $a_r \in R$; this operator induces as partial ordering on $R \cup N \cup P$: $\forall a, b \in P \cup N \cup R$, $a \leq b$ if and only if $a + b = b$;
- $\times : (N \cup P \cup R)^2 \to (N \cup P \cup R)$ is an operator that satisfies the following properties:
  - $a_{pn} \times r \leq r$, $\forall a_{pn} \in N \cup P$, $a_r \in R$;
  - commutativity: $a \times b = b \times a$, $\forall a, b, c \in N \cup P \cup R$;
  - monotonicity: if $a \leq b$, then $a \times c \leq b \times c$, $\forall a, b, c \in N \cup P \cup R$.

In the following, we will write $+_r$ instead of $+_{|_R}$ and $\times_r$ instead of $\times_{|_R}$. R is the set of negative preferences that cannot be compensated and we call them strong negative preferences. Every element of R is worse than any element of N (and thus also of P). As P and N have in common the indifference element $\Box$, N and R have in common element $\bot$, which is the worst negative preference that can participate in a compensation. In fact, the compensation operator, just as in a bipolar preference structure, is defined only on pairs involving positive and negative preferences but not pairs including strong negative preferences. The ordering among the elements of a strong bipolar preference structure, as well as the operators in each part of the structure, can be seen in Figure 3.

8. Positive *versus* negative preferences

Positive and negative preferences look so similar that, even though we know they need different combination operators, we could wonder why we need two different structures to handle them. Why cannot we just have one structure, for example the negative one, and transform each positive preference into a negative one? For example, if there are only two colours for cars, (i.e. red and blue), and we like only blue, instead of saying this using positive preferences (i.e. we like blue with a certain positive preference), we could phrase it using negative preferences (i.e. we do not like red with a certain negative preference). In other words, instead of associating a positive preference to blue and indifference to red, we could give a negative preference to red and indifference to blue.

In this section we will show that, by doing this, we could modify the solution ordering, thus representing a different optimisation problem. Thus the two preference structures are needed to model the problems under consideration: using just one of them would not suffice.

It is easy to show that, by moving from a positive to a negative modelling of the same information, as we have done in the example above, all solutions get a lower preference value. In fact, in this transformation, a positive preference is replaced by indifference, or indifference is replaced by a negative preference. So, in any case, some preference is replaced by a lower one, and by monotonicity of the aggregation operators ($\times_n$, $\times_p$ and $\times$), the overall preference of the solutions is lower as well.

However, it is worth noting that this preference lowering might not preserve the ordering among solutions. That is, solutions that were ordered in a certain way before the modification can be ordered in the opposite way after it. This is due to the fact that aggregation of positive and negative preferences may behave differently. The following example shows this.

Example 9: Consider the bipolar preference structure $(R^-, R^+, \max, \times, -\infty, 0, +\infty)$, where $\times$ is such that $\times_p = \times_{np} = \text{sum}$ and $\times_n = \min$. This means that we want to maximise the sum of positive preferences and to maximise the minimal negative preference (thus negative preferences are handled as fuzzy constraints). Also, positive preferences are between 0 and $+\infty$ and negative preferences are between 0 and $-\infty$. Compensation is via algebraic sum, thus values v and $-v$ are compensated completely (i.e. the result of the compensation is 0), while the compensation of values v and $-v'$ is $v - v'$.

Consider now a bipolar CSP over this structure with four variables, say X, Y, Z, W, where each variable has a Boolean domain as follows: $D(X) = \{a, \bar{a}\}$, $D(Y) = \{b, \bar{b}\}$, $D(Z) = \{c, \bar{c}\}$ and $D(W) = \{d, \bar{d}\}$. Assume now that the preference of a is 2, of b is 1, of c is 2.4 and of d is 0.5, while the preference of the other elements is indifference (i.e. 0 in this example). This means that we have expressed all our statements in a positive form.

Consider now two solutions s and s' as follows: $s = (a, b, \bar{c}, \bar{d})$ and $s' = (\bar{a}, \bar{b}, c, d)$. By computing the preference of s, we get $(2+1) + \min(0, 0) = 3$, while for s' we get $\min(0, 0) + (2.4 + 0.5) = 2.9$. Thus s is better than s'.

Assume now to express the same statements in negative terms that if we like an assignment t at level p, then we dislike $\bar{t}$ at the same level p. Hence, the preference of $\bar{a}$ is -2, of $\bar{b}$ is -1, of $\bar{c}$ is -2.4 and of $\bar{d}$ is -0.5, while the preference of the other elements is 0. Now the preference of s is $(0+0) + \min(-2.4, -0.5) = -2.4$, while the preference of s' is $\min(-2, -1) + (0+0) = -2$. Thus s' is better than s.

It is, however, possible to prove that, under some assumptions, moving from a positive to a negative modelling of the same information, or vice versa, does not change the ranking of the solutions.

