Bipolar Representations in Reasoning, Knowledge Extraction and Decision Processes

Didier Dubois and Henri Prade

IRIT, 118 route de Narbonne
31062 Toulouse Cedex, France
{dubois, prade}@irit.fr

Abstract. This paper surveys various areas in information engineering where an explicit handling of positive and negative sides of information is appropriate. Three forms of bipolarity are laid bare. They can be instrumental in logical representations of incompleteness, rule representation and extraction, argumentation, and decision analysis.

1 Introduction

Bipolarity refers to the propensity of the human mind to reason and make decisions on the basis of positive and negative affects. It expresses the fact that beyond ranking pieces of information or acts in terms of plausibility, utility and so on, the human mind also relies on absolute landmarks with positive and negative flavor, plus a third landmark expressing neutrality or indifference, corresponding to the boundary between positive and negative zones. For instance people make choices by checking the good sides and the bad sides of alternatives separately. Then they choose according to whether the good or the bad sides are stronger. Results in cognitive psychology have pointed out the importance of bipolar reasoning in human cognitive activities [19] [6] [22]. It even seems that positive and negative affects are not processed in the same area of the brain.

The presence of absolute landmarks in the way humans apprehend information creates limitations in some well-established theories of knowledge representation and reasoning. For instance, probability theory handles certainty and impossibility in a very rigid manner, leaving no room for the state of ignorance. Classical (Von-Neumann-Savage) utility theory builds interval scales for utilities regardless of positive and negative values, since a utility function is invariant with respect to increasing affine transformations. More generally ranking alternatives in a purely ordinal way cannot account for bipolarity in a straightforward manner. In decision theory, the first formal account of bipolarity is Cumulative Prospect Theory [23]. In quite a different matter, the fact that material implication does not provide a good model of if-then rules can be explained in terms of neglecting the bipolar nature of such rules, which have both examples and counter-examples.

The aim of this paper is to briefly survey some areas where bipolarity seems to be present and play a major role. The first section lays bare three forms of bipolarity. The subsequent sections are devoted to various cognitive tasks that naturally involve bipolar ingredients: uncertainty representations, conjointly exploiting knowledge and data, learning, expressing conditional information, and finally decision-making.

2 A Typology of Bipolarity

There are several forms of bipolarity according to the strength of the link between the positive and the negative aspects; in the most constrained form, the positive is just the mirror image of the negative and they are mutually exclusive. A looser form of bipolarity considers a possible coexistence between positive and negative evaluations, while a duality relation between them is maintained. In the loosest form, the positive and the negative sides express pieces of information of a different nature.

2.1 Bipolar scales

A bipolar scale $(L, >)$ is a totally ordered set with a prescribed interior element $\mathbf{0}$ called *neutral*, separating the positive evaluations $\lambda > \mathbf{0}$ from the negative ones $\lambda < \mathbf{0}$. Mathematically, if the scale is equipped with a binary operation $\star$ (an aggregation operator), $\mathbf{0}$ is an idempotent element for $\star$, possibly acting as an identity.

Examples :

- The most obvious quantitative bipolar scale is the (completed) real line equipped with the standard addition, where 0 is the neutral level. Isomorphic to it is the unit interval equipped with an associative uninorm like $\frac{xy}{xy+(1-x)(1-y)}$. Then the neutral point is 0.5, 0 plays the same role as $-\infty$ and 1 as $+\infty$ in the real line. Also the interval $[-1, 1]$ is often used as a bipolar scale;
- The simplest qualitative bipolar scale contains three elements: $\{-, \mathbf{0}, +\}$.

In such a bipolar scale, the negative side of the scale is the inverse mirror of the positive one. An object is evaluated on such a bipolar scale as being either positive or negative or neutral. It cannot be positive and negative at the same time. This is called a *univariate bipolar* framework.

Another type of bipolar framework uses two distinct totally ordered scales L^+ and L^- for separately evaluating the positive and the negative information. This is the *bivariate unipolar* framework. Here each scale is unipolar in the sense that the neutral level is at one end of the scale. In a *positive* scale the bottom element is neutral. In a *negative* scale the top element is neutral. A bipolar scale can be viewed as the union of a positive and a negative scale $L^+ \cup L^-$ extending the ordering relations on each scale so $\forall \lambda^+ \in L^+, \lambda^- \in L^-, \lambda^+ > \lambda^-$. The symmetrisation of finite unipolar scales is incompatible with associative operations [14] : only infinite bipolar scales seem to support such operations!

