Logical representation of preferences for group decision making

Céline Lafage **Jérôme Lang**

IRIT - Université Paul Sabatier

118, route de Narbonne
31062 Toulouse Cedex (France)
{lafage,lang}@irit.fr

Abstract

We propose two logical approaches for representing preferences in a group decision making context. The first one is based on weighted logics: each agent expresses her preferences by means of logical formulas weighted by importance degrees, and preferred decisions are those that satisfy best the collectivity of agents (according to some criteria). Aggregation is made in two steps: first the utility of each agent is computed from her individual preferences (independently from the other agents), and then individual utilities are aggregated into a collective utility function. We study the particular case corresponding to the pair of aggregation functions (sum,max) and show that some equity problems arise when the preferences expressed by the agents are not constrained enough. The second approach makes use of distances. We show that connections exist between both approaches.

1 Introduction

Group decision

Group decision making mainly consists of aggregating preferences of several agents into a collective preference ordering. The collective decision will then be a decision that is maximal with respect to the collective preference ordering (see, e.g., [20] or [19]). It is important to point out that the group eventually takes a unique, *collective decision* (contrariwise to game theory where each agent of a group acts separately). Group decision theory is composed of several classes of problems, well studied from a normative point of view in the field of economy. In a first class of problems, individual preferences of the agents are expressed quantitatively by means of *utility functions*; the collective preference ordering can then be represented by a collective utility function. In a second class of problems, individual preferences are expressed in an ordinal way by means of a *preference relation* over the set of possible solutions (or candidates); under this assumption, a *vote rule* maps a collection of individual preference relations into a single selected candidate and a *preference aggregation method* maps a collection of individual preference relations into a collective preference relation. Arrow's impossibility theorem is defined in the latter context: it tells that if an aggregation procedure satisfies a given series of rational postulates then the aggregation must be dictatorial, i.e., the collective preference relation must be equal to the preference relation of one of the agents.

Logical representation of preference

From the point of view of representation and computation, usual social choice methods (aggregation functions, voting rules etc.) can be *directly* applicable only when the respective numbers of agents (voters) and solutions (candidates) are reasonably low with respect to the available computational resources. This is always true when agents and candidates are individuals. Now, more generally, one may reasonably continue assuming that agents are still individuals, but we cannot any longer continue assuming that the number of candidates (or solutions) is not prohibitive. This happens in particular when the decision to take consists in giving a value to each of the variables of a set of n decision variables: here, the set of candidates is equal to the set of feasible assignments and thus grows exponentially with the number n of propositional variables. Therefore, it is obviously unreasonable to require from the agents the specification of an explicit utility function or preference ordering (under the form of a table or a list) on the set of all solutions. This argues for the need of a *compact* (or *succinct*) representation language for preferences. Furthermore, such a language for preference representation should be as *expressive* as possible, and as close as possible to the intuition (ideally, it should be easily translated from a specification in natural language of the preferences of an agent). Lastly, it should be equipped with decision procedures, as efficient as possible, so as to enable the automation of the search for an optimal collective decision.

The KR community has contributed a lot to the study of such succint and expressive languages for representing *knowledge*[1]. Now, for the last few years, there has been several approaches dedicated to the representation of *preferences* of an agent, to be used in a decision making process. It is sometimes the case that a logic initially developed for reasoning about knowledge can be used efficiently for representing and handling preferences. This is for instance the case for default logic which has been used for decision-theoretic issues [24][5] or for conditional logic [4]. More basically, the same remark applies to standard propositional logic, which can be seen as a (very) rough tool for representing preferences: the agent expresses her desires by specifying some logical formulas; feasible decisions are worlds, and optimal decisions consist of the worlds satisfying all formulas (worlds that violate any of the formulas being equally bad). This principle can be refined in the two following ways:

1. by associating to each formula a positive weight, namely, the penalty for not satisfying this formula; then, the penalty (or disutility) of a world being the sum of the weights of the formulas it violates [22, 15, 9]. If the penalties of the violated formulas are aggregated with *max* instead of sum, then we obtain a framework equivalent (up to a change of the valuation scale) to possibilistic logic [7]. We call *weighted logics* these logics built on this principle (propositional formulas + weights to be aggregated with a given function).

2. by grading the violation of a formula using a *distance* between worlds; given a formula φ representing a specific desire of the agent, and given a distance d, then the closer a world w to a model of φ, the better w (the optimal case being when w satisfies φ).

Outline of the paper

In this paper we develop the two above-mentioned approaches. For each of both, we start with the represention of individual preferences, and then we consider the case of a group of agents. We devote Section 2 to weighted logics for group decision making; in particular, we prove that given a set of reasonable postulates on how a utility function should be derived from a collection of individual preferences, the aggregation should be done in two successive steps (first at the individual level and then at the group level), thus leading to two (not necessarily identical) aggregation functions which respectively satisfy a number of specific properties. We study in further detail the case where the pair of aggregation functions is $\langle +, \max\rangle$, point on specific problems concerning equity between agents, and propose some solutions. In Section 3 we study distance-based logics for preference representation and we point out connections and intertranslations with weighted logics. Lastly, we position our work with respect to related work.

2 Weighted logics

Let $\mathcal{L}$ be a propositional language built from a finite set of variables and the usual connectives. Ω the set of interpretations (or worlds) associated with $\mathcal{L}$. Interpretations are denoted w, w_1 etc. formulas of $\mathcal{L}$ by Greek letters φ, ψ etc., and real numbers by x, y etc. $Mod(\varphi)$ is the set of all models of φ.

2.1 Individual preferences

Preferences are encoded in a compact and structured way under the form of elementary preference items; each of these items is a pair (formula, valuation) where the formula expresses a constraint the agent wishes to see satisfied and the valuation expresses the importance of this constraint. We choose $[0, +\infty]$ as valuation scale, once noted that this choice does not induce a loss of generality. Any closed infinite real interval could have been chosen instead, for instance $[0, 1]$ or $[-\infty, +\infty]$ since all infinite closed real intervals are isomorphic (as it will be shown further, this choice of $[0, +\infty]$ is merely motivated by considerations of intuition). The valuation $+\infty$ associated to a formula means that the constraint corresponding to this formula is absolute, while the valuation 0 expresses pure indifference towards the satisfaction of the constraint.

These weighted formulas defined above are concerned with the relative preference on some worlds over other ones, not with the feasibility of worlds (an impossible world should not be confused with an unacceptable world). To express feasibility, we introduce (integrity) *constraints.*

Definition 1 (individual preference profile) *An* individual preference profile *is a pair* $\langle P, K\rangle$ *where*

1. P *(*preference base*) is a multiset*[2] $\{(\varphi^j, \alpha^j)\}_j$ *with*

[1] Here we reserve the term "knowledge" for its restrictive meaning "information about the real world", excluding thus preferences, goals and desires.

