Reasoning with Prioritized Data by Aggregation of Distance Functions

Ofer ARIELI

Department of Computer Science, The Academic College of Tel-Aviv, Israel

Abstract. We introduce a general framework for reasoning with prioritized data by aggregation of distance functions, study some basic properties of the entailment relations that are obtained, and relate them to other approaches for maintaining uncertain information.

1. Introduction

Reasoning with prioritized data is at the heart of many information systems. Eminent examples for this are, e.g., database systems, where integrity constraints are superior to raw data [1,2,3], ranked knowledge-bases, where information is graded according to its reliability or accuracy [4,5,6], and (iterated) belief revision, where more recent data gets higher precedence than older one [7,8,9,10]. There is no wonder, therefore, that reasoning with prioritized information is a cornerstone of many formalisms for maintaining uncertainty, such as annotated logic [11], possibilistic logic [12], and System Z [13].

In this paper we handle prioritized data by a possible-world semantics, derived by distance considerations. To illustrate this, consider the following example:

Example 1 Let Γ be a set of formulas, consisting of the following subtheories:

$$\Gamma_1 = \{\, \texttt{bird(x)} \to \texttt{fly(x)},\ \texttt{color\_of(Tweety,Red)} \,\},$$
$$\Gamma_2 = \{\, \texttt{bird(Tweety)},\ \texttt{penguin(Tweety)} \,\},$$
$$\Gamma_3 = \{\, \texttt{penguin(x)} \to \neg\texttt{fly(x)} \,\}.$$

Intuitively, Γ is a theory with three priority levels, where precedence is given to formulas that belong to subtheories with higher indices, that is, for $1 \le i < j \le 3$, each formula in Γ_j is considered more important (or more reliable) than the formulas in Γ_i.

A justification for the representation above may be the following: the highest level (Γ_3) consists of integrity constraints that should not be violated. In our case, the single rule in this level specifies that a characteristic property of penguins is that they cannot fly, and there are no exceptions for that. The intermediate level (Γ_2) contains some known facts about the domain of discourse, and the lowest level (Γ_1) consists of default assumptions about this domain (in our case, a bird can fly unless otherwise stated), and facts with lower certainty.

Note that as a 'flat' set of assertions (i.e., when all the assertions in Γ have the same priority), this theory is classically inconsistent, therefore everything follows from it, and so the theory is useless. However, as Γ is prioritized, one would like to draw the following conclusions from it:

1. Conclude $\texttt{bird}(\texttt{Tweety})$ and $\texttt{penguin}(\texttt{Tweety})$ (but not their negations), as these facts are explicitly stated in a level that is consistent with the higher priority levels.
2. Conclude $\neg\texttt{fly}(\texttt{Tweety})$ (and do not conclude $\texttt{fly}(\texttt{Tweety})$), as this fact follows from the two top priority levels, while its complement is inferrable by a lower level (which, moreover, is inconsistent with the higher levels).
3. Conclude $\texttt{color\_of}(\texttt{Tweety}, \texttt{Red})$ (but not its negation), since although this fact appears in the lowest level of Γ, and that level is inconsistent with the other levels, it is not contradicted by any consistent fragment of Γ, so there is no reason to believe that $\texttt{color\_of}(\texttt{Tweety}, \texttt{Red})$ does not hold.

The kind of reasoning described above is obtained in our framework by the following two principles:

- A distance-based preference relation is defined on the space of interpretations, so that inferences are drawn according to the most preferred interpretations. In the example above, for instance, interpretations in which $\texttt{fly}(\texttt{Tweety})$ is false will be 'closer' to Γ, and so more plausible, than interpretations in which $\texttt{fly}(\texttt{Tweety})$ is true, (thus item (2) above is obtained).
- Priorities are considered as extra-logical data that is exploited by an iterative process that first computes interpretations that are as close as possible to higher-levelled subtheories, and then makes preference among those interpretations according to their closeness to lower-levelled subtheories.

The goal of our work is to examine these principles and to show that they reflect a variety of methods for maintaining imprecise information. [1]

2. Distance-based Entailments for Prioritized Theories

Let $\mathcal{L}$ be a propositional language with a finite set Atoms of atomic formulas. The space of all the two-valued interpretations on Atoms is denoted by Λ_{Atoms}. A *theory* Γ in $\mathcal{L}$ is a (possibly empty) finite multiset of formulas in $\mathcal{L}$. The set of atomic formulas that occur in the formulas of Γ is denoted $\mathsf{Atoms}(\Gamma)$ and the set of models of Γ (that is, the interpretations on $\mathsf{Atoms}(\Gamma)$ in which every formula in Γ is true) is denoted $mod(\Gamma)$.

Definition 2 An *n-prioritized theory* is a theory $\Gamma^{\langle n\rangle}$ in $\mathcal{L}$, partitioned into $n \geq 1$ pairwise disjoint sub-theories Γ_i ($1 \leq i \leq n$). Notation: $\Gamma^{\langle n\rangle} = \Gamma_1 \oplus \Gamma_2 \oplus \ldots \oplus \Gamma_n$.

In what follows we shall usually write Γ instead of $\Gamma^{\langle n\rangle}$. Intuitively, formulas in higher levels are preferred over those in lower levels, so if $1 \leq i < j \leq n$ then a formula $\psi \in \Gamma_j$ overtakes any formula $\phi \in \Gamma_i$. Note that in this writing the precedence is *righthand increasing*.

[1] Due to a lack of space some proofs are reduced or omitted.

Definition 3 Let $\Gamma = \Gamma_1 \oplus \ldots \oplus \Gamma_n$ be an n-prioritized theory.