Theorem 4: *Consider three bipolar CSPs Q, Q_1 and Q_2 over the bipolar preference structure $\langle N, P, +, \times, \bot, \Box, \top \rangle$, with variables with Boolean domains of the form $\{d, \bar{d}\}$. Assume that Q, Q_1 and Q_2 differ only in the preference values associated to domain values as follows*:

- *Q associates to all domain values $d_i = \bar{d}$ a negative preference n_i and to all domain values $d_i = d$ a positive preference $p_i = f(n_i)$, where f is bijection from N to P, which reverses the ordering, and which associates to every value $n_i \in N$ a value $f(n_i) = p_i$ in P;*
- *Q_1 is like Q, except that it associates the indifference element to all the domain values $\bar{d}$,* i.e. *Q_1 is the problem obtained from Q if we express all the things in positive terms;*
- *Q_2 is like Q, except that it associates the indifference element to all the domain values d,* i.e. *Q_2 is the problem obtained from Q if we express all the things in negative terms.*

Consider now two different solutions of Q_1 and Q_2, say $s = (d_1, \ldots, d_n)$ and $s' = (d'_1, \cdots, d'_n)$, and let us call $pref_1(s)$ (resp., $pref_1(s')$) and $pref_2(s)$ (resp., $pref_2(s')$) the preference of solution s (resp., s') in Q_1 and Q_2.

If

- *$\forall a, b \in N$, $f(a \times_n b) = f(a) \times_p f(b)$ and*
- *$\times_p$ and $\times_n$ are strictly monotonic,*

then $pref_1(s) > pref_1(s')$ if and only if $pref_2(s) > pref_2(s')$, i.e. *s is better than s' in the positive modelling iff s is better than s' in the negative modelling.*

Proof: Without loss of generality, we order the assignments of s and s' such that the first k ($k = 1, \ldots, n$) assignments are equal in both s and s'. Then,

- $pref_1(s) = \Pi^k_{i=0} p_i \times_p \Pi^n_{i=k+1} p_i$;
- $pref_1(s') = \Pi^k_{i=0} p'_i \times_p \Pi^n_{i=k+1} p'_i$;
- $pref_2(s) = \Pi^k_{i=0} n_i \times_n \Pi^n_{i=k+1} n_i$ and
- $pref_2(s') = \Pi^k_{i=0} n'_i \times_n \Pi^n_{i=k+1} n'_i$.

By construction, $\Pi^k_{i=0} p_i = \Pi^k_{i=0} p'_i$ and $\Pi^k_{i=0} n_i = \Pi^k_{i=0} n'_i$. We will denote the first value with K_p, and the second one with K_n.

Moreover, $\forall i = k+1, \ldots, n$, if d_i has preference n_i in Q_2, then it has preference $\Box$ in Q_1 and $p_i = f(n_i)$ in Q. Then, since d_i is different from d'_i and domains are Boolean, we have that d'_1 has preference $\Box$ in Q_2 and preference $p_i = f(n_i)$ in Q_1. Thus,

- $pref_1(s) = K_p \times_p \Pi^n_{i=k+1} f(n'_i)$ and
- $pre\, f_1(s') = K_p \times_p \Pi^n_{i=k+1} f(n_i)$.

We will show that, if $pref_2(s) > pref_2(s')$, then $pref_1(s) > pref_1(s')$. The other case can be shown similarly.

If $pref_2(s) > pref_2(s')$, then $K_n \times_n \Pi^n_{i=k+1} n_i > K_n \times_n \Pi^n_{i=k+1} n'_i$. This implies that $\Pi^n_{i=k+1} n_i > \Pi^n_{i=k+1} n'_i$. In fact, if we assume that $\Pi^n_{i=k+1} n_i \leq \Pi^n_{i=k+1} n'_i$ then, by monotonicity of the operator $\times_n$, we obtain $K_n \times_n \Pi^n_{i=k+1} n_i \leq K_n \times_n \Pi^n_{i=k+1} n_i$, which is a contradiction. Since f is an order reversing map, we have that $f(\Pi^n_{i=k+1} n_i) < f(\Pi^n_{i=k+1} n'_i)$. Since, by hypothesis, $f(a \times_n b) = f(a) \times_p f(b)$ $\forall a, b \in N$, then we have that $\Pi^n_{i=k+1} f(n_i) < \Pi^n_{i=k+1} f(n'_i)$. By strict monotonicity of the operator $\times_p$, and since K_p is a positive preference, then $K_p \times_p \Pi^n_{i=k+1} f(n_i) < K_p \times_p \Pi^n_{i=k+1} f(n'_i)$, i.e. $pref_1(s') < pref_1(s)$.

Notice that in Example 9 the ranking of the solutions changes when we move from the positive to the negative modelling, since the two assumptions of Theorem 4 do not hold. The first condition of Theorem 4 is not satisfied since, if we take two different negative preferences a and b which are not the indifference element, we have that $f(a \times_n b) = -\min(a, b) > -(a+b) = (-a) + (-b) = f(a) \times_p f(b)$, and so $f(a \times_n b) \neq f(a) \times_p f(b)$. Moreover, the second condition of Theorem 4 does not hold since $\times_n$, that is the minimum operator, is not strictly monotone.

9. Associativity of preference compensation

In many cases, the compensation operator $\times$ may be not associative. Here, we list some sufficient conditions for the non-associativity of the $\times$ operator.

