Distance-Based Non-Deterministic Semantics

Ofer ARIELI [a], Anna ZAMANSKY [b,1]
[a] *Department of Computer Science, The Academic College of Tel-Aviv, Israel*
[b] *Department of Computer Science, Tel-Aviv University, Israel*

Abstract. Representing uncertainty and reasoning with dynamically evolving systems are two related issues that are in the heart of many information systems. In this paper we show that these tasks can be successfully dealt with by incorporating distance semantics and non-deterministic matrices. The outcome is a general framework for capturing the principle of minimal change and providing a non-deterministic view of the domain of discourse. We investigate some properties of the entailment relations that are induced by this framework and demonstrate their usability in some test-cases.

1. Introduction

Distance-based semantics is a common technique for reflecting the principle of minimal change in different scenarios where information is dynamically evolving, such as belief revision, integration of independent data sources, and planning systems. By nature, the underlying data in such cases is often incomplete or inconsistent. Yet, this fact is not always representable in terms of standard truth functions, and so other alternatives must be looked for. One such alternative is to borrow the idea of *non-deterministic computations* from automata and computability theory. This idea leads to a quest for structures, where the value assigned by a valuation to a complex formula might be chosen non-deterministically from a certain (non-empty) set of options (see [7]). The advantage of combining distance-based semantics and non-deterministic computations is demonstrated in the following example:

Example 1 Suppose that a reasoner wants to discover some properties of an unknown Boolean function (e.g., determine a sequence of inputs for which the output is known). The reasoner may have some idea on the structure of an electronic circuit that implements the unknown function, but this information may be partial or even unreliable. Two common problems in this respect are the following:

- It might happen that the behaviour of an electronic gate in the circuit is not coherent and therefore cannot be predicted (e.g., because of the presence of disturbing noise sources on or off chip).

[1]Supported by the Israel Science Foundation, grant No. 809–06.

- The reasoner may get conflicting evidence on the behaviour of the circuit from different sources (e.g., due to unreliable indicators, adversary third parties, erroneous communication among sources, etc.).

Situations like that of the first item can be handled by non-deterministic structures, and distance-based considerations are helpful to narrow the gap among contradictory sources as in the second item. However, any system aspiring to general intelligence should be able to deal with both of these types of uncertainty at the same time. For this, we introduce a general framework that combines non-determinism with distance semantics, demonstrate its usefulness by applying it to a number of case studies, and study some of the basic properties of the entailment relations that are obtained.

2. Preliminaries

2.1. Distance-Based Semantics

Distance semantics is a cornerstone behind many paradigms of handling incomplete or inconsistent information, such as belief revision [8,11,15,19] database integration systems [1,5,9,16], and social choice theory [13,18]. In [2,3] this approach is described in terms of entailment relations. The idea is simple: Given a distance function d on a space of valuations, reasoning with a given set of premises Γ is based on those valuations that are *'d-closest'* to Γ (called the *most plausible valuations* of Γ). For instance, it is intuitively clear that valuations in which q is true should be closer to $\Gamma = \{p, \neg p, q\}$ than valuations in which q is false, and so q should follow from Γ while $\neg q$ should *not* follow from Γ, although Γ is not consistent. The formal details are given in [2,3] and are sketched in what follows.

Suppose that $\mathcal{L}$ is a fixed propositional language with a finite set Atoms of atomic formulas. We denote by Γ a *finite multiset* of $\mathcal{L}$-formulas, for which $\mathsf{Atoms}(\Gamma)$ and $\mathsf{SF}(\Gamma)$ denote, respectively, the atomic formulas that occur in Γ and the subformulas of Γ. The set of models of Γ (that is, the valuations that satisfy every formula in Γ) is denoted $mod(\Gamma)$.

Definition 2 A *pseudo-distance* on a set U is a total function $d : U \times U \to \mathbb{R}^+$, satisfying the following conditions:

- *symmetry:* for all $\nu, \mu \in U$ $d(\nu, \mu) = d(\mu, \nu)$,
- *identity preservation:* for all $\nu, \mu \in U$ $d(\nu, \mu) = 0$ iff $\nu = \mu$.

A pseudo-distance d is a *distance* (metric) on U if it has the following property:

- *triangular inequality:* for all $\nu, \mu, \sigma \in U$ $d(\nu, \sigma) \leq d(\nu, \mu) + d(\mu, \sigma)$.

Example 3 It is easy to verify that the following two functions are distances on the space Λ_{Atoms} of two-valued valuations on Atoms:

- *The drastic distance:* $d_U(\nu, \mu) = 0$ if $\nu = \mu$ and $d_U(\nu, \mu) = 1$ otherwise.
- *The Hamming distance:* $d_H(\nu, \mu) = |\{p \in \mathsf{Atoms} \mid \nu(p) \neq \mu(p)\}|$.

Definition 4 A *numeric aggregation function* is total function f whose argument is a multiset of real numbers and whose values are real numbers, such that: (i) f is non-decreasing in the value of its argument, (ii) $f(\{x_1,\ldots,x_n\}) = 0$ iff $x_1 = x_2 = \ldots x_n = 0$, and (iii) $f(\{x\}) = x$ for every $x \in \mathbb{R}$.

Definition 5 Given a theory $\Gamma = \{\psi_1,\ldots,\psi_n\}$, a two-valued valuation $\nu \in \Lambda_{\mathsf{Atoms}}$, a pseudo-distance d and an aggregation function f, define:

- $d(\nu,\psi_i) = \begin{cases} \min\{d(\nu,\mu) \mid \mu \in mod(\psi_i)\} & \text{if } mod(\psi_i) \neq \emptyset, \\ 1 + \max\{d(\mu_1,\mu_2) \mid \mu_1,\mu_2 \in \Lambda_{\mathsf{Atoms}}\} & \text{otherwise.} \end{cases}$
- $\delta_{d,f}(\nu,\Gamma) = f(\{d(\nu,\psi_1),\ldots,d(\nu,\psi_n)\})$.

