Distance-Based Semantics for Multiple-Valued Logics

Ofer Arieli
Department of Computer Science,
The Academic College of Tel-Aviv
4 Antokolski street, Tel-Aviv 61161, Israel.
oarieli@mta.ac.il

Abstract

We show that the incorporation of distance-based semantics in the context of multiple-valued consequence relations yields a general, simple, and intuitively appealing framework for reasoning with incomplete and inconsistent information.

Introduction

Reasoning with distance functions is a common way of giving semantics to formalisms that are non-monotonic in nature. The basic intuition behind this approach is that, given a set of possible worlds (alternatively, interpretations) that represent the reasoner's epistemic states or the information content of different data sources, the similarity between those worlds can be expressed quantitatively (that is, in terms of distance measurements), and thus can be evaluated by corresponding distance operators. In this respect, there is no wonder that distance semantics has played a prominent role in different paradigms for (non-monotonic) information processing. Two remarkable examples for this are the following:

- Formalisms for modeling belief revision, in which distance minimization corresponds to the idea that the difference between the reasoner's new states of belief and the old one should be kept as minimal as possible, that is, restricted only to what is really implied by the new information (see, e.g., (Lehmann, Magidor, & Schlechta 2001; Peppas, Chopra, & Foo 2004; Delgrande 2004)).
- Database integration systems (Arenas, Bertossi, & Chomicki 1999; 2003; Lin & Mendelzon 1999) and merging operators for independent data-sources (Konieczny, Lang, & Marquis 2002; Konieczny & Pino Pérez 2002), where the basic idea is that the amalgamated information should be kept coherent and at the same time as close as possible to the collective information as it is depicted by the distributed sources.

The goal of this paper is to introduce similar distance considerations in the context of *paraconsistent logics*, that is: formalisms that tolerate inconsistency and do not become trivial in the presence of contradictions (see (da Costa 1974) and (Priest 2002); some collections of papers on this topic appear, e.g., in (Batens *et al.* 2000; Carnielli, Coniglio, & Dóttaviano 2002)). One could identify at least four parties with different philosophical attitudes to such logics: the *traditionalists* defend classical logics and deny any need of paraconsistent logics. On the other extreme, the *dialetheists* contend that the world is fundamentally inconsistent and hence the true logic should be paraconsistent. The *pluralists* view inconsistent structures as fundamental but provisional, and favour their replacement, at least in empirical domains, by consistent counterparts. Finally, the *reformists* defend consistency in ontological matters, but argue that human knowledge and thinking necessarily requires inconsistency, and hence that classical logic should be replaced by a paraconsistent counterpart. The underlying theme here, following the reformists, is that conflicting data is unavoidable in practice, but it corresponds to inadequate information about the real world, and therefore it should be minimized. As we show below, this intuition is nicely and easily expressed in terms of distance semantics. Indeed, the incorporation of distance-based semantics in the context of multiple-valued consequence relations yields a framework in which a variety of paraconsistent multiple-valued logics are definable. These logics are naturally applied in many situations where uncertainty is involved.

The principle of uncertainty minimization by distance semantics is in fact a preference criterion among different interpretations of the premises. In this respect, the formalisms that are defined here may be considered as a certain kind of *preferential logics* (Shoham 1987; 1988; Makinson 1994). In particular, the intuition and the motivation behind this work is closely related to other extensions to multiple-valued semantics of the theory of preferential reasoning (see for instance (Arieli & Avron 1998; 2000; Konieczny & Marquis 2002; Arieli & Denecker 2003; Ben Naim 2005; Arieli 2004; 2006)).

The rest of this paper is organized as follows: in the next section we set up the framework; we consider basic multiple-valued entailments and define their distance-based variants. Then we consider different distance metrics and investigate some of the properties of the induced consequence relations. Finally, we discuss a generalization of the distance-based entailments to prioritized theories and show its usefulness for modeling belief revision and for consistent query answering in database systems. In the last section we conclude.

The Framework

Basic Multiple-Valued Entailments

Definition 1 Let $\mathcal{L}$ be an arbitrary propositional language. A *multiple-valued structure* for $\mathcal{L}$ is a triple $\langle \mathcal{V}, \mathcal{O}, \mathcal{D} \rangle$, where $\mathcal{V}$ is set of elements ("truth values"), $\mathcal{O}$ is a set of operations on $\mathcal{V}$ that correspond to the connectives in $\mathcal{L}$, and $\mathcal{D}$ is a nonempty proper subset of $\mathcal{V}$.

The set $\mathcal{D}$ consists of the *designated* values of $\mathcal{V}$, i.e., those that represent true assertions. In what follows we shall assume that $\mathcal{V}$ contains at least the classical values $\mathsf{true}, \mathsf{false}$, and that $\mathsf{true} \in \mathcal{D}$, $\mathsf{false} \notin \mathcal{D}$.

Definition 2 Let $\mathcal{S} = \langle \mathcal{V}, \mathcal{O}, \mathcal{D} \rangle$ be a multiple-valued structure for a propositional language $\mathcal{L}$.

a) A (multiple-valued) *valuation* ν is a function that assigns an element of $\mathcal{V}$ to each atomic formula in $\mathcal{L}$. Extensions to complex formulae are done as usual. In what follows we shall sometimes write $\nu = \{p_1 : x_1, \ldots, p_n : x_n\}$ to denote that $\nu(p_i) = x_i$ for $i = 1, \ldots, n$. The set of valuations on $\mathcal{V}$ is denoted by $\Lambda^{\mathcal{V}}$.

b) A valuation ν *satisfies* a formula ψ if $\nu(\psi) \in \mathcal{D}$.

c) A valuation ν is a *model* of a set Γ of formulae in $\mathcal{L}$, if ν satisfies every formula in Γ. The set of the models of Γ is denoted by $mod^{\mathcal{S}}(\Gamma)$.

Definition 3 Let $\mathcal{S} = \langle \mathcal{V}, \mathcal{O}, \mathcal{D} \rangle$ be a multiple-valued structure for a language $\mathcal{L}$. A *basic $\mathcal{S}$-entailment* is a relation $\models^{\mathcal{S}}$ between sets of formulae in $\mathcal{L}$ and formulae in $\mathcal{L}$, defined as follows: $\Gamma \models^{\mathcal{S}} \psi$ if every model of Γ satisfies ψ.

Example 4 In many cases the underlying semantical structure of a multiple-valued logic is a lattice, and so it is usual to include in $\mathcal{O}$ (at least) the basic lattice operations. In such cases a conjunction in $\mathcal{L}$ is associated with the join, the disjunction corresponds to the meet, and if the lattice has a negation operator, it is associated with the negation of the language. In what follows we use these definitions for the operators in $\mathcal{O}$. Now, the two-valued structure TWO is defined by the two-valued lattice, and is obtained by taking $\mathcal{V} = \{\mathsf{true}, \mathsf{false}\}$ and $\mathcal{D} = \{\mathsf{true}\}$. The corresponding entailment is denoted $\models^2$. For three-valued structures we take $\mathcal{V} = \{\mathsf{true}, \mathsf{false}, \mathsf{middle}\}$, the lattice operators in $\mathcal{O}$ are defined with respect to the total order $\mathsf{false} < \mathsf{middle} < \mathsf{true}$, and $\mathcal{D}$ is either $\{\mathsf{true}\}$ or $\{\mathsf{true}, \mathsf{middle}\}$. The structure with $\mathcal{D} = \{\mathsf{true}\}$ is denoted here by $\mathsf{THREE}_{\bot}$. The associated entailment, $\models^{3_\bot}$, corresponds to Kleene's three-valued logic (Kleene 1950). The other three-valued structure, $\mathsf{THREE}_{\top}$, corresponds to Priest's logic LP (Priest 1989; 1991).[1] Note that by different choices of the operators in $\mathcal{O}$ other three-valued logics are obtained, line weak Kleene logic, strong Kleene logic, and Łukasiewicz's logic (see, e.g., (Fitting 1990; Avron 1991)). In the four-valued case there are usually two middle elements, denoted here by both and $\mathsf{neither}$.[2] In this context it is usual to take true and both as the designated values. The corresponding structure is known as Belnap's bilattice (see (Belnap 1977a; 1977b) as well as (Arieli & Avron 1998)), and it is denoted here by FOUR. Its entailment is denoted by $\models^4$. Entailments in which $\mathcal{V}$ is the unit interval and $\mathcal{D} = \{1\}$ are common in the context of fuzzy logic (see, e.g., (Hájek 1998)). In this context it is usual to consider different kinds of operations on the unit interval (T-norms, T-conorms, residual implications, etc.), and this is naturally supported in our framework as well. The simplest case is obtained by associating $\wedge$ and $\vee$ with the meet and the join operators on the unit interval, which in this case are the same as the minimum and the maximum functions (respectively), and relating negation to the involutive operator $\neg$, defined for every $0 \le x \le 1$ by $\neg x = 1 - x$. In what follows we denote the corresponding structure ($\mathcal{S}$) by $[0, 1]$.

