Propositional distances and preference representation

Céline Lafage and Jérôme Lang
IRIT - Université Paul Sabatier, 31062 Toulouse Cedex (France)
{lafage,lang}@irit.fr

No Institute Given

Abstract

Distances between possible worlds play an important role in logic-based knowledge representation (especially in belief change, reasoning about action, belief merging and similarity-based reasoning). We show here how they can be used for representing in a compact and intuitive way the preference profile of an agent, following the principle that given a goal G, then the closer a world w to a model of G, the better w. We give an integrated logical framework for preference representation which handles weighted goals and distances to goals in a uniform way. Then we argue that the widely used Hamming distance (which merely counts the number of propositional symbols assigned a different value by two worlds) is generally too rudimentary and too syntax-sensitive to be suitable in real applications; therefore, we propose a new family of distances, based on Choquet integrals, in which the Hamming distance has exactly a position very similar to that of the arithmetic mean in the class of Choquet integrals for multi-criteria decision making.

1 Introduction

The specification of a decision making or a planning process necessarily includes the goals, desires and preferences of the agent (we will make use of the generic term "preference" for all of these notions – a goal will merely be a specific, strong form of preference). At the object level, the preference structure can be either a utility function assigning each possible consequence to a numerical value, or a weak order relation (possibly allowing for incomparibility in addition to indifference). Now, once it has been fixed what a preference structure is, the question of *how it should be represented*, or in other terms, how it should be encoded so as to be "computationally friendly", arises. A straightforward possibility consists in writing it down explicitly, namely by listing all possible consequences together with their utility values or by listing all pairs of consequences in the preference relation. Clearly, this explicit mode of representation is possible in practice only when the number of possible consequences is reasonably low with respect to the available computational resources. This is often not the case, 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 consequences is equal to the set of feasible assignments and thus grows exponentially with n.

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. They are briefly recalled in Section 2. We particularly focus on two families of representation languages, whose induced preference structure is a utility function: first, weighted logics, which consist in associating to each formula (representing an elementary goal) a positive weight representing the penalty for not satisfying this goal (together with a way of aggregating penalties coming from different violated goals); second, distance-based logics, where the violation of a goal is graded using a *distance* between worlds: given a goal (encoded as a propositional formula) G and given a distance d between worlds, then the closer a world w to a model of G, the better w (the optimal case being when w satisfies φ).

In [15] we showed that in some specific cases, a goal base written within the framework of distance-based logics could be translated into a goal base written within the framework of weighted logics and vice-versa. Here we go further by showing that this translation can be done independently of the chosen distance and aggregation function, and we propose an integrated logical framework which enables both the expression of weighted goals and distance-labelled goals in a uniform way. This is done in Section 3.

In Section 4 we focus on the practical choice of a distance between worlds. We first argue that the widely used distance in the Knowledge Representation community, namely the Hamming distance, is very poor with respect to expressivity and robustness to small changes in the propositional language. This leads us to propose alternative and more general families of propositional distances: we recast the specification or propositional distance in the framework of multicriteria decision making and we introduce the class of Choquet propositional distances, defined as a Choquet integral. We will show that the position of the Hamming distance in this family (to which is belongs) is the same as the position of the arithmetic mean in the family of Choquet integrals; namely, both are central but "degenerate" members of their respective families.

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

2 Logical representation of preferences

2.1 Notations, definitions and basic assumptions

From now on, $\mathcal{L}$ denotes a propositional language built from a finite set of propositional symbols PS, the usual connectives, and two symbols $\top$ and $\bot$ denoting respectively tautology and contradiction. A *literal* is either a propositional symbol or its negation. A *clause* (resp. a *cube*) is a disjunction (resp. a conjunction) of literals. A formula is under CNF (resp. DNF) iff it is a conjunction of clauses (resp. a disjunction of cubes).

$\Omega = 2^{PS}$ is the set of truth assignments (called interpretations, or *worlds*) for $\mathcal{L}$. A world w on PS is denoted by listing the set of literals it satisfies: for instance, if $PS = \{a, b, c, d\}$, then the world in which a and c are false and b and d are true is written as $(\neg a, b, \neg c, d)$. If T is a nonempty subset of VAR, then $w^{\downarrow T}$ denotes the projection of w on T. For instance, $(\neg a, b, \neg c, d)^{\downarrow\{b,c\}}$ is the partial world in which b is true and c is false, written as $(b, \neg c)$.