Note that in the two extreme degenerate cases, when ψ is either a tautology or a contradiction, all the valuations are equally distant from ψ. In any other case, the valuations that are closest to ψ are its models and their distance to ψ is zero. This also implies that $\delta_{d,f}(\nu,\Gamma) = 0$ iff $\nu \in mod(\Gamma)$ (see [3]).

Definition 6 The *most plausible valuations* of Γ (with respect to a pseudo distance d and an aggregation function f) are defined as follows:

$$\Delta_{d,f}(\Gamma) = \begin{cases} \{\nu \in \Lambda_{\mathsf{Atoms}} \mid \forall \mu \in \Lambda_{\mathsf{Atoms}}\ \delta_{d,f}(\nu,\Gamma) \leq \delta_{d,f}(\mu,\Gamma)\} & \text{if } \Gamma \neq \emptyset, \\ \Lambda_{\mathsf{Atoms}} & \text{otherwise.} \end{cases}$$

Definition 7 Denote: $\Gamma \models_{d,f} \psi$, if $\Delta_{d,f}(\Gamma) \subseteq mod(\psi)$. That is, conclusions should follow from *all* of the most plausible valuations of the premises.

Example 8 Consider $\Gamma = \{p, \neg p, q\}$ together with the Hamming distance and the summation function. By classical logic, everything follows from Γ, including $\neg q$. In contrast, since $\Delta_{d_H,\Sigma}(\Gamma)$ consists only of valuations in which q is true, we have that $\Gamma \models_{d_H,\Sigma} q$ while $\Gamma \not\models_{d_H,\Sigma} \neg q$, as intuitively expected.

2.2. Non-Deterministic Matrices

According to the classical principle of assigning truth values to formulas, the truth value assigned to a complex formula is uniquely determined by the truth values of its subformulas. This approach is no longer appropriate in the real world, were incomplete, imprecise or even inconsistent information is involved. To cope with this, Avron and Lev [7] introduced the notion of *non-deterministic matrices*, in which the value assigned by a valuation to a complex formula can be chosen *non-deterministically* out of a certain nonempty set of options. Below, we recall the basic definitions behind this approach.

Definition 9 A *non-deterministic matrix* (henceforth, *Nmatrix*) for a propositional language $\mathcal{L}$ is a tuple $\mathcal{M} = \langle \mathcal{V}, \mathcal{D}, \mathcal{O} \rangle$, where $\mathcal{V}$ is a non-empty set of truth values, $\mathcal{D}$ is a non-empty proper subset of $\mathcal{V}$, and for every n-ary connective $\diamond$ of $\mathcal{L}$, $\mathcal{O}$ includes an n-ary function $\widetilde{\diamond}$ from $\mathcal{V}^n$ to $2^{\mathcal{V}} - \{\emptyset\}$.

Definition 10 An $\mathcal{M}$-*valuation* is a function $\nu : \mathcal{L} \to \mathcal{V}$ that satisfies the following condition for every n-ary connective $\diamond$ of $\mathcal{L}$ and $\psi_1, \ldots, \psi_n \in \mathcal{L}$:

$$\nu(\diamond(\psi_1, \ldots, \psi_n)) \in \widetilde{\diamond}(v(\psi_1), \ldots, v(\psi_n)).$$

We denote by $\Lambda_{\mathcal{M}}$ the space of all the $\mathcal{M}$-valuations.

It is important to stress that in Nmatrices the truth-values assigned to $\psi_1, ..., \psi_n$ do not uniquely determine the truth-value assigned to $\diamond(\psi_1, \ldots, \psi_n)$, as ν makes a non-deterministic choice out of the set of options $\tilde{\diamond}(\nu(\psi_1), \ldots, \nu(\psi_n))$. Thus, the non-deterministic semantics is non-truth-functional, as opposed to the deterministic case.

Example 11 Let $\mathcal{M} = \langle \{t, f\}, \{t\}, \mathcal{O} \rangle$, where $\mathcal{O}$ consists of the following operators:

	$\neg$
t	$\{f\}$
f	$\{t\}$

		$\wedge$
t	t	$\{t, f\}$
t	f	$\{f\}$
f	t	$\{f\}$
f	f	$\{f\}$

Let $p, q \in \mathsf{Atoms}$ and $\nu_1, \nu_2 \in \Lambda_{\mathcal{M}}$, such that $\nu_1(p) = \nu_2(p) = \nu_1(q) = \nu_2(q) = t$. While ν_1 and ν_2 coincide on $\neg p$ and $\neg q$, their value for $p \wedge q$ may *not* be the same.

Definition 12 A valuation $\nu \in \Lambda_{\mathcal{M}}$ is a *model* of (or *satisfies*) a formula ψ in $\mathcal{M}$ (notation: $\nu \models_{\mathcal{M}} \psi$) if $\nu(\psi) \in \mathcal{D}$. ν is a *model* in $\mathcal{M}$ of a set Γ of formulas (notation: $\nu \models_{\mathcal{M}} \Gamma$) if it satisfies every formula in Γ. A formula ψ is $\mathcal{M}$-*satisfiable* if it is satisfied by a valuation in $\Lambda_{\mathcal{M}}$. ψ is an $\mathcal{M}$-*tautology* if it is satisfied by every valuation in $\Lambda_{\mathcal{M}}$.

Notation 13 Let $\mathcal{M}$ be an Nmatrix for $\mathcal{L}$, ψ a formula, and Γ a set of formulas in $\mathcal{L}$. Denote: $mod_{\mathcal{M}}(\psi) = \{\nu \in \Lambda_{\mathcal{M}} \mid \nu(\psi) \in \mathcal{D}\}$ and $mod_{\mathcal{M}}(\Gamma) = \cap_{\psi \in \Gamma}\, mod_{\mathcal{M}}(\psi)$.

