Polynomial-time Updates of Epistemic States in a Fragment of Probabilistic Epistemic Argumentation

Nico Potyka[1], Sylwia Polberg[2], and Anthony Hunter[3]

[1] Institute of Cognitive Science, University of Osnabrück, Germany,
[2] School of Computer Science and Informatics, Cardiff University, UK
[3] Department of Computer Science, University College London, UK

Abstract. Probabilistic epistemic argumentation allows for reasoning about argumentation problems in a way that is well founded by probability theory. Epistemic states are represented by probability functions over possible worlds and can be adjusted to new beliefs using update operators. While the use of probability functions puts this approach on a solid foundational basis, it also causes computational challenges as the amount of data to process depends exponentially on the number of arguments. This leads to bottlenecks in applications such as modelling opponent's beliefs for persuasion dialogues. We show how update operators over probability functions can be related to update operators over much more compact representations that allow polynomial-time updates. We discuss the cognitive and probabilistic-logical plausibility of this approach and demonstrate its applicability in computational persuasion.

1 Introduction

Probabilistic epistemic argumentation [40, 13, 20, 18] is an extension of Dung's classical argumentation framework [7]. While the original framework allows only for talking about attacks and accepting or rejecting arguments, probabilistic epistemic argumentation also allows more general relationships between arguments like support [4, 29, 5] and allows expressing more fine-grained beliefs by means of probabilities. Recent experiments give empirical evidence that these extensions are, in particular, beneficial when it comes to modelling human decision making [28]. One large application area of probabilistic epistemic argumentation is computational persuasion [15, 16]. Computational persuasion aims at convincing the user of a persuasion goal such as giving up bad habits or living a healthier lifestyle. In order to derive persuasion strategies autonomously, we require a user model that represents the user's beliefs and simulates belief changes when new arguments are presented to the user. The user's epistemic state can be represented by a probability function and different update operators have been studied that can be used to adapt the current beliefs [15, 19, 17].

Probability theory provides a strong foundational basis for probabilistic epistemic argumentation, but also comes with computational limitations. Without further assumptions, probability functions grow exponentially with the number of arguments. However, sometimes we are only interested in atomic beliefs in arguments, so that the full power of probability functions may not be required. For instance, we can consider the
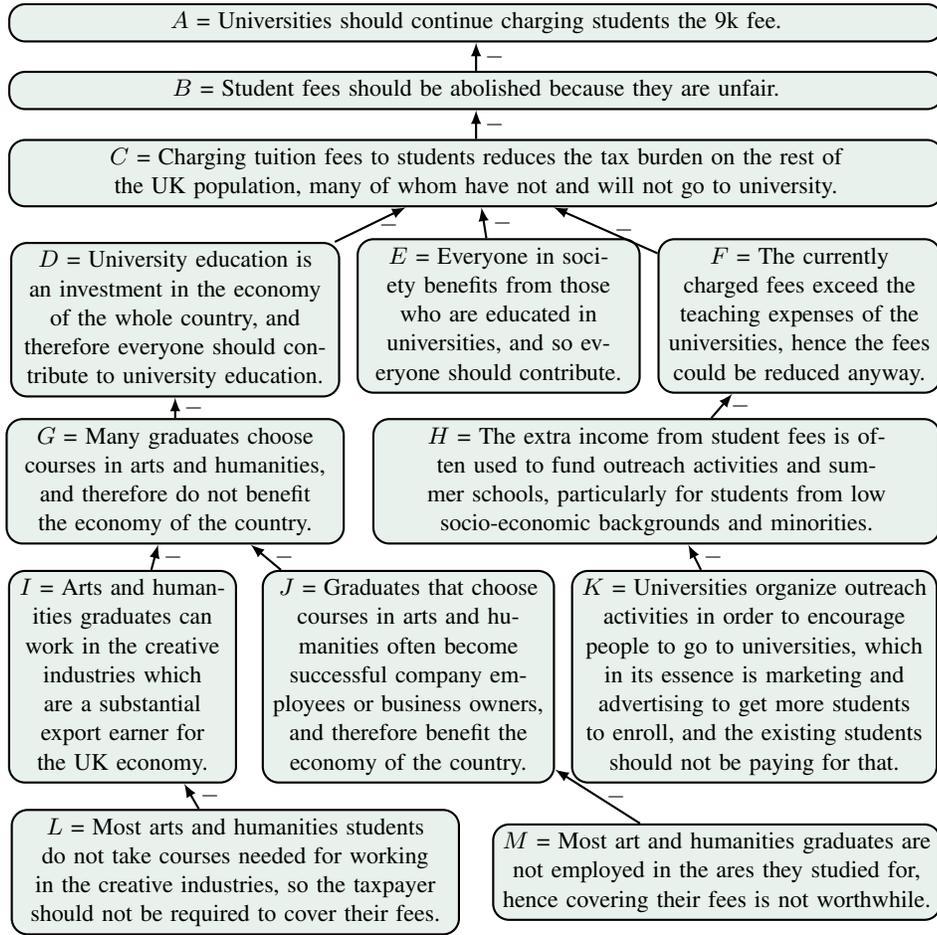
Fig. 1. Study fee dialogue.

graph depicted in Figure 1 induced by a dialogue between an automated dialogue system and a human participant that occurred in the empirical study considered in [11]. There are various constraints that could be attached to such a graph, as we will discuss further in Section 5. For instance, we could use postulates from the classical epistemic approach [40, 13] such as *coherence*, which bounds the belief in an argument based on the belief of its attacker. Formally, in our scenario, for every argument-attacker pair X and Y this would create a constraint of the form $\pi(X) + \pi(Y) \leq 1$, where $\pi(\alpha)$ should be read as the probability of α.

