

Updating Belief in Arguments in Epistemic Graphs *

Anthony Hunter, Sylwia Polberg, Nico Potyka

Department of Computer Science, University College London, UK,
Institute of Cognitive Science, University of Osnabrück, Germany

Abstract

Epistemic graphs are a recent generalization of epistemic probabilistic argumentation. Relations between arguments can be supporting, attacking, as well as neither supporting nor attacking. These interdependencies are represented by epistemic constraints, and the semantics of epistemic graphs are given in terms of probability distributions satisfying these constraints. We investigate the behaviour of epistemic graphs in a dynamic setting where a given distribution can be updated once new constraints are presented. Our focus is on update methods that minimize the change in probabilistic beliefs. We show that all methods satisfy basic commonsense postulates, identify fragments of the epistemic constraint language that guarantee the existence of well-defined solutions, and explain how the problems that arise in more expressive fragments can be treated either automatically or by user support. We demonstrate the usefulness of our proposal by considering its application in computational persuasion.

1 Introduction

Abstract argumentation is a rich research area and many of the recent developments have been focused on fine-grained approaches towards argument acceptability (Thimm 2012; Hunter 2013; Hunter and Thimm 2014; Cayrol and Lagasque-Schiex 2005; Leite and Martins 2011; Rago et al. 2016; Bonzon et al. 2016; Amgoud and Ben-Naim 2017; Brewka et al. 2018). Indeed, the empirical study carried out in (Polberg and Hunter 2018) shows the need for such methods when considering argumentation involving human participants. Amongst the aforementioned works we can find epistemic probabilistic argumentation (Thimm 2012; Hunter 2013; Hunter and Thimm 2014), which allows us to express the degree to which a given argument is believed or disbelieved in terms of probabilities. However, it has been primarily developed in the context of argumentation graphs involving only the attack relation, and it does not meet the need for bipolar settings highlighted by the empirical study. Therefore, (Hunter, Polberg, and Thimm 2018) recently introduced epistemic graphs as a generalization of epistemic probabilistic argumentation.

*This research was partly funded by the EPSRC Framework for Computational Persuasion project.
Copyright © 2015, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In epistemic graphs, an argument can be believed or disbelieved to a given degree, and how other arguments influence a given argument is expressed by epistemic constraints. Epistemic graphs are capable of handling **positive, negative, as well as neither positive nor negative relations** between arguments. They are **context sensitive** in the sense that the same graph structure can have different constraints associated with them. Assume an argument A, its attacker B and supporter C. For one instantiation of these arguments, A may be believed as long as C is believed more than B, whereas for another instantiation, A may be believed only if B is disbelieved. We can model these different situations using constraints. Epistemic graphs also allow **modelling of different perspectives**. The same graph structure with the same instantiations for the arguments can have different constraints, thus allowing for different agents to give their views on the relations between arguments. The graphs also allow for **modelling of imperfect agents** and **of incomplete situations**. For example, it is possible to model an agent who believes both an argument and its attacker, or who disbelieves an unattacked argument. These features are important in predicting how certain real-world agents reason, modelling agents which might be unable or unwilling to provide their counterarguments, and dealing with enthymeme arguments that can be decoded differently by the agents.

Semantics of epistemic graphs are given in terms of probability distributions satisfying the constraints associated with a given graph. A natural question is how a given distribution should be updated when new constraints are presented. In this work, we identify a number of intuitive properties that a reasonable update method should satisfy and present several distance minimizing update approaches satisfying them. This includes methods that consider distances between beliefs in sets of arguments as well as distances between the beliefs in arguments themselves. Guarantees for the existence and uniqueness of solutions depend on the expressiveness of our constraints. We thus identify the fragments of our epistemic language that have the desired properties and propose approaches for handling the parts that do not. Finally, we demonstrate the usefulness of our proposal by considering its application in computational persuasion.

The paper is organized as follows: Section 2 introduces epistemic graphs. We present postulates for epistemic update functions in Section 3 and propose several distance

minimizing methods that meet our requirements in Sections 4 and 5. Section 6 demonstrates how our approach can be harnessed in user modelling in persuasion dialogues. We discuss related work in Section 7 and conclude in Section 8.

2 Epistemic Graphs

An argument graph informs us what are the arguments and relations between them. It can be seen as a directed graph in which nodes represent arguments and arcs represent relations. We associate each arc with a label representing its nature. We use a positive label to denote a positive influence (i.e. support), a negative label to denote a negative influence (i.e. attack), and a star label to denote an influence that is neither strictly positive nor negative (i.e. dependency).

Definition 2.1. Let $\mathcal{G} = (V, R)$, where V is a set of nodes and $R \subseteq V \times V$ is a set of arcs, be a directed graph. A **labelled graph** is a tuple $X = (\mathcal{G}, \mathcal{L})$ where $\mathcal{L} : R \rightarrow 2^\Omega$ is labelling function and Ω is a set of possible labels. X is **fully labelled** iff for every $\alpha \in R$, $\mathcal{L}(\alpha) \neq \emptyset$.

Here, we assume that $\Omega = \{+, -, *\}$ and that the graph is fully labelled. $\text{Nodes}(\mathcal{G}) = V$ denotes the nodes and $\text{Arcs}(\mathcal{G}) = R$ denotes the arcs in \mathcal{G} . The parents of a node $B \in \text{Nodes}(\mathcal{G})$ are $\text{Parent}(B) = \{A \mid (A, B) \in \text{Arcs}(\mathcal{G})\}$.

The epistemic language introduced in (Hunter, Polberg, and Thimm 2018) consists of Boolean combinations of inequalities involving statements about probabilities of formulae built out of arguments. In this work we assume a simple refinement of this approach which we define next:

Definition 2.2. Let \mathcal{G} be a directed graph. The epistemic language based on \mathcal{G} is defined as follows:

- a **term** is a Boolean combination of arguments¹. $\text{Terms}(\mathcal{G})$ denotes all the terms that can be formed from the arguments in \mathcal{G} .
- a **linear operational formula** is a formula $\sum_{i=1}^k c_i \cdot p(\alpha_i)$ where all $\alpha_i \in \text{Terms}(\mathcal{G})$ and $c_i \in \mathbb{Q}$ are rational numbers. $\text{LOFormulae}(\mathcal{G})$ is the set of all linear operational formulae of \mathcal{G} and we read $p(\alpha)$ as “probability of α ”.
- a **linear epistemic atom** is of the form $\alpha \# x$ where $\# \in \{=, \neq, \geq, \leq, >, <\}$, $x \in \mathbb{Q}$ and $\alpha \in \text{LOFormulae}(\mathcal{G})$.
- a **linear epistemic formula** is a Boolean combination of linear atoms. $\text{LFormulae}(\mathcal{G})$ denotes the set of all possible linear epistemic formulae of \mathcal{G} .

For $\alpha \in \text{Terms}(\mathcal{G})$, $\text{Args}(\alpha)$ denotes the set of all arguments appearing in α and for a set of terms $\Gamma \subseteq \text{Terms}(\mathcal{G})$, $\text{Args}(\Gamma)$ denotes the set of all arguments appearing in Γ . Given a formula $\psi \in \text{LFormulae}(\mathcal{G})$, let $\text{FTerms}(\psi)$ denote the set of terms appearing in ψ and let $\text{FArgs}(\psi) = \text{Args}(\text{FTerms}(\psi))$ be the set of arguments appearing in ψ .

Example 2.3. For $A, B, C, D \in \text{Nodes}(\mathcal{G})$, $\psi : p(A \wedge B) - p(C) - p(D) > 0$ is an example of a linear epistemic formula. The terms of that formula are $\text{FTerms}(\psi) = \{A \wedge B, C, D\}$, the arguments appearing in them are $\text{FArgs}(\psi) = \{A, B, C, D\}$.