$Diff(w, w')$ denotes the set of propositional symbols that are assigned a different value by w and w'. For instance, if $w = (\neg a, b, \neg c, d)$ and $w' = (\neg a, \neg b, c, d)$ then $Diff(w, w') = \{b, c\}$. $Mod(\varphi)$ is the set of all models of φ. We write $w \models \varphi$ when $w \in Mod(\varphi)$, i.e., φ is true in w, and $\varphi \models \psi$ when ψ is a logical consequence of φ, i.e., $Mod(\varphi) \subseteq Mod(\psi)$.

Throughout this paper we make the assumption that the set of feasible worlds is $Mod(K)$ where K is a propositional formula expressing integrity constraints. Therefore, the preference structure can be defined on $Mod(K)$. By default we assume $K = \top$.

Lastly, we introduce the notions of [weak][pseudo-] distances between propositional worlds.

Definition 1
A (propositional) weak distance *is a mapping* $d : \Omega \times \Omega \to \mathbb{N}$ *satisfying*

Sep $\forall w, w', d(w, w') = 0 \Leftrightarrow w = w'$
Sym $\forall w, w', d(w, w') = d(w', w)$

If d *satisfies furthermore*

TI $\forall w, w', w'', d(w, w'') \leq d(w, w') + d(w', w'')$

then d *is a* distance.

If φ *is a formula and* $w, w' \in \Omega$ *then* $d(w, \varphi) = \min_{w' \models \varphi} d(w, w')$ *and* $d(\varphi, \psi) = \min_{w \models \varphi, w' \models \psi} d(w, w')$

The *Hamming* distance d_H is defined as the number of propositional symbols that are assigned different values in the two worlds w, w', i.e., $d_H(w, w') = |Diff(w, w')|$. The choice of the propositional Hamming distance is by far the most frequent one in the Knowledge Representation community.

The *drastic* (or *Dirac*) distance d_δ is defined by $d_\delta(w, w') = \begin{cases} 0 \text{ if } w = w' \\ 1 \text{ if } w \neq w' \end{cases}$

2.2 Logical representation of preference: brief overview

A common way for an agent of specifying its preferences consists in listing a set of goals, each consisting of a propositional formula, and possibly some extra information such as weights, priorities, contexts or distances. How the preference relation or the utility function is determined from the set of goals depends on how these goals are understood. We list here (by increasing complexity) the most common constructions of preference structures from logical specifications. In the sequel, the G_i's and the C_i's denote propositional formulas, the α_i's denote numbers in a numerical scale (for example $\mathbb{N}$, or $\mathbb{R}$, or $[0,1]$, or a finite scale etc.), and the d_i's denote weak distances. GB is a "goal base" (by analogy with "knowledge base") and u_{GB} (resp. $\geq_{GB}$) denotes the utility function (resp. the preference relation) induced by GB.

1. $GB = \{G_1, ..., G_n\}$ and $u_{GB}(w) = \begin{cases} 1 \text{ if } w \models G_1 \wedge \ldots \wedge G_n \\ 0 \text{ otherwise} \end{cases}$. This very rough representation assumption (the preference structure – a binary utility – cannot be more elementary) considers crisp goals, interpreted conjunctively and in a dependent way (i.e., the violation of any of the goals cannot be compensated by the satisfaction of another one).
2. $GB = \{G_1, ..., G_n\}$ and $u_{GB}(w) = |\{i | w \models G_i\}|$. This refinement of 1 considers crisp but independent goals and allows for compensation.
3. $GB = \{G_1, ..., G_n\}$ and $w \geq_{GB} w'$ if and only if $\{i | w \models G_i\} \supseteq \{i | w' \models G_i\}$. This (partial) preorder on worlds is nothing but the Pareto ordering.
4. $GB = \langle \{G_1, ..., G_n\}, > \rangle$ where $>$ is a priority order on $\{1, ..., n\}$; $\geq_{GB}$ is then computed in a way extending (2) and (3) so as to take the priority relation between goals into account. See for instance [6].
5. $GB = \{\langle \alpha_1, G_1 \rangle, ..., \langle \alpha_n, G_n \rangle\}$ and
$$u_{GB}(w) = F_1(F_2(\{\alpha_i | w \models G_i\}, F_3(\{\alpha_j | w \models \neg G_j\}))$$
i.e., $u_{GB}(w)$ is a function of the weights of the goals satisfied by w and of the weights of the goals violated by w. When only relative utilities (i.e., differences of utilities between worlds) are important (which is usually the case), we can make the simplifying assumption that only the violated goals count (*negative* representation of preference [24] [14])[2], which leads to this simplified expression: $u_{GB}(w) = F(\{\alpha_i | w \models \neg G_i\}))$. F should of course satisfy some specific properties (see [15]).
6. $GB = \{C_1 : G_1, ..., C_n : G_n\}$. C_i is the *context* of G_i. For this representation by means of *conditional* goals, several ways of constructing the preorder $\geq_{GB}$ have been considered (see [5], [23], [16], [?] as well as extensions to weighted and/or polarized conditional desires [24] [17]).
7. $GB = \langle \{G_1, ..., G_n\}, d \rangle$ where d is a weak distance; and
$$u_{GB}(w) = F(d(w, G_1), ..., d(w, G_n))$$

