A Preferential Framework for Trivialization-Resistant Reasoning with Inconsistent Information

Anna Zamansky*

Vienna University of Technology

Abstract. Paraconsistent entailments based on more than two truth-values are useful formalisms for handling inconsistent information in large knowledge bases. However, such entailments suffer from two major drawbacks: they are often too cautious to allow intuitive classical inference, and are trivialization-prone. Two preferential mechanisms have been proposed to deal with these two problems, but they are formulated in different terms, and are hard to combine. This paper is a step towards a systematization and generalization of these approaches. We define an abstract framework, which allows for incorporating various preferential criteria into paraconsistent entailments in a modular way. We show that many natural cases of previously studied entailments can be simulated within this framework. Its usefulness is also demonstrated using a concrete domain related to ancient geography.

1 Introduction

Handling inconsistent information in large knowledge bases is an important practical problem, that has been recently drawing a lot of attention. The main drawback of classical logic (CL) in this context is that it fails to accommodate the fact that knowledge bases containing contradictory data may still produce useful answers to queries. This is because in CL a single inconsistency leads to trivialization of the whole knowledge base. The traditional approach to handling inconsistency in knowledge bases has been that of *revision*: in case of an inconsistency, some pieces of information must be abandoned in order to maintain consistency. In many natural cases, however, inconsistency is an inherent, and even a very important part of information systems, and should not be treated as an undesirable phenomenon (see [12] for further discussion).

As a running example in this paper, we use a domain, in which inconsistent information plays an important role, and as such should often be preserved, rather than attempted to be discarded. The domain is taken from a concrete practical problem, pointed out to the author in a personal communication[1] by

* The author is supported by funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 252314.

[1] The author would like to thank Ekaterina Ilyuschetchkina for her valuable contribution to this paper.

a researcher of descriptive ancient geography. Due to a vast amount of information extracted from various ancient texts, there is a need for automated methods for keeping track and reasoning about geographic information. More specifically, texts of ancient authors (such as Eratosthenes, Strabo, Ptolemy, Pausanias, etc.) are used in the study of ancient geography to extract information about the ancient perception of the world. Such texts contain a lot of conflicting and contradictory information. The explanations for the contradictions vary: the primary objective of such texts is not scientific and thus they are often imprecise, the authors rely on information from different sources (the identity of which is usually not known), etc. One example of a text containing several controversial contradictions is a description of the habitable world written in a terse and elegant style by Dionysios Periegetes[2] ([15]). The text begins with describing a "bird's-eye" view of the world. In it, Periegetes describes schematically the continents of Africa, Europe and Asia as triangles (of different forms), and notes that the eastern border of Africa is defined by the Nile river. When later describing the "regional picture" of the continent of Africa, however, he states that on the east Africa borders with the Arabian Gulf[3] (the modern Red Sea). When describing in details the Nile river, he repeats, however, that Nile is the eastern border of Africa. Another example of (perhaps a more subtle) contradiction is that in his "bird's-eye" description of the world, Periegetes describes the continent of Africa as having the form of a triangle. However, when describing the regional picture, he states that is has the form of a trapezoid. But taking into account the mathematical theories developed far before his time, it is commonly believed that Periegetes makes a clear distinction between a triangle and a trapezoid.

Attempting to collect the pieces of information described in the above examples in a traditional knowledge base obviously results in its trivialization, i.e., everything (classically) follows from it. However, discarding any of the facts from the knowledge base would lead to loss of important information. The ideal solution is finding a way to keep all pieces of information in the knowledge base, while preventing its trivialization. This naturally leads to the need for *paraconsistent entailment relations*, which would allow to make useful inferences from inconsistent theories.

One of the most common ways of defining useful paraconsistent entailments is using three-valued (or in general many-valued) matrices. The idea is to add a new truth value $\top$ (or more), with the intuitive meaning of being "inconsistent", or "both true and false". The entailments induced by such matrices are "well-behaved" consequence relations (crs), which enjoy a simple and intuitive semantics and have a well-developed proof theory. However, such crs suffer from two major drawbacks. The first is that in many cases they are too cautious to

[2] The lifedates of Dionysios Periegetes (also known as Dionysios of Alexandria) are estimated around the 2nd century A.D. His poem enjoyed great popularity in ancient times. Although he is a poet and not a scientist, his texts today are a useful source for extracting information on geographic knowledge in the ancient world.

[3] More precisely, he speaks of the isthmus between the Arabian Gulf and the Mediterranean Sea.

allow natural (and seemingly harmless) classical inferences. Interestingly, such problem obtains even for those paraconsistent logics which are maximal (that is, any extension leads to a logic that is no longer paraconsistent), which means that extending the logic cannot be a satisfactory solution. One solution, proposed in [17] in the spirit of Shoham's preferential semantics ([21]), is to focus on preferential refinements of the basic inference relation, applying the principle of inconsistency minimization: the interpretations which are as close as possible to classical interpretations are preferred, resulting in a better approximation of classical reasoning.

The second drawback of paraconsistent many-valued logics (induced by some matrix $\mathcal{M}$), which becomes crucial in the context of practical applications, is that despite the fact that they tolerate classical inconsistency, they may still be trivialized in case of inconsistency in $\mathcal{M}$. A database repair method was proposed in [9, 20] for two particular paraconsistent logics. However, in our context repair implies loss of information, so instead of refining the knowledge base, we would prefer to refine our entailment relation. A distance-based approach for defining trivialization-tolerant variants for logics induced by denotational semantics was proposed in [2, 4] (see also [16]). The idea is to apply distance-based minimization of non-satisfiability for "approximation" of $\mathcal{M}$-models of an $\mathcal{M}$-inconsistent knowledge base. The resulting entailment coincides with the original logic in case that the knowledge base is $\mathcal{M}$-consistent, but does not trivialize otherwise. This method, however, does not apply to the above mentioned preferential refinements, which are, as observed in [17], trivialization-prone as well.

