

Measuring Inconsistency through Minimal Inconsistent Sets

Anthony Hunter

Department of Computer Science
University College London, UK
a.hunter@cs.ucl.ac.uk

Sébastien Konieczny

CRIL - CNRS
Université d'Artois, Lens, France
konieczny@cril.fr

Abstract

In this paper, we explore the links between measures of inconsistency for a belief base and the minimal inconsistent subsets of that belief base. The minimal inconsistent subsets can be considered as the relevant part of the base to take into account to evaluate the amount of inconsistency. We define a very natural inconsistency value from these minimal inconsistent sets. Then we show that the inconsistency value we obtain is a particular Shapley Inconsistency Value, and we provide a complete axiomatization of this value in terms of five simple and intuitive axioms. Defining this Shapley Inconsistency Value using the notion of minimal inconsistent subsets allows us to look forward to a viable implementation of this value using SAT solvers.

Introduction

The need to develop robust, but principled, logic-based techniques for analysing inconsistent information is increasingly recognized as an important research area for artificial intelligence in particular, and for computer science in general (Bertossi, Hunter, & Schaub 2004). This interest stems from the recognition that the dichotomy between consistent and inconsistent sets of formulae that comes from classical logics is not sufficient for describing inconsistent information.

A number of proposals have been made for measuring the degree of information of a belief base in the presence of inconsistency (Lozinskii 1994; Wong & Besnard 2001; Knight 2003; Konieczny, Lang, & Marquis 2003), and for measuring the degree of inconsistency of a belief base (Grant 1978; Knight 2001; Hunter 2002; Knight 2003; Konieczny, Lang, & Marquis 2003; Hunter 2004; 2003; Grant & Hunter 2006; Hunter & Konieczny 2006; Grant & Hunter 2008). For a review see (Hunter & Konieczny 2004).

These measures are potentially important in diverse applications in artificial intelligence, such as belief revision, belief merging, negotiation, multi-agent systems, decision-support, and software engineering tools. Already, measuring inconsistency has been seen to be a useful tool in analysing a diverse range of information types including news reports (Hunter 2006), integrity constraints (Grant & Hunter 2006), information merging (Qi, Liu, & Bell 2005), databases (Martinez *et al.* 2007), ontologies (Ma *et al.*

2007), software specifications (Barragáns-Martínez, Pazos-Arias, & Fernández-Vilas 2004; Mu *et al.* 2005), and e-commerce protocols (Chen, Zhang, & Zhang 2004).

Each of the current proposals for measuring inconsistency can be described as being one of the following two approaches.

The first approach involves “counting” the minimal number of formulae needed to produce the inconsistency in a set of formulae. The more formulae needed to produce the inconsistency, the less inconsistent the set (Knight 2001). This idea is an interesting one, but it rejects the possibility of a more fine-grained inspection of the (content of the) formulae. In particular, if one looks to singleton sets only, one is back to the initial problem, with only two values: consistent or inconsistent.

The second approach involves looking at the proportion of the language that is touched by the inconsistency in a set of formulae. This allows us to look *inside* the formulae (Hunter 2002; Konieczny, Lang, & Marquis 2003; Grant & Hunter 2006; 2008). This means that two formulae (singleton sets) can have different inconsistency measures. In these proposals one can identify the set of formulae with its conjunction (i.e. the set $\{\varphi, \varphi'\}$ has the same inconsistency measure as the set $\{\varphi \wedge \varphi'\}$). Whilst the lack of syntax sensitivity may be appropriate for some applications, it does mean that the distribution of the contradiction among the formulae is not taken into account.

It seems difficult to build measures that take these two dimensions into account. But, in (Hunter & Konieczny 2006) a unified framework was proposed with this aim. The main point was to define inconsistency values that give the inconsistency of each formula of the base, in contrast to the above inconsistency measures that give the inconsistency of the whole base. This allows us to draw a more precise picture of the inconsistencies of the base. The idea is to start from one of the measures considered in the two approaches above and use it to assign a measure of inconsistency to a set of formulae, and then to use a technique based on cooperative game theory: the Shapley value (Shapley 1953). This then allows us to identify the blame/responsibility of each formula in the inconsistency of the belief base (Hunter & Konieczny 2006). This means for example that we can use a measure from the second approach which considers the proportion of the language touched by inconsistency, and then

using the Shapley value, apportion the blame for the inconsistency in the set to the individual formulae in a principled way.

Against this background, it is interesting to note that the use of minimal inconsistent subsets of a belief base has received much less attention as the basis for defining inconsistency measures. So in this paper, we explore the nature of some interesting measures of inconsistency based on minimal inconsistent subsets of the belief base, and we consider how these new measures relate to the family of Shapley Inconsistency Values.

It has been known for a long time that minimal inconsistent subsets of the base are a cornerstone of analysing inconsistencies. For instance, to recover consistency, one has just to remove one formula from each minimal inconsistent subset (Reiter 1987). For conflict resolution, where the syntactic representation of the information is important, measuring inconsistency in terms of the minimal inconsistent subsets is intuitive and it is informative for deciding how to change the set of formulae through a process such as negotiation, compromise, or resolution. So it should be natural to study how these minimal inconsistent subsets could be used to define measures of inconsistency.

The idea to analyse minimal inconsistent subsets of the belief base was followed in order to define scoring functions that, for each subset K' of a set of formulae K , gives a score that is the number of minimal inconsistent subsets of K that would be eliminated if K' were removed from K . Then this score is used to compare different subsets K' (Hunter 2004).

Obviously, these approaches are syntax sensitive, which for some applications, is necessary. Consider capturing requirements for a new corporate computer system in a process where a user may present his or her requirements in the form of a set of propositional formulae (Hunter & Nuseibeh 1998). Here presenting the set of requirements $\{\alpha, \beta\}$ should be treated differently to the set of requirements $\{\alpha \wedge \beta\}$ since the first set says that there are two requirements, the first being α and the second β , whereas in the second set says that is one requirement, namely $\alpha \wedge \beta$.