¹We use \vee , \wedge and \neg as connectives in the usual way, and can derive secondary connectives, such as implication \rightarrow , as usual.

Having defined the syntax of our language, let us now focus on its semantics. A **belief distribution** on arguments is a function $P : 2^{\text{Nodes}(\mathcal{G})} \rightarrow [0, 1]$ s.t. $\sum_{\Gamma \subseteq \text{Nodes}(\mathcal{G})} P(\Gamma) = 1$. With $\text{Dist}(\mathcal{G})$ we denote the set of all belief distributions on $\text{Nodes}(\mathcal{G})$. Each $\Gamma \subseteq \text{Nodes}(\mathcal{G})$ corresponds to an interpretation of arguments. We say that Γ *satisfies an argument* A and write $\Gamma \models A$ iff $A \in \Gamma$. The satisfaction relation is extended to complex terms as usual. For instance, $\Gamma \models \neg\alpha$ iff $\Gamma \not\models \alpha$ and $\Gamma \models \alpha \wedge \beta$ iff $\Gamma \models \alpha$ and $\Gamma \models \beta$.

The probability of a term is defined as the sum of the probabilities (beliefs) of its models:

$$P(\alpha) = \sum_{\Gamma \subseteq \text{Nodes}(\mathcal{G}) \text{ s.t. } \Gamma \models \alpha} P(\Gamma).$$

We say that an agent believes a term α to some degree if $P(\alpha) > 0.5$, disbelieves α to some degree if $P(\alpha) < 0.5$, and neither believes nor disbelieves α when $P(\alpha) = 0.5$.

Definition 2.4. Let φ be a linear atom $\sum_{i=1}^k c_i \cdot p(\alpha_i) \# b$. The **satisfying distributions** of φ are defined as $\text{Sat}(\varphi) = \{P' \in \text{Dist}(\mathcal{G}) \mid \sum_{i=1}^k c_i \cdot P'(\alpha_i) \# b\}$.

The set of satisfying distributions for a linear formula is as follows where ϕ and ψ are linear formulae:

- $\text{Sat}(\phi \wedge \psi) = \text{Sat}(\phi) \cap \text{Sat}(\psi)$;
- $\text{Sat}(\phi \vee \psi) = \text{Sat}(\phi) \cup \text{Sat}(\psi)$; and
- $\text{Sat}(\neg\phi) = \text{Sat}(\top) \setminus \text{Sat}(\phi)$.

For a set of linear formulae $\Phi = \{\phi_1, \dots, \phi_n\}$, the set of satisfying distributions is $\text{Sat}(\Phi) = \text{Sat}(\phi_1) \cap \dots \cap \text{Sat}(\phi_n)$.

Epistemic constraints are epistemic formulae that contain at least one argument. Epistemic graphs are labelled graphs equipped with a set of such constraints:

Definition 2.5. A **linear epistemic constraint** is a linear epistemic formula $\psi \in \text{LFormulae}(\mathcal{G})$ s.t. $\text{FArgs}(\psi) \neq \emptyset$. An **epistemic graph** is a tuple $(\mathcal{G}, \mathcal{L}, \mathcal{C})$ where $(\mathcal{G}, \mathcal{L})$ is a labelled graph, and $\mathcal{C} \subseteq \text{LFormulae}(\mathcal{G})$ is a set of epistemic constraints associated with the graph.

Example 2.6. Consider a graph with nodes $\{A, B, C, D\}$ and the constraint $\psi : p(A \wedge B) - p(C) - p(D) > 0 \wedge p(D) > 0$. A probability distribution P_1 with $P_1(A \wedge B) = 0.7$, $P_1(C) = 0.1$ and $P_1(D) = 0.1$ is in $\text{Sat}(\psi)$. However, a distribution P_2 with $P_2(A \wedge B) = 0$ cannot satisfy ψ and so $P_2 \notin \text{Sat}(\psi)$.

The semantics of epistemic graphs are given in terms of probability distributions. A range of semantics have been proposed in (Hunter, Polberg, and Thimm 2018). In the context of this work it suffices to focus on the simplest one, demanding that the constraints of the graph are satisfied:

Definition 2.7. Let $X = (\mathcal{G}, \mathcal{L}, \mathcal{C})$ be an epistemic graph. An **epistemic semantics** associates X with a set $\mathcal{R} \subseteq \text{Dist}(\mathcal{G})$. A distribution $P \in \text{Dist}(\mathcal{G})$ meets the **satisfaction semantics** iff $P \in \text{Sat}(\mathcal{C})$.

We say that a framework is constraint consistent iff $\text{Sat}(\mathcal{C}) \neq \emptyset$, i.e. the satisfaction semantics produces at least one distribution for this graph.

Example 2.8. Consider the labelled graph in Figure 1 and imagine a passenger named Terry. We model Terry’s opinions on how the arguments interact in the following manner.




Figure 1: Labelled graph for Example 2.8. Edges labelled – denote attack and edges labelled + denote support.

Since A is attacked by B and C, we want that the belief in A is bounded from above by the average belief in B and C. This can be described with the formula $p(A) + \frac{1}{2}p(B) + \frac{1}{2}p(C) \leq 1$. Since D supports A, we also want to impose a lower bound on the belief in A. This lower bound can be decreased by the average belief in B and C. We capture this intuition with the formula $p(A) + \frac{1}{2}p(B) + \frac{1}{2}p(C) - p(D) \geq 0$. Finally, Terry is a regular on this line and believes that the train normally arrives late. We model this by the formula $p(B) \geq 0.65$. Probability distributions P s.t. $P(A) = 0.45$, $P(B) = 0.65$, $P(C) = 0.2$ and $P(D) = 0.5$, and P' s.t. $P'(A) = 0.425$, $P'(B) = 0.65$, $P'(C) = 0.5$ and $P'(D) = 0.5$, are examples of satisfying distributions of this graph.

3 Postulates for Epistemic Update Functions

Whether we consider dialogical or monological argumentation, learning new information calls for an update in our beliefs. We are therefore interested in developing epistemic update functions, which take our current epistemic state (i.e. current belief distribution) and an epistemic formulae representing new information, and return a set of candidates for the next epistemic state.

Definition 3.1 (Update Function). An update function for an epistemic graph $(\mathcal{G}, \mathcal{L}, \mathcal{C})$ is a function $U : \text{Dist}(\mathcal{G}) \times \mathcal{L} \rightarrow \mathcal{L}$.

In order to guarantee meaningful update functions, we consider properties similar to (Hunter and Potyka 2017).

- **Uniqueness:** $|U(P, \Psi)| \leq 1$.
- **Completeness:** If $\text{Sat}(\mathcal{C} \cup \Psi) \neq \emptyset$, then $|U(P, \Psi)| \geq 1$.
- **Epistemic Consistency:** $U(P, \Psi) \subseteq \text{Sat}(\mathcal{C})$.
- **Success:** $U(P, \Psi) \subseteq \text{Sat}(\Psi)$.
- **Tautology:** If $\text{Sat}(\Psi) = \text{Sat}(\top)$ then $U(P, \Psi) = \{P\}$.
- **Contradiction:** If $\text{Sat}(\mathcal{C} \cup \Psi) = \emptyset$ then $U(P, \Psi) = \emptyset$.
- **Representation Invariance:** If Ψ_1, Ψ_2 are equivalent, i.e., $\text{Sat}(\Psi_1) = \text{Sat}(\Psi_2)$, then $U(P, \Psi_1) = U(P, \Psi_2)$.
- **Idempotence:** If $U(P, \Psi) = \{P^*\}$ then $U(P^*, \Psi) = \{P^*\}$.