In this paper we propose a *systematic* approach to incorporate preferential criteria into paraconsistent entailments based on many-valued matrices. The approach is based on the observation, that despite the fact that the methods of [17] for handling cautiousness and the methods of [2, 4] for handling trivialization are formulated in different terms, they are in fact different facets of minimization according to "epistemic" criteria. We therefore define an abstract framework for specifying such criteria in a *uniform* way, which also allows for a modular combination of these criteria. This allows for defining trivialization-resistant variants of various preferential paraconsistent entailments based on a matrix $\mathcal{M}$, which coincide with the original entailment whenever the premises are $\mathcal{M}$-consistent, but do not trivialize otherwise. This is done via *lexicographic aggregation of preference*, which is a common technique in choice theory [1]. The framework also captures many previously studied paraconsistent entailments (including the preferential relations on [17], summation-based entailments of [2, 4] and maximal-consistency based cautious semantics of [13]). We demonstrate the usefulness of the framework using the above mentioned domain of descriptive ancient geography.

2 Preliminaries

In what follows, $\mathcal{L}$ denotes a propositional language with a countable set $\mathsf{Atoms} = \{p, q, r \ldots\}$ of atomic formulas and a (countable) set $\mathcal{W}_{\mathcal{L}} = \{\psi, \phi, \sigma, \ldots\}$ of wellformed formulas. A theory Γ is a finite set of formulas. The set of all theories of $\mathcal{L}$ is denoted by $\mathcal{T}_{\mathcal{L}}$.

Definition 1. – An *entailment relation* $\vdash$ for $\mathcal{L}$ is a binary relation between theories from $\mathcal{T}_{\mathcal{L}}$ and formulas in $\mathcal{W}_{\mathcal{L}}$.

- An entailment relation $\vdash$ for $\mathcal{L}$ is a *(Tarskian) consequence relation (tcr)* if it has the following properties: (i) *Reflexivity*: if $\psi \in \Gamma$ then $\Gamma \vdash \psi$, (ii) *Monotonicity*: if $\Gamma \vdash \psi$ and $\Gamma \subseteq \Gamma'$, then $\Gamma' \vdash \psi$, and (iii) *Transitivity*: if $\Gamma \vdash \psi$ and $\Gamma', \psi \vdash \varphi$ then $\Gamma, \Gamma' \vdash \varphi$.
- A *propositional logic* is a pair $\langle \mathcal{L}, \vdash \rangle$, where $\vdash$ is a structural[4] consequence relation for $\mathcal{L}$.

The most standard semantic way of defining logics is by many-valued matrices:

Definition 2. A (many-valued) *matrix* for a language $\mathcal{L}$ is a triple $\mathcal{M} = \langle \mathcal{V}, \mathcal{D}, \mathcal{O} \rangle$, where $\mathcal{V}$ is a non-empty set of truth values, $\mathcal{D}$ is a non-empty proper subset of $\mathcal{V}$, containing the *designated* elements of $\mathcal{V}$, and $\mathcal{O}$ includes an n-ary function $\widetilde{\diamond}_{\mathcal{M}} : \mathcal{V}^n \to \mathcal{V}$ for each n-ary connective $\diamond$ of $\mathcal{L}$.

Example 1. The simplest example of a many-valued matrix is the (standard) classical two-valued matrix for $\mathcal{L}_{cl}$, which we shall denote by $\mathcal{M}_{cl}$.
The three-valued matrix $\mathbf{M_{B_0}} = \{\{t, \top, f\}, \{t, \top\}, \mathcal{O}\}$ for $\mathcal{L}_{cl} = \{\neg, \vee, \wedge, \supset\}$ is used in [17], and is defined as follows:

$\tilde{\wedge}$	t	f	$\top$
t	t	f	$\top$
f	f	f	f
$\top$	$\top$	f	$\top$

$\tilde{\vee}$	t	f	$\top$
t	t	t	t
f	t	f	$\top$
$\top$	t	$\top$	$\top$

$\tilde{\to}$	t	f	$\top$
t	t	f	$\top$
f	t	t	t
$\top$	t	f	$\top$

	$\tilde{\neg}$
t	f
f	t
$\top$	$\top$

Definition 3. Let $\mathcal{M} = \langle \mathcal{V}, \mathcal{D}, \mathcal{O} \rangle$ be a matrix for $\mathcal{L}$.

- An $\mathcal{M}$*-valuation* for $\mathcal{L}$ is a function $\nu : \mathcal{W}_{\mathcal{L}} \to \mathcal{V}$ such that for every n-ary connective $\diamond$ of $\mathcal{L}$ and every $\psi_1, \ldots, \psi_n \in \mathcal{W}_{\mathcal{L}}$, $\nu(\diamond(\psi_1, \ldots, \psi_n)) = \widetilde{\diamond}_{\mathcal{M}}(\nu(\psi_1), \ldots, \nu(\psi_n))$. We denote the set of all the $\mathcal{M}$-valuations by $\Lambda_{\mathcal{M}}$.
- A valuation $\nu \in \Lambda_{\mathcal{M}}$ is an $\mathcal{M}$*-model* of a formula ψ, if it belongs to the set $mod_{\mathcal{M}}(\psi) = \{\nu \in \Lambda_{\mathcal{M}} \mid \nu(\psi) \in \mathcal{D}\}$. The $\mathcal{M}$-models of a theory Γ are the elements of the set $mod_{\mathcal{M}}(\Gamma) = \cap_{\psi \in \Gamma}\, mod_{\mathcal{M}}(\psi)$.
- A formula ψ (a theory Γ) is $\mathcal{M}$*-consistent* if $mod_{\mathcal{M}}(\psi) \neq \emptyset$ ($mod_{\mathcal{M}}(\Gamma) \neq \emptyset$).
- The entailment $\vdash_{\mathcal{M}}$ *induced by* $\mathcal{M}$, is defined by $\Gamma \vdash_{\mathcal{M}} \psi$ if $mod_{\mathcal{M}}(\Gamma) \subseteq mod_{\mathcal{M}}(\psi)$.

It is easy to see that for every matrix $\mathcal{M}$ for $\mathcal{L}$, $\langle \mathcal{L}, \vdash_{\mathcal{M}} \rangle$ is a propositional logic. This notion, however, is in many cases too restrictive. In the context of non-monotonic reasoning it is usual to consider the following weaker notion, that still guarantees in many cases that the induced entailment is "well-behaved" (see, e.g., [3, 18, 19]):

[4] An entailment $\vdash$ is structural if $\Gamma \vdash \varphi$ implies $\sigma(T) \vdash \sigma(\varphi)$ for every substitution σ.