Theorem 5: *Given a bipolar preference structure* $(P, N, +, \times, \bot, \Box, \top)$, *operator* $\times$ *is not associative if at least one of the following two conditions is satisfied*:

(1) $\top \times \bot = c \in (N \cup P) - \{\top, \bot\}$;
(2) $\exists p \in P - \{\top, \Box\}$ *and* $n \in N - \{\bot, \Box\}$ *such that* $p \times n = \Box$ *and at least one of the following conditions holds*:
   (a) $\times_p$ *or* $\times_n$ *is idempotent*;
   (b) $\exists p' \in P - \{p, \top\}$ *such that* $p' \times n = \Box$ *or* $\exists n' \in N - \{n, \bot\}$ *such that* $p \times n' = \Box$;
   (c) $\top \times \bot = \bot$ *and* $\exists n' \in N - \{\bot\}$ *such that* $n \times n' = \bot$;
   (d) $\top \times \bot = \top$ *and* $\exists p' \in P - \{\top\}$ *such that* $p \times p' = \top$;
   (e) $\exists a, c \in N \cup P$ *such that* $a \times p = c$ *iff* $c \times n \neq a$ (*or* $\exists a, c \in N \cup P$ *such that* $a \times n = c$ *iff* $c \times p \neq a$).

Proof:

(1) If $c \in P - \{\top\}$, then $\top \times (\top \times \bot) = \top \times c = \top$, while $(\top \times \top) \times \bot = \top \times \bot = c$. If $c \in N - \{\bot\}$, then $\bot \times (\bot \times \top) = \bot \times c = \bot$, while $(\bot \times \bot) \times \top = \bot \times \top = c$.
(2) Assume that $\exists p \in P - \{\top, \Box\}$ and $n \in N - \{\bot, \Box\}$ such that $p \times n = \Box$.
   (a) If $\times_p$ is idempotent, then $p \times (p \times n) = p \times \Box = p$, while $(p \times p) \times n = p \times n = \Box$. Similarly if $\times_n$ is idempotent.
   (b) If $\exists p' \in P - \{p, \top\}$ such that $p' \times n = \Box$, then $(p \times n) \times p' = p'$ while $p \times (n \times p') = p$. Analogously, if $\exists n' \in N - \{n, \bot\}$ such that $p \times n' = \Box$.
   (c) If $\top \times \bot = \bot$, then, by Theorem 3, $p \times \bot = \bot$. If $\exists n' \in N - \{\bot\}$ such that $n \times n' = \bot$, then $(p \times n) \times n' = \Box \times n' = n'$ while $p \times (n \times n') = p \times \bot = \bot \neq n'$.
   (d) If $\top \times \bot = \top$, then, by Theorem 3, $n \times \top = \top$. If $\exists p' \in P - \{\top\}$ such that $p \times p' = \top$, then $(n \times p) \times p' = \Box \times p' = p'$ while $n \times (p \times p') = n \times \top = \top \neq p'$.
   (e) If $c \times n \neq a$, then $(a \times p) \times n = c \times n \neq a$, but $a \times (p \times n) = a \times \Box = a$. Analogously if $c \times p \neq a$.

Notice that the above sufficient conditions refer to various aspects of a bipolar preference structure: the properties of its operators, the shape of P and N orderings, the relation between $\times$ and the other operators. Since some of these conditions often occur in practice, it is not reasonable to always require associativity of $\times$.

Sometimes, it could be desirable to be able to build bipolar preference structures where compensation is associative to have a less complex bipolar preference structure to handle.

Moreover, the associativity of the compensation operator could be useful to exploit solving techniques similar to the ones used for classical soft constraint problems, that are based on the associativity of the combination operator of the considered semiring.

It is obvious that, if we are free to choose any positive and any negative preference structure when building the bipolar framework, we will never be able to assure associativity of the compensation operator. Thus, if we want to assure this, we must pose some restrictions on the way a bipolar preference structure is built.

9.1. *One way to assure associativity of* $\times$

We now describe a procedure to build positive preferences as inverses of negative preferences, which assures that the resulting bipolar preference structure has an associative compensation operator. To do that, $\times_n$ must be non-idempotent. The methodology is called *localisation* and represents a standard systematic technique for adding multiplicative inverses to a (semi)ring (Bruns and Herzog 1998).

Given a (semi)ring with carrier set N (representing, in our context, a negative preference structure), and a subset $S \subseteq N$, we can construct another structure with carrier set P (representing, for us, a positive preference structure) and a mapping from N to P which makes all elements in the image of S invertible in P. The localisation of N by S is also denoted by $S^{-1}N$.

We can select any subset S of N. However, it is usual to select a subset S of N which is closed under $\times_n$, such that $\mathbf{1} \in S$ ($\mathbf{1}$ is the unit for $\times_n$, which represents indifference), and $\mathbf{0} \notin S$.

Given N and S, let us consider the quotient field of N w.r.t. S. This is denoted by $Quot(N, S)$, and will represent the carrier set of our bipolar structure. One can construct $Quot(N, S)$ by just taking the set of equivalence classes of pairs (n, d), where n and d are elements of N and S, respectively, and the equivalence relation is: $(n, d) \equiv (m, b) \Leftrightarrow n \times_n b = m \times_n d$. We can think of the class of (n, d) as the fraction $\frac{n}{d}$.