- For any $1 \leq i \leq n$, denote the $n-i+1$ highest levels of Γ by $\Gamma_{\geq i}$, that is, $\Gamma_{\geq i} = \Gamma_i \oplus \ldots \oplus \Gamma_n$.
- Denote by $\overline{\Gamma}_{\geq i}$ the 'flat' (1-prioritized) theory, obtained by taking the union of the priority levels in $\Gamma_{\geq i}$, that is, $\overline{\Gamma}_{\geq i} = \Gamma_i \cup \ldots \cup \Gamma_n$. Also, $\overline{\Gamma} = \overline{\Gamma}_{\geq 1}$.
- The *consistency level* con of Γ is the minimal value $i \leq n$ such that $\overline{\Gamma}_{\geq i}$ is consistent. If there is no such value, let $\mathsf{con} = n+1$ (and then $\Gamma_{\geq \mathsf{con}} = \emptyset$).

Definition 4 A total function $d : U \times U \to \mathbb{R}^+$ is called *pseudo distance* on U if it satisfies the following two properties:

Symmetry: $\forall u, v \in U \; d(u,v) = d(v,u)$.

Identity Preservation: $\forall u, v \in U \; d(u,v) = 0$ iff $u = v$.

A *distance function* on U is a pseudo distance on U with the following property:

Triangulation: $\forall u, v, w \in U \; d(u,v) \leq d(u,w) + d(w,v)$.

Example 5 The following two functions are distances on Λ_{Atoms}.

- *The drastic distance:* $d_U(\nu, \mu) = 0$ if $\nu = \mu$ and $d_U(\nu, \mu) = 1$ otherwise.
- *The Hamming distance:* $d_H(\nu, \mu) = |\{p \in \mathsf{Atoms} \mid \nu(p) \neq \mu(p)\}|$. [2]

Definition 6 A *numeric aggregation function* f is a total function that accepts a multiset of real numbers and returns a real number. In addition, f is non-decreasing in the values of its argument[3], $f(\{x_1, \ldots, x_n\}) = 0$ iff $x_1 = \ldots = x_n = 0$, and $\forall x \in \mathbb{R} \; f(\{x\}) = x$.

Definition 7 An aggregation function is called *hereditary*, if $f(\{x_1, ..., x_n\}) < f(\{y_1, ..., y_n\})$ implies that $f(\{x_1, ..., x_n, z_1, ..., z_m\}) < f(\{y_1, ..., y_n, z_1, ..., z_m\})$.

In the sequel, we shall apply aggregation functions to distance values. As distances are non-negative, summation, average, and maximum are all aggregation functions. Note, however, that while summation and average are hereditary, the maximum function is not.

Definition 8 A pair $\mathcal{P} = \langle d, f \rangle$, where d is a pseudo distance and f is an aggregation function, is called a (distance-based) *preferential setting*. Given a theory $\Gamma = \{\psi_1, \ldots, \psi_n\}$, an interpretation ν, and a preferential setting $\langle d, f \rangle$, define:

- $d(\nu, \psi_i) = \min\{d(\nu, \mu) \mid \mu \in mod(\psi_i)\}$, [4]
- $\delta_{d,f}(\nu, \Gamma) = f(\{d(\nu, \psi_1), \ldots, d(\nu, \psi_n)\})$.

[2] That is, $d_H(\nu, \mu)$ is the number of atoms p such that $\nu(p) \neq \mu(p)$. This function is also known as the Dalal distance [14].

[3] I.e., the function does not decrease when a multiset element is replaced by a bigger one.

[4] Below, we exclude classical contradictions from a theory. Alternatively, if ψ is not satisfiable, one may let $d(\nu, \psi) = 1 + \max\{d(\nu, \mu) \mid \nu, \mu \in \Lambda_{\mathsf{Atoms}}\}$.

Definition 9 A (pseudo) distance d is *unbiased*, if for every formula ψ and interpretations ν_1, ν_2, if $\nu_1(p) = \nu_2(p)$ for every $p \in \mathsf{Atoms}(\psi)$, then $d(\nu_1, \psi) = d(\nu_2, \psi)$.

The last property assures that the 'distance' between an interpretation and a formula is independent of irrelevant atoms (those that do not appear in the formula). Note, e.g., that the distances in Example 5 are unbiased.

For a preferential setting $\mathcal{P} = \langle d, f \rangle$ we now define an operator $\Delta_{\mathcal{P}}$ that introduces, for every n-prioritized theory Γ, its 'most plausible' interpretations, namely: the interpretations that are $\delta_{d,f}$-closest to Γ.

Definition 10 Let $\mathcal{P} = \langle d, f \rangle$ be a preferential setting. For an n-prioritized theory $\Gamma = \Gamma_1 \oplus \Gamma_2 \oplus \ldots \oplus \Gamma_n$ consider the following n sets of interpretations:

- $\Delta^n_{\mathcal{P}}(\Gamma) = \big\{\nu \in \Lambda_{\mathsf{Atoms}} \mid \forall \mu \in \Lambda_{\mathsf{Atoms}}\ \delta_{d,f}(\nu, \Gamma_n) \leq \delta_{d,f}(\mu, \Gamma_n)\big\}$,
- $\Delta^{n-i}_{\mathcal{P}}(\Gamma) = \big\{\nu \in \Delta^{n-i+1}_{\mathcal{P}}(\Gamma) \mid \forall \mu \in \Delta^{n-i+1}_{\mathcal{P}}(\Gamma)\ \delta_{d,f}(\nu, \Gamma_{n-i}) \leq \delta_{d,f}(\mu, \Gamma_{n-i})\big\}$ for every $1 \leq i < n$.