Definition 4. An entailment $\vdash$ is a *cautious consequence relation* (with respect to $\mathcal{M}$) if it has the following properties: (i) *Cautious Reflexivity (with respect to $\mathcal{M}$):* if Γ is $\mathcal{M}$-consistent and $\psi \in \Gamma$ then $\Gamma \vdash \psi$, (ii) *Cautious Monotonicity* [11]: if $\Gamma \vdash \psi$ and $\Gamma \vdash \phi$ then $\Gamma, \psi \vdash \phi$, and (iii) *Cautious Transitivity* [18]: if $\Gamma \vdash \psi$ and $\Gamma, \psi \vdash \phi$ then $\Gamma \vdash \phi$.

3 Paraconsistent Entailments

The standard notion of paraconsistency with respect to $\neg$ (defined in [8, 5] for propositional logics), can be extended to the context of (less restrictive) entailments as follows:

Definition 5. Let $\mathcal{L}$ be a language which includes a unary connective $\neg$. An entailment $\vdash$ for $\mathcal{L}$ is *paraconsistent* (with respect to $\neg$) if for every theory Γ and every $\mathcal{L}$-formula ψ, such that $\psi, \neg\psi \in \Gamma$, there is some $\mathcal{L}$-formula ϕ, such that $\Gamma \nvdash \phi$.

There is a vast amount of different paraconsistent entailments that have been suggested and investigated over the years. A natural question arises: what are the properties that a "well-behaved" entailment should satisfy? This question was considered in [5] for the case of propositional logics. The intuitive answer is that an ideal paraconsistent logic should follow the original intention of one of the founders of paraconsistent logics, Newton da Costa ([10]), by "retaining as much of classical logic as possible", while still allowing non-trivial inconsistent theories. This rather vague notion of "ideal logics" was defined in [5] in precise terms, and it includes the following basic requirements[5]: (i) containment of classical logic, (ii) absolute maximal paraconsistency (in the sense that any extension of the logic results in a non-paraconsistent logic), (iii) maximal paraconsistency with respect to classical logic (in the sense that extending the set of theorems of the logic results in classical logic), and (iv) reasonable language (that is, having a natural implication and conjunction). It was also shown that practically all reasonable three-valued paraconsistent logics are ideal in the defined sense.

In addition to having a natural implication and conjunction, a useful feature of the language of a paraconsistent logic is an internalization of the notion of consistency. A well-known example of logics which have this feature is the family of paraconsistent logics motivated by da Costa's approach, known as *Logics of Formal (In)consistency* (LFIs, [8]). In these logics we are able to make a distinction in the language between consistent and inconsistent propositions (in other words, "normal" ones – which cannot be both true and false, and "abnormal" ones for which such case is possible), and use this distinction to restrict classical rules of inference only to the "normal" ones. This distinction is often done using a primitive (or defined) unary connective $\circ$, where the meaning of $\circ\varphi$ is "φ is consistent".

[5] We refer the reader to [5] for the formal definitions and further details.

Example 2. Consider the matrix $\mathcal{M}_{\mathbf{B}}$ for $\mathcal{L} = \{\wedge, \vee, \supset, \neg, \circ\}$ obtained by extending $\mathcal{M}_{\mathbf{B_0}}$ from Example 1 by the following interpretation of $\circ$: $\tilde{\circ}t = \tilde{\circ}f = t$, $\tilde{\circ}\top = f$. The logic $\mathbf{B} = \langle \mathcal{L}, \vdash_{\mathcal{M}_{\mathbf{B}}} \rangle$ was introduced in [9, 20] for repairing evolutionary databases (it was called **LFI1** there), and shown in [5] to be ideal in the sense explained above. Note that the explosion principle of classical logic does not hold in $\mathbf{B}$: $\psi, \neg\psi \not\vdash_{\mathbf{B}} \varphi$, and so $\mathbf{B}$ is paraconsistent. However, a weakened version of the explosion principle does hold: $\circ\psi, \psi, \neg\psi \not\vdash_{\mathbf{B}} \varphi$. Intuitively, this means that a consistent proposition cannot be both true and false.

Example 3. Let us utilize the logic **B** in the context of the ancient geography domain. The text of Periegetes starts with a "bird's-eye" description of the world, in which (among others) the following facts are mentioned: (i) The earth consists of three continents: Africa, Europe and Asia, (ii) Africa and Europe have the form of a triangle, (iii) The eastern border of Africa is the Nile river. In the regional description of Africa, a new fact appears: (iv) The eastern border of Africa is the Arabian Gulf. We can capture this information in the knowledge base EK (of explicit knowledge of Periegetes):

Continent(africa)
Continent(europe)
Continent(asia)
EastBorder (africa, nile)
Triangle(europe)
Triangle(africa)
EastBorder (africa, arabgulf)

We will also maintain an additional set of formulas which reflects implicit assumptions related to the geographic concepts (for instance, that a continent has only one geographical object on its eastern border), to the additional background knowledge Periegetes is assumed to have (for instance, that triangle is a different shape from a trapezoid), etc. Suppose we impose the following set of assumptions[6]: $\mathcal{IA} = \{\forall x \forall y \forall z \text{ EastBorder}(x,y) \wedge \text{EastBorder}(x,z) \to y = z, \ \mathsf{nile} \neq \mathsf{arabgulf} \ \}$. Clearly, $KB = EK \cup \mathcal{IA}$ is classically inconsistent, and is trivialized in classical logic. However, this is not the case if a three-valued paraconsistent logic is used instead. Using **B**, however, trivialization is avoided: (i) $KB \not\vdash_{\mathbf{B}}$ $\mathsf{EastBorder(europe, arabgulf)}$, (ii) $KB \not\vdash_{\mathbf{B}} \mathsf{Triangle(asia)}$.

4 The Problems of Cautiousness and Trivialization

Despite the usefulness of three-valued paraconsistent logics for dealing with inconsistency, such entailments suffer from two major drawbacks. We discuss them in further details below.