2.2 Symmetric, Dual and Loose Variants of Bipolarity

Three forms of bipolarity can be found at work in the literature, we call types I, II, III for simplicity.

- **Type I : Symmetric bipolarity**. It relies on the use of bipolar scales. Generally, positive and negative evaluations are comparable and sometimes can even add. Of course, the two truth-values *true* and *false* of classical logic offer a basic view of bipolarity. However, the neutral value only appears in three-valued logics. Note that the neutral truth value must be interpreted as *half-true*, and not as modalities such as *unknown* or *possible*. Probability theory exhibits a type I bipolarity as the probability of an event is clearly living on a bipolar scale $[0, 1]$ whose top means *totally sure* and bottom *impossible* (not to be confused with *true* and *false*). The neutral value is 0.5 and refers to the total uncertainty about whether an event or its contrary occurs (not to be confused with *half-true*). In decision theory, utility theory does not exhibit bipolarity as utility functions only encode an ordering relation between decisions. On the contrary, Tverski-Kahneman's Cumulative Prospect Theory uses the real line as a bipolar scale. It is numerical, additive, and bipolar. It measures the importance of positive affects and negative affects *separately*, by two monotonic set functions σ^+, σ^- and finally computes a net predisposition $N = \sigma^+ - \sigma^-$.
- **Type II : Homogeneous bivariate bipolarity.** It works with two separate positive and negative scales related via a duality relation. Here, an item is judged according to two independent evaluations : a positive one (in favor of the item), a negative one (in disfavor of the item). However positive and negative strengths are computed similarly on the basis of the same data. The point is here that the positive and the negative sides do not exhaust all possibilities. Part of the data may neither favor nor disfavor the evaluated item. Well-known examples of such a bipolarity can be found in formal frameworks for argumentation where reasons for asserting a proposition and reasons for refuting it are collected. In decision theory, one may compare decisions using pairs of positive and negative evaluations according to several criteria.
  Apart from the positive evaluation, a weak positive evaluation, gathering data not in disfavor of the item can be used. For instance, working with intervals on a (type I) bipolar univariate scale (in the case of an ill-known evaluation) comes down to a type II bipolarity. There is a duality relation relating the weak evaluation and the strong positive evaluation, if each item has a "contrary" : the weak evaluation of an item is the complement of the positive evaluation of the "contrary" item when the latter makes sense. This is typical of uncertainty theories leaving room for incomplete information. Namely, the confidence in some event A is evaluated by two set functions $C(A)$ and $\Pi(A)$ reflecting their certainty and plausibility respectively. They are related by the inequality $C(A) \leq \Pi(A)$, so that the certainty of A is expressed by $C(A) = \Pi(A) = 1$, the impossibility of A by $C(A) = \Pi(A) = 0$ while the neutral state of ignorance is when $C(A) = 0; \Pi(A) = 1$. Clearly, $C(A)$ lives on a positive scale, while $\Pi(A)$ lives on a negative one. The duality

relation expresses that $C(A) = 1 - \Pi(A^c)$ where A^c is the complement of A. A good example of certainty/plausibility pairs displaying this kind of bipolarity are belief and plausibility functions of Shafer. In the case of possibility/necessity measures, type II bipolarity is also present in the sense that necessity degrees live on a positive scale while possibility degrees live on a negative scale. However, the two scales are tightly related by the constraint stating that positive necessity degree implies a maximal possibility degree for a given event. So in this case the bivariate setting is degenerated and the pair $(C(A), \Pi(A))$ can be mapped in a one-to-one way to a symmetric type I bipolar scale.

