

Reasoning with Inconsistent Knowledge using the Epistemic Approach to Probabilistic Argumentation

Anthony Hunter

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

Abstract

Structured argumentation involves drawing inferences from knowledge in order to construct arguments and counterarguments. Since knowledge can be uncertain, we can use a probabilistic approach to representing and reasoning with the knowledge. Individual arguments can be constructed from the knowledge, with the belief in each argument determined just from the belief in the formulae appearing in the argument. However, if the original knowledgebase is inconsistent, this does not take into account the counterarguments that can be constructed. We therefore need a wider perspective that revises the belief in individual arguments in order to take into account the counterarguments. To address this need, we present a framework for probabilistic argumentation that uses relaxation methods to give a coherent view on the knowledge, and thereby revises the belief in the arguments that are generated from the knowledge.

1 Introduction

Argumentation is an important human ability for handling incomplete and inconsistent information (Atkinson and *et al.* 2017). Computational models of argument aim to formalize aspects of this ability. In abstract argumentation, a set of arguments and counterarguments is represented by an argument graph. This is a graph where each node denotes an argument and each arc denotes an attack by one argument on another. Acceptable arguments can then be identified using a **dialectical approach** (such as grounded or preferred semantics) that identifies coalitions of arguments that can defend themselves from attack (Dung 1995).

An alternative to the dialectical approach is the **epistemic approach** which assumes a probability assignment to each argument. For a probability function P , and argument A , $P(A)$ is the degree of belief that A is acceptable. When $P(A) > 0.5$, then the argument is believed to be acceptable, whereas when $P(A) \leq 0.5$, then the argument is not believed to be acceptable. For the set of arguments $\text{Args}(G)$ appearing in argument graph G , and probability function P , an **epistemic extension** is the set $\{A \in \text{Args}(G) \mid P(A) > 0.5\}$. So the extension is determined from the probability function rather the structure of the graph.

Example 1. For the argument graph G in Figure 1, if $P(A_1) = 0.1$, $P(A_2) = 0.9$, $P(A_3) = 0.1$, $P(A_4) = 0.4$, and $P(A_5) = 0.8$, then the epistemic extension is $\{A_2, A_5\}$.

	A_1	A_2	A_3	A_4	A_5	Rational	Coherence
P_1	0.6	0.9	0.4	0.6	0.7	No	No
P_2	0.3	0.9	0.3	0.1	0.8	Yes	No
P_3	0.9	0.1	0.1	0.8	0.2	Yes	Yes

Table 1: Examples of belief in arguments in Figure 1.




Figure 1: Example of an argument graph.

The epistemic approach provides a finer grained assessment of an argument graph than given by Dung’s definition of extensions. By adopting constraints on the function, the epistemic approach subsumes Dung’s approach (Thimm 2012; Hunter and Thimm 2017). However, there is also a need for a non-standard view where we adopt alternative constraints on the function. For instance, we may wish to represent disbelief in arguments even when they are unattacked (Polberg and Hunter 2018). Nonetheless, for the non-standard view we may want the probabilities to respect the structure of the graph in some sense (Hunter 2013; Hunter and Thimm 2017; Polberg and Hunter 2018). For example, when argument A attacks argument B , the **rational constraint** ensures that if A is believed (i.e. $P(A) > 0.5$), then B is not believed (i.e. $P(B) \leq 0.5$), and the **coherence constraint** ensures that the belief in A and B sum to less than or equal to 1 (i.e. $P(A) + P(B) \leq 1$). So with constraints, we can manage how the structure of the graph is reflected in the probability function (and vice versa), and thereby combine the belief in the composition of an argument, and the belief in acceptability of its counterarguments.

The literature on the epistemic approach has assumed that the argument graph is given. In this paper, we investigate how we can construct argument graphs from uncertain knowledge. Essentially, the problems that we address in this paper are: (1) how to generate individual arguments and counterarguments from uncertain knowledge, and (2) how to revise each argument so that the belief in each argument takes into the account the belief in the counterarguments.

We will tackle these problems by assuming a language based on **belief statements**. Each belief statement expresses belief in a propositional formula, and each belief formula is a Boolean combination of belief statements. For example, for the proposition b , denoting *bird*, a couple of belief statements are $p(b) = 0.75$, which represents the belief in the proposition b is 0.75, and $p(b) \geq 0.5$, which represents the belief in the proposition b is greater than or equal to 0.5. A **probability distribution** is over the models of the propositional formulae, and we will use these probability distributions as models of the belief formulae. From belief formulae, we will construct **deductive arguments** where the premises are consistent and minimal for entailing the claim (where we use a probabilistic logic for entailment). For example, we will see that the following is an argument where the first item is a set of premises (composed from atoms b denoting *bird* and f denoting *flies*) and the second item is a claim that follows from the premises.

$$\{\{p(b) \geq 0.75, p(b \rightarrow f) \geq 0.75\}, p(f) \geq 0.5\}$$

We will use the belief in the claim as the **belief in the argument**, since this is a function of the belief in the premises and the inference process. So for the above, the belief in the argument is greater than or equal to 0.5.

The above gives belief in individual arguments. But we assume that our knowledgebase can be inconsistent, and if so, we will generate counterarguments. And if the knowledge is inconsistent, and the knowledge is in the form of belief statements, there is not a probability distribution that satisfies all the knowledge, and hence, we are not able to obtain a probability function for all the arguments. For example, we may have the knowledge $\{p(a) > 0.75, p(a) < 0.25\}$, and so arguments including $\{\{p(a) > 0.75\}, p(a) > 0.75\}$ and $\{\{p(a) < 0.25\}, p(a) < 0.25\}$, but there is no probability distribution that satisfies the knowledgebase (or the arguments from it taken together).

Our solution to this problem is to use a form of **belief relaxation**. This means that we change the belief statements until we get consistency. We introduce some methods for this that ensure that the process is measured (in the sense that we do not unnecessarily lose information) and fair to the sources (in the sense that the loss is spread fairly over the items of knowledge). The result of the relaxation process is that there is a probability distribution that satisfies all the revised knowledge, and thereby provides a consistent belief assignment to all the arguments. So we use this relaxed probability distribution to get the probability of acceptability for each argument, and hence the epistemic extension of the argument graph. Hence, in this paper, we regard the probability distribution as primary, and the knowledge, and arguments constructed from the knowledge, as secondary.

The language of belief statements being based on propositional logic with probabilistic qualification of uncertainty offers an expressive but practical representation for diverse applications such as intelligence analysis, medicine, and science. For instance, belief statements can be used to represent uncertainty in propositional formulae that capture logical relationships between symptoms, diseases, and treatments. Since this scientific or clinical knowledge can come

from empirical studies, the belief can be quantified, but since the knowledge may come from different studies and sources, there may be inconsistencies. So the ability to construct arguments, and then relax the beliefs until there is a consistent probability distribution, provides a potentially valuable tool for analyzing the inconsistent knowledge.

We proceed as follows: (1) We review a language for belief formulae; (2) We use the principle of maximum entropy for reasoning with belief formulae; (3) We propose an approach to generating argument graphs instantiated with deductive arguments obtained using maximum entropy reasoning with belief formulae; (4) We propose relaxation of belief formulae in order to find a coherent set of belief statements; (5) We propose the use of confidence in belief statements to provide a context-dependent relaxation of belief statements; (6) We discuss the related literature; and (7) We summarize and discuss our proposal.

2 Language

We assume classical propositional logic for describing aspects of the world. We then assume that we model uncertainty about the formulae using a probability distribution over models of the propositional logic. For this, we will draw on an established proposal for capturing probabilistic belief in propositional formulae (Paris 1994).

2.1 Probabilistic Belief in Propositional Formulae

We assume that the propositional language \mathcal{L} is finite. Given a language \mathcal{L} , the set of models (i.e. interpretations) of the language is denoted \mathcal{M} . Each **model** in \mathcal{L} is an assignment of true or false to the formulae of the language defined in the usual way for classical logic. We assume the usual machinery of classical propositional logic including the consequence relation \vdash and the satisfaction/entailment relation \models . For $\phi \in \mathcal{L}$, $\text{Models}(\phi)$ denotes the set of models of ϕ (i.e. $\text{Models}(\phi) = \{m \in \mathcal{M} \mid m \models \phi\}$). For $\alpha, \beta \in \mathcal{L}$, $\alpha \equiv \beta$ denotes $\text{Models}(\alpha) = \text{Models}(\beta)$.

Definition 1. Let \mathcal{M} be the models of the language \mathcal{L} . A **probability distribution** P on \mathcal{M} is a function $P : \mathcal{M} \rightarrow [0, 1]$ such that $\sum_{m \in \mathcal{M}} P(m) = 1$. Let \mathcal{P} denote the set of all probability distributions on \mathcal{M} .

We use a probability distribution to quantify uncertainty in a propositional formula as follows.

Definition 2. Let \mathcal{M} be the models of the language \mathcal{L} , and let P be a probability distribution on \mathcal{M} . The **belief in a formula** $\phi \in \mathcal{L}$ w.r.t. P is $P(\phi) = \sum_{m \in \text{Models}(\phi)} P(m)$.