[6] To simplify the presentation in the sequel, we would like to restrict our discussion to the propositional level. Thus we interpret here the atomic formulas as propositional atoms, and the universally quantified formulas as the set of their substitution instances.

The first drawback, pointed out in [17], is that three-valued paraconsistent logics are often too cautious to make intuitive classical entailments. Intuitively, e.g., when a knowledge base contains evidence that ψ is true and does not contain any evidence that $\neg\psi$ is true, one cannot infer that $\neg\psi$ is false, although it seems plausible.

Example 4. Let $\mathbf{B}_\perp$ be the logic obtained from $\mathbf{B}$ by adding the bottom element $\perp$ (which is assigned f by all valuations). Now note that in KB from Example 3, we have evidence that $\mathsf{Continent(europe)}$ is true, but *no* evidence that its negation is true. Despite this, we cannot infer that its negation is false: $KB \not\vdash_{\mathbf{B}_\perp} \neg\mathsf{Continent(europe)} \rightarrow \perp$.

One of the solutions to the problem of cautious inference, proposed in [17], is to focus on preferential refinements of the basic inference relation, applying the principle of inconsistency minimization. The definitions in [17] apply only to the three-valued matrix $\mathcal{M}_{\mathbf{B_0}}$ from Example 1, but can be adapted as follows to the context of arbitrary finite-valued matrices:

Definition 6. *For a set of elements U and a well-founded (partial) order $\leq$ on U, we define: $min(\leq, U) = \{u \in U \mid \neg\exists u' \in U.u' \leq u \text{ and } u \not\leq u'\}$.*

Definition 7. *Let $\mathcal{M}$ be a matrix and $\leq$ a well-founded (partial) order on $\Lambda_{\mathcal{M}}$. Then $\Gamma \vdash^{\leq}_{\mathcal{M}} \psi$ if $min(\leq, mod_{\mathcal{M}}(\Gamma)) \subseteq mod_{\mathcal{M}}(\psi)$.*

Example 5. The following are two examples of order relations on $\Lambda_{\mathcal{M}_{\mathbf{B_0}}}$ proposed in [17, 7], where the set of atoms in $\mathcal{L}$ is assumed to be finite (this restriction will be eliminated in our framework, see the discussion in Remark 1 below). For a valuation μ and a theory Γ, denote by $\mu!$ the set of atoms in $\mathcal{L}$, which are assigned $\top$ by μ. Let $\mu \leq_P \nu$ iff $\mu! \subseteq \nu!$, and $\mu \leq_{CP} \nu$ iff $|\nu!| \leq |\mu!|$.

Example 6. Let us return to Example 4 and see how using, e.g., $\vdash^{\leq_P}_{\mathbf{M}_{\mathbf{B}_\perp}}$ instead of $\vdash_{\mathbf{M}_{\mathbf{B}_\perp}}$ solves the problem discussed there. It is easy to see that for every $\nu \in min(mod_{\mathcal{M}_{\mathbf{B}}}(\Gamma), \leq_P)$, $\nu(\mathsf{Continent(europe)}) = t$. It now follows that $KB \vdash^{\leq_P}_{\mathbf{B}_\perp} \neg\mathsf{Continent(europe)} \rightarrow \perp$.

The second major drawback of paraconsistent entailments is that they are prone to trivialization. Of course, such entailments are no longer trivialized in the face of a classical inconsistency. However, the danger of trivialization remains, as it may happen that a knowledge base is no longer $\mathcal{M}$-consistent, as demonstrated by the following example.

Example 7. Consider the following toy knowledge base $EB_0 = \{\mathsf{EastBorder(africa, nile)}, \mathsf{Triangle(europe)}, \mathsf{Triangle(africa)}, \neg\mathsf{Triangle(africa)}\}$. Suppose also that the researchers come to believe that all geographical facts specified in the "bird's-eye" of the world should be taken as consistent, in the sense that they do not have reason to expect to find any information contradicting this part of the text. This includes the facts that Europe and Africa have the form of a triangle. This can be captured by the following set of implicit assumptions (using the consistency operator available in the logic **B**): $\mathcal{IA}_0 = \{\circ\mathsf{Triangle}(\mathsf{africa}), \circ\mathsf{Triangle}(\mathsf{europe})\}$. However, the knowledge base $KB_0 = EB_0 \cup \mathcal{IA}_0$ is $\mathcal{M}_{\mathbf{B}}$-inconsistent, resulting in its trivialization.

A method for avoiding trivialization for logics induced by denotational semantics was proposed in [2, 4]. The idea is to apply distance-based minimization of non-satisfiability for "approximation" of $\mathcal{M}$-models of an $\mathcal{M}$-inconsistent knowledge base. The resulting entailment coincides with the original logic in case that the knowledge base is $\mathcal{M}$-consistent, and does not trivialize otherwise. The problem of trivialization, however, obtains also for the preferential entailments from Definition 7, to which the above method is not applicable. Indeed, preference is made there over models of the knowledge base, and as soon as the set of models is empty, we are again facing trivialization. In the next section we combine distance-based and preferentical approaches in a unified framework, allowing for trivialization-resistant preferential entailments.

5 Trivialization-Resistant Preferential Framework

In what follows we define a framework for incorporating preference criteria into paraconsistent entailments based on *arbitrary* many-valued matrices. The framework is based on the key notion of *order generators*. The intuitive idea is that the preference criteria are knowledge base dependent, that is each theory induces its own order relation on the space of valuations. This dependency is encapsulated in an order generator, which generates a different order for each theory[7]:

Definition 8. *Let $\mathcal{M}$ be a matrix.*

- *An order generator* $\mathbf{O}$ *for* $\mathcal{M}$ *is a function which for every theory* Γ *returns a well-founded partial order* $\leq_\Gamma$ *on* $\Lambda_\mathcal{M}$.
- *For some order generator* $\mathbf{O}$ *for* $\mathcal{M}$, *we write* $\Gamma \vdash^{\mathbf{O}}_{\mathcal{M}} \psi$ *if* $min(\mathbf{O}(\Gamma), \Lambda_\mathcal{M}) \subseteq mod_\mathcal{M}(\psi)$.