In the rest of this paper, we study how to use minimal inconsistent subsets in order to define inconsistency values, that will allow us to define the inconsistency of each formula of the base (like with Shapley Inconsistency Values (Hunter & Konieczny 2006)). To this end we introduce the family of MIV (for MinInc Inconsistency Value), and focus on a very intuitive one: MIV_C . As these values have the same aim as the family of Shapley Inconsistency Values, it is interesting to study the relationship between these two families. Moreover, we show a surprising result: the value MIV_C is in fact a Shapley Inconsistency Value. This result is very interesting as it allows us to state interesting logical properties, and to find a way to implement the Shapley Inconsistency Values using existing automated reasoning technology. The additional interest of this value is that its intuitive logical properties have led us to a complete axiomatization through five intuitive axioms.

Preliminaries

We consider a propositional language $\mathcal{L}$ built from a finite set of propositional symbols $\mathcal{P}$. We use $a, b, c, \dots$ to denote the propositional variables, and Greek letters $\alpha, \beta, \varphi, \dots$ to denote the formulae. An interpretation is a total function from $\mathcal{P}$ to $\{0, 1\}$. The set of all interpretations is denoted $\mathcal{W}$. An interpretation ω is a model of a formula φ , denoted $\omega \models \varphi$, if and only if it makes φ true in the usual truth-functional way. $Mod(\varphi)$ denotes the set of models of the formula φ , i.e. $Mod(\varphi) = \{\omega \in \mathcal{W} \mid \omega \models \varphi\}$. We will use $\subseteq$ to denote the set inclusion, and we will use $\subset$ to denote the strict set inclusion, i.e. $A \subset B$ iff $A \subseteq B$ and $B \not\subseteq A$. We will denote the set of natural numbers by $\mathbb{N}$ and the set of real numbers by $\mathbb{R}$.

Let A and B be two subsets of C , we note $C = A \oplus B$ if A and B form a partition of C , i.e. $C = A \oplus B$ iff $C = A \cup B$ and $A \cap B = \emptyset$.

A *belief base* K is a finite set of propositional formulae. More exactly, as we will need to identify the different formulae of a belief base in order to associate them with their inconsistency value, we will consider belief bases K as vectors of formulae. For logical properties we will need to use the set corresponding to each vector, so we suppose that we have a function such that for each vector $K = (\alpha_1, \dots, \alpha_n)$, $\overline{K}$ is the set $\{\alpha_1, \dots, \alpha_n\}$. As it will never be ambiguous, in the following we will omit the $\overline{}$ and write K as both the vector and the set. We use $\mathcal{K}_{\mathcal{L}}$ to denote the set of belief bases definable from formulae of the language $\mathcal{L}$.

A belief base is consistent if there is at least one interpretation that satisfies all its formulae. If a belief base K is not consistent, then one can define the minimal inconsistent subsets of K as:

$$MI(K) = \{K' \subseteq K \mid K' \vdash \perp \text{ and } \forall K'' \subset K', K'' \not\vdash \perp\}$$

If one wants to recover consistency from an inconsistent base K , then the minimal inconsistent subsets can be considered as the purest form of inconsistency, since to recover consistency, one has to remove at least one formula from each minimal inconsistent subset (Reiter 1987).

A *free formula* of a belief base K is a formula of K that does not belong to any minimal inconsistent subset of the belief base K . This means that this formula has nothing to do with the conflicts of the base.

Inconsistency values defined from minimal inconsistent sets

Apart from (Hunter & Konieczny 2006), existing inconsistency measures allow us to evaluate the amount of inconsistency of a whole base, but not to evaluate the amount of inconsistency of each formula of the base. In other words, existing inconsistency measures do not allow us to evaluate the responsibility of each formula in the inconsistency of the base. Yet it is possible to define some such measures using minimal inconsistent sets. Furthermore, as these minimal inconsistent sets are the parts of the base where inconsistencies lie, it should be natural to use only the minimal inconsistent sets for evaluating the amount of inconsistency of the bases. This is the motivation for the following definition where f

is some function that takes as input a formula α and the set of minimal inconsistent subsets for a belief base K .

Definition 1 A *MinInc Inconsistency Value (MIV)* is a function $MIV : \mathcal{K}_{\mathcal{L}} \times \mathcal{L} \rightarrow \mathbb{R}$ such that $MIV(K, \alpha) = f(\alpha, \text{MI}(K))$ where f is a function of α and $\text{MI}(K)$.

Instances of a MIV (such as those given in Definitions 2 and 4) depend on the choice of function f .

So this definition states exactly the fact that the inconsistency value only takes into account the minimal inconsistent subsets of the base. In particular two different bases K and K' with exactly the same minimal inconsistent sets will have the same amount of inconsistency.

The simplest types of MIV one can define are the following ones:

Definition 2 MIV_D and $MIV_{\#}$ are defined as follows:

- $MIV_D(K, \alpha) = \begin{cases} 1 & \text{if } \exists M \in \text{MI}(K) \alpha \in M \\ 0 & \text{otherwise} \end{cases}$
- $MIV_{\#}(K, \alpha) = |\{M \in \text{MI}(K) \mid \alpha \in M\}|$

The first value is the drastic one, that takes value one if the formula belongs to a minimal inconsistent subset, and zero otherwise. The second one is a cardinality value, that counts the number of minimal inconsistent subsets the formula belongs to.

The first value is of little interest, since it allows us just to make a distinction between free formulas and the other ones. The second one is more useful, and allows us to find more interesting results. Let us check this on the following example.

Example 1 Let $K_1 = \{a, \neg a, \neg a \wedge c, a \vee d, \neg d, b \wedge \neg b, e\}$, so the minimal inconsistent subsets of K_1 are

$$\text{MI}(K_1) = \{\{b \wedge \neg b\}, \{a, \neg a\}, \{a, \neg a \wedge c\}, \{\neg a, a \vee d, \neg d\}\}$$

and the $MIV_{\#}$ value gives as a result:

$$\begin{aligned} MIV_{\#}(K_1, a) &= 2 & MIV_{\#}(K_1, \neg d) &= 1 \\ MIV_{\#}(K_1, \neg a) &= 2 & MIV_{\#}(K_1, b \wedge \neg b) &= 1 \\ MIV_{\#}(K_1, \neg a \wedge c) &= 1 & MIV_{\#}(K_1, e) &= 0 \\ MIV_{\#}(K_1, a \vee d) &= 1 & & \end{aligned}$$

It is easy to check that $MIV_{\#}$ is just a scoring function (Hunter 2004) applied uniquely on formulae:

Definition 3 Let $K \in \mathcal{K}_{\mathcal{L}}$. Let S be the **scoring function** for K defined as follows, where $S : \wp(K) \mapsto \mathbb{N}$ and $K' \in \wp(K)$

$$S(K') = |\text{MI}(K)| - |\text{MI}(K - K')|$$

For a belief base K , a scoring function S gives the number of minimal inconsistent subsets of K that would be eliminated if the subset K' was removed from K . See (Hunter 2004) for more details on the use of these scoring functions. So if $\alpha \in K$, it is straightforward to see that we have $MIV_{\#}(\alpha, K) = S(\{\alpha\})$.