Uniqueness guarantees that there is at most one candidate for the next epistemic state, and *Completeness* states that there is at least one if the update is consistent. If both properties are satisfied, the next epistemic state is uniquely

defined whenever the update is consistent. *Epistemic Consistency* demands that the constraints in our graph are maintained and *Success* demands that the next state satisfies the beliefs that we updated with. *Tautology* states that updating with a tautological set of constraints should not change anything and *Contradiction* that an inconsistent update should yield the empty set. *Representation invariance* guarantees that changing the syntactic representation of updates does not change the outcome of the update. Finally, *Idempotence* demands that a repeated update does not change beliefs.

Throughout the next sections, we introduce some update functions and investigate under which conditions they satisfy our desiderata. Subsequently, we consider how these update functions can be harnessed in dialogical approaches.

4 Distance Minimizing Update Functions

We now focus on the update functions that minimize some notion of distance to the prior epistemic state. Our distance functions may not necessarily be metrics, but we assume that they satisfy the properties explained below:

Definition 4.1 (Epistemic Distance Function). An epistemic distance function is a function $d : \text{Dist}(\mathcal{G}) \times \text{Dist}(\mathcal{G}) \rightarrow \mathbb{R}$ that satisfies

1. **Positive Definiteness:** $d(P, P') \geq 0$ and $d(P, P') = 0$ iff $P = P'$.
2. **Continuity:** d is continuous in the second argument.
3. **Strict Convexity:** d is strictly convex in the second argument.

Continuity and convexity are defined as usual (Rudin 1976). Intuitively, continuity guarantees that probability distributions that assign similar probabilities to subsets of arguments have a low distance value. Strict convexity guarantees that there often is a unique solution and no non-global local minima when we minimize the distance. Popular examples of epistemic distance functions are the Least Squares distance and the KL-divergence.

• **Least Squares Distance:**

$$d_2(P, P') = \sum_{X \subseteq \text{Nodes}(\mathcal{G})} (P(X) - P'(X))^2$$

• **KL-divergence:**

$$d_{KL}(P, P') = \sum_{X \subseteq \text{Nodes}(\mathcal{G})} P(X) \cdot \log \frac{P(X)}{P'(X)}$$

The KL-divergence is an example of an epistemic distance function that is not a metric (it does not satisfy symmetry and the triangle-inequality). However, it still has some intuitive geometric properties and is a popular measure to compare probability distributions (Csiszar 1975). We focus on update functions that minimize some epistemic distance function to a prior belief state. In the definition of optimization problems, “ $\min f(x)$ ” denotes the minimum function value that f takes over the feasible region and “ $\arg \min f(x)$ ” denotes the set of points where f takes this value.

Definition 4.2 (Distance-minimizing Update Function). Given some epistemic distance function d , the distance-minimizing update function w.r.t. d is defined by

$$U_d(P, \Psi) = \arg \min_{P' \in \text{Sat}(\mathcal{C} \cup \Psi)} d(P, P')$$

for all finite sets of formulae $\Psi \subseteq \mathcal{L}$.

To begin with, we note that every distance-minimizing update functions necessarily satisfies all our desiderata other than uniqueness and completeness.

Proposition 4.3. *Every distance-minimizing update functions satisfies Epistemic Consistency, Success, Tautology, Contradiction, Representation Invariance and Idempotence.*

Uniqueness and and Completeness are more subtle and depend on the numerical nature of our constraints. Throughout the next sections, we will investigate which fragments of the language of linear epistemic formulae guarantee uniqueness and completeness. The following proposition translates a useful standard result from the theory of numerical optimization to our framework.

Lemma 4.4. *If $\text{Sat}(\mathcal{C} \cup \Psi)$ is non-empty, convex and compact, then $U_d(P, \Psi) = \{P^*\}$ is a singleton set and can be computed by means of convex programming techniques.*

For the formal definition of convex and compact sets, we refer again to (Rudin 1976). Roughly speaking, convexity means that the set is closed under weighted combinations of probability distributions. A compact set is bounded and contains all points on its boundary similar to a compact interval.

We will now look at some fragments of the language of linear epistemic formulae that give us additional guarantees.

Updating with Non-strict Epistemic Atoms

To begin with, we restrict ourselves to atoms that contain only non-strict inequalities and equality.

Definition 4.5 (Non-strict Epistemic Atom). *A **non-strict epistemic atom** is a linear atom $\sum_{i=1}^n c_i \cdot p(\alpha_i) \# b$ where $\# \in \{\leq, =, \geq\}$.*

If \mathcal{C} and Ψ consist only of non-strict epistemic atoms, $\text{Sat}(\mathcal{C} \cup \Psi)$ is always well-behaved in the following sense.

Proposition 4.6. *If $\mathcal{C} \cup \Psi$ contains only non-strict epistemic atoms, then $\text{Sat}(\mathcal{C} \cup \Psi)$ is compact and convex.*

Hence, for the fragment of non-strict epistemic axioms, Proposition 4.3 guarantees that distance-minimizing update functions satisfy our remaining desiderata.

Theorem 4.7. *In the fragment of non-strict epistemic axioms, every distance-minimizing update function satisfies Uniqueness and Completeness.*

Let us note again that both the least-squares distance and the KL-divergence satisfy the assumptions on d . For the following examples, we used the least-squares distance and computed solutions with IBM CPLEX².

Example 4.8. *Note that the constraints in Example 2.8 are non-strict epistemic atoms. The first row in Table 1 shows the beliefs in arguments for our initial epistemic state P_1 . Suppose that Terry can access the info app and learns that the train is indeed on time. Furthermore, he also strongly believes that the train is travelling at its usual speed. We can model this by an update with $\Psi_1 = \{p(D) = 1, p(C) = 0\}$. The second row in Table 1 shows the beliefs in arguments*

| P | $P(A)$ | $P(B)$ | $P(C)$ | $P(D)$ |
|------------------------------|--------|--------|--------|--------|
| P_1 | 0.45 | 0.65 | 0.2 | 0.5 |
| $P_2 = U_{d_2}(P_1, \Psi_1)$ | 0.67 | 0.65 | 0 | 1 |
| $P_3 = U_{d_2}(P_2, \Psi_2)$ | 0.175 | 0.65 | 1 | 0.5 |

Table 1: Returning to Example 4.8, beliefs in arguments before and after updating the epistemic state with new knowledge $\Psi_1 = \{p(D) = 1, p(C) = 0\}$ and $\Psi_2 = \{p(C) = 1\}$.

after updating P_1 to $P_2 = U_{d_2}(P_1, \Psi_1)$. Assume that a little bit later, the train has to slow down because of bad weather conditions. The third row in Table 1 shows the beliefs in arguments after updating P_2 with $\Psi_2 = \{p(C) = 1\}$ to $P_3 = U_{d_2}(P_2, \Psi_2)$. Note that the fact that the train slows down not only decreases Terry’s belief in A, but also indirectly leads to a decrease in D, which can be seen as Terry no longer being sure that the app showing the train on time indeed means the train will arrive on time.

Updating with Non-strict Epistemic Formulae

We now extend our fragment by allowing connecting epistemic atoms via logical conjunction and disjunction.

Definition 4.9 (Non-strict Epistemic Formulae). *The set of non-strict epistemic formulae is the closure of non-strict epistemic atoms under the logical connectives \wedge and \vee .*

Including conjunctions of non-strict epistemic atoms in our fragment does not cause any difficulties because adding a conjunction $\varphi_1 \wedge \varphi_2$ to our set of constraints \mathcal{C} is equivalent to adding both φ_1 and φ_2 . However, while conjunction extends our language syntactically, we do not gain anything semantically. Allowing disjunctions improves the expressiveness of our language greatly, but at a considerable cost. In general, distance-minimizing updates might not be well-defined anymore.