The sequence $\Delta^n_{\mathcal{P}}(\Gamma), \ldots, \Delta^1_{\mathcal{P}}(\Gamma)$ is clearly non-increasing, as sets with smaller indices are subsets of those with bigger indices. This reflects the intuitive idea that higher-levelled formulas are preferred over lower-levelled formulas, thus the interpretations of the latter are determined by the interpretations of the former. Since the relevant interpretations are derived by distance considerations, each set in the sequence above contains the interpretations that are $\delta_{d,f}$-closest to the corresponding subtheory among the elements of the preceding set in the sequence.

Denote by $\Delta_{\mathcal{P}}(\Gamma)$ the last set obtained by this sequence (that is, $\Delta_{\mathcal{P}}(\Gamma) = \Delta^1_{\mathcal{P}}(\Gamma)$). The elements of $\Delta_{\mathcal{P}}(\Gamma)$ are the *most plausible interpretations* of Γ. These are the interpretations according to which the Γ-conclusions are drawn:

Definition 11 Let $\mathcal{P} = \langle d, f \rangle$ be a preferential setting. A formula ψ *follows* from an (n-prioritized) theory Γ, if every interpretation in $\Delta_{\mathcal{P}}(\Gamma)$ satisfies ψ (That is, if $\Delta_{\mathcal{P}}(\Gamma) \subseteq mod(\psi)$). We denote this by $\Gamma \models_{\mathcal{P}} \psi$.

Example 12 Consider again Example 1, and let $\mathcal{P} = \langle d_H, \Sigma \rangle$. Then:

$$\begin{array}{ll} \Gamma \models_{\mathcal{P}} \texttt{bird(Tweety)}, & \Gamma \models_{\mathcal{P}} \texttt{penguin(Tweety)}, \\ \Gamma \models_{\mathcal{P}} \texttt{color\_of(Tweety,Red)}, & \Gamma \models_{\mathcal{P}} \neg\texttt{fly(Tweety)}, \\ \\ \Gamma \not\models_{\mathcal{P}} \neg\texttt{bird(Tweety)}, & \Gamma \not\models_{\mathcal{P}} \neg\texttt{penguin(Tweety)}, \\ \Gamma \not\models_{\mathcal{P}} \neg\texttt{color\_of(Tweety,Red)}, & \Gamma \not\models_{\mathcal{P}} \texttt{fly(Tweety)}, \end{array}$$

as intuitively expected. In fact, using the results in the next section, one can show that the conclusions regarding bird(Tweety), penguin(Tweety), and fly(Tweety) hold in *every* setting (Proposition 19); The conclusions about the color of Tweety hold whenever d is unbiased and f is hereditary (see Proposition 24 and Note 25).

Note 13 The entailment relations in Proposition 11 generalize some other settings considered in the literature. For instance, $\models_{\langle d_U, \Sigma \rangle}$ corresponds to the merging operator in [15] (see also [8]). Also, if every Γ_i in Γ is a singleton, the iterated process of computing distances with respect to the prioritized subtheories is actually a linear revision in the sense of [16]. Other related formalisms are considered in Section 4.

3. Reasoning with $\models_{\mathcal{P}}$

In this section, we consider some basic properties of $\models_{\mathcal{P}}$. First, we examine 'flat' theories, that is: multisets in which all the assertions have the same priority. Proposition 16 recalls the main characteristics of reasoning with such theories.

Definition 14 Denote by $\models$ the standard classical entailment, that is: $\Gamma \models \psi$ if every model of Γ is a model of ψ.

Definition 15 Two sets of formulas Γ_1 and Γ_2 are called *independent* (or disjoint), if $\mathsf{Atoms}(\Gamma') \cap \mathsf{Atoms}(\Gamma'') = \emptyset$. Two independent theories Γ_1 and Γ_2 are a *partition* of a theory Γ, if $\Gamma = \Gamma_1 \cup \Gamma_2$.

Proposition 16 [17] *Let $\mathcal{P} = \langle d, f \rangle$ be a preferential setting and Γ a 1-prioritized theory. Then:*

- *$\models_{\mathcal{P}}$ is the same as the classical entailment with respect to consistent premises: if Γ is consistent, then for every ψ, $\Gamma \models_{\mathcal{P}} \psi$ iff $\Gamma \models \psi$.*
- *$\models_{\mathcal{P}}$ is weakly paraconsistent: inconsistent premises do not entail every formula (alternatively, for every Γ there is a formula ψ such that $\Gamma \not\models_{\mathcal{P}} \psi$).*
- *$\models_{\mathcal{P}}$ is non-monotonic: the set of the $\models_{\mathcal{P}}$-conclusions does not monotonically grow in the size of the premises.*

If d is unbiased, then

- *$\models_{\mathcal{P}}$ is paraconsistent: if ψ is independent of Γ then $\Gamma \not\models_{\mathcal{P}} \psi$.*

If, in addition, f is hereditary, then

- *$\models_{\mathcal{P}}$ is rationally monotonic* [18]*: if $\Gamma \models_{\mathcal{P}} \psi$ and ϕ is independent of $\Gamma \cup \{\psi\}$, then $\Gamma, \phi \models_{\mathcal{P}} \psi$.*
- *$\models_{\mathcal{P}}$ is adaptive* [19,20]*: if $\{\Gamma_1, \Gamma_2\}$ is a partition of Γ, and Γ_1 is classically consistent, then for every formula ψ that is independent of Γ_2, if $\Gamma_1 \models \psi$ then $\Gamma \models_{\mathcal{P}} \psi$.*

The arrangement of the premises in a stratified structure of priority levels allows to refine and generalize the results above. As a trivial example, it is clear that the 1-prioritized theory $\{p, \neg p\}$ is totally different than the 2-prioritized theory $\{p\} \oplus \{\neg p\}$, as in the latter the symmetry between p and $\neg p$ breaks up.