The embedding of N in $Quot(N, S)$ is given by the mapping $f(n) = (n, \mathbf{1})$, thus the (semi)ring N is a subring of $S^{-1}N$ via the identification $f(a) = \frac{a}{1}$.

The next step is to define the $+$ and $\times$ operator in $Quot(N, S)$, as function of the operators $+_n$ and $\times_n$ of N. We define $(n, d) + (m, b) = ((n \times_n b) +_n (m \times_n d),\ d \times_n b)$ and $(n, d) \times (m, b) = (m \times_n n, d \times_n b)$. By using the fraction representation we obtain the usual form where the addition and the multiplication of the formal fractions are defined according to the natural rules: $\frac{a}{s} + \frac{b}{t} = \frac{(a \times_n t) +_n (b \times_n s)}{s \times_n t}$ and $\frac{a}{s} \times \frac{b}{t} = \frac{a \times_n b}{s \times_n t}$.

It can be shown that the structure $(P, +_p, \times_p, \frac{1}{1}, \frac{1}{0})$, where $P = \{\frac{1}{a}$ such that $a \in (S \cup \{\mathbf{0}\})\}$, $+_p$ and $\times_p$ are the operators $+$ and $\times$ restricted over $\frac{1}{S} \times \frac{1}{S}$, $\frac{1}{1}$ is the bottom element in the induced order (notice that the element coincide with $\mathbf{1}$), and $\frac{1}{0}$ is the top element of the structure (this element is introduced *ad hoc* because $\mathbf{0}$ is not an unit and cannot be used to build its inverse) is a positive preference structure. Moreover, $Quot(N, S) = P \cup N$, and it is the carrier of a bipolar preference structure $\langle P, N, +, \times, 0, \frac{1}{1}, \frac{1}{0} \rangle$ where $\times$ is an associative compensation operator by construction.

Notice that the first example of Table 1 in Section 5, as well as the third example restricted to rational numbers, can be obtained via the localisation procedure.

10. Solving bipolar preference problems

Bipolar problems are NP-complete since they generalise both classical and soft constraints, which are already known to be difficult problems (Bistarelli et al. 1997; Bistarelli 2004). In this section we will consider how to adapt some usual techniques for solving soft constraints to bipolar problems.

10.1. *Branch and bound*

Preference problems based on c-semirings can be solved via a BB technique, possibly augmented via soft constraint propagation, which may lower the preferences and thus allow for the computation of better bounds (Bistarelli et al. 1997; Bistarelli 2004).

In bipolar CSPs, we have both positive and negative preferences. However, if the compensation operator is associative, standard BB can be applied. Thus bipolar preferences can be handled without additional effort.

However, if compensation is not associative, the upper bound computation has to be slightly changed to avoid performing compensation before all the positive and the negative preferences have been collected. More precisely, each node of the search tree is associated to a positive and a negative preference, say p and n, which are obtained by aggregating all preferences of the same type obtained in the instantiated part of the problem. Next, all the best preferences (which may be positive or negative) in the uninstantiated part of the problem are considered. By aggregating those of the same type, we get a positive and a negative preference, say p' and n' which can be combined with the ones associated to the current node. This produces the upper bound $ub = (p \times_p p') \times (n \times_n n')$, where $p' = p_1 \times_p \cdots \times_p p_w, n' = n_1 \times_n \cdots \times_n n_s$, with $w + s = r$, where r is the number of uninstantiated variables/constraints. Notice that ub is computed via $r - 1$ aggregation steps and one compensation step.

On the other hand, when compensation is associative, we do not need to postpone compensation until all constraints have been considered. Thus, ub can be computed as $ub = a_1 \times \cdots \times a_{p+r}$, where $a_i \in N \cup P$ is the best preference found in a constraint of either the instantiated part of the problem (first p elements) or the uninstantiated part of the problem (last r elements). Thus ub can be computed via at most $p + r - 1$ steps among which there can be many compensation steps. This is useful since the compensation steps can reduce the total number of steps. In fact, a compensation can generate the indifference element $\square$, which is the unit element for the compensation operator, and thus, when $\square$ is generated, the successive computation step can be avoided.

10.2. *Bipolar preference propagation*

When looking for an optimal solution in a soft constraint problem, BB can be helped by some form of partial or full constraint propagation. To see whether this can be done when solving bipolar problems as well, we must first understand what constraint propagation means in such problems. For sake of simplicity, we will focus here on arc-consistency (Bistarelli et al. 1997; Bistarelli 2004; Rossi et al. 2006).

Given any bipolar constraint, let us first define its negative version $neg(c)$, which is obtained by just replacing the positive preferences with indifference. Similarly, the positive version $pos(c)$ is obtained by replacing negative preferences with indifference.

A binary bipolar constraint c is then said negatively arc-consistent (NAC) iff $neg(c)$ is soft arc-consistent. More precisely,

Definition 5 (NAC): Consider a binary constraint that connects variables X and Y, let us call it c_{XY}, and let us call c_X the soft domain of X and c_Y the soft domain of Y. The constraint c_{XY} is NAC iff $neg(c_X) = (neg(c_X) \times_n neg(c_Y) \times_n neg(c_{XY})) \Downarrow_X$ and $neg(c_Y) = (neg(c_X) \times_n neg(c_Y) \times_n neg(c_{XY})) \Downarrow_Y$.

