

Germane Conflicts: Desirable Properties for Localising Inconsistency

Glauber De Bona¹, Anthony Hunter²

¹Polytechnic School of the University of São Paulo, Av. Prof. Luciano Gualberto 158, São Paulo, Brazil

²Department of Computer Science, University College London, Gower Street, London, UK
glauberbona@gmail.com, anthony.hunter@ucl.ac.uk

Abstract

Inconsistency is a common problem in knowledge, and so there is a need to analyse it. Inconsistency measures assess its severity, but there is a more basic question: “where is the inconsistency?”. Typically, not all subsets of a knowledgebase are causing the inconsistency, and minimal inconsistent sets have been the standard way to localise the *germane* ones, even though there are shortcomings in some scenarios. Recently, \star -conflicts were proposed as a more suitable definition to localise inconsistency when considering a method to repair it. But in general there is no way to tell what is a sensible definition to capture the germane conflicts. This work provides a set of desirable properties to assess definitions for germane conflicts. Also, a new conflict definition, based on substitution, is presented and evaluated via the proposed properties, and the related computational complexity is analysed.

1 Introduction

Inconsistency is common when formally representing data, facts about the world, beliefs, etc. Even though paraconsistent logics (Belnap 1977; Priest 2002; Carnielli, Coniglio, and Marcos 2007) allow for conclusions to be derived from inconsistent premises, this is typically not the case in classical inference systems, which calls for ways to analyse the inconsistency. Given a set of inconsistent pieces of information, one might ask where is the inconsistency and how severe it is. Inconsistency measures have been proposed to tackle the latter (Knight 2002; Hunter 2002; Hunter and Konieczny 2005; Doder et al. 2010; Grant and Hunter 2011; McAreavey, Liu, and Miller 2014; Thimm 2016; Thimm and Wallner 2016; Grant and Hunter 2017, 2023), but the former seems underdeveloped, receiving little attention (De Bona and Hunter 2017; Mu 2024).

In classical logic, typically it is not the case that all inconsistent subsets of a knowledgebase are causing the inconsistency, hence the need to determine the *germane* ones. Traditionally, inconsistency has been localised via minimal inconsistent/unsatisfiable (sub)sets – or *minimal conflicts* –, which have been called the purest form of inconsistency and the causes of it (Hunter and Konieczny 2008, 2010). Conversely, formulas that do not take part in minimal conflicts are said to be *free* of contradiction. The central role of

these minimal conflicts is justified by the fact that removing a formula from each such subset restores the consistency of the whole set, and no free formula needs to be withdrawn. This is at the core of Reiter’s (1987) diagnosis problem and the AGM belief revision paradigm (Alchourrón, Gärdenfors, and Makinson 1985), for instance.

For instance, consider the following sentences: “Alice is a dog”, “Alice is not a dog;” and “It is raining”. The triple is inconsistent, but it intuitively seems that “It is raining” has nothing to do with the inconsistency. In fact, the only minimal conflict is the first pair, and the third sentence is free. That is, the inconsistency is somehow localised in the first pair, which would be the only subset germane to the inconsistency of the knowledgebase.

Despite their ubiquity and usefulness, minimal conflicts are not suitable for capturing the subsets that are germane to the inconsistency of the knowledgebase in all contexts and scenarios. In particular, if the inconsistency can be repaired via weakening formulas, minimal conflicts can fail to capture problematic formulas, as illustrated in the following example (De Bona and Hunter 2017):

Example 1 *The police are investigating a robbery at a jewellery shop, and the employees give the following testimonies:*

- *salesperson: “I did not open the safe, and the criminals carried no guns!”*
- *security chief: “Only the manager or the salesperson could have opened the safe, and the criminals carried guns.”*
- *manager: “I did not open the safe.”*

The set is inconsistent, as the security chief and the salesperson are clearly contradicting each other, forming the only minimal conflict, while the manager’s testimony is free. Assuming this is the only germane conflict, the police could focus on the salesperson and the security chief in the search for the complicit employees, for at least one appears to be lying, ignoring the manager. The investigators could meet the salesperson and the security chief in order to discuss whether or not the criminals carried guns. The salesperson could admit the possibility of not having noticed them, updating their testimony to “I did not open the safe”. But the new set of testimonies is now a minimal conflict, and maybe both the salesperson and the security chief have always been

telling the truth about the safe opening. That would mean the salesperson did not open it, but the manager did, implying the manager was the one lying, and being complicit.

In the example above, we have a type of *iceberg inconsistency* (De Bona and Hunter 2017), hidden by the presence of a smaller minimal conflict. Only when the salesperson weakens their testimony that conflict is revealed, indicating that the set containing all three original testimonies was also germane to the inconsistency. This is because minimal conflicts capture all problematic formulas when consistency is to be restored by removing, not by weakening formulas.

Another type of iceberg inconsistency, not involving free formulas, is illustrated in the following example:

Example 2 *A dad is planning a birthday party for his triplet children, called Albert, Bob and Charlie. To pick a theme for the party, he wants to hear his kids' suggestions and asks them to think about it. Albert wants a party about cars, and Bob, superheroes; and both could be pleased with a superhero car theme. Charlie, however, is tired of past parties about cars and superheroes, and just wants something different (not cars nor superheroes).*

The three requirements are inconsistent, with Charlie's contradicting each of the other two, which are consistent between themselves. In other words, the constraints form two minimal conflicts: Albert's and Charlie's; and Bob's and Charlie's. To solve both minimal conflicts, Charlie accepts a party that does not involve both cars and superheroes at the same time, being possibly either about cars or about superheroes. Even though both minimal conflicts are resolved, the triple is still inconsistent.

Charlie's requirement in the example above is being weakened from a conjunction (not cars and no superheroes) to a disjunction (not cars or no superheroes). Again, minimal conflicts by themselves could not capture the germaneness of the initial requirements to the inconsistency.