– **Type III : Heterogeneous bipolarity.** In this form of bipolarity, the negative part of the information does not refer to the same kind of source as as the positive part. So positive and negative information are of a different nature, while in type II bipolarity only the polarity is different. Especially, in the case of information merging, negative and positive pieces of information will not be aggregated using the same principles. The positive side is not a mirror image of the negative side either. Nevertheless, positive and negative information cannot be completely unrelated. They must obey minimal consistency requirements. In uncertainty modeling or knowledge representation heterogeneous bipolarity corresponds to the pair (knowledge, data). Knowledge is negative information in the sense that it expresses constraints on how the world behaves, by ruling out impossible or unlikely relations: laws of physics, common sense background knowledge (claims like "birds fly"). On the contrary, data represent positive information because it represents examples, actual observations on the world. A not yet observed event is not judged impossible; observing it is a positive token of support. Accumulating negative information leads to ruling out more possible states of the world (the more constraints, the less possible worlds). Accumulating positive information enlarges the set of possibilities as being guaranteed by empirical observation. In decision making, heterogeneous bipolarity concerns the opposition between constraints (possibly flexible ones) that state which solutions to a problem are unfeasible, and goals or criteria, that state which solutions are preferred.

3 Bipolarity in Logical Representations of Belief

As said above, bipolarity appears in logic in two forms, one pertaining to the truth or the falsity of propositions, and the other pertaining to a (sincere) agent's capability to assert a proposition or its contrary. It is important to notice the existence of two scales: one that measures truth, one that measures belief. A truth- scale is type I bipolar and, when many-valued, it enables propositional variables and propositions whose truth is a matter of degree to be modelled. The neutral point in the scale is *half-true*. Working with Boolean or non-Boolean propositions is a matter of modelling convention, not a matter of how much knowledge is available. So $[0, 1]$-valued membership functions of fuzzy sets are type I bipolar (non-membership 0 being negative, 0.5 being neutral). In the Boolean case, the truth-scale is reduced to the pair $\{0, 1\}$.

Another issue is the belief scale. It is positive unipolar in the sense that while believing p is a positive piece of information, not believing p is non-committal, because it differs from believing $\neg p$, the negation of p. There is a companion negative unipolar plausibility scale whose bottom expresses impossibility and whose top has a neutral value for expressing non-committal statements of the form p is possible. In classical logic, beliefs are represented by propositions assumed true and forming a belief base K. Belief is Boolean : either p is believed (when $K \vdash p$) or not. Moreover p is believed if and only if $\neg p$ is impossible, indicating that this is type II bipolarity. Clearly in the case of incomplete belief bases, the epistemic state of a proposition is ternary in classical logic even if truth is 2-valued: one may either believe p, believe $\neg p$, or believe neither due to ignorance.

There are temptations to use belief states or belief values as truth values (a set of the form {*True, Unknown, False*}) and build a 3-valued logic on it. This is basically what the so-called "partial logic" [5] does. Its truth-tables use an implicit order whereby *Unknown* is less true than *True*, more true than *False*. But this approach runs into paradoxes related to the excluded-middle law [11]. Adopting truth-tables for conjunction and disjunction, one must assign a truth-value to $p \vee q$ when p and q are both unknown, which clearly depends on whether p and q are logically independent or not. The point is that ultimately, in the Boolean framework p is true or false, so that $p \vee \neg p$ must be a tautology, even if the truth-value of p is not known. So *Unknown* is not a truth-value in the usual sense: it does not prevent 0 and 1 from being exhaustive and mutually exclusive as truth-values. *Unknown* lives on the belief /plausibility bivariate scale. Just like *Unknown*, *True* and *False*, understood as above, are not truth-values, they are epistemic states because they stand for *certainly* 1, and *certainly* 0, respectively. They can be modelled as disjunctive subsets of the truth scale: *Unknown* $= \{0, 1\}$, *True* $= \{1\}$, *False* $= \{0\}$. Belnap so-called "4-valued logic" [3] supposedly adds a fourth "truth-value" expressing the contradiction to {*True, Unknown, False*}. However it is subject to the same criticism as above, as to what this 4-valued logic means, regardless of the fact that a multivalued logic based on such a kind of truth-set (a bilattice) can be devised and enjoys nice properties.

One reason for this confusion between truth and certainty of truth is that the language of classical logic does not support the expression of unknown propositions: only believed propositions can be written in the knowledge base. It becomes clearer when prefixing each believed proposition in K with a necessity-like belief modality C. Then a possibility-like modality Π, such that Πp may stand for $\neg C \neg p$. It can be shown that the proper logic here is the KD45 modal logic. Then *True* can be interpreted, in some sense, as a truth-value of Cp, not of p. *Unknown* is encoded as $\Pi p \wedge \Pi \neg p$. It applies when $Cp \vee C\neg p$ is false, as clearly Cp is not the negation of $C\neg p$. So, the presence of the epistemic state *Unknown* does not question the excluded middle law at all. Casting propositional logic into an epistemic modal logic lays bare the type II bipolarity of reasoning in classical logic. In fact it can be proved [8] that, denoting $CK = \{Cp, p \in K\}$,

$K \vdash p$ in classical logic if and only if $CK \vdash Cp$ in KD45. Note that this kind of embedding is not the usual one of propositional logic into modal logic: it says that the fragment of KD45 made of classical propositions prefixed by C behaves like classical logic, which justifies the name "belief base" for a set of classical propositions.