[2] A *multiset* (or *bag*) is a set whose elements can appear

$\varphi^j \in \mathcal{L},\ \alpha^j \in [0, +\infty]$;

2. *K is a consistent set of formulas (*integrity constraints*).*

A world w is *possible* iff $w \models K$. Thus, the set of possible worlds is $Mod(K)$. In the rest of the paper we make the assumption that the decisions available to the collectivity of agents are precisely the possible worlds.This assumption is consistent with Boutilier's qualitative decision theory [4] when all atoms are controllable. Preference between worlds will be defined between possible worlds only. Namely, from P and K we build a *disutility* function, from $Mod(K)$ to $[0, +\infty]$, computed by aggregating the penalties of the constraints violated by a world:

Definition 2 (individual disutility)
The individual disutility function $disu_P$ associated to the individual preference profile $\langle K, P\rangle$ is the mapping from $Mod(K)$ to $[0, +\infty]$ defined by[3]
$\forall w \in Mod(K) \quad disu_P(w) = *\{\alpha^i \mid w \models \neg\varphi^i\}$
where $*$ is an operation from $[0, +\infty] \times [0, +\infty]$ to $[0, +\infty]$ satisfying commutativity, non-decreasingness, associativity and $\forall a,\ a * 0 = a$.

This disutility function, expressing the preferences of an agent under a numerical form, is thus computed by aggregating the valuations of all *violated* constraints. The justification of this computation scheme and of the four listed properties is discussed in Section 2.3. Noticeably, they imply $a * +\infty = +\infty$: indeed, since $*$ is non-decreasing, $a * 0 = a$ implies $a * b \geq \max(a, b)$ for all a, b, from which it follows that $+\infty$ is absorbent for $*$. If $disu(w) = +\infty$ then the possible world w is unacceptable by the agent.

Before justifying axiomatically the four listed properties, we comment them from an informal point of view. Commutativity is needed because the order in which the weighted formulas appear in P should be irrelevant. The need for non-decreasingness is obvious. 0 being neutral element of $*$ expresses that formulas with a null weight are irrelevant to the disutility function, and $+\infty$ being absorbent expresses that nothing can compensate the violation of an absolute constraint (whatever the other satisfied constraints, this world will have a maximal disutility). Associativity is needed by the requirement that the way the disutility of an agent is updated when taking a new constraint into account does not depend on where her initial disutility came from (this is the Decomposition postulate, cf. Section 2.3.). Noticeably, up to a change of scale from $[0, 1]$ to $[0, +\infty]$ and an order reversal, $*$ enjoys the same properties as *triangular norms*[4].

The most natural choices for $*$ are $+$ and max[5]. The sum is an *utilitarianist* aggregation mode, allowing for compensations between violated formulas (as in penalty calculus [22]) while the maximum is an *egalitarianist* aggregation mode, where the disutility degree is computed relatively to the most important of the violated formulas (as in possibilistic logic [7]).

A disutility function induces in the usual way a weak total order, i.e., a transitive, reflexive and total relation, on the set of possibles worlds: a possible world w_1 is preferred to the possible world w_2 (for a given agent) iff $disu_P(w_1) \leq disu_P(w_2)$.

2.2 Collective preference aggregation

Definition 3
A multi-agent preference base (MAPB) is a collection $(P_1, ..., P_n)$ of individual preference bases where $P_i = \{(\varphi_i^j, \alpha_i^j)\}_j$. A multi-agent preference profile consists of a MAPB and a set of integrity constraints K.

The collective disutility function is computed in two steps: first, the individual disutility functions are computed independently (as above), and then these functions are aggregated into a collective disutility function that best conveys the preferences of the group (see figure 2.2).

Definition 4 (collective disutility)
The collective disutility function of a MAPB P is defined as:
$\forall w \models K \quad disu_P(w) = \diamond\{disu_i(w) \mid i = 1, ..., n\}$
where

- *$disu_i$ is the disutility function of agent i computed from P_i.*
- *$\diamond$ is a function from $[0, +\infty]^n$ to $[0, +\infty]$ satisfying non-decreasingness in each of its arguments and commutativity.*

ω^* *is* optimal *w.r.t. P iff* $disu_P(w^*) = \min_{w \in \Omega} disu_P(w)$.

---

several times.

[3] We should in principle mention K as well, i.e., we should note $disu_{K,P}$ but since when it is not ambiguous we drop K for making notations lighter.

[4] A triangular norm is a mapping $T : [0, 1]\times[0, 1] \to [0, 1]$ satisfying commutativity, monotonicity, associativity and with 1 as neutral element. A common bijective transformation f from $[0, +\infty]$ to $[0, 1]$ is $f(x) = e^{-x}$; note that the ordering is then reversed ($a \leq b$ iff $f(a) \geq f(b)$, $f(0) = 1$ and $f(+\infty) = 0$).

[5] On $[0, 1]$, using the transformation given in the previous footnote, they correspond respectively to the product and the minimum. The nilpotent triangular norm $x, y \mapsto \max(o, x + y - 1)$ corresponds on $[0, +\infty]$ to an "infinipotent" function.

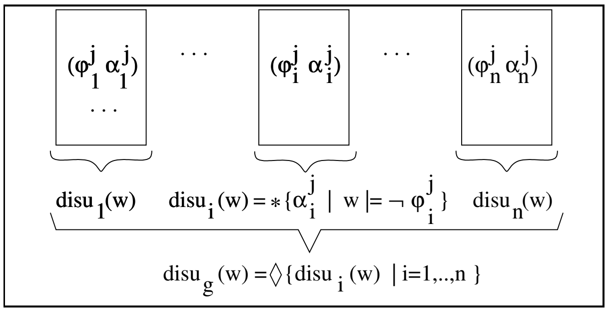

Figure 1: Preferences bases

The only two properties required for $\Diamond$ are non-decreasingness in each of its arguments and commutativity. Commutativity guarantees *anonymity* between agents because a permutation of preference bases of agents do not change global disutility. Non-decreasingness is required for obvious reasons. Among the possible $\Diamond$ operators, we again find max and $+$ that still correspond to egalitarianism and utilitarianism (in a more usual meaning than for $*$ – because these notions primarily come from group decision theory). min would be theoretically acceptable, but it would not be intuitively satisfactory since it compares decisions with respect to the *most* satisfied agent. An interesting class of acceptable operators consists of the "Ordered Weighted Average" operators (OWA) [27]. An OWA is an application $F : \mathrm{IR}^n \rightarrow \mathrm{IR}$ associated to a vector of weights $\vec{x} = [x_1, x_2, ..., x_n]$ with $x_i \in [0,1]$ and $\sum_{i=1}^{n} x_i = 1$, such that $F(a_1, a_2, ..., a_n) = \sum_{i=1}^{n} x_i b_i$ where $(b_1, ..., b_n)$ is a reordering of $(a_1, ..., a_n)$ such that $b_1 \geq b_2 ... \geq b_n$. Distinguished OWAs are min, max and the mean, associated respectively to the vectors $(0, ..., 0, 1)$, $(1, 0, ..., 0)$ and $(\frac{1}{n}, ..., \frac{1}{n})$. More refined OWAs express a compromise between min and max to a given degree. Note that OWAs are generally not associative.