Definition 14 The *consequence relation induced by the Nmatrix* $\mathcal{M}$ is defined by: $\Gamma \models_{\mathcal{M}} \varphi$ if $mod_{\mathcal{M}}(\Gamma) \subseteq mod_{\mathcal{M}}(\psi)$.

We note, finally, that the use of Nmatrices has the benefit of preserving all the advantages of logics with ordinary finite-valued semantics (in particular, decidability and compactness in the propositional case), while Nmatrices are applicable to a much larger family of logics (see [6,7]).

Henceforth, we concentrate on two-valued Nmatrices with $\mathcal{V} = \{t, f\}$ and $\mathcal{D} = \{t\}$, and denote by $\mathcal{M}$ such an Nmatrix.

3. Distance-Based Semantics for Non-Deterministic Matrices

In this section we generalize distance-based semantics to the context of Nmatrices. In doing so, we have to take into consideration some issues that follow from the non-deterministic character of our framework. The main problem is that, unlike the deterministic case, for computing distances between valuations it is no longer sufficient to consider their values on atomic formulas, since two valuations for an Nmatrix can agree on all the atoms of a formula, but still assign two different values to that formula (see Example 11). It follows that for computing distances between valuations one has to take into account also the truth-values assigned by those valuations to complex formulas. This implies that even under the assumption that the set of atoms is finite, there are infinitely many complex formulas to consider. To handle this, the distances computations are *context dependent*, that is: restricted to a certain set of relevant formulas. This allows us to generalize the notion of distance-based semantics to non-deterministic matrices as described in this section.

Definition 15 A *context* C is a finite set of $\mathcal{L}$-formulas that is closed under subformulas. Now,

- The *restriction to* C of a valuation $\nu \in \Lambda_{\mathcal{M}}$ is a valuation $\nu^{\downarrow\mathsf{C}}$ on C, such that $\nu^{\downarrow\mathsf{C}}(\psi) = \nu(\psi)$ for every ψ in C.
- The *restriction to* C of $\Lambda_{\mathcal{M}}$ is the set $\Lambda_{\mathcal{M}}^{\downarrow\mathsf{C}} = \{\nu^{\downarrow\mathsf{C}} \mid \nu \in \Lambda_{\mathcal{M}}\}$, that is, $\Lambda_{\mathcal{M}}^{\downarrow\mathsf{C}}$ consists of all the $\mathcal{M}$-valuations on C.

Example 16 Consider the following functions on $\Lambda_{\mathcal{M}}^{\downarrow\mathsf{SF}(\Gamma)} \times \Lambda_{\mathcal{M}}^{\downarrow\mathsf{SF}(\Gamma)}$:

- $d_U^{\downarrow\mathsf{SF}(\Gamma)}(\nu,\mu) = \begin{cases} 0 & \text{if } \nu(\psi) = \mu(\psi) \text{ for every } \psi \in \mathsf{SF}(\Gamma) \\ 1 & \text{otherwise} \end{cases}$
- $d_H^{\downarrow\mathsf{SF}(\Gamma)}(\nu,\mu) = |\{\psi \in \mathsf{SF}(\Gamma) \mid \nu(\psi) \neq \mu(\psi)\}|$

Proposition 17 *$d_U^{\downarrow\mathsf{SF}(\Gamma)}$ and $d_H^{\downarrow\mathsf{SF}(\Gamma)}$ are distance functions on $\Lambda_{\mathcal{M}}^{\downarrow\mathsf{SF}(\Gamma)}$.*

Note 18 Recall that the language $\mathcal{L}$ has a finite set Atoms of atomic formulas. Thus, the distance functions of Section 2.1 can be represented in the non-deterministic case as distance functions on $\Lambda_{\mathcal{M}_c}^{\downarrow\mathsf{Atoms}}$, where $\mathcal{M}_c$ is the Nmatrix for the language $\{\neg, \wedge, \vee, \rightarrow\}$ with the classical interpretations of the connectives (i.e., $\mathcal{M}_c$ is similar to the classical deterministic matrix, except that its valuation functions return singletons of truth-values instead of truth-values).

The next definition captures our intention to consider distances between partial $\mathcal{M}$-valuations (i.e., distances between $\mathcal{M}$-valuations modulo a given context).

Definition 19 Let $2^{\Lambda_{\mathcal{M}}} = \bigcup_{\{\mathsf{C}=\mathsf{SF}(\Gamma) \mid \Gamma \in 2^{\mathcal{L}}\}} \Lambda_{\mathcal{M}}^{\downarrow\mathsf{C}}$ and d a function on $2^{\Lambda_{\mathcal{M}}} \times 2^{\Lambda_{\mathcal{M}}}$.

- The *restriction* of d to a context C is a function $d^{\downarrow\mathsf{C}}$ on $\Lambda_{\mathcal{M}}^{\downarrow\mathsf{C}} \times \Lambda_{\mathcal{M}}^{\downarrow\mathsf{C}}$, defined for every $\nu, \mu \in \Lambda_{\mathcal{M}}^{\downarrow\mathsf{C}}$ by $d^{\downarrow\mathsf{C}}(\nu,\mu) = d(\nu,\mu)$.
- d is a *generic distance on* $\Lambda_{\mathcal{M}}$, if for every context C, $d^{\downarrow\mathsf{C}}$ is a (pseudo) distance on $\Lambda_{\mathcal{M}}^{\downarrow\mathsf{C}}$.