The evaluation of the inconsistency value of each formula given by $MIV_{\#}$ is still very rough. In particular, it does not take into account the cardinalities of the minimal inconsistent subsets the formula belongs to. But, as explained in several works (see for example (Knight 2001)), the size of

the minimal inconsistent subset can have an impact on the evaluation of the inconsistency. The idea is that the smaller the value of the minimal inconsistent subset, the bigger is the inconsistency. To illustrate this we use the prototypical example of the lottery paradox given by Knight to motivate his approach.

Example 2 There are a number of lottery tickets with one of them being the winning ticket. Suppose w_i denotes ticket i will win, then we have the assumption $w_1 \vee \dots \vee w_n$. In addition, for each ticket i , we may pessimistically (or probabilistically if the number of tickets is important) assume that it will not win, and this is represented by the assumption $\neg w_i$. So the base K_L is:

$$K_L = \{\neg w_1, \dots, \neg w_n, w_1 \vee \dots \vee w_n\}$$

Clearly if there are three or two (or one!) tickets in the lottery, then this base is highly inconsistent. But if there are millions of tickets there is intuitively (nearly) no conflict in the base.

So it could prove better not to simply take the number of minimal inconsistent subsets a formula belongs to, but to take into account their cardinalities. This idea gives a third MIV value:

Definition 4 MIV_C is defined as follows:

$$MIV_C(K, \alpha) = \sum_{M \in \text{MI}(K) \text{ s.t. } \alpha \in M} \frac{1}{|M|}$$

This allows us to define a much more precise view of the inconsistency, as illustrated in the following example.

Example 3 Let $K_1 = \{a, \neg a, \neg a \wedge c, a \vee d, \neg d, b \wedge \neg b, e\}$ and the MIV_C value gives as a result:

$$\begin{aligned} MIV_C(K_1, a) &= 1 & MIV_C(K_1, \neg d) &= \frac{1}{3} \\ MIV_C(K_1, \neg a) &= \frac{5}{6} & MIV_C(K_1, b \wedge \neg b) &= 1 \\ MIV_C(K_1, \neg a \wedge c) &= \frac{1}{2} & MIV_C(K_1, e) &= 0 \\ MIV_C(K_1, a \vee d) &= \frac{1}{3} & & \end{aligned}$$

We can compare these obtained results with the ones of Example 1, where less distinction was possible, with only three different levels. We can notice that now we can make a distinction between $\neg a \wedge c$, and $a \vee d$, that both belong to only one minimal inconsistent subset, but the one of $a \vee d$ is bigger. We also note that an inconsistent formula has a high degree of inconsistency according to the MIV_C value. One can remark that with MIV_C the formula $b \wedge \neg b$ is evaluated as more conflicting than the formula $\neg a$ which belongs to two larger minimal inconsistent subsets, whereas the evaluation is the converse for $MIV_{\#}$.

We can see more clearly on a dedicated example why MIV_C gives a more precise view of the conflict brought by each formula than $MIV_{\#}$.

Example 4 Consider $K_2 = \{a \wedge \neg a, b \wedge \neg b\}$ and $K_3 = \{a \wedge \neg b, \neg a \wedge b\}$. Here we see that the $MIV_{\#}$ value assigns the same value to each formula, even though for instance $a \wedge \neg a$ is entirely responsible for an inconsistency, whereas $a \wedge \neg b$ is only partially responsible for an inconsistency.

$$\begin{aligned} MIV_{\#}(K_2, a \wedge \neg a) &= 1 & MIV_{\#}(K_3, a \wedge \neg b) &= 1 \\ MIV_{\#}(K_2, b \wedge \neg b) &= 1 & MIV_{\#}(K_3, \neg a \wedge b) &= 1 \end{aligned}$$

In contrast, the MIV_C value is more discriminating and so for instance $a \wedge \neg a$, which is entirely responsible for an inconsistency, has the maximum value of 1, whereas $a \wedge \neg b$, which is only half of the cause of an inconsistency, has a value of 1/2.

$$\begin{aligned} MIV_C(K_2, a \wedge \neg a) &= 1 & MIV_C(K_3, a \wedge \neg b) &= \frac{1}{2} \\ MIV_C(K_2, b \wedge \neg b) &= 1 & MIV_C(K_3, \neg a \wedge b) &= \frac{1}{2} \end{aligned}$$

More generally, we see that the MIV_C value is affected by the size of each minimal inconsistent subset: Returning to Example 2, we see that for K_L and for some $\alpha \in K_L$, the value of $MIV_C(K_L, \alpha)$ decreases as the cardinality of K_L increases.

We now give a few observations regarding the MIV_C definition. Other properties will be also derivable from later results of the paper.

Proposition 1

- If α is a free formula in K , then $MIV_C(K, \alpha) = 0$
- $MIV_C(K \cup K', \alpha) \geq MIV_C(K, \alpha)$
- If $\alpha \equiv \perp$, then $MIV_C(K, \alpha) = 1$
- If $\phi \vdash \psi$ and $\phi \not\vdash \perp$ then

$$MIV_C(K \cup \{\phi\}, \alpha) \geq MIV_C(K \cup \{\psi\}, \alpha)$$

So from the examples and observations in this section, it seems that MIV_C is an appealing and informative measure of inconsistency.

Shapley Inconsistency Values

Shapley Inconsistency Values were introduced in (Hunter & Konieczny 2006) in order to be able to define a measure of inconsistency for each formula, from a measure of inconsistency on belief bases.

The idea is to start from one basic measure of inconsistency from the literature, that allows us to evaluate the inconsistency of a belief base, to use this measure as the definition of a coalitional game, and to use a notion from cooperative game theory, the Shapley value (Shapley 1953), that allows us to define the merits of one individual in a given game. In our setting, with the scale reversal, this amounts to define the blame/responsability of one formula in a given base for the inconsistencies.