Iceberg inconsistencies can also be formed when the meaning of a proposition is instantiated, specified or replaced, as in the situation below:

Example 3 *Three companies (A, B and C) are being merged to form a new one. They have different rules about who should be on the board of directors:*

- Company A: "The CEO must not be on the board."
- Company B: "The CFO must be on the board."
- Company C: "If the CFO is on the board, then the CEO must also be."

The directors of the new company have been appointed and a board has to be formed. To save costs, the company will initially have the same person (Jane) as CFO and CEO. Thus, the guidelines about the board are instantiated as:

- Company A: "Jane must not be on the board."
- Company B: "Jane must be on the board."
- Company C: "If Jane is on the board, then Jane must also be."

They realise the board cannot meet all restrictions from the three companies, for the first two contradict each other, forming the only minimal conflict – and the third one is a

tautology. They let Jane on the board, thinking the problem would resolve itself when eventually a different CEO is appointed, which would not be on the board. Nonetheless, the conflict would remain, as a new minimal conflict would involve all three rules.

Company C's guideline in Example 3 is germane to the whole conflict, but this is hidden by the fact that the CFO and the CEO are the same person, when it becomes trivially satisfied. We would like a germane conflict definition that captures cases like this, when a substitution concentrates the inconsistency of a minimal conflict in a subset of it, creating a smaller minimal conflict and some free formulas.

More generally, we would like definitions of germane conflict that capture problematic sets of formulas, where all formulas are contributing to the inconsistency, like the triples in Examples 1, 2 and 3, and that leave out sets where clearly there are innocent formulas, like the triple in Alice's example. Nonetheless, a myriad of definitions for when a conflict is germane to the inconsistency of a knowledgebase can be devised in order to behave as expected in particular examples. Yet, we would like to employ definitions that generalise well to other cases. That is, we expect a germane conflict definition to exhibit some desirable properties.

In order to guide the choice of a conflict definition or the formulation of entirely new methods to localise inconsistency, this work puts forward a set of desirable properties for germane conflicts. Also, a new approach to localise inconsistency, based on a desirable property, is proposed to handle cases like Example 3. The new definition is then assessed with respect to the desirable properties for germane conflicts and with respect to its computational complexity.

The remainder of the paper is organised as follows: Section 2 covers the technical preliminaries; Section 3 reviews \star -conflicts; Section 4 presents a range of properties for germane conflicts; Section 5 formulates substitution conflicts; and Section 6 reviews related work.

2 Preliminaries

We focus on classical propositional logic, whose language \mathcal{L} is inductively built from a countable set of atoms $X = \{x_1, x_2, \dots\}$ (usually denoted by a, b, c, \dots) with the standard connectives (\neg, \wedge, \vee and \rightarrow), possibly with parentheses. A *clause* is a disjunction of *literals*, which are possibly negated atoms. A *knowledgebase* (KB) is a finite set $\Gamma \subseteq \mathcal{L}$, and \mathbb{K} denotes the set of all KBs.

A (letter-to-letter) *substitution* is a function $s : X \rightarrow \mathcal{L}$ ($s : X \rightarrow X$), which can be extended to $s : \mathcal{L} \rightarrow \mathcal{L}$ inductively: $s(\neg\phi) = \neg s(\phi)$ and $s(\phi \bowtie \psi) = s(\phi) \bowtie s(\psi)$ for $\bowtie \in \{\wedge, \vee, \rightarrow\}$. For any substitution s and KB Γ , $s(\Gamma)$ denotes the KB $\{s(\phi) \mid \phi \in \Gamma\}$. When we define a specific substitution, we assume $s(x_i) = x_i$ for all $x_i \in X$ unless stated otherwise.

An interpretation is a function $i : X \rightarrow \{0, 1\}$ (0 denotes FALSE; 1, TRUE), which can be extended to $i : \mathcal{L} \rightarrow \{0, 1\}$ as usual. A formula $\phi \in \mathcal{L}$ is *consistent/satisfiable* (a *tautology*) if $i(\phi) = 1$ for some (all) interpretations(s) i . For any $\phi \in \mathcal{L}$, when $i(\phi) = 1$, we say i is a *model* of (or *satisfies*) ϕ ; $\text{Models}(\phi)$ denotes the set of all models of ϕ . A KB Γ is

consistent if there is an interpretation i such that $i(\phi) = 1$ for all $\phi \in \Gamma$ – when we say i *satisfies* Γ .

3 \star -conflicts

The concept of \star -conflict (De Bona and Hunter 2017) is based on consolidation methods and was proposed to generalise minimal conflicts in order to tackle scenarios where the latter fall short. In this section, we review the main results of the definition.

Each definition for \star -conflicts is based on a specific consolidation procedure, formalised by a consequence operation $Cn^\star : \mathbb{K} \rightarrow \mathbb{K}$, which is a parameter of the framework. It is generally assumed that Cn^\star is Tarskian¹ and subclassical, meaning that any $\phi \in Cn^\star(\Gamma)$ must be a classical consequence of Γ ($\Gamma \cup \{\neg\phi\}$ is inconsistent). For simplicity, we will here focus on the case where Cn^\star is *modular*: $Cn^\star(\Gamma) = \bigcup \{Cn^\star(\phi) \mid \phi \in \Gamma\}$.

Definition 1 Given a modular consequence operation Cn^\star , a KB $\Gamma \in \mathbb{K}$ is a \star -conflict if there is a minimal conflict Δ and a surjective function $f : \Delta \rightarrow \Gamma$ such that $\phi \in Cn^\star(\{f(\phi)\})$ for all $\phi \in \Delta$. A formula $\phi \in \Psi$ is \star -free in the KB $\Psi \in \mathbb{K}$ if there is no \star -conflict $\Psi' \subseteq \Psi$ such that $\phi \in \Psi'$.