Example 20 Denote by $\mathsf{Dom}(\nu)$ the domain of a valuation $\nu \in 2^{\Lambda_{\mathcal{M}}}$ (that is, the formulas ψ in $\mathcal{L}$ for which $\nu(\psi)$ is defined). Now, consider the following functions on $2^{\Lambda_{\mathcal{M}}} \times 2^{\Lambda_{\mathcal{M}}}$:

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

- $d_H(\nu,\mu) = \begin{cases} |\{\psi \in \mathsf{Dom}(\nu) \mid \nu(\psi) \neq \mu(\psi)\}| & \text{if } \mathsf{Dom}(\nu) = \mathsf{Dom}(\mu) \\ \infty & \text{otherwise.} \end{cases}$

The restrictions of the two functions to $\mathsf{SF}(\Gamma)$ are given in Example 16. By Proposition 17, then, both of these functions are generic distances on $\Lambda_{\mathcal{M}}$ for every Nmatrix $\mathcal{M}$.

Note 21 In the notations of Example 20 (and Definition 19), the generic distances on $2^{\Lambda_{\mathcal{M}_c}}$ considered in Definition 3, are denoted $d_U^{\downarrow\mathsf{Atoms}}$ and $d_H^{\downarrow\mathsf{Atoms}}$.

Definition 22 A (distance-based, nondeterministic) *setting* for a language $\mathcal{L}$, is a triple $\mathcal{S} = \langle \mathcal{M}, d, f \rangle$, where $\mathcal{M}$ is a non-deterministic matrix for $\mathcal{L}$, d is a generic distance on $2^{\Lambda_{\mathcal{M}}}$, and f is an aggregation function.

The next three definitions are natural generalizations to the non-deterministic case of Definitions 5–7, respectively.

Definition 23 Given a setting $\mathcal{S} = \langle \mathcal{M}, d, f \rangle$ for a language $\mathcal{L}$, a valuation $\nu \in \Lambda_{\mathcal{M}}$, and a set $\Gamma = \{\psi_1, \ldots, \psi_n\}$ of formulas in $\mathcal{L}$, define:

- $d^{\downarrow\mathsf{SF}(\Gamma)}(\nu, \psi_i) =$

$$\begin{cases} \min\{d^{\downarrow\mathsf{SF}(\Gamma)}(\nu^{\downarrow\mathsf{SF}(\Gamma)}, \mu^{\downarrow\mathsf{SF}(\Gamma)}) \mid \mu \in mod_{\mathcal{M}}(\psi_i)\} & \text{if } mod_{\mathcal{M}}(\psi_i) \neq \emptyset, \\ 1 + \max\{d^{\downarrow\mathsf{SF}(\Gamma)}(\mu_1^{\downarrow\mathsf{SF}(\Gamma)}, \mu_2^{\downarrow\mathsf{SF}(\Gamma)}) \mid \mu_1, \mu_2 \in \Lambda_{\mathcal{M}}\} & \text{otherwise.} \end{cases}$$

- $\delta_{d,f}^{\downarrow\mathsf{SF}(\Gamma)}(\nu, \Gamma) = f(\{d^{\downarrow\mathsf{SF}(\Gamma)}(\nu, \psi_1), \ldots, d^{\downarrow\mathsf{SF}(\Gamma)}(\nu, \psi_n)\})$.

Definition 24 Given a setting $\mathcal{S} = \langle \mathcal{M}, d, f \rangle$, the *most plausible valuations* of Γ are defined as follows:

$$\Delta_{\mathcal{S}}(\Gamma) = \begin{cases} \{\nu \in \Lambda_{\mathcal{M}} \mid \forall \mu \in \Lambda_{\mathcal{M}}\ \delta_{d,f}^{\downarrow\mathsf{SF}(\Gamma)}(\nu, \Gamma) \leq \delta_{d,f}^{\downarrow\mathsf{SF}(\Gamma)}(\mu, \Gamma)\} & \text{if } \Gamma \neq \emptyset, \\ \Lambda_{\mathcal{M}} & \text{otherwise.} \end{cases}$$

Definition 25 For a setting $\mathcal{S} = \langle \mathcal{M}, d, f \rangle$, denote $\Gamma \models_{\mathcal{S}} \psi$ if $\Delta_{\mathcal{S}}(\Gamma) \subseteq mod_{\mathcal{M}}(\psi)$.

To give some examples of reasoning with $\models_{\mathcal{S}}$, we use the following notation:

Notation 26 Let Γ be a set of formulas, such that $\mathsf{SF}(\Gamma) = \{\psi_1, \psi_2, \ldots, \psi_n\}$. A valuation $\nu \in \Lambda_{\mathcal{M}}^{\downarrow\mathsf{SF}(\Gamma)}$ is represented by $\{\psi_1 : \nu(\psi_1), \psi_2 : \nu(\psi_2), \ldots, \psi_n : \nu(\psi_n)\}$.

Example 27 Let $\mathcal{S} = \langle \mathcal{M}, d_U, \Sigma \rangle$, where $\mathcal{M}$ is the Nmatrix considered in Example 11. Let $\Gamma = \{p, \neg p, q, \neg(p \wedge q)\}$. Then:

$$\Delta_{\mathcal{S}}(\Gamma) = \left\{ \begin{array}{l} \{p : t,\ \neg p : f,\ q : t,\ p \wedge q : f,\ \neg(p \wedge q) : t\}, \\ \{p : f,\ \neg p : t,\ q : t,\ p \wedge q : f,\ \neg(p \wedge q) : t\} \end{array} \right\}.$$

Thus, $\Gamma \models_{\mathcal{S}} q$ and $\Gamma \models_{\mathcal{S}} \neg(p \wedge q)$, while $\Gamma \not\models_{\mathcal{S}} p$ and $\Gamma \not\models_{\mathcal{S}} \neg p$.