So the idea is similar to the one that drove the definition of MIV in the last section. Therefore it is natural to wonder if there are some links between the two approaches.

Let us first give the background on Shapley Inconsistency Values (SIV). We will just give here the definitions needed for this paper and for the proofs, for more details see (Hunter & Konieczny 2006).

First we recall the standard definitions of games in coalitional form and of the Shapley value (Aumann & Hart 2002).

Definition 5 Let $N = \{1, \dots, n\}$ be a set of n players. A game in coalitional form is given by a function $v : 2^N \rightarrow \mathbb{R}$, with $v(\emptyset) = 0$.

This framework defines games in a very abstract way, focusing on the possible coalition formations. A coalition is just a subset of N . This function gives what payoff can be achieved by each coalition in the game v when all its members act together as a unit.

A natural notion of solution for this kind of game is to try to define the payoff that can be expected by each player i for the game v . This is what is called a *value*.

Definition 6 A value is a function that assigns to each game v a vector of payoff $S(v) = (S_1, \dots, S_n)$ in $\mathbb{R}^n$, where S_i is the payoff for player i .

Despite these very abstract definitions of the game and of the value, it is possible to define a notion of solution, i.e. a value, that gives for any game the expected payoff (merits) of each player. The first (and main) such value has been defined by Shapley (Shapley 1953).

Basically the idea can be explained as follows: considering that the coalitions form according to some order (a first player enters the coalition, then another one, then a third one, etc), and that the payoff attached to a player is its marginal utility (i.e. the utility that it brings to the existing coalition), so if C is a coalition (subset of N) not containing i , player's i marginal utility is $v(C \cup \{i\}) - v(C)$. As one can not make any hypothesis on which order is the correct one, we may suppose that each order is equally probable. This leads to the following formula:

Let σ be a permutation on N , with σ_n denoting all the possible permutations on N . We need the following notation:

$$p_\sigma^i = \{j \in N \mid \sigma(j) < \sigma(i)\}$$

That means that p_σ^i represents all the players that precede player i for a given order σ .

Definition 7 Let $i \in N$ be a player, and n be the number of players. The Shapley value of a game v is defined as.

$$S_i(v) = \frac{1}{n!} \sum_{\sigma \in \sigma_n} v(p_\sigma^i \cup \{i\}) - v(p_\sigma^i)$$

The Shapley value can be directly computed from the possible coalitions (without looking at the permutations) using the following expression:

$$S_i(v) = \sum_{C \subseteq N} \frac{(c-1)!(n-c)!}{n!} (v(C) - v(C \setminus \{i\}))$$

where c is the cardinality of C .

Besides the fact that this definition gives very sensible results, its legitimacy is also given by a nice characterization result:

Proposition 2 (Shapley 1953) The Shapley value is the only value that satisfies all of Efficiency, Symmetry, Dummy and Additivity.

- $\sum_{i \in N} S_i(v) = v(N)$ (Efficiency)
- If i and j are such that for all C s.t. $i, j \notin C$, $v(C \cup \{i\}) = v(C \cup \{j\})$, then $S_i(v) = S_j(v)$ (Symmetry)

- If i is such that $\forall C \ v(C \cup \{i\}) = v(C)$, then $S_i(v) = 0$ **(Dummy)**
- $S_i(v + w) = S_i(v) + S_i(w)$ **(Additivity)**

This result supports several variations: there are other equivalent axiomatizations of the Shapley value, and there are some different values that can be defined by relaxing some of the above axioms. See (Aumann & Hart 2002).

So the idea is to consider an inconsistency measure (that allows us to evaluate the inconsistency of a belief base) as a game in coalitional form, and to compute the corresponding Shapley value, in order to be able to define the inconsistency value of each formula of the base.

We ask some properties for the underlying inconsistency measure:

Definition 8 An inconsistency measure I is called a **basic inconsistency measure** if it satisfies the following properties, $\forall K, K' \in \mathcal{K}_{\mathcal{L}}, \forall \alpha, \beta \in \mathcal{L}$:

- $I(K) = 0$ iff K is consistent **(Consistency)**
- $I(K \cup K') \geq I(K)$ **(Monotony)**
- If α is a free formula of $K \cup \{\alpha\}$, then $I(K \cup \{\alpha\}) = I(K)$ **(Free Formula Independence)**
- If $\alpha \vdash \beta$ and $\alpha \not\vdash \perp$, then $I(K \cup \{\alpha\}) \geq I(K \cup \{\beta\})$ **(Dominance)**

In (Hunter & Konieczny 2006), a Normalization property was also presented as an optional property. We do not consider that property in this paper.

Definition 9 Let I be a basic inconsistency measure. We define the corresponding **Shapley Inconsistency Value (SIV)**, noted S_I , as the Shapley value of the coalitional game defined by the function I , i.e. let $\alpha \in K$:

$$S_{\alpha}^I(K) = \sum_{C \subseteq K} \frac{(c-1)!(n-c)!}{n!} (I(C) - I(C \setminus \{\alpha\}))$$

where n is the cardinality of K and c is the cardinality of C .

Note that this SIV gives a value for each formula of the base K . This definition allows us to define to what extent a formula inside a belief base is concerned with the inconsistencies of the base. It allows us to draw a precise picture of the contradiction of the base.

So, from a SIV, one can define an inconsistency measure for the whole belief base:

Definition 10 Let K be a belief base,

$$\hat{S}^I(K) = \max_{\alpha \in K} S_{\alpha}^I(K)$$

There are alternatives to Definition 10, see the discussion in (Hunter & Konieczny 2006).

MI Shapley Inconsistency Value

Since minimal inconsistent subsets of a base can be considered as fundamental features in characterizing inconsistency, we use the notion here as the basic inconsistency measure.

Definition 11 The **MI inconsistency measure** is defined as the number of minimal inconsistent sets of K , i.e. :

$$I_{MI}(K) = |\text{MI}(K)|$$

Example 5 $K = \{a, \neg a, \neg a \wedge c, a \vee d, \neg d, b \wedge \neg b, e\}$

Hence, we get the following:

$$\begin{aligned} I_{MI}(K) &= 4 & I_{MI}(\{a, \neg a, \neg a \wedge c\}) &= 2 \\ I_{MI}(\{b \wedge \neg b, e\}) &= 1 & I_{MI}(\{\neg a, \neg a \wedge c\}) &= 0 \end{aligned}$$

Proposition 3 The MI inconsistency measure I_{MI} is a basic inconsistency measure, i.e. it satisfies the properties of Consistency, Monotonicity, Free Formula Independence, and Dominance.