Example 2. Let \mathcal{L} be the usual propositional formulae that can be formed from $\{a, b\}$. Now suppose $P(\{a, b\}) = 0.8$ and $P(\{a\}) = 0.2$. Hence, $P(a) = 1$, $P(a \wedge b) = 0.8$, $P(b \vee \neg b) = 1$, $P(a \wedge \neg b) = 0.2$, etc.

Example 3. Let \mathcal{L} be the usual propositional formulae that can be formed from $\{a, b\}$. For the probability assignments $P(a) = 0.8$, $P(\neg a \vee \neg b) = 0.8$, and $P(b) = 0.8$, there is no probability distribution P such that these hold together.

As shown in (Paris 1994), P is a probability distribution on \mathcal{M} iff (1) if $\models \alpha$, then $P(\alpha) = 1$; and (2) if $\models \neg(\alpha \wedge \beta)$, then $P(\alpha \vee \beta) = P(\alpha) + P(\beta)$.

2.2 Belief Formulae

We bring the use of probability into the language so that we can have formulae expressing uncertainty.

Definition 3. A **belief statement** is of the form $p(\alpha)\#x$ where $\alpha \in \mathcal{L}$ is a propositional formula, $\# \in \{=, \geq, \leq, >, <\}$, and $x \in [0, 1]$. A **belief formula** is a Boolean combination of belief statements (i.e. if ϕ is a belief statement, then it is a belief formula, and if ϕ and ψ are belief formulae, then each of $\phi \wedge \psi$, $\phi \vee \psi$ and $\neg\phi$ is a belief formula). Let \mathcal{B} be the set of belief formulae.

Example 4. For $a, b \in \mathcal{L}$, $(p(a \wedge b) > 0.9) \vee (p(\neg a \wedge \neg b) < 0.5)$ is an example of a belief formula.

We assume equivalences, denoted \equiv , between belief formulae: (1) $p(\alpha) \geq x \equiv (p(\alpha) = x) \vee (p(\alpha) > x)$, (2) $p(\alpha) \leq x \equiv (p(\alpha) = x) \vee (p(\alpha) < x)$, (3) $p(\alpha) \neq x \equiv \neg(p(\alpha) = x)$, (4) $p(\alpha) \not> x \equiv \neg(p(\alpha) > x)$, and (5) $p(\alpha) \not< x \equiv \neg(p(\alpha) < x)$.

Definition 4. The **satisfying distributions** for a belief statement $p(\alpha)\#x$ is $\text{Sat}(p(\alpha)\#x) = \{P' \in \mathcal{P} \mid P'(\alpha)\#x\}$, where $\# \in \{=, \geq, \leq, >, <\}$. For belief formulae ϕ and ψ , (1) $\text{Sat}(\phi \wedge \psi) = \text{Sat}(\phi) \cap \text{Sat}(\psi)$; (2) $\text{Sat}(\phi \vee \psi) = \text{Sat}(\phi) \cup \text{Sat}(\psi)$; and (3) $\text{Sat}(\neg\phi) = \text{Sat}(\top) \setminus \text{Sat}(\phi)$. Also, for a set of belief formulae Δ , let $\text{Sat}(\Delta) = \bigcap_{\phi \in \Delta} \text{Sat}(\phi)$.

Example 5. For \mathcal{L} composed from $\{a, b\}$, if $P_1(\{a, b\}) = 1$ and $P_2(\{c\}) = 1$, then $P_1, P_2 \in \text{Sat}(p(a \wedge b) = 1) \vee p(\neg a \wedge \neg b) = 1$. For \mathcal{L} composed from $\{c\}$, if $P_3(\{c\}) = 0.5$ and $P_4(\{c\}) = 0.6$, then $P_3 \notin \text{Sat}(p(c) > 0.5)$ and $P_4 \in \text{Sat}(p(c) > 0.5)$.

Proposition 1. (1) For $x \in (0, 1]$, $\text{Sat}(p(\perp) = x) = \emptyset$. (2) $\text{Sat}(p(\top) = 1) = \mathcal{P}$. (3) For any proposition α , $\text{Sat}(p(\alpha) \leq 1) = \mathcal{P}$ and $\text{Sat}(p(\alpha) \geq 0) = \mathcal{P}$. (4) When $x \neq y$, $\text{Sat}(p(\alpha) = x \wedge p(\alpha) = y) = \emptyset$. (5) When $\alpha \equiv \beta$, $\text{Sat}(p(\alpha) = x) = \text{Sat}(p(\beta) = x)$.

So belief formulae are a simple way of talking about belief in propositional formulae in the object language.

2.3 Restricted Value Sets

For certain applications a restricted set of probability distributions can be used where the probability values come from a finite set of values (Hunter, Polberg, and Thimm 2020). This may be appropriate if we want to represent probability values as in a Likert scale (Likert 1931). It also has the benefit of always producing a finite set of distributions. However, for the approach to be coherent, this set should be closed under addition and subtraction (assuming the resulting value is in the $[0, 1]$ interval) and it should contain 1.

Definition 5. A **finite set of rational numbers from the unit interval** Π is a **reasonable restricted value set** iff $1 \in \Pi$ and for any $x, y \in \Pi$ it holds that if $x + y \leq 1$, then $x + y \in \Pi$, and if $x - y \geq 0$, then $x - y \in \Pi$.

Since, we will only consider reasonable restricted value sets, we will refer to them as value sets. Examples include $\{0, 1\}$, $\{0, 0.5, 1\}$, and $\{0, 0.25, 0.5, 0.75, 1\}$.

A probability distribution P for a value set Π is a probability distribution such that for each model $m \in \mathcal{M}$,

$P(m) \in \Pi$. We will assume that all our probability distributions are with respect to a given value set. Also let $\text{Sat}(\phi, \Pi)$ (respectively $\text{Sat}(\Delta, \Pi)$) be the set of probability distributions w.r.t. Π that satisfy the belief formula ϕ (respectively the set of belief formulae Δ).

3 Maximum Entropy Reasoning

For a set of belief formulae, there is potentially a wide range of probability distributions that satisfy the set. This creates problems if we want to draw commonsense inferences from the set. One established approach is to use maximum entropy reasoning which involves selecting a subset of the distributions that have maximum entropy. A distribution that satisfies a set of constraints with maximum entropy is a distribution that makes the least extra commitment beyond satisfying the constraints (Jaynes 1982).

Definition 6. For a set of models \mathcal{M} , the **entropy** of probability distribution $P \in \mathcal{P}$ is $\text{Entropy}(P) = -\sum_{m \in \mathcal{M}} P(m) \times \log_2(P(m))$.

In the following, we focus on probability distributions for a value set, and out of those that satisfy a set of belief formulae, we pick those that are maximal for entropy.

Definition 7. For a value set Π , the set of **maximum entropy distributions** for belief formulae $\Delta \subseteq \mathcal{B}$ is $\text{Me}(\Delta, \Pi) = \{P \in \text{Sat}(\Delta, \Pi) \mid \text{for all } P' \in \text{Sat}(\Delta, \Pi), \text{Entropy}(P) \leq \text{Entropy}(P')\}$.

The maximum entropy distribution have some valuable commonsense properties including indifference to the introduction of irrelevant information, indifference to renaming, and independence in the absence of explicit information to the contrary (Paris 1994).

There are various ways that maximum entropy distributions can be used to create an inference relation depending the logic language and the assumptions about the uncertainty, see for example (Cheeseman 1983; Nilsson 1995; Goldszmidt, Morris, and Pearl 1993; Paskin 2002). We use the following notion of maximum entropy entailment (which is a form of defeasible reasoning).

Definition 8. For belief formulae $\Delta \subseteq \mathcal{B}$, and a value set Π , Δ **entails by maximum entropy** the belief formula $\psi \in \mathcal{B}$, denoted $\Delta \models_{\text{me}}^{\Pi} \psi$, iff $\text{Me}(\Delta, \Pi) \subseteq \text{Sat}(\psi, \Pi)$.

Example 6. Let $\Pi = \{0, 0.25, 0.5, 0.75, 1\}$, and let b denote bird and f denote fly. The inference $\{p(b) = 1, p(b \rightarrow f) \geq 0.75\} \models_{\text{me}}^{\Pi} p(f) > 0.5$ holds because $\text{Me}(\{p(b) = 1, p(b \rightarrow f) \geq 0.75\}) = \{P_1\}$, where $P_1(\{b\}) = 0.25$ and $P_1(\{b, f\}) = 0.75$.

The consequence relation is not monotonic as shown by the following example.

Example 7. For proposition $\alpha \in \mathcal{L}$, and value set Π , we obtain $\emptyset \models_{\text{me}}^{\Pi} p(\alpha) = 0.5$. Now assume the belief statement $p(\alpha) = 1$, then $\{p(\alpha) = 1\} \not\models_{\text{me}}^{\Pi} p(\alpha) = 0.5$.

Next, we give some properties of the entailment relation.