[1] Also known as J_3, RM_3, and PAC (see (D'ottaviano 1985; Rozoner 1989; Avron 1991) and chapter IX of (Epstein 1990)).

[2] The names of the middle elements correspond to their intuitive meaning as representing conflicts ('both true and false') and incomplete information ('neither true nor false').

Distance-Based Entailments

By their definition, basic $\mathcal{S}$-entailments are monotonic. In addition, some of them are trivial in the presence of contradictions (e.g., $p, \neg p \models^2 q$ and $p, \neg p \models^{3_\bot} q$), or exclude classically valid rules (e.g., $p, \neg p \vee q \not\models^{3_\top} q$ and $p, \neg p \vee q \not\models^4 q$). Common-sense reasoning, on the other hand, is frequently non-monotonic and tolerant to inconsistency. For assuring such properties we consider in what follows distance-based derivatives of the basic entailments. In the sequel, unless otherwise stated, we shall consider *finite* sets of premises in the classical propositional language $\mathcal{L} = \{\neg, \wedge, \vee, \rightarrow\}$, the operators of which correspond, respectively, to a negation, meet, join, and the material implication on the underlying lattice.

Definition 5 A total function $d : U \times U \rightarrow \mathbb{R}^+$ is called *pseudo distance* on U if it is symmetric (that is, $\forall u, v \in U \; d(u, v) = d(v, u)$) and preserves identity ($\forall u, v \in U \; d(u, v) = 0$ iff $u = v$). A *distance function* on U is a pseudo distance on U that satisfies the triangular inequality ($\forall u, v, w \in U \; d(u, v) \le d(u, w) + d(w, v)$).

Definition 6 An *aggregation function* f is a total function that accepts arbitrarily many real numbers[3] and returns a real number. In addition, the following conditions should be satisfied: (a) f is non-decreasing in each of its arguments, (b) $f(x_1, \ldots, x_n) = 0$ if $x_1 = \ldots = x_n = 0$, and (c) $\forall x \in \mathbb{R}, f(x) = x$.

Definition 7 An *$\mathcal{S}$-distance metric* is a quadruple $\mathfrak{D} = \langle \mathcal{S}, d, f, g \rangle$, where $\mathcal{S} = \langle \mathcal{V}, \mathcal{O}, \mathcal{D} \rangle$ is a multiple-valued structure, d is a pseudo distance on the space of the $\mathcal{V}$-valued interpretations $\Lambda^{\mathcal{V}}$, and f and g are aggregation functions.

Definition 8 Given a theory $\Gamma = \{\psi_1, \ldots, \psi_n\}$, a $\mathcal{V}$-valued interpretation ν, and an $\mathcal{S}$-distance metric $\mathfrak{D} = \langle \mathcal{S}, d, f, g \rangle$, define:

- $d_f(\nu, \psi_i) = f_{\mu \in mod^{\mathcal{S}}(\psi_i)} \, d(\mu, \nu)$
- $d_g(\nu, \Gamma) = g(d_f(\nu, \psi_1), \ldots, d_f(\nu, \psi_n))$

[3] This can be formally handled by associating f with the set $\{f_n : \mathbb{R}^n \rightarrow \mathbb{R} \mid n \in \mathbb{N}\}$ of n-ary functions.

It is common to define f as the minimum function, so that a distance between an interpretation ν to a formula ψ is the minimal distance between ν and some model of ψ. Frequent choices of g are the summation function (over the distances to the formulae in Γ) and the maximal value (among those distances).

Note 9 Let $\mathfrak{D} = \langle \mathcal{S}, d, f, g \rangle$ be an $\mathcal{S}$-distance metric. As distances are non-negative numbers, by conditions (a) and (b) in Definition 6, d_f is a non-negative function for every choice of an aggregation function f. This implies that d_g is obtained by applying an aggregation function g on non-negative numbers, and so d_g is non-negative as well.

Definition 10 An $\mathcal{S}$-distance metric $\mathfrak{D} = \langle \mathcal{S}, d, f, g \rangle$ is called *normal*, if: (a) $d_f(\nu, \psi) = 0$ for every $\nu \in mod^{\mathcal{S}}(\psi)$, and (b) $g(x_1, \ldots, x_n) = 0$ only if $x_1 = \ldots = x_n = 0$.

As easily verified, the standard choices of f and g mentioned above preserve the conditions in Definition 10. Thus, for instance, for every multi-valued structure $\mathcal{S}$ and a pseudo distance d, $\mathfrak{D} = \langle \mathcal{S}, d, \min, g \rangle$ is a normal metric for each $g \in \{\Sigma, \max, \text{avg}, \text{median}\}$.[4]

Definition 11 Given a finite theory Γ and an $\mathcal{S}$-distance metric $\mathfrak{D} = \langle \mathcal{S}, d, f, g \rangle$, define:

$$\Delta^{\mathfrak{D}}(\Gamma) = \big\{ \nu \in \Lambda^{\mathcal{V}} \mid \forall \mu \in \Lambda^{\mathcal{V}}\ d_g(\nu, \Gamma) \leq d_g(\mu, \Gamma) \big\}.$$

Proposition 12 *Let $\mathfrak{D} = \langle \mathcal{S}, d, f, g \rangle$ be a normal metric. If $mod^{\mathcal{S}}(\Gamma) \neq \emptyset$ then $\Delta^{\mathfrak{D}}(\Gamma) = mod^{\mathcal{S}}(\Gamma)$.*

Proof. If ν is a model of $\{\psi_1, \ldots, \psi_n\}$, then as $\mathfrak{D}$ is normal, $d_f(\nu, \psi_i) = 0$ for every $1 \leq i \leq n$. Thus, as g is an aggregation function, by condition (b) in Definition 6, $d_g(\nu, \Gamma) = 0$. Since $d_g(\mu, \Gamma) \geq 0$ for every $\mu \in \Lambda^{\mathcal{V}}$ (Note 9), it follows that $\nu \in \Delta^{\mathfrak{D}}(\Gamma)$.

For the converse, consider the following lemma:

Lemma 13 *In every normal metric $\langle \mathcal{S}, d, f, g \rangle$ the function g is strictly positive whenever it has at least one strictly positive argument and the rest of its arguments are non-negative.*

Lemma 13 follows from the fact that $g(x_1, \ldots, x_n) = 0$ iff $x_1 = \ldots = x_n = 0$ (by conditions (b) in Definitions 6 and 10) together with the requirements that g is non-decreasing in each of its arguments (condition (a) in Definitions 6).

