

Argument-based Belief in Topological Structures

Chenwei Shi Sonja Smets Fernando R. Velázquez-Quesada

Institute for Logic, Language and Computation
Amsterdam, The Netherlands

{C.Shi, S.J.L.Smets, F.R.VelazquezQuesada}@uva.nl

This paper combines two studies: a topological semantics for epistemic notions and abstract argumentation theory. In our combined setting, we use a topological semantics to represent the structure of an agent’s collection of evidence, and we use argumentation theory to single out the relevant sets of evidence through which a notion of *beliefs grounded on arguments* is defined. We discuss the formal properties of this newly defined notion, providing also a formal language with a matching modality together with a sound and complete axiom system for it. Despite the fact that our agent can combine her evidence in a ‘rational’ way (captured via the topological structure), argument-based beliefs are not closed under conjunction. This illustrates the difference between an agent’s reasoning abilities (i.e. the way she is able to combine her available evidence) and the closure properties of her beliefs. We use this point to argue for why the failure of closure under conjunction of belief should not bear the burden of the failure of rationality.

1 Introduction

The concept of belief has been extensively studied in philosophy, logic and computer science. The representations range from purely qualitative structures including the relational models with serial, transitive and Euclidean relations (the *KD45* representation in doxastic logic; [23, 30]), plausibility models [10, 13] and topological models [3, 4, 5] to quantitative structures including the ranking-based plausibility representation [35], Bayesian models using subjective probability functions and conditional probabilistic spaces designed to represent conditional belief [9].

Despite the success of the early qualitative and quantitative approaches, most of them had one important disadvantage: they represent only what the agent believes, without dealing with the reasons, justifications or evidence on which such beliefs are grounded.¹ To remedy this shortcoming, the last decade witnessed a number of different frameworks that can portray the evidence or justifications on which different epistemic attitudes are based. For instance, the work in [8, 7], combines dynamic epistemic logic with justification logic in the tradition of [1] and [2] where justifications are syntactic terms. In [12] and [11] a semantic approach is adopted, representing evidence as a set of possible worlds so beliefs are defined by the maximally consistent ways in which evidence can be combined (so called *evidence models*). The work in [5] follows the latter direction, adding a topological structure which represents the different ways in which the available pieces of evidence can be combined, and then using topological notions to single out relevant sets of *combined* pieces of evidence.

The present contribution follows the mentioned topological approach [5], where we view arguments as special types of evidence that an agent possesses. By combining the topological approach to evidence logic with the tools from abstract argumentation theory [18], we can single out the sets of pieces of evidence (or arguments) on which a new notion of *belief grounded on arguments* can be defined. This new

¹An important exception are the so called *truth maintenance systems* [17, 26] which, inspired by the typical forms of arguments in systems for inference in natural deduction, keep track of *syntactic* justifications.

concept of *grounded belief* provides the bridge between an agent’s doxastic attitude and the arguments she has at her disposal. While argumentation theorists such as [18] informally talk about the relation between beliefs and arguments, the formalization of the details has only recently gained more attention (e.g., [33, 22]). The approach in this paper is related to the line of work which studies abstract argumentation theory by using the tools and techniques of modal logic [21, 14], and to the work by [22] which introduces a notion of justified belief by combining doxastic logic and abstract argumentation theory. Yet in contrast to the work on relational models for modal logic, we do turn to the topological semantics of [5].

Due to its semantic nature, one might expect for this paper’s notion of grounded belief to be very much like the notions of belief studied in relational and evidence models, but actually this is not the case. Grounded beliefs turn out to be not closed under conjunction, which makes them more similar to the probabilistic notions of belief as studied in [29] and [36]. In the final part of this paper we come back to this analysis and provide a comparison between our concept of grounded belief and its quantitative counterpart. From a philosophical perspective the lack of closure raises a number of interesting issues. It shows how, in our framework, there is a clear difference between the agent’s reasoning abilities (i.e., the way she is able to combine her available evidence, represented by the topological structure) and the closure properties of her beliefs (related to the strategy she uses to select the sets of combined evidence on which her beliefs will be based). This relates to a point which is raised by [20] in the context of the Lottery paradox. [20] describes a dividing line between processes of deductive reasoning (which are closed under conjunction) and those which involve beliefs and are not closed under conjunction. While we agree with Foley that there are essential differences in such reasoning patterns, our paper does show how the two can be made formally precise and can moreover be united in a formal framework that defines grounded beliefs on the basis of deductive reasoning processes for combining evidence.

Outline. Section 2 recalls the work on evidence and belief that our proposal is based on [12, 5]. Section 3 introduces our framework, defining our notion of *belief grounded on arguments* and comparing it with the topological notion of evidence-based belief in [5]. Section 4 provides a sound and complete logic characterizing the properties of our notion of belief. Section 5 compares our setting with the probabilistic notion of belief, re-examining the criticism faced by notions of belief which are not closed under conjunction. We conclude in Section 6, listing several directions for future work.