Therefore, a weighted logic will be associated to a pair of operators $\langle *, \Diamond \rangle$. Choosing $*$ is independent from choosing $\Diamond$. Noticeably, if $*$ and $\Diamond$ coincide then preferences of agents 1, ..., n are gathered as if there were only one agent, which comes down to a mono-agent weighted logic (thus, $\langle +, + \rangle$ corresponds to penalty logic). Now, when $*$ and $\Diamond$ do not coincide, some associations are more satisfactory than others from the point of view of intuition. In particular, egalitarianism is intuitively more relevant between agents than between the constraints expressed by a single agent, which makes the choice $\langle +, \max \rangle$ a priori more natural than the choice $\langle \max, + \rangle$. The former ($\langle +, \max \rangle$) will be studied in more detail in Section 3.

2.3 Axiomatic justification

Let $P = \{P_1, ..., P_n\}$ be a MAPB and K a set of integrity constraints. We introduce the following notations:

- $Add(P, i, (\varphi\ \alpha))$ is the MAPB obtained from P by replacing P_i by $P_i \cup \{(\varphi\ \alpha)$;
- $Delete(P, i, (\varphi\ \alpha))$ is obtained from P by deleting $(\varphi\ \alpha)$ from P_i (and if $(\varphi\ \alpha)\ \ /inP$ then $Delete(P, i, (\varphi\ \alpha)) = P$).
- $Replace(P, i, (\varphi\ \alpha), (\varphi'\ \alpha'))$ is obtained from P by substituting $(\varphi\ \alpha)$ by $(\varphi'\ \alpha')$ in P_i;
- $Neutralize(P, i) = \{P_1, ..., P_{i-1}, \emptyset, P_{i+1}, ..., P_n\}$ is obtained from P by forgetting the preferences of agent i.

Now, we define the following postulates. The first one is distinguished from the following ones because it pertains to the integrity constraints.

Insensitivity towards integrity constraints:
Let K_1, K_2 two sets of integrity constraints; then for any $w \models K_1 \wedge K_2$, we have
$disu_{\langle P, K_1 \rangle}(w) = disu_{\langle P, K_2 \rangle}(w)$. (IC)

Thus, when (IC) is satisfied we can merely note $disu_P(w)$ instead of $disu_{\langle P, \emptyset \rangle}(w)$. We implicitely assume (IC) when writing the following axioms, thus enabling the fixing of K and using the above simplified notation.

Anonymity:
Let $P_\sigma = (P_{\sigma(1)}, ..., P_{\sigma(n)})$ where σ is a permutation of $\{1, ..., n\}$; then
$\forall w \models K \quad disu_P(w) = disu_{P_\sigma}(w)$. (A)

Neutrality:
Let $w, w' \in Mod(K)$ be two possible worlds such that for any weighted formula $(\varphi_i^j\ \alpha_i^j)$ appearing in one of the P_i's, the following equivalence holds: $(w \models \varphi_i^j) \Leftrightarrow (w' \models \varphi_i^j)$. Then $disu_P(w) = disu_P(w')$. (N)

Irrelevance of the syntax:
Let $P' = Replace(P, i, (\varphi_i^j\ \alpha_i^j), (\varphi_i'^j\ \alpha_i^j))$, with $\varphi_i^j \wedge K \equiv \varphi_i'^j \wedge K$; then
$\forall w \models K,\ disu_P(w) = disu_{P'}(w)$. (IS)

Insensitivity to satisfied formulas:
Let $P' = Replace(P, i, (\varphi\ \alpha), (\varphi\ \beta))$; then
$\forall w \models K \wedge \varphi,\ disu_{P'}(w) = disu_P(w)$. (ISF)

Monotonicity:
Let $P' = Replace(P, i, (\varphi\ \alpha), (\varphi\ \beta))$ with $\beta > \alpha$; $\forall w \models K \wedge \neg\varphi,\ disu_{P'}(w) \geq disu_P(w)$. (MON)

Strict monotonicity:
as above, replacing $disu_{P'}(w) \geq disu_P(w)$ by $disu_{P'}(w) > disu_P(w)$. (MONS)

Insensitivity to constraint ordering:
If P' is obtained from P by exchanging the ordering of the constraints $(\varphi_i^j\ \alpha_i^j)$ in one of the P_i then $disu_P = disu_{P'}$. (ICO)

Priority to hard constraints:
If one of the P_i contains $(\varphi_i^j\ +\infty)$ then $\forall w \models K \wedge \neg\varphi_i^j$ we have $disu_P(w) = +\infty$. (H)

Insensitivity to 0:
Let $P' = Add(P, i, (\varphi\ 0))$, then $disu_P = disu_{P'}$. (I0)

Decomposition:
Let $P' = Add(P, i, (\varphi\ \alpha))$ and $Q' = Add(Q, i, (\varphi\ \alpha))$, if $\forall j \neq i \quad P_j = \emptyset\ and\ Q_j = \emptyset$ then
$(disu_P = disu_Q) \Rightarrow (disu_{P'} = disu_{Q'})$. (DEC)

Transfer:
Let $P' = Add(Delete(P, i, (\varphi_i^j\ \alpha_i^j)), k, (\varphi_i^j\ \alpha_i^j))$, i.e. P' is obtained from P by transferring $(\varphi_i^j\ \alpha_i^j) \in P_i$ from P_i to P_k, then $disu_P = disu_{P'}$. (TR)

Separation:
If $disu_{Neutralize(P,i)} = disu_{Neutralize(P',i)}$ and $P_i = P_i'$ then $disu_P = disu_{P'}$. (SEP)

Initialization:
If $P = \{\emptyset, ..., \emptyset\}$ then $\forall w \models K,\ disu_P(w) = 0$. (INIT)

Insensitivity towards integrity constraints mean that the way the disutility of a possible world is computed does not depend on where the possible worlds are. This postulate is reminiscent of the independence from irrelevant alternatives assumption is social theory [20].

Anonymity and neutrality are two central properties in social choice theory [20]. *Anonymity* ensures equality between agents, while *neutrality* ensures initial equality between possible worlds. An aggregation rule that would not be neutral is the following one: let us consider a collection of agents in a bus expressing their preferences about opening or closing the window. An aggregation rule such as "if $w \models \neg windowopen$ then $disu_P(w) = \max_i disu_i(w)$ else $disu_P(w) = \frac{\sum_i disu_i(w)}{n}$" is anonymous but not neutral since it favours worlds in which the window is closed. An aggregation rule such as $\forall w \models K, disu_P(w) = \sum_i \lambda_i disu_i(w)$ where λ_i increases with the age of agent i is neutral but not anonymous.

The following group of postulates deal specifically with the influence of valuations on the computation of the disutility function. *Monotonicity* is very natural and does not need any comments. *Strict monotonicity* is much more demanding since it rules out some natural $\diamond = max$. *Priority to hard constraints* is rather natural and can be ignored simply by forbidding agents to specify constraints with infinite penalties. *Insensitivity to 0* expresses the neutrality of the valuation 0. Of course, this neutrality is a consequence of the assumption about negative preferences expressed by (ISF) (see next paragraph). Whether 0 should be neutral or not has been discussed a lot in social decision theory. For practical applications in KR it does not seem to be harmful.