Proposition 2. For belief formulae $\Delta \subseteq \mathcal{B}$, $\phi, \psi \in \mathcal{B}$, and value set Π , the $\models_{\text{me}}^{\Pi}$ entailment relation satisfies reflexivity (REF), left logical equivalence (LEQ), right weakening

(RW), and (AND), and cautious monotonicity (CM).

$$\begin{aligned}
\Delta \cup \{\phi\} &\models_{\text{me}}^{\Pi} \phi && \text{(REF)} \\
\Delta \cup \{\phi\} &\models_{\text{me}}^{\Pi} \psi \text{ if } \Delta \cup \{\phi'\} \models_{\text{me}}^{\Pi} \psi \text{ and } \phi \equiv \phi' && \text{(LEQ)} \\
\Delta &\models_{\text{me}}^{\Pi} \phi \text{ if } \Delta \models_{\text{me}}^{\Pi} \psi \text{ and } \text{Sat}(\psi, \Pi) \subseteq \text{Sat}(\phi, \Pi) && \text{(RW)} \\
\Delta &\models_{\text{me}}^{\Pi} \phi \wedge \psi \text{ if } \Delta \models_{\text{me}}^{\Pi} \phi \text{ and } \Delta \models_{\text{me}}^{\Pi} \psi && \text{(AND)} \\
\Delta \cup \{\phi\} &\models_{\text{me}}^{\Pi} \psi \text{ if } \Delta \models_{\text{me}}^{\Pi} \phi \text{ and } \Delta \models_{\text{me}}^{\Pi} \psi && \text{(CM)}
\end{aligned}$$

Proof. We show each property as follows. (REF) $\text{Me}(\Delta \cup \{\phi\}, \Pi) \subseteq \text{Sat}(\Delta \cup \{\phi\}, \Pi)$ and $\text{Sat}(\Delta \cup \{\phi\}, \Pi) \subseteq \text{Sat}(\{\phi\}, \Pi)$, imply $\text{Me}(\Delta \cup \{\phi\}, \Pi) \subseteq \text{Sat}(\{\phi\}, \Pi)$. (LEQ) $\phi \equiv \phi'$ implies $\text{Me}(\Delta \cup \{\phi\}, \Pi) = \text{Me}(\Delta \cup \{\phi'\}, \Pi)$. So $\text{Me}(\Delta \cup \{\phi'\}, \Pi) \subseteq \text{Sat}(\phi', \Pi)$ implies $\text{Me}(\Delta \cup \{\phi\}, \Pi) \subseteq \text{Sat}(\phi, \Pi)$. (RW) Follows from the fact that when $\text{Sat}(\psi, \Pi) \subseteq \text{Sat}(\phi, \Pi)$ holds, then $\text{Me}(\Delta, \Pi) \subseteq \text{Sat}(\psi, \Pi)$ implies $\text{Me}(\Delta, \Pi) \subseteq \text{Sat}(\phi, \Pi)$. (AND) Assume $\Delta \models_{\text{me}}^{\Pi} \phi$ and $\Delta \models_{\text{me}}^{\Pi} \psi$. So $\text{Me}(\Delta, \Pi) \subseteq \text{Sat}(\phi, \Pi)$ and $\text{Me}(\Delta, \Pi) \subseteq \text{Sat}(\psi, \Pi)$. So $\text{Me}(\Delta, \Pi) \subseteq \text{Sat}(\phi, \Pi) \cap \text{Sat}(\psi, \Pi)$. So $\text{Me}(\Delta, \Pi) \subseteq \text{Sat}(\phi \wedge \psi, \Pi)$. So $\Delta \models_{\text{me}}^{\Pi} \phi \wedge \psi$. (CM) Since $\text{Me}(\Delta, \Pi) \subseteq \text{Sat}(\Delta, \Pi)$ and $\text{Me}(\Delta, \Pi) \subseteq \text{Sat}(\phi, \Pi)$, then $\text{Me}(\Delta, \Pi) \subseteq \text{Sat}(\Delta) \cap \text{Sat}(\{\phi\}, \Pi)$, and so $\text{Me}(\Delta, \Pi) \subseteq \text{Sat}(\Delta \cup \{\phi\}, \Pi)$. Therefore, $\text{Me}(\Delta, \Pi) = \{P \in \text{Sat}(\Delta \cup \{\phi\}, \Pi) \mid \text{for all } P' \in \text{Sat}(\Delta \cup \{\phi\}, \Pi), \text{Entropy}(P) \leq \text{Entropy}(P')\}$. Also, $\text{Me}(\Delta \cup \{\phi\}, \Pi) = \{P \in \text{Sat}(\Delta \cup \{\phi\}, \Pi) \mid \text{for all } P' \in \text{Sat}(\Delta \cup \{\phi\}, \Pi), \text{Entropy}(P) \leq \text{Entropy}(P')\}$. So $\text{Me}(\Delta, \Pi) = \text{Me}(\Delta \cup \{\phi\}, \Pi)$. So, $\text{Me}(\Delta, \Pi) \subseteq \text{Sat}(\psi, \Pi)$, implies $\text{Me}(\Delta \cup \{\phi\}, \Pi) \subseteq \text{Sat}(\psi, \Pi)$. \square

The entailment relation is trivializable (i.e. if the premises are inconsistent, then any formula follows). This is a property in common with classical propositional logic. As we will use the entailment relation to construct arguments with consistent premises, this is not a problem for our purposes.

Proposition 3. For belief formulae $\Delta \subseteq \mathcal{B}$, and value set Π , if $\text{Sat}(\Delta, \Pi) = \emptyset$, then for all $\phi \in \mathcal{B}$, $\Delta \models_{\text{me}}^{\Pi} \phi$.

Proof. Assume $\text{Sat}(\Delta, \Pi) = \emptyset$. So $\text{Me}(\Delta, \Pi) = \emptyset$. So for all ϕ , $\text{Me}(\Delta, \Pi) \subseteq \text{Sat}(\{\phi\}, \Pi)$. So $\Delta \models_{\text{me}}^{\Pi} \phi$. \square

There are algorithmic approaches to finding maximum entropy distributions that can be used for logical reasoning (e.g. (Nilsson 1995; Paskin 2002; Kern-Isberner and Lukasiewicz 2004)). Using restricted values has the advantage that all probability distributions can be enumerated in order to identify those with maximum entropy (though uniqueness is not guaranteed), and thereby is a simple route to implementation for smaller examples, and potentially could be scaled by harnessing SAT technology.

4 Generating Argument Graphs

In this section, we adapt definitions from deductive argumentation for constructing arguments and counterarguments based on belief formulae and maximum entropy reasoning. We assume our knowledge can be represented by a tuple. So each item comes from some source. The ordering in this sequence has no significance.

Definition 9. A **knowledge tuple** is a tuple $\Phi = [\phi_1, \dots, \phi_n]$ where each $\phi_i \in \mathcal{B}$ is a belief formula.

In a knowledge tuple, we can have multiple occurrences of the same formula (i.e. it is possible for $i, j \in \{1, \dots, n\}$ such that $i \neq j$ and $\phi_i = \phi_j$). When we want to consider the formulae in the knowledge tuple as a set, we use $\text{Set}(\Phi)$, which ensures that there are no duplicate occurrences of a formula. A knowledge tuple Φ is **consistent** iff $\text{Sat}(\text{Set}(\Phi), \Pi) \neq \emptyset$. For a knowledge tuple $\Phi = [\phi_1, \dots, \phi_n]$, the **length** of Φ is $\text{Len}(\Phi) = n$.

Definition 10. For a set of belief formulae $\Gamma \subseteq \mathcal{B}$, and a belief statement $\phi \in \mathcal{B}$, $\langle \Gamma, \phi \rangle$ is a **argument** iff $\Gamma \models_{\text{me}}^{\Pi} \phi$ and Γ is consistent (i.e. $\text{Sat}(\Gamma, \Pi) \neq \emptyset$) and there is no proper subset Γ' of Γ such that $\Gamma' \models_{\text{me}}^{\Pi} \phi$. The set of arguments from a knowledge tuple Φ is $\text{Args}(\Phi, \Pi) = \{\langle \Gamma, \phi \rangle \mid \langle \Gamma, \phi \rangle \text{ is an argument and } \Gamma \subseteq \text{Set}(\Phi)\}$.

So an argument satisfies both minimality (in order to avoid irrelevant premises) and consistency (in order to avoid the useless inferences that come with trivialization).

Example 8. Let $\Pi = \{0, 0.25, 0.5, 0.75, 1\}$. For the set of belief formulae $\{p(b) = 0.75, p(b \rightarrow f) = 0.75, p(p) = 0.75, p(p \rightarrow \neg f) = 1\}$, where b denotes bird, f denotes fly and p denotes penguin, the following are examples of arguments from these belief statements.