By definition, for all Tarskian Cn^\star , any minimal conflict is also a \star -conflict, thus \star -free formulas are always free. Furthermore, if Cn^\star is the identity function $Cn^{Id} : \mathbb{K} \rightarrow \mathbb{K}$, \star -conflicts are exactly the minimal conflicts.

Example 4 Recall the situation of Example 1 (De Bona and Hunter 2017) and consider the following atomic propositions:

- s stands for “the salesperson opened the safe”;
- m stands for “the manager opened the safe”;
- g stands for “the criminals carried guns”.

Now we can formalise the testimonies:

- *salesperson*: $\phi = \neg s \wedge \neg g$;
- *security chief*: $\psi = (s \vee m) \wedge g$;
- *manager*: $\theta = \neg m$.

Note that $\Delta = \{\neg s \wedge \neg g, (s \vee m) \wedge g\}$ is the only minimal conflict in the set of testimonies $\Gamma = \Delta \cup \{\neg m\}$, thus $\theta = \neg m$ is free in Γ . However, if Cn^\star “breaks” conjunctions ($\neg s \in Cn^\star(\{\neg s \wedge \neg g\})$), we have that the minimal conflict $\Gamma' = \{\neg s, (s \vee m) \wedge g, \neg m\}$ (the updated set of testimonies) is formed by taking a \star -consequence of each formula in Γ , and the latter is a \star -conflict. Consequently, the manager’s testimony is not \star -free.

The example above presents a \star -conflict for a consequence operation that can break conjunctions. This can be captured by a modular consequence operation $Cn^\wedge : \mathbb{K} \rightarrow \mathbb{K}$ such that $Cn^\wedge(\{\phi\})$ contains the conjuncts of ϕ . Formally, $\psi \in Cn^\wedge(\{\phi\})$ if there are $\alpha, \beta \in \mathcal{L}$ such that $\phi = \alpha \wedge \psi$, or $\phi = \psi \wedge \beta$ or $\phi = \alpha \wedge \psi \wedge \beta$ or $\phi = \psi$. Now, for any knowledgebase $\Gamma \in \mathbb{K}$, modularity implies

¹For all KBs $\Gamma, \Delta \in \mathbb{K}$: $\Gamma \subseteq Cn^\star(\Gamma)$; $Cn^\star(Cn^\star(\Gamma)) \subseteq Cn^\star(\Gamma)$; and $\Gamma \subseteq \Delta$ implies $Cn^\star(\Gamma) \subseteq Cn^\star(\Delta)$.

$Cn^\wedge(\Gamma) = \bigcup \{Cn^\wedge(\phi) \mid \phi \in \Gamma\}$. Note that Cn^\wedge is subclassical and Tarskian, as any formula is a conjunct of itself.

Returning to Example 2, let the atoms c (cars) and h (superheroes) encode Albert’s and Bob’s requirements. The original requirement by Charlie then is $\neg c \wedge \neg h$. The knowledgebase $\Gamma = \{c, h, \neg c \wedge \neg h\}$ is inconsistent, with two minimal conflicts: $\{c, \neg c \wedge \neg h\}$ and $\{h, \neg c \wedge \neg h\}$. If $\neg c \vee \neg h \in Cn^\star(\{\neg c \wedge \neg h\})$, Γ is a \star -conflict. However, note that this is not the case for $Cn^\star = Cn^\wedge$.

To capture conflicts like the one in Example 2, we can employ, for instance, a Cn^\star based on dilation (Bloch et al. 2023). We start by defining the (Manhattan) distance between two interpretations i, i' as $d(i, i') = \sum_{i=1}^n |i(x_i) - i'(x_i)|$. For a $k \in \mathbb{N}$ and a formula $\phi \in \mathcal{L}$, the k -dilation of the set of models of ϕ can then be defined as the set of interpretations whose distance to $Models(\phi)$ is at most k ; formally, $M^k(\phi) = \{i \mid \inf_{i' \in Models(\phi)} d(i, i') \leq k\}$ ². Now the modular consequence operation Cn^{dil} can be defined as $Cn^{dil}(\{\phi\}) = \{\psi \mid Models(\psi) = M^k(\phi) \text{ for some } k \in \mathbb{N}\}$ for all $\phi \in \mathcal{L}$ and, for all $\Gamma \in \mathbb{K}$, modularity implies $Cn^{dil}(\Gamma) = \bigcup \{Cn^{dil}(\phi) \mid \phi \in \Gamma\}$.

The \star -conflicts implied by $Cn^\star = Cn^{dil}$ are determined purely via the models of the formulas involved. For instance, consider the knowledgebase from Example 2, $\Gamma = \{c, h, \neg c \wedge \neg h\}$. Since $\neg c \vee \neg h \in Cn^{dil}(\{\neg c \wedge \neg h\})$, Γ is a \star -conflict. Rewriting $\neg c \wedge \neg h$ as $\neg c \wedge \neg h \wedge (\neg c \vee \neg h)$ (or anything equivalent to $\neg c \wedge \neg h$) does not prevent Γ from being a \star -conflict, as $Cn^\star(\phi) = Cn^{dil}(\phi)$ depends only on the interpretations satisfying ϕ .

The strongest modular, subclassical, Tarskian consequence operation corresponds to the classical consequence relation \vdash . We call this consequence operation $Cn^{mod} : \mathbb{K} \rightarrow \mathbb{K}$, as it is the modular version of the classical Cn . It can be formally defined, for any formula $\phi \in \mathcal{L}$, as $Cn^{mod}(\{\phi\}) = Cn(\{\phi\})$, and its definition for knowledgebases follows from its modularity. The corresponding \star -conflict definition is the most inclusive, capturing irrelevant inconsistent knowledgebases. For instance, formalising Alice’s example, let a denote “Alice is a dog” and b denote “it is raining”. The sentences in that example translate to the knowledgebase $\Gamma = \{a, \neg a, b\}$. Even though b seems not involved in the inconsistency, $Cn^\star = Cn^{mod}$ makes the whole Γ a \star -conflict, for $\{a \vee \neg b, \neg a, b\}$ is a minimal conflict and $a \vee \neg b \in Cn^{mod}(\{a\})$. This is in fact a particular case of a stronger result:

Proposition 1 ((De Bona and Hunter 2017)) Assuming $Cn^\star = Cn^{mod}$, a knowledgebase $\Gamma \in \mathbb{K}$ is a \star -conflict iff it is inconsistent and contains no tautologies.