And in fact the following result shows that this MI Shapley Inconsistency Value is exactly the MIV_C measure of Definition 4:

Proposition 4 $S_{\alpha}^{I_{MI}}(K) = MIV_C(K, \alpha)$

Proof: Let us first show the following lemma that will be useful in the proof.

Lemma 1 If a simple game in coalitional form on a set of players $N = \{1, \dots, n\}$ is defined by a single winning coalition $C' \subseteq N$, i.e.:

$$v(C) = \begin{cases} 1 & \text{if } C' \subseteq C \\ 0 & \text{otherwise} \end{cases}$$

Then the corresponding Shapley value is:

$$S_i(v) = \begin{cases} 0 & \text{if } i \notin C' \\ \frac{1}{|C'|} & \text{if } i \in C' \end{cases}$$

Proof of Lemma 1 : The proof is direct using the logical properties of the Shapley value given in Proposition 2. Since by (Dummy) we get that if $i \notin C'$, then $S_i(v) = 0$. By (Efficiency) we know that the outcome of the grand coalition N must be shared in the sum of the Shapley values of the players: $\sum_{i \in N} S_i(v) = 1$. Since for players $i \notin C'$ we know that $S_i(v) = 0$, it means that it has to be split between members of C' . So $\sum_{i \in C'} S_i(v) = 1$. Now by (Symmetry) we get that for all $i, j \in C'$, we have $S_i(v) = S_j(v)$. So this implies that if $i \in C'$, then $S_i(v) = \frac{1}{|C'|}$. $\square$

Let us now state the result. First suppose that α is a free formula of K , then we have immediately by (Minimality) that $S_{\alpha}^{I_{MI}}(K) = 0$. We also have immediately by definition that $MIV_C(K, \alpha) = 0$. So the equality is satisfied in this case.

Now suppose that α is not a free formula of K . First remark that I_{MI} can be decomposed in $I_{MI}(C) = \sum_{M \in \text{MI}(K)} \hat{M}(C)$, where $\hat{M}$ is the following characteristic function

$$\hat{M}(C) = \begin{cases} 1 & \text{if } M \subseteq C \\ 0 & \text{otherwise} \end{cases}$$

Let us denote by $\hat{M}(K)$ the game in coalitional form defined from K and the characteristic function $\hat{M}$.

So now let us start from the MI Shapley Inconsistency Value:

$$\begin{aligned}
& S_{\alpha}^{I_{MI}}(K) \\
&= \sum_{C \subseteq K} \frac{(c-1)!(n-c)!}{n!} (I_{MI}(C) - I_{MI}(C \setminus \{\alpha\})) \\
&= \sum_{C \subseteq K} \frac{(c-1)!(n-c)!}{n!} \left(\sum_{M \in \text{MI}(K)} \hat{M}(C) \right. \\
&\quad \left. - \sum_{M \in \text{MI}(K)} \hat{M}(C \setminus \{\alpha\}) \right) \\
&= \sum_{C \subseteq K} \frac{(c-1)!(n-c)!}{n!} \left(\sum_{M \in \text{MI}(K)} (\hat{M}(C) - \hat{M}(C \setminus \{\alpha\})) \right) \\
&= \sum_{C \subseteq K} \sum_{M \in \text{MI}(K)} \frac{(c-1)!(n-c)!}{n!} (\hat{M}(C) - \hat{M}(C \setminus \{\alpha\})) \\
&= \sum_{M \in \text{MI}(K)} \sum_{C \subseteq K} \frac{(c-1)!(n-c)!}{n!} (\hat{M}(C) - \hat{M}(C \setminus \{\alpha\})) \\
&= \sum_{M \in \text{MI}(K)} S_{\alpha}(\hat{M}(K))
\end{aligned}$$

Now note that by Lemma 1 we have $S_{\alpha}(\hat{M}(K)) = \frac{1}{|M|}$.

That gives $S_{\alpha}^{I_{MI}}(K) = \sum_{M \in \text{MI}(K)} \frac{1}{|M|} = MIV_C(K, \alpha)$.

□

This proposition is interesting for several reasons. First, it confirms the appeal of Shapley Inconsistency Values, since the very natural measure MIV_C is a special case of these measures. Second, it gives a simpler definition of $S_{\alpha}^{I_{MI}}$ than the one using the Shapley value. In general, obtaining a Shapley value is computationally demanding (Deng & Papadimitriou 1994). However, in the case of $S_{\alpha}^{I_{MI}}$, the above proposition hints at the possibility of computationally viable implementations for calculating these values. Finally, this equality is useful to state the logical properties of this value, as done in the next Section.

Logical Properties

It is quite difficult to state logical properties about inconsistency handling (and measure of inconsistency) in a purely classical framework, i.e. without adding too many hypotheses, by using an (arbitrary) paraconsistent logic to do so.

In (Hunter & Konieczny 2006) some logical properties for inconsistency measures are defined, as well as specific ones for Shapley Inconsistency Values. But there was no characterization theorem in that paper. We provide such a theorem below.

Let us first strengthen the condition on the basic inconsistency measure:

Definition 12 A *MinInc Separable basic inconsistency measure (MSBIM)* I is a basic inconsistency measure that satisfies this additional property:

• If $\text{MI}(K \cup K') = \text{MI}(K) \oplus \text{MI}(K')$, then $I(K \cup K') = I(K) + I(K')$ (**MinInc Separability**)

This property basically expresses the fact that the inconsistency measure depends on the minimal inconsistent sub-

sets, so that if we can partition the belief base in two subbases without “breaking” any minimal inconsistent subset, then the global inconsistency measure is the sum of the inconsistency measure of the two subbases. Clearly, the MI inconsistency measure satisfies this property.