- $\langle \{p(b) = 1, p(b \rightarrow f) = 0.75\}, p(f) \geq 0.75 \rangle$
- $\langle \{p(p) = 0.75, p(p \rightarrow \neg f) = 1\}, p(\neg f) \geq 0.75 \rangle$

Given the richness of this logic (in terms of language and inferences), there is a number of natural ways that we can define counterarguments. We give some options in Definition 11 after some subsidiary definitions: For a belief statement $p(\alpha)\#x$, let $\text{Prop}(p(\alpha)\#x) = \alpha$; For a set of belief statements $\Gamma = \{\phi_1, \dots, \phi_n\}$, where we assume some arbitrary ordering over the formulae, let $\text{Prop}(\Gamma) = \text{Prop}(\phi_1) \wedge \dots \wedge \text{Prop}(\phi_n)$; And for an argument $A = \langle \Gamma, \phi \rangle$, let $\text{Support}(A) = \Gamma$, and $\text{Claim}(A) = \phi$.

Definition 11. Let A and B be two arguments. We define the following types of **attack**. A is an **undercut** of B if $\exists \Psi \subseteq \text{Support}(B)$ s.t. $\text{Prop}(\text{Claim}(A)) \equiv \neg \text{Prop}(\Psi)$. A is a **direct undercut** of B if $\exists \phi \in \text{Support}(B)$ s.t. $\text{Prop}(\text{Claim}(A)) \equiv \neg \text{Prop}(\phi)$. A is a **rebuttal** of B if $\text{Prop}(\text{Claim}(A)) \equiv \neg \text{Prop}(\text{Claim}(B))$.

Example 9. In Example 8, the arguments rebut each other, since $f \equiv \neg(\neg f)$ holds. In Figure 2, the root argument is undercut by the other argument, since $\neg n \equiv \neg n$ holds.

So a rebuttal occurs when the propositions in the claims of the two arguments cannot be satisfied together, and an undercut occurs when the proposition in the claim of the attacker cannot be satisfied together with some propositions in the premises of an attackee.

From the belief in the formulae in the premise, the entailment relation gives the belief in the claim. So the belief in the claim reflects the belief in the premises. We therefore use the belief in the claim as the belief in the argument. So we extend the use of probability distributions over \mathcal{M} as probability functions over arguments.

Definition 12. For an argument $A = \langle \Gamma, p(\alpha)\#v \rangle$, where $\# \in \{<, \leq, =, \geq, >\}$ and $v \in [0, 1]$, the **belief in argument** A is denoted $P(A)\#v$, where P is a probability distribution such that $P(\alpha)\#v$ holds.




Figure 2: An argument graph constructed from the knowledge tuple in Example 11. The atoms are $s = \text{symptoms of cold}$, $c = \text{has cold}$, $n = \text{no evidence of a serious virus}$, and $e = \text{there is a flu epidemic}$, and $\Pi = \{0, 0.01, 0.02, 0.03, \dots, 0.98, 0.99, 1\}$. Note, the root argument A_1 is a strong argument for the patient having a cold ($P(A_1) \geq 0.8$), and this is undercut by a weaker argument A_2 against there being no evidence of serious virus ($P(A_2) \leq 0.5$). In other words, the second argument raises a doubt about there being no evidence of a serious virus, and it may be useful to be aware of this doubt, even though only A_1 is in the epistemic extension.

Example 10. Continuing Example 8 let $A = \langle \{p(b) = 1, p(b \rightarrow f) = 0.75\}, p(f) \geq 0.75 \rangle$ and $B = \langle \{p(p) = 0.75, p(p \rightarrow \neg f) = 1\}, p(\neg f) \geq 0.75 \rangle$. So $P(A) \geq 0.75$ and $P(B) \geq 0.75$.

Using Definitions 10 and 11, we can instantiate an argument graph. Since we are not using the argument graph to determine which arguments are acceptable, (because we use a probability distribution to determine the epistemic extension), we can be selective in the arguments we present, as proposed in (Hunter 2020). So the role of the argument graph is to provide a presentation of the conflicts in the knowledge. For the purposes of this paper, we will restrict graphs to involve just an argument with an undercut. However, the framework in this paper supports any selection criteria, such as those in (Hunter 2020), including exhaustive presentation of arguments.

A **local view** on a knowledge tuple Φ is an argument obtained from Φ (i.e. each argument of the form $\langle \Gamma, p(\alpha) \# v \rangle$ where $\Gamma \subseteq \text{Set}(\Phi)$, and $\# \in \{<, \leq, =, \geq, >\}$ and $v \in [0, 1]$) and so each argument in an argument graph is a local view. Furthermore, each local view is a **global view** when Φ is consistent (as illustrated in Example 11).

We now return to the epistemic constraints on probability distributions. If a knowledge tuple Φ is consistent (i.e. $\text{Sat}(\text{Set}(\Phi), \Pi) \neq \emptyset$) then each probability distribution that satisfies the tuple is rational with respect to the arguments that can be constructed from the knowledge tuple.

Proposition 4. For a consistent knowledge tuple Φ , a value set Π , and $P \in \text{Sat}(\Phi, \Pi)$. If $A, B \in \text{Args}(\Phi, \Pi)$, and A is a counterargument to B , then P satisfies the rational constraint (i.e. $P(A) > 0.5$ implies $P(B) \leq 0.5$).

Proof. Since Φ is consistent, there is a $P \in \text{Sat}(\Phi, \Pi)$. Let A be of the form $p(\alpha) \#_a v_a$ and B be of the form $p(\beta) \#_b v_b$. Suppose A rebuts B . So $\alpha \equiv \neg\beta$. So $P(\alpha) + P(\beta) \leq 1$. So if $p(\alpha) \#_a v_a$ satisfies $P(A) > 0.5$ (i.e. $\text{Sat}(p(\alpha) \#_a v_a, \Pi) \subseteq \text{Sat}(p(\alpha) > 0.5, \Pi)$), then $p(\beta) \#_b v_b$ satisfies $P(B) \leq 0.5$ (i.e. $\text{Sat}(p(\beta) \#_b v_b, \Pi) \subseteq \text{Sat}(p(\beta) \leq 0.5, \Pi)$). Hence, P satisfies the rational constraint. The case is similar for undercut. \square

Example 11. The consistent knowledge tuple $[p(s) \geq 0.9, p(n) \geq 0.9, p(s \wedge n \rightarrow c) \geq 0.9, p(e) \leq 0.4, p(e \rightarrow \neg n) = 0.7, p(\neg n) \leq 0.5]$ is used for Figure 2 which satisfies the rational constraint. Each argument is a local and global view.




Figure 3: Argument graph constructed from an inconsistent knowledge tuple (i.e. there is no probability distribution that satisfies the union of the premises from the two arguments) where b denotes *bird*, f denotes *fly*, n denotes *normal bird*, and p denotes *penguin*. The knowledge tuple is composed of the five formulae appearing in the support of the arguments. Each argument is a local view. For global views, see Figures 4 and 5.

$\neg n) = 0.6]$ is used for Figure 2 which satisfies the rational constraint. Each argument is a local and global view.

When a knowledge tuple is not consistent, then there is not a probability distribution that satisfies all the arguments from the knowledge tuple (e.g. Figure 3). And so, if we want to have a coherent perspective on the uncertainty (i.e. the same probability distribution for all the arguments), and to be normative (e.g. satisfying the rational constraint), we need to relax belief in some of the premises. In other words, we need to relax belief in some of the formulae in the knowledge tuple.

5 Belief Relaxation

The aim of relaxation is to change the formulae in the knowledge tuple so that it becomes consistent. We will focus on belief statements for reasons of space in this paper. We assume that we do not relax the propositions (e.g. $p(a \wedge b) = 0.9$ is not relaxed to $p(a) = 0.9$). So we keep the propositions as they are but not the beliefs in them. Rather, we introduce a method for relaxation that is based on changing the values and/or the comparators in the belief formulae until the tuple is consistent.

Example 12. Consider $[p(a) > 0.75, p(a) = 0.75]$ with $\Pi = \{0, 0.25, 0.5, 0.75, 1\}$.

- *Valued-based relaxation results in the $[P(a) > 0.5, P(a) = 0.75]$ and $[p(a) > 0.75, p(a) = 1]$ relaxed knowledge tuples.*
- *Comparator-based relaxation results in the $[p(a) \geq 0.75, p(a) = 0.75]$ and $[p(a) > 0.75, p(a) \geq 0.75]$ relaxed knowledge tuples.*

We will proceed by identifying good relaxations which are the knowledge tuples obtained by minimal change.

Definition 13. For a knowledge tuple $\Phi = [p(\alpha_1) \#_1 v_1, \dots, p(\alpha_n) \#_n v_n]$, a **value tuple** is a tuple $[v_1, \dots, v_n]$, and a **comparator tuple** is a tuple $[\#_1, \dots, \#_n]$.

We change the values and/or comparators that appear in the formulae in the knowledge tuple using the following notion of an update function.

Definition 14. Let $\Phi = [p(\alpha_1) \#_1 v_1, \dots, p(\alpha_n) \#_n v_n]$ be a knowledge tuple, let $V = [v'_1, \dots, v'_n]$ be a value tuple, and let $C = [\#'_1, \dots, \#'_n]$ be a comparator tuple. The **update function**, denoted $\text{Update}(\Phi, V, C)$, is as follows.

$$\text{Update}(\Phi, V, C) = [p(\alpha_1) \#'_1 v'_1, \dots, p(\alpha_n) \#'_n v'_n]$$

Proposition 5. Let $\Phi = [p(\alpha_1)\#_1v_1, \dots, p(\alpha_n)\#_nv_n]$ be a knowledge tuple. If V is the value tuple for Φ , and C is the comparator tuple for Φ , then $\text{Update}(\Phi, V, C) = \Phi$.