In the rest of this section we examine how preferences determine the set of conclusions. The first, trivial observation, is that even if the set of premises is not consistent, the set of its $\models_{\mathcal{P}}$-conclusions remains classically consistent:

Proposition 17 *For every setting $\mathcal{P}$, prioritized theory Γ, and formula ψ, if $\Gamma \models_{\mathcal{P}} \psi$ then $\Gamma \not\models_{\mathcal{P}} \neg\psi$.*

Proof. Otherwise, $\Delta_{\mathcal{P}}(\Gamma) \subseteq mod(\psi)$ and $\Delta_{\mathcal{P}}(\Gamma) \subseteq mod(\neg\psi)$. Since $mod(\psi) \cap mod(\neg\psi) = \emptyset$, we get a contradiction to the fact that $\Delta_{\mathcal{P}}(\Gamma) \neq \emptyset$ (as Λ_{Atoms} is finite, there are always interpretations that are minimally $\delta_{d,f}$-distant from Γ). □

Another clear characteristic of $\models_{\mathcal{P}}$ is that priorities do have a primary role in the reasoning process; conclusions of higher levelled observations remain valid when the theory is augmented with lower-levelled observations.[5]

Proposition 18 *Let Γ be an n-prioritized theory. Then for every $1 \leq i < j \leq n$, if $\Gamma_{\geq j} \models_{\mathcal{P}} \psi$ then $\Gamma_{\geq i} \models_{\mathcal{P}} \psi$.*

Proof. If $\Gamma_{\geq j} \models_{\mathcal{P}} \psi$ then $\Delta^j_{\mathcal{P}}(\Gamma) \subseteq mod(\psi)$. But $\Delta^i_{\mathcal{P}}(\Gamma) \subseteq \Delta^j_{\mathcal{P}}(\Gamma)$, and so $\Delta^i_{\mathcal{P}}(\Gamma) \subseteq mod(\psi)$ as well. Thus, $\Gamma_{\geq i} \models_{\mathcal{P}} \psi$. □

Proposition 18 implies, in particular, that anything that follows from a subtheory that consists of the higher levels of a prioritized theory, also follows from the whole theory. Next we show that when the subtheory of the higher levels is classically consistent, we can say more than that: anything that can be *classically inferred* from the highest consistent levels of a prioritized theory is also deducible from the whole theory (even when lower-levelled subtheories imply the converse). To see this we suppose, then, that at least the most preferred level of Γ is classically consistent (that is, $\mathsf{con} \leq n$).

Proposition 19 *For every setting $\mathcal{P} = \langle d, f \rangle$ and for every n-prioritized theory Γ with a consistency level $\mathsf{con} \leq n$, if $\overline{\Gamma}_{\geq \mathsf{con}} \models \psi$ then $\Gamma \models_{\mathcal{P}} \psi$.*

Proof (outline). Note, first, that for every preferential setting $\mathcal{P} = \langle d, f \rangle$ and n-prioritized theory Γ with $\mathsf{con} \leq n$, $\Delta_{\mathcal{P}}(\Gamma_{\geq \mathsf{con}}) = mod(\overline{\Gamma}_{\geq \mathsf{con}})$. By the definition of $\Delta_{\mathcal{P}}$, then, $\Delta_{\mathcal{P}}(\Gamma) \subseteq \Delta_{\mathcal{P}}(\Gamma_{\geq \mathsf{con}}) = mod(\overline{\Gamma}_{\geq \mathsf{con}})$. Now, if $\overline{\Gamma}_{\geq \mathsf{con}} \models \psi$, then ψ is true in every element of $mod(\overline{\Gamma}_{\geq \mathsf{con}})$, and so ψ holds in every element of $\Delta_{\mathcal{P}}(\Gamma)$. Thus $\Gamma \models_{\mathcal{P}} \psi$. □

Note 20 Consider again the three-levelled theory of Example 1. Proposition 19 guarantees the satisfaction of the first two items discussed in that example (the third item is considered in Note 25 below).

Proposition 21 *For every setting $\mathcal{P} = \langle d, f \rangle$ and n-prioritized theory Γ with $\mathsf{con} \leq n$, we have that $\Gamma_{\geq \mathsf{con}} \models_{\mathcal{P}} \psi$ iff $\overline{\Gamma}_{\geq \mathsf{con}} \models \psi$.*

Proof (outline). Follows from the fact that for every n-prioritized theory Γ with $\mathsf{con} \leq n$ it holds that $\Delta_{\mathcal{P}}(\Gamma_{\geq \mathsf{con}}) = mod(\overline{\Gamma}_{\geq \mathsf{con}})$. □

In particular, then, $\models_{\mathcal{P}}$ coincides with the classical entailment with respect to consistent sets of premises:

Corollary 22 *If $\overline{\Gamma}$ is consistent, then $\Gamma \models_{\mathcal{P}} \psi$ iff $\overline{\Gamma} \models \psi$.*

Proof. By Proposition 21, since if $\overline{\Gamma}$ is consistent then $\mathsf{con} = 1$, and so $\Gamma_{\geq \mathsf{con}} = \Gamma$ and $\overline{\Gamma}_{\geq \mathsf{con}} = \overline{\Gamma}$. □

In the general case, we have the following relation between $\models_{\mathcal{P}}$ and $\models$:

[5] In [8] this is called 'the principle of prioritized monotonicity'.

Corollary 23 *If* $\Gamma \models_{\mathcal{P}} \psi$ *then* $\overline{\Gamma} \models \psi$.