If a binary bipolar constraint is not NAC, we can make it NAC by modifying the soft domains of its two variables such that the two equations above hold. The modifications required can only decrease some preference values. Thus some negative preferences can become more negative than before. If operator $\times_n$ is idempotent, then such modifications generate a new constraint which is equivalent to the given one.

Let us now consider the positive version of a constraint. Let us also define an operation $\Uparrow_X$, which, taken any constraint c_S over variables S such that $X \in S$, computes a new constraint over X as follows: for every value a in the domain of X, its preference is computed by taking the greatest lower bound of all preferences given by c_S to tuples containing $X = a$. Then we say that a binary bipolar constraint is positively arc-consistent (PAC) iff the following holds.

Definition 6 (PAC): Consider a binary constraint that connects variables X and Y, let us call it c_{XY}, and let us call c_X the soft domain of X and c_Y the soft domain of Y. The constraint c_{XY} is PAC iff $pos(c_X) = (pos(c_X) \times_p pos(c_Y) \times_p pos(c_{XY})) \Uparrow_X$ and $pos(c_Y) = (pos(c_X) \times_p pos(c_Y) \times_p pos(c_{XY})) \Uparrow_Y$.

As in the NAC case, if a binary bipolar constraint is not PAC, we can make it PAC by modifying opportunely the soft domains of its two variables. However, while for NAC the needed modifications can decrease some preference values, for PAC these modifications can only increase such values. This is due to the different behaviour of the combination operator of the positive preferences. Moreover, as in the NAC case, if the combination operator (that is, $\times_p$ in the PAC case) is idempotent, then these modifications generate a new constraint which is equivalent to the given one.

Finally, we define when a binary bipolar constraint is bipolar arc-consistent (BAC).

Definition 7 (BAC): A binary bipolar constraint is BAC iff it is both NAC and PAC.

A bipolar constraint problem is BAC iff all its constraints are BAC.

If a bipolar constraint problem is not BAC, we can consider its negative and positive versions and achieve PAC and NAC on them. If both $\times_n$ and $\times_p$ are idempotent, this can be seen as the application of functions which are monotone, inflationary and idempotent on a suitable partial order. Thus usual algorithms based on chaotic iterations (Apt 2003) can be used, with the assurance of terminating and having a unique equivalent result which is independent of the order in which constraints are considered. However, this can generate two versions of the problem (of which one is NAC and the other one is PAC) which could be impossible to reconcile into a single bipolar problem. In fact, the two new problems could associate both a positive preference and a negative one to the same domain element. Consider, for example, a domain value that has a positive preference p in the given bipolar problem. Then, the positive version will consider p, and the negative version will consider $\square$. Then, the two PAC and NAC algorithms could transform, respectively, p into p' (higher than p) and $\square$ into n (smaller than $\square$). So this domain value will have both p' and n, while by definition each domain value has just one preference. Taking their compensation would return a problem which is not equivalent to the given one, since the compensation operator is not associative.

This problem can be solved by achieving only partial forms of PAC and NAC in a bipolar problem. The basic idea is to consider the given bipolar problem, apply the NAC and PAC algorithms to its negative and positive versions, and then modify the preferences of the original problem only when the two new versions can be reconciled that is, when at least one of the two new preferences is the indifference element. In fact, this means that, in one of the two consistency algorithms, no change has been made. If this holds, the other preference is used to modify the original one. This algorithm achieves a partial form of BAC, that we call p-BAC, and assures equivalence. More precisely, we give the following proposition.

Proposition 1: *Consider a bipolar preference problem Q, its negative version neg(Q), and its positive version pos(Q). Apply NAC to neg(Q) and PAC to pos(Q), and denote with NAC(neg(Q)) and PAC(pos(Q)) the new bipolar problems. Consider also the bipolar preference problem PBAC(Q) obtained from NAC(neg(Q)) and PAC(pos(Q)) as follows: every tuple associated with the indifference element in PAC (pos(Q)) (resp., NAC(neg(Q))) and a negative (resp., positive) preference n (resp., p) in NAC(neg(Q)) (resp., PAC(pos(Q))) has preference n (resp., p) in PBAC(Q), and every tuple associated with a positive preference p in PAC(pos(Q)) and a negative preference n in NAC(neg(Q)) has the same preference as in Q.*

Then, PBAC(Q) is equivalent to Q.

Proof: In order to show that PBAC(Q) is equivalent to Q, we have to show that for every solution of Q, the preference of this solution in PBAC(Q) coincides with the one in Q. Assume that s is a solution of Q with preference $p \times n$, where p is a positive preference, n is a negative preference and $\times$ is the compensation operator. Then, by construction of *pos*(Q) and *neg*(Q), s has preference p in *pos*(Q) and n in *neg*(Q). Moreover, since $\times_n$ is idempotent, NAC(*neg*(Q)) is equivalent to *neg*(Q). Thus also in NAC(*neg*(Q)) the preference of s is n. Similarly, since $\times_p$ is idempotent, PAC(*pos*(Q)) is equivalent to *pos*(Q) and so also in PAC(*pos*(Q)) the preference of s is p. By construction of PBAC(Q), we have that the preference of s in PBAC(Q) is $p \times n$.