Example 4.10. *Consider an epistemic graph over a single argument A with the constraint $p(A) \leq 0.3 \vee p(A) \geq 0.7$ (the constraint says that the probability of A should be bounded away from 0.5 by at least 0.2). Let P_0 be the uniform distribution from $\text{Dist}(\mathcal{G})$. Then $U_{d_2}(P_0, \emptyset)$ is not well-defined because both the distribution P_1 with $P_1(A) = 0.3$ and P_2 with $P_2(A) = 0.7$ minimize the least-squares distance from P_0 among the distributions in $\text{Sat}(\{p(A) \leq 0.3 \vee p(A) \geq 0.7\})$.*

The general problem of disjunctions is that we may lose convexity of $\text{Sat}(\mathcal{C})$. This can cause the existence of multiple optimal solutions and can also cause the existence of non-global local minima which complicates the computational problem in practice. However, we can deal with this problem as we explain below.

Consider a set of formulae from the fragment of non-strict epistemic formulae. Two formulae φ_1, φ_2 in this set can be equivalently replaced with $\varphi_1 \wedge \varphi_2$. Hence, we can represent the complete set by a single formula φ . In particular, we can assume that φ is in disjunctive normal form (DNF)³.

²<https://www.ibm.com/analytics/data-science/prescriptive-analytics/cplex-optimizer>

³Normal forms of epistemic formulae are analogous to propo-

Example 4.11. Consider an epistemic graph over arguments A, B, C and suppose \mathcal{C} contains the formulae $p(A) \geq 0.6 \vee p(B) \geq 0.6$ and $p(B) \leq 0.4 \vee p(C) \leq 0.4$. We can replace both expressions equivalently by $(p(A) \geq 0.6 \vee p(B) \geq 0.6) \wedge (p(B) \leq 0.4 \vee p(C) \leq 0.4)$, which in turn is equivalent to the following expression in disjunctive normal form $(p(A) \geq 0.6 \wedge p(B) \leq 0.4) \vee (p(A) \geq 0.6 \wedge p(C) \leq 0.4) \vee (p(B) \geq 0.6 \wedge p(C) \leq 0.4)$.

We can find all distance-minimizing probability distributions by solving an independent convex optimization problem for every conjunction in the DNF. In particular, the number of minimal solutions can be bounded by the number of conjunctions in the DNF and is therefore always finite.

Proposition 4.12. Suppose that $\mathcal{C} \cup \Psi$ is from the fragment of non-strict epistemic formulae. Let $\bigvee_{i=1}^k \Gamma_i$ denote a DNF representation of $\mathcal{C} \cup \Psi$. Let P_i^* be the unique solution of $U_d(P, \Gamma_i)$ and let $m^* = \min\{d(P, P_i^*) \mid 1 \leq i \leq k\}$. Then $U_d(P, \Psi) = \{P_i^* \mid d(P, P_i^*) = m^*\}$ and $|U_d(P, \Psi)| \leq k$.

Of course, the number of conjunctions in the disjunctive normal form can be exponential in the size of \mathcal{C} . However, this is only the worst case and there are many interesting argumentation problems of moderate size. In summary, for the fragment of non-strict epistemic formulae, we have a general completeness, but not a general uniqueness guarantee.

Theorem 4.13. In the fragment of non-strict epistemic formulae, every distance-minimizing update functions satisfies Completeness. Uniqueness can be violated, but $|U_d(P, \Psi)| \leq k^*$, where k^* is the smallest number of conjunctions in all DNF representations of $\mathcal{C} \cup \Psi$.

Example 4.14. Consider again the initialization problem from Example 4.10. Γ_1 corresponds to $p(A) \leq 0.3$, Γ_2 to $p(A) \geq 0.7$. We have $P_1^*(A) = 0.3$ and $P_2^*(A) = 0.7$ as possible solutions.

If we have multiple solutions, we may define tie-breaking rules to select a preferred distribution, let the user decide which distribution to use or relax our uniqueness condition and work with multiple candidates at the same time.

Updating with General Epistemic Formulae

Let us now consider the full language of epistemic formulae. Negation and strict inequalities come at a price. Whereas we may lose the convexity of $\text{Sat}(\mathcal{C})$ due to disjunctions, we may lose closedness due to negations. This is again a problem for the well-definedness of distance-minimizing updates because an optimal solution might not exist anymore.

Example 4.15. Consider again an epistemic graph over a single argument A with the constraint $\neg(p(A) = 0.5)$ which is equivalent to $p(A) \neq 0.5$. Let P_0 be the uniform distribution from $\text{Dist}(\mathcal{G})$. Then there is no distribution in $\text{Sat}(\{p(A) \neq 0.5\})$ that minimizes the least-squares distance to P_0 because the distributions P_ϵ with $P_\epsilon(A) = 0.5 + \epsilon$ come arbitrarily close to P_0 as $\epsilon \rightarrow 0$.

sitional normal forms where epistemic atoms are treated as propositional atoms.

Let us first focus on computing solutions if they exist. A given set Γ of epistemic formulae can be expressed as a single formula representing a conjunction of all $\varphi \in \Gamma$. This formula can then be transformed to negation normal form in the usual manner, and each negative literal can be turned into a positive one by observing that negation changes only the (in-)equality relation. For instance, negating $=$ yields \neq and negating \leq yields $>$. We can thus assume that the formula contains again only disjunction and conjunction. The difference is that we can now have strict inequalities with $<$, \neq and $>$. Given that these relations may cost us the existence of update solutions, an interesting question is what do we lose when relaxing strict inequalities. We therefore propose the following definition.

Definition 4.16 (Relaxed DNF). Consider a general epistemic formula ϕ in DNF where all negations have been eliminated. The relaxed DNF of ϕ is obtained by first deleting all contradictory conjunctions in ϕ and then replacing all appearances of $<$ with \leq , $>$ with \geq and replacing all atoms containing \neq with the tautology $p(\top) = 1$ afterwards.

Example 4.17. Consider the general epistemic formula $(p(A) = 1 \rightarrow p(B) = 1) \wedge (p(B) = 1 \rightarrow p(C) = 0)$ and the corresponding DNF $\Gamma_1 \wedge \Gamma_2 \wedge \Gamma_3 \wedge \Gamma_4$ where $\Gamma_1 : p(A) \neq 1 \wedge p(B) \neq 1$, $\Gamma_2 : p(A) \neq 1 \wedge p(C) = 0$, $\Gamma_3 : p(B) = 1 \wedge p(B) \neq 1$, $\Gamma_4 : p(B) = 1 \wedge p(C) = 0$. Note that Γ_3 is contradictory. The corresponding relaxed DNF is $\Gamma'_1 \wedge \Gamma'_2 \wedge \Gamma'_4$ where $\Gamma'_1 : p(\top) = 1 \wedge p(B) \neq 1$, $\Gamma'_2 : p(\top) = 1 \wedge p(C) = 0$, $\Gamma'_4 : p(B) = 1 \wedge p(C) = 0$.

Notice that the relaxed disjunctive normal form is in the fragment of non-strict epistemic formulae. As the following proposition explains, the relaxed disjunctive normal form allows us to reduce the update problem for general epistemic formulae to the one for the fragment of non-strict epistemic formulae. The solutions of the update under the relaxed formula that satisfy the strict constraints are exactly the solutions of the update under the strict formula. In particular, if none of the solutions of the update under the relaxed formula satisfy the strict constraints, the strict update is not well-defined because the minimum does not exist.