To complete the proof of Proposition 12, suppose then that ν is not a model of $\{\psi_1, \ldots, \psi_n\}$. As such, it does not satisfy ψ_k for some $1 \leq k \leq n$, and so $d_f(\nu, \psi_k) > 0$. By Lemma 13, $d_g(\nu, \Gamma) > 0$ as well. On the other hand, we have shown that $d_g(\mu, \Gamma) = 0$ for every $\mu \in mod^{\mathcal{S}}(\Gamma)$, thus $\nu \notin \Delta^{\mathfrak{D}}(\Gamma)$. □

Now we are ready to define distance-based entailments:

Definition 14 For a metric $\mathfrak{D}$, define $\Gamma \models^{\mathfrak{D}} \psi$ if every valuation in $\Delta^{\mathfrak{D}}(\Gamma)$ is a model of ψ.

[4]Note that the arguments of g are non-negative numbers, and so letting g be the summation, average, or median of such numbers preserves condition (b) in Definition 10.

Example 15 Consider $\Gamma = \{p, \neg q, r, p \to q\}$, and let $\mathfrak{D}_2 = \langle \mathsf{TWO}, d_H, \min, \Sigma \rangle$ be a (normal) distance metric, where d_H is the Hamming distance between two-valued valuations[5]. The distances between the relevant two-valued valuations and Γ are given in the following table:

model	p	q	r	d_Σ
ν_1	true	true	true	1
ν_2	true	true	false	2
ν_3	true	false	true	1
ν_4	true	false	false	2
ν_5	false	true	true	2
ν_6	false	true	false	3
ν_7	false	false	true	1
ν_8	false	false	false	2

Thus, $\Delta^{\mathfrak{D}_2}(\Gamma) = \{\nu_1, \nu_3, \nu_7\}$, and so, for instance, $\Gamma \models^{\mathfrak{D}_2} r$, while $\Gamma \not\models^{\mathfrak{D}_2} p$ and $\Gamma \not\models^{\mathfrak{D}_2} q$. This can be intuitively explained by the fact that, unlike p and q, the atomic formula r is not related to the contradictory fragment of Γ, thus it is a reliable information that can be safely deduced from Γ.

Proposition 16 *Let $\mathfrak{D}$ be a normal $\mathcal{S}$-distance metric, and let Γ be a set of formulas in $\mathcal{L}$ such that $mod^{\mathcal{S}}(\Gamma) \neq \emptyset$. Then for every formula ψ in $\mathcal{L}$, $\Gamma \models^{\mathcal{S}} \psi$ iff $\Gamma \models^{\mathfrak{D}} \psi$.*

Proof. Immediately follows from Proposition 12. □

Some important particular cases of Proposition 16 are the following:

Corollary 17 *Let $\mathfrak{D}$ be a normal distance metric in* TWO*. For every classically consistent set of formulas Γ and for every formula ψ, $\Gamma \models^{2} \psi$ iff $\Gamma \models^{\mathfrak{D}} \psi$.*

Proof. By Proposition 16, since every classically consistent theory has a model. □

Corollary 18 *Let $\mathfrak{D}$ be a normal $\mathcal{S}$-distance metric.*
a) If $\mathcal{S} = \mathsf{THREE}_\top$ then $\Gamma \models^{3_\top} \psi$ iff $\Gamma \models^{\mathfrak{D}} \psi$.
b) If $\mathcal{S} = \mathsf{FOUR}$ then $\Gamma \models^{4} \psi$ iff $\Gamma \models^{\mathfrak{D}} \psi$.

Proof. By Proposition 16, since in $\mathsf{THREE}_\top$ and in FOUR, a valuation that assigns the designated middle element to every atom is a model of every theory in the classical propositional language. □

Example 19 Consider again the distance metric $\mathfrak{D}_2$ of Example 15. By Corollary 17, $\models^{\mathfrak{D}_2}$ is the same as $\models^{2}$ with respect to classically consistent sets of premises, but unlike the basic two-valued entailment, it does not become trivial in the presence of contradictions. On the contrary, as Example 15 shows, $\models^{\mathfrak{D}_2}$ allows to draw conclusion from inconsistent theories in a non-trivial way, and so $\models^{\mathfrak{D}_2}$ (as well as many other distance-based relations that are induced by Definition 14; see Proposition 22 below) is a *paraconsistent* consequence relation.

Consider now $\mathfrak{D}^{3_\perp} = \langle \mathsf{THREE}_\perp, d_H, \min, \Sigma \rangle$. The induced entailment, $\models^{\mathfrak{D}^{3_\perp}}$, is again paraconsistent, and with respect to consistent set of premises it coincides with Kleene's logic, $\models^{3_\perp}$ (note that the latter relation is *not* paraconsistent, so in general $\models^{3_\perp}$ and $\models^{\mathfrak{D}^{3_\perp}}$ are *not* the same). By

[5]I.e., $d_H(\nu, \mu)$ is the number of atomic formulas p such that $\nu(p) \neq \mu(p)$; see also the next section.

Corollary 18, the three-valued entailment, $\models^{\mathfrak{D}^{3\top}}$, induced by $\mathfrak{D}^{3\top} = \langle \mathsf{THREE}_\top, d_H, \min, \Sigma\rangle$, and the four-valued entailment $\models^{\mathfrak{D}_4}$, induced by $\mathfrak{D}_4 = \langle \mathsf{FOUR}, d_H, \min, \Sigma\rangle$, are paraconsistent consequence relations that coincide with the consequence relation of Priest's logic LP and with the consequence relation of Belnap's four-valued logic, respectively.

Note that the above observations still hold when the summation function in the metrics is replaced, e.g., by maximum, average, or the median function.

Reasoning with Distance-based Semantics

Distance Functions

A major consideration in the definition of the entailment relations considered in the previous section is the choice of the distance functions. In this section we consider some useful definitions of distances in the context of multiple-valued semantics. For this, we need the following notation.

Notation 20 Denote by Atoms the set of atomic formulas of the language $\mathcal{L}$ and by Atoms(Γ) the set of the atomic formulae that appear in some formula of Γ.

Many distance definitions have been considered in the literature as quantitative measurements of the level of similarity between given interpretations. For instance, the *drastic distance*, considered in (Konieczny, Lang, & Marquis 2002), is defined by

$$d_D(\nu,\mu) = \begin{cases} 0 & \text{if } \nu=\mu, \\ 1 & \text{otherwise.}\end{cases}$$

Another common measurement of the distance between two-valued interpretations is given by the *Hamming distance* that counts the number of atomic formulae that are assigned different truth values by these interpretations (see also (Dalal 1988)):

$$d_H(\nu,\mu) = |\{p \in \mathsf{Atoms} \mid \nu(p) \neq \mu(p)\}|.$$

For three-valued logics (such as Kleene's and Priest's logics considered above) it is possible to apply the same distance measurements, or to use a natural extension of the Hamming distance that considers the distance between the extreme elements true and false as strictly bigger than the distances between each one of them and the middle element. In this case, true is associated with the value 1, false is associated with 0, and the middle element corresponds to $\frac{1}{2}$. The generalized Hamming distance is then defined as follows:

$$d^3_H(\nu,\mu) = \sum_{p\in\mathsf{Atoms}} |\nu(p) - \mu(p)|.$$

This function is used, e.g., in (de Amo, Carnielli, & Marcos 2002) as part of the semantics behind (three-valued) database integration systems.