[2] Or, symetrically, that only the satisfied goals count (*positive* representation of preference).

where F is non-increasing in each of its arguments. The intuition is that ideally, the world w satisfies G, and when it is not the case, then the further to G the world w is, the less preferred it is w.r.t. the satisfaction of G. In a more general framework, the distances do not have to be constant, *i.e.*, we may start with $GB = \{\langle d_1, G_1\rangle, ..., \langle d_n, G_n\rangle\}$ and induce u by $u_{GB}(w) = F(d_1(w, G_1), ..., d_n(w, G_n))$.

In this paper we focus on items 5 (weighted logics) and 7 (distance-based logics). It has been shown in [15] that these two representation languages can be translated from one to each other (under some assumptions). Here we go a bit further and propose a representation language which integrate both, namely, where preference items are triples composed of a weight, a distance and a goal.

3 An integrated framework for the logical representation of weighted and distance-graded goals

In the sequel, we assume that individual goals are specified together with a weight and/or a distance.

Definition 2
A WD goal base *is a finite collection of triples* $GB = \{\langle \alpha_1, d_1, G_1\rangle, ..., \langle \alpha_n, d_n, G_n\rangle\}$ *where, for each* i, $\alpha_i \in I\!N$, d_i *is a weak propositional distance and* G_i *is a consistent propositional formula.*

We have now to choose how a preference structure should be induced from G. The preference structure will be numerical, i.e., a utility function. By convention, we will work with disutility functions on $I\!N$ (or equivalently, utility functions on Z^-): elementary disutilities, or *penalties*, are induced by violation of goals, and ideal worlds, if any, have a null disutility, and therefore satisfy all goals. A disutility function will be denoted by $disu$, with the implicit assumption $disu(w) = -u(w)$.

Definition 3 *Let* $GB = \{\langle \alpha_1, d_1, G_1\rangle, ..., \langle \alpha_n, d_n, G_n\rangle\}$ *be a WD goal base. Then*

$$disu_{GB}(w) = F(H(\alpha_1, d(w, G_1)), ..., (H(\alpha_n, d(w, G_n))$$

where

- F *is a commutative, associative, non-decreasing operator on* $I\!N$, *and has* 0 *as neutral element*[3]*;*
- $H : I\!N \times I\!N \to I\!N$ *is non-decreasing in both arguments and satisfies* $H(\alpha, x) = 0 \Leftrightarrow x = 0$.

[3] Note that F has the properties of *triangular conorms*, except that it is defined on $I\!N$ instead of $[0, 1]$.

Common choices for F are max and $+$ (see [15], especially for an explanation why associativity is required). The operator H, whose role is to aggregate the weight of the goal (reflecting its importance) with the dissatisfaction stemming from the distance to the goal, does not have to be commutative. Monotonicity is natural; the second property means that there is no penalty only if the goal is fully satisfied, i.e., w satisfies it. We consider three possible choices for H, which differ in their interpretation of the weight of a goal:

1. $H_1(\alpha, x) = \min(\alpha, x)$: α is an upper bound of the penalty; once the distance to the goal is at least α, the goal is considered as completely violated: no distinction is made between a distance of α and a larger distance.
2. $H_2(\alpha, x) = \begin{cases} 0 & \text{if } x = 0 \\ \alpha + x - 1 & \text{if } x > 0 \end{cases}$: α is the minimal penalty for a violation (even small) of the goal; after this penalty is taken into account, penalty grows linearly with distance to the goal;
3. $H_3(\alpha, x) = \alpha.x$: α determines the rate of growth of the disutility function.