Proposition 4.18. Let $\mathcal{C} \cup \Psi$ contain arbitrary epistemic formulae. Let $\bigvee_{i=1}^k \Gamma_i$ denote a DNF of $\mathcal{C} \cup \Psi$ where negations have been eliminated. Let $\bigvee_{i=1}^{k'} \Gamma'_i$ denote the corresponding relaxed DNF. For $i = 1, \dots, k'$, let P_i^* be the unique solution of the optimization problem

$$\min_{P' \in \text{Sat}(\Gamma'_i)} d(P, P')$$

corresponding to the i -th conjunction Γ'_i in the relaxed DNF. Let $m^* = \min\{d(P, P_i^*) \mid 1 \leq i \leq k'\}$ be the minimum distance obtained among all P_i^* . Then $U_d(P, \Psi)$ equals

$$\{P_i^* \mid 1 \leq i \leq k', d(P, P_i^*) = m^* \text{ and } P_i^* \in \text{Sat}(\Gamma_i)\}.$$

In general, we have neither completeness nor uniqueness guarantees for the full language of epistemic formulae. However, as Proposition 4.18 explains, the update problem for the full language can be reduced to the update problem for the fragment of non-strict epistemic formulae: given

a general epistemic formula, we compute the corresponding disjunctive normal form, eliminate negation by changing (in-)equalities appropriately and then consider the corresponding relaxed disjunctive normal form. Proposition 4.12 explains how to compute the solutions and guarantees that there is only a finite number of them. We then restrict the solutions to those that satisfy the strict constraints and obtain the solutions (possibly none) for the original update as Proposition 4.18 explains. The following two simple examples show cases in which a solution exists and does not exist.

Example 4.19. Consider an epistemic graph over two arguments A, B . Let $C = \{p(A) = 1 \rightarrow p(B) = 0\}$ and let P_0 be the uniform distribution from $\text{Dist}(\mathcal{G})$. Consider the update $\Psi = \{p(A) = 1\}$. We have $\Gamma_1 : p(A) \neq 1 \wedge p(A) = 1$ and $\Gamma_2 : p(B) = 0 \wedge p(A) = 1$. Since Γ_1 is contradictory, the relaxed DNF is just $\Gamma'_2 = \Gamma_2$. For the corresponding optimal solution P_2^* , we have $P_2^*(A) = 1$ and $P_2^*(B) = 0$ as desired. Notice that deleting Γ_1 is crucial because relaxing the contradictory formula Γ_1 would yield the satisfiable formula $\Gamma'_1 : p(\top) = 1 \wedge p(A) = 1$.

Example 4.20. Consider again the update from Example 4.15. Γ_1 corresponds to $p(A) \neq 0.5$ and Γ'_1 is the tautology $p(\top) = 1$. We have $P_1^*(A) = 0.5$ and P_1^* violates the strict constraint, so we know that no solution exists.

Note that if there is no solution for the strict DNF, but the relaxed DNF has a solution, then there must be a P_i^* with $d(P, P_i^*) = m^*$ and $P_i^* \notin \text{Sat}(\Gamma_i)$. Hence, a strict inequality in Γ_i causes a non-existence problem. In this way, we can identify problematic strict inequalities and suggest to the user to replace strict inequalities temporarily with slightly relaxed inequalities. In our example above, we may replace $P_1^*(A) = 0.5$ with the relaxed constraint $p(A) \leq 0.49 \vee p(A) \geq 0.51$. In this way, we can circumvent the non-existence problem in practice.

5 Atomic Distance Minimization

While a distribution updated with the previously discussed methods will meet the epistemic constraints, minimizing the least squares distance does not necessarily change the probability of arguments in an intuitive manner. Consider the following example:

Example 5.1. Let us come back to Example 4.8 and the graph from Figure 1. We consider the following constraints describing Terry's current impressions of the trip $\{p(A) + \frac{1}{2}p(B) + \frac{1}{2}p(C) \leq 1, p(A) + \frac{1}{2}p(B) + \frac{1}{2}p(C) - p(D) \geq 0, p(B) \geq 0.65\}$ and obtain the probability distribution described in the first row of Table 1. Suppose another traveller tells Terry that the info app states that the train is indeed on time and that Terry trusts his statement. We describe this with $\Psi_1 = \{p(D) = 1\}$. By performing a least-squares update, we obtain a distribution P_2 s.t. $P_2(A) = 0.5, P_2(B) = 0.65, P_2(C) = 0.36$ and $P_2(D) = 1$. Although this distribution satisfies our constraints and minimizes the overall change in probability mass, the increase in belief in C might not be considered justified.

The reason for the behaviour observed in the example is that we change a probability distribution over sets of arguments, not over atomic arguments. Probability distributions

are defined in this way because probabilistic reasoning is not truth functional. That is, the probability of complex formulae cannot be computed from the probabilities of atomic formulae without making further assumptions. Thus, it is not sufficient to define probabilities only for atomic arguments.

However, since humans tend to think in terms of probabilities of atomic arguments, a natural idea to make updates more intuitive is to measure the difference in probability mass assigned to arguments. Such distance measures have been called atomic distances in (Hunter and Potyka 2017). Unfortunately, they satisfy neither positive definiteness nor strict convexity and, thus, are not epistemic distance functions. The main problem is that many different distributions can assign the same beliefs to arguments. However, we can use such a measure as a preprocessing step. We first restrict the models in $\text{Sat}(C \cup \Psi)$ to those that leave the probability mass of arguments as close as possible to the prior state. Among those distributions that are left, we then minimize the distance to the prior state with respect to an epistemic distance function. In order to do this, we consider a weighted atomic distance measure.

Definition 5.2 (Weighted Atomic Distance). Let $S \subseteq \text{Nodes}(\mathcal{G})$ be a set of arguments and let $w : S \rightarrow \mathbb{R}_0^+$ be a weight function assigning a non-negative weight to each argument in S . The **weighted atomic distance** with respect to w is defined as $d_{\text{At}}^w(P, P') = \sum_{A \in S} w(A) \cdot |P(A) - P'(A)|$.

The weight function allows us to control which atomic beliefs should change first. For instance, when updating arguments in a set S , the weight function could assign to each argument the minimal distance in the graph from arguments in S . As a result, the probabilities of arguments close to S will change first. We will discuss such approaches in more detail in future work. Let us note that the weighted atomic distance still has some useful properties.

Lemma 5.3. d_{At}^w is a continuous and convex pseudometric.

As we explained in the previous sections, we can break down all update problems to ones that contain only a disjunction of conjunctions of non-strict atoms. We can then solve independent optimization problems for each conjunction to address our original problem. The following lemma explains that restricting $\text{Sat}(C \cup \Psi)$ to those distributions that minimize a weighted atomic distance measure d_{At}^w produces a set that still maintains useful properties. In particular, we can identify the restricted set through linear programming techniques.

Lemma 5.4. If Γ is a conjunction of non-strict epistemic atoms, then the solutions of

$$\arg \min_{P' \in \text{Sat}(\Gamma)} d_{\text{At}}^w(P, P')$$

correspond to the solution of the linear program

$$\begin{aligned} \min \quad & \sum_{A \in S} (\delta_A^+ + \delta_A^-) \\ \text{s.t.} \quad & P' \in \text{Sat}(\Gamma) \\ & w(A) \cdot (P(A) - P'(A)) = \delta_A^+ - \delta_A^- \quad \text{for all } A \in S \\ & \delta_A^+, \delta_A^- \in \mathbb{Q}_0^+ \quad \text{for all } A \in S \end{aligned} \tag{1}$$

and form a compact and convex set.

| P | P(A) | P(B) | P(C) | P(D) |
|--------------------------------|-------|------|------|------|
| P_1 | 0.45 | 0.65 | 0.2 | 0.5 |
| $P_2 = U_{d_2}^w(P_1, \Psi_1)$ | 0.575 | 0.65 | 0.2 | 1 |
| $P_3 = U_{d_2}^w(P_2, \Psi_2)$ | 0.175 | 0.65 | 1 | 1 |

Table 2: Beliefs in arguments before and after updating the epistemic state with new knowledge $\Psi_1 = \{p(D) = 1\}$ and $\Psi_2 = \{p(C) = 1\}$ in Example 5.7.