Rough set theory [20] also displays a form of type II (symbolic) homogeneous bipolarity, since a set is approximated by a pair of subsets, respectively containing elements surely belonging to it (lower approximation), and elements surely not belonging to it. The so-called upper approximation of the set is again the complement of the lower approximation of its complement. This can be represented using ordered pairs of truth-values from $\{0, 1\}$, viewed as an elementary unipolar scale, assigning $(1, 1)$ to elements surely belonging to A, $(0, 0)$ to elements surely belonging to A^c, and $(0, 1)$ to elements whose membership is unknown. However, it does not lead to a truth-functional three-valued logic on a (type I) bipolar scale, since the lower (resp. upper) approximation of a union (resp. intersection) of sets is not the union (resp. intersection) of their lower (resp. upper) approximations. Yet, links between three-valued logics and rough sets have been explored in the literature (e.g. Banerjee [2]).

4 Heterogeneous Bipolar Information: Knowledge vs. Data, and Learning

In the previous section, bipolarity in knowledge representation was due to incomplete information. There is a very different kind of bipolarity, this time heterogeneous, opposing background knowledge and empirical data. Background knowledge takes the form of generic statements, integrity constraints, laws, necessary conditions, and point out what cannot be possibly observed. On the contrary, data is made of observed cases that are positive pieces of information. Beware that positive knowledge may not just mirror what is not impossible. Indeed what is not impossible, not forbidden, does not coincide with what is explicitly possible or permitted. So, a situation that is not impossible (i.e., possible) is not necessarily guaranteed possible (i.e., positive) if it is not explicitly permitted, observed or given as an example.

Possibility theory is a suitable framework for modelling and reasoning about this kind of bipolar information[12][8]. Negative and positive information is represented by two separate possibility distributions, denoted by π and δ, yielding possibility and guaranteed possibility measures respectively. A possibility distribution π encodes a total pre-order on a set S of interpretations or possible states. It associates to each interpretation s a real number $\pi(s) \in [0, 1]$, which represents the compatibility of the interpretation s with the available knowledge on the real world (in case of uncertain knowledge), or equivalently to what extent s is not impossible. The less $\pi(s)$, the more impossible s is. The second possibility distribution δ should be understood differently. The degree $\delta(s) \in [0, 1]$ estimates to what extent the presence of s is supported by evidence, and $\delta(s) = 0$ just means that s has not been observed yet. In the crisp case, the set I of impossible situations is $I = \{s \in S, \pi(s) = 0\}$, and the set GP of guaranteed possible situations is $GP = \{s \in S, \delta(s) = 1\}$.

A characteristic property of heterogeneous bipolarity is the fact that the sets of guaranteed possible (the support GP of δ) and impossible (I) situations should be disjoint and generally do not cover all the referential. This is expressed by the coherence condition $GP \subseteq I^c$. This condition means that what is guaranteed possible should be not impossible. When uncertainty is graded, this coherence condition now reads: $\delta \leq \pi$.

Example : Assume for instance one has some information about the opening hours and prices of a museum M. We may know that museum M is open from 2 pm to 4 pm, and certainly closed at night (from 9 pm to 9 am). Note that nothing forbids museum M to be open in the morning although there is no positive evidence supporting it. Its ticket fare is neither less than 2 euros nor more than 8 euros (following legal regulations), prices between 4 and 5 euros are guaranteed to be possible (they are prices actually proposed by the museum).