Let us now enumerate the properties that we expect inconsistency values to satisfy:

- $\sum_{\alpha \in K} S_{\alpha}^I(K) = I(K)$ (**Distribution**)
- If $\exists \alpha, \beta \in K$ s.t. for all $K' \subseteq K$ s.t. $\alpha, \beta \notin K'$, $I(K' \cup \{\alpha\}) = I(K' \cup \{\beta\})$, then $S_{\alpha}^I(K) = S_{\beta}^I(K)$ (**Symmetry**)
- If α is a free formula of K , then $S_{\alpha}^I(K) = 0$ (**Minimality**)
- If $\alpha \vdash \beta$ and $\alpha \not\vdash \perp$, then $S_{\alpha}^I(K) \geq S_{\beta}^I(K)$ (**Dominance**)
- If $\text{MI}(K \cup K') = \text{MI}(K) \oplus \text{MI}(K')$, then $S_{\alpha}^I(K \cup K') = S_{\alpha}^I(K) + S_{\alpha}^I(K')$ (**Decomposability**)

The first four properties were already discussed in (Hunter & Konieczny 2006). The first three of these are closely related to original Shapley’s properties. The distribution property states that the inconsistency values of the formulae sum to the total amount of inconsistency in the base ($I(K)$). The symmetry property ensures that only the amount of inconsistency brought by a formula matters for computing the inconsistency value. As one could expect, a formula that is not embedded in any contradiction (i.e. does not belong to any minimal inconsistent subset) will not be blamed by the inconsistency value. This is what is expressed in the minimality property. The dominance property states that logically stronger formulae bring (potentially) more conflicts. It was shown in (Hunter & Konieczny 2006) that every Shapley Inconsistency Value satisfies these four properties.

The Decomposability property is related to Shapley’s Additivity property. We explain in (Hunter & Konieczny 2006) that a direct translation of this Additivity property makes little sense because it is not meaningful to add different (basic) inconsistency measures. But one can consider another translation of the additivity property, by looking to the “addition” of two different bases: the set union. So direct translation of that meaning leads to

$$S_{\alpha}^I(K \cup K') = S_{\alpha}^I(K) + S_{\alpha}^I(K')$$

This formulation is not satisfactory because it forgets the fact that new conflicts can appear when making the union of the two bases. So we want this property to hold only when joining two bases does not create any new inconsistencies. That is ensured by the condition of the Decomposability property. Note that this possibility of interaction between the two subgames that is not taken into account in the usual Additivity condition, is one of the criticisms about this condition. Let us quote for instance the following paragraph from (Luce & Raiffa 1957):

The last condition is not nearly so innocent as the other two. For although $v + w$ is a game composed from v and w , we cannot in general expect it to be played as if it were the two separate games. It will have its own structure which will determine a set of equilibrium outcomes which may be different from those for v and w .

Therefore, one might very well argue that its a priori value should not necessarily be the sum of the values of the two component games. This strikes us as a flaw in the concept value, but we have no alternative to suggest.

In our framework the interaction between the two bases is simply the new logical conflicts that appears when joining the bases, that allows us to say when this addition can hold, and when it is not sensible.

For setting the characterization result we have to ask one additional property that states that each minimal inconsistent subset brings the same amount of conflict:

- If $M \in \text{MI}(K)$, then $I(M) = 1$ **(MinInc)**

So now we reach the wanted characterization result

Proposition 5 *An inconsistency value satisfies Distribution, Symmetry, Minimality, Decomposability and MinInc if and only if it is the MI Shapley Inconsistency Value S_{α}^{IMI} .*

Proof: To prove that the MI Shapley Inconsistency Value satisfy the logical properties is easy. **(Distribution)**, **(Symmetry)**, **(Minimality)** are satisfied by all Shapley Inconsistency Values (Proposition 3 of (Hunter & Konieczny 2006)).

So it remains to show **(Decomposability)** and **(MinInc)**. **(MinInc)** is satisfied by definition since $I_{MI}(M) = |\text{MI}(M)| = 1$ for any $M \in \text{MI}(K)$.

For **(Decomposability)**, by definition $S_{\alpha}^I(K \cup K') = \sum_{C \subseteq K \cup K'} \frac{(c-1)!(n-c)!}{n!} (I(C) - I(C \setminus \{\alpha\}))$. Now split C on K and K' , i.e. define $H = C \cap K$ and $H' = C \cap K'$. It is easy to check that $C = H \cup H'$, and from the hypothesis that $\text{MI}(K \cup K') = \text{MI}(K) \oplus \text{MI}(K')$ we deduce that $\text{MI}(H \cup H') = \text{MI}(H) \oplus \text{MI}(H')$, so as $S_{\alpha}^{IMI} = \text{MIV}_C$ satisfies **(MinInc Separability)** we have that $I(C) = I(H) + I(H')$. So using this in the definition we have

$$\begin{aligned} S_{\alpha}^I(K \cup K') &= \sum_{C \subseteq K \cup K'} \frac{(c-1)!(n-c)!}{n!} (I(C) - I(C \setminus \{\alpha\})) \\ &= \sum_{C \subseteq K \cup K'} \frac{(c-1)!(n-c)!}{n!} (I(H) + I(H') - I(H \setminus \{\alpha\}) - I(H' \setminus \{\alpha\})) \\ &= \sum_{C \subseteq K \cup K'} \frac{(c-1)!(n-c)!}{n!} (I(H) - I(H \setminus \{\alpha\})) \\ &\quad + \sum_{C \subseteq K \cup K'} \frac{(c-1)!(n-c)!}{n!} (I(H') - I(H' \setminus \{\alpha\})) \\ &= \sum_{H \subseteq K} \frac{(c-1)!(n-c)!}{n!} (I(H) - I(H \setminus \{\alpha\})) \\ &\quad + \sum_{H' \subseteq K'} \frac{(c-1)!(n-c)!}{n!} (I(H') - I(H' \setminus \{\alpha\})) \\ &= S_{\alpha}^I(K) + S_{\alpha}^I(K') \end{aligned}$$

For the converse implication suppose that we have an inconsistency value that satisfies **(Distribution)**, **(Symmetry)**, **(Minimality)**, **(Decomposability)** and **(MinInc)**. We want to show that it is the MI Shapley Inconsistency Value.