Example 1, Revisited

Let's consider again Example 1. Suppose that the reasoner's information about the circuit is given below:

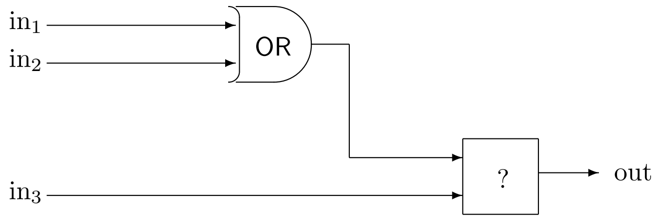

The reasoner knows, then, the structure of the circuit and that it contains two gates, one of which is an OR gate. In addition, the second gate does not behave coherently when its two inputs are not equal. In this case, one can use an Nmatrix $\mathcal{M}$ for the language $\mathcal{L} = \{\neg, \vee, \rightarrow, \Diamond\}$, in which $\vee, \neg, \rightarrow$ are interpreted standardly, and $\Diamond$ is evaluated as follows:

$\Diamond$	t	f
t	$\{t\}$	$\{t, f\}$
f	$\{t, f\}$	$\{f\}$

Accordingly, a representation of the circuit above may be given by:

$$\Gamma_0 = \{out \leftrightarrow (in_1 \vee in_2) \Diamond in_3\}$$

The reasoner may also consult with several external sources for detecting (some properties of) the underlying function. Let's suppose, for instance, that there are two other sources whose indications are represented by the following theories:

$$\Gamma_1 = \{(\neg in_1 \wedge \neg in_2) \rightarrow out, \quad \neg in_3 \rightarrow out\},$$
$$\Gamma_2 = \{(in_1 \vee in_2) \rightarrow \neg out, \quad in_3 \rightarrow \neg out\}.$$

Note that each source is consistent ($\mathcal{M}$-satisfiable) with the reasoner's belief as represented by Γ_0, but taken altogether, the indications of the sources contradict what the reasoner believes about the circuit. A proper way of integrating the information above may be in the context of non-deterministic distance-based semantics, using the following natural extension of Definitions 22 and 23:

Definition 28 An *extended setting* is a quadruple $\mathcal{S} = \langle \mathcal{M}, d, f, g \rangle$, where $\langle \mathcal{M}, d, f \rangle$ is a setting (in the sense of Definition 22) and g is an aggregation function.[2] Given a set $\overline{\Gamma} = \{\Gamma_1, \ldots, \Gamma_n\}$ of n finite theories and a valuation $\nu \in \Lambda_{\mathcal{M}}$, define:

$$\delta_{d,f,g}^{\downarrow \mathsf{SF}(\overline{\Gamma})}(\nu, \overline{\Gamma}) = g\{\delta_{d,f}^{\downarrow \mathsf{SF}(\overline{\Gamma})}(\nu, \Gamma_1), \ldots, \delta_{d,f}^{\downarrow \mathsf{SF}(\overline{\Gamma})}(\nu, \Gamma_n)\}.$$

[2] Intuitively, f is used for computing distances inside each source, while g aggregates distances among different sources.

The most plausible valuations of (the integration of the elements in) $\overline{\Gamma}$ are defined, like before, by $\Delta_{\mathcal{S}}(\overline{\Gamma}) = \{\nu \in \Lambda_{\mathcal{M}} \mid \forall \mu \in \Lambda_{\mathcal{M}}\ \delta_{d,f,g}^{\downarrow \mathsf{SF}(\overline{\Gamma})}(\mu, (\overline{\Gamma})) \leq \delta_{d,f,g}^{\downarrow \mathsf{SF}(\overline{\Gamma})}(\nu, (\overline{\Gamma}))\}$.

The entailment relation that is induced by $\mathcal{S}$ is now defined as follows:

$$\overline{\Gamma} \models_{\mathcal{S}} \psi \text{ iff } \Delta_{\mathcal{S}}(\overline{\Gamma}) \subseteq mod_{\mathcal{M}}(\psi).$$

Consider now the extended setting $\mathcal{S} = \langle \mathcal{M}, d_U, \Sigma, \Sigma \rangle$. In our example, the relevant distances of the valuations for $\overline{\Gamma} = \{\Gamma_0, \Gamma_1, \Gamma_2\}$ are given below, where $\psi = (in_1 \vee in_2) \diamond in_3$, and $\delta(\nu, \Gamma_i)$ abbreviates $\delta_{d,f}^{\downarrow \mathsf{SF}(\overline{\Gamma})}(\nu, \Gamma_i)$, for $i = 0, 1, 2$.

	in_1	in_2	in_3	out	ψ	$\delta(\nu,\Gamma_0)$	$\delta(\nu,\Gamma_1)$	$\delta(\nu,\Gamma_2)$	$\delta(\nu,\overline{\Gamma})$
ν_1	t	t	t	t	t	0	0	2	2
ν_2	t	t	t	f	t	1	0	0	1
ν_3	t	t	f	t	t	0	0	1	1
ν_4	t	t	f	t	f	1	0	2	3
ν_5	t	t	f	f	t	1	1	0	2
ν_6	t	t	f	f	f	0	1	0	1
ν_7	t	f	t	t	t	0	0	2	2
ν_8	t	f	t	f	t	1	0	0	1
ν_9	t	f	f	t	t	0	0	1	1
ν_{10}	t	f	f	t	f	1	0	1	2
ν_{11}	t	f	f	f	t	1	1	0	2
ν_{12}	t	f	f	f	f	0	1	0	1
ν_{13}	f	t	t	t	t	0	0	2	2
ν_{14}	f	t	t	f	t	1	0	0	1
ν_{15}	f	t	f	t	t	0	0	1	1
ν_{16}	f	t	f	t	f	1	0	1	2
ν_{17}	f	t	f	f	t	1	1	0	2
ν_{18}	f	t	f	f	f	0	1	0	1
ν_{19}	f	f	t	t	t	0	0	1	1
ν_{20}	f	f	t	t	f	1	0	1	2
ν_{21}	f	f	t	f	t	1	1	0	2
ν_{22}	f	f	t	f	f	0	1	0	1
ν_{23}	f	f	f	t	f	1	0	0	1
ν_{24}	f	f	f	f	f	1	2	0	3

Thus, out of the 24 possible valuations in this case, 12 are the most plausible ones. The reasoner may conclude, then, that when all the three input lines have the same value, the output line has the opposite value. This conclusion, which is not consistent with the original belief of the reasoner about the circuit, exemplifies how inconsistency is maintained in our framework. This is further addressed in the next section.