Since observations accumulate, while increasing background knowledge eliminate new possible worlds, positive information aggregate disjunctively, and negative information aggregate conjunctively. This can be understood in our setting in the following way. A constraint like *the value of X is restricted by A_i* is encoded by a possibility distribution π s. t. $\pi \leq \mu_{A_i}$. Several such constraints are thus equivalent to $\pi \leq \min_i \mu_{A_i}$. By the principle of minimal commitment (anything not declared impossible is possible), it leads to choose the greatest possibility distribution $\pi = \min_i \mu_{A_i}$ compatible with the constraints. Hence a conjunctive combination. In the case of positive information X *is* A_i is equivalent to $\delta \geq \mu_{A_i}$, since it reflects empirical support. Then several such observations are equivalent to $\delta \geq \max_i \mu_{A_i}$. By closed world assumption (anything not observed as actually possible is not considered), one gets $\delta = \max_i \mu_{A_i}$. Hence a disjunctive combination.

Given a pair of possibility distributions (π, δ), we can define: the possibility degree of an event A, $\Pi(A) = \max\{\pi(s) : s \in A\}$, the dual necessity degree $N(A) = 1 - \Pi(A^c)$ and the guaranteed possibility degree $\Delta(A) = \min\{\delta(s) : s \in A\}$ (let alone the dual degree of potential necessity $1 - \Delta(A^c)$). Note that set function Π underlies an existential quantifier since $\Pi(A)$ is high as soon as some $s \in A$ is plausible enough. It agrees with the negative nature of information, since A is impossible, i. e. $\Pi(A) = 0 \iff N(A^c) = 1$, corresponds to the non-existence of an interpretation $s \in A$ having a non-zero degree of possibility $\pi(s)$. In contrast, Δ underlies a universal quantifier since $\Delta(A)$ is high as soon as all $s \in A$ be supported by evidence. It agrees with the positive nature of information encoded by δ, since $\Delta(A) = 1$ requires that all states where A occurs be maximally supported by evidence. The duality between N and Δ ($\Delta(A) = N^c(A^c)$ where N^c is the necessity measure based on $\pi^c = 1 - \delta$) is different from the one (characteristic of type II bipolarity) between N and Π.

Merging bipolar information [12], by disjunctive (resp. conjunctive) combination of positive (resp. negative) information, may create inconsistency when the upper and lower possibility distributions, which represent the negative part and the positive part of the information respectively, fail to satisfy the consistency condition $\pi \geq \delta$. Then, since empirical observations are generally regarded as more solid information than prior knowledge, the latter must be revised for instance as $\pi' = \max(\pi, \delta)$, so as to account for unexpected evidence.

Learning processes turn data into knowledge, hence positive into negative information in the sense of type III bipolarity: situations that are often observed are eventually considered as normal and those never observed are considered as impossible. Recently [21], it has been shown that Mitchell's version space concept learning, based on an explicit set of examples and counterexamples, can be reformulated in the language of possibility theory under heterogeneous bipolarity. Distributions π (induced by counterexamples) and δ (induced by examples) respectively become the most general and the most specific hypotheses explaining the data. The theory also explains how these hypotheses are progressed as new data come in.

5 Bipolarity and If-then Rules

An if-then rule is not a two-valued entity, it is a three valued one. To see it, consider a database containing descriptions of items in a set S. If a rule *if A then B* is to be evaluated in the face of this database, it clearly creates a 3-partition of S, namely:

1. the set of examples of the rule: $A \cap B$,
2. its set of counter-examples: $A \cap B^c$,
3. the set of irrelevant items for the rule: A^c.

Each situation should be encoded by means of a different truth-value. This view of a rule is at odds with the logical tradition, for which it is a material implication. The two first situations corresponding to the usual truth-values 1 (true) and 0 (false) respectively. The third case corresponds to a third truth-value that must be be interpreted as *irrelevant* as the rule does not apply. This idea of a rule as a *tri-event* actually goes back to De Finetti in the 1930's. This framework for modelling a rule produces a precise bipolar mathematical model: a rule is modeled as a pair of disjoint sets representing the examples and the counter-examples of a rule, namely $(A \cap B, A \cap B^c)$.