2 Preliminaries

In this section we recall the basic concepts and definitions of the framework in [12] designed to present evidence logic as well as the further topological development presented in [5].

Beliefs in Evidence Models The evidence an agent has for supporting her beliefs can be represented in different ways. The proposal in [12] takes a semantic perspective and uses neighborhood models [34, 31], a generalisation of relational models, in which a so-called *neighborhood function* assigns a family of subsets of the domain to each possible world.^{2 3}

²For our purposes, we work only with the so-called *uniform* evidence models in [12], in which the neighborhood function is world-invariant.

³As Definition 2.1 shows, the concept of evidence represented in evidence models is rather abstract: a piece of evidence is understood as a (non-empty) set of possible worlds. As it will be discussed, these pieces of evidence can be understood as information the agent receives from external sources (observations, communication), and the closure properties the evidence set might have can be understood as the different ways she can ‘extract’ further information from what she gets. For an in-depth study about the different interpretations the term *evidence* and its ‘cousin’ *justification* might have, the reader is referred to [25]

Definition 2.1 (Evidence model [12]). *An evidence model is a tuple $\mathcal{M} = (X, E_0, V)$ where $X \neq \emptyset$ is a set of possible worlds, $E_0 \subseteq 2^X - \{\emptyset\}$ is a family of non-empty sets called pieces of evidence, and $V : \text{Prop} \rightarrow 2^W$ is a valuation function for a given set Prop of atomic propositions. The family E_0 is required to include X (i.e. $X \in E_0$).*

Note that the above definition mentions two explicit constraints on evidence models: X is always in the neighborhood (the whole domain is always a piece of evidence) but \emptyset is never there (evidence per se is never contradictory). As these are the only possible constraints, it can happen that two pieces of evidence in E_0 are in conflict with each other (i.e., there might be $e_1, e_2 \in E_0$ such that $e_1 \cap e_2 = \emptyset$). Moreover, in this setting E_0 is not required to satisfy any closure property, e.g. closure under finite/arbitrary unions and/or intersections. Still, the lack of closure properties in *the family* of pieces of evidence does not imply that *the agent* is not able to combine the evidence pieces in a meaningful way. Following the terminology in [6, 5], we introduce the notion of a body of evidence:

Definition 2.2 (Body of evidence). *Let $\mathcal{M} = (X, E_0, V)$ be an evidence model. A body of evidence $F \subseteq E_0$ is a subfamily of E_0 which has the finite intersection property.⁴ A body of evidence is maximal if it cannot be properly extended to any other body of evidence.*

Next we indicate when a body of evidence supports a proposition P [6, 5]:

Definition 2.3 (Support of a proposition). *Let $\mathcal{M} = (X, E_0, V)$ be an evidence model. A body of evidence $F \subseteq E_0$ supports a proposition P if and only if $\bigcap F \subseteq P$.*

The notion of belief à la van Benthem and Pacuit is based on the agent's body of evidence. In [12], an agent *believes a proposition* $P \subseteq X$ if and only if every *maximal* body of evidence $F \subseteq E_0$ supports P (i.e., $\bigcap F \subseteq P$). According to this definition in evidence logic, the agent is able to combine her available evidence in a maximally consistent way. Yet as explicitly indicated in [5], this does not mean that all her beliefs will be consistent. Indeed, while the mentioned concept of belief works well in the finite case, it can yield inconsistent beliefs in the infinite case:

Example 2.1. *Consider the evidence model $(\mathbb{N}, E_0 = \{[n, +\infty) \mid n \in \mathbb{N}\}, \emptyset)$. Note how E_0 itself is a body of evidence and, moreover, is the unique maximal one. But $\bigcap E_0 = \emptyset$, and thus the agent believes \emptyset .*

Such examples have been the main motivation in [5] to provide an improved semantics for evidence-based beliefs which changes the above given definition of van Benthem and Pacuit. Their revised concept of belief is provided in the context of a topological semantics:

Topological Semantics for Beliefs The *topological* approach of [5] generalizes the setting of [12], as it introduces a concept of belief which is always consistent. Moreover, the definitions it provides for knowledge and other related epistemic notions have both epistemic and topological significance. We limit ourselves here to the basic doxastic concept:

Definition 2.4 (Topological evidence model [5]). *A topological evidence model $\mathcal{M} = (X, E_0, \tau_{E_0}, V)$ extends an evidence model (X, E_0, V) (Definition 2.1) with τ_{E_0} , the topology over X generated by E_0 .⁵ For simplicity, and when no confusion arises, τ_{E_0} will be denoted simply by τ .*

(for the first) and [32] (for the second).

⁴A family $F \subseteq 2^X$ has the *finite intersection property* when the intersection of every finite subset of F is non-empty.

⁵A *topology* over $X \neq \emptyset$ is a family $\tau \subseteq 2^X$ containing both X and \emptyset , and closed under finite intersections and arbitrary unions. The elements of a topology are called *open sets*. The *topology generated by* $E \subseteq X$ is the smallest topology τ_E over X such that $E \subseteq \tau_E$.

