

Measuring Inconsistency in Probabilistic Knowledge Bases

Glauber De Bona¹

¹Thesis approved at the Institute of Mathematics and Statistics – USP
Supervisor: Marcelo Finger, Institute of Mathematics and Statistics – USP

Abstract. *In AI, inconsistency measures have been proposed as a way to manage inconsistent knowledge bases. This work investigates inconsistency measuring in probabilistic logic. We show that previously existing rationality postulates for inconsistency measures in probabilistic knowledge bases are themselves incompatible and introduce a new way of localising inconsistency to reconcile these postulates. We then show the equivalence between distance-based inconsistency measures, from the AI community, and incoherence measures, from philosophy, that are based on the disadvantageous gambling behaviour entailed by incoherent probabilistic beliefs (via Dutch books). This provides a meaningful interpretation to the former and efficient methods to compute the latter.*

1. Motivation

Representing real-world knowledge and performing inference usually demand formalisms that can cope with uncertainty. Probabilistic logics combine the deductive power of logical systems with the well-founded Theory of Probability to attend to this need. Probabilities have been assigned to logical formulas for more than a century [Boole 1854, de Finetti 1930], and more recently probabilistic logics entered into the Artificial Intelligence (AI) community [Nilsson 1986, Hansen and Jaumard 2000, Halpern 2003].

Most approaches to probabilistic logic rely on the consistency of the set of premises in order to perform inference, which is a standard practice throughout logical reasoning. Nonetheless, many are the possible sources of inconsistency in a probabilistic knowledge base: it may contain statistical data from different samples, it could have been formed by the opinion of different experts, or even a single expert could lack the resources to check his own consistency while building the base, etc. To restore consistency in such cases, the inconsistency may be analysed, which calls for a way to measure it. This work mainly investigates measures of inconsistency for knowledge bases over probabilistic logic.

Example 1.1. Consider we are devising an expert system to assist medical diagnosis. Suppose a group of experts on a given disease D is required to quantify the relationship between D and its symptoms. Suppose three conditional probabilities (among others) are presented:

- the probability of a patient of disease D exhibiting both symptom S_1 and symptom S_2 is at least 60%;
- the probability of a patient of disease D exhibiting symptom S_1 but not symptom S_2 is at least 50%;
- the probability of a patient of disease D exhibiting symptom S_1 is at most 80%.

A knowledge engineer, while checking those facts, finds that they are inconsistent: according to the first two items, the probability of symptom S_1 , given disease D , should be at least 110%. He does not even know where each probability came from, but plans to adjust them, since consistency is a requirement. How should he proceed? Which probabilities is the degree of inconsistency most sensitive to? Once chosen which number to change, should it be raised or lowered in order to approximate consistency? These are the kind of questions an inconsistency measure can help to answer. \square

The problem of measuring inconsistency in probabilistic knowledge bases has recently been tackled via distance minimisation [Thimm 2013, Potyka 2014]. Thimm has adapted to probabilistic logic the rationality postulates for measuring inconsistency in classical propositional logic [Hunter and Konieczny 2008]. Informally, an inconsistency measure is a function taking knowledge bases (sets of formulas) to non-negative numbers. The first postulate is (*Consistency*), which claims that an inconsistency measurement is zero if, and only if, the corresponding base is consistent. Another desirable property suggested is (*Independence*), stating that the withdrawal of a free formula of the base — *i.e.*, a formula that does not belong to any minimal inconsistent set — should not change the inconsistency measurement. Thimm brought these postulates, among others, to the probabilistic context, adding (*Continuity*) to the list, which intuitively says that small changes in probabilities lead to small changes at the value of the inconsistency measure. We prove that (*Consistency*), (*Independence*) and (*Continuity*) cannot hold together, and some of these postulates must be abandoned or exchanged for jointly satisfiable ones. Since (*Independence*) is based on minimal inconsistent sets, a problem related to the postulates reconciliation is how to characterise “atomic” inconsistencies in probabilistic logic.