Proof. Follows directly from Definitions 13 and 14. \square

The update function allows for semantically unacceptable formulae to be obtained. For example for the knowledge tuple $\Phi = [p(a) \leq 0.25, p(b) \geq 0.75]$, and the tuples $[<, >]$ and $[0, 1]$, we get $[p(a) < 0, p(b) > 1]$ as the update. To obviate this problem, we use a valid update.

Definition 15. Let Π be a value set. A **valid update** for Π and n is a pair (V, C) such that $V = [v_1, \dots, v_n] \in \Pi^n$ and $C = [\#_1, \dots, \#_n] \in \{<, \leq, =, \geq, >\}^n$ and for each i , if $\#_i$ is $>$, then $v_i \neq 1$, and if $\#_i$ is $<$, then $v_i \neq 0$. The set of valid updates for Π and n is $\text{ValidTuples}(\Pi, n)$.

We will chose the relaxed version of the knowledge tuple from the following set of candidates.

Definition 16. Let Π be a value set, and let Φ be a knowledge tuple, where $\text{Len}(\Phi) = n$. The set of **candidates**, denoted $\text{Candidates}(\Phi, \Pi)$, is $\text{Candidates}(\Phi, \Pi) = \{\text{Update}(\Phi, V, C) \mid (V, C) \in \text{ValidTuples}(\Pi, n)\}$.

Example 13. Consider $[p(\alpha) > 0.75, p(\alpha) = 0.75]$. The set of candidates includes the following

$$\begin{aligned} & [p(\alpha) > 0.75, p(\alpha) = 0.75] \quad [p(\alpha) > 0.75, p(\alpha) > 0.75] \\ & [p(\alpha) \leq 0.75, p(\alpha) = 0.75] \quad [p(\alpha) > 0.75, p(\alpha) < 0.75] \\ & [p(\alpha) > 0.5, p(\alpha) \leq 1] \quad [p(\alpha) = 0.5, p(\alpha) = 0.75] \end{aligned}$$

Proposition 6. If Π is a value set, where $|\Pi| = m$, and Φ is a knowledge tuple, where $\text{Len}(\Phi) = n$, then $|\text{Candidates}(\Phi, \Pi)| = (5m - 2)^n$.

Proof. There are five comparators and m restricted values for each belief statement (i.e. $5m$ possibilities). But two of these are not semantically valid (viz. < 0 and > 1). So for each formula, we have $5m - 2$ options in the update. So with n statement, there are $(5m - 2)^n$ knowledge tuples. \square

A belief statement specifies the values in the value set for which it holds. We get these values with the next function.

Definition 17. Let Π be a value set, and let $p(\alpha)\#_iv$ be a belief statement. The **interval** function, denoted $\text{Interval}(p(\alpha)\#_iv, \Pi)$, is defined as follows

$$\text{Interval}(p(\alpha)\#_iv, \Pi) = \{x \in \Pi \mid x\#_iv\}$$

Example 14. Let $\Pi = \{0, 0.25, 0.5, 0.75, 1\}$. $\text{Interval}(p(a) > 0.5, \Pi) = \{0.75, 1\}$.

Proposition 7. Let Π be a value set, and let $p(\alpha)\#_iv$ be a belief statement. If $v' \in \text{Intervals}(p(\alpha)\#_iv, \Pi)$, then there is a $P \in \text{Sat}(p(\alpha)\#_iv, \Pi)$ such that $P(\alpha) = v'$.

Proof. By definition $\text{Sat}(p(\alpha)\#_iv, \Pi) = \{P' \in \mathcal{P} \mid P'(\alpha)\#_iv\}$. So for each $P' \in \text{Sat}(p(\alpha)\#_iv, \Pi)$, $P'(\alpha) \in \Pi$, and $P'(\alpha) \in \text{Intervals}(p(\alpha)\#_iv, \Pi)$. \square

Next, we consider the possible belief that could be assigned to a proposition given the satisfying distributions of a candidate for the knowledge tuple. In other words, for an item $p(\alpha)\#_iv$ in the original knowledge tuple, we get the possible probability assignment for α if the original knowledge tuple is replaced by the candidate knowledge tuple.

Definition 18. For a knowledge tuple Ψ , and value set Π , the set of **qualified values** is $\text{Qualified}(p(\alpha)\#_iv, \Psi, \Pi) = \{v'' \in \Pi \mid P''(\alpha) = v'' \text{ and } P'' \in \text{Sat}(\text{Set}(\Psi), \Pi)\}$.

Example 15. Consider the formula $\phi = p(\alpha) > 0$ and $\Pi = \{0, 0.25, 0.5, 0.75, 1\}$. If the candidate is $\Psi = [p(\alpha) > 0.5, p(\alpha) < 1]$, then $\text{Qualified}(\phi, \Psi, \Pi)$ is $\{0.75\}$.

If the candidate is inconsistent, then there are no values for which the formula holds, and if it is satisfied by all the probability distributions, then the interval for the formula is contained in the set of qualified values.

Proposition 8. Let Ψ be a knowledge tuple, and let Π be a value set. (1) If $\text{Sat}(\text{Set}(\Psi), \Pi) = \emptyset$, then $\text{Qualified}(p(\alpha)\#_iv, \Psi, \Pi) = \emptyset$. (2) If $\text{Sat}(\text{Set}(\Psi), \Pi) = \mathcal{P}$, then $\text{Interval}(p(\alpha)\#_iv, \Pi) \subseteq \text{Qualified}(p(\alpha)\#_iv, \Psi, \Pi)$.

Proof. (1) Assume $\text{Sat}(\text{Set}(\Psi), \Pi) = \emptyset$. So for any $p(\alpha)\#_iv$, there is no $v'' \in \Pi$ s.t. $P''(\alpha) = v''$. (2) Assume $\text{Sat}(\text{Set}(\Psi), \Pi) = \mathcal{P}$. So $\text{Qualified}(p(\alpha)\#_iv, \Psi, \Pi) = \mathcal{P}$. So $\text{Interval}(p(\alpha)\#_iv, \Pi) \subseteq \text{Qualified}(p(\alpha)\#_iv, \Psi, \Pi)$. \square

Now we define a difference measure between a formula and a set of formula as the nearest point in the two intervals. We start with the Hausdorff distance which is a metric, i.e. satisfies symmetry, positivity, and triangle inequality, (see (Conci and Kubrusly 2017) for a review).

Definition 19. Let X and Y be two sets subsets of a value set Π . The **Hausdorff distance**, denoted $H(X, Y)$, is

$$H(X, Y) = \max\{\max_{x \in X} \min_{y \in Y} |x - y|, \max_{y \in Y} \min_{x \in X} |x - y|\}$$

The following definition gives the Hausdorff distance between the values that the original belief statement can take (given by the interval function) and the values that the revised belief statement can take (given by the qualified function). The revised belief statement is taken from the candidate knowledge tuple.

Definition 20. Let Π be a value set, let Ψ be a knowledge tuple, and let $\phi = p(\alpha)\#_iv$ be a belief statement. The **diff function** is defined as follows: If $\text{Sat}(\text{Set}(\Psi), \Pi) = \emptyset$, then $\text{Diff}(\phi, \Psi) = \infty$, otherwise $\text{Diff}(\phi, \Psi) = H(\text{Interval}(\phi, \Pi), \text{Qualified}(\phi, \Psi, \Pi))$.

Example 16. Let $\Pi = \{0, 0.25, 0.5, 0.75, 1\}$, let $\phi = p(a) < 0.5$, and let $\Psi = [p(a \wedge b) = 0.5, p(a) = 0.5]$. So $\text{Interval}(\phi, \Pi) = \{0, 0.25\}$, and $\text{Qualified}(\phi, \Psi, \Pi) = \{0.5\}$. So $\text{Diff}(\phi, \Psi) = 0.5$.

We explain the properties in Proposition 9 as follows: (1) A form of reflexivity (i.e. a formula that is updated to itself has zero difference); (2) A form of positivity (i.e. there is a set of values that satisfy the formula, and there is a set of values that satisfy its updated form, and the difference is then calculated as the Hausdorff distance between these two sets); and (3) A form of monotonicity (i.e. expanding the knowledge tuple reduces the set of satisfying distributions, and hence the Hausdorff distance is the same or greater).