A topological evidence model extends an evidence model with a topology: a structure describing the different ways in which the available pieces of evidence can be combined. Note that pieces of evidence are modeled as opens in the topology; since a topology is closed under *finite* intersection, this fits our intuition about the agent's limited ability for combining evidence. Further reasons for modeling evidence as opens can be found in [37, Chapter 2], [24] and [5]. In this setting, Baltag et al. provide an improved version of evidence-based belief:

Definition 2.5 (Evidence-based belief [5]). *Let \mathcal{M} be a topological evidence model (X, E_0, τ, V) . The agent believes a proposition $P \subseteq X$ if and only if every non-empty open set can be 'strengthened' into an open set supporting P (i.e., for any $t \in \tau \setminus \{\emptyset\}$ there is $t' \in \tau \setminus \{\emptyset\}$ such that $t' \subseteq t$ and $t' \subseteq P$).*

This notion of belief, which we call *evidence-based belief* (Bel), coincides with the concept of belief à la van Benthem and Pacuit in evidence models when the domain is finite. However, Bel does remain consistent in the infinite case (i.e., in topological evidence models, $Bel\emptyset = \emptyset$).

Based on these ingredients, the next section will extend the topological evidence models by adding tools from formal argumentation theory [18].

3 Topological Argumentation Models

Recall how, in the above evidence models of [12], pieces of evidence may be in conflict (i.e., there may be $e_1, e_2 \in E_0$ such that $e_1 \cap e_2 = \emptyset$). This is a natural assumption, as an agent may collect different pieces of evidence (from, e.g., different sources) supporting contradictory facts. The crucial issue is then to find reasonable ways for combining these pieces of evidence or, in other words, to single out meaningful subsets of E_0 that can be used to define a notion of belief. In the above mentioned approach, this role is played by maximal bodies of evidence.

Something similar happens in the topological approach of [5]: even though *combined* pieces of evidence may be in conflict (there may be $t_1, t_2 \in \tau$ such that $t_1 \cap t_2 = \emptyset$), the authors singled out those whose complement is *nowhere dense* in order to define their topological notion of evidence-based belief.⁶ This suggests that different doxastic concepts can be provided by singling out alternative (but, of course, reasonable) combined pieces of evidence.

Let us now focus on specific pieces of evidence that we can think of as examples of an agent's arguments. As readers familiar with the abstract argumentation theory of [18] know, its framework and the various concepts that are defined upon it, can be understood as a collection of tools that allow us to single out acceptable arguments (in our context, amounting to 'acceptable' pieces of evidence) from all given arguments. In order to combine such tools with the semantic setting of topological evidence models, one only needs to provide a reasonable definition of what it means for a piece of evidence (or 'argument') to attack another, where a piece of evidence is represented as a set of worlds. It is natural to say that two pieces of evidence attack each other when they are in conflict, that is, when their intersection is empty. For example, in a given topological evidence model,

$$\text{for any } t_1, t_2 \in \tau \setminus \{\emptyset\}: t_1 \text{ and } t_2 \text{ attack each other iff } t_1 \cap t_2 = \emptyset.$$

Note that in this example the attack is mutual (i.e., the attack relation is symmetric), yet that does not always need to be the case.

From Conflict to Attack In this paragraph we use an attack relation that is not necessarily symmetric, but it is still related to the mentioned notion of conflict.

⁶See [5] for the definition of *dense*, *nowhere dense* and other topological notions, and also for their epistemic interpretation.

Definition 3.1 (Topological argumentation model). A topological argumentation model $\mathcal{M} = (X, E_0, \tau_{E_0}, \leftarrow, V)$ extends a topological evidence model (X, E_0, τ, V) (Definition 2.4) with a relation $\leftarrow \subseteq (\tau \times \tau)$, called the attack relation on τ (with $t_1 \leftarrow t_2$ read as “ t_2 attacks t_1 ”), which is required to satisfy the following conditions:

1. for any $t_1, t_2 \in \tau$: $t_1 \cap t_2 = \emptyset$ if and only if $t_1 \leftarrow t_2$ or $t_2 \leftarrow t_1$;
2. for any $t, t_1, t'_1 \in \tau$: if $t_1 \leftarrow t$ and $t'_1 \subseteq t_1$, then $t'_1 \leftarrow t$;
3. for every $t \in \tau \setminus \{\emptyset\}$: $\emptyset \leftarrow t$ and $t \not\leftarrow \emptyset$.