We observe that the aforementioned formulas operate on probabilities of single arguments rather than on complex logical expressions. Consequently, the detailed information contained in a full probability function can be seen as excessive. In such a situation, probability functions can sometimes be replaced by probability labellings that

assign probabilities to arguments directly without changing the semantics [33]. In our case, this would decrease the number of probabilities that need to be processed from 8,192 (i.e. 2^{13}) to 13, which has obvious computational benefits.

In this paper, we are interested in the relationship between epistemic states represented by probability functions and those represented by probability labellings. Formally, probability labellings can be related to equivalence classes of probability functions that assign the same atomic beliefs to arguments [33]. In order to establish an interesting relationship, update operators must respect this equivalence relation. We define such an operator in Section 3 and show in Section 4 that it satisfies our desiderata. In particular, updates can be computed in polynomial time In this approach, epistemic states correspond to sets of probability functions that satisfy the same atomic beliefs and updates are performed by satisfying the new beliefs while minimizing the required changes. We will argue that this approach is not only computationally attractive, but can also result in cognitively more plausible updates. We illustrate our method with an application in computational persuasion in Section 5. All proofs for the results in this article can be found in the corresponding technical report [35].

2   Basics

We consider *bipolar argumentation frameworks (BAFs)* $(\mathcal{A}, \mathcal{R}, \mathcal{S})$ consisting of a set of arguments $\mathcal{A}$, an attack relation $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ and a support relation $\mathcal{S} \subseteq \mathcal{A} \times \mathcal{A}$. $\Omega = \{w \mid w \subseteq \mathcal{A}\}$ denotes the set of *possible worlds*. Intuitively, each $w \in \Omega$ contains the arguments that are accepted in a particular state of the world. We represent beliefs by probability functions $P : \Omega \to [0, 1]$ such that $\sum_{w \in \Omega} P(w) = 1$. $\mathcal{P}_{\mathcal{A}}$ denotes the set of all probability functions over $\mathcal{A}$. The probability of an argument $A \in \mathcal{A}$ under P is defined by adding the probabilities of all worlds in which A is accepted, that is, $P(A) = \sum_{w \in \Omega, A \in w} P(w)$. $P(A)$ can be understood as a degree of belief, where $P(A) = 1$ means complete acceptance and $P(A) = 0$ means complete rejection[4].

The epistemic probabilistic argumentation approach developed in [40, 13, 20, 18] defines semantics of attack and support relations by means of constraints over probability functions. Some constraints can be automatically derived from the relations between arguments. For example, the *coherence constraint* demands that if A attacks B, we must have $P(B) \leq 1 - P(A)$, that is, the belief in an attacked argument B is bounded from above by the belief in an attacker A. However, it is also possible to design individual constraints manually. For example, if B is attacked by three related arguments A_1, A_2, A_3, we may want to bound the belief in B by the average belief in these attackers via $P(B) \leq 1 - \frac{1}{3} \sum_{i=1}^{3} P(A_i)$. To allow this flexibility, a general constraint language has been introduced in [18, 17]. We will focus on the fragment of *linear atomic constraints* here because it is sufficiently expressive for most of the constraints considered in [40, 13, 20] and sometimes allows polynomial-time computations [33].

Formally, a *linear atomic constraint* over a set of arguments $\mathcal{A}$ is an expression of the form $\sum_{i=1}^{n} c_i \cdot \pi(A_i) \leq c_0$, where $A_i \in \mathcal{A}$ and $c_i \in \mathbb{Q}$. π is just a syntactic symbol

---

[4] Note that $P(A)$ denotes the probability of argument A (the sum of probabilities of all possible worlds that accept A), while $P(\{A\})$ denotes the probability of the possible world $\{A\}$.

that can be read as "the probability of". We let $\mathcal{C}_\mathcal{A}$ denote the set of all linear atomic constraints over $\mathcal{A}$. A probability function P *satisfies* such a linear atomic constraint iff $\sum_{i=1}^{n} c_i \cdot P(A_i) \leq c_0$. P satisfies a set of linear atomic constraints C, denoted as $P \models C$, iff it satisfies all $l \in C$. In this case, we call C *satisfiable*. We let $\mathrm{Sat}_\Pi(C) = \{P \in \mathcal{P}_\mathcal{A} \mid P \models C\}$ denote the set of all probability functions that satisfy C. We call sets of constraints C_1, C_2 equivalent and write $C_1 \equiv C_2$ iff they are satisfied by the same probability functions, that is, $\mathrm{Sat}_\Pi(C_1) = \mathrm{Sat}_\Pi(C_2)$

Note that constraints with $\geq$ and $=$ can be expressed as well in our language. For $\geq$, just note that $\sum_{i=1}^{n} c_i \cdot \pi(A_i) \leq c_0$ is equivalent to $\sum_{i=1}^{n} -c_i \cdot \pi(A_i) \geq -c_0$. For $=$, note that $\sum_{i=1}^{n} c_i \cdot \pi(A_i) \leq c_0$ and $\sum_{i=1}^{n} c_i \cdot \pi(A_i) \geq c_0$ together are equivalent to $\sum_{i=1}^{n} c_i \cdot \pi(A_i) = c_0$. In particular, we can express probability assignments of the form $\pi(A) = p$ or probability bounds of the form $l \leq \pi(A) \leq u$.