We can now minimize an epistemic distance function exactly as before with the only difference that we restrict ourselves to those satisfying distributions that minimize a weighted atomic distance. Thus, what we consider is a two-stage procedure. In the first stage, we minimize the atomic distance to the prior distribution to restrict to those probability distributions that do not change the beliefs in arguments more than necessary. This stage usually produces an infinite number of candidates with equal atomic distance. Then, in the second stage, we pick from these candidates the one(s) that minimize an epistemic distance to the prior distribution.

Definition 5.5 (Atomic Distance-minimizing Update Function). *Given some epistemic distance function d and some weight function $w : S \rightarrow \mathbb{R}_0^+$ over a subset of arguments $S \subseteq \text{Nodes}(\mathcal{G})$, the atomic distance-minimizing update function $U_d^w(P, \Psi)$ with respect to d is defined by the set of minimal solutions of the optimization problem*

$$\begin{aligned}
\min \quad & d(P, P') \\
\text{s.t.} \quad & P' \in \text{Sat}(\mathcal{C} \cup \Psi) \\
& \sum_{A \in S} (\delta_A^+ + \delta_A^-) = m^* \\
& w(A) \cdot (P(A) - P'(A)) = \delta_A^+ - \delta_A^- \quad \text{for all } A \in S \\
& \delta_A^+, \delta_A^- \in \mathbb{Q}_0^+ \quad \text{for all } A \in S,
\end{aligned}$$

where m^* is the minimum of all minima of (1) computed for all conjunctions of a relaxed DNF of $\mathcal{C} \cup \Psi$.

Since restricting the feasible region in the first stage maintains all useful properties of $\text{Sat}(\mathcal{C} \cup \Psi)$, the following results can be proved completely analogously as before.

Theorem 5.6. *Every atomic distance-minimizing update functions satisfies Epistemic Consistency, Success, Tautology, Contradiction, Representation Invariance and Idempotence. In the fragment of non-strict epistemic formulae, Completeness is satisfied as well and $U_d^w(P, \Psi)$ is guaranteed to be finite. In the fragment of non-strict epistemic atoms, Uniqueness is also satisfied.*

Example 5.7. *Let us come back to Example 5.1 and this time carry out the updates using the atomic distance-minimizing update approach s.t. for every argument X , $w(X) = 1$ (i.e. all atomic distances are treated equally). By performing an update with the formula $\Psi_1 = \{p(D) = 1\}$, we obtain the probability function P_2 visible in the second row of Table 2. The new probability assignments to arguments are now more intuitive than in Example 5.1.*

6 Updates in Persuasion Dialogues

In this section we show how the methods we have discussed can be harnessed in dialogical argumentation with epistemic graphs. We focus on asymmetric dialogues, but note that our approach can be adapted to further methods as well.

An *argument dialogue* can be seen as a sequence of moves $D = [m_1, \dots, m_k]$. Possible *moves* include positing arguments, making claims, conceding a claim, providing premises for an argument and more. We can differentiate between *symmetric* and *asymmetric dialogues*. In the latter, certain moves may only be available to some of the dialogue participants, or the moves made by one participant can be restricted by another. In contrast, the former allow all participants equal freedom (Prakken 2006; Hunter 2015). A *dialogue protocol* states the rules of the dialogue, such as what moves can follow other moves, and can include requirements such as a persuasion goal being the first posited argument and more.

Asymmetric dialogues can be harnessed by automated persuasion systems so that they do not need to be equipped with natural language processing capabilities. The system (*persuader*) can posit arguments (denoted $A!$), request reasons for disagreeing with an argument (denoted $A^-?$), request reasons for agreeing with an argument (denoted $A^{+?}$), and end the dialogue (denoted \perp_S). The system only terminates the dialogue when there are no arguments left to consider. The user (*persuadee*) reacts to a posit with a statement of belief (denoted $A : x$, where $x \in [0, 1]$) and to reasons for agreeing and disagreeing with statements of beliefs to the attackers/supporters of a given argument that are supplied by the system (denoted $[A_1 : x_1, \dots, A_n : x_n]$, where $x_i \in [0, 1]$ and arguments A_i are made available by the system and depend on the request). Finally, the user is allowed to terminate the dialogue at any point in time (denoted \perp_U). We assume that the dialogue participants take turns when performing moves, i.e. if m_i is a move belonging to the system, then m_{i+1} is associated with the user.

Example 6.1. *Consider the graph in Figure 2 and imagine the following exchange. The system (S) puts forward A , to which the user named Morgan (M) disagrees and assigns a belief of 0.125. S asks for the reasons for this and presents B , C and D as possible options. M agrees with B and D (beliefs 0.875) and strongly disagrees with C (belief 0). In response to that, S first posits E , to which M agrees with belief 0.75, and then posits F , to which M somewhat disagrees (belief 0.375). S follows up on that with G , to which M once more disagrees with belief 0.125. This dialogue can be encoded as a sequence $[A!; A : 0.125; A^-?; [B : 0.875, C : 0, D : 0.875], E!; E : 0.875; F!; F : 0.375; G!; G : 0.125; \perp_S]$.*

In epistemic approaches to persuasion, the model of the user is typically understood as a belief distribution over the graph that the dialogue is based on. It is meant to reflect the opinions that the user has about the arguments. As the dialogue progresses, the beliefs of the user can change, and therefore so should the belief distribution (see Figure 3).

Given the asymmetric setup of our dialogue, it is sufficient for us to consider only the updates caused by the moves belonging to the user. In our setup, every user move

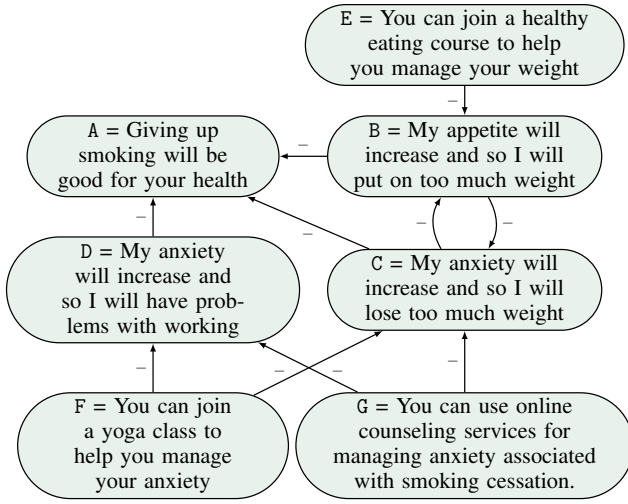


Figure 2: Example of argument graph for persuading someone to give up smoking. Edges labelled – represent attack.




Figure 3: Schematic representation of a dialogue $D = [m_1, \dots, m_n]$ and user models P_i . Each user model P_i is obtained from P_{i-1} and move m_i using an update method.

can be mapped to a non-strict epistemic formula. For example, the belief statement $A : x$ for $x \in [0, 1]$ becomes $p(A) = x$, and a series of statements $[A_1 : x_1, \dots, A_n : x_n]$, where $x_i \in [0, 1]$, is associated with the formula $p(A_1) = x_1 \wedge \dots \wedge p(A_n) = x_n$. Let us consider how the exchange from Example 6.1 can be used to update the user model:

Example 6.2. *Let us continue Example 6.1 and assume that the dialogue was preceded by a profiling stage, where the constraints describing Morgan’s opinions on the interactions of arguments are obtained from the constraints crowdsourced from other participant’s with similar profiles. Assume we also deduced that Morgan is probably more afraid of weight gain than weight loss or anxiety affecting his work, independently of his willingness to discuss these topics. We model our assumptions with an epistemic graph that contains the following set of constraints:*