For four-valued interpretations there is also a natural generalization of the Hamming distance. The idea here is that each one of the four truth values is associated with a pair of two-valued components as follows: true $= (1,0)$, false $= (0,1)$, neither $= (0,0)$, both $= (1,1)$. This pairwise representation preserves Belnap's original four-valued structure (see (Arieli & Denecker 2003; Arieli 2004; 2006)), and so it is a valid rewriting of the truth values. Now, the distance between two values $x = (x_1, x_2)$, $y = (y_1, y_2)$ in this pairwise representation is given by

$$d^4(x,y) = \frac{|x_1 - y_1| + |x_2 - y_2|}{2},$$

so the graphic representation of d^4 on the four-valued structure is the following:

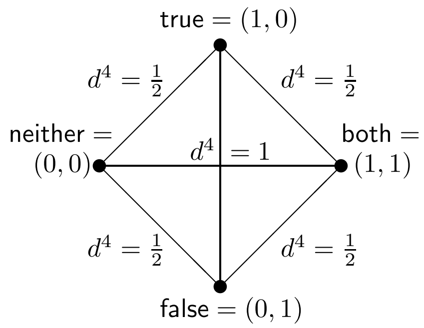

Now, the generalized Hamming distance between two four-valued interpretations ν, μ is defined by:

$$d^4_H(\nu,\mu) = \sum_{p\in\mathsf{Atoms}} d^4\big(\nu(p), \mu(p)\big).$$

Clearly, this definition may be applied on any lattice whose elements have a pairwise representation (see (Arieli 2004; 2006)).

It is not difficult to verify that all the functions defined above satisfy the conditions in Definition 5. Below are some further observations on these distance functions:

1. Given two interpretations ν, μ into $\{\mathsf{true}, \mathsf{false}\}$, it holds that $d^4_H(\nu,\mu) = d^3_H(\nu,\mu) = d_H(\nu,\mu)$, thus d^4_H and d^3_H indeed generalize the standard Hamming distance.
2. As the following example shows, the choice of the distance function (as well as the choice of the other components of a distance metric) has a great impact on the induced entailment.

   Example 21 Consider the following two metrics:

   $$\mathfrak{D}' = \langle \mathsf{THREE}_\bot, d_H, \min, \Sigma\rangle,$$
   $$\mathfrak{D}'' = \langle \mathsf{THREE}_\bot, d^3_H, \min, \Sigma\rangle.$$

   For $\Gamma = \{p, \neg p\}$, we have

   $$\Delta^{\mathfrak{D}'}(\Gamma) = \big\{\{p:\mathsf{true}\}, \{p:\mathsf{false}\}\big\},$$
   $$\Delta^{\mathfrak{D}''}(\Gamma) = \big\{\{p:\mathsf{true}\}, \{p:\mathsf{false}\}, \{p:\mathsf{middle}\}\big\}.$$

   Thus, for instance, $\Gamma \models^{\mathfrak{D}'} p \vee \neg p$, while $\Gamma \not\models^{\mathfrak{D}''} p \vee \neg p$.[6]
3. In (Konieczny, Lang, & Marquis 2002) it is shown that the choice of the distance function has also a major affect on the computational complexity of the underlying formalism. See Section 4 of that paper for some complexity results of distance-based operators when $\mathcal{S} = \mathsf{TWO}$.

[6]This is so, since $\nu(p \vee \neg p) = \mathsf{middle}$ when $\nu(p) = \mathsf{middle}$, and in $\mathsf{THREE}_\bot$ the middle element is not designated.

Basic Properties of $\models^{\mathfrak{D}}$

Paraconsistency. In what follows we consider some characteristic properties of the distance-based entailments. We begin with the ability to reason with inconsistent theories in a non-trivial way. The following proposition shows that this property is common to many distance-based logics that are definable within our framework.

Proposition 22 *The consequence relations $\models^{\mathfrak{D}}$, induced by the following metrics, are all paraconsistent:*

a) *$\mathfrak{D} = \langle \mathsf{TWO}, d, \min, g\rangle$, where d is the drastic distance (d_D) or the Hamming distance (d_H) and g is either a summation or a maximum function.*

b) *$\mathfrak{D} = \langle \mathsf{THREE}_\bot, d, \min, g\rangle$, where $d \in \{d_D, d_H, d^3_H\}$ and g is either a summation or a maximum function.*

c) *$\mathfrak{D} = \langle \mathsf{THREE}_\top, d, \min, g\rangle$, where $d \in \{d_D, d_H, d^3_H\}$ and g is either a summation or a maximum function.*

d) *$\mathfrak{D} = \langle \mathsf{FOUR}, d, \min, g\rangle$, where d is any distance function of those considered in the previous section and g is either a summation or a maximum function.*

e) *$\mathfrak{D} = \langle [0,1], d, \min, g\rangle$, where d is the drastic distance or the Hamming distance and g is either a summation or a maximum function.*

Proof. For any of the items above we shall show that $p, \neg p \not\models^{\mathfrak{D}} q$, and so it is *not* the case that any formula follows from an inconsistent theory. Indeed, in item (a) we have that $\{p:\mathsf{true},\ q:\mathsf{false}\}$ (as well as $\{p:\mathsf{false},\ q:\mathsf{false}\}$) is in $\Delta^{\mathfrak{D}}(\{p,\neg p\})$, thus q does not follow from $\{p,\neg p\}$. For item (b) note that although different distance functions induce different sets of preferred models of $\{p,\neg p\}$ (see Example 21), it is easy to verify that whenever g is the summation function then $\{p:\mathsf{true},\ q:\mathsf{false}\}$ is, e.g., an element of $\Delta^{\mathfrak{D}}(\{p,\neg p\})$, and whenever g is the maximum function $\{p:\mathsf{middle},\ q:\mathsf{false}\}$ is an element of $\Delta^{\mathfrak{D}}(\{p,\neg p\})$. Thus, in both cases, q does not follow from $\{p,\neg p\}$. Part (c) holds since by Proposition 12 we have that $\Delta^{\mathfrak{D}}(\{p,\neg p\}) = mod^{3_\top}(\{p,\neg p\})$, and so $\{p:\mathsf{middle}, q:\mathsf{false}\}$ is an element in $\Delta^{\mathfrak{D}}(\{p,\neg p\})$ (recall that in $\mathsf{THREE}_\top$ the middle element is designated, and so $\{p:\mathsf{middle}\}$ is a model of $\{p,\neg p\}$). We therefore again have that $p,\neg p \not\models^{\mathfrak{D}} q$. The proof of part (d) is similar to that of part (c) with the obvious adjustments to the four-valued case. Part (e) is similar to part (a) replacing, respectively, true and false by 1 and 0. ☐

Monotonicity. Next we consider monotonicity, that is: whether the set of $\models^{\mathfrak{D}}$-conclusions is non-decreasing in terms of the size of the premises. As the next two propositions show, this property is determined by the multi-valued structure and the distance metric at hand:

Proposition 23 *Let $\mathfrak{D}$ be a normal distance metric for FOUR. Then the corresponding distance-based entailment, $\models^{\mathfrak{D}}$, is monotonic.*

Proof. By Corollary 18(b), $\models^{\mathfrak{D}}$ is the same as the basic four-valued entailment $\models^4$ of Belnap's logic. The proposition now follows from the monotonicity of the latter (see (Arieli & Avron 1996, Theorem 3.10) and (Arieli & Avron 1998, Proposition 19)). ☐

Proposition 24 *Let $\mathfrak{D} = \langle \mathsf{TWO}, d, \min, g\rangle$ be a normal distance metric such that $g(x_1,\ldots,x_n) \leq g(y_1,\ldots,y_m)$ if $\{x_1,\ldots,x_n\} \subseteq \{y_1,\ldots,y_m\}$.*[7] *Then the corresponding distance-based entailment, $\models^{\mathfrak{D}}$, is non-monotonic.*