Given that even the most inclusive instance of \star -conflict (with $Cn^\star = Cn^{mod}$) never contains a tautology, it is expected that no \star -conflict does, which is indeed the case:

Proposition 2 ((De Bona and Hunter 2017)) For any modular, subclassical, Tarskian consequence operation Cn^\star , \star -conflicts do not contain tautologies.

²We use infimum for the distance to the empty set to be well defined (as ∞).

If we formalise with propositional logic the restrictions in Example 3 about Jane being on the board, company C's rule ("If Jane is on the board, then Jane is on the board.") yields a tautology. Hence, Proposition 2 implies that no \star -conflict can contain that formula, no matter the underlying consequence operation Cn^* . Thus, an entirely new approach is needed to detect that conflict, as we will see in Section 5.

There is a myriad of \star -conflict definitions, corresponding to the infinitely many possible choices for the underlying Cn^* . If one is to repair the inconsistency of a knowledgebase by weakening formulas with a consequence operation Cn^* , that is the obvious choice, as \star -free formulas could then be ignored during the consolidation (De Bona and Hunter 2017). Nevertheless, without an underlying consolidation procedure, it is not clear which \star -conflict definition should be employed, and other, perhaps new conflict definitions could be used. In this broader setting, although, we would lack a way to evaluate whether the conflict definition in hands is reasonable. To address these questions, a set of desirable properties for conflict definition can ground our assessment of particular instances of \star -conflicts and any other conflict definitions to be proposed.

4 Desirable Properties

We are looking for properties of definitions of conflicts that can capture when a subset of a knowledgebase is somewhat problematic, or relevant in some sense to its inconsistency. We will call these subsets *germane conflicts*. Given an arbitrary definition for germane conflicts, we want to assess it via the set of desirable properties it satisfies. Depending on the context and on the application, some properties of germane conflicts might have higher priority than others, so these desiderata can guide the choice of the conflict definition to be employed. Similarly, some intended behaviour for germane conflicts in a given scenario, characterised via the properties they should hold, can drive the formulation of new definitions to tackle that specific situation.

We are interested in analysing when a given set of formulas $\Delta \in \mathbb{K}$ forms a germane conflict, regardless of any larger context $\Gamma \in \mathbb{K}$ where Δ might appear. This is because classical logic is monotonic, and the extra formulas in $\Gamma \setminus \Delta$ should not impact the conflict Δ .

Any well-posed conflict definition yields a set of knowledgebases regarded as germane conflicts. The properties to be proposed for the definition will constrain this set of all germane conflicts, denoted by \mathcal{C} . For presentation purposes, the properties are clustered into five categories:

- **Core Properties** are the most basic ones, expected to hold for any reasonable conflict definition;
- **Set-theoretical Properties** are based on set operations;
- **Semantic-based Properties** capture behaviours related to the semantics of the formulas, such as logical equivalence or implication;
- **Properties based on Inconsistency Measures** states relations between those measures and conflict definitions;
- **Atoms-related Properties** are related to the atoms in a formula and to substitutions applied to atoms.

Core Properties

The most basic thing to expect from a germane conflict is that it is indeed inconsistent.

Property 1 (Inconsistency) *If $\Gamma \in \mathcal{C}$, then Γ is inconsistent.*

We also expect that any inconsistent knowledgebase includes a germane conflict as a subset. That is, germane conflicts can be seen as "causes" of inconsistency, for their removal restores consistency.

Property 2 (Inclusion) *If $\Delta \in \mathbb{K}$ is inconsistent, then there is a $\Gamma \in \mathcal{C}$ such that $\Gamma \subseteq \Delta$.*

This simple pair of properties already implies the central role of minimal conflicts, as the following proposition points out.

Proposition 3 *If (Inconsistency) and (Inclusion) hold, then every minimal conflict in \mathbb{K} is also in \mathcal{C} .³*

Another basic property that germane conflicts should exhibit deals with their indifference regarding the atoms employed. For instance, if we replace all x_i with x_{i+1} in a germane conflict, it should remain a germane conflict.

Property 3 (Atom-indifference) *If $\Gamma \in \mathcal{C}$ and $s : X \rightarrow X$ is a bijective letter-to-letter substitution, then $s(\Gamma) \in \mathcal{C}$.*

Note that the substitution in the property above is bijective, so different letters should continue to be different after the substitution. The property of (Atom-indifference) also implies that a knowledgebase that is not a germane conflict will not become one if atoms are renamed:

Proposition 4 *Assuming (Atom-indifference), given any bijective letter-to-letter substitution $s : X \rightarrow X$, if $\Gamma \in \mathbb{K} \setminus \mathcal{C}$, then $s(\Gamma) \in \mathbb{K} \setminus \mathcal{C}$.*

In order to satisfy (Inconsistency), (Inclusion) and (Atom-indifference), one could spuriously assume that every inconsistent knowledgebase is a germane conflict. This possibility is ruled out by the following property.

Property 4 (Non-Triviality) *There exists some inconsistent $\Gamma \in \mathbb{K} \setminus \mathcal{C}$.*

This basic set of properties is satisfied by \star -conflicts and, as a particular case, by minimal conflicts.