Example 1.2. Recall the situation in Example 1.1. Now, instead of adjusting these probabilities by himself, this knowledge engineer decides that the experts who elicit them should do the job. It happens that each of these probabilistic assessments has come from a different expert. The engineer then intends to schedule a meeting among the physicians responsible for the inconsistency in order for them to reassess their assignments in a consistent way. These experts are very busy, and their time, expensive, thus the knowledge engineer wants to invite for the meeting only the very physicians whose probabilistic assessments are collaborating, or causing, the inconsistency of the whole base. Who should the engineer call? Or, which pieces of information can be “blamed” for causing the inconsistency? It seems clear that the doctors responsible for the two first statements should be invited, but what about the third? These questions can be answered by *localising* the inconsistency in the base. \square

While computer scientists are investigating the problem of measuring inconsistency in probabilistic knowledge bases, statisticians and philosophers have been interested in evaluating the incoherence of formal agents that assign probability to events or propositions. In Bayesian epistemology, probabilities are usually construed as an agent’s degrees of belief, which can be operationally defined as the relative prices she would be willing to pay for gambles. Under these assumptions, the incoherence of an agent by can be quantified via the sure loss she would be exposed to via a disadvantageous set of bets — a *Dutch book* [Schervish et al. 2002]. To the best of our knowledge, these proposals for measuring incoherence of Bayesian agents have been ignored within the AI community, even though they correspond to measures on probabilistic knowledge bases.

2. Objectives

The first main objective of this work is to reconcile the rationality postulates for measuring the inconsistency of probabilistic knowledge bases. To achieve this, the characterisation of problematic sets of formulas in a probabilistic knowledge base — those causing the inconsistency — must be analysed, since these sets are the basis of some incompatible desirable properties. That is, before fixing the postulates, we have to understand how to localise inconsistency in probabilistic logic.

As a second major aim, this work intends to link different approaches from different communities to very similar problems: on the one hand, distance-based methods to measure the inconsistency of probabilistic knowledge bases in Artificial Intelligence; on the other hand, sure losses via Dutch books to quantify the incoherence of agents in Bayesian statistics and formal epistemology.

3. Main Contributions

Identifying and fixing the incompatibility of the rationality postulates for inconsistency measures in probabilistic logic is the first major contribution of this work. We prove that the desirable properties are not jointly satisfiable and provide suitable for one them. A derived achievement of this work is the proposal of alternative forms of localising the inconsistency in probabilistic bases, which are employed to reconcile the postulates.

In the second main contribution of the thesis here summarised, we prove a formal equivalence between inconsistency measures via distance minimisation and incoherence measure via Dutch books. This provides the former with a meaningful interpretation, through the betting behaviour induced by the probabilities, and latter with efficient computation methods. Furthermore, our results open the path for the reciprocal interchange of ideas and techniques between the corresponding communities.

4. List of Publications

Part of this thesis' results, as well as related research done by the author, has been appearing in major international journals and conferences. The list of publications shows that this research has had impact in its surrounding areas, providing both a general framework for localising inconsistency in classical logic and a foundation for probabilistic consolidation on the AGM framework of belief revision, for instance. Also, the techniques developed were applied in more distant fields, like formal epistemology in philosophy and natural language processing. Additionally, the most recent papers are evidence of international collaboration: with Anthony Hunter, from the University College London, where the author has done post-doctoral research, and with Julia Staffel, German philosopher based at the Washington University in St. Louis.

De Bona, G. and Hunter, A. (2017). Localising iceberg inconsistencies. *Artificial Intelligence*. In press. Qualis-Capes: A1.

De Bona, G. and Staffer, J. (2017). Graded incoherence for accuracy-firsters. *Philosophy of Science*. In press. InCites-JCR ranking: #7 out of 44 in History and Philosophy of Science.

Finger, M. and De Bona, G. (2017). Algorithms for deciding counting quantifiers over unary predicates. In *AAAI*. In press. Qualis-Capes: A1.

De Bona, G., Finger, M., Ribeiro, M., Santos, Y., and Wassermann, R. (2016). Consolidating probabilistic knowledge bases via belief contraction. In *Proceedings, Fifteenth International Conference on the Principles of Knowledge Representation and Reasoning (KR 2016)*, pages 125–134. Qualis-Capes: A2.

De Bona, G. and Finger, M. (2015). Measuring inconsistency in probabilistic logic: rationality postulates and Dutch book interpretation. *Artificial Intelligence*, 227:140–164. Qualis-Capes: A1.

de Morais, E., De Bona, G., and Finger, M. (2015). Non-markovian logic-probabilistic modeling and inference. In *IJCAI-2015 Workshop on Weighted Logics for Artificial Intelligence (WL4AI-2015)*, pages 48–55.

Finger, M. and De Bona, G. (2015). Probabilistic satisfiability: algorithms with the presence and absence of a phase transition. *Annals of Mathematics and Artificial Intelligence*, pages 1–39. Qualis-Capes: B1.

De Bona, G., Cozman, F., and Finger, M. (2015). Generalized probabilistic satisfiability through integer programming. *Journal of the Brazilian Computer Society*, 21(1):1–14. Qualis-Capes: B2.

De Bona, G., Cozman, F., and Finger, M. (2014). Towards classifying propositional probabilistic logics. *Journal of Applied Logic*, 12(3):349–368. Qualis-Capes: B1.