Proposition 9. The Diff measure satisfies the following properties where Φ and $\Phi + \Psi$ (i.e. the concatenation of Φ and Ψ) are consistent knowledge tuples:

1. $\text{Diff}(p(\alpha)\#v, \Phi) = 0$ when $\Phi = [p(\alpha)\#v]$
2. $\text{Diff}(p(\alpha)\#v, \Phi) \geq 0$
3. $\text{Diff}(p(\alpha)\#v, \Phi) \leq \text{Diff}(p(\alpha)\#v, \Phi + \Psi)$

Proof. (1) If Φ is $[p(\alpha)\#v]$, then $\text{Interval}(p(\alpha)\#v', \Pi) = \text{Qualified}(p(\alpha)\#v, \Phi, \Pi)$. So $\text{Diff}(p(\alpha)\#v, \Phi) = 0$. (2) Follows from Definition 20 and the fact that the Hausdorff measure is a metric function. (3) First observe that $\text{Qualified}(p(\alpha)\#v, \Phi + \Psi, \Pi) \subseteq \text{Qualified}(p(\alpha)\#v, \Phi, \Pi)$ holds in general. Furthermore, observe that if $Y \subseteq Y'$, then $H(X, Y) \leq H(X, Y')$. Hence, $\text{Diff}(p(\alpha)\#v, \Phi) \leq \text{Diff}(p(\alpha)\#v, \Phi + \Psi)$ holds. \square

We will access a knowledge tuple by an index in order to retrieve the formula at that index position. So for a knowledge tuple $\Phi = [p(\alpha_1)\#_1v_1, \dots, p(\alpha_n)\#_nv_n]$, we have $\Phi[i] = p(\alpha_i)\#_iv_i$.

In general, for knowledge tuples Φ, Ψ and Ψ' , we have various options for defining Ψ is more (respectively equally) relaxed than Ψ' w.r.t. Φ , which we denote $\Psi \prec_{\Phi}^{\Psi'} \Psi'$ (respectively $\Psi \simeq_{\Phi}^{\Psi'} \Psi'$). Also, we let $\Psi \preceq_{\Phi}^{\Psi'} \Psi'$ hold iff $\Psi \prec_{\Phi}^{\Psi'} \Psi'$ or $\Psi \simeq_{\Phi}^{\Psi'} \Psi'$. We consider two options for this.

Definition 21. Let Φ be a knowledge tuple, where $\text{Len}(\Phi) = n$, and let $\Psi, \Psi' \in \text{Candidates}(\Phi, \Pi)$. The **max ranking**, denoted \preceq_{\max}^{Φ} , and the **sum ranking**, denoted $\preceq_{\text{sum}}^{\Phi}$, are defined as follows.

$$\Psi \preceq_{\max}^{\Phi} \Psi' \text{ iff } \max_{i \in \{1, \dots, n\}} \text{Diff}(\Phi[i], \Psi) \leq \max_{i \in \{1, \dots, n\}} \text{Diff}(\Phi[i], \Psi')$$

$$\Psi \preceq_{\text{sum}}^{\Phi} \Psi' \text{ iff } \sum_{i \in \{1, \dots, n\}} \text{Diff}(\Phi[i], \Psi) \leq \sum_{i \in \{1, \dots, n\}} \text{Diff}(\Phi[i], \Psi')$$

Proposition 10. For knowledge tuple tuples Φ, Ψ and Ψ' , if Ψ is consistent and Ψ' is inconsistent, then $\Psi \prec_{\max}^{\Phi} \Psi'$ and $\Psi \prec_{\text{sum}}^{\Phi} \Psi'$.

Proof. In Definition 20, $\text{Sat}(\text{Set}(\Psi'), \Pi) = \emptyset$, when the difference is ∞ . Therefore, for all Ψ s.t. $\text{Sat}(\text{Set}(\Psi), \Pi) \neq \emptyset$, $\max_{i \in \{1, \dots, n\}} \text{Diff}(\Phi[i], \Psi) < \infty$ and $\sum_{i \in \{1, \dots, n\}} \text{Diff}(\Phi[i], \Psi) < \infty$. Therefore, $\Psi \prec_{\max}^{\Phi} \Psi'$ and $\Psi \prec_{\text{sum}}^{\Phi} \Psi'$. \square

Proposition 11. The max ranking and the sum ranking are pre-ordering relations.

Proof. We show this for max ranking. (Reflexivity) We can assume that $\max_{i \in \{1, \dots, n\}} \text{Diff}(\Phi[i], \Psi) \leq \max_{i \in \{1, \dots, n\}} \text{Diff}(\Phi[i], \Psi)$. Hence, $\Psi \preceq_{\max}^{\Phi} \Psi$ holds. (Transitivity) Assume $\Psi \preceq_{\max}^{\Phi} \Psi'$ and $\Psi' \preceq_{\max}^{\Phi} \Psi''$. The first above (respectively the second above) implies $\max_{i \in \{1, \dots, n\}} \text{Diff}(\Phi[i], \Psi) \leq \max_{i \in \{1, \dots, n\}} \text{Diff}(\Phi[i], \Psi')$ (respectively $\max_{i \in \{1, \dots, n\}} \text{Diff}(\Phi[i], \Psi') \leq \max_{i \in \{1, \dots, n\}} \text{Diff}(\Phi[i], \Psi'')$). From these, we obtain the following $\max_{i \in \{1, \dots, n\}} \text{Diff}(\Phi[i], \Psi) \leq \max_{i \in \{1, \dots, n\}} \text{Diff}(\Phi[i], \Psi'')$. Since this is equivalent to $\Psi \preceq_{\max}^{\Phi} \Psi''$, we have shown transitivity. We can show these properties in the same way for sum ranking. \square

Based on the definition for ranking, we use the following definitions for good relaxations.

Definition 22. Let Π be a value set. For knowledge tuple $\Phi, \Psi \in \text{Candidates}(\Phi, \Pi)$ is a **good max relaxation** (resp. a **good sum relaxation**) of Φ iff there is no $\Psi' \in \text{Candidates}(\Phi, \Pi)$ s.t. $\Psi' \prec_{\max}^{\Phi} \Psi$ (resp. $\Psi' \prec_{\text{sum}}^{\Phi} \Psi$).

Example 17. Consider $\Phi = [p(\alpha) > 0.75, p(\alpha) = 0.75]$ with $\Pi = \{0, 0.25, 0.5, 0.75, 1\}$. The good sum (resp. max) relaxations are the first two (resp. all three) below.

Candidate knowledge tuple	Max	Sum
$[p(\alpha) > 0.75, p(\alpha) = 1]$	0.25	0.25
$[p(\alpha) \geq 0.75, p(\alpha) = 0.75]$	0.25	0.25
$[p(\alpha) \geq 0.75, p(\alpha) \geq 0.75]$	0.25	0.5

Example 18. Consider $\Phi = [p(\alpha) = 0.5, p(\alpha \wedge \beta) = 1]$ with $\Pi = \{0, 0.25, 0.5, 0.75, 1\}$. The good sum (resp. max) relaxations are the first three (resp. first) below.

Candidate knowledge tuple	Max	Sum
$[p(\alpha) = 0.75, p(\alpha \wedge \beta) = 0.75]$	0.25	0.5
$[p(\alpha) = 1, p(\alpha \wedge \beta) = 1]$	0.5	0.5
$[p(\alpha) = 0.5, p(\alpha \wedge \beta) = 0.5]$	0.5	0.5
$[p(\alpha) = 0.75, p(\alpha \wedge \beta) = 0.5]$	0.5	0.75
$[p(\alpha) = 1, p(\alpha \wedge \beta) = 0.75]$	0.5	0.75

We can always get a good relaxation of a knowledge tuple by flipping the comparators and/or values.

Proposition 12. For all knowledge tuples Φ , there is a good max relaxation, and a good sum relaxation, of Φ that is consistent.

Proof. Let $\Phi = [p(\alpha_1)\#_1v_1, \dots, p(\alpha_n)\#_nv_n]$ be a knowledge tuple. Let $\alpha_1, \dots, \alpha_n \in \mathcal{L}$ be the propositional formulae appearing in the knowledge tuple. For the propositional language, consider a probability distribution $P' : \wp(\mathcal{M}) \rightarrow [0, 1]$. For each $i \in \{1, \dots, n\}$, if $P'(\alpha_i) = v'_i$, then let $\Psi = [p(\alpha_1) = v'_1, \dots, p(\alpha_n) = v'_n]$. So Ψ is a consistent knowledge tuple. Now the comparators and values can be flipped exhaustively to identify a candidate that is minimal in the max ranking or the sum ranking. \square

So a good relaxation provides a measured (i.e. minimal) and fair (w.r.t. to the items in the tuple) way of relaxing the different items in a knowledge tuple in order to obtain arguments that are global views on the knowledge.

6 Ideal Relaxation

Good relaxation does not favour some items over others. However, often we do have extra information to enable us to prefer some items over others. So in order to make a better choice of relaxation, we introduce a probability distribution over subsets of the knowledge tuple to represent the confidence that the set of formulae is correct. So for the marginal for a formula, the higher the marginal confidence, the higher the inertia in changing it.

Definition 23. Let Π be a value set. A **confidence distribution** for a knowledge tuple $\Phi = [\phi_1, \dots, \phi_n]$ is a probability distribution $Pr : \wp(\{1, \dots, n\}) \rightarrow \Pi$ such that $\sum_{S \subseteq \{1, \dots, n\}} Pr(S) = 1$.

A confidence distribution is intended to be meat-information that represents the reliability of the belief statements and/or the quality of the sources of the belief statements. This could be learned over time.