The first condition states that attack implies conflict (right to left), but also that, while conflict implies attack, the implied attack does not need to be mutual (left to right). The second condition asks that, if t attacks t_1 , then it should also attack any stronger t'_1 . The last condition establishes that the empty set is attacked by all non-empty opens, and that it does not attack any of them.⁷

3.1 Grounded Semantics for Argument-based Belief

The attack relation \leftarrow defines, together with its domain τ , an attack graph $\mathcal{A}_\tau = (\tau, \leftarrow)$. Given such a structure, it is possible to use the tools from abstract argumentation theory in order to single out ‘good’ families of combined evidence (or arguments), which in turn will allow us to define a notion of belief grounded in arguments. In order to do that, here is first a quick recollection of the basic concepts of argumentation theory put forward in our framework:

Definition 3.2 (Characteristic (defense) function). Let \mathcal{M} be a topological argumentation model, and let $\mathcal{A}_\tau = (\tau, \leftarrow)$ be the attack graph on τ . A subset $T \subseteq \tau$ is said to defend $t \in \tau$ if and only if any open t' attacking t (i.e., for all $t' \in \tau$ such that $t \leftarrow t'$) is itself attacked by some open in T (i.e., there is $t'' \in T$ such that $t' \leftarrow t''$). The characteristic function of \mathcal{A}_τ , denoted by d_τ and also called the defense function, receives a set of opens $T \subseteq \tau$ and returns the set of opens that are defended by T :

$$d_\tau(T) := \{t \in \tau \mid t \text{ is defended by } T\}$$

When $t \in d_\tau(T)$, it is said that t is acceptable with respect to T .

The characteristic function d_τ is monotonic [18, Lemma 19], and therefore it has a least fixed point LFP_τ (i.e., LFP_τ is the smallest subset of τ satisfying $\text{LFP}_\tau = d_\tau(\text{LFP}_\tau)$). Note, then, how this least fixed point is a set that can defend all (\subseteq) and only (\supseteq) its members against any attack. Moreover, it can be proved that the set is also conflict-free (i.e., there are no $t, t' \in \text{LFP}_\tau$ such that $t \leftarrow t'$). Due to these properties, LFP_τ provides an excellent candidate for the relevant family of open sets in τ over which beliefs can be defined. (In abstract argumentation, LFP_τ is called *the grounded extension*.) There are, of course, other alternatives on which we can base beliefs: abstract argumentation theory defines also, e.g., the notions of an *admissible set* (a set that is conflict-free and all its arguments are acceptable with respect to it), a *complete set* (one that is admissible and contains every acceptable argument with respect to it), a *preferred set* (a maximal admissible set) and a *stable set* (a conflict-free set that attacks every argument not in it). One of the reasons why the grounded set has been chosen is that, while the other alternatives provide more than one set (there is, in general, more than one admissible/complete/preferred/stable set), the grounded set is always unique; this guarantees that beliefs will be consistent, a property that is typically desirable.⁸ One possible drawback of the grounded extension is that, in general, it may be

⁷In fact, as the first condition implies, it only attacks itself.

⁸Of course, consistency can be also achieved when some of the other mentioned sets are used, but this would require the use of an external mechanism, and that might not be desirable for simplicity reasons.

empty; however, this is not the case in our framework, as X is never attacked (it is in conflict only with the empty set, which does not attack anybody) and thus it is always in LFP_τ . Moreover, the grounded extension is closely related to the topological notion of belief defined in [5], as we will show later. Thus we propose the following definition of belief grounded on arguments:

Definition 3.3 (Grounded Belief). *Let \mathcal{M} be the topological argumentation model $(X, E_0, \tau_{E_0}, \leftarrow, V)$. The agent believes a proposition $P \subseteq X$ (notation: $\mathcal{B}P$) if and only if there is an open set in LFP_τ supporting P , that is*

$$\mathcal{B}P \text{ if and only if } \exists f \in \text{LFP}_\tau : f \subseteq P$$

The definition says that the agent has grounded belief of P if and only if the agent has a “good” argument for P which is a member of the least fixed point.

Properties of grounded beliefs From its definition (Definition 3.3) we observe that grounded beliefs are upward-closed ($\mathcal{B}P$ and $P \subseteq Q$ imply $\mathcal{B}Q$). But a stronger claim also holds: LFP_τ itself is closed upwards.

Proposition 3.1. *Given a topological argumentation model, if $f \in \text{LFP}_\tau$ and $f' \in \tau$ is such that $f \subseteq f'$, then $f' \in \text{LFP}_\tau$.*

However, grounded beliefs are not closed under conjunction. To illustrate this fact, consider the topological argumentation model

$$\mathcal{M} = (\{1, 2, 3\}, E_0 = \{\{1\}, \{2\}, \{3\}, \{1, 2\}, \{2, 3\}\}, \tau = 2^X, \leftarrow, V) \quad (1)$$

with \leftarrow an attack relation in which singletons attack each other and $\{3\} \leftarrow \{1, 2\}$, $\{1\} \leftarrow \{2, 3\}$, $\{2\} \leftarrow \{1, 3\}$ and $\{1, 3\} \leftarrow \{2\}$, and the empty set is attacked by any sets in τ , as shown in Figure 3.1.⁹ Thus, according to the definition, LFP_τ is $\{\{1, 2\}, \{2, 3\}, \{1, 2, 3\}\}$, a set that is not closed under intersection.