This definition has several consequences. First, it justifies the claim made by De Finetti that a conditional probability $P(B \mid A)$ is the probability of a particular entity denoted by $B \mid A$ that can be called *a conditional event*. Indeed it is obvious to see that the probability $P(B \mid A)$ is entirely defined by $P(A \cap B)$ and $P(A \cap B^c)$. Moreover it precisely shows that material implication only partially captures the intended meaning of an if-then rule. It is obvious that the set of items where the material implication $A^c \cup B$ is true is the complement of the set of counter-examples of a rule. Hence the usual logical view only emphasizes the negative side of the rule. It does not single out its examples. This is clearly in agreement with the fact that propositions in classical logic represent negative information. On the other hand, the set of examples of a rule is $A \cap B$ and clearly represents positive information. Thus, the three-valued representation of an if-then rule also strongly suggests that a rule contains both positive and negative information. Note that in data mining, the merit of an association rule $A \Rightarrow B$ extracted from a database is evaluated by two indices: the support and the confidence degrees, respectively corresponding to the probability $P(A \cap B)$ and the conditional probability $P(B \mid A) = \frac{P(A \cap B)}{P(A \cap B) + P(A \cap B^c)}$. This proposal may sound ad hoc. However the deep reason why two indices are necessary to evaluate the quality of a rule is because the rule generates a 3-partition of the database, and two evaluations are needed to picture their relative importance. In fact the primitive quality indices of an association rule are the proportion of its examples and the proportion of its counter-examples. All other indices derive from these basic evaluations.

It is intuitively satisfying to consider that a rule R1 = "if A then B" entails a rule R2 = "if C then D", if R2 has more examples and less counterexamples than R1 (in the sense of inclusion). R2 is safer than R1. This entailment relation (denoted $\models$) can be formally written as

$$B \mid A \models D \mid C \text{ if and only if } A \cap B \subseteq C \cap D \text{ and } C \cap D^c \subseteq A \cap B^c.$$

It is non-monotonic. Indeed, it has been shown [10] that the three-valued semantics of rules provide a representation for the calculus of conditional assertions of Kraus, Lehmann and Magidor [17], which is the main principled approach to nonmonotonic reasoning.

Lastly, the bipolar view has been also applied to fuzzy rules "'if A then B" (when A and/or B are fuzzy sets). It is clear that the usual modeling of fuzzy rules in fuzzy control, based on the fuzzy conjunction of A and B corresponds to the positive information contained in rules, while the less usual approach based on many-valued implications views rules as constraints and better fits classical logic. The bipolar view can be exploited for building a typology of fuzzy if-then rules, based on multivalued implications or conjunctions, where each type of fuzzy rules serves a specific purpose [13]. It emphasizes the advantages of using conjointly implicative rules (encoding negative information) and conjunctive rules (encoding positive information) in the same rule-based system. Finally the bipolar view is instrumental in rigourously extending the support and the confidence degrees to fuzzy association rules [9].

6 Bipolarity and Decision

Decision processes are pervaded with bipolar notions. All types of bipolarity are involved. Type I bipolar decision-making stems from evaluating decision on a bipolar scale, thus providing an explicit account of whether a decision is good or bad. An automatic procedure ranking decisions from the best to the worst does not prevent the best ranked decision from being bad (the other ones being worse), nor, for another case, the worst decision from still being reasonably good. It is useful to propose absolute evaluations, at least a qualitative advice about what is good and what is bad. Using a bipolar scale is clearly instrumental, due to the presence of the neutral point separating good grades from bad ones. This type of bipolarity is especially used in Cumulative Prospect Theory and more recently by Grabisch and Labreuche [18].

Type II bipolarity occurs when faced with several criteria, and evaluating separately the criteria where the ratings of a decision are good and the criteria where the ratings are bad. Each criterion can be evaluated on a type I bipolar scale and the global evaluation on a bivariate unipolar scale, hence pairs (how good, how bad) of evaluations are compared. Or, in a more complex and more expressive setting, each criterion can be itself rated on a bivariate unipolar scale, as done by Greco et al. [16].

In the bipolar setting the importance of criteria cannot be assessed as usual using set functions g like capacities, $g(C)$ evaluating the importance of the group C of criteria. So-called bicapacities [15] are of the form $g(C^+, C^-)$ where C^+ (resp. C^-) is a set of criteria where the decision performance is good (resp. bad). If criteria are rated individually on a bipolar scale, $C^+ \cap C^- = \emptyset$. The overall evaluation is performed using a variant of Choquet integral adapted to bicapacities. In the more expressive model, criteria importance is evaluated by so-called bipolar capacities[16]. The idea is to use two measures, a measure of positiveness (that increases with the addition of positive arguments and the deletion of negative arguments) and a measure of negativeness (that increases with the addition of negative arguments and the deletion of positive arguments), without combining them.