Proposition 5 *For any Tarskian, modular, sublassical consequence operation $Cn^* : \mathbb{K} \rightarrow \mathbb{K}$, the set of \star -conflicts satisfy (Inconsistency), (Inclusion), (Atom-indifference) and (Non-Triviality).*

Set-theoretical Properties

Since conflicts are sets of formulas, we can apply set theory tools to state some (un)desirable behaviour for them. For instance, one might argue that germane conflicts should not be properly included in other germane conflicts:

Property 5 (Minimality) *If $\Gamma \in \mathcal{C}$, then there is no $\Delta \in \mathcal{C}$ such that $\Delta \subsetneq \Gamma$.*

³Proofs can be found at <https://tinyurl.com/22jtjah2>.

This property actually characterises germane conflicts as minimal conflicts, given the basic properties:

Proposition 6 *If (Inconsistency), (Inclusion) and (Minimality) hold, then $\Gamma \in \mathcal{C}$ iff Γ is a minimal conflict.*

Consequently, any \star -conflict definition that captures a set that is not a minimal conflict violates (Minimality). That is, (Minimality) is only satisfied by the set of \star -conflicts with $Cn^\star = Cn^{Id}$.

Intuitively, when a conflict is germane, all of its formulas are somehow involved in the inconsistency of the knowledgebase. Thus, the union of the germane conflicts of a knowledgebase is the set of all formulas involved in the inconsistency, which could be seen as a germane conflict. This is captured by the following property:

Property 6 (\cup -closure) *If $\Gamma, \Delta \in \mathcal{C}$, then $\Gamma \cup \Delta \in \mathcal{C}$.*

Minimal conflicts and most \star -conflicts violate (\cup -closure), though, as they are meant to characterise sets where all formulas are contributing to inconsistency.

Proposition 7 *The set of \star -conflicts violates (\cup -closure) for any $Cn^\star \in \{Cn^{Id}, Cn^{dil}, Cn^\wedge\}$ and satisfies it for $Cn^\star = Cn^{mod}$.*

Semantic-based Properties

As minimal conflicts cannot capture hidden iceberg conflicts, as illustrated in Example 1, (Minimality) is too strong a property. Nonetheless, they exhibit a semantic behaviour that might be desirable in most scenarios. Specifically, minimal conflicts are defined entirely via the consistency/inconsistency of subsets, which entails a series of semantic-based properties. For instance, minimal conflicts are not dependent on the actual syntax of the formulas, and that is captured by the property below.

Property 7 (Syntax-Robustness) *Given a knowledgebase $\Gamma \in \mathcal{C}$, if $\phi \in \Gamma$ is replaced by a logically equivalent $\phi' \in \mathcal{L}$, the resulting knowledgebase is also in \mathcal{C} .*

The set of minimal conflicts clearly satisfies (Syntax-Robustness). For \star -conflicts in general, the formula's syntax matters if it does for the corresponding Cn^\star . For instance, a Cn^\star that “breaks” conjunctions is syntactical, but a consequence operation that dilates the models of the formulas is purely semantic.

Proposition 8 *The set of \star -conflicts satisfies (Syntax-Robustness) for any $Cn^\star \in \{Cn^{Id}, Cn^{dil}, Cn^{mod}\}$ and violates it for $Cn^\star = Cn^\wedge$.*

A related semantic property forbids equivalent formulas in a germane conflict.

Property 8 (No-Equivalences) *If Γ is a germane conflict, then there are no formulas $\phi, \psi \in \Gamma$ such that ψ is logically equivalent to ϕ .*

As logical equivalence is a special case of logical implication, a stronger property can be formulated:

Property 9 (Non-Redundancy) *If $\Gamma \in \mathcal{C}$, then there are no formulas $\phi, \psi \in \Gamma$ such that ψ implies ϕ .*

Intuitively, if ϕ logically implies ψ , the latter would seem superfluous in a germane conflict where the former is; the information of ψ would be redundant.

Proposition 9 *The set of \star -conflicts satisfies (Non-Redundancy) for $Cn^\star = Cn^{Id}$, satisfies (No-Equivalences) for $Cn^\star = Cn^{dil}$, violates (Non-Redundancy) for $Cn^\star = Cn^{dil}$ and violates (No-Equivalences) for any $Cn^\star \in \{Cn^\wedge, Cn^{mod}\}$.*

Considering still only the semantics of a formula, we would expect that tautologies do not take part in germane conflicts.

Property 10 (Tautology-Freeness) *If $\Gamma \in \mathcal{C}$, then there is no tautology $\phi \in \Gamma$.*

Tautologies are always free formulas, not appearing in minimal conflicts. Desirably, this property is inherited by all \star -conflicts via Proposition 2.

Even though not containing tautologies is usually a desirable property of a germane conflict, a definition that includes all inconsistent knowledgebases with no tautologies would be rather dull. We call this undesirable behaviour *quasi-triviality*, leading to the following property:

Property 11 (Non-Quasi-Triviality) *There exists some inconsistent $\Gamma \in \mathbb{K} \setminus \mathcal{C}$ that has no tautologies.*

The property of (Non-Triviality) is entailed by (Non-Quasi-Triviality), as the latter is strictly stronger. Although, for any Cn^\star , the set of all \star -conflicts is non-trivial, it quasi-trivialises for a $Cn^\star = Cn^{mod}$, the modular version of the classical consequence operation, due to Proposition 1.