4. Reasoning with $\models_{\mathcal{S}}$

In this section, we consider some basic properties of the entailments that are induced by our framework.[3] First, we show the relations between standard non-

[3] Due to a lack of space, some proofs are omitted or outlined.

deterministic entailments (Definition 14) and distance-based ones (Definition 25). Below, unless otherwise stated, $\mathcal{S} = \langle \mathcal{M}, d, f \rangle$ is a general setting with an Nmartix $\mathcal{M}$, a pseudo distance d, and an aggregation function f.

Proposition 29 *If* Γ *is* $\mathcal{M}$*-satisfiable, then* $\Gamma \models_{\mathcal{M}} \psi$ *iff* $\Gamma \models_{\mathcal{S}} \psi$.

Proposition 29 immediately follows from the following result:

Proposition 30 Γ *is* $\mathcal{M}$*-satisfiable iff* $\Delta_{\mathcal{S}}(\Gamma) = mod_{\mathcal{M}}(\Gamma)$.

Proof. If $\Delta_{\mathcal{S}}(\Gamma) = mod_{\mathcal{M}}(\Gamma)$ then since $\Delta_{\mathcal{S}}(\Gamma)$ consists of minimal elements over a finite set, it is never empty, and so $mod^{\mathcal{M}}(\Gamma) \neq \emptyset$. Thus Γ is $\mathcal{M}$-satisfiable. The converse follows from the fact that, as $d^{\downarrow \mathsf{SF}(\Gamma)}$ is a pseudo-distance on $\Lambda_{\mathcal{M}}^{\downarrow \mathsf{SF}(\Gamma)}$, $d^{\downarrow \mathsf{SF}(\Gamma)}(\nu, \psi) = 0$ iff $\nu \models_{\mathcal{M}} \psi$ and $\delta_{d,f}^{\downarrow \mathsf{SF}(\Gamma)}(\nu, \Gamma) = 0$ iff $\nu \models_{\mathcal{M}} \Gamma$. Thus,

$$\begin{aligned} \nu \in mod_{\mathcal{M}}(\Gamma) \iff & \delta_{d,f}^{\downarrow \mathsf{SF}(\Gamma)}(\nu, \Gamma) = 0, \\ \iff^{4} & \forall \mu \in \Lambda_{\mathcal{M}}\ \delta_{d,f}^{\downarrow \mathsf{SF}(\Gamma)}(\nu, \Gamma) \leq \delta_{d,f}^{\downarrow \mathsf{SF}(\Gamma)}(\mu, \Gamma), \\ \iff & \nu \in \Delta_{\mathcal{S}}(\Gamma). \quad \Box \end{aligned}$$

Next, we check to what extent the entailment relations of our framework are paraconsistent (that is, inconsistent information is tolerated in a non-trivial way), and non-monotonic (conclusions may be revised). Both these properties are related to the question of handling contradictory information. A common way of representing such information is by the standard negation operator.

Definition 31 We say that $\mathcal{M} = \langle \{t, f\}, \{t\}, \mathcal{O} \rangle$ is an Nmatrix *with negation*, if there is a unary function $\tilde{\neg}$ in $\mathcal{O}$ such that $\tilde{\neg}(t) = \{f\}$ and $\tilde{\neg}(f) = \{t\}$.

Distance-based entailments that correspond to Nmatrices with negation preserve the consistency of their conclusions:

Proposition 32 *Let* $\mathcal{S} = \langle \mathcal{M}, d, f \rangle$ *be a setting, where* $\mathcal{M}$ *is with negation. Then for every* Γ *and every* ψ*, if* $\Gamma \models_{\mathcal{S}} \psi$ *then* $\Gamma \not\models_{\mathcal{S}} \neg\psi$.

Proof. Suppose that there is a formula ψ such that $\Gamma \models_{\mathcal{S}} \psi$ and $\Gamma \models_{\mathcal{S}} \neg\psi$. Then $\Delta_{\mathcal{S}}(\Gamma) \subseteq mod_{\mathcal{M}}(\psi)$ and $\Delta_{\mathcal{S}}(\Gamma) \subseteq mod_{\mathcal{M}}(\neg\psi)$. But $mod_{\mathcal{M}}(\psi) \cap mod_{\mathcal{M}}(\neg\psi) = \emptyset$, and so $\Delta_{\mathcal{S}}(\Gamma) = \emptyset$, a contradiction to the fact that $\Delta_{\mathcal{S}}(\Gamma) \neq \emptyset$ for every Γ. $\Box$

Corollary 33 *Let* $\mathcal{S} = \langle \mathcal{M}, d, f \rangle$ *be a setting such that* $\mathcal{M}$ *is with negation. Then* $\models_{\mathcal{S}}$ *is non-monotonic.*

Proof. Clearly, $p \models_{\mathcal{S}} p$ and $\neg p \models_{\mathcal{S}} \neg p$. By Proposition 32, on the other hand, either $p, \neg p \not\models_{\mathcal{S}} p$ or $p, \neg p \not\models_{\mathcal{S}} \neg p$ (or both). Hence, the set of conclusions does not monotonically grow with respect to the size of the premises, and so $\models_{\mathcal{S}}$ is non-monotonic. $\Box$

Definition 34 A consequence relation $\models$ is *weakly paraconsistent* if for every theory Γ there is some ψ such that $\Gamma \not\models \psi$.

[4]Note that the direction $\Leftarrow$ follows from the fact that Γ is $\mathcal{M}$-satisfiable.