Example 1: In Figure 4 it is shown how to make a bipolar constraint partially BAC. Part (*a*) shows a bipolar preference problem with three constraints c_X, c_Y and c_{XY}. Preferences are modelled by the bipolar preference structure ($N=[-1,0]$, $P=[0,1]$, $+=\max$, $\times$, $\bot=-1$, $\Box=0$, $\top=1$), where $\times$ is such that $\times_p=\max$, $\times_n=\min$ and $\times_{np}=\text{sum}$. Since preferences are given independently in c_X, c_Y and c_{XY}, it is possible to give a low positive preference for a value of X (e.g. $X=b$) in c_X, a negative preference for a value of Y (for example, $X=b$) in C_Y, but an high positive preference for the combination of such values in c_{XY}. In Part (*b*) we present the positive version of c_{XY}, that becomes PAC by increasing the positive preference associated to $X=b$ from +0.1 to +0.6. Part (*c*) presents the negative version of c_{XY}, that becomes NAC by decreasing the negative preference associated to $X=$a from −0.2 to −0.4. In Part (*d*) we show how to achieve p-BAC of c_{XY}: we must reconcile the modified preferences obtained in Part (*b*) and in Part (*c*) when it is possible. Since in this example it is always possible to reconcile such preferences, we obtain a bipolar constraint which is not only p-BAC, but also BAC.

In this approach we require idempotency of $\times_p$ and $\times_n$. However, when such operators are not idempotent, we can follow the approach used in the extended version of arc-consistency presented in Larrosa and Schiex (2003); Cooper and Schiex (2004) and Bistarelli and Gadducci (2006).

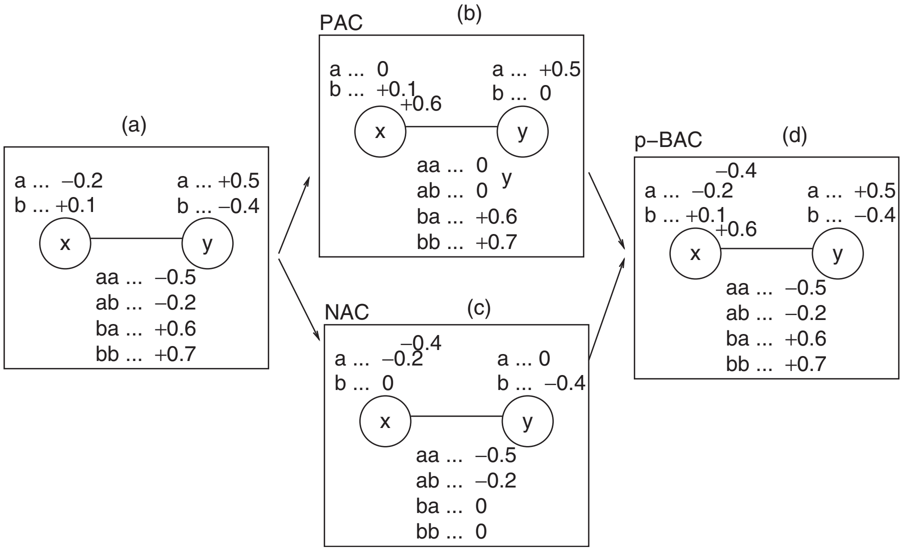

Figure 4. How to make a bipolar constraint p-BAC.

Notice that our algorithm will possibly decrease some negative preferences and increase some positive preferences. Therefore, if we use constraint propagation to improve the bounds in a BB algorithm, it will actually sometimes produce worse bounds, due to the increase of the positive preferences. We will thus use only the propagation of negative preferences (that is, NAC) within a BB algorithm. Since the upper bound is just a combination of several preferences, and since preference combination is monotonic, lower preferences give a lower, and thus better, upper bound.

Even if PAC is not useful in a BB procedure, it could be useful (alone or with NAC) when we do not want to follow a BB approach, but when we want to obtain preference values in the variables' domain that are closer to the ones that can be really obtained. When we have bipolar preference problems with only positive preferences, it is useful to perform only PAC, while when also negative preferences are present, it is useful to perform both PAC and NAC, and to reconcile their results if it is possible.

11. Related and future work

Bipolarity is an important topic in several fields, such as psychology (Osgood and Tannenbaum 1957; Tversky and Kahneman 1992; Cacioppo and Berntson 1997; Slovic, Finucane and Mag-Gregor 2002) and multi-criteria decision making (Grabisch and Labreuche 2005). Also, it has recently attracted interest in the AI community, especially in argumentation (Amgoud and Prade 2005) and qualitative reasoning (Benferhat et al. 2002, 2006; Dubois and Fargier 2005, 2006). These works consider how two alternatives should be compared, given for each a set of positive arguments and a set of negative ones, but they do not analyse the question of combinatorial choice.