Proof. If $\overline{\Gamma}$ is consistent then by Corollary 22 $\Gamma \models_{\mathcal{P}} \psi$ iff $\overline{\Gamma} \models \psi$. If $\overline{\Gamma}$ is not classically consistent, then for *every* formula ψ, $\overline{\Gamma} \models \psi$. □

Next we show that in many cases we can go beyond the result of Propositions 18 and 19: Not only that one may deduce from the whole theory everything that is included in its highest levels, but also lower-levelled assertions are deducible from the whole theory, provided that no higher-levelled information contradicts them. This shows that our formalism avoids the so called *drowning effect*, that is: formulas with low priority are not inhibited just due to the fact that the information at higher levels is contradictory. Prevention of the grounding effect is very important, e.g., in the context of belief revision, as it implies that anything that has no relation to the new information need not be revised.

Proposition 24 *Let* $\mathcal{P} = \langle d, f \rangle$ *be a setting where* d *is non-biased and* f *is hereditary. If a prioritized theory* Γ *can be partitioned to a consistent theory* Γ' *and a (possible inconsistent) theory* Γ''*, then* $\Gamma \models_{\mathcal{P}} \psi$ *for every* $\psi \in \Gamma'$.

Note 25 If $\Gamma' \subseteq \Gamma_{\geq \mathsf{con}}$, then Proposition 24 is a straightforward consequence of Proposition 19. Yet, Proposition 24 is useful in cases where Γ' contains elements that are *below* the inconsistency level of Γ, and then the claim assures that the drowning effect is not imposed on these elements. Tweety dilemma Γ, considered in Example 1, is a good example for this. It can be partitioned to $\Gamma' = \{\mathtt{color\_of}(\mathtt{Tweety}, \mathtt{Red})\}$ and $\Gamma'' = \Gamma \setminus \Gamma'$. In this representation the conditions of Proposition 24 are satisfied for every preferential setting $\mathcal{P} = \langle d, f \rangle$ where d is unbiased and f is hereditary. In this case, then, $\Gamma \models_{\mathcal{P}} \mathtt{color\_of}(\mathtt{Tweety}, \mathtt{Red})$, as indeed suggested in the third item of Example 1. Note, however, that $\Gamma \not\models_{d_U, \max}$ $\mathtt{color\_of}(\mathtt{Tweety}, \mathtt{Red})$, which shows that the condition in Proposition 24, that the aggregation function should be hereditary, is indeed necessary.

Example 26 According to the possibilitic revision operator introduced in [5,6], a formula ψ is a consequence of a prioritized (possibilistic) theory Γ if it follows from all the formulas above the consistency level of Γ. In our notations, then, ψ follows from Γ iff $\overline{\Gamma}_{\geq \mathsf{con}} \models \psi$,[6] and so this formalism has the drowning effect, which prevents the drawing of any conclusion that resides below the consistency level. In other formalisms for handling prioritized theories, such as those in [4,21,22], the drowning effect is avoided by using a similar policy as ours, namely: the elements of the revised theory are constructed in a stepwise manner, starting with the highest priority level and selecting from each level as many formulas as possible without violating consistency (see also [8]).

4. Related Areas and Applications

In this section we consider in greater detail two paradigms in which priorities are exploited to determine consequences.

[6]Note that by Proposition 19 this implies that every possiblistic conclusion of Γ may be inferred also by our formalisms.

4.1. Iterated Belief Revision

Belief revision, the process of changing beliefs in order to take into account new pieces of information, is perhaps closest in spirit to the basic ideas behind our framework. A widely accepted rationality criterion in this context is the success postulate that asserts that a new item of information is always accepted. In our case, this means that new data should have a higher priority over older one. Thus, assuming that Γ represents the reasoner's belief, the revised belief state in light of new information ψ may be represented by $\Gamma \oplus \{\psi\}$. Consequently, a revision by a sequence of (possibly conflicting) observations $\psi_1, \ldots, \psi_m$ may be expressed by $\Gamma \oplus \{\psi_1\} \oplus \ldots \oplus \{\psi_m\}$.

The well-known set of rationality postulates, introduced in [23] by Alchourrón, Gärdenfors, and Makinson (AGM) for belief revision in the non-prioritized case, is often considered as the starting point in this area. These postulates were rephrased by Katsuno and Mendelzon [24] in terms of order relations as follows:

Proposition 27 *Let Γ be a set of formulas in a propositional language $\mathcal{L}$. A revision operator $\circ$ satisfies the AGM postulates if and only if there is a faithful order $\leq_\Gamma$, such that $mod(\Gamma \circ \psi) = \min(mod(\psi), \leq_\Gamma)$.*[7]

In light of this result, one may represent revision in our framework in terms of minimization of a preferential (ranking) order. For this, we consider the following adjustment, to the context of prioritized theories, of faithful orders.

Definition 28 Let $\mathcal{P}$ be a preferential setting and Γ a prioritized theory. A total preorder $\leq^{\mathcal{P}}_{\Gamma}$ on Λ_{Atoms} is called (preferentially) *faithful*, if the following conditions are satisfied:

1. If $\nu, \mu \in \Delta_{\mathcal{P}}(\Gamma)$ then $\nu <^{\mathcal{P}}_{\Gamma} \mu$ does *not* hold.
2. If $\nu \in \Delta_{\mathcal{P}}(\Gamma)$ and $\mu \notin \Delta_{\mathcal{P}}(\Gamma)$ then $\nu <^{\mathcal{P}}_{\Gamma} \mu$.