- $\psi_1 : p(A) = 1 - \frac{3}{5}p(B) - \frac{1}{5}p(C) - \frac{1}{5}p(D)$
- $\psi_2 : p(B) \leq 1 - p(C) \wedge p(B) \leq 1 - p(E)$
- $\psi_3 : p(C) \leq 1 - p(F) \wedge p(C) \leq 1 - p(G)$
- $\psi_4 : p(D) \leq 1 - p(F) \wedge p(D) \leq 1 - p(G)$

Our initial belief distribution is the uniform distribution P_1 that satisfies all constraints. Recall that the dialogue was encoded as sequence $[A!; A : 0.125; A^-?; [B : 0.875, C : 0, D : 0.875], E!; E : 0.875; F!; F : 0.375; G!; G : 0.125; \perp_S]$. We thus create the following formulae corresponding to

| P | A | B | C | D | E | F | G |
|---------------------------------------|-------|-------|-----|-------|------|-------|-------|
| P_1 | 0.5 | 0.5 | 0.5 | 0.5 | 0.5 | 0.5 | 0.5 |
| $P_2 = U_{d_2}^w(P_1, \{\varphi_1\})$ | 0.25 | 1 | 0 | 0.750 | 0 | 0.25 | 0.25 |
| $P_3 = U_{d_2}^w(P_2, \{\varphi_2\})$ | 0.3 | 0.875 | 0 | 0.875 | 0 | 0.125 | 0.125 |
| $P_4 = U_{d_2}^w(P_3, \{\varphi_3\})$ | 0.675 | 0.25 | 0 | 0.875 | 0.75 | 0.125 | 0.125 |
| $P_5 = U_{d_2}^w(P_4, \{\varphi_4\})$ | 0.725 | 0.25 | 0 | 0.625 | 0.75 | 0.375 | 0.125 |
| $P_6 = U_{d_2}^w(P_5, \{\varphi_5\})$ | 0.725 | 0.25 | 0 | 0.625 | 0.75 | 0.375 | 0.125 |

Table 3: Updates to the belief distribution describing Morgan’s belief throughout the dialogue from Example 6.2.

Morgan’s moves: $\{\varphi_1 : p(A) = 0.125, \varphi_2 : p(B) = 0.875 \wedge p(C) = 0 \wedge p(D) = 0.875, \varphi_3 : p(E) = 0.875, \varphi_4 : p(F) = 0.375, \varphi_5 : p(G) = 0.125\}$. In order to compute updates, we use the atomic distance-minimizing update function $U_{d_2}^w$ with uniform weights (all weights are 1) and least-squares distance minimization. The results of updating Morgan’s user model throughout the dialogue are visible in Table 3. We observe that positing argument E resolves Morgan’s concerns about B. However, while the system manages to cast some doubt concerning D, it does not address the issue completely. Nevertheless, given the fact that Morgan’s more pressing problem was satisfactorily discussed, the dialogue ends with A being believed.

Updating the belief distribution is an important, but not the only component that needs to be considered in automated persuasion systems. Another crucial aspect is the update of the constraints describing the user’s reasoning patterns. Although initial constraints can be obtained through, for example, crowdsourcing opinions from participant’s that have profiled similarly according to given criteria, they may need to be updated during the discussion. This may be a result of the user changing his or her opinions concerning the relations between arguments as well as the system assigning inappropriate initial constraints. If, for instance, in the analyzed examples, it has turned out that Morgan would prefer to discuss B and C rather than B and D, the formula generated for this move would lead to an inconsistency with the constraint ψ_2 . This could possibly be addressed by not including the constraint on the B and C dependency. Nevertheless, learning and updating user constraints is a deeper problem that requires a separate analysis and is left for future work.

7 Related Work

Epistemic graphs are a generalization of epistemic probabilistic argumentation to a setting with more advanced relations between arguments. In (Hunter, Polberg, and Thimm 2018) it was shown how the epistemic postulates (Thimm 2012; Hunter 2013; Hunter and Thimm 2014; Polberg and Hunter 2018) and abstract dialectical frameworks (Brewka et al. 2013; Linsbichler et al. 2018), which themselves generalize a wide range of existing argumentation formalisms (Polberg 2016), can be expressed within epistemic graphs. The ability to represent constraints that are not limited to arguments that are directly connected in the graph also allows epistemic graphs to handle constrained argumentation frameworks (Coste-Marquis, Devred, and Marquis 2006).

Given the fine-grained nature of the epistemic approach, it is natural to compare our proposal to the graded and

ranking-based semantics proposed for a number of argumentation frameworks (Cayrol and Lagasque-Schiex 2005; Leite and Martins 2011; Rago et al. 2016; Bonzon et al. 2016; Amgoud and Ben-Naim 2017; Brewka et al. 2018). Although in most of these approaches what the semantics produce can be seen as “assigning numbers from $[0, 1]$ ” to arguments (either as a side or end product), probabilities in the epistemic approach are interpreted as belief, while in the remaining works they are typically left abstract. Thus, many of the postulates set out in the aforementioned methods are, by design, counter-intuitive in the epistemic approach, even though they can be perfectly applicable in other scenarios. Furthermore, with the exception of (Brewka et al. 2018), the patterns set out by the graded and ranking-based semantics have to be global, while in our case we can choose to define the way parents of an argument affect it differently for every argument. A more in-depth analysis can be found in (Hunter, Polberg, and Thimm 2018).

The epistemic approach is not the only form of probabilistic argumentation. One can also name the constellation approach (Li, Oren, and Norman 2011; Hunter 2013), in which we consider a probability distribution over subgraphs of a given graph. The probability of each subgraph is interpreted as its chances of being the “real graph”, which is quite distinct from the belief interpretation of the epistemic approach. Hence, despite the fact that both formalisms focus on probabilities, there are significant differences between how they model and use them. Further analysis can be found in (Hunter 2013; Polberg, Hunter, and Thimm 2017).

Applying the standard epistemic approach to modelling persuadee’s beliefs in arguments has produced methods for updating beliefs during a dialogue (Hunter 2015; Hunter and Potyka 2017). However, these methods are not equipped to handle epistemic graphs, and, in particular, do not consider positive relations between arguments. Our current work can be seen as a successor to the previous approaches to a more general setting inspired by the empirical studies we have carried out in (Polberg and Hunter 2018; Hunter and Polberg 2017).

Our update postulates have some resemblance to postulates considered for belief change (Gärdenfors 1988; Darwiche and Pearl 1997; Kern-Isberner 2001). For lack of space, we must omit a detailed discussion here. Updating probabilistic belief states has been considered in many related areas (Chan and Darwiche 2005; Beierle and Kern-Isberner 2009; Rens and Meyer 2015). However, constraints in these areas are usually restricted to linear constraints over the probability of formulae. Since we allow connecting such constraints with logical connectives, our setting is more general and special care needs to be taken in order to handle disjunctions and negation.

8 Conclusions and Future Work

In this paper we addressed the problem of updating belief distributions of epistemic graphs with new constraints. We presented a number of intuitive properties that a reasonable update method should satisfy together with several distance minimizing update approaches satisfying them. Finally, we

demonstrated the usefulness of our proposal by considering its application in persuasion with asymmetric dialogues.

Computationally, all of the involved optimization problems are convex and can be solved in polynomial time w.r.t. the number of probabilities in question. However, with belief distributions, the number of probabilities that need to be considered grows exponentially with the number of arguments. Still, a simple proof-of-concept implementation with IBM CPLEX, which minimizes least-squares distance naively, shows that problems with up to 15 arguments can be solved in a few seconds. The runtime then increases rapidly as the number of arguments goes up. Problems with up to 20 arguments can still be solved in under one minute, but in order to scale-up much further, we need to apply more sophisticated ideas. In the future, we can consider exploiting conditional independencies as explained in the theory of Markov random fields (Koller and Friedman 2009; Potyka, Beierle, and Kern-Isberner 2015; Wilhelm et al. 2017) or apply ideas such as column generation (Hansen and Perron 2008; Finger and De Bona 2011; Cozman and di Ianni 2013) in order to scale up.