Several postulates deal more specifically with representational issues.

Insensitivity to constraint ordering tells that the ordering in which a given agent has expressed her constraints does not matter. *Irrelevance of the syntax* tells that the logical content only of formulas is relevant when computing disutilities.

Insensitivity to satisfied formulas is a strong assumption. It says that the disutility degree of a world is computed from the constraints whose formulas are violated by this world, only. Or, equivalently, that given a world w, adding to the preference base of an agent any constraint satisfied by w does not change her disutility. This central assumption means that *preferences are represented negatively*, i.e., by a set of constraints (whose violation induces a loss of utility) — in particular, an agent cannot be strictly more satisfied than when she does not specify any constraint. Positively represented preferences would consist in *goals*, i.e., formulas whose satisfaction *increases the utility.* This assumption of negative preferencesd is present in all penalty-based approaches [22] [9] [15] and is often well-suited to practical applications. In [26], this distinction between constraints and goals (as well as the possibility to have preference items which are partly constraints and partly goals) is extensively discussed, using the notion of *polarity.* An extension of our axiomatic system to *positive preferences (goals) only* can be inferred easily by replacing insensitivity to satisfied formulas by insensitivity to violated formulas, and replacing disutilities by (positive) utilities. An extension to preferences with mixed or intermediate polarities is left to further research. *Initialization* is concerned both with the neutral role of 0 and the negative representation of preference, since, together with monotonicity and (ISF), it entails that 0 is the minimal disutility, and that this minimal idusility is reached in particular when agents do not express any constraints.

Decomposition and *Separation* deal with the preservation of the indifference (i.e., the utility equivalence) of two MAPBs when, respectively, a new constraint is added to both of them (to the same agent and with the same weight in both) and when a new preference base is added to both of them.

Lastly, *Transfer* is a very strong assumption (and often undesirable) since, by allowing to tranfer any constraint from one preference base into any other one, it implies that agents are not considered as individuals any longer and that the MAPB becomes equivalent to the single-agent preference base obtained by gathering the constraints expressed by the agents.

Proposition 1
A function $\Phi : \langle P, K\rangle \mapsto disu_{\langle P,K\rangle}$ mapping each pair $\langle P, K\rangle$ into a disutility function $disu_{\langle P,K\rangle}$ on $Mod(K)$ satisfies (A), (N), (MON), (IS), (ISF), (ICO), (I0), (DEC), (SEP), (INIT) and (H) iff there exists a function $* : [0, +\infty] \times [0, +\infty] \rightarrow [0, +\infty]$ commutative, associative, non-decreasing, such that $\alpha * 0 = \alpha$ and $\alpha * +\infty = +\infty$, and a function $\diamond : [0, +\infty]^n \rightarrow [0, +\infty]$, commutative and non-decreasing in each of its arguments, such that $disu_P(w) = \diamond\{disu_1(w), disu_2(w), ..., disu_n(w)\}$ where $disu_i(w) = *\{\alpha_i^j \mid w \models \neg\varphi_i^j\}$.

Proposition 1bis: *A function $\Phi : \langle P, K\rangle \mapsto disu_{\langle P,K\rangle}$ satisfies (A), (N), (MONS), (IS), (ISF), (ICO), (I0), (DEC), (SEP), (INIT), (H) iff there exists two functions $*$ and $\diamond$ as previously and also verifying: $*$ and $\diamond$ satisfy also strict monotonicity.*

The following more specific result also holds:

Proposition 2 *A function $\Phi : \langle P, K\rangle \mapsto disu_{\langle P,K\rangle}$ satisfies (A), (N), (MON), (IS), (ISF), (ICO), (I0), (DEC), (SEP), (INIT), (H) and transfer (TR) iff there exists a function $* : [0, +\infty]^2 \rightarrow [0, +\infty]$ commutative, associative, non-decreasing, such that $\alpha * 0 = \alpha$ and $\alpha * +\infty = +\infty$, and such that:*

$$disu_P(w) = *\{\alpha_i^j \mid w \models \neg\varphi_i^j,\ i = 1..n\}.$$

Proofs are in [16]. Note that strict monotonicity (in Proposition 2) excludes the idempotent operators *min* and *max*. Proposition 3 characterizes weighted logics for which $*$ and $\diamond$ coincide.

2.4 Comments

2.4.1 Weighted logics and Arrow's impossibility theorem

It is well-known that collective choice is a difficult task, especially because numerous impossibility theorems exist; each of these express a set of objectively reasonable postulates and show its inconsistency, which means that there cannot be any procedure for collective choice satisfying these postulates. The most well-known of these theorems is Arrow's [1], which (roughly) states the following: given a society of n agents and a set of $p \geq 3$ candidates, there is no mapping aggregating any collection of individual preference relations consisting of *total orderings over candidates* into a collective preference consisting of a total preordering (indifference are allowed) over candidates satisfying the three following conditions: (i) *unanimity*: if every agent prefers a candidate a to a candidate b then a is preferred to b is collectively preferred to b; (ii) *independence of irrelevant alternatives (IIA)*: for any two candidates a and b, the collective preference between a and b depends only upon individual preferences over that pair; (iii) *non-dictatorship*: there is no individual such that for any two candidates a and b, a is collectively preferred to b whenever this individual prefers a to b.

There are three main ways of escaping Arrow's theorem. The first one consists of relaxing one of the conditions (the condition that is almost always relaxed in the independence of irrelevant alternatives). The second one consists in restricting the domain of possible collections of individual preference relations (or, equivalently, in dropping the requirement that the mapping should be defined for any collection of individual preferences). The third one consists in changing the nature of the input, *i.e.*, dropping the assumption that individual preference should be total orderings over candidates. Clearly enough, weighted logics, by enabling agents expressing preferences *quantitatively* and not under the form of purely ordinal preferences, use the third os the above ways to escape Arrow's theorem. As explained in [13], things are similar for belief merging [23]. The main consequence of this choice is that *preferences should be expressed quantitatively and not in a purely ordinal way*. One the one hand, this has the further advantage to enable agent expressing intensities of preferences, which is a very good point. On the other hand, this leads to another well-known problem in social choice theory, namely, *interpersonal comparison of utilities*. This issue is discussed in the next paragraph.

2.4.2 Utility assessment and interpersonal comparison of utilities

Encoding preferences by means of formulas weighted by a numerical disutility weight raises several comments (see the Example in Section 3.1 for an illustration). In particular, one may wonder about *how this weights are assessed. Why these ones and not other ones?.* A first (but unsufficient) element of answer is the following: we assume that a preliminary (upstream) task of *preference elicitation* consisting in asking each agent questions so as to find how prioritary each other her desires are, and also so as to assess independence relations between the satisfaction of her different desires (see Section 3.1). Now, there is a more tricky question, namely, the problem of *interpersonal comparison of utilities.* This issue is central to group decision theory: as explained for instance in [19], since the utility function representing the true preferences of an (utilitarianist) agent is not unique[6], it is *a priori* meaningless comparing utilities of different agents; however, quoting [19], *in certain contexts, (...) it is reasonable to assume that a commensurating unit and a common origin of measurement exist.* This is indeed the case for most practical problems in restricted domains.