The choice of H may be further guided by the following consideration: we would like our framework to degenerate (a) into the framework of weighted logics when all distances are taken to be the drastic distance δ and (b) into the framework of distance-based logics when all weights are 1. This imposes the further constraints (a) $H(\alpha, 1) = \alpha$ and (b) $H(1, x) = x$, and therefores excludes H_1. *From now on we assume (unless explicitly stated) that* $H \in \{H_2, H_3\}$.

The simultaneous presence of both weights and distances makes the framework somewhat complex. Furthermore, when GB is a simple base of weighted goals (without distances) $GB = \{\langle\alpha_1, G_1\rangle, ..., \langle\alpha_n, G_n\rangle\}$, we know that $disu_{GB}(w)$ can be computed in polynomial time, whereas it is not the case with distance-based goals, because in general $d(w, G_i)$ cannot be computed in polynomial time.

In the sequel we show that a WD goal base can be translated into a base of weighted goals, while preserving the disutility function. This transformation is generally not polynomial; however, in some cases it is, especially when goals are expressed throw DNF formulas.

For this we start by giving a translation from distance-based weighted goals to standard weighted goals (thus getting rid of distances). We first need the following definitions of a *discs [15] and circles around a formula*:

Definition 4 *Let* $\varphi \in \mathcal{L}$*,* d *a weak distance and* $i \in IN$*. The* disc of radius i around G *is the formula (unique up to logical equivalence)* $D(G, i, d)$ *such that* $Mod(D(G, i, d)) = \{w \in \Omega | d(w, G) \leq i\}$ *and the* circle of radius i around G *is the formula (unique up to logical equivalence)* $C(G, i, d)$ *such that* $Mod(C(G, i, d)) = \{w \in \Omega | d(w, G) = i\}$

Clearly, we have (1) $D(G, 0, d) \equiv C(G, 0, d) \equiv G$; (2) $D(G, i, d) \models D(G, i + 1, d)$; (3) $D(G, dmax, d) \equiv \top$, where $dmax = \max_{w, w' \in \Omega} d(w, w')$ (the latter quantity is defined because Ω is finite); (4) $D(G_1 \vee ... \vee G_p, k, d) \equiv D(G_1, k, d) \vee \ldots \vee D(G_n, k, d)$; (5) $C(G, i, d) \equiv D(G, i, d) \wedge \neg D(G, i - 1, d)$. (6) $D(G, i, d) \equiv C(G, 0, d) \vee ... \vee C(G, i, d)$.

Proposition 1 *Let $GB = \{\langle \alpha, d, G\rangle\}$ be a singleton goal base, and let $\Phi_1(GB)$, $\Phi_2(GB)$ and $\Phi_3(GB)$ be the three distance-free goal bases defined as follows.*

- *let $\Phi_1(GB) = \{\langle 1, \neg C(G,1,d)\rangle, \langle 2, \neg C(G,2,d)\rangle, \ldots, \langle \alpha, \neg C(G,\alpha,d)\rangle\}$. If $H = H_1$ then $disu_{GB} = disu_{\Phi_1(GB)}$.*
- *let $\Phi_2(GB) = \{\langle \alpha, \neg C(G,1,d)\rangle, \langle \alpha+1, \neg C(G,2,d)\rangle, \ldots, \langle \alpha + dmax, \neg C(G, \alpha, dmax)\rangle\}$. If $H = H_2$ then $disu_{GB} = disu_{\Phi_2(GB)}$.*
- *let $\Phi_3(GB) = \{\langle \alpha, \neg C(G,1,d)\rangle, \langle 2\alpha, \neg C(G,2,d)\rangle, \ldots, \langle dmax.\alpha, \neg C(G, dmax, d)\rangle\}$. If $H = H_3$ then $disu_{GB} = disu_{\Phi_3(GB)}$.*

Thus, for $H \in \{H_1, H_2, H_3\}$, a WD goal base can be rewritten equivalently into a distance-free goal base. Unfortunately, there is no guarantee that the size of $\Phi_1(GB)$, $\Phi_2(GB)$, $\Phi_3(GB)$ should be polynomial, because in general $C(G,i,d)$ has a size exponential in $|G|$. However, if G is under DNF and $d = d_H$ then we have the following result:

Proposition 2 *[15]*
Let $d = d_H$ and $G = \gamma_1 \vee \ldots \vee \gamma_p$ be a propositional formula under DNF, where, for each i, $\gamma_i = l_1^i \wedge l_2^i \wedge \ldots \wedge l_{q_i}^i$ is a consistent conjunction of literals. Then

$$D(G,k,d_H) = \bigvee_{i=1\ldots p}\left(\bigvee_{I\subseteq\{1,\ldots,q_i\}, |I|=q_i-k}\left(\bigwedge_{i\in I} l_i\right)\right)$$

Note that, written this way, $D(G,k,d_H)$ still has an exponential size. However, it *does not have to be written this way*: enriching the syntax of propositional logic with so-called cardinality-formulas [4], the size of $D(G,k,d_H)$ is linear in the size of G. Therefore, if G is initially under DNF and d is taken to be the Hamming distance, computing $D(G,k,d_H)$, and therefore, computing $DC(G,k,d_H)$ and $d(w,G)$, can all be done in polynomial time. Moreover we get easily an upper bound of the $d_H(w,G)$ since $D(G,k,d_H) = \top$ as soon as k reaches the number of literals of the longest of the γ_i's, i.e. as soon as $k \geq \max_i |\gamma_i|$.

Unfortunately, what we gain by using the Hamming distance is lost in expressivity and sensitivity to the choice of the propositional symbols. In the following Section we discuss briefly the cons of the Hamming distance and propose a much more general family of weak distances (containing d_H) retaining this friendly computational properties of the Hamming distance.

4 Deconstructing Hamming: topic-decomposable and Choquet-decomposable weak distances

4.1 Against Hamming

It is surprising to see how popular the Hamming distance is in the Knowledge Representation community[4] – not only it is widely used, but it is almost the only

[4] In this research community it is sometimes called Dalal's distance after it was used in [7] for belief revision; it was used as well by Forbus [11] for belief update, and by many authors for belief merging [18] [19] [12] [2].

serious propositional distance considered. However, there are good reasons not to choose it. The Hamming distance assumes not only that propositional symbols are *equally relevant* to determining a distance between worlds (which can be fixed easily by assigning weights to propositional symbols), but also that they are *independent from each other* and that *nothing else is relevant* to the determination of the distance between worlds. These assumptions are extremely compelling and give the Hamming distance very little flexibility. The more serious of its drawbacks is its extreme *sensitivity to the syntax*, or more precisely, to the choice of the propositional language. It is obvious that there is an infinity of possible choices for the set of propositional symbols to be used for representing a certain piece of knowledge. Now, the point with the Hamming distance is that it is extremely sensitive to this choice. Let us consider an example. Assume that we talk about a variable X which can take any integer value between 0 and 7. The most efficient propositional encoding of this variable uses 3 propositional variables x_1, x_2, x_3 being respectively the 1st, 2nd and 3rd bit of the binary representation of the value of X: for instance, $X = 6$ is represented by the propositional formula $x_1 \wedge x_2 \wedge \neg x_3$. The Hamming distance is extremely counterintuitive here: let w_1, w_2 and w_3 be three worlds where X takes respectively the values 3, 4 and 7, all other things being equal. Then $d_H(w_1, w_2) = 3$ while $d_H(w_1, w_3) = 1$. The introduction of weights on the symbols do not help much ($d_H(w_1, w_2)$ will be still be maximal). The less efficient choice of representing X with 8 propositional variables $X = 0$, ..., $X = 7$ is not better since it gives $d_H(w_1, w_2) = d_H(w_1, w_3) = 1$. Finding a representation such that $d_H(w, w') = |v - v'|$ if w and w' differ only in the values v and v' they assign to X is far from being obvious[5].

In the rest of this Section we propose much more general classes of distances containing the Hamming distance and its weighted variants, and which retain as much as possible its computationally friendly properties discussed in Section 3.

4.2 Topic-decomposable weak distances

It is often the case that the set of propositional variables intuitively consists of the union of sets of variables, each of which corresponds to a specific *topic* (see e.g., [20]). It is not necessary to require that topics should be disjoint.

Definition 5 *Let $\mathcal{T} = \{T_1, ..., T_p\}$ be a collection of nonempty subsets of PS* (topics) *such that $\bigcup_i T_i = PS$. A weak distance d is $\mathcal{T}$-decomposable if and only if there exist p weak distances $d_1 : 2^{T_1} \times 2^{T_1} \rightarrow I\!N$, ..., $d_p : 2^{T_p} \times 2^{T_p} \rightarrow I\!N$, and an aggregation function $F : I\!N^p \rightarrow I\!N$, such that for all $w, w' \in \Omega$,*

$$d(w, w') = F(d_1(w^{\downarrow T_1}, w'^{\downarrow T_1}), \ldots, d_p(w^{\downarrow T_p}, w'^{\downarrow T_p}))$$

A topic can be considered as a criterion, hence for the choice of the function F it is relevant to make use of aggregation functions from the MCDM literature.