Note, we can assign all confidence to the empty set, which we means that we have no confidence in any formula, or assign all confidence to the whole set, which would mean we have total confidence in all the formulae. These extremes just mean that we are unable to discriminate between the formulae.

Definition 24. Let Pr be a confidence distribution for a knowledge tuple $\Phi = [\phi_1, \dots, \phi_n]$. For a formula $\phi \in \Phi$, the **marginal confidence** of the i th formulae, denoted $Pr(i)$, is $\sum_{S \subseteq \{1, \dots, n\} \text{ s.t. } i \in S} Pr(S)$.

Example 19. Let $[p(b) = 0.75, p(b \rightarrow f) = 0.75, p(p) = 0.75, p(p \rightarrow \neg f) = 1]$ be the knowledge tuple, where b denotes bird, p denotes penguin, and f denotes fly, with confidence $Pr(\{1, 3, 4\}) = 0.6$, $Pr(\{1, 2\}) = 0.3$, and $Pr(\{1, 3\}) = 0.1$. So the marginals are $Pr(1) = 1$, $Pr(4) = 0.6$, $Pr(3) = 0.7$, and $Pr(2) = 0.3$.

We treat belief and confidence as orthogonal. Because of the separation of the two dimensions at the formal level, we do not have the possibility to infer confidence in a belief statement from confidence in other belief statements. For example, for $\Phi = [p(\alpha) = 1, p(\neg\alpha) = 0]$, we are unable to get $Pr(1) = 1$ implies $Pr(2) = 1$ or vice versa. Nonetheless, we can define a confidence distribution so that it does satisfy some reasonable constraints such as the following which says that if ϕ implies ψ , then the confidence in ψ is greater than ϕ .

Definition 25. Let Φ be a knowledge tuple, and let Pr be a confidence distribution. The **special constraint** on Pr and Φ is the following

$$\text{If } \text{Sat}(\Phi[i], \Pi) \subseteq \text{Sat}(\Phi[j], \Pi), \text{ then } Pr(i) \leq Pr(j)$$

The following proposition gives some of the desirable consequences of assuming the special constraint.

Proposition 13. Let Pr be a confidence distribution satisfying the special constraint, and assume $\{\alpha\} \vdash \beta$ holds. (1) If $\Phi[i] = p(\alpha) > v$ and $\Phi[j] = p(\beta) > v$ then $Pr(i) \leq Pr(j)$. (2) If $\Phi[i] = p(\alpha) < v$ and $\Phi[j] = p(\beta) < v$ then $Pr(j) \leq Pr(i)$. (3) If $\Phi[i] = p(\alpha) > v$ and $\Phi[j] = p(\beta) > w$ and $v < w$ then $Pr(i) \leq Pr(j)$.

Proof. (1) Assume $\Phi[i] = p(\alpha) > v$ and $\Phi[j] = p(\beta) > v$. For all P , $P(\alpha) > v$ implies $P(\beta) > v$. Therefore, $Pr(i) \leq Pr(j)$. (2) Now assume $\Phi[i] = p(\alpha) < v$ $\Phi[j] = p(\beta) < v$. For all P , $P(\beta) < v$ implies $P(\alpha) < v$. Therefore, $Pr(j) \leq Pr(i)$. (3) Now assume $\Phi[i] = p(\alpha) > v$ and $\Phi[j] = p(\beta) > w$ and $v > w$. For all P , $P(\alpha) > v$ implies $P(\beta) > w$. Therefore, $Pr(i) \leq Pr(j)$. \square

We now define the weighted version of the ranking relations we saw before. The confidence for each belief statement is used as a weight for that formula in the distance calculation. At the extremes, if the confidence is 1, then any change has maximal effect, whereas if the confidence is 0, then the change is not taken into account.




Figure 4: Argument graph for Example 21.

Definition 26. Let Π be a value set, let Φ be a knowledge tuple, where $\text{Len}(\Phi) = n$, let $\Psi, \Psi' \in \text{Candidates}(\Phi, \Pi)$, and let Pr be a confidence distribution. The **weighted max ranking**, denoted $\preceq_{\max}^{\Phi, Pr}$, and the **weighted sum ranking**, denoted $\preceq_{\text{sum}}^{\Phi, Pr}$, are defined as follows.

$$\Psi \preceq_{\max}^{\Phi, Pr} \Psi' \text{ iff } \max_{i \in \{1, \dots, n\}} Pr(i) \times \text{Diff}(\Phi[i], \Psi) \leq \max_{i \in \{1, \dots, n\}} Pr(i) \times \text{Diff}(\Phi[i], \Psi')$$

$$\Psi \preceq_{\text{sum}}^{\Phi, Pr} \Psi' \text{ iff } \sum_{i \in \{1, \dots, n\}} Pr(i) \times \text{Diff}(\Phi[i], \Psi) \leq \sum_{i \in \{1, \dots, n\}} Pr(i) \times \text{Diff}(\Phi[i], \Psi')$$

The following follows from a simple generalization of the proof for Proposition 11.

Proposition 14. The weighted max ranking and the weighted sum ranking are pre-ordering relations.

Based on the definition for weighted ranking, it is straightforward to provide the following definitions for ideal relaxations.

Definition 27. Let Π be a value set. For a knowledge tuple Φ , $\Psi \in \text{Candidates}(\Phi, \Pi)$ is an **ideal max relaxation** (resp. an **ideal sum relaxation**) of Φ iff there is no $\Psi' \in \text{Candidates}(\Phi, \Pi)$ s.t. $\Psi' \prec_{\max}^{\Phi, Pr} \Psi$ (resp. $\Psi' \prec_{\text{sum}}^{\Phi, Pr} \Psi$).

Example 20. Consider the knowledge tuple $[p(a) = 1, p(a) = 0]$ with $\Pi = \{0, 0.25, 0.5, 0.75, 1\}$ and with $Pr(p(a) = 1) = 0.9$ and $Pr(p(a) = 0) = 0.1$. The ideal relaxation is $[p(a) = 1, p(a) = 1]$ where the max and sum distances are 0.1.

Note, if it were the case that $Pr(1) = Pr(2)$ in the above example, then there would be multiple ideal relaxations. In general when there multiple relaxations, we could make an arbitrary choice, or we could summarize the range of relaxations by using intervals.

Example 21. For value set $\Pi = \{0, 0.25, 0.5, 0.75, 1\}$, and knowledge tuple $[p(b) = 1, p(n) = 0.75, p(b \wedge n \rightarrow f) = 0.75, p(p) = 0.75, p(p \rightarrow \neg n) = 1]$, where b is bird, n is normal bird, f is fly, and p is penguin. Each argument in Figure 3 is a local view based on this knowledge. For confidence $Pr(1) = 1$, $Pr(2) = 0.2$, $Pr(3) = 0.6$, $Pr(4) = 0.6$, and $Pr(5) = 1$, the ideal relaxation is $[p(b) = 1, p(n) = 0.25, p(b \wedge n \rightarrow f) = 0.75, p(p) = 0.75, p(p \rightarrow \neg n) = 1]$, and the resulting global views are given in Figure 4 where the undercut (and not the root) is in the epistemic extension.

Example 22. Consider Example 21 but with the confidence $Pr(1) = 1$, $Pr(2) = 0.8$, $Pr(3) = 1$, $Pr(4) = 0.2$, and $Pr(5) = 0.8$. The ideal relaxation is $[p(b) = 1, p(n) = 0.75, p(b \wedge n \rightarrow f) = 0.75, p(p) = 0.5, p(p \rightarrow \neg n) = 0.75]$, and resulting global views given in Figure 5 where the root (and not the undercut) is in the epistemic extension.




Figure 5: Argument graph for Example 22.

Proposition 15. *For all knowledge tuples Φ , and all confidence distributions Pr , there is an ideal max relaxation, and an ideal sum relaxation, of Φ that is consistent.*

This proposition follows from a simple generalization of the proof for Proposition 12.

7 Comparison with the Literature

Two important approaches to probabilistic abstract argumentation are the constellations and the epistemic approaches (Hunter 2013): In the constellations approach, there is uncertainty over the structure of the graph (see (Dung and Thang 2010; Li, Oren, and Norman 2012; Hunter 2012) for instances of the approach); Whereas in the epistemic approach (which is the approach used in this paper), the topology of the argument graph is fixed, but there is uncertainty about whether an argument is believed (Thimm 2012; Hunter 2013; Baroni, Giacomin, and Vicig 2014; Gabbay and Rodrigues 2015; Hunter and Thimm 2017). A further approach is based on labellings for arguments using *in*, *out*, and *undecided*, from (Caminada and Gabbay 2009), augmented with *off* for arguments not occurring in the graph (Riveret and Governatori 2016). A probability distribution over labellings gives a form of probabilistic argumentation that overlaps with the constellations and epistemic approaches.