Now assume that we are given an epistemic state represented as a probability function $P \in \mathcal{P}_\mathcal{A}$. Given some new evidence represented as a set of linear atomic constraints (and possibly some existing constraints that we want to preserve), we want to update P. To this end, different update operators have been studied in [15, 19, 17]. Here, we are interested in update operators of the following type.

Definition 1 (Epistemic Update Operator). *An* epistemic update operator *is a function* $\mathcal{U} : \mathcal{P}_\mathcal{A} \times \mathcal{C}_\mathcal{A} \to \mathcal{P}_\mathcal{A} \cup \{\bot\}$ *that satisfies the following properties:*

- **Success:** *If* $C \subseteq \mathcal{C}_\mathcal{A}$ *is satisfiable, then* $\mathcal{U}(P, C) \in \mathrm{Sat}_\Pi(C)$.
- **Failure:** *If* $C \subseteq \mathcal{C}_\mathcal{A}$ *is not satisfiable, then* $\mathcal{U}(P, C) = \bot$.
- **Representation Invariance:** *If* $C_1 \equiv C_2$, *then* $\mathcal{U}(P, C_1) = \mathcal{U}(P, C_2)$.
- **Idempotence:** *If* $C \subseteq \mathcal{C}_\mathcal{A}$ *is satisfiable, then* $\mathcal{U}(\mathcal{U}(P, C), C) = \mathcal{U}(P, C)$.

Success and failure guarantee a well-defined update. That is, if the constraints are satisfiable, the update operator will return a new epistemic state that satisfies the constraints. If the constraints are not satisfiable, $\bot$ will be returned to indicate an inconsistency. Representation invariance guarantees that the result is independent of the syntactic representation of the evidence. Finally idempotence guarantees that applying the same update twice does not change the outcome.

3 The Two-stage Least-squares Update Operator

Several update operators in [19, 17] are based on the idea of satisfying new evidence by changing the current epistemic state in a minimal way. The distance between two probability functions is determined by looking at the probabilities that they assign to possible worlds. For example, one can use the least-squares distance $d_2(P, P') = \sum_{w \in \Omega} (P(w) - P'(w))^2$ or the KL-divergence $d_{KL}(P, P') = \sum_{w \in \Omega} P(w) \cdot \log \frac{P(w)}{P'(w)}$. While this makes perfect sense from a probability-theoretical point of view, the resulting belief changes may be intuitively implausible.

Example 1. Consider a BAF $(\{A, B\}, \emptyset, \emptyset)$ with two unrelated arguments A, B. Suppose our current epistemic state is P_1 as defined in Table 1. Then we have $P_1(A) = 0.6$ and $P_1(B) = 0.7$. Now suppose that we want to update the belief in A to 1. A distance-minimizing update w.r.t. d_2 (i.e. update returning a probability distribution satisfying

w	P_1	P_2	P_3	P_4	P_5	P_6	P_7	P_8	P_9	P_{10}
$\emptyset$	0.1	0	0	0.4	0.45	0.26	0	0.3	0.15	0.35
$\{A\}$	0.2	0.4	0.33	0.05	0	0.19	0.3	0	0.15	0
$\{B\}$	0.3	0	0	0.15	0.1	0.29	0	0.1	0.35	0.15
$\{A, B\}$	0.4	0.6	0.67	0.4	0.45	0.26	0.7	0.6	0.35	0.5

Table 1. Some probability functions over possible worlds used in Example 1.

$\pi(A) = 1$ that is minimally different from P_1 w.r.t. d_2) yields the new epistemic state P_2 from Table 1. Now we have $P_2(A) = 1$ as desired. However, we also have $P_2(B) = 0.6 < 0.7 = P_1(B)$. Similarly, updating with respect to d_{KL} yields P_3 from Table 1 with $P_3(B) = \frac{2}{3} < 0.7 = P_1(B)$. This behaviour is rather counterintuitive in this context, since A and B are completely unrelated. Therefore, we should have $P_1(B) = P_2(B) = P_3(B)$.

In order to bring our model closer to humans' intuition, a two-stage minimization process has been proposed in [17]. In stage 1, we identify all probability distributions that minimize an atomic distance measure. Instead of comparing probability functions elementwise on possible worlds, atomic distance measures compare probability functions only based on the probabilities that they assign to arguments [19]. We consider a quadratic variant here that will allow us to compute some updates in polynomial time.

Definition 2 (Atomic Least-squares Distance (ALS)). *The* **ALS distance measure** *is defined as* $d_{\mathrm{At}}^2(P, P') = \sum_{A \in \mathcal{A}} (P(A) - P'(A))^2$ *for all* $P, P' \in \mathcal{P}_{\mathcal{A}}$.