Figure 3.1: Grounded beliefs are not closed under conjunction.

To get a better grasp of the whole framework, consider the following scenario which can be modeled by the above topological argumentation model.

Example 3.1. *The zoo in Tom’s town bought a new animal and had it under exhibition. Tom was curious about what the animal is, so he asked his colleagues. However, he got different answers from them. Some told him that the animal was a penguin ($\{1\}$), some told him that the animal was a pterosaur ($\{2\}$) and some told him that the animal was a bat ($\{3\}$). Moreover, two other colleagues he really trusts told him that the animal could fly ($\{2, 3\}$) and the animal was not a mammal ($\{1, 2\}$). After receiving all these*

⁹Attack edges involving the empty set are not drawn.

pieces of information, Tom was very puzzled. After all, it was very hard to imagine that there could be a pterosaur (although genetic technology has developed fast). So in such a situation, Tom believed that the animal could fly and the animal was not a mammal, but he did not believe that the animal was a pterosaur.

The lack of closure under conjunction does not indicate that our agent lacks formal reasoning abilities, as is usually the case in weaker epistemic notions. Indeed, as the topology shows, our agent can put her pieces of evidence together in a meaningful way. In section 5, we further elaborate on this distinction between the agent's reasoning abilities and the closure properties of her grounded beliefs. For now we focus first on the reasons behind the failure of this closure property, by identifying additional conditions under which the property holds.

Proposition 3.2. *Let \mathcal{M} be a topological argumentation model and $\mathcal{A}_\tau = (\tau, \leftarrow)$ its attack graph.*

- *If \leftarrow is transitive (i.e. for any $t_1, t_2, t_3 \in \tau$, $t_1 \leftarrow t_2$ and $t_2 \leftarrow t_3$ imply $t_1 \leftarrow t_3$), then LFP_τ is closed under intersections;*
- *If \leftarrow is unambiguous (i.e. for any $t_1, t_2, t_3 \in \tau$, if $t_1 \leftarrow t_2$ and $t_2 \leftarrow t_3$, then $t_1 \not\leftarrow t_3$ and $t_3 \not\leftarrow t_1$), then LFP_τ is closed under intersections;*
- *If \leftarrow is symmetric (i.e. for any $t_1, t_2 \in \tau$, if $t_1 \leftarrow t_2$ then $t_2 \leftarrow t_1$), then LFP_τ is closed under intersections.*

Note here how when \leftarrow is symmetric, our notion of belief boils down to the evidence-based belief in topological evidence models of [5], and thus it is closed under conjunction. We will elaborate on this point below. With respect to transitivity, one may wonder whether closure under conjunction still holds under a weaker requirement that asks for the transitivity of the attack relation only when the involved sets are in conflict (i.e., for any $t_1, t_2, t_3 \in \tau$, if $t_1 \leftarrow t_2$, $t_2 \leftarrow t_3$ and $t_1 \cap t_3 = \emptyset$, then $t_1 \leftarrow t_3$). We defer the answer to this question till Section 5, where we compare our concept of grounded belief with a probabilistic notion of belief.

We conclude this section with a comparison between our concept of grounded beliefs and the evidence-based beliefs under the topological semantics of [5]. Recall that, in the latter, the agent believes P (denoted by $\text{Bel}P$) if and only if for any $t \in \tau \setminus \{\emptyset\}$ there is $t' \in \tau \setminus \{\emptyset\}$ such that $t' \subseteq t$ and $t' \subseteq P$ (Definition 2.5). Moreover,

Definition 3.4 (Family of Bodies of evidence, Combined Evidence). *Following [5],*

- *let \mathcal{F} denote the family of all bodies of evidence (see Definition 2.2);*
- *let $\mathcal{F}^{\text{finite}}$ denote the family of all finite bodies of evidence;*
- *define E as the family of the combined evidence given by finite bodies of evidence:*

$$E := \left\{ \bigcap F \subseteq X \mid F \in \mathcal{F}^{\text{finite}} \right\}.$$

According to [5, Proposition 2], evidence-based beliefs can be equivalently defined as $\text{Bel}P$ if and only if there is an open that supports P and is also consistent with any combined evidence given by finite bodies of evidence (i.e., if and only if there is $t \in \tau$ such that $t \subseteq P$ and $t \cap e \neq \emptyset$ for any $e \in E$).

Now, given a topological argumentation model \mathcal{M} , let $J_{\mathcal{M}}$ denote the set of opens that are consistent with any combined evidence given by finite bodies, that is

$$J_{\mathcal{M}} := \{t \in \tau \mid \forall e \in E : e \cap t \neq \emptyset\} \quad 10$$

¹⁰In [5], the elements of $J_{\mathcal{M}}$ are called *justifications*.

Note that if \leftarrow in the topological argumentation model \mathcal{M} is symmetric, then $\text{LFP}_\tau = J_\mathcal{M}$.

The following result shows that the notion of grounded belief is weaker than the evidence-based belief of [5].