At the structured level, Haenni (Haenni 1998) considered a restricted form of probabilistic argumentation in which pros and cons are generated from a classical logic knowledgebase, and then a probability distribution over models of the language are used to assign a belief in each argument. Subsequently, this was generalized by Hunter to arbitrary argument graphs (Hunter 2013) in which various kinds of counterargument can be accommodated. More recently, Prakken considered a similar approach for ASPIC+ (Prakken 2018). In other logic-based proposals, Verheij has combined probabilities with non-monotonic inference (Verheij 2012), and separately, he has combined qualitative reasoning in terms of reasons and defeaters (adapting Pollock’s definitions (Pollock 1995)), with quantitative reasoning using argument strength, modeled as the conditional probability of the conclusions given the premises (Verheij 2014). In these proposals, the language is a form of propositional language, rather than a probabilistic language, the inference relation is for rules or propositional formulae, rather than defeasible inference for the probabilistic formulae, and there is no consideration of how the knowledge could be revised so that we could get a coherent view on it, and thereby on the probability of acceptability of the arguments. In a rule-based system for dialogical argumentation, the belief in the premises of an argument is used to calculate the belief in the

argument, though the nature of this belief is not investigated (Riveret et al. 2007).

Bayesian networks can be used to model argumentative reasoning (Vreeswijk 2004; Grabmair, Gordon, and Walton 2010). Conversely, arguments can be generated from a Bayesian network, and used to explain the network (Timmer et al. 2015). Argumentation can also be used to construct Bayesian networks (Bex and Renooij 2016; Wieten et al. 2019), and it can be used to combine multiple Bayesian networks (Nielsen and Parsons 2007). However, none of the above use a probabilistic logic for representing and reasoning with uncertain knowledge to generate arguments, and none offer methods for resolving conflict to give a consistent knowledgebase for constructing arguments.

8 Discussion

In this paper, we have provided a framework for reasoning with inconsistent knowledge using probabilistic argumentation with the following features: (1) A probabilistic logic for generating arguments from belief formulae; and (2) Belief relaxation to find a coherent knowledge tuple. Whilst the probabilistic logic, the maximum entropy reasoning, the distance measure, etc are harnessed from the literature, the novelty in our approach is in putting them together for computational argumentation.

This proposal means that we can harness uncertain formulae in argumentation. There are many applications where uncertainty in propositional formulae can be quantified. This may be because there are experts available to assign subjective probabilities to the items of knowledge, or because the propositional formulae have been obtained by statistical analysis of data (e.g. from scientific studies or from surveys with participants), or by machine learning (and so for each belief statement, both the proposition and the belief assignment can be automatically obtained). Then using relaxation we obtain a consistent knowledge tuple that means there is at least one probability distribution that satisfies all the premises of the arguments, and satisfies normative properties such as the rational constraint.

In future work, we will consider alternative monotonic and non-monotonic entailment relations for belief formulae, we will investigate alternative constraints on confidence distribution, we will generalize relaxation to belief formulae, and we will develop methods for harnessing SAT technology for entailment and relaxation. We will also investigate incorporating further kinds of probabilistic reasoning including formalisms for handling inconsistent probabilistic knowledge (e.g. (Finthammer, Kern-Isberner, and Ritterkamp 2007; Potyka and M.Thimm 2015)).

References

- Atkinson, K., and *et al.* 2017. Towards artificial argumentation. *AI Magazine* 38(3):25–36.
- Baroni, P.; Giacomin, M.; and Vicig, P. 2014. On rationality conditions for epistemic probabilities in abstract argumentation. In *Proceedings of COMMA’14*. IOS Press.
- Bex, F., and Renooij, S. 2016. From arguments to

- constraints on a Bayesian network. In *Proceedings of COMMA'16*, 95–106. IOS Press.
- Caminada, M., and Gabbay, D. 2009. A logical account of formal argumentation. *Studia Logica* 93:109–145.
- Cheeseman, P. 1983. A method of computing generalized Bayesian probability values for expert systems. In *Proceedings IJCAI'83*, 198–202. IJCAI.
- Conci, A., and Kubrusly, C. 2017. Distance between sets: A survey. *Advances in Mathematical Sciences and Applications* 26:1–18.
- Dung, P., and Thang, P. 2010. Towards (probabilistic) argumentation for jury-based dispute resolution. In *Proceedings of COMMA'10*, 171–182. IOS Press.
- Dung, P. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77(2):321–358.
- Finthammer, M.; Kern-Isberner, G.; and Ritterskamp, M. 2007. Resolving inconsistencies in probabilistic knowledge bases. In *Proceedings of KI'07*, volume 4667 of *LNCS*, 114–128. Springer.
- Gabbay, D., and Rodrigues, O. 2015. Probabilistic argumentation: an equational approach. *Logica Universalis* 9(3):345–382.
- Goldszmidt, M.; Morris, P.; and Pearl, J. 1993. A maximum entropy approach to nonmonotonic reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15:220–232.
- Grabmair, M.; Gordon, T.; and Walton, D. 2010. Probabilistic semantics for the carneades argument model using Bayesian networks. In *Proceedings of COMMA'10*, 255–266. IOS Press.
- Haenni, R. 1998. Modelling uncertainty with propositional assumptions-based systems. In *Applications of Uncertainty Formalisms*, volume 1455 of *LNCS*, 446–470. Springer.
- Hunter, A., and Thimm, M. 2017. Probabilistic reasoning with abstract argumentation frameworks. *Journal of Artificial Intelligence Research* 59:565–611.
- Hunter, A.; Polberg, S.; and Thimm, M. 2020. Epistemic graphs for representing and reasoning with positive and negative influences of arguments. *Artificial Intelligence* 281:103236.
- Hunter, A. 2012. Some foundations for probabilistic abstract argumentation. In *Proceedings of COMMA'12*, 117–128. IOS Press.
- Hunter, A. 2013. A probabilistic approach to modelling uncertain logical arguments. *International Journal of Approximate Reasoning* 54(1):47–81.
- Hunter, A. 2020. Generating instantiated argument graphs from probabilistic information. In *Proceedings of ECAI'20*. IOS Press. in press.
- Jaynes, E. 1982. On the rationale of maximum entropy methods. *Proceedings of the IEEE* 70(9):939–952.
- Kern-Isberner, G., and Lukasiewicz, T. 2004. Combining probabilistic logic programming with the power of maximum entropy. *Artificial Intelligence* 157:139–202.
- Li, H.; Oren, N.; and Norman, T. J. 2012. Probabilistic argumentation frameworks. In *Proceedings of TAFE'11*, volume 7132 of *LNCS*, 1–16. Springer.
- Likert, R. 1931. A technique for the measurement of attitudes. *Archives of Psychology* 140:1–55.
- Nielsen, S., and Parsons, S. 2007. An application of formal argumentation: fusing Bayesian networks in multi-agent systems. *Artificial Intelligence* 171(10-15):754–775.
- Nilsson, N. 1995. Probabilistic logic. *Artificial Intelligence* 28:71–87.
- Paris, J. B. 1994. *The Uncertain Reasoner's Companion – A Mathematical Perspective*. Cambridge University Press.
- Paskin, M. 2002. Maximum entropy probabilistic logic. Technical Report UCB/CSD-01-1161, University of California Berkeley.
- Polberg, S., and Hunter, A. 2018. Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches. *International Journal of Approximate Reasoning* 93:487–543.
- Pollock, J. 1995. *Cognitive Carpentry*. MIT Press.
- Potyka, N., and M.Thimm. 2015. Probabilistic reasoning with inconsistent beliefs using inconsistency measures. In *Proceedings of IJCAI'2015*, 3156–3163. AAAI Press.
- Prakken, H. 2018. Probabilistic strength of arguments with structure. In *Proceedings of KR'18*, 158–167. AAAI Press.
- Riveret, R., and Governatori, G. 2016. On learning attacks in probabilistic abstract argumentation. In *Proceedings of AAMAS'16*, 653–661. IFAAMAS.
- Riveret, R.; Rotolo, A.; Sartor, G.; Prakken, H.; and Roth, B. 2007. Success chances in argument games: a probabilistic approach to legal disputes. In *Proceedings of JURIX'07*, 99–108. IOS Press.
- Thimm, M. 2012. A probabilistic semantics for abstract argumentation. In *Proceedings of ECAI'12*, 750–755. IOS Press.
- Timmer, S. T.; Meyer, J.-J. C.; Prakken, H.; Renooij, S.; and Verheij, B. 2015. Explaining Bayesian networks using argumentation. In *Proceedings of ECSQARU'15*, volume 9161 of *LNCS*, 83–92. Springer.
- Verheij, B. 2012. Jumping to conclusions: A logico-probabilistic foundation for defeasible rule-based arguments. In *Proceedings of JELIA'12*, volume 7519 of *LNCS*, 411–423. Springer.
- Verheij, B. 2014. Arguments and their strength: revisiting Pollock's anti-probabilistic starting points. In *Proceedings of COMMA'14*, 433–444. IOS Press.
- Vreeswijk, G. 2004. Argumentation in Bayesian belief networks. In *Proceedings of ArgMAS'04*, volume 3366 of *LNCS*, 111–129. Springer.
- Wieten, R.; Bex, F.; Prakken, H.; and Renooij, S. 2019. Supporting discussions about forensic Bayesian networks using argumentation. In *Proceedings of ICAIL'19*, 143–152. ACM Press.