The key property of entailments of the form defined above is that for any matrix satisfying the following natural normality condition from [4], they are completely trivialization-resistant:

Definition 9. *A matrix* $\mathcal{M}$ *for* $\mathcal{L}$ *is* normal *if for every* $\nu \in \Lambda_\mathcal{M}$, *there is a* $\mathcal{L}$*-formula* ψ, *such that* $\nu \not\models_\mathcal{M} \psi$.

It is easy to see, e.g., that any matrix $\mathcal{M}$ in which a bottom element is definable, is normal. Moreover, so is any three-valued paraconsistent matrix for a language with $\neg$ and $\circ$, where these unary connectives are interpreted like in $\mathcal{M}_\mathbf{B}$ from Example 2.

Proposition 1 (non-trivialization). *Let* $\mathcal{M}$ *be a normal matrix for* $\mathcal{L}$ *and* $\mathbf{O}$ *an order generator. Then for every theory* Γ, *there is some* ψ, *such that* $\Gamma \not\vdash^{\mathbf{O}}_{\mathcal{M}} \psi$.

[7] See also the discussion in Remark 1 on order generators.

Proof: Let Γ be a theory. Since $\mathbf{O}(\Gamma)$ is a well-founded relation, $min(\mathbf{O}(\Gamma), \Lambda_{\mathcal{M}})$ is non-empty. Let $\mu \in min(\mathbf{O}(\Gamma), \Lambda_{\mathcal{M}})$. By the normality of $\mathcal{M}$, there is some ψ, such that $\mu \notin mod_{\mathcal{M}}(\psi)$. Hence $\Gamma \not\vdash^{\mathbf{O}}_{\mathcal{M}} \psi$.

Order generators will be constructed using (any number of) *profile settings*, consisting of two ingredients. The first is a *profile*, assigned to each valuation μ, which contains information to judge how "well-behaved" this valuation is with respect to a given knowledge base KB. Making this idea more concrete, a possible profile is a set of formulas, on which μ is "well-behaved" (the latter can mean different things: satisfaction of a formula, assigning a particular truth-value to a formula, etc.). To keep this set finite, we only use formulas which are "relevant" (this notion can also vary, and is defined below in precise terms) for KB. The second ingredient of a profile setting is some ordering on profiles (or finite sets of formulas). Since a profile is theory-dependent, we encapsulate this dependency (like in order generators) by defining *profile matching*:

Definition 10. Let $\mathcal{M}$ be a matrix for $\mathcal{L}$.

- A *profile matching* for $\mathcal{M}$ is a function $\mathcal{J}$, which given an $\mathcal{M}$-valuation μ and a theory Γ, returns a finite set of formulas, and which satisfies the boundedness condition: for every theory Γ, there is some number n_Γ, such that for all $\mu \in \Lambda_{\mathcal{M}}$, $|\mathcal{J}(\mu, \Gamma)| \leq n_\Gamma$.
- A *profile setting* for $\mathcal{M}$ is a pair $\mathcal{S} = \langle \mathcal{J}, \leq \rangle$, where $\mathcal{J}$ is a profile matching for $\mathcal{M}$ and $\leq$ a partial well-founded order on $\mathcal{T}_{\mathcal{L}}$.

A profile setting can then be used to define natural orderings among valuations. Moreover, we can define more fine-grained orderings by applying a common technique in choice theory based on lexicographic aggregation of orderings ([1]):

Definition 11. *Let $\mathcal{S}_1 = \langle \mathcal{J}_1, \leq_1 \rangle, \ldots, \mathcal{S}_n = \langle \mathcal{J}_n, \leq_n \rangle$ be profile settings for $\mathcal{M}$. $\mathbf{O}_{\mathcal{S}_1,\ldots,\mathcal{S}_n}$ is defined inductively as follows:*

- $\mathbf{O}_{\mathcal{S}_1}(\Gamma)(\mu, \nu)$ *iff* $\mathcal{J}_1(\mu, \Gamma) \leq_1 \mathcal{J}_1(\nu, \Gamma)$.
- *For* $n > 1$, $\mathbf{O}_{\mathcal{S}_1,\ldots,\mathcal{S}_n}(\Gamma)(\mu, \nu)$ *iff* $\mathbf{O}_{\mathcal{S}_1,\ldots,\mathcal{S}_{n-1}}(\Gamma)(\mu, \nu)$ *and if also* $\mathbf{O}_{\mathcal{S}_1,\ldots,\mathcal{S}_{n-1}}(\Gamma)(\nu, \mu)$, *then* $\mathcal{J}_n(\mu, \Gamma) \leq_n \mathcal{J}_n(\nu, \Gamma)$.

The intuitive idea behind lexicographic aggregation is that the preference expressed by $\mathcal{S}_{i-1}$ is more important than the one expressed by $\mathcal{S}_i$. Below we shall use this idea to capture the fact that approximating the behaviour of a model of the knowledge base is more important than approximating a classical behaviour. Note that the fact that profile settings are based on well-founded orders ensures that $\mathbf{O}_{\mathcal{S}_1,\ldots,\mathcal{S}_n}$ is indeed an order generator.

Now we provide some concrete examples of the two ingredients of a profile setting. Some natural examples of orders on $\mathcal{T}_{\mathcal{L}}$ are (i) $\Gamma_1 \leq_c \Gamma_2$ if $|\Gamma_1| \leq |\Gamma_2|$, and (ii) $\Gamma_1 \leq_i \Gamma_2$ if $\Gamma_1 \subseteq \Gamma_2$. To define profile matchings, we will need the notion of contexts from [4], capturing the intuitive idea of "relevance" for a given knowledge base.

Definition 12. A *context* is a finite set of formulas (i.e., an element of $\mathcal{T}_{\mathcal{L}}$). A *context generator* (for $\mathcal{L}$) is a function $\mathcal{G} : \mathcal{T}_{\mathcal{L}} \to \mathcal{T}_{\mathcal{L}}$, producing a context for every theory.