Our bipolar framework is in particular related to the usual ways to represent a bipolar scale in psychology: a bipolar univariate method (Osgood and Tannenbaum 1957) and a unipolar bivariate method (Cacioppo and Berntson 1997). The first approach (Osgood and Tannenbaum 1957), which is the most similar to our approach, considers a scale with a central element ranging from negative elements, which are lower than the neutral element, to positive ones, which are higher than negative elements. The second approach (Cacioppo and Berntson 1997), instead, uses two independent unipolar scales for representing positive and negative aspects.

Although our approach is similar to the method in Osgood and Tannenbaum (1957), the scale we use for representing positive and negative preferences is not a unique scale with the neutral element in the exact middle. The positive and the negative scales may have a different granularity, since we want to let the user free to choose whatever structure he wants, if it satisfies the properties mentioned in the definition of a bipolar preference structure.

Another work regarding bipolarity handling in constraint-based reasoning is presented in Fargier and Wilson (2007). However, differently from our approach, it defines an algebraic structure to model bipolarity that follows the approach in Cacioppo and Berntson (1997) and not the one in Osgood and Tannenbaum (1957). The bipolar structure considered in Fargier andWilson (2007) allows a combination between positive and negative preferences that may produce a preference that is neither positive nor negative, that is not allowed in our approach. Moreover, the operator to combine positive and negative preferences is assumed always to be associative, while we do not force the user to choose only associative operators. However, if the user want to be sure to have a bipolar preference structure with an associative combination operator, in our article we have shown how to obtain it.

The approach in Cacioppo and Berntson (1997) allows one to express, for the same feature of an object, both a positive and a negative preference. Our current framework does not model this. However, tuples of preferences could be considered, and this would allow having more than one preference value for each feature.

Bipolar reasoning and preferences have been considered also in the context of qualitative reasoning. In Benferhat et al. (2002, 2006) a bipolar preference model based on a fuzzy-possibilistic approach is described where fuzzy preferences are considered and negative preferences are interpreted as violations of constraints. Precedence is given to negative preference optimisation and positive preferences are only used to distinguish among the optimals found in the first phase, thus not allowing for compensation.

Another work related to bipolar scales is Grabisch de Baets and Fodor (2003), which considers only totally ordered unipolar and bipolar preference scales. When the preferences are totally ordered, our operators $\times_n$ and $\times_p$ correspond respectively to the *t-norm* and *t-conorm* used in Grabisch et al. (2003). Moreover, in Grabisch et al. (2003) it is defined an operator, the *uninorm*, which can be seen as a restricted form of compensation and is forced to always be associative.

As future work, we plan to develop a solver for bipolar CSPs, which should be flexible enough to accommodate for both associative and non-associative compensation operators. The outlined algorithms for BB, NAC, PAC and p-BAC will also be implemented and tested over classes of bipolar problems. We also intend to consider the presence of uncertainty in bipolar problems, possibly using possibility theory to model such uncertainty, and to develop solving techniques for such scenarios. A first study in this direction has been presented in Bistarelli Pini, Rossi and Venable (2007b). Another line of future research is the generalisation of other preference formalisms, such as CP-nets, to deal with bipolar preferences.

Acknowledgements

This work has been supported by Italian MIUR PRIN project 'Constraints and Preferences' (No. 2005015491).

References

Amgoud, L., and Prade, H. (2005), 'An Argumentation-based Approach to Multiple Criteria Decision', *Proceedings of ECSQARU-05*, Berlin: Springer, pp. 10–17.

Apt, K.R. (2003), *Principles of Constraint Programming*, Cambridge: Cambridge University Press.

Bacchus, F., Grove, A.J., 'Utility Independence in a Qualitative Decision Theory', in *Proceedings of Kr-96*, pp. 542–552.

Benferhat, S., Dubois, D., Kaci, S., and Prade, H. (2002), 'Bipolar Representation and Fusion of Preferences in the Possibilistic Logic Framework', in *Proceedings of KR-02*, San Francisco: Morgan Kaufmann.

Benferhat, S., Dubois, D., Kaci, S., and Prade, H. (2006), 'Bipolar Possibility Theory in Preference Modelling: Representation, Fusion and Optimal Solutions', *Information Fusion: An International Journal on Multi-Sensor, Multi-Source Information Fusion*, 7, 135–150.

Benferhat, S., Dubois, D., and Prade, H. (2001), 'Towards a Possibilistic Logic Handling of Preferences', *Applied Intelligence*, 14, 303–317.

Bentham, J. (1789), *An Introduction to the Principles of Morals and Legislation*, New York: Oxford University Press.

Bistarelli, S. (2004), *Semirings for Soft Constraint Solving and Programming*, Vol. 2962 of LNCS, Heidelberg: Springer.

Bistarelli, S., and Gadducci, F. (2006), 'Enhancing Constraints Manipulation in Semiring-Based Formalisms', in *Proceedings of ECAI-06*, Amsterdam: IOS Press, pp. 63–67.

Bistarelli, S., Montanari, U., and Rossi, F. (1997), 'Semiring-based Constraint Solving and Optimization', *Journal of the ACM*, 44, 201–236.

Bistarelli, S., Pini, M.S., Rossi, F., and Venable, K.B. (2006), 'Bipolar Preference Problems', in *Proceedings of ECAI-06*, Amsterdam: IOS Press, pp. 705–706.