Proposition 3.3. *Given any topological argumentation model \mathcal{M} , $J_\mathcal{M} \subseteq d_\tau(\emptyset)$.*

It follows from this proposition that $J_\mathcal{M} \subseteq \text{LFP}_\tau$, and thus *BelP* implies *BP* in any given topological argumentation model. Therefore, grounded beliefs are weaker in the sense of being less restrictive and therefore allow more formulas to be believed by the agent. In the next section we study the logic of beliefs grounded on arguments.

4 The Logic of Beliefs Grounded on Arguments (ABBL)

This section studies the logic of the notion of belief grounded on arguments. In order to provide a complete axiomatic characterization of the logic of grounded belief, we work for now with a language that extends the propositional part with a belief operator.

Definition 4.1 (Language). *The language \mathcal{L} of ABBL is generated by the following grammar, where $p \in \text{Prop}$.*

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mathcal{B}\varphi$$

Definition 4.2 (Semantics). *Given a topological argumentation model $\mathcal{M} = (X, E_0, \tau_{E_0}, \leftarrow, V)$ and a possible world $x \in X$, the truth condition of formulas in \mathcal{L} is defined as usual for atomic propositions and Boolean operators. For the grounded belief modality,*

$$\mathcal{M}, x \models \mathcal{B}\varphi \quad \text{iff} \quad \text{there exists } f \in \text{LFP}_\tau \text{ such that } f \subseteq \llbracket \varphi \rrbracket_\mathcal{M}$$

with $\llbracket \varphi \rrbracket_\mathcal{M} := \{x \in X \mid \mathcal{M}, x \models \varphi\}$ the set of φ -worlds in \mathcal{M} .

Axiom System Table 1 shows the axiom system ABBS, which below (Theorem 4.1) will be shown to be sound and complete for the language ABBL with respect to topological argumentation models.

Table 1: Axiom system ABBS for ABBL

Propositional Tautologies and Modus Ponens	
4: $\mathcal{B}\varphi \rightarrow \mathcal{B}\mathcal{B}\varphi$	5: $\neg\mathcal{B}\varphi \rightarrow \mathcal{B}\neg\mathcal{B}\varphi$
RE: from $\varphi \leftrightarrow \psi$ infer $\mathcal{B}\varphi \leftrightarrow \mathcal{B}\psi$	D: $\mathcal{B}\varphi \rightarrow \neg\mathcal{B}\neg\varphi$
M: $\mathcal{B}(\varphi \wedge \psi) \rightarrow \mathcal{B}\varphi \wedge \mathcal{B}\psi$	N: $\mathcal{B}\top$

The axiom system ABBS extends the well-known axiom system EMN (the propositional fragment, the RE rule and axioms M and N) with axioms 4, 5 and D. It is known that EMN is sound and complete with respect to the class of neighborhood models that are *supplemented* (i.e., closed under conjunction elimination or, equivalently, closed under supersets) and contain the unit (i.e., the domain is in the neighborhood) [15]. This suggests a detour for proving the completeness result for our logic with respect to topological argumentation models.

Theorem 4.1. *For any $\varphi \in \mathcal{L}$ and $\Phi \subseteq \mathcal{L}$,*

$$\Phi \vdash_{\text{ABBS}} \varphi \quad \text{if and only if} \quad \Phi \models \varphi$$

where $\Phi \vdash_{\text{ABBS}} \varphi$ indicates that φ is derivable from Φ in the axiom system ABBS and $\Phi \models \varphi$ indicates that, for any topological argumentation model \mathcal{M} and any $x \in X$, if $\mathcal{M}, x \models \Phi$ (i.e. $\mathcal{M}, x \models \phi$ for all $\phi \in \Phi$) then $\mathcal{M}, x \models \varphi$.

Proof sketch. Soundness is straightforward; for completeness, the proof uses a modal equivalence result. First, define a *belief neighborhood model* M as a uniform neighborhood model (X, N_B, V) where the neighborhood function $N_B \subseteq 2^X$ satisfies the following conditions:

- $X \in N_B$ (N_B contains the unit);
- if $b \in N_B$, then $b' \in N_B$ for any b' such that $b \subseteq b'$ (N_B is closed under supersets);
- if $b \in N_B$, then $X \setminus b \notin N_B$ (N_B does not contain the complement of any of its elements).

In such structures, the semantic interpretation of a modality for this *neighborhood-based belief* is given by

$$M, x \models \mathbf{B}\varphi \quad \text{iff} \quad \llbracket \varphi \rrbracket_M \in N_B$$

The following lemma states that every topological argumentation model gives rise to a modally equivalent belief neighborhood model, and vice-versa. Since the axiom system ABBS is sound and complete with respect to belief neighborhood models (see, e.g., [15]), it follows that the system is also sound and complete with respect to topological argumentation models.