A purely ordinal setting for bipolar decision-making was recently proposed by Dubois and Fargier [7]. Each criterion is rated on the basic qualitative bipolar scale $\{-, \mathbf{0}, +\}$. The set $\mathcal{C}$ of criteria is mapped on a unipolar positive scale, for instance $[0, 1]$, where 0 indicates no importance. Let $\pi(c)$ be the importance of criterion c. The weight of a subset C of criteria is supposed to be $\Pi(A)$, using a possibility measure; the idea is to focus on the most important affect when making a choice. For a decision a, the evaluation of criterion c is either positive or negative or zero. Let $A^+ = \{c, c(a) = +\}$, and $A^- = \{c, c(a) = -\}$ be the positive reasons for a and the negative reasons against a, respectively. Comparing decisions a and b in the type II bipolar framework is based on evaluations $\Pi(A^-)$, $\Pi(A^+)$, $\Pi(B^-)$, and $\Pi(B^+)$. Several decision rules can be proposed. The first one is a Pareto-based comparison of pairs $(N((A^-)^c), \Pi(A^+))$ and $(N((B^-)^c), \Pi(B^+))$. It is a transitive partial ordering. It is perhaps too partial: for instance, when $\Pi(A^-) > \Pi(A^+)$, it concludes that a is incomparable with b where $B^+ = B^- = \emptyset$. In this case, one would rather say that a is worse than an indifferent b. Another drawback is observed when $\Pi(A^+) > \Pi(B^+)$ and $\Pi(A^-) = \Pi(B^-)$: this enforces preference of a over b, even if $\Pi(A^+)$ is very weak w.r.t the order of magnitude of the negative arguments — in the latter case, a rational decider would examine the negative arguments in details before concluding.

The other decision rule is a complete preorder that assumes commensurability between positive and negative evaluations, counting a reason against b as a reason for a:

$$a \succeq^{Biposs} b \Leftrightarrow \max(\Pi(A^+), \Pi(B^-)) \geq \max(\Pi(B^+), \Pi(A^-))$$

Only the strict part of the generated ordering is transitive. This rule focuses on the most salient affects pertaining to a and b. $a \succeq^{Pareto} b$ implies $a \succeq^{Biposs} b$. It is also clear that $\succeq^{Biposs}$ is a bipolar generalisation of a possibility measure. However, $\succeq^{Biposs}$ and Pareto are very rough rules that may be not decisive enough. Lexicographic refinements of $\succeq^{Biposs}$ offer more decisive and actually realistic decision rules. One such rule checks how many reasons for a and for b there are at each importance level and decides on the basis of the most important such discriminating level. It can be simulated by Cumulative Prospect Theory. This kind of qualitative bipolar setting can be useful in formal argumentation for the evaluation and comparison of arguments[1].

Quite another form of bipolarity in decision refers to the (in some sense more classical) opposition between constraints and goals. It is a form of heterogeneous bipolarity. A decision problem on a solution space is then modelled by two possibility-like distributions π and δ [4]. However, now $\pi(s)$ evaluates to what extent a solution s is feasible, not rejected. It is a matter of degree in the face of soft constraints. On the contrary δ is an objective function and $\delta(s)$ evaluates the extent to which s is fully satisfactory. All formal considerations pertaining to type III bipolarity apply here, especially the consistency condition between π and δ. Distribution π is generally the conjunctive aggregation of local soft constraints. Distribution δ is generally the disjunctive or additive aggregation of several objective functions. This approach can be expressed in possibilistic logic using a constraint base (containing negative information as in classical logic) and a goal base (containing positive information and behaving like in a data-driven logic [8]). Several strategies for defining best solutions can be devised. The most natural scheme is to first check consistency between constraints and goals, possibly modifying goals if necessary, then define a set of feasible solutions that achieves a compromise between soft constraints, and finally finding the best feasible solutions according to δ inside this set.

7 Conclusion

This paper suggests that bipolarity is naturally present in cognitive and decision processes. Bipolarity lays bare the presence of absolute landmarks in evaluation scales, having positive or negative flavor, thus revealing a cognitive limitation of purely ordinal representations. Modelling bipolarity in an explicit manner is useful in many areas of information engineering such as knowledge representation, learning, decision analysis, inconsistency handling, argumentation, question-answering systems.