To begin with, we use the ALS distance to define a naive update operator which does not satisfy our desiderate from Definition 1 yet.

Definition 3 (Naive Least-squares Update Operator). *The* naive LS update operator $u_{\mathrm{At}} : \mathcal{P}_{\mathcal{A}} \times \mathcal{C}_{\mathcal{A}} \to 2^{\mathcal{P}_{\mathcal{A}}}$ *is defined by* $u_{\mathrm{At}}(P, C) = \arg \min_{P' \in \mathrm{Sat}_\Pi(C)} d_{\mathrm{At}}^2(P, P')$.

u_{At} yields those probability functions that satisfy C and minimize the ALS distance to P. However, there is not necessarily a unique solution.

Example 2. Consider P_1 from Table 1. Suppose we recognize a conflict between A and B and want to update with the constraint $l_1 : \pi(A) + \pi(B) \leq 1$. We have $P_1(A) = 0.6$ and $P_1(B) = 0.7$. The cheapest way to satisfy the constraint with respect to the ALS distance is to decrease both probabilities by 0.15. That is, a solution P' must satisfy $P'(A) = P'(\{A\}) + P'(\{A, B\}) = 0.45$ and $P'(B) = P'(\{B\}) + P'(\{A, B\}) = 0.55$. P_4 and P_5 from Table 1 show two minimal solutions from the set $u_{\mathrm{At}}(P_1, \{l_1\})$.

The second stage of the minimization process from [17] deals with the uniqueness problem. Among those probability functions that minimize the atomic distance, we pick the unique one that minimizes a sufficiently strong second distance measure. Here, we will consider again the least-squares distance for stage 2.

Definition 4 (Two-stage Least-squares Update Operator (2LS)). *The* 2LS update operator $\mathcal{U}_{\mathrm{At}}^2 : \mathcal{P}_{\mathcal{A}} \times \mathcal{C}_{\mathcal{A}} \to \mathcal{P}_{\mathcal{A}} \cup \{\bot\}$ *is defined by*

$$\mathcal{U}_{\mathrm{At}}^2(P, C) = \begin{cases} \arg \min_{P' \in u_{\mathrm{At}}(P,C)} \sum_{w \in \Omega} (P(w) - P'(w))^2, & \text{if } u_{\mathrm{At}}(P, C) \neq \emptyset \\ \bot & \text{otherwise.} \end{cases}$$

Before looking at an example, we note that $\mathcal{U}^2_{\mathrm{At}}$ is an epistemic update operator as defined in Definition 1.

Proposition 1. *The 2LS update operator is an epistemic update operator.*

Example 3. Consider again P_1 and the constraint l_1 from Example 2. P_6, shown in Table 1, is the unique solution that minimizes the least-squares distance to P_1 among those distributions that minimize the ALS distance to P_1. That is, $\mathcal{U}^2_{\mathrm{At}}(P_1, \{l_1\}) = P_6$.

Example 4. As another example, we consider again the scenario from Example 1 where a one-stage update changed the belief in B in an implausible way. We get $\mathcal{U}^2_{\mathrm{At}}(P_1, \{\pi(A) = 1\}) = P_7$ shown in Table 1. In particular, we have $P_7(B) = 0.7 = P_1(B)$ as desired.

Intuitively, stage 1 determines which atomic beliefs in arguments have to be changed in order to satisfy the new constraints. This avoids the counterintuitive behaviour of elementwise minimization over the possible worlds, but does not yield a unique solution. Therefore, stage 2 performs an elementwise minimization over the possible worlds to pick a best solution among the ones that minimize the change in atomic beliefs.

4   Updates over Probability Labellings

The two-stage minimization process solves our semantical problems, but we are still left with a considerable computational problem. This is because we consider probability functions over possible worlds whose number grows exponentially with the number of arguments in our framework. However, as illustrated in our previous examples, human reasoning may be guided by atomic beliefs in arguments rather than by beliefs in possible worlds. Therefore, a natural question is, what changes semantically when considering belief functions over arguments rather than over possible worlds? As shown in [33], probability functions over possible worlds can sometimes just be replaced with *probability labellings* $L : \mathcal{A} \to [0, 1]$ that assign beliefs to atomic arguments directly without changing the semantics. We let $\mathcal{L}_\mathcal{A}$ denote the set of all probability labellings.

Formally, probability functions can be related to probability labellings via an equivalence relation [33]. Two probability functions P_1, P_2 are called *atomically equivalent*, denoted as $P_1 \equiv P_2$, iff $P_1(A) = P_2(A)$ for all $A \in \mathcal{A}$. As usual, $[P] = \{P' \in \mathcal{P}_\mathcal{A} \mid P' \equiv P\}$ denotes the equivalence class of P and $\mathcal{P}_\mathcal{A}/\!\equiv\, = \{[P] \mid P \in \mathcal{P}_\mathcal{A}\}$ denotes the set of all equivalence classes. As shown in [32], there is a one-to-one relationship between $\mathcal{P}_\mathcal{A}/\!\equiv$ and $\mathcal{L}_\mathcal{A}$.