2.5 Computational issues

Although this paper does not focus on how to compute effectively optimal group decisions, we evoke here briefly some questions regarding computational issues.

We first identify the computational complexity of finding an optimal group decision given a MAPB and two aggregation functions $*$ and $\diamondsuit$. The necessary background (omitted for considerations of space) can be found in Papadimitriou's textbook [21], especially in Chapter 17.

When $* = \diamond = \max$ (which makes the framework equivalent to possibilistic logic whose complexity issues are investigated in [17]), finding an optimal assignment for a given multi-agent preference profile can be done using $\mathcal{O}(\log n)$ queries to a satisfiability or unsatisfiability problem. This result is due to these two facts: (i) for any α, the problem of determining whether there exists a possible world w such that $disu_P(w) \leq \alpha$ is coNP-complete; (ii) the use of *max* as aggregation function does not create any new valuations, i.e., $disu_P(w) = x$ iff there is a $(\varphi\ x)$ in one of the P_i's. Therefore, when $* = \diamond = \max$, finding an optimal assignment is in $\mathsf{FP}^{\mathsf{NP}[\mathcal{O}(\log n)]}$. This result holds as well for the general case when penalties appearing in the preference bases are all identical (for instance, unitary), and for the same reasons (i) and (ii) as above. In the general case, however, this result does not hold and what we have is that finding an optimal assignment is in $\mathsf{FP}^{\mathsf{NP}[\mathcal{O}(n)]}$. The complexity of more general problems concerning weighted logics is investigated in [18].

[6] Namely, if u is a utility function representing the preferences of a given agent, then this is also the case of any positive linear combination $au + b$ ($a > 0$, $b \leq 0$).

3 A case study: $\langle +, \max \rangle$

We now choose that preference aggregation between constraints of a given individual is utilitarianist and that aggregation of preferences coming from different agents is utilitarianist. An intuitive argumentation supporting this choice has been given at the end of Section 2.3.

3.1 An example

Example 1:
Three friends (named 1,2,3) are organizing their next vacation. If they decide to spend their vacation together then they have to choose when to go (in summertime S or in wintertime W) and where (in a hot place H or not, in the mountains M or not). 1 likes very much the idea of common vacation, and then prefers to go in summertime and prefers hot places. 2 would really like to go in the mountains and preferably in wintertime. 3 desires very much having vacation and slightly prefers going to the sea in summer. The company permits at most one vacation period (therefore it is not possible to go in vacation both in wintertime and in summertime); there is no hot place in wintertime; furthermore, it is never hot in the mountains.

These specifications are encoded by
$K = \{V \rightarrow (S \leftrightarrow \neg W), W \vee M \rightarrow \neg H\}$,
$P_1 = \{(V\ 60), (V \rightarrow S\ 20), (V \rightarrow H\ 20)\}$,
$P_2 = \{V\ 50), (V \rightarrow M \wedge W\ 25), (V \rightarrow M\ 25)\}$,
$P_3 = \{(V\ 80), (\neg M\ 10), (S\ 10)\}$.

This encoding raises several comments. In particular, one may wonder about *how this weights are assessed.* This question has been already evoked in Section 2.4.2. Let us consider the preferences expressed by agent 2. The weights imply (and should come from) that 2 is indifferent between having vacation in a disliked place at a disliked period and not having vacation, and that the difference of utility between "mountain in winter" and "moutain in summer" is the same as between "moutain in summer" and "not in mountain". As to agent 3, the choice of translating that he slightly prefers going to the sea in summer by $\{(\neg M\ 10), (S\ 10)\}$ expresses independence between these two desires, contrariwise to $\{\neg M \wedge S\ 10)\}$.

Let $w_1 = (V, S, M, \neg H)$, $w_2 = (V, S, \neg M, H)$ and $w_3 = (V, W, M, \neg H)$. We have $disu_1(w_1) = 20$, $disu_2(w_1) = 25$, $disu_3(w_1) = 10$, thus $disu(w) = \max(20, 25, 10) = 25$; $disu(w_2) = \max(0, 50, 0) = 50$ and $disu(w_3) = \max(40, 0, 20) = 40$. It can be checked that w_1 is optimal (for the pair of operators $\langle +, \max\rangle$). If we had chosen $\langle +, +\rangle$ (resp. $\langle \max, \max\rangle$) then the optimal world would have been w_2 (resp. w_3). Note that the three agents are indifferent between any two worlds satisfying both $\neg V$ (which reflects that the variables S,W,H,M are irrelevant when there is no vacation). For any such w we have $disu_1(w) = 60$, $disu_2(w) = 50$, $disu_3(w) = 80$.

3.2 Fairness issues

Expressing preferences in logic in a group decision theory context raises some problems dealing with fairness. We discuss and illustrate here these problems in the special case of the weighted logic associated to $\langle +, \max\rangle$ but we believe that similar problems would hold not only with other aggregation operators but more generally with any logical preference representation language.

3.2.1 Scale sensitivity

Let $P_1 = \{(A\ 60), (B\ 40)\}$ and $P_2 = \{(\neg A\ 6), (\neg B\ 4)\}$. Obviously, specifying larger weights favours agent 1 since the optimal world is (A, B). This is easily solved by requiring the sum of the penalties of each agent to be constant (this was the case in the previous example). Let us now focus on something less obvious.

3.2.2 Sensitivity to constrainedness

Example 2: Let
$P_1 = \{(A \wedge B \wedge C \wedge \neg D\ 100)\}$,
$P_2 = \{(\neg A\ \wedge \neg B\ 50), (\neg C\ 25), (\neg D\ 25)\}$ *and*
$P_3 = \{(\neg A \vee \neg B\ 50), (\neg C \vee \neg D\ 50)\}$.

The disutility function of each of these 3 agents is shown on figure 3.2.3. We notice that the only world that does not induce a maximal disutility to agent 1 is $w_2 = (A, B, C, \neg D)$ and thus it is the only $\langle +, \max\rangle$-optimal world (since the global disutility degree of a world is the maximum of the disutility degrees of each agent for this world). Consequently, agent 1 imposes her preferences to agents 2 and 3. What favours her is not the sum of the weights she assessed but their repartition: first, she chose to put the whole weight on one formula only, so that every world violating this formula has a maximum disutility degree; then, more importantly, this formula has only one model, which means that the whole weight was used to express that *one world is fully acceptable, and all other worlds are unacceptable.* In contrast, agents 2 and 3 have more balanced disutility distributions. For agent 2, one world satifies all constraints, two worlds are rather acceptable and only three have maximum disutility. Agent 3 is even more helpful: nine worlds fully satisfy her, six are reasonably good to her (disutility 50) and only one has maximum disutility.