Proposition 29 *A preferential setting $\mathcal{P}$ is characterized by faithful orders: For every prioritized theory Γ there is a faithful order $\leq^{\mathcal{P}}_{\Gamma}$ (depending on $\mathcal{P}$ and Γ), such that $\Delta_{\mathcal{P}}(\Gamma) = \{\nu \in \Lambda_{\mathsf{Atoms}} \mid \forall \mu \in \Lambda_{\mathsf{Atoms}}\ \nu \leq^{\mathcal{P}}_{\Gamma} \mu\}$.*

Proposition 30 *Let $\mathcal{P}$ be a preferential setting and Γ a prioritized theory on $\mathcal{L}$. Then there is a faithful order $\leq^{\mathcal{P}}_{\Gamma}$ (depending on $\mathcal{P}$ and Γ), such that, for every non-contradictory formula ψ in $\mathcal{L}$, $\Delta_{\mathcal{P}}(\Gamma \oplus \psi) = \min\big(mod(\psi), \leq^{\mathcal{P}}_{\Gamma}\big)$.*

In terms of entailments, the last two propositions may be rewritten as follows:

Corollary 31 *Let $\mathcal{P}$ be a preferential setting, Γ a prioritized theory, and ψ a non-contradictory formula in $\mathcal{L}$. Then there is a faithful order $\leq^{\mathcal{P}}_{\Gamma}$, such that, for every formula ϕ in $\mathcal{L}$,*

1. $\Gamma \models_{\mathcal{P}} \phi$ *iff ϕ is satisfied by every $\leq^{\mathcal{P}}_{\Gamma}$-minimal element of Λ_{Atoms}.*
2. $\Gamma \oplus \psi \models_{\mathcal{P}} \phi$ *iff ϕ is satisfied by every $\leq^{\mathcal{P}}_{\Gamma}$-minimal element of $mod(\psi)$.*

[7] The reader is referred, e.g., to [7,9] for detailed discussions on this result and its notions.

Note 32 In [24] a belief base Γ is represented by a single formula which is the conjunction of the elements in Γ. In the prioritized setting this is, of-course, not possible, as different formulas in Γ have different priorities. Also, in [24] the faithful property is defined in terms of $mod(\Gamma)$ rather than $\Delta_{\mathcal{P}}(\Gamma)$. This distinction follows again from the fact that in the non-prioritized case the formula that represents a belief set Γ is consistent and as such it always has models, while in our case a prioritized theory $\Gamma = \bigoplus_i \Gamma_i$ is different than the 'flat' theory $\bigcup_i \Gamma_i$ that may not even be consistent.

Proposition 30 refers to a single revision. For successive revisions one may follow Darwiche and Pearl's approach [7], extending the AGM postulates with four additional ones. As it turns out, three of these postulates hold in our context:

Definition 33 Denote by $\Gamma \equiv_{\mathcal{P}} \Gamma'$ that Γ and Γ' have the same $\models_{\mathcal{P}}$-conclusions.

Proposition 34 *For every preferential setting $\mathcal{P} = \langle d, f \rangle$, prioritized theory Γ, and satisfiable formulas ψ, ϕ,*

C1: If $\psi \models \phi$ then $\Gamma \oplus \{\phi\} \oplus \{\psi\} \equiv_{\mathcal{P}} \Gamma \oplus \{\psi\}$.

C3: If $\Gamma \oplus \{\psi\} \models_{\mathcal{P}} \phi$ then $\Gamma \oplus \{\phi\} \oplus \{\psi\} \models_{\mathcal{P}} \phi$.

C4: If $\Gamma \oplus \{\psi\} \not\models_{\mathcal{P}} \neg\phi$ then $\Gamma \oplus \{\phi\} \oplus \{\psi\} \not\models_{\mathcal{P}} \neg\phi$.

Proof. We show C1; the proof of C3 and C4 is similar. If $\psi \models \phi$ then $mod(\psi) \subseteq mod(\phi)$, which implies that $\forall \nu \in mod(\psi)\ d(\nu, \phi) = 0$. Thus, $\forall \nu \in \Delta_{\mathcal{P}}(\{\psi\})\ d(\nu, \phi) = 0$, and so $\Delta_{\mathcal{P}}(\{\psi\}) = \Delta_{\mathcal{P}}(\{\phi\} \oplus \{\psi\}) = mod(\psi)$. It follows that $\Delta_{\mathcal{P}}(\Gamma \oplus \{\psi\}) = \Delta_{\mathcal{P}}(\Gamma \oplus \{\phi\} \oplus \{\psi\})$, and therefore $\Gamma \oplus \{\phi\} \oplus \{\psi\} \equiv_{\mathcal{P}} \Gamma \oplus \{\psi\}$. □

The forth postulate in [7], namely

C2: If $\psi \models \neg\phi$ then $\Gamma \oplus \{\phi\} \oplus \{\psi\} \equiv_{\mathcal{P}} \Gamma \oplus \{\psi\}$

is the most controversial one (see, e.g., [3,9]) and indeed in our framework it is falsified. To see this, let $\Gamma = \emptyset$, $\psi = p$, $\phi = \neg p \wedge \neg q$, and $\mathcal{P} = \langle d_H, f \rangle$ for arbitrary aggregation function f.[8] Clearly, $\psi \models \neg\phi$. However, as $\Delta_{\mathcal{P}}(\{\psi\})$ consists of interpretations that assign t to p regardless of their assignments to q, while the interpretations in $\Delta_{\mathcal{P}}(\{\phi\} \oplus \{\psi\})$ assign t to p and f to q, it follows that $\{\phi\} \oplus \{\psi\}$ and $\{\psi\}$ are *not* $\models_{\mathcal{P}}$-equivalent.