Lemma 1 ([32]). *The function $r : \mathcal{P}_\mathcal{A}/\!\equiv\, \to \mathcal{L}_\mathcal{A}$ defined by $r([P]) = L_P$, where $L_P(A) = P(A)$ for all $A \in \mathcal{A}$ is a bijection.*

Intuitively, r determines a compact representation of the equivalence class $[P]$, namely the probability labelling $L_P = r([P])$. Since r is a bijection, every probability labelling can also be related to a set of probabity functions $r^{-1}(L_P) = [P] = \{P' \in \mathcal{P}_\mathcal{A} \mid P' \equiv P\}$. Intuitively, $r^{-1}(L)$ is just the set of probability functions that satify the atomic beliefs encoded in L. We say that a probability labelling L satisfies a linear atomic constraint $\sum_{i=1}^n c_i \cdot \pi(A_i) \leq c_0$ iff $\sum_{i=1}^n c_i \cdot L(A_i) \leq c_0$. The set of probability

labellings that satisfy a set of such constraints C is denoted by $\mathrm{Sat}_A(C)$. The following observations from [32] are helpful to simplify computational problems by replacing probability functions with probability labellings.

Lemma 2 ([32]). *The following statements are equivalent: (1) P satisfies a linear atomic constraint l; (2) All $P' \in [P]$ satisfy l; (3) $L_P = r([P])$ satisfies l.*

For example, in order to decide whether a set of linear atomic constraints C is satisfiable by a probability function (of exponential size), we can just check whether it can be satisfied by a probability labelling (of linear size) [32]. If such a labelling L exists, all probability functions in $r^{-1}(L)$ satisfy C. Conversely, if some probability function P satisfies C, then $L = r([P])$ satisfies C as well.

In order to perform updates more efficiently, we could represent epistemic states by probability labellings. However, we should ask, what is the relationship between updates over probability functions and updates over probability labellings? We first note that update operators $\mathcal{U}_W$ that simply minimize the distance over possible worlds are not necessarily compatible with atomic equivalence. That is, given a set of linear atomic constraints C and two probability functions P_1 and P_2 such that $P_1 \equiv P_2$, we do not necessarily have $\mathcal{U}_W(P_1, C) \equiv \mathcal{U}_W(P_2, C)$.

Example 5. Consider P_1 and P_8 in Table 1. We have $P_1(A) = 0.6 = P_8(A)$ and $P_1(B) = 0.7 = P_2(B)$, that is, $P_1 \equiv P_8$. Suppose, we update with $C = \{\pi(A) = 0.5\}$ and update by just minimizing the least-squares distance to P_1. Then $\mathcal{U}_W(P_1, C) = P_9$ and $\mathcal{U}_W(P_8, C) = P_{10}$, where P_9, P_{10} are again shown in Table 1. We have $P_9(B) = 0.7 \neq 0.65 = P_{10}(B)$, that is, $P_9 \not\equiv P_{10}$.

Update operators based on atomic distance measures give us compatibility guarantees that we explain in the following proposition.

Proposition 2. *Let $P_1, P_2 \in \mathcal{P}_\mathcal{A}$ and let $C \subset \mathcal{C}_\mathcal{A}$ be a finite set of linear atomic constraints. If $P_1 \equiv P_2$, then*

1. *$d^2_{\mathrm{At}}(P_1, P) = d^2_{\mathrm{At}}(P_2, P)$ for all $P \in \mathcal{P}_\mathcal{A}$,*
2. *$u_{\mathrm{At}}(P_1, C) = u_{\mathrm{At}}(P_2, C)$,*
3. *$P_1' \equiv P_2'$ for all $P_1', P_2' \in u_{\mathrm{At}}(P_1, C)$,*
4. *$\mathcal{U}^2_{\mathrm{At}}(P_1, C) \equiv \mathcal{U}^2_{\mathrm{At}}(P_2, C)$.*

Item 1 says that the ALS distance is invariant under atomically equivalent probability functions. This implies that the updates that minimize the ALS distance are invariant as well (item 2). As we demonstrated in Example 2, such updates do not necessarily yield a unique solution. However, when using the ALS distance, we can guarantee that all solutions are atomically equivalent (item 3). This implies that the 2LS update operator is invariant under atomically equivalent probability functions in the sense that it yields equivalent results when the prior probability functions are equivalent (item 4).

Hence, when updating with respect to linear atomic constraints, there is a well defined relationship between probability functions and probability labellings. If we start with an epistemic state represented by a probability labelling L, L can be understood as a compact representation of the set of probability functions $r^{-1}(L)$ that satisfy the
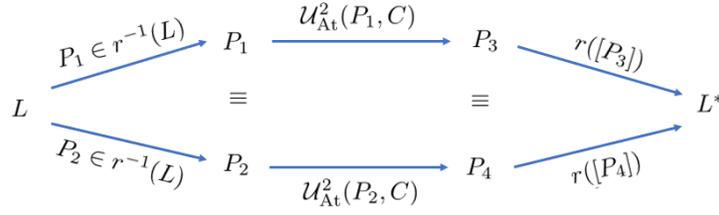
Fig. 2. The 2LS update operator $\mathcal{U}^2_{\mathrm{At}}$ respects atomic equivalence.

atomic beliefs encoded in L. The 2LS update operator is compatible with this representation. That is, no matter which probability functions from $r^{-1}(L)$ we choose, an update with linear atomic constraints will always lead to the same equivalence class and therefore to a well defined next probability labelling L^*. We illustrate this in Figure 2.