Example 8. Common examples for context generators are, e.g., the following functions defined for every theory Γ by $\mathcal{G}^{\mathsf{At}}(\Gamma) = \mathsf{Atoms}(\Gamma)$ (denoting the atoms of Γ), $\mathcal{G}^{\mathsf{ID}}(\Gamma) = \Gamma$, and $\mathcal{G}^{\mathsf{SF}}(\Gamma) = \mathsf{SF}(\Gamma)$ (denoting the subformulas of Γ).

Now we are ready to define the profile matching $\mathcal{J}^{\mathcal{G}}_{\mathcal{Y}}$ induced by a context generator $\mathcal{G}$ and a subset $\mathcal{Y}$ of truth-values of $\mathcal{M}$. The intuition behind the captured preference is that we want to minimize in some sense (according to the chosen order relation) the formulas "relevant" to our knowledge base (according to the generator $\mathcal{G}$) which are assigned values from $\mathcal{Y}$. For instance, by choosing $\mathcal{Y}$ to be the non-designated truth-values of $\mathcal{M}$, we "maximize the satisfaction" of the knowledge base, preferring valuations which approximate models of it. By choosing $\mathcal{Y}$ to contain the inconsistent truth-value $\top$, we "maximize classical reasoning", preferring valuations which best approximate classical valuations. The most natural approach seems to us having "model approximation" as a primary objective, which can then be refined by other criteria.

Definition 13. *Let $\mathcal{M} = \langle \mathcal{V}, \mathcal{D}, \mathcal{O} \rangle$ be a matrix, $\mathcal{Y} \subseteq \mathcal{V}$, and $\mathcal{G}$ - a context generator. The profile maching $\mathcal{J}^{\mathcal{G}}_{\mathcal{Y}}$ is defined as follows: $\mathcal{J}^{\mathcal{G}}_{\mathcal{Y}}(\mu, \Gamma) = \{\psi \in \mathcal{G}(\Gamma) \mid \mu(\psi) \in \mathcal{Y}\}$.*

Example 9. Let $\mathcal{M} = \langle \mathcal{V}, \mathcal{D}, \mathcal{O} \rangle$ be a matrix. Denote $\overline{\mathcal{D}} = \mathcal{V} \setminus \mathcal{D}$. Then $\mathcal{J}_{mod} = \mathcal{J}^{\mathcal{G}^{\mathsf{ID}}}_{\overline{\mathcal{D}}}$ is a profile matching for $\mathcal{M}$. Moreover, in the case that $\mathcal{V} = \{t, f, \top\}$, so are $\mathcal{J}_{con} = \mathcal{J}^{\mathcal{G}^{\mathsf{Atoms}}}_{\{\top\}}$ and $\mathcal{J}_{true} = \mathcal{J}^{\mathcal{G}^{\mathsf{ID}}(\Gamma)}_{\{f,\top\}}$.

Example 10. Let $\Gamma = \{p, \neg p, \circ p, q, \circ q, \neg(p \wedge q)\}$. Let $\mathcal{M}$ be the matrix $\mathcal{M}_{\mathbf{B}}$ from Example 2. To demonstrate the difference between various profile settings, consider the cases $\mathcal{S}^{con}_c = \langle \mathcal{J}_{con}, \leq_c \rangle$, $\mathcal{S}^{mod}_i = \langle \mathcal{J}_{mod}, \leq_i \rangle$, and $\mathcal{S}^{mod}_c = \langle \mathcal{J}_{mod}, \leq_c \rangle$. Note that the "relevant" formulas for $\mathcal{J}_{mod}$ are all the formulas from Γ, while the "relevant" formulas for $\mathcal{J}_{con}$ are only the atoms occurring in Γ. Computing the profiles of the $\mathcal{M}$-valuations, we obtain:

ν_i	p	q	$\neg(p \wedge q)$	$\neg p$	$\circ p$	$\circ q$	$\mathcal{J}_{mod}(\nu_i, \Gamma)$	$\mathcal{J}_{con}(\nu_i, \Gamma)$
ν_1	t	t	f	f	t	t	$\{\neg(p \wedge q), \neg p\}$	$\emptyset$
ν_2	t	$\top$	$\top$	f	t	f	$\{\neg p, \circ q\}$	$\{q\}$
ν_3	t	f	t	f	t	t	$\{q, \neg p\}$	$\emptyset$
ν_4	$\top$	t	$\top$	$\top$	f	t	$\{\circ p\}$	$\{p\}$
ν_5	$\top$	$\top$	$\top$	$\top$	f	f	$\{\circ p, \circ q\}$	$\{p, q\}$
ν_6	$\top$	f	t	$\top$	f	t	$\{q, \circ p\}$	$\{p\}$
ν_7	f	t	t	t	t	t	$\{p\}$	$\emptyset$
ν_8	f	$\top$	t	t	t	f	$\{p, \circ q\}$	$\{q\}$
ν_9	f	f	t	t	t	t	$\{p, q\}$	$\emptyset$

So we have $min(\mathbf{O}_{\mathcal{S}^{mod}_i}, \Lambda_{\mathcal{M}}) = \{\nu_1, \nu_2, \nu_3, \nu_4, \nu_7\}$, while $min(\mathbf{O}_{\mathcal{S}^{mod}_c}, \Lambda_{\mathcal{M}}) = \{\nu_4\}$. Hence, e.g., $\Gamma \nvdash^{\mathbf{O}_{\mathcal{S}^{mod}_i}} \circ q$, while $\Gamma \vdash^{\mathbf{O}_{\mathcal{S}^{mod}_c}} \circ q$. The cautiousness of

$\vdash^{\mathbf{O}_{\mathcal{S}_i^{mod}}}$ can be recovered by using an aggregation-based order. For instance, $min(\mathbf{O}_{\mathcal{S}_i^{mod},\mathcal{S}_c^{con}}, \Lambda_{\mathcal{M}}) = \{\nu_1, \nu_3, \nu_7\}$, and so $\Gamma \vdash^{\mathbf{O}_{\mathcal{S}_i^{mod},\mathcal{S}_c^{con}}} \circ q$.

Below we show some basic properties of the defined entailments. First of all, for $\mathcal{M}$-consistent knowledge bases, the two basic entailments coincide with the logic $\vdash_{\mathcal{M}}$ (while for inconsistent ones they behave differently, as we have just seen in Example 10 above).