Proof. Consider, e.g., $\Gamma = \{p, \neg p \vee q\}$. By Corollary 17, $\Gamma \models^{\mathfrak{D}} q$. On the other hand, consider $\Gamma' = \Gamma \cup \{\neg p\}$, and let ν_t and ν_f be two-valued valuations that respectively assign true and false to p. By the assumption on g we have that

$$\begin{aligned} d_g(\nu_t,\Gamma') &= g\big(d_{\min}(\nu_t,\neg p), d_{\min}(\nu_t, \neg p \vee q)\big) \\ &\geq g\big(d_{\min}(\nu_t,\neg p)\big) \\ &= d_{\min}(\nu_t,\neg p) \\ &= d_{\min}(\nu_f, p) \\ &= g\big(d_{\min}(\nu_f,p)\big) \\ &= d_g\big(\nu_f,\Gamma'\big). \end{aligned}$$

It follows, then, that every two-valued valuation ν_f that assigns false to p is in $\Delta^{\mathfrak{D}}(\Gamma')$, no matter what value it assigns to q (as $d_g(\nu_f,\Gamma')$ is not affected by $\nu_f(q)$). In particular, $\Delta^{\mathfrak{D}}(\Gamma')$ contains valuations that assign false to q, and so $\Gamma' \not\models^{\mathfrak{D}} q$. ☐

Rationality. In (Lehmann & Magidor 1992), Lehmann and Magidor consider some properties that a "rational" non-monotonic consequence relation should satisfy. One property that is considered as particularly important assures that a reasoner will not have to retract any previous conclusion when learning about a new fact that has no influence on the existing set of premises. Consequence relations that satisfy this property are called *rational*. Next we show that many distant-based entailments are indeed "rational".

Notation 25 An aggregation function f is called *hereditary*, if $f(x_1,\ldots,x_n,z_1,\ldots,z_m) < f(y_1,\ldots,y_n,z_1,\ldots,z_m)$ whenever $f(x_1,\ldots,x_n) < f(y_1,\ldots,y_n)$.[8]

Proposition 26 Let $\mathfrak{D} = \langle \mathcal{S}, d, f, g\rangle$ be an $\mathcal{S}$-distance metric with a hereditary function g. If $\Gamma \models^{\mathfrak{D}} \psi$ then $\Gamma, \phi \models^{\mathfrak{D}} \psi$ for every ϕ such that $\mathsf{Atoms}(\Gamma \cup \{\psi\}) \cap \mathsf{Atoms}(\phi) = \emptyset$.

Intuitively, the condition on ϕ in Proposition 26 guarantees that ϕ is 'irrelevant' for Γ and ψ. The intuitive meaning of Proposition 26 is, therefore, that the reasoner does not have to retract ψ when learning that ϕ holds.

Proof of Proposition 26. Let $\mu \in \Lambda^{\mathcal{V}}$ be a valuation that does not satisfy ψ. As $\Gamma \models^{\mathfrak{D}} \psi$ while $\mu(\psi) \notin \mathcal{D}$, necessarily μ is not in $\Delta^{\mathfrak{D}}(\Gamma)$, and so there is a valuation ν in $\Delta^{\mathfrak{D}}(\Gamma)$, for which $d_g(\nu,\Gamma) < d_g(\mu,\Gamma)$. Again, since $\Gamma \models^{\mathfrak{D}} \psi$, $\nu(\psi) \in \mathcal{D}$. Assuming that $\Gamma = \{\psi_1,\ldots,\psi_n\}$, we have that

$$g(d_f(\nu,\psi_1),\ldots,d_f(\nu,\psi_n)) < g(d_f(\mu,\psi_1),\ldots,d_f(\mu,\psi_n)).$$

Now, consider a valuation σ, defined for every atom p as follows:

$$\sigma(p) = \begin{cases} \nu(p) & \text{if } p \in \mathsf{Atoms}(\Gamma \cup \psi) \\ \mu(p) & \text{otherwise} \end{cases}$$

[7] As the arguments of g are non-negative, summation, maximum, and many other aggregation functions satisfy this property.

[8] Note that heredity, unlike monotonicity, is defined by strict inequalities. Thus, for instance, the summation is hereditary, while the maximum function is not.

Note that $\sigma(p) = \nu(p)$ for every $p \in \mathsf{Atoms}(\psi)$, and so $\sigma(\psi) \in \mathcal{D}$ as well. As $\mathsf{Atoms}(\Gamma \cup \{\psi\}) \cap \mathsf{Atoms}(\phi) = \emptyset$ and since g is hereditary, we have that

$$\begin{aligned} d_g(\sigma, \Gamma \cup \{\phi\}) &= g(d_f(\sigma,\psi_1),\ldots,d_f(\sigma,\psi_n),d_f(\sigma,\phi)) \\ &= g(d_f(\nu,\psi_1),\ldots,d_f(\nu,\psi_n),d_f(\mu,\phi)) \\ &< g(d_f(\mu,\psi_1),\ldots,d_f(\mu,\psi_n),d_f(\mu,\phi)) \\ &= d_g(\mu, \Gamma \cup \{\phi\}). \end{aligned}$$

Thus, for every valuation μ such that $\mu(\psi) \notin \mathcal{D}$ there is a valuation σ such that $\sigma(\psi) \in \mathcal{D}$ and $d_g(\sigma, \Gamma \cup \{\phi\}) < d_g(\mu, \Gamma \cup \{\phi\})$. It follows that the elements of $\Delta^{\mathfrak{D}}(\Gamma \cup \{\phi\})$ must satisfy ψ, and so $\Gamma, \phi \models^{\mathfrak{D}} \psi$. □

Adaptivity. The ability to handle theories with contradictions in a nontrivial way and at the same time to presuppose a consistency of all sentences 'unless and until proven otherwise', is called *adaptivity* (Batens 1989; 1998). Consequence relations with this property *adapt* to the *specific* inconsistencies that occur in the theories. For instance, a plausible inference mechanism should *not* apply the Disjunctive Syllogism for concluding that q follows from $\{p, \neg p, \neg p \vee q\}$. On the other hand, in the case of $\{p, \neg p, r, \neg r \vee q\}$, applying the Disjunctive Syllogism to r and $\neg r \vee q$ may be justified by the fact that the subset of formulae to which the Disjunctive Syllogism is applied should not be affected by the inconsistency of the whole theory, therefore inference rules that are classically valid can be applied to it.

The following proposition shows that in many cases distance-based entailments are adaptive. If a given theory can be split up to a consistent and an inconsistent parts, then every assertion that is not related to the inconsistent part, and which classically follows from the consistent part, must be entailed by the whole theory.

Proposition 27 Let $\mathfrak{D} = \langle \mathcal{S}, d, f, g \rangle$ be a normal $\mathcal{S}$-distance metric with a hereditary function g. Suppose that Γ is a theory that can be represented as $\Gamma' \cup \Gamma''$, where $mod^{\mathcal{S}}(\Gamma') \neq \emptyset$ and $\mathsf{Atoms}(\Gamma') \cap \mathsf{Atoms}(\Gamma'') = \emptyset$. Then for every formula ψ such that $\mathsf{Atoms}(\psi) \cap \mathsf{Atoms}(\Gamma'') = \emptyset$, it holds that if $\Gamma' \models^{\mathcal{S}} \psi$ then $\Gamma \models^{\mathfrak{D}} \psi$.