If we are only interested in atomic beliefs, it would be convenient if we could move directly from L to L^* in Figure 2 without generating (exponentially large) probability functions in the process. We can do this indeed in polynomial time for the 2LS update operator. In order to show this, we first define an update operator on labellings.

Definition 5 (Least-squares Labelling Update Operator (2LS)). *The* LS labelling update operator $LU^2_\lambda : \mathcal{L}_\mathcal{A} \times \mathcal{C}_\mathcal{A} \to \mathcal{L}_\mathcal{A} \cup \{\bot\}$ *is defined by*

$$LU^2_\lambda(L,C) = \begin{cases} \arg\min_{L' \in \mathrm{Sat}_\Lambda(C)} \sum_{A \in \mathcal{A}} (L(A) - L'(A))^2, & \text{if } \mathrm{Sat}_\Lambda(C) \neq \emptyset \\ \bot & \text{otherwise.} \end{cases}$$

As we explain in the following theorem, LU^2_λ provides us with a direct path from L to L^* and can be computed in polynomial time.

Theorem 1. *Let $C \subset \mathcal{C}_\mathcal{A}$ be a finite and satisfiable set of linear atomic constraints and let $L \in \mathcal{L}_\mathcal{A}$. Then $LU^2_\lambda(L,C) = L^*$ is well-defined and can be computed in polynomial time. Furthermore, $L^* = r([\mathcal{U}^2_{\mathrm{At}}(P,C)])$ for all $P \in r^{-1}(L)$.*

Hence, when we are only interested in atomic beliefs, we can use probability labellings to represent epistemic states and use the least-squares labelling update operator for updates. Semantically, this is equivalent to regarding epistemic states as sets of probability functions that satisfy the same atomic beliefs and updating with respect to the 2LS update operator. The benefit of the labelling representation is that we can perform updates in polynomial time.

5 Application Example

In this section we come back to the graph in Figure 1 and analyze a scenario that, while being hypothetical, uses the data from an empirical study in [11]. In this study, the user's belief in argument A changed from 0 to 0.19 during the dialogue.[5]

---

[5] We note that the study data contained examples of dialogues that resulted in a bigger belief change, however, we have chosen this one due to its interesting structure.

L	A	B	C	D	E	F	G	H	I	J	K	L	M
L_0	0	1	0	1	1	1	0	0	1	1	1	0	0
$L_1 = LU_{\bar\lambda}^2(L, C \cup \Phi)$	0.19	0.81	0.19	0.505	0.975	0.95	0.495	0.05	0.92	0.09	0.95	0.08	0.91

Table 2. Probability labelings before and after the dialogue from Section 5.

The graph in Figure 1 is generated from an existing dialogue that involved an automated dialogue system and a human user. Arguments at even depth (starting from A) are system arguments (A, C, G, H, L and M), while the ones at odd depth are user arguments. The agents take turns in uttering their arguments (starting with A), and arguments at the same depth are uttered at the same point by a given party. We observe that not all user arguments are met with a system response (see arguments E and K). Despite this fact, the presented arguments have led to a positive change in belief in A, contrary to what would be the intuition from the classical Dungean approaches. It is possible that if all of the user's counterarguments were addressed, then the belief increase would be even more prominent.

We can try to provide an explanation for the belief change observed in [11] by modeling the reasoning process in our framework. Let us assume that the constraints representing the user's reasoning demand that the belief in an argument is dual to the belief in the average of its attackers. That is, we assume $P(X) = 1 - \frac{1}{|Att(X)|} \sum_{Y \in Att(X)} P(Y)$, where $Att(X) = \{Y \in \mathcal{A} \mid (Y, X) \in \mathcal{R}\}$). This assumption leads to the following set of constraints:

$$C = \{\pi(A) + \pi(B) = 1, \pi(B) + \pi(C) = 1, \pi(D) + \pi(G) = 1, \pi(F) + \pi(H) = 1,$$
$$\pi(C) + 0.33\pi(D) + 0.34\pi(E) + 0.33\pi(F) = 1, \pi(C) + 0.5\pi(I) + 0.5\pi(J) = 1,$$
$$\pi(H) + \pi(K) = 1, \pi(I) + \pi(L) = 1, \pi(J) + \pi(M) = 1\}$$

Let us further assume that the user initially completely accepts his or her own arguments and completely rejects the system's arguments. This belief state is represented by the labeling L_0 shown in Table 2. We now consider a possible persuasion system which, once a given dialogue branch is exhausted, asks the user about his or her beliefs in the unattacked arguments. In our case, the user states that he or she believes L, M, E and K with the degrees 0.08, 0.91, 0.975 and 0.95 respectively. This produces constraints $\Phi = \{\pi(L) = 0.08, \pi(M) = 0.91, \pi(E) = 0.975, \pi(K) = 0.95\}$. We can use this information along with C to update L_0 without asking the user his or her beliefs in all possible arguments. The resulting labeling $L_1 = LU_{\bar\lambda}^2(L, C \cup \Phi)$ is shown in Table 2.