First note that for any K such that $\text{MI}(K) = \{M_1, \dots, M_n\}$, if one chooses a sequence $M_1, \dots, M_n$, then for all i where $1 \leq i < n$, the following holds:

$$\text{MI}(M_1 \cup \dots \cup M_i \cup M_{i+1}) = \text{MI}(M_1 \cup \dots \cup M_i) \oplus \text{MI}(M_{i+1})$$

Hence, there is a sequence of the minimal inconsistent subsets of K , such that by use of **(Minimality)** and successive use of **(Decomposability)** we have that

$$S_{\alpha}^I(K) = \sum_{M \in \text{MI}(K)} S_{\alpha}^I(M)$$

Now for each M if $\alpha \notin M$ we have by **(Minimality)** that $S_{\alpha}^I(M) = 0$. And if $\alpha \in M$ then we have by **(Distribution)** $\sum_{\alpha \in M} S_{\alpha}^I(M) = I(M)$. And by **(Symmetry)** we have that $\forall \alpha, \beta \in M, S_{\alpha}^I(M) = S_{\beta}^I(M)$. So we obtain that

$$\forall \alpha \in M, S_{\alpha}^I(M) = \frac{I(M)}{|M|}$$

and therefore

$$S_{\alpha}^I(K) = \sum_{M \in \text{MI}(K) \text{ s.t. } \alpha \in M} \frac{I(M)}{|M|}$$

Now by **(MinInc)** we know that for all $M \in \text{MI}(K)$, $I(M) = 1$. That gives

$$S_{\alpha}^I(K) = \sum_{M \in \text{MI}(K) \text{ s.t. } \alpha \in M} \frac{1}{|M|}$$

That is the definition of MI Shapley Inconsistency Value. $\square$

This result means that the Shapley Inconsistency Value S_{α}^{IMI} is completely characterized by five simple and intuitive axioms.

Note that Dominance, although satisfied by SIV, is not required for stating this proposition.

Towards Implementation of Inconsistency Values

The development of SAT solvers has made impressive progress in recent years, allowing, despite the computational complexity of the problem, to practically solve a number of intractable problems (Kautz & Selman 2007).

Based on SAT solvers, some techniques have been aimed at the identification of minimal inconsistent subsets (called in these works Minimally Unsatisfiable Subformulas or MUS). Although the identification problem is computationally hard, since checking whether a set of clauses is a MUS or not is DP-complete, and checking whether a formula belongs to the set of MUSes of a base, is in Σ_2^P (Eiter & Gottlob 1992); it seems that finding each MUS can be practically feasible (Grégoire, Mazure, & Piette 2007; 2008).

Thanks to Proposition 4, we then can define an easy algorithm to compute the MI Shapley Inconsistency Value

of the formulae of a base:

Input: A belief base $K = \{\alpha_1, \dots, \alpha_n\}$
Output: A profile of values $(S_{\alpha_1}^I(K), \dots, S_{\alpha_n}^I(K))$
1- **For** i **from** 1 **to** n
 $S_i \leftarrow 0$
2- Compute $MI(K)$
3- **For each** $C \in MI(K)$
 For each $\alpha_i \in C$
 $S_i \leftarrow S_i + \frac{1}{|C|}$
4- **Return** $(S_1, \dots, S_n)$

The hard step in the above algorithm is step 2. But if it can be viably computed, then the rest of the algorithm is just polynomial in the size of the $MI(K)$ (but of course the size of $MI(K)$ can be exponential in the size of K).

In future work, we plan to implement such an algorithm using an existing algorithm to identify MUS (Grégoire, Mazure, & Piette 2007; 2008). This would give us a practical tool to measure inconsistency.

A possible approach to ameliorate the cost of entailment in finding minimal inconsistent subsets is to use approximate entailment: Proposed in (Levesque 1984), and developed in (Schaerf & Cadoli 1995), classical entailment is approximated by two sequences of entailment relations. The first is sound but not complete, and the second is complete but not sound. Both sequences converge to classical entailment. For a set of propositional formulae Δ , a formula α , and an approximate entailment relation $\models_i$, the decision of whether $\Delta \models_i \alpha$ holds or $\Delta \not\models_i \alpha$ holds can be computed in polynomial time. Approximate entailment has been developed for anytime coherence reasoning (Koriche 2001; 2002), and in future work, we will investigate its potential for an approximate version of the MI Shapley Inconsistency Value.

By focussing on subsystems of classical logic, such as description logics, there appears to be much potential in harnessing existing specialized reasoning systems for finding minimal inconsistent subsets of a belief base (for example by using the Pellet reasoning system for description logics (Kalyanpur *et al.* 2005; Parsia, Sirin, & Kalyanpur 2005)). Furthermore, measuring inconsistency in description logic ontologies offers a potentially interesting and worthwhile application problem (Qi & Hunter 2007).

Another application area for inconsistency measures is in supporting reasoning with inconsistent databases (for example when using maximally consistent subsets of the database (Bertossi & Bravo 2005)). Again, this is an application where language restrictions and specialized reasoning systems offer the potential for viable means for finding minimal inconsistent subsets of a belief base, and thereby finding the MI Shapley Inconsistency Value.

Conclusion

As discussed in the introduction, there are a number of proposals for measures of inconsistency. The main novel contributions provided by this paper are :

- A first discussion on the definition of inconsistency values based on minimal inconsistent subsets of belief bases, and this leads to the definition of the family of Minimal Inconsistent Values;
- An equivalence between a particular Shapley Inconsistency Value (the MI Shapley Inconsistency Value) and a simple Minimal Inconsistent Value which gives an additional argument to support the Shapley Inconsistency Value definition since it captures this very natural value as particular case;
- The first (as far as we know) axiomatization of an inconsistency value;
- And finally, the characterization of the MI Shapley Inconsistency Value in terms of the MIV_C measure opens the possibility for computationally viable calculation of inconsistency values.

In future work, we would like to analyse the computational complexity of using the MI Shapley Inconsistency Value, develop algorithms and implementations (possibly based on approximation techniques), and undertake case studies of applications of this value.

Acknowledgements

This work has been partly supported by a bilateral project CNRS - Royal Society.