Proposition 2. *Let $\mathcal{M}$ be a matrix and Γ - an $\mathcal{M}$-consistent theory. Let $\mathcal{S} = \langle \mathcal{J}_{mod}, \leq \rangle$, where $\leq$ is either $\leq_i$ or $\leq_c$. Then $\Gamma \vdash_{\mathcal{M}} \psi$ iff $\Gamma \vdash_{\mathcal{M}}^{\mathbf{O}_{\mathcal{S}}} \psi$*

The basic entailment induced by $\mathcal{J}_{mod}$ is particularly well-behaved:

Proposition 3. *Let $\mathcal{S}_1 = \langle \mathcal{J}_1, \leq_1 \rangle$, where $\mathcal{J}_1 = \mathcal{J}_{mod}$ and $\leq_1$ is either $\leq_i$ or $\leq_c$. Then $\vdash_{\mathcal{M}}^{\mathbf{O}_{\mathcal{S}}}$ is a cautious consequence relation (with respect to $\mathcal{M}$).*

Proposition 4 (decidability). *We say that a profile setting $\mathcal{S} = \langle \mathcal{J}, \leq \rangle$ for a finite matrix $\mathcal{M}$ is* simple *if $\mathcal{J}$ has the form from Definition 13, and $\leq$ is either $\leq_c$ or $\leq_i$. For simple profile settings $\mathcal{S}_1 = \langle \mathcal{J}_1, \leq_1 \rangle, \ldots, \mathcal{S}_n = \langle \mathcal{J}_n, \leq_n \rangle$ for a finite matrix $\mathcal{M}$, the question whether $\Gamma \vdash_{\mathcal{M}}^{\mathbf{O}_{\mathcal{S}_1,\ldots,\mathcal{S}_n}} \psi$ is decidable.*

Remark 1. In addition to defining new paraconsistent entailments, many natural cases of previously studied paraconsistent formalisms can be simulated within our framework, which is perhaps an indication for the naturality of the above definitions. Below are some examples:

1. For $\mathcal{S}_c = \langle \mathcal{J}_{mod}, \leq_c \rangle$ and any matrix $\mathcal{M}$, $\vdash_{\mathcal{M}}^{\mathbf{O}_{\mathcal{S}}}$ coincides with the distance-based paraconsistent entailment $\mid\!\sim_{\langle \mathsf{Atoms},\mathsf{d},\boldsymbol{\Sigma} \rangle}$ from [4], where d is the standard distance on $\{t, \top, f\}$ (where, e.g., $\mathsf{d}(t, f) = 1$ and $\mathsf{d}(t, \top) = \mathsf{d}(f, \top) = 0.5$).
2. For $\mathcal{S}_i = \langle \mathcal{J}_{mod}, \leq_i \rangle$ and the classical matrix $\mathcal{M}_{cl}$, $\vdash_{\mathcal{M}}^{\mathbf{O}_{\mathcal{S}}}$ coincides with the (propositional fragment) of the cautious entailment based on maximally consistent subsets of [13, 14].
3. Let $\mathcal{M}$ be any paraconsistent three-valued matrix. Let $\mathcal{S}_0$ be either $\mathcal{S}_c$ or $\mathcal{S}_i$. Then the following holds for any $\mathcal{M}$-consistent theory Γ and any ψ: (i) For $\mathcal{S}_1 = \langle \mathcal{J}_{\top}^{\mathsf{At}}, \leq_i \rangle$, $\Gamma \vdash_{\mathcal{M}}^{\mathbf{O}_{\mathcal{S}_0,\mathcal{S}_1}} \psi$ iff $\Gamma \models_P \psi$, (ii) For $\mathcal{S}_1 = \langle \mathcal{J}_{\top}^{\mathsf{At}}, \leq_c \rangle$, $\Gamma \vdash_{\mathcal{M}}^{\mathbf{O}_{\mathcal{S}_0,\mathcal{S}_1}} \psi$ iff $\Gamma \models_{CP} \psi$, (iii) For $\mathcal{S}_1 = \langle \mathcal{J}_{\{\top,f\}}^{\mathsf{ID}}, \leq_i \rangle$, $\Gamma \vdash_{\mathcal{M}}^{\mathbf{O}_{\mathcal{S}_0,\mathcal{S}_1}} \psi$ iff $\Gamma \models_{BS}^{\Gamma} \psi$, (iv) For $\mathcal{S}_1 = \langle \mathcal{J}_{\{\top,f\}}^{\mathsf{ID}}, \leq_c \rangle$, $\Gamma \vdash_{\mathcal{M}}^{\mathbf{O}_{\mathcal{S}_0,\mathcal{S}_1}} \psi$ iff $\Gamma \models_{CBS}^{\Gamma} \psi$, where $\models_P, \models_{CP}, \models_{BS}^{\Gamma}, \models_{CBS}^{\Gamma}$ are the entailments defined in [17, 7] (see also Example 5), and we assume that the set of atoms of $\mathcal{L}$ is finite and is equal to $\mathsf{Atoms}(\Gamma)$.
   It should be noted that there are a number of aspects, in which the above entailments differ from their corresponding counterparts $\models_P, \models_{CP}, \models_{BS}^{\Gamma}$ and $\models_{CBS}^{\Gamma}$. First of all, while the former coincide with the latter for $\mathcal{M}$-consistent premises, they do not trivialize in the presence of $\mathcal{M}$-inconsistency. Secondly, due to the incorporation of context generators, the former do not depend on the (rather restrictive) assumption of finiteness of the underlying language. Finally, the entailments $\models_{BS}^{\Gamma}$ and $\models_{CBS}^{\Gamma}$ are called in [17] non-standard relations, as their definition depends on the belief base Γ under consideration.

The incorporation of the notion of order generators, however, solves this problem, as the dependency of the belief base is encapsulated in it, so the corresponding relations can be thought of as standard in the sense of [17].