Proposition 10 *The set of \star -conflicts satisfies (Non-Quasi-Triviality) for any $Cn^\star \in \{Cn^{Id}, Cn^\wedge, Cn^{dil}\}$ and violates it for $Cn^\star = Cn^{mod}$.*

Properties based on Inconsistency Measures

Localising and measuring inconsistency can be very closely related. Some inconsistency measures proposed in the literature are based on an underlying notion of conflict. Conversely, we could constrain germane conflicts using an inconsistency measure. For instance, if removing a formula from a knowledgebase decreases its inconsistency measurement, that formula seems to be involved somehow in a kind of conflict. This is formalised in the property below:

Property 12 (\mathcal{I} -Respect) *Given an inconsistency measure $\mathcal{I} : \mathbb{K} \rightarrow \mathbb{R}$ and a knowledgebase $\Gamma \in \mathbb{K}$, if $\mathcal{I}(\Gamma \setminus \{\phi\}) < \mathcal{I}(\Gamma)$ for some $\phi \in \Gamma$, then there is a $\Delta \in \mathcal{C}$ such that $\Delta \subseteq \Gamma$ and $\phi \in \Delta$.*

For instance, consider the drastic measure \mathcal{I}_d , which assigns 1 to inconsistent knowledgebases and 0 to consistent ones. The corresponding property (\mathcal{I}_d -Respect) entails that formulas in minimal conflicts are also in germane conflicts – which, given (Inconsistency), implies that minimal conflicts are germane conflicts.

This property can be seen as restricting inconsistency measures, given a conflict definition. For example, if germane conflicts are simply minimal conflicts, then (\mathcal{I} -Respect) implies that removing free formulas does not decrease the inconsistency measurement, which is a version of

the property called Free-Formula Independence (Hunter and Konieczny 2010).

Atoms-related Properties

A property similar to (Tautology-Freeness) precludes germane conflicts from having formulas that are intuitively “innocent”. Given a knowledgebase $\Gamma \in \mathbb{K}$, with $|\Gamma| \geq 2$, a set of formulas $\Delta \subset \Gamma$ is said to be *safe* in Γ if it is non-empty, consistent and shares no atoms with $\Gamma \setminus \Delta$. For instance $\{x_2, x_3\}$ is safe in $\{x_1, \neg x_1, x_2, x_3\}$.

Property 13 (Safe-Set-Freeness) *If $\Gamma \in \mathcal{C}$, then there is no safe set $\Delta \subset \Gamma$.*

Proposition 11 *The set of \star -conflicts satisfies (Safe-set-Freeness) for any $Cn^* \in \{Cn^{Id}, Cn^\wedge, Cn^{dil}\}$ and violates (Safe-Set-Freeness) for $Cn^* = Cn^{mod}$.*

Another property related to the atoms in a knowledgebase deals with replacing atoms. A general substitution does not need to be bijective, as required in (Atom-indifference), and could replace for instance two different letters with the same. Yet, as the general structure is kept, one might expect that applying a substitution to a germane conflict yields another germane conflict. To see the possible effects of a substitution in a conflict, consider the following example:

Example 5 *Recall the situation in Example 3. Let e and f be atomic propositions denoting “The CEO must be on the board” and “The CFO must be on the board”, respectively. The general guidelines about board members from companies A, B and C can then be formalised in the knowledgebase $\Gamma = \{\neg e, f, f \rightarrow e\}$, which is a minimal conflict. Let j encode the proposition “Jane must be on the board”. Being Jane both the CEO and the CFO, we can replace both e and f by j in Γ . Formally, consider a substitution $s : \mathcal{X} \rightarrow \mathcal{L}$ such that $s(e) = s(f) = j$, so that $s(\Gamma) = \{\neg j, j, j \rightarrow j\}$. Note that the first pair is a minimal conflict, and the third formula is a tautology.*

The general board guidelines in the example above (Γ) form a minimal conflict and, given (Inconsistency) and (Inclusion), they are also a germane conflict. Replacing both “CFO” and “CEO” by “Jane” should intuitively keep the germaneness of the conflict. That is, germane conflicts should be robust to substitution, which is captured by the following property:

Property 14 (Substitution-Robustness) *Given a substitution $s : \mathcal{X} \rightarrow \mathcal{L}$, if Γ is a germane conflict, then so is $s(\Gamma)$.*

When applying a substitution to a minimal conflict, new, smaller minimal conflicts might appear, as illustrated in Example 5, and the whole knowledgebase stops being a minimal conflict. More generally, this sensitivity to substitution holds for all \star -conflicts definitions.

Proposition 12 *For any modular, subclassical, Tarskian consequence operation Cn^* , the set of all \star -conflicts violates (Substitution-Robustness).*

The proposition above is a corollary of a stronger result:

Theorem 1 *There is no set $\mathcal{C} \subseteq \mathbb{K}$ satisfying (Inconsistency), (Inclusion), (Tautology-Freeness) and (Substitution-Robustness).*

As (Substitution-Robustness) is not satisfied by any type of \star -conflict, nor by minimal conflicts, its intuitive appeal calls for an entirely new approach to localise inconsistency, investigated in the next section.

5 Substitution Conflicts

The motivation behind \star -conflicts is to detect conflicts that are not minimal but somehow problematic. That is, a \star -conflict aims to reveal sets of formulas that are in conflict in some way but that are obfuscated or ignored due to a smaller minimal conflict within it. A different approach to detecting conflicts hidden by minimal conflicts employs substitutions from atoms to formulas. Intuitively, if we depart from a minimal conflict, a substitution yields some sort of conflict, as the general structure of the formula is kept fixed. Nevertheless, a substitution applied to a minimal conflict might cause some proper subset to become inconsistent; and the new knowledgebase ceases to be a minimal conflict, as shown in Example 5. Thus, applying a substitution to a minimal conflict might indeed hide the original conflict.

We can formally define the conflicts resulting from applying a substitution to a minimal conflict.

Definition 2 *A KB $\Gamma \in \mathbb{K}$ is a substitution conflict if there is a minimal conflict $\Delta \in \mathbb{K}$ and a substitution $s : \mathcal{X} \rightarrow \mathcal{L}$ such that $s(\Delta) = \Gamma$.*

The rules about Jane being on the board, formalised in Example 5, form a substitution conflict. This is due to the fact that $\Gamma = \{\neg j, j, j \rightarrow j\}$ is the result of applying a substitution to the minimal conflict $\Delta = \{\neg e, f, f \rightarrow e\}$.