De Bona, G., Cozman, F., and Finger, M. (2013). Generalized probabilistic satisfiability. In *Intelligent Systems (BRACIS), 2013 Brazilian Conference on*, pages 182–188. IEEE.

Finger, M. and De Bona, G. (2011). Probabilistic satisfiability: Logic-based algorithms and phase transition. In *IJCAI*, pages 528–533. Qualis-Capes: A1.

5. Relevant Results

A knowledge base (KB) is a set of (un)conditional (possibly imprecise) probabilities assigned to propositional formulas. In Examples 1.1 and 1.2 the KB encoding the three statements is $\Delta = \{P(S_1 \wedge S_2) \geq 0.6, P(S_1 \wedge \neg S_2) \geq 0.5, P(S_1) \leq 0.8\}$. An inconsistency measure is a function $\mathcal{I} : \mathbb{K} \rightarrow [0, \infty)$ intended to capture the inconsistency degree of a KB, where \mathbb{K} is the set of all KBs.

A direct approach to measure inconsistency in the unconditional case is via the distance between the probabilities in an KB and the closest consistent ones. For instance, to measure the inconsistency of Δ using some distance $d(\cdot, \cdot)$, we minimise $d(\langle 0.6, 0.5, 0.8 \rangle, \langle x, y, z \rangle)$ subject to $\{P(S_1 \wedge S_2) \geq x, P(S_1 \wedge \neg S_2) \geq y, P(S_1) \leq z\}$ being consistent. If d is the Manhattan/Absolute distance, the corresponding inconsistency measure, denoted by \mathcal{I}_1 , is such that $\mathcal{I}_1(\Delta) = |0.6 - 0.4| + |0.5 - 0.4| + |0.8 - 0.9| = 0.3$, since $\{P(S_1 \wedge S_2) \geq 0.5, P(S_1 \wedge \neg S_2) \geq 0.4, P(S_1) \leq 0.9\}$ is a closest consistent KB. Similarly, using Chebyshev distance, we have the inconsistency measure \mathcal{I}_∞ , which is such that $\mathcal{I}_\infty(\Delta) = \max\{|0.6 - 0.4|, |0.5 - 0.4|, |0.8 - 0.9|\} = 0.1$. These measures have

been generalized to the conditional case in a way that keeps computational practicability [Potyka 2014], and by \mathcal{I}_1 and \mathcal{I}_∞ we denote such generalisation. Any p -norm distance yields an inconsistency measure \mathcal{I}_p , but \mathcal{I}_1 and \mathcal{I}_∞ are the most efficiently computable ones, employing linear programming. These measures have been criticised for lacking a meaningful interpretation, being only geometrical.

The formulation of inconsistency measures has been guided by rationality postulates requiring basic, intuitive properties. Our first relevant result proves the incompatibility of three of these postulates, informally presented in the introduction¹:

Theorem 5.1. *(Consistency), (Independence) and (Continuity) are incompatible.*

To fix the postulates, we investigate (Independence), which is the least appealing to the intuition. This postulate is based in the concept of minimal inconsistent set (MIS). A $\Gamma \in \mathbb{K}$ is a MIS if it is inconsistent and every $\Gamma' \subset \Gamma$ is consistent. For instance, $\Theta = \{P(S_1 \wedge S_2) \geq 0.6, P(S_1 \wedge \neg S_2) \geq 0.5\}$ is the only MIS contained in Δ .

Recall the situation in Example 1.2. Suppose that only the two experts responsible for the probabilities in Θ are called to reconcile their assignments, via relaxing the probability bounds, yielding the consistent KB $\Theta' = \{P(S_1 \wedge S_2) \geq 0.5, P(S_1 \wedge \neg S_2) \geq 0.5\}$. Nevertheless, adding the third assignment we have $\Delta' = \{P(S_1 \wedge S_2) \geq 0.5, P(S_1 \wedge \neg S_2) \geq 0.5, P(S_1) \leq 0.8\}$, which is still inconsistent – indeed, Δ' is a MIS. Note that $P(S_1 \wedge S_2) \geq 0.5$ and $P(S_1 \wedge \neg S_2) \geq 0.5$ entail $P(S_1) \geq 1$, contradicting $P(S_1) \leq 0.8$.