The simplest possible topic is a singleton and the "most" decomposable distances are those which are decomposable with respect to the set of singletons.

[5] Nevertheless, it is possible, but this representation is not intuitive at all.

Definition 6 *a weak distance d is* symbol-decomposable *if and only if (1) it is $\mathcal{ST}$-decomposable where $\mathcal{ST} = \{\{s_1\}, \{s_2\}, \ldots, \{s_n\}\}$ and (2) for each $i \in 1 \ldots n$, $d_i((s_i), (s_i)) = d_i((\neg s_i), (\neg s_i)) = 0$ and $d_i((s_i), (\neg s_i)) = 1$.*

More generally, a topic may be the set of propositional symbols used for encoding the possible values of a non-binary variable. For instance, if the domain of variable v is $\{1, 2, 3\}$ then we may use the propositional symbols $\{v = 1, v = 2, v = 3\}$[6] together with the integrity constraint stating that one and only one among $v = 1, v = 2, v = 3$ is true[7]. Still more generally, a topic may concern several variables and thus be the union of propositional symbols they generate.

In the definition of a topic-decomposable distance, nothing was said on the aggregation function F. Clearly, how d_1,...,d_p should be aggregated into d is a typical issue of multicriteria decision making. In this paper we focus on the particular class of topic-decomposable distances for which the aggregation function is a *Choquet integral* [22][21].

4.3 Choquet-decomposable weak distances

Definition 7 (fuzzy measures)
A (unnormalized, integer-valued) fuzzy measure *on $\mathcal{T} = \{T_1, ..., T_p\}$ is a mapping $\mu : 2^{\mathcal{T}} \rightarrow I\!N$ such that (1) $\mu(\emptyset) = 0$; (2) for any subsets A and B of $\mathcal{T}$, $A \subseteq B$ implies $\mu(A) \leq \mu(B)$[8]. A fuzzy measure is said* strictly positive *if and only if for all $X \subseteq \mathcal{T}$, $X \neq \emptyset$ implies $\mu(X) > 0$.*

Definition 8 (Choquet integrals) *[22][21]*
Let μ be a fuzzy measure on $2^{\mathcal{T}}$ and $\boldsymbol{x} = (x_1, ..., x_p)$ a vector of real numbers (with $p = |\mathcal{T}|$). Let σ be any permutation of $\{1, ..., p\}$ such that $x_{\sigma(1)} \leq ... \leq x_{\sigma(p)}$. The Choquet integral *of $(x_1, ..., x_p)$ with respect to μ is defined as*

$$C_\mu(x_1, ..., x_p) = \sum_{i=1}^{p} (\mu(\{\sigma(i), ..., \sigma(p)\}) - \mu(\{\sigma(i+1), ..., \sigma(p)\}).x_{\sigma(i)}$$

Definition 9 *A weak distance d is a $\mathcal{T}$-*Choquet-decomposable *if and only if it is $\mathcal{T}$-decomposable and the associated aggregation function F is a Choquet integral.*

1em[6] Note that in spite of the symbol = appearing in them, they are atomic propositional symbols.

1em[7] Note that in practice, it would be more efficient to encode a 3-valued variable with two propositional symbols only: it is is well-known that a variable with a value domain of cardinality k requires $\lceil \ln_2(k) \rceil$, i.e., the upper integer part of the logarithm in base 2 of k.

1em[8] Usually fuzzy measures are also require to satisfy $\mu(\mathcal{T}) = 1$; however this requirement has little impact; for the sake of homogeneity with Section 3, we prefer weak distances to remain integer-valued, which makes us choose integer-valued, and therefore unnormalized, fuzzy measures.

In other terms, a Choquet-decomposable weak distance d is defined by the collection of elementary weak distances d_i and a fuzzy measure $\mu : 2^{\mathcal{T}} \rightarrow [0,1]$ such that for any w, w', we have $d(w,w') = C_\mu(d_1(w,w'), ..., d_p(w,w'))$.