There are many interesting questions that we would like to investigate in future work. In particular, we intend to analyze different weight functions for atomic distance-minimizing update functions. For example, we can use the weights derived from the distance in the graph between a given argument and the arguments contained in the constraint used to perform the update, where the distance can be computed respecting or ignoring the directions of edges in order to achieve different behaviour. For instance, we could model a water-like ripple effect, in which a bigger change in a closer argument may be preferred to a smaller change in an argument further away. Another important topic are methods for dealing with non-uniqueness and non-existence of update solutions. We will investigate different approaches how to select the best distribution from a finite set of candidates. Furthermore, we will elaborate on how to identify those strict inequalities that cause non-existence problems and how to transform them in the most innocuous way.

In the future we plan to incorporate priorities in our framework in order to resolve conflicting information. For instance, we could assign higher priority to user updates than to constraints in the epistemic graph. Another application is to assign higher priority to particular constraints in the epistemic graph or in the update individually. This allows, in particular, expressing specialized beliefs like general beliefs about traveling by plane (low priority) and more specialized beliefs about traveling with particular airlines (higher priority). Roughly speaking, we can handle these priorities by solving a sequence of optimization problems that first satisfy high priority constraints and later minimize the violation of low priority constraints (Potyka 2015). These methods, combined with constraint learning and constraint updating investigations, would lead to further developments in handling user models in automated persuasion systems.

References

Amgoud, L., and Ben-Naim, J. 2017. Evaluation of arguments in weighted bipolar graphs. In *Proc. of ECSQARU’17*,

- volume 10369 of *LNCS*, 25–35. Springer.
- Beierle, C., and Kern-Isberner, G. 2009. A conceptual agent model based on a uniform approach to various belief operations. In *Proc. of KI'09*, volume 5803 of *LNCS*, 273–280. Springer.
- Bonzon, E.; Delobelle, J.; Konieczny, S.; and Maudet, N. 2016. A comparative study of ranking-based semantics for abstract argumentation. In *Proc. of AAI'16*, 914–920. AAAI Press.
- Brewka, G.; Ellmauthaler, S.; Strass, H.; Wallner, J. P.; and Woltran, S. 2013. Abstract dialectical frameworks revisited. In *Proc. of IJCAI'13*, 803–809. AAAI Press.
- Brewka, G.; Strass, H.; Wallner, J. P.; and Woltran, S. 2018. Weighted abstract dialectical frameworks. In *Proc. of AAI'18*, 1779–1786. AAAI Press.
- Cayrol, C., and Lagasque-Schiex, M. 2005. Graduality in argumentation. *Journal of Artificial Intelligence Research* 23:245–297.
- Chan, H., and Darwiche, A. 2005. A distance measure for bounding probabilistic belief change. *International Journal of Approximate Reasoning* 38(2):149–174.
- Coste-Marquis, S.; Devred, C.; and Marquis, P. 2006. Constrained argumentation frameworks. In *Proc. of KR'06*, 112–122. AAAI Press.
- Cozman, F. G., and di Ianni, L. F. 2013. Probabilistic satisfiability and coherence checking through integer programming. In *Proc. of ECSQARU'13*, volume 7958 of *LNCS*, 145–156. Springer.
- Csiszar, I. 1975. I-Divergence Geometry of Probability Distributions and Minimization Problems. *The Annals of Probability* 3(1):146–158.
- Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artificial intelligence* 89(1-2):1–29.
- Finger, M., and De Bona, G. 2011. Probabilistic satisfiability: Logic-based algorithms and phase transition. In *Proc. of IJCAI'11*, 528–533. AAAI Press.
- Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books. MIT Press.
- Hansen, P., and Perron, S. 2008. Merging the local and global approaches to probabilistic satisfiability. *International Journal of Approximate Reasoning* 47(2):125 – 140.
- Hunter, A., and Polberg, S. 2017. Empirical methods for modelling persuadees in dialogical argumentation. In *Proc. of ICTAI'17*. IEEE. In Press.
- Hunter, A., and Potyka, N. 2017. Updating probabilistic epistemic states in persuasion dialogues. In *Proc. of ECSQARU'17*, volume 10369 of *LNCS*, 46–56. Springer.
- Hunter, A., and Thimm, M. 2014. Probabilistic argumentation with incomplete information. In *Proc. of ECAI'14*, volume 263 of *FAIA*, 1033–1034. IOS Press.
- Hunter, A.; Polberg, S.; and Thimm, M. 2018. Epistemic graphs for representing and reasoning with positive and negative influences of arguments. *ArXiv*.
- Hunter, A. 2013. A probabilistic approach to modelling uncertain logical arguments. *International Journal of Approximate Reasoning* 54(1):47–81.
- Hunter, A. 2015. Modelling the persuadee in asymmetric argumentation dialogues for persuasion. In *Proc. of IJCAI'15*, 3055–3061. AAAI Press.
- Kern-Isberner, G. 2001. *Conditionals in nonmonotonic reasoning and belief revision: considering conditionals as agents*, volume 2087 of *LNCS*. Springer.
- Koller, D., and Friedman, N. 2009. *Probabilistic graphical models: principles and techniques*. MIT press.
- Leite, J., and Martins, J. 2011. Social abstract argumentation. In *Proc. of IJCAI'11*, 2287–2292. AAAI Press.
- Li, H.; Oren, N.; and Norman, T. J. 2011. Probabilistic argumentation frameworks. In *Proc. of TFA'11*, volume 7132 of *LNCS*, 1–16. Springer.
- Linsbichler, T.; Maratea, M.; Niskanen, A.; Wallner, J.; and Woltran, S. 2018. Novel algorithms for abstract dialectical frameworks based on complexity analysis of subclasses and sat solving. In *Proc. of IJCAI'18*, 1905–1911. AAAI Press.
- Nocedal, J., and Wright, S. J. 2006. *Numerical Optimization*. New York: Springer, 2nd edition.
- 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.
- Polberg, S.; Hunter, A.; and Thimm, M. 2017. Belief in attacks in epistemic probabilistic argumentation. In *Proc. of SUM'17*, volume 10564 of *LNCS*, 223–236. Springer.
- Polberg, S. 2016. Understanding the Abstract Dialectical Framework. In *Proc. of JELIA'16*, volume 10021 of *LNCS*, 430–446. Springer.
- Potyka, N.; Beierle, C.; and Kern-Isberner, G. 2015. A concept for the evolution of relational probabilistic belief states and the computation of their changes under optimum entropy semantics. *Journal of Applied Logic* 13(4):414–440.
- Potyka, N. 2015. Reasoning over linear probabilistic knowledge bases with priorities. In *Proc. of SUM'17*, volume 10564 of *LNCS*, 121–136. Springer.
- Prakken, H. 2006. Formal systems for persuasion dialogue. *The Knowledge Engineering Review* 21(2):163188.
- Rago, A.; Toni, F.; Aurisicchio, M.; and Baroni, P. 2016. Discontinuity-free decision support with quantitative argumentation debates. In *Proc. of KR'16*, 63–73. AAAI Press.
- Rens, G. B., and Meyer, T. 2015. A new approach to probabilistic belief change. In *Proc. of FLAIRS'15*, 581–587. AAAI Press.
- Rudin, W. 1976. *Principles of mathematical analysis*, volume 3. McGraw-hill New York.
- Thimm, M. 2012. A probabilistic semantics for abstract argumentation. In *Proc. of ECAI'12*, volume 242 of *FAIA*, 750–755. IOS Press.
- Wilhelm, M.; Finthammer, M.; Kern-Isberner, G.; and Beierle, C. 2017. First-order typed model counting for probabilistic conditional reasoning at maximum entropy. In *Proc. of SUM'17*, volume 10564 of *LNCS*. Springer.