Lemma 4.1. *For any topological argumentation model $\mathcal{M} = (X, E_0, \tau, \leftarrow, V)$ there is a belief neighborhood model $M = (X, N_B, V)$ with the same domain and atomic valuation such that \mathcal{M} and M are point-wise modally equivalent with respect to the language \mathcal{L} (i.e. for any $\varphi \in \mathcal{L}$ and any $x \in X$, $\mathcal{M}, x \models \varphi$ if and only if $M, x \models \varphi$.) And vice versa.*

For showing modal equivalence, the key is the following lemma:

Lemma 4.2. $\text{LFP}_\tau = N_B$

Hence, M and \mathcal{M} are modally equivalent. □

Thus, our notion of grounded belief can be completely characterized by the axiom system ABBS.

5 Failure of Closure and Rational Belief

One of the distinctive features of our notion of grounded belief is its lack of closure under conjunction. This feature is also a distinctive mark of the well-known *probabilistic* notion of belief, where belief is interpreted as high enough probability (typically, above threshold $\frac{1}{2}$). It is thus worthwhile to compare the two approaches.

Definition 5.1 (Probabilistic Belief Model [29, 36]). *A probabilistic belief model is a tuple $\mathbf{M} = (X, \mu, V)$ where X is a countable space and $\mu : 2^X \rightarrow [0, 1]$ is a probability function satisfying Kolmogorov's Axioms: (1) $\mu(Q) \geq 0$ for any $Q \subseteq X$, (2) $\mu(X) = 1$, and (3) for any countable sequence of disjoint sets S_1, S_2, \dots , $\mu(\bigcup_{i=1}^{\infty} S_i) = \sum_{i=1}^{\infty} \mu(S_i)$.*

In this probabilistic belief model, μ is interpreted as agent's credence function which measures the agent's credence in each proposition. Hence credible enough propositions can be taken to represent the agent's beliefs, where "credible enough" means "above some threshold". In order to illustrate the relation between our concept of grounded belief and this quantitative counterpart, we consider the threshold to be 0.5. Thus, the concept of *probabilistic belief* that we consider in this paper is defined as

$$\mathbf{M}, x \models \mathbf{B}\varphi \quad \text{iff} \quad \mu(\llbracket \varphi \rrbracket_{\mathbf{M}}) > 0.5$$

It is important to note that the axiom system ABBS is sound with respect to the probabilistic belief model. Moreover, for any probabilistic belief model we can construct a modally equivalent belief neighborhood model (see the proof of Theorem 4.1): it is enough to put in the neighbourhood N_B all subsets of the domain whose probability is greater than 0.5, i.e. $N_B = \{S \subseteq X \mid \mu(S) > 0.5\}$. Given Kolmogorov's axioms, the resulting structure indeed satisfies the belief neighborhood requirements. Even more: by Lemma 4.1, we can build a topological argumentation model from the belief neighborhood model, which is modally equivalent to the belief neighborhood model and thus modally equivalent to the probabilistic belief model. However, not every topological argumentation model has a modally equivalent probabilistic belief model. This is because the axiom system ABBS is sound but not complete with respect to the probabilistic belief model (c.f. [27, Section 3.1], [29] and [36]).

Thus, probabilistic belief can be seen as a special form of grounded belief, in the sense that the former can be defined through a special attack relation on 2^X : for any $e, e' \in 2^X$, define $e \leftarrow_{\mu} e'$ if and only if $e \cap e' = \emptyset$ and $\mu(e) \leq \mu(e')$. This attack relation satisfies extra conditions, for example, for any $e_1, e_2, e_3 \in 2^X$, if $e_1 \leftarrow e_2$, $e_2 \leftarrow e_3$ and $e_1 \cap e_3 = \emptyset$, then $e_1 \leftarrow e_3$. Since the probabilistic notion of belief is not closed under conjunction either, this extra condition does not save the failure of closure under conjunction, which answers the question we posed in Section 2.

Philosophical Debate. Let us re-examine one of the most important criticisms faced by probabilistic beliefs from the perspective of our topological argumentation models. The criticism is based on the *conjunction rule* which states that if it is 'rational' for an agent to believe φ and it is 'rational' to believe ψ , then it is 'rational' to believe φ and ψ . As stated in [20],

there are [...] prima facie worries associated with rejecting the rule, the most fundamental of which is that if we are not required on pains of irrationality to believe the conjunction of propositions that we rationally believe, we might seem to lose some of our most powerful argumentative and deliberative tools. (p.42)

Thus, according to this (which [20] actually criticizes), a rational agent should follow the conjunction rule. Since probabilistic beliefs violate the rule, philosophers have debated whether probabilistic beliefs are good candidates for representing a rational concept of belief. In this discussion it is essential to note that if "rational" is viewed as the willingness and capability of following logical rules, it implicitly burdens belief with a requirement about the agent's ability of deductive reasoning.¹¹ In [20], the author argues not only that this worry is misplaced, but also that beliefs should not be burdened with the agents' deductive reasoning abilities. Note that Foley proposes a clear distinction between two reasoning processes, one involving an agent's beliefs (not closed under conjunction) and deductive processes (which are closed under conjunction). What our proposal shows is that the two processes are distinct but can be united in one formal framework.