Example 11. Revisiting the knowledge base KB_0 from Example 7, let us now apply, e.g., the entailment $\vdash^{\mathbf{O}_{\mathcal{S}_0,\mathcal{S}_1}}_{\mathcal{M}_{\mathbf{B}_\perp}}$ for $\mathcal{S}_0 = \langle \mathcal{J}_{mod}, \leq_c \rangle$ and $\mathcal{S}_1 = \langle \mathcal{J}_{con}, \leq_i \rangle$ (where $\mathcal{M}_{\mathbf{B}_\perp}$ is the matrix from Example 4). The entailment is trivialization-resistant, since despite the fact that KB_0 is not $\mathcal{M}_{\mathbf{B}_\perp}$-consistent, we cannot infer from it $\mathsf{Triangle(africa)}$ (while we can still infer $\mathsf{Triangle(europe)}$). Moreover, the entailment is less cautious than $\vdash_{\mathcal{M}_{\mathbf{B}_\perp}}$ or $\vdash^{\mathbf{O}_{\mathcal{S}_0}}_{\mathcal{M}_{\mathbf{B}_\perp}}$, since, as opposed to the latter, we can infer in it, e.g., $\neg\mathsf{EastBorder(africa, nile)} \to \perp$ from KB_0 (i.e, $\mathsf{EastBorder(africa, nile)}$ is consistently true).

6 Summary and Further Research

In this paper we have proposed an abstract framework for incorporating various "epistemic" preference criteria into paraconsistent entailments, which are useful for overcoming problems such as cautiousness and trivialization. The framework has a number of attractive properties. Most importantly, it allows for defining *trivialization-resistant* variants of paraconsistent entailments based on a many-valued matrix, including the case of preferential ones, which to the best of our knowledge has not been treated before in this context. Moreover, it is *general*, as nothing is assumed about the underlying matrix, and the preference criteria definable in it also have a general form, which can also be further generalized. The framework is also *modular*, as it has a built-in mechanism for easily combining any finite number of preference criteria. We believe that this framework can be a useful tool for the design of paraconsistent knowledge bases. The reason is that for a problem coming from a new, not yet studied domain (like the domain of descriptive ancient geography), it is often hard to choose the right paraconsistent approach (e.g., fixing the matrix, choosing the most natural preference criteria, etc.). In many cases the right choice becomes evident only after some experimental data is available. Our framework, which combines many paraconsistent approaches and allows for a smooth transition between them, can be used for their comparative study.

The most immediate research directions include investigating the logical properties of the entailments defined in our framework, as well as of their computational complexity (from the results of [17] for some special cases, it is clear that intractability is expected for the general case). It should be noted, however, that the current framework aims at generality rather than efficiency, while for improving the latter, entailment-specific optimizations can be employed. For practical applications, the results should also be extended to the first-order case. Concerning the considered domain of ancient geography, it seems promising to combine our approach for handling inconsistent knowledge bases with methods for handling vague geographic concepts along the lines of [6].

References

1. H. Andréka, M. Ryan, and P.-Y. Schobbens. Operators and laws for combining preference relations. *Journal of Logic and Computation*, 12(1):13–53, 2002.
2. O. Arieli. Distance-based paraconsistent logics. *International Journal of Approximate Reasoning*, 48(3):766–783, 2008.
3. 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.
4. O. Arieli and A. Zamansky. Inconsistency-tolerance in knowledge-based systems by dissimilarities. In *Proc. 7th International Symposium on Foundations of Information and Knowledge Systems (FoIKS'12)*, pages 34 – 50. Springer, 2012.
5. A. Avron, O. Arieli, and A. Zamansky. Ideal paraconsistent logics. *Studia Logica*, 99(1-3):31–60, 2011.
6. B. Bennett, D. Mallenby, and A. Third. An ontology for grounding vague geographic terms. In C. Eschenbach and M. Gruninger, editors, *Proceedings of the 5th International Conference on Formal Ontology in Information Systems (FOIS-08)*. IOS Press, 2008.
7. P. Besnard and T. Schaub. Circumscribing inconsistency. In *Proc. of IJCAI'97*, pages 150 – 155, 1997.
8. W. A. Carnielli, M. E. Coniglio, and J. Marcos. Logics of formal inconsistency. In D. M. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume 14, pages 15–107. Springer, 2007. Second edition.
9. W. A. Carnielli, J. Marcos, and S. de Amo. Formal inconsistency and evolutionary databases. *Logic and logical philosophy*, 8:115–152, 2000.
10. N. C. A. da Costa. On the theory of inconsistent formal systems. *Notre Dame Journal of Formal Logic*, 15:497–510, 1974.
11. D. Gabbay. Theoretical foundation for non-monotonic reasoning, Part II: Structured non-monotonic theories. In *Proc. SCAI'91*. IOS Press, 1991.
12. D. Gabbay and A. Hunter. Making inconsistency respectable: A logical framework for inconsistency in reasoning. In *In Fundamentals of Artificial Intelligence*, page pages. Springer, 1992.
13. J. Grant and V.S. Subrahmanian. Reasoning about inconsistent knowledge bases. In *IEEE TKDE*, volume 7(1), pages 177–189, 1995.
14. J. Grant and V.S. Subrahmanian. Applications of paraconsistency for data and knowledge bases. *Synthese*, 125(1/2):121–132, 2000.
15. E. Ilyuschetchkina. Das Weltbild des Dionysios Periegetes. In *Die Vermessung der Oikumene / Mapping the Oikumene*. Berlin, De Gruyter, 2012. forthcoming.
16. S. Konieczny, J. Lang, and P. Marquis. DA2 merging operators. *Artificial Intelligence*, 157(1–2):49–79, 2004.
17. S. Konieczny and P. Marquis. Three-valued logics for inconsistency handling. In *Proc. of JELIA'02*, pages 332 – 344, 2002.
18. S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1–2):167–207, 1990.
19. D. Makinson. General patterns in nonmonotonic reasoning. In *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3, pages 35–110. 1994.
20. De Amo S., W. A. Carnielli, and J. Marcos. A logical framework for integrating inconsistent information in multiple databases. In *International Symposium on Foundations of Information and Knowledge Systems*, pages 67–84. Springer, 2002.
21. Y. Shoham. *Reasoning about Change*. MIT Press, 1988.