Hence, MISs do not capture all the inconsistency in a KB. To circumvent that, we introduce the concept of inescapable conflict. Informally, a KB is an inescapable conflict when its probability bounds can be relaxed to form a MIS. Thus, Δ is an inescapable conflict, for Δ' is a MIS. We propose then a new version of (Independence), called (*i*-Independence), demanding that removing a probabilistic assignment from a KB should not decrease its inconsistency measurement if such assignment does not take part in any inescapable conflict in the KB. Our second relevant result follows:

Theorem 5.2. *(Consistency), (*i*-Independence) and (Continuity) are compatible.*

The next main result deals with inconsistency measures based on bet transactions². Imagine a situation where Alice thinks the probability of a given coin landing heads up is 60% and the probability of such coin landing tails up (not heads up) is also 60%. The KB $\Delta = \{P(H) = 0.6, P(\neg H) = 0.6\}$ captures Alice's epistemic state, which is clearly inconsistent. Now suppose there is two bet tickets: one that pays \$100 (the bet prize) if the coin lands heads up and the other pays \$100 if the coin lands tails up. Based on her probabilistic beliefs, Alice think it is fair to buy these bet tickets for \$60 each. Thus, she is willing to spend \$120 on these gambles, but will receive only \$100 back, no matter the coin outcome, losing \$20 for sure. This set of bet transactions with guaranteed loss is called a Dutch book.

It is a well-known result that an agent is vulnerable to a Dutch book iff she is incoherent; *i.e.*, if the KB corresponding to her epistemic state is inconsistent. Hence, a natural idea proposed in philosophy is to measure the degree of an agent's incoherence through the size of the maximum sure loss she is exposed to via a Dutch book. The intuition is that the more incoherent an agent is, the greater the guaranteed loss she is exposed

¹In this section, the technical results are only sketched, and their precise mathematical formulation can be found in the thesis here summarised.

²To explain this concept and the related measures, we focus on the unconditional case, for simplicity.

to via a Dutch book. Nevertheless, if bet prizes are not limited, the sure loss in a Dutch book can be arbitrarily scaled up. Hence, the maximum guaranteed loss is normalised by either the greatest bet prize or the sum the prizes, yielding the incoherence measures \mathcal{I}_{SSK}^{max} and \mathcal{I}_{SSK}^{sum} [Schervish et al. 2002], respectively. For instance, if Alice’s loss is divided by the maximum prize, we have $\mathcal{I}_{SSK}^{max} = 20/100 = 0.2$, whereas normalising by the prize’s sum, we have $\mathcal{I}_{SSK}^{sum} = 20/200 = 0.1$.

Although apparently unrelated at first, the most efficient distance-based measures and the incoherence measures based on Dutch-books are equivalent, giving a meaningful interpretation to the former and efficient computation methods to the latter:

Theorem 5.3. *For any knowledge base $\Gamma \in \mathbb{K}$, $\mathcal{I}_1(\Gamma) = \mathcal{I}_{SSK}^{max}(\Gamma)$ and $\mathcal{I}_\infty(\Gamma) = \mathcal{I}_{SSK}^{sum}(\Gamma)$.*

6. Concluding Remarks

In the thesis here summarised we studied different ways of measuring inconsistency in probabilistic knowledge bases. The incompatibility among (Consistency), (Independence) and (Continuity) was proved and fixed, with a new form of characterising conflicts being develop along the way. Furthermore, we showed how distance-based inconsistency measures correspond to Dutch-book-based incoherence measures, fostering cross-pollination between computer science and philosophy. The ongoing impact of this research is attested by the list of publications, with the framework here developed being generalised and applied to different areas.

References

- Boole, G. (1854). *An Investigation of the Laws of Thought: on which are Founded the Mathematical Theories of Logic and Probabilities*. Walton and Maberly.
- de Finetti, B. (1930). Problemi determinati e indeterminati nel calcolo delle probabilità. *Rendiconti Reale Accademia dei Lincei*, 6:367–373.
- Halpern, J. (2003). Reasoning about uncertainty.
- Hansen, P. and Jaumard, B. (2000). Probabilistic satisfiability. *Handbook of Defeasible Reasoning and Uncertainty Management Systems: Algorithms for uncertainty and defeasible reasoning*, page 321.
- Hunter, A. and Konieczny, S. (2008). Measuring inconsistency through minimal inconsistent sets. In *11th International Conference on Principles of Knowledge Representation and Reasoning (KR’08)*, pages 358–366.
- Nilsson, N. (1986). Probabilistic logic. *Artificial Intelligence*, 28(1):71–87.
- Potyka, N. (2014). Linear programs for measuring inconsistency in probabilistic logics. In *Fourteenth International Conference on Principles of Knowledge Representation and Reasoning (KR-14)*. AAAI.
- Schervish, M., Seidenfeld, T., and Kadane, J. (2002). Measuring incoherence. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 561–587.
- Thimm, M. (2013). Inconsistency measures for probabilistic logics. *Artificial Intelligence*, 197:1–24.