References

1. L. Amgoud and H. Prade. Comparing decisions on the basis of a bipolar typology of arguments . In A. Hunter J. Dix, editor, *Proc. 11th Workshop on Nonmonotonic Reasoning*, pages 426–432, Windermere U.K., 2006.

2. M. Banerjee. Rough sets and 3-valued Lukasiewicz logic. *Fundamenta Informaticae*, 32:213–220, 1997.
3. N. D. Belnap. A useful four-valued logic. In J.M. Dunn and G. Epstein, editors, *Modern Uses of Multiple-Valued Logic*, pages 8–37. D.Reidel, Dordrecht, The Netherlands, 1977.
4. S. Benferhat, D. Dubois, S. Kaci, and H. Prade. Bipolar possibility theory in preference modeling: Representation, fusion and optimal solutions. *Information Fusion*, 7:135–150, 2006.
5. S. Blamey. Partial logic. In D. Gabbay and F. Guentner, editors, *Handbook of Philosophical Logic, 2d Ed.*, volume 5, pages 261–353. Kluwer Academic Publ., Dordrecht, The Netherlands, 1998.
6. J. T. Cacioppo, W. L. Gardner, and G. G. Berntson. Beyond bipolar conceptualizations and measures: The case of attitudes and evaluative space. *Personality and Social Psychology Review*, 1:3– 25, 1997.
7. D. Dubois and H. Fargier. Qualitative decision making with bipolar information. In *Proc. of Int. Conf. on Principles of Knowledge Representation and Reasoning KR'06, Windermere, UK*, pages 175–186, Menlo Park, Ca., 2006. AAAI Press.
8. D. Dubois, P. Hajek, and H. Prade. Knowledge-driven versus data-driven logics. *J. of Logic, Language, and Information*, 9:65–89, 2000.
9. D. Dubois, E. Huellermeier, and H. Prade. A systematic approach to the assessment of fuzzy association rules. *Data Mining and Knowledge Discovery, to appear*, 2006.
10. D. Dubois and H. Prade. Conditional objects as non-monotonic consequence relationships. *IEEE Trans. on Systems, Man and Cybernetics*, 24:1724–1739, 1994.
11. D. Dubois and H. Prade. Possibility theory, probability theory and multiple-valued logics: A clarification. *Ann. Math. and Artificial Intelligence*, 32:35–66, 2001.
12. D. Dubois, H. Prade, and P. Smets. "Not impossible" vs. "guaranteed possible" in fusion and revision. In S. Benferhat and Besnard P., editors, *Proceedings of ECSQARU'01 - LNAI 2143*, pages 522–531, 2001.
13. D. Dubois, H. Prade, and L. Ughetto. A new perspective on reasoning with fuzzy rules. *Int. J. of Intelligent Systems*, 18:541–567, 2003.
14. M. Grabisch. The Moebius transform on symmetric ordered structures and its application to capacities on finite sets. *Discrete Math.*, 28(1-3):17–34, 2004.
15. M. Grabisch and Ch. Labreuche. Bi-capacities — parts I and II. *Fuzzy Sets and Systems*, 151(2):211–260, 2005.
16. S. Greco, B. Matarazzo, and R. Slowinski. Bipolar Sugeno and Choquet integrals. In *EUROFUSE Workshop on Information Systems*, pages 191–196, Varenna, 2002.
17. S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1-2):167–207, 1990.
18. Ch. Labreuche and M. Grabisch. Generalized Choquet-like aggregation functions for handling bipolar scales. *Eur. J. of Operational Research*, 172(3):931–955, 2006.
19. C. E. Osgood, G.J. Suci, and P. H. Tannenbaum. *The Measurement of Meaning*. Univ. of Illinois Press, Chicago, 1957.
20. Z. Pawlak. *Rough Sets - Theoretical Aspects of Reasoning about Data*. Kluwer Academic Publ., Dordrecht, The Netherlands, 1991.
21. H. Prade and M. Serrurier. Version space learning for possibilistic hypotheses. In *Proc. of 17th Europ. Conf. on Art. Intel., ECAI'06*, Riva del Garda, Italy, 2006.
22. P. Slovic, M. Finucane, E. Peters, and D.G. MacGregor. Rational actors or rational fools? implications of the affect heuristic for behavioral economics. *The Journal of Socio-Economics*, 31:329–342, 2002.
23. A. Tversky and D. Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5:297–323, 1992.