We observe that the belief in A in and L_0 and L_1 match the expected beliefs 0 and 0.19 based on the data in [11].

6 Related Work

There is a large variety of other probabilistic argumentation approaches [8, 25, 37, 14, 6, 27, 41, 24, 38, 42, 39], which basically differ in the level of detail (e.g., structured or abstract argumentation), in the way how uncertainty is introduced (e.g. possible worlds

correspond to argument interpretations or the graph structure) and in the nature of uncertainty (e.g., uncertainty about the acceptance state or uncertainty about the nature of a relation between arguments).

One limitation when restricting to probability labellings is that we cannot compute the probabilities of complex formulas over arguments anymore without adding further assumptions. However, as we demonstrated, we can sometimes do without complex formulas. In this context, probability labellings can be seen as an alternative to weighted argumentation frameworks that also assign a strength value between 0 and 1 to arguments [3, 36, 2, 26, 31]. What makes probability labellings an interesting alternative is their well-defined relationship to probability functions and probability theory.

The problem of adapting an epistemic state with respect to new knowledge has been studied extensively in the belief revision literature that evolved from the AGM theory developed in [1]. An up-to-date discussion of the main ideas can be found in [12]. Our postulates are inspired by AGM postulates. For example, *Success* and *Representation Invariance* can be seen as the counterparts of the *Closure* and *Extensionality* postulates in AGM theory. The closest relative to our setting is probably the probabilistic belief change framework from [21]. For a discussion of relationships between classical and probabilistic belief changes, see [21] and [22].

Other equivalence relations have been studied in order to improve the computational performance of probabilistic reasoning algorithms [10, 23, 9, 30]. However, usually, these equivalence relations are introduced over possible worlds, not over probability functions. They can be applied to more expressive reasoning formalisms (they are not restricted to atomic beliefs), but identifying compact representatives for the corresponding equivalence classes remains intractable in general [34].

7 Conclusions

We demonstrated that, in the fragment of linear atomic constraints, it is possible to relate updates over probability labellings to equivalent updates over classes of probability functions. This is interesting from a cognitive, a probabilistic-logical and a computational perspective. Atomic beliefs are often easier to understand for humans. If we can relate these beliefs to probability functions, we get a strong foundational basis. Finally, they can be stored much more compactly and give us polynomial runtime guarantees. Our results can probably be generalized to other two-stage update operators. However, the building blocks for the two stages have to be chosen carefully in order to guarantee that the update operator respects atomic equivalence. For example, it may not be possible to relate the two-stage update process considered in [17], Section 5, to an update operator over probability labellings in a meaningful way. However, we may be able to construct similar relationships by replacing the least-squares distance with KL-divergence or more general classes of distance measures. An implementation of our update operator is available in the Java library *ProBabble*[6].

---

[6] https://sourceforge.net/projects/probabble/

References

1. Alchourrón, C., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. Journal of Symbolic Logic 50(2), 510–530 (1985)
2. Amgoud, L., Ben-Naim, J.: Evaluation of arguments in weighted bipolar graphs. In: Proc. of ECSQARU'17. LNCS, vol. 10369, pp. 25–35. Springer (2017)
3. Baroni, P., Romano, M., Toni, F., Aurisicchio, M., Bertanza, G.: Automatic evaluation of design alternatives with quantitative argumentation. Argument & Computation 6(1), 24–49 (2015)
4. Cayrol, C., Lagasquie-Schiex, M.C.: Bipolarity in argumentation graphs: Towards a better understanding. International Journal of Approximate Reasoning 54(7), 876–899 (2013)
5. Cohen, A., Gottifredi, S., García, A.J., Simari, G.R.: A survey of different approaches to support in argumentation systems. Knowledge Engineering Review 29(5), 513–550 (2014)
6. Doder, D., Woltran, S.: Probabilistic argumentation frameworks–a logical approach. In: Proc. of SUM'14. LNCS, vol. 8720, pp. 134–147. Springer (2014)
7. Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. Artificial intelligence 77(2), 321–357 (1995)
8. Dung, P.M., Thang, P.M.: Towards (probabilistic) argumentation for jury-based dispute resolution. In: Proc. of COMMA'10. FAIA, vol. 216, pp. 171–182. IOS Press (2010)
9. Finthammer, M., Beierle, C.: Using equivalences of worlds for aggregation semantics of relational conditionals. In: Proc. of KI'12. LNCS, vol. 7526, pp. 49–60. Springer (2012)
10. Fischer, V.G., Schramm, M.: Tabl-a tool for efficient compilation of probabilistic constraints. Tech. Rep. Technical Report TUM-19636, Technische Universitaet Muenchen (1996)
11. Hadoux, E., Hunter, A., Polberg, S.: Strategic argumentation dialogues for persuasion: Framework and experiments based on modelling the beliefs and concerns of the persuadee. Tech. rep., University College London (2019)
12. Hansson, S.: Logic of belief revision. In: The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University, winter 2017 edn. (2017)
13. Hunter, A.: A probabilistic approach to modelling uncertain logical arguments. International Journal of Approximate Reasoning 54(1), 47–81 (2013)
14. Hunter, A.: Probabilistic qualification of attack in abstract argumentation. International Journal of Approximate Reasoning 55(2), 607–638 (2014)
15. Hunter, A.: Modelling the persuadee in asymmetric argumentation dialogues for persuasion. In: Proc. of IJCAI'15. pp. 3055–3061. AAAI Press (2015)
16. Hunter, A.: Computational persuasion with applications in behaviour change. In: Proc. of COMMA'16. FAIA, vol. 287, pp. 5–18. IOS Press (2016)
17. Hunter, A., Polberg, S., Potyka, N.: Updating belief in arguments in epistemic graphs. In: Proc. of KR'18. pp. 138–147. AAAI Press (2018)
18. Hunter, A., Polberg, S., Thimm, M.: Epistemic graphs for representing and reasoning with positive and negative influences of arguments. arXiv preprint arXiv:1802.07489v1 (2018)
19. Hunter, A., Potyka, N.: Updating probabilistic epistemic states in persuasion dialogues. In: Proc. of ECSQARU'17. LNCS, vol. 10369, pp. 46–56. Springer (2017)
20. Hunter, A., Thimm, M.: On partial information and contradictions in probabilistic abstract argumentation. In: Proc. of KR'16. pp. 53–62. AAAI Press (2016)
21. Kern-Isberner, G.: Conditionals in nonmonotonic reasoning and belief revision, LNAI, vol. 2087. Springer (2001)
22. Kern-Isberner, G.: Linking iterated belief change operations to nonmonotonic reasoning. In: Proc. of KR'08. pp. 166–176. AAAI Press, Menlo Park, CA (2008)