The problem is rather complex and concerns both (a) the *repartition* of weights on formulas and (b) the *constrainedness* of the formulas — namely, specifying formulas with a small set of models such as large conjunctions (c.f. agent 1) gives an advantage over specifying formulas with a larger set of models. This encourages agents to be unsincere, i.e. to hide their true preferences: thus, an agent would tend to express $\{(X \wedge Y\ 50)\}$ instead of $(X\ 25), (Y\ 25)\}$ even if her desires for X and Y are independent. For instance, it is not unreasonable to assume that agent 1 in unsincere when she expresses P_1: it may actually be the case that in reality, her preferences for A, B, C and $\neg D$ are independent, which corresponds $P_1' = \{(A\ 25), (B\ 25), (C\ 25), (\neg D\ 25)\}$. Clearly, P_1' does not favour agent 1 as much as P_1.

3.2.3 Towards solutions

The abovementioned fairness problems are the consequence of the *freedom* given to agents to specify preference bases as they like. To remedy them, two general classes of solutions can be considered:

(1) *Imposing a format for preference specification.*

Under this assumption, agents are required to respect some syntactical constraints (i) on how the weights are distributed and, especially, (ii) on the formulas allowed (for instance, we may allow for literals only, or disjonctions only, or allowing for a few conjunctions with a bounded number of literals, etc.). This suffers from the severe drawback of loosing expressivity, even if this solution can be satisfactory in many practical situations. For instance, suppose we forbid conjunctions then some desires cannot be expressed, especially in case of dependencies: how would an agent express without conjunctions that she wants to go skiing if the language is composed only of the variables M(ountain) and W(inter)? Clearly, expressing M and W in two separate constraints means something different.

(2) *Postprocessing the disutility distribution.*

Under this assumption, agents are completely free to specify any kind of constraints but once the disutility function computed, it is *postprocessed* so as to homogeneize the disutility function between agents. For instance, normalizing the sum of the penalties distributed is a straightforward postprocessing task (which is equivalent to imposing a fixed penalty amount); as shown above, this normalization is not sufficient. Let us consider Example 2 again. What is needed is a reequilibration of the disutility distribution. To measure how demanding an agent is (i.e., how constrained her disutility function is), a possible approach consists in *measuring the repartition of the distility with an entropy index* and then *transform the distributions so as to make their entropy close — ideally, equal.* This calls for some technical comments. First, entropy is usually computed for a probability distribution, not for a (dis)utility distribution. In order to make things homogeneous, these disutilities have to be redefined so as to be technically equivalent to probability distributions. Then, knowing that the extreme cases are (a) equidisutility and (b) one fully acceptable world only and all other ones fully unacceptable, it is obvious that we first have to transform disutilities into (normalized) utilities. This transformation is the following: let $maxd_i$ be the sum of all penalties given by agent i, and $maxd = \max_{i=1..n} maxd_i$. Clearly enough, no world can have a collective disutility larger then $maxd$ and for any agent i, no world can have a disutility larger than $maxd_i$. Then, for any agent i and any possible world w, let

$$u_i(w) = \frac{maxd_i - disu_i(w)}{\sum\limits_{w' \models K} (maxd_i - disu_i(w'))}$$

Clearly, we have $u_i(w) \in [0, 1]$ and $\sum_{w \models K} u(w) = 1$, which means that u is, technically speaking, a probability distribution. We are now in position to compute its entropy index:

$$H(u_i) = - \sum_{w \models K u_i(w) \neq 0} u_i(w) \ln u_i(w)$$

Note that H increases with the "utility dispersion"; in particular, the entropy is minimal (= 0) in the extreme case when one world is acceptable and all the other ones fully inacceptable (cf. agent 1 in Example 2) and maximal (= $\ln |Mod(K)|$) when all worlds are equally preferred. In example 2, we have $H(u_1) = 0, H(u_2) = 2.44, H(u_3) = 2.66$. The maximal disutility index possible would be $\ln 16 = 2.77$.

Now, the postprocessing task consist of renormalizing the disutility $disu_i$ (or equivalently the utility u_i) of each agent so as to make the entropy indices tending to be equal – ideally, all equal to a fixed quantity x such that $0 < x < \ln |Mod(K)|$. For (demanding) agents with low entropy, the renormalization will decrease the disutility of very disliked worlds. Symetrically, for (non-demanding) agents whose entropy is too high, the disutility of disliked worlds is decreased. There are several possible choices for doing such a transformation (the simplest one consisting of homothetic transformations) and we do not go give them in detail. It is nevertheless important to emphasize that the candidate transformations should not be order-preserving and 0-preserving (worlds with an initial disutility degree of zero should keep this null utility). It may be the case that for some very little demanding agents, their (dis)utility cannot be modified so as to make the entropy index qual to the fixed quantity. This is for instance the case for an agent i such that u_i is constant ("equi-utility"). But this is not a problem, since there is no reason to create arbitrary preferences between worlds for an agent who has none!

Another (drastic but efficient) trick for renormalizing disutility functions consists of fixing a minimum number of worlds with a null disutility, and if an agent does not respect this condition, then one decreases down to 0 the disutility degrees of his preferred worlds; consequently, being too much demanding, for instance all the weight on a single world such, becomes a risky behaviour (this seems rather fair, since the agent is not very cooperative).

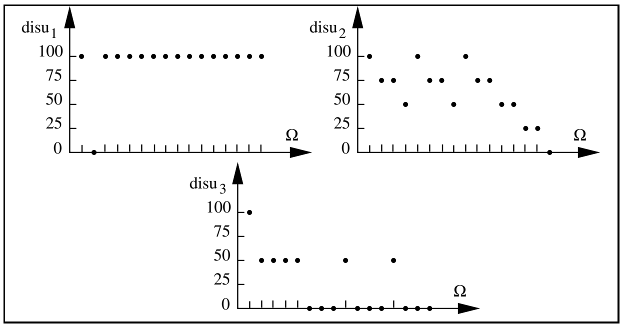

Figure 2: Example 2

4 Distance-based logics

The notion of distance plays an important role in knowledge representation, especially for belief revision, belief update, merging, and similarity-based reasoning (see [3] for a discussion). We show here that it also plays a role for preference representation and group decision making. This has been already (implicitly) remarked in [23]: although their primary motivation is the merging of *beliefs*, some of their examples show that their framework can also be used for aggregating preferences in a group decision context. We now explore this issue further and show a formal parallel between disutilities and distances; the intuition being that if a formula G represents some goal an agent wants to be satisfied, then the closer to G the world w is, the more preferred it is w.r.t. the satisfaction of G.

Definition 5 *A* distance *is a mapping* $d : \Omega \times \Omega \to IN$ *such that*

(i) $\forall w, w' \in \Omega, d(w, w') = 0 \Leftrightarrow w = w'$; *(ii)* $\forall w, w' \in \Omega, d(w, w') = d(w', w)$; *(iii)* $\forall w, w', w'' \in \Omega$,
$$d(w, w'') \leq d(w, w') + d(w', w'').$$

A weak distance *satisfies (ii), (iii) and (i')* $\forall w \in \Omega, d(w, w) = 0$.
Let φ *be a formula and* w, w' *two worlds; we define*
$d(w, \varphi) = \min_{w' \in Mod(\varphi)} d(w, w')$ *and*
$d(\varphi, \psi) = \min_{w \models \varphi, w' \models \psi} d(w, w')$.