References

- Aumann, R., and Hart, S., eds. 2002. *Handbook of Game Theory with Economic Applications*, volume 3. North-Holland.
- Barragáns-Martínez, B.; Pazos-Arias, J.; and Fernández-Vilas, A. 2004. On measuring levels of inconsistency in multi-perspective requirements specifications. In *Proceedings of the 1st Conference on the Principles of Software Engineering (PRISE'04)*, 21–30.
- Bertossi, L., and Bravo, L. 2005. Consistent query answers in virtual data integration systems. In *Inconsistency Tolerance*, volume 3300 of *LNCS*, 42–83. Springer.
- Bertossi, L.; Hunter, A.; and Schaub, T. 2004. Introduction to inconsistency tolerance. In *Inconsistency Tolerance*, volume 3300 of *LNCS*, 1–16. Springer.
- Chen, Q.; Zhang, C.; and Zhang, S. 2004. A verification model for electronic transaction protocols. In *Advanced Web Technologies and Applications*, volume 3007 of *LNCS*, 824–833. Springer.
- Deng, X., and Papadimitriou, C. H. 1994. On the complexity of cooperative solution concepts. *Mathematics of Operations Research* 19(2):257–255.
- Eiter, T., and Gottlob, G. 1992. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence* 57(2-3):227.
- Grant, J., and Hunter, A. 2006. Measuring inconsistency in knowledgebases. *Journal of Intelligent Information Systems* 27:159–184.

- Grant, J., and Hunter, A. 2008. Analysing inconsistent first-order knowledge bases. *Artificial Intelligence* 172:1064–1093.
- Grant, J. 1978. Classifications for inconsistent theories. *Notre Dame Journal of Formal Logic* 19:435–444.
- Grégoire, E.; Mazure, B.; and Piette, C. 2007. Boosting a complete technique to find mss and mus thanks to a local search oracle. In *International Joint Conference on Artificial Intelligence (IJCAI'07)*, 2300–2305.
- Grégoire, E.; Mazure, B.; and Piette, C. 2008. Using local search to find msses and muses. *European Journal of Operational Research (EJOR)*. to appear.
- Hunter, A., and Konieczny, S. 2004. Approaches to measuring inconsistent information. In *Inconsistency Tolerance*, volume 3300 of *LNCS*. Springer. 189–234.
- Hunter, A., and Konieczny, S. 2006. Shapley inconsistency values. In *Proceedings of the 10th International Conference on Knowledge Representation (KR'06)*, 249–259. AAAI Press.
- Hunter, A., and Nuseibeh, B. 1998. Managing inconsistent specifications: Reasoning, analysis and action. *ACM Transactions on Software Engineering and Methodology* 7:335–367.
- Hunter, A. 2002. Measuring inconsistency in knowledge via quasi-classical models. In *Proceedings of the National Conference on Artificial Intelligence (AAAI'02)*, 68–73. MIT Press.
- Hunter, A. 2003. Evaluating significance of inconsistencies. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI'03)*, 468–473.
- Hunter, A. 2004. Logical comparison of inconsistent perspectives using scoring functions. *Knowledge and Information Systems Journal* 6:528–543.
- Hunter, A. 2006. How to act on inconsistent news: Ignore, resolve, or reject. *Data and Knowledge Engineering* 57:221–239.
- Kalyanpur, A.; Parsia, B.; Sirin, E.; and Hendler, J. 2005. Debugging unsatisfiable classes in owl ontologies. *Journal of Web Semantics* 3(4):268–293.
- Kautz, H., and Selman, B. 2007. The state of SAT. *Discrete Applied Mathematics* 155(12):1514–1524.
- Knight, K. 2001. Measuring inconsistency. *Journal of Philosophical Logic* 31:77–98.
- Knight, K. 2003. Two information measures for inconsistent sets. *Journal of Logic, Language and Information* 12:227–248.
- Konieczny, S.; Lang, J.; and Marquis, P. 2003. Quantifying information and contradiction in propositional logic through epistemic actions. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI'03)*, 106–111.
- Koriche, F. 2001. On anytime coherence-based reasoning. In *Quantitative and Qualitative Approaches to Reasoning with Uncertainty*, volume 2143 of *LNCS*, 556–567. Springer.
- Koriche, F. 2002. Approximate coherence-based reasoning. *Journal of Applied Non-classical Logics* 12(2):239–258.
- Levesque, H. 1984. A logic of implicit and explicit belief. In *Proceedings of the National Conference on Artificial Intelligence (AAAI'84)*, 198–202.
- Lozinskii, E. 1994. Information and evidence in logic systems. *Journal of Experimental and Theoretical Artificial Intelligence* 6:163–193.
- Luce, R. D., and Raiffa, H. 1957. *Games and Decision*. Wiley.
- Ma, Y.; Qi, G.; Hitzler, P.; and Lin, Z. 2007. Measuring inconsistency for description logics based on paraconsistent semantics. In *Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU'07)*, volume 4724 of *LNCS*, 30–41. Springer.
- Martinez, M.; Pugliese, A.; Simari, G.; Subrahmanian, V.; and Prade, H. 2007. How dirty is your relational database? an axiomatic approach. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'07)*, volume 4724 of *Lecture Notes in Computer Science*, 103 – 114. Springer.
- Mu, K.; Jin, Z.; Lu, R.; and Liu, W. 2005. Measuring inconsistency in requirements engineering. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'05)*, volume 3571 of *LNAI*, 440–451. Springer.
- Parsia, B.; Sirin, E.; and Kalyanpur, A. 2005. Debugging owl ontologies. In *Proceedings of the 14th international conference on World Wide Web (WWW'05)*, 633–640. ACM.
- Qi, G., and Hunter, A. 2007. Measuring incoherence in description logic-based ontologies. In *The Semantic Web, 6th International Semantic Web Conference (ISWC'07)*, volume 4825 of *LNCS*, 381–394. Springer.
- Qi, G.; Liu, W.; and Bell, D. 2005. Measuring conflict and agreement between two prioritized belief bases. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI'05)*, 552–557.
- Reiter, R. 1987. A theory of diagnosis from first principles. *Artificial Intelligence* 32:57–95.
- Schaerf, M., and Cadoli, M. 1995. Tractable reasoning via approximation. *Artificial Intelligence* 74:249–310.
- Shapley, L. S. 1953. A value for n-person games. In Kuhn, H. W., and Tucker, A. W., eds., *Contributions to the Theory of Games II (Annals of Mathematics Studies 28)*, 307–317. Princeton University Press.
- Wong, P., and Besnard, P. 2001. Paraconsistent reasoning as an analytic tool. *Journal of the Interest Group in Propositional Logic* 9:233–246.



2008

This paper was published in the *Proceedings, Eleventh International Conference on Principles of Knowledge Representation and Reasoning*, edited by Gerhard Brewka and Jérôme Lang (Menlo Park, California: AAAI Press).

The 2008 KR conference was held in Sydney, Australia, 16–19 September, 2008.