Proof. If $\Gamma' \models^{\mathcal{S}} \psi$, then by Proposition 16, $\Gamma' \models^{\mathfrak{D}} \psi$. Now, as $\mathsf{Atoms}(\Gamma' \cup \{\psi\}) \cap \mathsf{Atoms}(\Gamma'') = \emptyset$, we have, by Proposition 26, that $\Gamma \models^{\mathfrak{D}} \psi$. □

Distance-based Entailments for Prioritized Theories

$\models^{\mathfrak{D}}$, Generalized

We now extend the distance-based semantics of the previous section to *prioritized theories*. An n-prioritized theory is a theory $\Gamma = \Gamma_1 \cup \ldots \cup \Gamma_n$, where the sets Γ_i ($1 \leq i \leq n$) are pairwise disjoint. Intuitively, when $i < j$ the formulas in Γ_i are preferred than those in Γ_j. A common situation in which theories are prioritized is, e.g., when data-sources are augmented with integrity constraints. In such cases the corresponding theory has two priority levels, as the integrity constraints must always be satisfied, while the data facts may be revised in case of conflicts.

To formalize the existence of different levels of priority in prioritized theories, we consider the following sequence of sets: for a metric $\mathfrak{D} = \langle \mathcal{S}, d, f, g \rangle$ and an n-prioritized theory $\Gamma = \Gamma_1 \cup \ldots \cup \Gamma_n$, define:

- $\Delta^{\mathfrak{D}}_1(\Gamma) = \{\nu \in \Lambda^{\mathcal{V}} \mid \forall \mu \in \Lambda^{\mathcal{V}}\ d_g(\nu, \Gamma_1) \leq d_g(\mu, \Gamma_1)\}$
- for every $1 < i \leq n$, let
  $\Delta^{\mathfrak{D}}_i(\Gamma) = \{\nu \in \Delta^{\mathfrak{D}}_{i-1}(\Gamma) \mid$
  $\forall \mu \in \Delta^{\mathfrak{D}}_{i-1}(\Gamma)\ d_g(\nu, \Gamma_i) \leq d_g(\mu, \Gamma_i)\}$

Definition 28 Given an $\mathcal{S}$-distance metric $\mathfrak{D}$, define for every n-prioritized theory Γ and formula ψ, $\Gamma \models^{\mathfrak{D}} \psi$ if every valuation in $\Delta^{\mathfrak{D}}_n(\Gamma)$ satisfies ψ.

Note that the last definition is a conservative extension of Definition 14, since for non-prioritized theories (i.e., when $n = 1$) the two definitions coincide.

Example 29 Consider the following puzzle, known as the Tweety dilemma:

$$\Gamma = \left\{ \begin{array}{l} \mathtt{bird(x)} \to \mathtt{fly(x)}, \\ \mathtt{penguin(x)} \to \mathtt{bird(x)}, \\ \mathtt{penguin(x)} \to \neg\mathtt{fly(x)}, \\ \mathtt{bird(Tweety)}, \\ \mathtt{penguin(Tweety)} \end{array} \right\}$$

As this theory is not consistent, everything classically follows from it, including, e.g., `fly(Tweety)`, which seems a counter-intuitive conclusion in this case, as penguins should not fly, although they are birds. The reason for this anomaly is that all the formulas above have the same importance, in contrast to the intuitive understanding of this case. Indeed,

1. The confidence level of strict facts (`bird(Tweety)` and `penguin(Tweety)` in our case) is usually at least as high as the confidence level of general rules (implications).
2. As penguins *never* fly, and this is a characteristic feature of penguins (without exceptions), one would probably like to attach to the assertion $\mathtt{penguin(x)} \to \neg\mathtt{fly(x)}$ a higher priority than that of $\mathtt{bird(x)} \to \mathtt{fly(x)}$, which states only a default property of birds.[9]

Consider now the metric $\mathfrak{D} = \langle \mathsf{TWO}, d_H, \min, \Sigma \rangle$ and regard Γ as a prioritized theory in which the two considerations above are satisfied. It is easy to verify that the unique valuation in $\Delta^{\mathfrak{D}}_n(\Gamma)$ (where $n > 1$ is the number of priority levels in Γ) assigns true to `bird(Tweety)`, true to `penguin(Tweety)`, and false to `fly(Tweety)`. Thus, e.g., $\Gamma \models^{\mathfrak{D}} \neg\mathtt{fly(Tweety)}$, as intuitively expected.

Applications

In this section we show how the generalized distance-based semantics for prioritized theories, introduced in the previous section, can be naturally applied in related areas. Below we consider two representative examples: database query systems and belief revision theory.

[9] The third assertion, $\mathtt{penguin(x)} \to \mathtt{bird(x)}$, could have an intermediate priority, as again there are no exceptions to the fact that every penguin is a bird, but still penguins are not *typical* birds, thus they shouldn't inherit all the properties we expect birds to have.

A. Consistent Query Answering in Database Systems

A particularly important context in which reasoning with prioritized theories naturally emerges is consistency handling in database systems. In such systems, it is of practical importance to enforce the validity of the data facts by a set of integrity constraints. In case of any violation of some integrity constraint, the set of data-facts is supposed to be modified in order to restore the database consistency. It follows, then, that integrity constraints are superior than the facts themselves, and so the underlying theory is a prioritized one. This also implies that consistent query answering from possibly inconsistent databases (Arenas, Bertossi, & Chomicki 1999; 2003; Greco & Zumpano 2000; Bravo & Bertossi 2003; Eiter 2005) or constraint data-sources (Konieczny, Lang, & Marquis 2002; Konieczny & Pino Pérez 2002) may be defined in terms of distance-based entailments on prioritized theories. Moreover, as our framework is tolerant to different semantics, such methods of query answering, which are traditionally two-valued ones, may be related to other formalisms that are based on many-valued semantics like those considered in (Subrahmanian 1994) and (de Amo, Carnielli, & Marcos 2002).

Let $\mathcal{L}$ be a propositional language with Atoms its underlying set of atomic propositions. A (propositional) *database instance* $\mathcal{I}$ is a finite subset of Atoms. The semantics of a database instance is given by the conjunction of the atoms in $\mathcal{I}$, augmented with the *closed world assumption* ($\mathsf{CWA}(\mathcal{I})$) (Reiter 1978) that assures that each atom which is not explicitly mentioned in $\mathcal{I}$ is false.

Definition 30 A *database* is a pair $(\mathcal{I}, \mathcal{C})$, where $\mathcal{I}$ is a database instance, and $\mathcal{C}$ — the set of *integrity constraints* — is a finite and consistent set of formulae in $\mathcal{L}$. A database $\mathcal{DB} = (\mathcal{I}, \mathcal{C})$ is *consistent* if every formula in $\mathcal{C}$ follows from $\mathcal{I}$, that is, there is no integrity constraint that is violated in $\mathcal{I}$.

Given a database $\mathcal{DB} = (\mathcal{I}, \mathcal{C})$, the theory $\Gamma_{\mathcal{DB}}$ that is associated with it contains the components of $\mathcal{DB}$ and imposes the closed word assumption on $\mathcal{I}$. In addition, this theory should reflect the fact that the integrity constraints in $\mathcal{C}$ are of higher priority than the rest of the data. That is, $\Gamma_{\mathcal{DB}}$ should be a two-leveled theory, in which $\Gamma_1 = \mathcal{C}$ and $\Gamma_2 = \mathcal{I} \cup \mathsf{CWA}(\mathcal{I})$. Now, query answering with respect to $\mathcal{DB}$ may be defined in terms of a distance-based entailment on $\Gamma_{\mathcal{DB}}$.