23. Kern-Isberner, G., Lukasiewicz, T.: Combining probabilistic logic programming with the power of maximum entropy. Artificial Intelligence 157(1-2), 139–202 (2004)

24. Kido, H., Okamoto, K.: A bayesian approach to argument-based reasoning for attack estimation. In: Proc. of IJCAI'17. pp. 249–255. AAAI Press (2017)

25. Li, H., Oren, N., Norman, T.J.: Probabilistic argumentation frameworks. In: Proc. of TAFA'11. LNCS, vol. 7132, pp. 1–16. Springer (2011)

26. Mossakowski, T., Neuhaus, F.: Modular semantics and characteristics for bipolar weighted argumentation graphs. arXiv preprint arXiv:1807.06685 (2018)

27. Polberg, S., Doder, D.: Probabilistic abstract dialectical frameworks. In: Proc. of JELIA'14. LNCS, vol. 8761, pp. 591–599. Springer (2014)

28. Polberg, S., Hunter, A.: Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches. International Journal of Approximate Reasoning 93, 487–543 (2018)

29. Polberg, S., Oren, N.: Revisiting support in abstract argumentation systems. In: Proc. of COMMA'14. FAIA, vol. 266, pp. 369–376. IOS Press (2014)

30. Potyka, N.: Solving Reasoning Problems for Probabilistic Conditional Logics with Consistent and Inconsistent Information. Ph.D. thesis (2016)

31. Potyka, N.: Continuous dynamical systems for weighted bipolar argumentation. In: Proc. of KR'18. pp. 148–157. AAAI Press (2018)

32. Potyka, N.: A polynomial-time fragment of epistemic probabilistic argumentation (technical report). arXiv preprint arXiv:1807.06685 (2018)

33. Potyka, N.: A polynomial-time fragment of epistemic probabilistic argumentation (extended abstract). In: Proc. of AAMAS'19. p. (to appear). IFAAMAS (2019)

34. Potyka, N., Beierle, C., Kern-Isberner, G.: 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 (2015)

35. Potyka, N., Polberg, S., Hunter, A.: Polynomial-time updates of epistemic states in a fragment of probabilistic epistemic argumentation (technical report). arXiv preprint arXiv:1906.05066 (2019)

36. Rago, A., Toni, F., Aurisicchio, M., Baroni, P.: Discontinuity-free decision support with quantitative argumentation debates. In: Proc. of KR'16. pp. 63–73. AAAI Press (2016)

37. Rienstra, T.: Towards a probabilistic Dung-style argumentation system. In: Proc. of AT'12. pp. 138–152 (2012)

38. Rienstra, T., Thimm, M., Liao, B., van der Torre, L.: Probabilistic abstract argumentation based on scc decomposability. In: Proc. of KR'18. pp. 168–177. AAAI Press (2018)

39. Riveret, R., Baroni, P., Gao, Y., Governatori, G., Rotolo, A., Sartor, G.: A labelling framework for probabilistic argumentation. Annals of Mathematics and Artificial Intelligence 83(1), 21–71 (2018)

40. Thimm, M.: A probabilistic semantics for abstract argumentation. In: Proc. of ECAI'12. FAIA, vol. 242, pp. 750–755. IOS Press (2012)

41. Thimm, M., Baroni, P., Giacomin, M., Vicig, P.: Probabilities on extensions in abstract argumentation. In: Proc. of TAFA'17. LNCS, vol. 10757, pp. 102–119. Springer (2017)

42. Thimm, M., Cerutti, F., Rienstra, T.: Probabilistic graded semantics. In: Proc. of COMMA'18. FAIA, vol. 305, pp. 369–380. IOS Press (2018)