Corollary 35 *For every setting $\mathcal{S} = \langle \mathcal{M}, d, f \rangle$ where $\mathcal{M}$ is with negation, $\models_{\mathcal{S}}$ is weakly paraconsistent.*

Proof. Consider a theory Γ and a formula ψ. If $\Gamma \not\models_{\mathcal{S}} \psi$ we are done. Otherwise, by Proposition 32, $\Gamma \not\models_{\mathcal{S}} \neg\psi$. $\square$

Next we define a family of settings in which one can show stronger results.

Definition 36 A setting $\mathcal{S} = \langle \mathcal{M}, d, f \rangle$ is *unbiased* if for every $\nu_1, \nu_2 \in \Lambda_{\mathcal{M}}$, every context $\mathsf{C} = \mathsf{SF}(\Gamma)$, and every $\psi \in \Gamma$, if $\nu_1(\varphi) = \nu_2(\varphi)$ for every $\varphi \in \mathsf{C}$, then $d^{\downarrow \mathsf{C}}(\nu_1, \psi) = d^{\downarrow \mathsf{C}}(\nu_2, \psi)$.

Intuitively, unbiasedness means that distances between valuations and formulas are not affected (biased) by irrelevant formulas (those that are not part of the relevant context).

Example 37 Clearly, $\langle \mathcal{M}, d_H, \Sigma \rangle$ is unbiased for every Nmatrix $\mathcal{M}$. It is also easy to verify that for any Nmatrix $\mathcal{M}$ and every aggregation function f, $\langle \mathcal{M}, d_U, f \rangle$ is unbiased. More generally, it is possible to show that every uniform setting is unbiased, where $\mathcal{S} = \langle \mathcal{M}, d, f \rangle$ is called uniform if for every context $\mathsf{C} = \mathsf{SF}(\Gamma)$ there is a $k_{\mathsf{C}} > 0$, such that for all $\psi \in \Gamma$ and $\nu \in \Lambda_{\mathcal{M}}$,

$$d^{\downarrow \mathsf{C}}(\nu, \psi) = \begin{cases} 0 & \nu \in mod^{\mathcal{M}}(\psi), \\ k_{\mathsf{C}} & \text{otherwise.} \end{cases}$$

Unbiased settings satisfy a stronger notion of paraconsistency than that of Corollary 35: As Proposition 39 shows, even if a theory is not consistent, it does not entail any irrelevant non-tautological formula.

Definition 38 Two sets of formulas Γ_1 and Γ_2 are called *independent* (or disjoint), if $\mathsf{Atoms}(\Gamma') \cap \mathsf{Atoms}(\Gamma'') = \emptyset$.

Proposition 39 *Let $\mathcal{S} = \langle \mathcal{M}, d, f \rangle$ be an unbiased setting. For every Γ and every ψ such that Γ and $\{\psi\}$ are independent, $\Gamma \models_{\mathcal{S}} \psi$ iff ψ is an $\mathcal{M}$-tautology.*

Proof. One direction is clear: if ψ is an $\mathcal{M}$-tautology, then for every $\nu \in \Delta_{\mathcal{S}}(\Gamma)$, $\nu(\psi) = t$ and so $\Gamma \models_{\mathcal{S}} \psi$. For the converse, suppose that ψ is not an $\mathcal{M}$-tautology. Then there is some $\mathcal{M}$-valuation σ, such that $\sigma(\psi) = f$. Let $\nu \in \Delta_{\mathcal{S}}(\Gamma)$. If $\nu(\psi) = f$, we are done. Otherwise, since Γ and $\{\psi\}$ are independent, there is an $\mathcal{M}$-valuation μ such that $\mu(\varphi) = \nu(\varphi)$ for every $\varphi \in \mathsf{SF}(\Gamma)$ and $\mu(\varphi) = \sigma(\varphi)$ for $\varphi \in \mathsf{SF}(\psi)$. Since $\mathcal{S}$ is unbiased, $d^{\downarrow \mathsf{SF}(\Gamma)}(\nu, \gamma) = d^{\downarrow \mathsf{SF}(\Gamma)}(\mu, \gamma)$ for every $\gamma \in \Gamma$. Thus, $\delta_{d,f}^{\downarrow \mathsf{SF}(\Gamma)}(\nu, \Gamma) = \delta_{d,f}^{\downarrow \mathsf{SF}(\Gamma)}(\mu, \Gamma)$ and $\mu \in \Delta_{\mathcal{S}}(\Gamma)$. But $\mu(\psi) = \sigma(\psi) = f$ and so $\Gamma \not\models_{\mathcal{S}} \psi$. $\square$

Corollary 40 *If $\mathcal{S}$ is unbiased, then $\models_{\mathcal{S}}$ is weakly paraconsistent.*

Proof. Given a set Γ of formulas, let $p \in \mathsf{Atoms} \setminus \mathsf{SF}(\Gamma)$ (if there is no such an atom, extend the language with a new atomic symbol p). As Γ and $\{p\}$ are independent, by Proposition 39, $\Gamma \not\models_{\mathcal{S}} p$. $\square$

Regarding the non-monotonic nature of the $\models_{\mathcal{S}}$-entailments, it turns out that in spite of Corollary 33, in unbiased settings one can specify conditions under which the entailment relations have some monotonic characteristics. Next we consider such cases. For this, we need the following property of aggregation functions:

Definition 41 An aggregation function is called *hereditary*, if $f(\{x_1, ..., x_n\}) < f(\{y_1, ..., y_n\})$ entails that $f(\{x_1, ..., x_n, z_1, ..., z_m\}) < f(\{y_1, ..., y_n, z_1, ..., z_m\})$.

Example 42 The aggregation function Σ is hereditary, while max is not.

The following proposition shows that in unbiased settings, in light of new information that is unrelated to the premises, previously drawn conclusions should not be retracted.[5]

Proposition 43 *Let $\mathcal{S} = \langle \mathcal{M}, d, f \rangle$ be an unbiased setting in which f is hereditary. If $\Gamma \models_{\mathcal{S}} \psi$, then $\Gamma, \phi \models_{\mathcal{S}} \psi$ for every ϕ such that $\Gamma \cup \{\psi\}$ and $\{\phi\}$ are independent.*

The discussion above, on the non-monotonicity of $\models_{\mathcal{S}}$, brings us to the question to what extent these entailments can be considered as consequence relations.