Substitution conflicts can capture some \star -conflicts as well. For instance, recall Example 4, where we have $\Gamma = \{\neg s \wedge \neg g, (s \vee m) \wedge g, \neg m\}$. Taking the minimal conflict $\Psi = \{\neg s \wedge a, (s \vee m) \wedge g, \neg m\}$ and the substitution $s(a) = \neg g$, we have that $s(\Psi) = \Gamma$ is a substitution conflict.

Besides being robust to substitutions, substitution conflicts have several desirable properties:

Proposition 13 *The set of all substitution conflicts satisfies (Inconsistency), (Inclusion), (Atom-indifference), (Non-Triviality), (Non-Quasi-Triviality), (Safe-Set-Freeness) and (Substitution-Robustness).*

To satisfy (Substitution-Robustness) together with (Inconsistency) and (Inclusion), the property of (Tautology-Freeness) must be given up. Other desirable properties, particularly the semantic ones, are violated, as a substitution is syntactical.

Proposition 14 *The set of all substitution conflicts violates (Minimality), (\cup -closure), (Syntax-Robustness), (No-Equivalences), (Non-Redundacy) and (Tautology-Freeness).*

The properties of (Inconsistency) and (Inclusion) being satisfied by the set of substitution conflicts imply that every minimal conflict is also a substitution conflict. This can be seen also by considering the trivial substitution that replaces each x_i with the same x_i . This implies that substitution conflicts respect syntactic inconsistency measures, which are a function of the inconsistency graph of the knowledgebase (De Bona et al. 2019), which encodes the structure of its minimal conflicts.

Proposition 15 *For any syntactic inconsistency measure \mathcal{I} , the set of all substitution conflicts satisfy (\mathcal{I} -Respect).*

Since substitution conflicts are defined via substitutions applied to minimal conflicts, if all minimal conflicts are germane, then (Substitution-Robustness) entails that all substitution conflicts are germane as well.

Proposition 16 *If (Inconsistency), (Inclusion) and (Substitution-Robustness) hold, then every substitution conflict is a germane conflict.*

The properties of (Inconsistency), (Inclusion) and (Substitution-Robustness) can be satisfied by definitions that also include knowledgebases that are not substitution conflicts. To characterise the set of substitution conflicts, we have the following result:

Proposition 17 *The smallest set $\mathcal{C} \subseteq \mathbb{K}$ satisfying (Inconsistency), (Inclusion) and (Substitution-Robustness) is the set of all substitution conflicts.*

The proposition above states that substitution conflicts are the most economical way, in the sense of capturing fewer conflicts, to satisfy (Substitution-Robustness) without violating (Inconsistency) or (Inclusion).

In order to assess the computation complexity of detecting substitution conflicts, we make use of a useful characterisation of them. Substitution conflicts could be equivalently defined using only letter-to-letter substitutions, according to the following result:

Theorem 2 *A KB $\Gamma \in \mathbb{K}$ is a substitution conflict iff there is a minimal conflict $\Delta \in \mathbb{K}$ and a letter-to-letter substitution $s : X \rightarrow X$ such that $s(\Delta) = \Gamma$.*

This result can now be employed to give a computational complexity bound to the problem of detecting substitution conflicts. This is due to the fact that, if Γ is substitution conflict, there is a Δ with the same size such that $s(\Delta) = \Gamma$ for some substitution.

Proposition 18 *The problem of deciding whether a given knowledgebase $\Gamma \in \mathbb{K}$ is a substitution conflict is in Σ_2^P .*

This is the same computational complexity of detecting \star -conflicts, given some conditions on Cn^* (De Bona and Hunter 2017).

The problem of detecting \star -conflicts for some particular Cn^* can be given tighter computational complexity bounds. Indeed, when $Cn^* = Cn^{Id}$, \star -conflicts are minimal conflicts and their detection is in $D^P \subseteq \Sigma_2^P$. Furthermore, to detect a \star -conflict when $Cn^* = Cn^{mod}$, by Proposition 1, we can simply check the inconsistency of the whole knowledgebase and the falsifiability of each formula, yielding the same bound.

6 Related Work

Methods for localising inconsistency often appear as intermediate tools for measuring inconsistency, as in the first two related works we discuss. The first related work (Jabbour et al. 2014) defines a conflict based on prime implicates. The second related work (Grant and Hunter 2023) localises inconsistency in atoms, not formulas, and the third (Mu 2019)

uses a similar, atom-based method to localise the problematic/innocent formulas in a knowledgebase. The fourth (Mu 2024) uses minimal conflicts and 3-valued logic to analyse inconsistency in all subsets of a knowledgebase.

Jabbour et al. (2014) have proposed an inconsistency measure based on counting a type of conflict in knowledgebases. The conflict definition they put forward is a pair (Δ, Ψ) , where Δ is a subset of the knowledgebase to be assessed, and Ψ is a minimal conflict derived from Δ . Roughly speaking, Δ should be formed by weakening each formula in Ψ exactly once via discarding prime implicates, which are the strong clauses implied. This is quite similar to a particular instance of the \star -conflicts framework (De Bona and Hunter 2017), though in the latter each formula can be weakened more than once to form a minimal conflict.

While we are interested in localising inconsistency in subsets of the knowledgebase, Grant and Hunter (2023) explore localising it to sets of problematic atoms in order to measure inconsistency. They employ a set of 3-valued logics that extend classical logic while allowing for atoms to be assigned the non-classical truth value B . Their approach is based on minimal atomic subsets, which are minimal sets of atoms that are assigned B in some interpretation satisfying the knowledgebase in hands.

Mu (2019) has proposed an inconsistency localising method to rework the independence postulate for inconsistency measures, which requires that adding free formulas does not alter the measurement. His approach is based specifically on Priest’s 3-valued logic (Priest 1991), also relying on minimal atomic subsets. In a nutshell, Mu suggests that a formula is involved in the inconsistency of a knowledgebase if removing it changes the set of minimal atomic subsets, proposing a corresponding independence postulate. Note that no conflict is defined, but only the problematic portion of a knowledgebase.