Note that if for all i, d_i is a weak distance on T_i, then it is not necessarily the case that the aggregated distance d is still a weak distance, because it may fail to be separable. The necessary and sufficient condition is that any nonempty subset of T has a non-null measure:

Proposition 3 *Let d_1, ..., d_p be weak distances on T_1, ..., T_p respectively, μ a fuzzy measure on T and d defined by $d(w,w') = C_\mu(d_1(w,w'), ..., d_p(w,w'))$. Then d is a weak distance if and only if μ is strictly positive.*

The problem of whether the triangular inequality (TI) is transferred from the d_i's to d is more complex. In general, it is not true. A sufficient (but non necessary) condition on μ under which d satisfies (TI) as soon as d_1, ..., d_p do is the following:

$$\forall(x_1, ..., x_p)\forall(y_1, ..., y_p),\ C_\mu(x_1, ..., x_p) + C_\mu(y_1, ..., y_p) \geq C_\mu(x_1 + y_1, ..., x_p + y_p)$$

We now focus on a particularly interesting class of weak distances.

4.4 SC-decomposable weak distances

Intuitively, d is SC-decomposable iff it both symbol-decomposable and Choquet-decomposable.

Definition 10 *A weak distance d is a* SC-decomposable *if and only if (1) it is symbol-decomposable and (2) the associated aggregation function F is a Choquet integral.*

These weak distances can easily be characterized:

Proposition 4 *d is SC-decomposable if and only if there exists a fuzzy measure μ such that $d(w,w') = \mu(Diff(w,w'))$.*

From now on we denote by d_μ the STC-decomposable function defined by $d(w,w') = \mu(Diff(w,w'))$.

A special case of interest is when μ is *additive*. In this case we have $d_\mu(w,w') = \sum_{x_i \in Diff(w,w')} \mu(\{i\})$. Therefore, d_μ is a *weighted Hamming distance*. If furthermore we have $\mu(\{i\}) = \mu(\{j\}) = 1$ for all i, j, then $d_\mu(w,w') = |Diff(w,w')|$, i.e., d_μ is the Hamming distance d_H.

This sheds some light of what assumptions are implicitly done when choosing the Hamming distance, namely:

1. symbol-decomposability: criteria or topics are identified with single propositional symbols.
2. 1-additivity: propositional symbols are fully independent.
3. neutrality: propositional symbols have equal importance.

Therefore, the position of the Hamming distance in the set of distances is very similar to the position of the arithmetic mean in the set of aggregation functions: namely, it is central, but the choice of d_H assumes a lot of very strong assumptions without which it is far too arbitrary to be a good choice.

Apart of additive measures, another interesting subclass of measures is the set of measures μ such that $\mu(X)$ depends only on the cardinality of $|X|$. In this case, the Choquet integral associated with μ is an *Ordered Weighted Average* (OWA) [25]. This is equivalent to say that an (integer-valued, unnormalized) OWA is characterized by a vector $\boldsymbol{\alpha} = (\alpha_1, ..., \alpha_p)$[9] such that $C_\mu(x_1, ..., x_p) = \sum_{i=1}^{p} \alpha_i.x_{\sigma(i)}$[10] d_μ is a weak distance if and only if μ is strictly positive, which is equivalent to $\alpha_p > 0$.

The weak distances induced by OWAs include

- the drastic distance d_δ, associated with $\boldsymbol{\alpha} = (0, \ldots, 0, 1)$;
- the Hamming distance d_H, associated with $\boldsymbol{\alpha} = (1, \ldots, 1)$;
- the bounded Hamming distances d_H^k defined by $d_H^k(w, w') = \min(k, d_H(w, w'))$ where $k \in \{1, \ldots, p\}$ (remark that $d_H^1 = d_\delta$ and that $d_H^p = d_H$).

We end this Section by showing that STC-decomposable weak distances retain the nice computational property of the Hamming distances discussed at he end of Section??.

Proposition 5 *Let $d = d_\mu$ and $G = \gamma_1 \vee ... \vee \gamma_p$ be a propositional formula under DNF, where, for each i, $\gamma_i = l_1^i \wedge l_2^i \wedge \ldots \wedge l_{q_i}^i$ is a consistent conjunction of literals. Then*

$$D(G, k, d_\mu) = \bigvee_{i=1\ldots p} \left(\bigvee_{I \subseteq \{1,\ldots,q_i\}, \mu(T \setminus I) \leq k} \left(\bigwedge_{i \in I} l_i \right) \right)$$

Like discs around the DNF for d_H, $D(G, k, d_\mu)$ should not be explicitly written this way (or else it has an exponential size). Slightly generalizing cardinality-formulas by introducing the syntactical notation $(k, \mu) : \{l_1, ..., l_q\}$ – such a formula is satisfied if $\mu(\{s_1, ..., s_q\}) \leq k$, where s_i is the propositional symbol associated with l_i, $D(G, k, d_\mu)$ can be written with a size linear in the size of G.