A commonly used distance is the Hamming distance: $d_H(w, w')$ is defined as the number of propositional symbols that are assigned different values in the two worlds w, w'. It can be generalized assigning weights to each propositional symbol, the resulting distance being the weighted Hamming distance: $\forall w, w' \in \Omega \quad d_{WH}(w, w') = \sum_{v_i \in \mathcal{V}, w(v_i) \neq w(v_i)} p_i$.
If $\forall v_i \in \mathcal{V} \quad p_i \neq 0$ then d_{WH} is a distance, while if we allow the p_i to be null, i.e. the value of the corresponding propositional symbol has no importance, then d_{WH} is only a weak distance.
From now on we take by default $d = d_H$ (unless stated otherwise).

4.1 From distances to weighted formulas and vice-versa

A distance-based group decision making problem then consists of the data of a distance (assumed identical for all agents – but this constraint could be relaxed) and, for each agent i, a set of goals G_i^j. The disutility of agent i is then computed as $disu_i(w) = *\{d(w, G_i^j) | j\}$ where $*$ is an operator as in Section 2, and inter-agent aggregation is exactly as in Section 2. What remains for us is to exhibit a translation between the specification of a set of weighted formulas and the specification of a set of goals together with a distance. First of all, we need the following definition of a *disc around a formula*, obtained by gathering all worlds that are at most at a given distance from the formula. Formally: for $i \in$ IN , $D(\varphi, i)$ represents a formula such that $Mod(D(\varphi, i)) = \{w \mid d(w, \varphi) \leq i\}$ (as an example, see figure 4.1. Clearly, we have $D(\varphi, 0) = \varphi \models D(\varphi, 1) \models \ldots \models D(\varphi, m) = \top$, where m is the number of propositional variables in the language.

Clearly, we have $D(\varphi_1 \vee ... \vee \varphi_p, k) = D(\varphi_1, k) \vee \ldots \vee D(\varphi_n, k)$. Moreover, if $\gamma = l_1 \wedge \ldots \wedge l_q$ is a consistent conjunction of literals, then it can easily be shown that $\forall k \leq q \quad D(G, k) = \bigvee_{I \subseteq \{1, \ldots, q\}, |I| = q - k} \bigwedge_{i \in I} l_i$

In particular, we have
$D(\gamma, 0) = \gamma, \quad D(\gamma, 1) = \vee(\bigwedge_{i \neq j} l_i), \quad D(\gamma, q) = \top.$

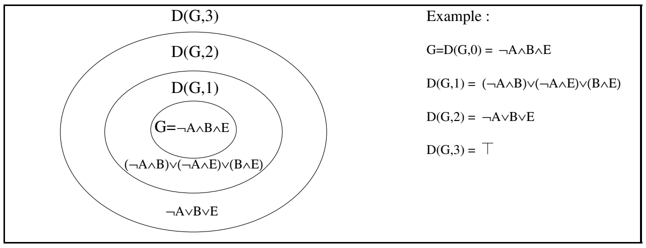

Figure 3: Goals and distance discs

These relations show that computing the disc around a DNF formula is easy:

Proposition 3
For $\varphi = \gamma_1 \vee ... \vee \gamma_p$ *in DNF,*
$$D(\varphi, k) = \bigvee_{\substack{i = 1 \ldots p \\ |I| = q - k}} \bigvee_{\substack{L \subseteq Lit(\gamma_i) \\ |L| = |\gamma_i| - k}} \bigwedge_{l \in L} l$$

Therefore, when φ is under DNF, (a) $D(\varphi, k)$ can be computed in time $\mathcal{O}(n^k)$ and (b) $D(\varphi, k) = \top$ as soon as k reaches the number of literals of the longest of the γ_i's, i.e. as soon as $k \geq \max_{\gamma_i} |\gamma_i|$.

We are now able to establish the formal connection between distance-based and weighted constraint based representation of preferences. With distances, the input consist of a set of goals and a predefined distance; with weighted formulas, it consists of a set of pairs (formula, weight). We now briefly show how both approaches relate. For the sake of clarity we give the transformation in the particular case of penalties ($* =$ sum)[7].

From distances to penalties: to a pair $\langle G, d \rangle$ we associate the set of weighted formulas $F = \{(G\ 1), (D(G, 1)\ 1), \ldots (D(G, k)\ 1)\}$ where k is the largest integer such that $d(G, k) \neq \top$. Then we have $\forall w,\ disu_F(w) = d(w, G)$. This transformation is polynomial when G is under DNF.

From penalties to distances: in this direction, the

[7] A similar work could be done with other aggregation functions; it would however be more complicated and less natural, except in the case of max.

transformation is less easy and need the use of a weak weighted Hamming distance. Let $P = \{(\varphi_1\ \alpha_1), ..., (\varphi_n\ \alpha_n)\}$. We define $K' = K \cup \{\varphi_i \leftrightarrow v_i, i = 1 \ldots n\}$ where the v_i's are new propositional symbols, and let d the weak weighted Hamming distance corresponding to the weights $p_i = \alpha_i$ for each v_i, and 0 to all other propositional symbols. Then we have $\forall w \models K'$, $\sum_i d(w, v_i) = disu_{\langle K,P \rangle}(w)$. This transformation is polynomial. It is however less general than in the other direction, because the distance used depends on the input.

4.2 Robustness

Suppose that a (individual or multi-agent) preference profile is "very consistent", i.e., after the aggregation is finished, there remains many possible worlds with a null disutility (i.e., satisfying all desires). Then, rather than picking any world from this set, we may choose to select one being *robust* in the sense that if new constraints or desires are added later on, this world has a chance to remain a good solution afterwards (a similar idea is the motivation for *supermodels* [12]). This amounts this time to computing discs *inside* (and not any longer *around*) a formula: the inner a world inside a formula, the further from the "boarder" of the formula and the more robust it is. Formally, we define $D(\varphi, -k) = K \wedge \neg D(\neg\varphi, k)$; we have $w \models D(\varphi, -k)$ iff $d \models K$ and $d(w, \neg\varphi) > k$. From what precedes we infer that $D(\varphi, -k)$ can be computed in polynomial time if $\neg\varphi$ is in DNF, i.e., if φ is in CNF. The *kernel* of a formula, defined by $Ker(\varphi) = D(\varphi, -k^*)$ where $k^* = \max\{k | D(\varphi, -k) \neq \bot\}$, contains the "most inner" (and thus preferred) worlds.

5 Related work and conclusion

Merging

Konieczny and Pino-Perez [23] propose some postulates for *merging* and show that an operator Δ (mapping beliefs sets, i.e. multisets of formulas, to knowledge bases) is a merging operator iff there exists an assignment δ mapping each knowledge set E to a preorder $\leq_E$ such that $Mod(\Delta(E)) = min(\leq_E)$. They show (Section 5 of [23]) that some specific merging operators can be induced by distances.