4.2. Prioritized Integration of Independent Data Sources

Information systems often have to incorporate *several* sources with possibly different preferences. In this section we show how this can be done in our framework. For this, we use two types of distance aggregations: *internal aggregations*, for prioritizing different formulas in the same theory, and *external aggregations*, for prioritizing different theories. As internal and external aggregations may reflect different kinds of considerations, they are represented by two different aggrega-

[8] f is irrelevant here since each priority level is a singleton.

tion functions, denoted f and g, respectively. Now, using the terminology and the notations of the previous sections, we can think of the underlying n-prioritized theory as follows:

$$\Gamma = \{\Gamma^1_1, \ldots, \Gamma^1_{k_1}\} \oplus \ldots \oplus \{\Gamma^n_1, \ldots, \Gamma^n_{k_n}\}, \tag{1}$$

where now each Γ^i_j is a different theory, theories with the same superscript have the same precedence, and Γ^i is preferred over Γ^j iff $i > j$. This can be formalized by the following generalizations of Definitions 8 and 10:

Definition 35 An *extended preferential setting* is a triple $\mathcal{E} = \langle d, f, g \rangle$, where d is a pseudo distance and f, g are aggregation functions. Given an n-prioritized theory $\Gamma = \{\Gamma^1_1, \ldots, \Gamma^1_{k_1}\} \oplus \ldots \oplus \{\Gamma^n_1, \ldots, \Gamma^n_{k_n}\}$ and an interpretation ν, define for every $1 \leq i \leq n$ and $1 \leq j \leq k_i$ the value of $\delta_{d,f}(\nu, \Gamma^i_j)$ just as in Definitions 8. Also, let

$$\delta_{\mathcal{E}}(\nu, \overline{\Gamma^i}) = \delta_{d,f,g}(\nu, \overline{\Gamma^i}) = g\big(\{\delta_{d,f}(\nu, \Gamma^i_1), \ldots, \delta_{d,f}(\nu, \Gamma^i_{k_i})\}\big).$$

Definition 36 Let $\mathcal{E} = \langle d, f, g \rangle$ be an extended preferential setting. Given an n-prioritized theory $\Gamma = \{\Gamma^1_1, \ldots, \Gamma^1_{k_1}\} \oplus \ldots \oplus \{\Gamma^n_1, \ldots, \Gamma^n_{k_n}\}$, consider the following n sets of interpretations:

- $\Delta^n_{\mathcal{E}}(\Gamma) = \big\{\nu \mid \forall \mu \ \delta_{\mathcal{E}}(\nu, \overline{\Gamma^n}) \leq \delta_{\mathcal{E}}(\mu, \overline{\Gamma^n})\big\}$,
- $\Delta^{n-i}_{\mathcal{E}}(\Gamma) = \big\{\nu \in \Delta^{n-i+1}_{\mathcal{E}}(\Gamma) \mid \forall \mu \in \Delta^{n-i+1}_{\mathcal{E}}(\Gamma) \ \delta_{\mathcal{E}}(\nu, \overline{\Gamma^{n-i}}) \leq \delta_{\mathcal{E}}(\mu, \overline{\Gamma^{n-i}})\big\}$ for every $1 \leq i < n$.

The *most plausible interpretations* of Γ (with respect to d, f, g) are the interpretations in $\Delta^1_{\mathcal{E}}(\Gamma)$ (henceforth denoted by $\Delta_{\mathcal{E}}(\Gamma)$).

The corresponding consequence relations are now defined as follows:

Definition 37 Let $\mathcal{E} = \langle d, f, g \rangle$ be an extended preferential setting. A formula ψ *follows* from an an n-prioritized theory Γ if every interpretation in $\Delta_{\mathcal{E}}(\Gamma)$ satisfies ψ. We denote this by $\Gamma \models_{\mathcal{E}} \psi$.

Clearly, Definition 37 generalizes Definition 11 in the sense that if in (1) above $k_i = 1$ for every $1 \leq i \leq n$, then for every g, $\models_{\mathcal{E}}$ (in the sense of Definition 37) is the same as $\models_{\mathcal{P}}$ (in the sense of Definition 11).[9]

Example: Constraint-Based Merging of Prioritized Data-Sources

Consider the following scenario regarding speculations on the stock exchange (see also [3]). An investor consults with four financial experts about their opinions regarding four different shares, denoted s_1, s_2, s_3 and s_4. The opinion of expert i is represented by a theory (data-source) Γ_i. Suppose that $\Gamma_1 = \Gamma_2 = \{s_1, s_2, s_3\}$, $\Gamma_3 = \{\neg s_1, \neg s_2, \neg s_3, \neg s_4\}$, and $\Gamma_4 = \{s_1, s_2, \neg s_4\}$. Thus, for instance, expert 4 suggests to buy shares s_1 and s_2, doesn't recommend to buy share s_4, and doesn't have an opinion about s_3.

[9]Alternatively, $\models_{\mathcal{E}}$ coincides with $\models_{\mathcal{P}}$ if in (1) each T^i_j is a singleton and $g = f$.

Suppose, in addition, that the investor has his own restrictions about the investment policy. For instance, if some share, say s_4, is considered risky, buying it may be balanced by purchasing at least two out of the three other shares, and vice-versa. This may be represented by the following integrity constraint: $\mathcal{IC} = \{s_4 \longleftrightarrow \big((s_1 \wedge s_2) \vee (s_2 \wedge s_3) \vee (s_1 \wedge s_3)\big)\}$. Assuming that all the expert are equally faithful, their suggestions may be represented by the 2-prioritized theory $\Gamma = \{\Gamma_1, \Gamma_2, \Gamma_3, \Gamma_4\} \oplus \{\mathcal{IC}\}$, in which the investor's constraint about the purchasing policy is of higher precedence than the experts' opinions. For the extended setting $\langle d_U, \Sigma, \Sigma\rangle$ we get that the most plausible interpretations of Γ are the elements of the following set:

$$\Delta_{d_U,\Sigma,\Sigma}(\Gamma) = \{\nu \in mod(\mathcal{IC}) \mid \forall \mu \in mod(\mathcal{IC}) \\ \delta_{d_U,\Sigma,\Sigma}(\nu, \{\Gamma_i \mid 1 \leq i \leq 4\}) \leq \delta_{d_U,\Sigma,\Sigma}(\mu, \{\Gamma_i \mid 1 \leq i \leq 4\})\}.$$