5 Conclusion

The contribution of this paper to the logical representation of preference in twofold. First, we proposed a unified framework enabling the representation of weighted goals and distance-labelled goals, extending the approach first proposed in [15]. Second, we criticized the frequent use of the Hamming distance for the representation of preference and proposed a much more general family of distances based on Choquet integrals.

1em[9] Note that the usual normalization condition $\sum_{i=1}^{p} \alpha_i = 1$ is not required here.

1em[10] The weights α_i come directly from μ by the following relation: since $\mu(X)$ depends only on $|X|$, let $\mu(X) = f(|X|)$; then $\alpha_p = f(1)$; $\ldots$; $\alpha_i = f(p - i + 1) - f(p - i)$; $\ldots$; $\alpha_1 = f(p) - f(p - 1)$.

References

1. O. Bailleux, P. Marquis. DISTANCE-SAT : complexity and algorithms. Actes de la conférence nationale d'Intelligence Artificielle *AAAI'99*, pages 642–647, 1999.
2. S. Benferhat, D. Dubois, S. Kaci and H. Prade. Encoding information fusion in possibilistic logic: a general framework for rational syntactic merging. *Proceedings of ECAI'2000*,3-7.
3. I. Bloch and J. Lang. Towards mathematical morpho-logics. *Proceedings of IPMU'2000.*
4. B. Benhamou, L. Sais and P. Siegel. Two proof procedures for a cardinality-based language in propositional calculus. *Proceedings of STACS'94.*
5. C. Boutilier. Towards a logic of qualitative decision theory. *Proceedings of KR'94*, p. 75-86, 1994.
6. C. Cayrol. Un modèle logique pour le raisonnement révisable (in French). *Revue d'Intelligence Artificielle* 6, 255-284, *Hermès*, 1992.
7. M. Dalal. Investigations into a theory of knowledge base revision: preliminary report. Proceedings of *AAAI'88*, p. 475-479, 1988.
8. 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.
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. K. D. Forbus. Introducing actions into qualitative simulation. *Proc. IJCAI'89*, 1273-1278.
12. S. Konieczny, R. Pino-Pérez, On the logic of merging. *Proceedings of KR'98*, 488-498, 1998.
13. S. Konieczny, R. Pino-Pérez. Merging with integrity constraints. *Proceedings of ECSQARU'99*, 233-244, Lecture Notes in Artificial Intelligence 1638, Springer Verlag, 1999.
14. C. Lafage. *Représentation logique de préférences. Application à la décision de groupe.* Thèse de l'Unievsrté Paul Sabatier, Toulouse, 2001.
15. C. Lafage, J. Lang. Logical representation of preferences for group decision making. *Proceedings of KR'2000*, 457-468, 2000.
16. J.Lang. Conditional desires and utilities - an alternative approach to qualitative decision theory. *Proceedings of ECAI'96*, 318–322, 1996.
17. J. Lang, L. van der Torre and E. Weydert. Utilitarian desires. To appear in *International Journal on Autonomous Agents and Multi-Agent Systems*, 2001.
18. P. Liberatore and M. Schaerf. Arbitration: a commutative operator for belief revision. *Proceedings of the World Conference on the Fundamentals of Artificial Research* , 1995, 217-228.
19. J. Lin. Integration of weighted knowledge bases. *Artificial Intelligence* 83 (1996), 363-378.
20. P. Marquis and N. Porquet. Decomposing propositional knowledge bases through topics (extended abstract). Personal communication.
21. T. Murofushi and M. Sugeno. An interpretation of fuzzy measures and the Choquet integral as an integral with respect to a fuzzy measure, *Fuzzy Sets and Systems* 29 (1989), 201-227.

22. D. Schmeidler. Integral representation without additivity. *Proc. Americ. Math. Soc.* 97 (1986), 255-261.
23. S.-W. Tan and J. Pearl, Specification and evaluation of preferences for planning under uncertainty, *Proc. KR'94*.
24. L. van der Torre and E. Weydert. Parameters for utilitarian desires in a qualitative decision theory, *Applied Intelligence*, 2001, to appear.
25. 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.