Bistarelli, S., Pini, M.S., Rossi, F., and Venable, K.B. (2007a), 'Bipolar Preference Problems: Framework, Properties and Solving Techniques', *Recent Advances in Constraints. Selected papers from 2006 ERCIM workshop on constraints* (Vol. 4651), Berlin: Springer LNAI.

Bistarelli, S., Pini, M.S., Rossi, F., and Venable, K.B. (2007b), 'Uncertainty in Bipolar Preference Problems', in *Proceedings of CP-07*, Vol. 4741 of LNCS, Berlin: Springer, pp. 782–789.

Boutilier, C., Brafman, R.I., Hoos, H.H., and Poole, D. (1999), 'Reasoning With Conditional Ceteris Paribus Preference Statements', in *Proceedings of UAI-99*, eds. K.B. Laskey and H. Prade, San Francisco: Morgan Kaufmann, pp. 71–80.

Bruns, W., and Herzog, J.C.-M. (1998), *Rings*, Cambridge: Cambridge University Press.

Cacioppo, J.T., and Berntson, G.G. (1999), 'The Affect System: Architecture and Operating Characteristics', *Current Directions in Psychological Science*, 8, 133–137.

Cacioppo, J.T., and Berntson, G.G. (1997), 'Beyond Bipolar Conceptualizations and Measures: the Case of Attitudes and Evaluative Space', *Personality and Social Psychology Review*, 1, 3–25.

Cooper, M., and Schiex, T. (2004), 'Arc Consistency for Soft Constraints', *AI Journal*, 154(1–2), 199–227.

Dechter, R. (2003), *Constraint Processing*, San Francisco: Morgan Kaufmann.

Doyle, J., Shoham, Y., and Wellman, M.P. (1991), 'A Logic of Relative Desire (Preliminary Report)', in *Proceedings of ISMIS-91*, Vol. 542 of LNCS, Berlin: Springer, pp. 16–31.

Dubois, D., and Fargier, H. (2005), 'On the Qualitative Comparison of Sets of Positive and Negative Affects', *Proceedings of ECSQARU-05*, Berlin: Springer, pp. 305–316.

Dubois, D., and Fargier, H. (2006), 'Qualitative Decision Making with Bipolar Information', *Proceedings of KR-06*, Menlo Park, CA: AAAI Press, pp. 175–186.

Ehrgott, M., and Gandibleux, X. (2002), *Multiple Criteria Optimization. State of the Art Annotated Bibliographic Surveys*, Dordrecht: Kluwer Academic.

Fargier, H., and Lang, J. (1993), 'Uncertainty in Constraint Satisfaction Problems: a Probalistic Approach', in *Proceedings of ECSQARU-93*, Vol. 747 of LNCS, Berlin: Springer, pp. 97–104.

Fargier, H., and Wilson, N. (2007), 'Algebraic Structures for Bipolar Constraint-Based Reasoning', *Proceedings of ECSQARU-07*, Berlin: Springer, pp. 623–634.

Grabisch, M., de Baets, B., and Fodor, J. (2004), 'The Quest for Rings on Bipolar Scales', *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 12, 449–512.

Grabisch, M., and Labreuche, C. (2005), 'Bi-capacities – Parts I and II', *Fuzzy Sets and Systems*, 151, 211–260.

Lang, J. (1996), 'Conditional Desires and Utilities: an Alternative Logical Approach to Qualitative Decision Theory', *Proceedings of ECAI-96*, Chichester, UK: John Wiley and Sons, pp. 318–322.

Larrosa, J., and Schiex, T. (2003), 'In the Quest of the Best Form of Local Consistency for Weighted CSP', in *Proceedings of IJCAI-03*, San Francisco: Morgan Kaufman, pp. 239–244.

Osgood, C.E., and Tannenbaum, P.H. (1957), *The Measurement of Meaning*, Chicago: University of Illinois Press.

Slovic, P., Finucane, E.P., and MagGregor, D.G. (2002), 'Rational Actors or Rational Fools? Implications of the Affect Heuristic for Behavioural Economics', *The Journal of Socio-Economics*, 31, 329–342.

Rossi, F., Van Beek, P., and Walsh, T. (2006), *Handbook of constraint programming*, New York: Elsevier.

Ruttkay, Z., (1994), 'Fuzzy Constraint Satisfaction', in *Proceedings of 1st IEEE Conference on Evolutionary Computing*, IEEE Press, Orlando, pp. 542–547.

Schiex, T. (1992), 'Possibilistic Constraint Satisfaction Problems or How to Handle Soft Constraints?', in *Proceedings of UAI-92*, San Francisco: Morgan Kaufmann, pp. 268–275.

Schiex, T., Fargier, H., and Verfaillie, G. (1995), 'Valued Constraint Satisfaction Problems: Hard and Easy Problems', in *Proceedings of IJCAI-95*, California: Morgan Kaufmann, pp. 631–639.

Tversky, A., and Kahneman, D. (1992), 'Advances in Prospect Theory: Cumulative Representation of Uncertainty', *Journal of Risk and Uncertainty*, 5, 297–323.

Wellman, M.P., and Doyle, J. (1991), 'Preferential Semantics for Goals', *Proceedings of AAAI*, Menlo Park, CA: AAAI Press, pp. 698–703.