Suppose, then, that $\mathfrak{D}$ is a normal $\mathcal{S}$-distance metric for some multiple-valued structure $\mathcal{S}$, and let $\mathcal{DB} = (\mathcal{I}, \mathcal{C})$ be a (possibly inconsistent) database. Its prioritized theory is

$$\Gamma_{\mathcal{DB}} = \Gamma_1 \cup \Gamma_2 = \mathcal{C} \cup (\mathcal{I} \cup \mathsf{CWA}(\mathcal{I})),$$

and $\mathcal{Q}$ is a consistent query answer if $\Gamma_{\mathcal{DB}} \models^{\mathfrak{D}} \mathcal{Q}$. Now, as $\mathcal{C}$ is classically consistent, by Proposition 12, $\Delta_1^{\mathfrak{D}}(\Gamma_{\mathcal{DB}}) = mod(\mathcal{C})$. It follows, therefore, that $\mathcal{Q}$ is a consistent query answer of $\mathcal{DB}$ if it is satisfied by every model of $\mathcal{C}$ with minimal distance (in terms of d_g) from $\mathcal{I} \cup \mathsf{CWA}(\mathcal{I})$.

Example 31 Let $\mathcal{DB} = (\{p, r\}, \{p \to q\})$. Here,

$$\mathcal{I} \cup \mathsf{CWA}(\mathcal{I}) = \mathcal{I} \cup \{\neg x \mid x \notin \mathcal{I}\} = \{p, \neg q, r\},$$

so the associated theory is

$$\Gamma_{\mathcal{DB}} = \{p \to q\} \cup \{p, \neg q, r\}.$$

This theory is the same as the one considered in Example 15, but with one major difference: now $p \to q$ is preferred over the other formulas, thus only its models are taken into account. Consider the same metric as that of Example 15. As valuations ν_3, ν_4 in the table of that example do not satisfy $\mathcal{C}$, they are excluded. Among the remaining valuations, ν_1 and ν_7 are the closest to $\mathcal{I} \cup \mathsf{CWA}(\mathcal{I})$, and so the consistent query answers of $(\mathcal{I}, \mathcal{C})$ are the formulas that are satisfied by both ν_1 and ν_7.

Note 32 Example 31 shows, in particular, that $\models^{\mathfrak{D}}$ is *not* reflexive, since for instance $\Gamma_{\mathcal{DB}} \not\models^{\mathfrak{D}} p$ although $p \in \Gamma_{\mathcal{DB}}$. This can be justified by the fact that one way of restoring the consistency of $\mathcal{DB}$ is by removing p from $\mathcal{I}$ (ν_7 corresponds to this situation), and so p does not hold in all the consistency 'repairs' of $\Gamma_{\mathcal{DB}}$.[10] Similarly, the fact that $\Gamma_{\mathcal{DB}} \not\models^{\mathfrak{D}} \neg q$ although $\neg q \in \Gamma_{\mathcal{DB}}$ may be justified by the alternative way of restoring the consistency of $\mathcal{DB}$, in which q is added to $\mathcal{I}$ (ν_1 corresponds to this situation). Note also that there is no reason to remove r from $\mathcal{I}$, as this will not contribute to the consistency restoration of $\mathcal{DB}$. This intuitively justifies the fact that for r (unlike the other atomic formulae in $\Gamma_{\mathcal{DB}}$), we do have that $\Gamma_{\mathcal{DB}} \models^{\mathfrak{D}} r$ (cf. Example 15). This is also to the intuition behind the query answering formalisms for inconsistent databases, considered e.g. in (Arenas, Bertossi, & Chomicki 1999; 2003; Greco & Zumpano 2000; Bravo & Bertossi 2003; Eiter *et al.* 2003; Arieli *et al.* 2004; 2006).

B. Modeling of Belief Revision

A belief revision theory describes how a belief state is obtained by the revision of a belief state $\mathcal{B}$ by some new information, ψ. A belief revision operator $\circ$ describes the kind of information change that should be made in face of the new (possibly contradicting) information depicted by ψ. The new belief state, denoted $\mathcal{B} \circ \psi$, is usually characterized by the closest worlds to $\mathcal{B}$ in which ψ holds. This criterion, often called *the principle of minimal change*, is one of the most widely advocated postulates of belief revision theory. Clearly, it is derived by distance considerations, so it is not surprising that this consideration can be expressed in our framework. Indeed, the intended meaning of the revision operator is to describe 'how to revise $\mathcal{B}$ in order to be consistent with ψ'. In our context the revised belief state corresponds to the (coherent) set of conclusions that can be derived from the prioritized theory $\{\psi\} \cup \mathcal{B}$, in which ψ is superior than $\mathcal{B}$. Indeed, suppose again that $\mathfrak{D}$ is a normal $\mathcal{S}$-distance metric for some multiple-valued structure $\mathcal{S}$, and consider $\Gamma = \Gamma_1 \cup \Gamma_2 = \{\psi\} \cup \mathcal{B}$. Again, by Proposition 12, $\Delta_1^{\mathfrak{D}}(\Gamma) = mod(\psi)$, and so the new belief state consists of the formulas that are satisfied by every model of ψ and that are minimally distant (in terms of d_g) from $\mathcal{B}$. In other words,

$$\mathcal{B} \circ \psi \; = \; \Delta_2^{\mathfrak{D}}(\Gamma), \qquad (1)$$

[10] Or, equivalently, p is involved in contradictions in $\Gamma_{\mathcal{DB}}$; see also the discussion in Example 15 above.

where $\Gamma = \Gamma_1 \cup \Gamma_2$, $\Gamma_1 = \{\psi\}$, and $\Gamma_2 = \mathcal{B}$.

Example 33 For $\mathfrak{D}_2 = \langle \mathsf{TWO}, d_H, \min, \Sigma \rangle$ define a belief revision operator $\circ$ by Equation (1) above. The revision operator that is obtained is the same as the one considered in (Dalal 1988). It is well-known that this operator satisfies the AGM postulates (Alchourrón, Gärdenfors, & Makinson 1985).

Conclusion

In this paper we have introduced a family of multiple-valued entailments, the underlying semantics of which is based on distance considerations. It is shown that such entailments can be incorporated in a variety of deductive systems, mediators of distributed databases, consistent query answering engines, and formalisms for belief revision.

A characteristic property of the entailments considered here is that, although being paraconsistent in nature, to a large extent they retain consistency. For instance, the entailments that are defined by normal distance metrics in a two-valued (respectively, $\mathcal{S}$-valued) semantics, are identical to classical two-valued entailment (respectively, are identical to the corresponding basic $\mathcal{S}$-entailment), as long as the set of premises is kept consistent. Moreover, even when the set of premises becomes inconsistent, the conclusions that are obtained from the fragment of the theory that is not related to the 'core' of the inconsistency, are the same as those obtained by the classical two-valued (respectively, the basic $\mathcal{S}$-valued) entailment, when only the consistent fragment is taken into account. In contrast to the classical entailment, however, our formalisms are not degenerated in the presence of contradictions, so the set of conclusions is not 'exploded' is such cases.

References

Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision function. *Journal of Symbolic Logic* 50:510–530.

Arenas, M.; Bertossi, L.; and Chomicki, J. 1999. Consistent query answers in inconsistent databases. In *Proc. 18th Symp. on Principles of Database Systems* (PODS'99), 68–79.

Arenas, M.; Bertossi, L.; and Chomicki, J. 2003. Answer sets for consistent query answering in inconsistent databases. *Theory and Practice of Logic Programming* 3(4–5):393–424.

Arieli, O., and Avron, A. 1996. Reasoning with logical bilattices. *Journal of Logic, Language, and Information* 5(1):25–63.

Arieli, O., and Avron, A. 1998. The value of the four values. *Artificial Intelligence* 102(1):97–141.

Arieli, O., and Avron, A. 2000. Bilattices and paraconsistency. In Batens, D.; Mortenson, C.; Priest, G.; and Van Bendegem, J., eds., *Frontiers of Paraconsistent Logic*, volume 8 of *Studies in Logic and Computation*. Research Studies Press. 11–27.