When we make the two following restrictions to our framework: (a) one goal only (G_i) for each agent i; (b) $K = \top$; then the following result holds:

For each distance d and each operator $\diamond$ [8]*, the operator $\oplus$ defined by $Mod(\oplus(G_1, \ldots, G_n)) = \{w \mid w$ optimal for $(G_1, ..., G_n)$ w.r.t. d and $\diamond\}$*

$(= \{w \mid \diamond\{d(w, G_1), \ldots, d(w, G_n)\}\ minimal\})$

is a merging operator.

which establishes a strong link between merging operators and distance-based logics. Using then the results of Section 3.1, we can establish a (less straightforward and more technical) link between merging and weighted logics. These connnections strengthen the (implicit) claim made in [23] that the logic of merging can be useful for preference representation and group decision making.

Supermodels

The idea of selecting a *robust* model comes from Ginsberg et al. [12]. If a and b are integers, then w is a (a, b)-supermodel of φ iff for any consistent conjunction γ of at most a literals, then there exists a world w' such that $w' \models \varphi \wedge \gamma$ and $d(w, w') \leq b$. There is a strong connection with our discs around a formula and *supermodels*[12]. Namely, when $K = \top$, the following result holds:

$w \models D(\varphi, k)$ iff w *is a* $(0, k)$*-supermodel of* φ
or equivalently, iff w *is not a* $(k, 0)$*-supermodel of* $\neg\varphi$

Distance-SAT, MAXSAT and computational issues

DISTANCE-SAT [2] is the following problem: "given a CNF formula φ, a partial world w and an integer k, does there exist a model w' of φ such that w' disagrees with w on at most k variables?". Clearly, the algorithms developed and experimented in [2] are a good starting point when it comes to design efficient algorithms for computing optimal decisions within the distance-based representation model. Similarly, algorithmic studies on MAXSAT (which can be seen as a particular case of weighted logics when $* = \diamond = +$ and all weights are unitary) and on *valued constraint satisfaction* [25] are a good starting point for computing optimal decisions in weighted logics.

Acknowledgments

We would like to thank the participants of the European working group "FUSION" supported by the European Commission for helpful discussions, with a special mention to Salem Benferhat and Isabelle Bloch. Thanks as well to Sébastien Konieczny, Pierre Marquis and Ramon Pino-Pérez.

References

[1] K. Arrow. *Social Choice and Individual Values*, 2nd ed. (1st ed., 1951), John Wiley, 1963.

[8] $*$ is useless because of assumption (a).

[2] O. Bailleux and P. Marquis. DISTANCE-SAT: complexity and algorithms. *Proceedings of AAAI'99*, p. 642-647, 1999.

[3] I. Bloch and J. Lang. Towards mathematical morphologics. To appear.

[4] C. Boutilier. Towards a logic of qualitative decision theory. Proceedings of *KR'94*, p. 75-86, 1994.

[5] R. I. Brafman, N. Friedman. On decision-theoretic foundations for default. Proceedings of *IJCAI'95*, p. 1458-1465, 1995.

[6] M. Dalal. Investigations into a theory of knowledge base revision: preliminary report. Proceedings of *AAAI'98*, p. 475-479, 1988.

[7] D. Dubois, J. Lang and H. Prade. Possibilistic logic. *Handbook of Logic in Artificial Intelligence and Logic Programming* (D.M. Gabbay, C.J. Hogger, J.A. Robinson, eds.), Vol. 3, 439-513, Oxford University Press. 1994.

[8] F. Dupin de Saint-Cyr. *Gestion de l'évolutif et de l'incertain en logiques pondérées* (in French). PhD thesis, Université Paul Sabatier, Toulouse, France, 1996.

[9] F. Dupin de Saint-Cyr, J. Lang and T. Schiex., Penalty logic and its link with Dempster-Shafer theory, *UAI'94*, 1994.

[10] J. Doyle and M. P. Wellman. Preferential semantics for goals. Proceedings of *AAAI'91*, p. 698-703, 1991.

[11] J. Fodor, S. Orlovski, P. Perny and M. Roubens. The use of fuzzy preference models in multiple criteria choice, ranking and uncertainty. Chapter 3 of *Fuzzy sets in decision analysis, operations research and statistics* (R. Slowinski, ed.), pp. 69-101, Kluwer Academic Publishers, 1998.

[12] M. L. Ginsberg, A. J. Parkes, A. Roy, Supermodels and Robustness. Proceedings of *AAAI'98*, p.334-339, 1998.

[13] S. Konieczny. *Sur la logique du changement : révision et fusion de bases de connaissances* (in French). PhD thesis, Université des Sciences et Technologies de Lille, 1999.

[14] M. Grabisch, S. A. Orlovski and R. R. Yager. Fuzzy aggregation of numerical preferences. Chapter 2 of *Fuzzy sets in decision analysis, operations research and statistics* (R. Slowinski, ed.), pp. 31-68, Kluwer Academic Publishers, 1998.

[15] P. Haddawy, S. Hanks, Representations for decision-theoretic planning: utility functions for deadline goals. Proceedings of *KR'92*, p. 71-82, 1992.

[16] C. Lafage, J. Lang. Logical representation of preferences for group decision making. Technical report, IRIT Université Paul Sabatier Toulouse. electronically available at `ftp://ftp.irit.fr/pub/IRIT/RPDMP/langlafa.ps.gz`

[17] J. Lang, Possibilistic logic: complexity and algorithms. To appear in *Handbook of Algorithms for Uncertainty and Defeasible Reasoning* (J. Kohlas, S. Moral, eds.), Kluwer Academic Publishers.

[18] J. Lang and P. Marquis. The complexity of weighted logics and related issues. In preparation.

[19] R.D. Luce, H. Raiffa, *Games and Decisions*, Wiley, New-York, 1967.

[20] H. Moulin, *Axioms of Cooperative Decision Making*, Cambridge University Press, Cambridge, 1988.

[21] C. H. Papadimitriou, *Computational complexity*, Addison-Wesley Publishing Company, 1993.

[22] G. Pinkas, Propositional nonmonotonic reasoning and inconsistency in symmetric neural networks. Proceedings of *IJCAI'91* , p.525-530, 1991.

[23] S. Konieczny, R. Pino-Pérez, On the logic of merging, *KR'98*, p.488-498, 1998.

[24] D. Poole. Decision-theoretic defaults. Proceedings of *UAI'92*, p. 190-197, 1992.

[25] T. Schiex, H. Fargier and G. Verfaillie. Valued constraint satisfaction problems: hard and easy problems, Proceedings of *IJCAI'95*, p. 631-637, 1995.

[26] L. van der Torre and E. Weydert. Parameters for utilitarian desires in a qualitative decision theory. To appear in: *Applied Intelligence*.

[27] R.R. Yager, On ordered weighted averaging aggregation operators in multi-criteria decision making, *IEEE Transactions on Systems, Man and Cybernetics* 18, p.183-190, 1988.