Definition 44 A Tarskian *consequence relation* [20] for a language $\mathcal{L}$ is a binary relation $\vdash$ between sets of formulas of $\mathcal{L}$ and formulas of $\mathcal{L}$ that satisfies the following conditions:

Reflexivity:	if $\psi \in \Gamma$, then $\Gamma \vdash \psi$.
Monotonicity:	if $\Gamma \vdash \psi$ and $\Gamma \subseteq \Gamma'$, then $\Gamma' \vdash \psi$.
Transitivity:	if $\Gamma \vdash \psi$ and $\Gamma', \psi \vdash \varphi$, then $\Gamma, \Gamma' \vdash \varphi$.

As follows from Example 27 and Corollary 33, entailments of the form $\models_{\mathcal{S}}$ are, in general, neither reflexive nor monotonic. To see that transitivity may not hold either, consider the propositional language and $\mathcal{S} = \langle M_c, d, f \rangle$ for any d and f.[6] If $p, \neg p \not\models_{\mathcal{S}} q$, transitivity is falsified since, by Proposition 29, $p \models_{\mathcal{S}} \neg p \to q$ and $\neg p, \neg p \to q \models_{\mathcal{S}} q$; Otherwise, if $p, \neg p \models_{\mathcal{S}} q$, then by Proposition 32, $p, \neg p \not\models_{\mathcal{S}} \neg q$, and this, together with the facts that $p \models_{\mathcal{S}} \neg p \to \neg q$ and $\neg p, \neg p \to \neg q \models_{\mathcal{S}} \neg q$ (Proposition 29 again) provide a counterexample for transitivity.

In the context of non-monotonic reasoning, however, it is usual to consider the following weaker conditions that guarantee a 'proper behaviour' of nonmonotonic entailments in the presence of inconsistency (see, e.g., [4,12,14,17]):

Definition 45 A *cautious* consequence relation for $\mathcal{L}$ is a relation $\mid\!\sim$ between sets of $\mathcal{L}$-formulas and $\mathcal{L}$-formulas, that satisfies the following conditions:

Cautious Reflexivity:	if Γ is $\mathcal{M}$-satisfiable and $\psi \in \Gamma$, then $\Gamma \mid\!\sim \psi$.
Cautious Monotonicity [10]:	if $\Gamma \mid\!\sim \psi$ and $\Gamma \mid\!\sim \phi$, then $\Gamma, \psi \mid\!\sim \phi$.
Cautious Transitivity [12]:	if $\Gamma \mid\!\sim \psi$ and $\Gamma, \psi \mid\!\sim \phi$, then $\Gamma \mid\!\sim \phi$.

[5]This type of monotonicity is kind of *rational monotonicity*, considered in detail in [14].
[6]$\mathcal{M}_c$ is the matrix considered in Note 18.

Proposition 46 *Let $\mathcal{S} = \langle \mathcal{M}, d, f \rangle$ be an unbiased setting in which f is hereditary. Then $\models_{\mathcal{S}}$ is a cautious consequence relation.*

Proof. A simple generalization to the non-deterministic case of the proof in [3]. $\Box$

References

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

[2] O. Arieli. Commonsense reasoning by distance semantics. In *Proc. TARK'07*, pages 33–41, 2007.

[3] O. Arieli. Distance-based paraconsistent logics. *International Journal of Approximate Reasoning*, 2008. Accepted.

[4] O. Arieli and A. Avron. General patterns for nonmonotonic reasoning: from basic entailments to plausible relations. *Logic Journal of the IGPL*, 8(2):119–148, 2000.

[5] O. Arieli, M. Denecker, and M. Bruynooghe. Distance semantics for database repair. *Annals of Mathematics and Artificial Intelligence*, 50(3–4):389–415, 2007.

[6] A Avron. Non-deterministic semantics for logics with a consistency operator. *International Journal of Approximate Reasoning*, 45:271–287, 2007.

[7] A. Avron and I. Lev. Non-deterministic multi-valued structures. *Journal of Logic and Computation*, 15:241–261, 2005.

[8] J. Ben Naim. Lack of finite characterizations for the distance-based revision. In *Proc. KR'06*, pages 239–248, 2006.

[9] J. Chomicki and J. Marchinkowski. Minimal-change integrity maintenance using tuple deletion. *Information and Computation*, 197(1–2):90–121, 2005.

[10] D. Gabbay. Theoretical foundation for non-monotonic reasoning, Part II: Structured non-monotonic theories. In *Proc. SCAI'91*. IOS Press, 1991.

[11] A. Grove. Two modellings for theory change. *Philosophical Logic*, 17:157–180, 1988.

[12] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1–2):167–207, 1990.

[13] C. Lafage and J. Lang. Propositional distances and preference representation. In *Proc. ECSQARU'01*, LNAI 2143, pages 48–59. Springer, 2001.

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

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

[16] A. Lopatenko and L. Bertossi. Complexity of consistent query answering in databases under cardinality-based and incremental repair semantics. In *Proc. ICDT'07*, LNCS 4353, pages 179–193. Springer, 2007.

[17] D. Makinson. General patterns in nonmonotonic reasoning. In D. Gabbay, C. Hogger, and J. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3, pages 35–110. 1994.

[18] G. Pigozzi. Two aggregation paradoxes in social decision making: the ostrogorski paradox and the discursive dilemma. *Episteme*, 2(2):33–42, 2005.

[19] W. Spohn. Ordinal conditional functions: a dynamic theory of epistemic states. In *Belief Change and Statistics*, volume II, pages 105–134. Kluwer Academic Pubnlishers, 1988.

[20] A. Tarski. *Introduction to Logic*. Oxford University Press, 1941.