The models of $\mathcal{IC}$ and their distances to $\overline{\Gamma} = \{\Gamma_1, \dots, \Gamma_4\}$ are given below.

	s_1	s_2	s_3	s_4	$\delta_{d_U,\Sigma,\Sigma}(\nu_i, \overline{\Gamma})$
ν_1	t	t	t	t	**5**
ν_2	t	t	f	t	7
ν_3	t	f	t	t	7
ν_4	t	f	f	f	7

	s_1	s_2	s_3	s_4	$\delta_{d_U,\Sigma,\Sigma}(\nu_i, \overline{\Gamma})$
ν_5	f	t	t	t	7
ν_6	f	t	f	f	6
ν_7	f	f	t	f	6
ν_8	f	f	f	f	8

Thus, $\Delta_{d_U,\Sigma,\Sigma}(\Gamma) = \{\nu_1\}$, and so the investor will purchase all the four shares.

Clearly, the experts could have different reputations, and this may affect the investor's decision. For instance, assuming that expert 4 has a better reputation than the other experts, his or her opinion may get a higher precedence, yielding the following 3-prioritized theory: $\Gamma' = \{\Gamma_1, \Gamma_2, \Gamma_3\} \oplus \{\Gamma_4\} \oplus \{\mathcal{IC}\}$. It is interesting to note that in this case the recommendation of the most significant expert (number 4) does not comply with the investor's restriction.

By using the same setting as before ($d = d_U,\ f = g = \Sigma$), the investor ends up with a different investment policy, according to the following table:

	s_1	s_2	s_3	s_4	$\delta_{d_U,\Sigma,\Sigma}(\nu_i, \Gamma_4)$	$\delta_{d_U,\Sigma,\Sigma}(\nu_i, \{\Gamma_1, \Gamma_2, \Gamma_3\})$
ν_1	t	t	t	t	**1**	0+0+4 = 4
ν_2	t	t	f	t	**1**	1+1+3 = 5
ν_3	t	f	t	t	2	N.A.
ν_4	t	f	f	f	**1**	**1+1+1 = 3**
ν_5	f	t	t	t	2	N.A.
ν_6	f	t	f	f	**1**	**1+1+1 = 3**
ν_7	f	f	t	f	2	N.A.
ν_8	f	f	f	f	2	N.A.

Here, $\Delta_{d_U,\Sigma,\Sigma}(\Gamma') = \{\nu_4, \nu_6\}$, and the decision would be to purchase either s_1 or s_2, *but not both*, which seems as a 'fair balance' between the investor's restriction and the recommendation of the most significant expert (taking into account also the other recommendations).

References

[1] O. Arieli, M. Denecker, B. Van Nuffelen, and M. Bruynooghe. Coherent integration of databases by abductive logic programming. *Artificial Intelligence Research*, 21:245–286, 2004.

[2] A. Calì, D. Calvanese, G. De Giacomo, and M. Lanzerini. Data integration under integrity constraints. *Information Systems*, 29(2):147–163, 2004.

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

[4] O. Arieli. Four-valued logics for reasoning with uncertainty in prioritized data. In B. Bouchon-Meunier, R. Yager, and L. Zadeh, editors, *Information, Uncertainty, Fusion*, pages 293–304. Kluwer, 1999.

[5] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, and H. Prade. Inconsistency management and prioritized syntax-based entailment. In *Proc. IJCAI'93*, pages 640–645, 1993.

[6] S. Benferhat, D. Dubois, and H. Prade. How to infer from inconsistent beliefs without revising? In *Proc. IJCAI'95*, pages 1449–1455, 1995.

[7] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.

[8] J. Delgrande, D. Dubois, and J. Lang. Iterated revision and prioritized merging. In *Proc. KR'06*, pages 210–220, 2006.

[9] Y. Jin and M. Thielscher. Iterated revision, revised. *Artificial Intelligence*, 171:1–18, 2007.

[10] M. Williams. Iterated theory base change: A computational model. In *Proc. IJCAI'95*, pages 1541–1547, 1995.

[11] V. S. Subrahmanian. Mechanical proof procedures for many valued lattice-based logic programming. *Journal of Non-Classical Logic*, 7:7–41, 1990.

[12] D. Dubois, J. Lang, and H. Prade. Possibilistic logic. In D. Gabbay, C. Hogger, and J. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3, pages 439–513. 1994.

[13] M. Goldszmidt and J. Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence*, 84:57–112, 1996.

[14] M. Dalal. Investigations into a theory of knowledge base revision. In *Proc. AAAI'88*, pages 475–479. AAAI Press, 1988.

[15] D. Lehmann. Another perspective on default reasoning. *Annals of Methematics and Artificial Intelligence*, 15(1):61–82, 1995.

[16] B. Nebel. Base revision operations and schemes: Semantics, representation, and complexity. In *Proc. ECAI'94*, pages 341–345, 1994.

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

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

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

[20] D. Batens. Inconsistency-adaptive logics. In E. Orlowska, editor, *Logic at Work*, pages 445–472. Physica Verlag, 1998.

[21] G. Brewka. Preferred subtheories: an extended logical framework for default reasoning. In *Proc. IJCAI'89*, pages 1043–1048, 1989.

[22] B. Nebel. Belief revision and default reasoning: Syntax-based aproaches. In *Proc. KR'01*, pages 417–428, 1991.

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

[24] Katsumo H. and A. O. Mendelzon. Propotisional knowldge base revision and minimal change. *Artificial Intelligence*, 52:263–294, 1991.