Arieli, O., and Denecker, M. 2003. Reducing preferential paraconsistent reasoning to classical entailment. *Logic and Computation* 13(4):557–580.

Arieli, O.; Denecker, M.; Van Nuffelen, B.; and Bruynooghe, M. 2004. Database repair by signed formulae. In Seipel, D., and Turull Torres, J. M., eds., *Proc. 3rd Symp. on Foundations of Information and Knowledge Systems* (FoIKS'02), number 2942 in LNCS, 14–30. Springer.

Arieli, O.; Denecker, M.; Van Nuffelen, B.; and Bruynooghe, M. 2006. Computational methods for database repair by signed formulae. *Annals of Mathematics and Artificial Intelligence* 46(1–2):4–37.

Arieli, O. 2004. Paraconsistent preferential reasoning by signed quantified Boolean formulae. In de Mántaras, R., and Saitta, L., eds., *Proc. 16th European Conference on Artificial Intelligence* (ECAI'04), 773–777. IOS Press.

Arieli, O. 2006. Paraconsistent reasoning and preferential entailments by signed quantified Boolean formulae. *ACM Transactions on Computational Logic*. Accepted.

Avron, A. 1991. Natural 3-valued logics: Characterization and proof theory. *Journal of Symbolic Logic* 56(1):276–294.

Batens, D.; Mortenson, C.; Priest, G.; and Bendegem, J. V., eds. 2000. *Frontiers of Paraconsistent Logic*. Research Studies Press.

Batens, D. 1989. Dynamic dialectical logics. In Priest, G.; Routely, R.; and Norman, J., eds., *Paraconsistent Logic. Essay on the Inconsistent*. Philosophia Verlag. 187–217.

Batens, D. 1998. Inconsistency-adaptive logics. In Orlowska, E., ed., *Logic at Work*. Physica Verlag. 445–472.

Belnap, N. D. 1977a. How a computer should think. In Ryle, G., ed., *Contemporary Aspects of Philosophy*. Oriel Press. 30–56.

Belnap, N. D. 1977b. A useful four-valued logic. In Dunn, J. M., and Epstein, G., eds., *Modern Uses of Multiple-Valued Logics*. Reidel Publishing Company. 7–37.

Ben Naim, J. 2005. Preferential and preferential-discriminative consequence relations. *Logic and Computation* 15(3):263–294.

Bravo, L., and Bertossi, L. 2003. Logic programming for consistently querying data integration systems. In Gottlob, G., and Walsh, T., eds., *Proc. 18th Int. Joint Conference on Artificial Intelligence* (IJCAI'03), 10–15.

Carnielli, W. A.; Coniglio, M. E.; and Dóttaviano, I., eds. 2002. *Paraconsistency: The Logical Way to the Inconsistent*, volume 228 of *Lecture Notes in Pure and Applied Mathematics*. Marcel Dekker.

da Costa, N. C. A. 1974. On the theory of inconsistent formal systems. *Notre Dame Journal of Formal Logic* 15:497–510.

Dalal, M. 1988. Investigations into a theory of knowledge base revision. In *Proc. National Conference on Artificial Intelligence* (AAAI'98), 475–479. AAAI Press.

de Amo, S.; Carnielli, W. A.; and Marcos, J. 2002. A logical framework for integrating inconsistent information in multiple databases. In *Proc. 2nd Int. Symp. on Foundations of Information and Knowledge Systems* (FoIKS'02), number 2284 in LNCS, 67–84. Springer.

Delgrande, J. 2004. Preliminary considerations on the modelling of belief change operators by metric spaces. In *Proc. Int. Workshop on Non-Monotonic Reasoning* (NMR'04), 118–125.

D'ottaviano, I. 1985. The completeness and compactness of a three-valued first-order logic. *Revista Colombiana de Matematicas* XIX(1–2):31–42.

Eiter, T.; Fink, M.; Greco, G.; and Lembo, D. 2003. Efficient evaluation of logic programs for querying data integration systems. In *Proc. 19th Int. Conf. on Logic Programming* (ICLP'03), number 2916 in LNCS, 163–177. Springer.

Eiter, T. 2005. Data integration and answer set programming. In Baral, C.; Greco, G.; Leone, N.; and Terracina, G., eds., *Proc. 8th Int. Conf. on Logic Programming and Nonmonotonic Reasoning* (LPNMR'05), number 3662 in LNCS, 13–25. Springer.

Epstein, R. L. 1990. *The semantic foundation of logic. Vol.I: propositional logics*. Kluwer.

Fitting, M. 1990. Kleene's logic, generalized. *Logic and Computation* 1:797–810.

Greco, S., and Zumpano, E. 2000. Querying inconsistent databases. In *Proc. Int. Conf. on Logic Programming and Automated Reasoning* (LPAR'2000), number 1955 in LNAI, 308–325. Springer.

Hájek, P. 1998. *Metamatematics of Fuzzy Logic*. Kluwer.

Kleene, S. C. 1950. *Introduction to Metamathematics*. Van Nostrand.

Konieczny, S., and Marquis, P. 2002. Three-valued logics for inconsistency handling. In Flesca, S.; Greco, S.; Leone, N.; and Ianni, G., eds., *Proc. 8th European Conference on Logics in Artificial Intelligence* (JELIA'02), number 2424 in LNAI, 332–344. Springer.

Konieczny, S., and Pino Pérez, R. 2002. Merging information under constraints: a logical framework. *Logic and Computation* 12(5):773–808.

Konieczny, S.; Lang, J.; and Marquis, P. 2002. Distance-based merging: A general framework and some complexity results. In *Proc 8th Int. Conf. on Principles of Knowledge Representation and Reasoning* (KR'02), 97–108.

Lehmann, D., and Magidor, M. 1992. What does a conditional knowledge base entail? *Artificial Intelligence* 55:1–60.

Lehmann, D.; Magidor, M.; and Schlechta, K. 2001. Distance semantics for belief revision. *Journal of Symbolic Logic* 66(1):295–317.

Lin, J., and Mendelzon, A. O. 1999. Knowledge base merging by majority. In *Dynamic Worlds: From the Frame Problem to Knowledge Management*. Kluwer.

Makinson, D. 1994. General patterns in nonmonotonic reasoning. In Gabbay, D.; Hogger, C.; and Robinson, J., eds., *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3. Oxford Science Publications. 35–110.

Peppas, P.; Chopra, S.; and Foo, N. 2004. Distance semantics for relevance-sensitive belief revision. In *Proc 9th Int. Conf. on Principles of Knowledge Representation and Reasoning* (KR'04), 319–328. AAAI Press.

Priest, G. 1989. Reasoning about truth. *Artificial Intelligence* 39:231–244.

Priest, G. 1991. Minimally inconsistent LP. *Studia Logica* 50:321–331.

Priest, G. 2002. Paraconsistent logic. In Gabbay, D., and Guenthner, F., eds., *Handbook of Philosophical Logic*, volume 6. Kluwer. 287–393.

Reiter, R. 1978. On closed world databases. In *Logic and Databases*. Plenum Press. 55–76.

Rozoner, L. I. 1989. On interpretation of inconsistent theories. *Information Sciences* 47:243–266.

Shoham, Y. 1987. A semantical approach to nonmonotonic logics. In Ginsberg, M. L., ed., *Readings in Non-Monotonic Reasoning*. Morgan Kaufmann Publishers. 227–249.

Shoham, Y. 1988. *Reasoning About Change: Time and Causation from the Standpoint of Artificial Intelligence*. MIT Press.

Subrahmanian, V. S. 1994. Amalgamating knowledge bases. *ACM Transactions on Database Systems* 19(2):291–331.