In another work, Mu (2024) analyses how a characterisation of inconsistency is spread over the subsets of a knowledgebase. Inconsistency is characterised via either the set of minimal conflicts or the set of minimal models, which are 3-valued interpretations with a minimal number of atoms with non-classical truth values. The interior of inconsistency for a KB Γ is then defined as the collection containing, for all inconsistent $\Gamma' \subseteq \Gamma$, either the set of minimal conflicts of Γ' or its set of minimal models. Although the former cannot identify conflicts with free formulas, the latter approach might reveal some conflicts hidden by minimal conflicts when the inconsistency involves different sets of atoms.

7 Conclusion and Future Work

Localising the inconsistency of a knowledgebase into its subsets can be tackled in a myriad of ways, and this work presents the first set of desirable properties for these conflict definitions. Additionally, substitution conflicts are proposed to address the need for a conflict that is robust to substitution. The proposals in this paper can be used to develop alternative, and in some cases, better ways of dealing with inconsistency, including inconsistency resolution, in applications. Future work includes investigating the use of 3-valued logics for defining conflicts.

Acknowledgments

GDB was financially supported by grant #2023/01165-5, São Paulo Research Foundation (FAPESP). AH was financially supported by Leverhulme Trust grant RPG-2021-182. We would like to thank David Makinson for suggesting the concept of substitution conflict as a way to localise inconsistency and proposing some important related results, such as Theorem 2, including a proof sketch.

References

- Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(2): 510–530.
- Belnap, N. 1977. A useful four-valued logic. In *Modern Uses of Multiple-valued Logic*, 8–37. Reidel.
- Bloch, I.; Lang, J.; Pérez, R. P.; and Uzcátegui, C. 2023. Morphologic for knowledge dynamics: revision, fusion and abduction. *Journal of Applied Non-Classical Logics*, 33(3-4): 421–466.
- Carnielli, W.; Coniglio, M. E.; and Marcos, J. 2007. Logics of formal inconsistency. In *Handbook of Philosophical Logic*, 1–93. Springer.
- De Bona, G.; Grant, J.; Hunter, A.; and Konieczny, S. 2019. Classifying inconsistency measures using graphs. *Journal of Artificial Intelligence Research*, 66: 937–987.
- De Bona, G.; and Hunter, A. 2017. Localising iceberg inconsistencies. *Artificial Intelligence*, 246: 118–151.
- Doder, D.; Rašković, M.; Marković, Z.; and Ognjanović, Z. 2010. Measures of inconsistency and defaults. *International Journal of Approximate Reasoning*, 51(7): 832–845.
- Grant, J.; and Hunter, A. 2011. Measuring the good and the bad in inconsistent information. In *International Joint Conference on Artificial Intelligence (IJCAI’11)*, 2632–2637.
- Grant, J.; and Hunter, A. 2017. Analysing inconsistent information using distance-based measures. *International Journal of Approximate Reasoning*, 89: 3–26.
- Grant, J.; and Hunter, A. 2023. Semantic inconsistency measures using 3-valued logics. *International Journal of Approximate Reasoning*, 156: 38–60.
- Hunter, A. 2002. Measuring Inconsistency in Knowledge via Quasi-Classical Models. In *Proceedings of AAAI’02*, 68–73. AAAI Press / The MIT Press.
- Hunter, A.; and Konieczny, S. 2005. Approaches to Measuring Inconsistent Information. In *Inconsistency Tolerance*, volume 3300 of *Lecture Notes in Computer Science*, 191–236. Springer Berlin Heidelberg. ISBN 978-3-540-24260-4.
- Hunter, A.; and Konieczny, S. 2008. Measuring Inconsistency through Minimal Inconsistent Sets. In *Principles of Knowledge Representation and Reasoning (KR’08)*, 358–366.
- Hunter, A.; and Konieczny, S. 2010. On the measure of conflicts: Shapley inconsistency values. *Artificial Intelligence*, 174(14): 1007–1026.
- Jabbour, S.; Ma, Y.; Raddaoui, B.; and Sais, L. 2014. On the Characterization of Inconsistency Measures: A Prime Implicates Based Framework. In *International Conference on Tools with Artificial Intelligence (ICTAI’14)*, 146–153.
- Knight, K. 2002. Measuring inconsistency. *Journal of Philosophical Logic*, 31(1): 77–98.
- McAreavey, K.; Liu, W.; and Miller, P. 2014. Computational approaches to finding and measuring inconsistency in arbitrary knowledge bases. *International Journal of Approximate Reasoning*, 55(8): 1659–1693.
- Mu, K. 2019. Formulas free from inconsistency: an atom-centric characterization in Priest’s minimally inconsistent LP. *Journal of Artificial Intelligence Research*, 66: 279–296.
- Mu, K. 2024. The interior of inconsistency in a knowledge base. *International Journal of Approximate Reasoning*, 166: 109127.
- Priest, G. 1991. Minimally inconsistent LP. *Studia Logica*, 50: 321–331.
- Priest, G. 2002. Paraconsistent Logic. In Gabbay, D. M.; and Guenther, F., eds., *Handbook of Philosophical Logic*, volume 6, 287–393. Dordrecht: Springer Netherlands. ISBN 978-94-017-0460-1.
- Reiter, R. 1987. A theory of diagnosis from first principles. *Artificial intelligence*, 32(1): 57–95.
- Thimm, M. 2016. On the expressivity of inconsistency measures. *Artificial Intelligence*, 234: 120–151.
- Thimm, M.; and Wallner, J. 2016. Some Complexity Results on Inconsistency Measurement. In *Principles of Knowledge Representation and Reasoning (KR’16)*, 114–123.