The topological argumentation model shows how the notion of grounded belief is based on the agents' topology of evidence and the attack relation on the topology. So on the one hand, the agent modelled by a topological argumentation model is a powerful logician: she can generate a full topology from her collections of pieces of evidence. On the other hand, this logician's grounded belief does not follow the conjunction rule. So even if the agent has a good argument/reason to believe φ and a good argument/reason to believe ψ , she may not have a good argument/reason to believe $\varphi \wedge \psi$. In contrast to many probabilistic belief models where these parts are hidden, our setting makes the different processes explicit.

¹¹Our discussion does not exclude the possibility that 'rationality' may well require more than simply reasoning in a logical way. But if extra constraints are imposed the formal setting is bound to change as well.

6 Conclusion and Future Work

This paper studies a notion of grounded belief. It does so by applying ideas from both abstract argumentation theory and the topological semantics for evidence.

From the perspective of abstract argumentation theory, our work can be seen as an attempt of giving structure to the otherwise abstract arguments in an attack graph. Here arguments are instantiated as sets of possible worlds, and thus we can use logical tools to define not only which propositions they support, but also when one of them attacks another. From the perspective of the topological semantics for evidence, our work provides a generalization of the conflict relation between pieces of evidence. The integration of these ingredients from two different fields brings to fruition a notion of *belief grounded on arguments* whose logical behavior is proved to be characterized by the axiom system ABBS. This notion is compared with other forms of belief, including the evidence-based belief of [5] and the probabilistic notion of belief in [29] and [36].

One of the most intriguing features of the introduced concept of grounded belief is its failure of closure under conjunction. In the philosophical literature, this failure is widely taken as a failure of rationality, especially when it comes down to the agent's beliefs. However, as our setting shows, even if the conjunction rule is adopted by the agent on the level of *reasoning* about evidence or arguments, it does not need to be adopted to form new *beliefs*.

In future work we extend our formal language to include operators that address the agent's arguments explicitly. While our current restricted language allows us to characterize the behavior of our notion of belief, it does lack the expressive power to talk about the relationship between beliefs and the arguments supporting it. Another dimension worthwhile to explore further is the understanding of argumentation not as a single agent's inner reasoning mechanism, but rather as a social process that involves several agents. There exists work in argumentation theory dedicated to this dimension, as the study of methods for merging different attack graphs [16] shows. This line of work can further be related to the recent developments in social choice theory [19]. In our context, the question of merging different agent's attack graphs can shed new light at group-related notions, in particular, distributed belief, which can then be understood from the perspective of abstract argumentation theory.

Finally, while we have restricted our attention to the concept of grounded belief, epistemologists will be interested to see a connection to the seminal studies of knowledge, belief, justifications and their relation (e.g., [28], where knowledge is understood as belief with a correct justification that cannot be defeated by any further true justification). There exist new proposals in this direction (e.g., [5]), which, with the help of argumentation theory can shed more light on different epistemic theories and, in particular, on the different notions of knowledge.

Acknowledgments We thank Johan van Benthem and Abyüke Özgün for feedback and discussions on the topic of this paper. We also thank the anonymous referees of TARK 2017 for their useful comments.

Appendix

PROOF OF PROPOSITION 3.1 Take any $f \in \text{LFP}_\tau$ and any $f' \in \tau$ such that $f \subseteq f'$. Suppose no one attacks f' ; then we are done as, by LFP_τ 's definition, every non-attacked element of τ should be in LFP_τ . Suppose otherwise, and let t be one of such opens attacking f' ; it is enough to find a $f'' \in \text{LFP}_\tau$ attacking t , as then f' would be defended by someone in LFP_τ and thus, by definition, f' would be in LFP_τ . Now, since t attacks f' , it should also attack the stronger f (as required by \leftarrow 's definition);

but then, since f is in LFP_τ , it should be defended by someone in LFP_τ , that is, there is a f'' in LFP_τ attacking t . This completes the proof.

PROOF OF PROPOSITION 3.2 The following lemma will be useful.

Lemma .1. *Let $\mathcal{M} = (X, E_0, \tau_{E_0}, \leftarrow, V)$ be a topological argumentation model. Then, for any $f_1, f_2 \in \text{LFP}_\tau$, $f_1, f_2 \in \text{LFP}_\tau$ implies $f_1 \cap f_2 \in \text{LFP}_\tau$ if and only if for any $t \in \tau$, if $f_1 \cap f_2 \leftarrow t$, then $t \cap f = \emptyset$ for some $f \in \text{LFP}_\tau$.*

Proof. (\Rightarrow) From left to right, take arbitrary $f_1, f_2 \in \text{LFP}_\tau$. Suppose there is an open $t \in \tau$ such that t attacks $f_1 \cap f_2$ but is not in conflict with anybody in LFP_τ . From the latter it follows that nobody in LFP_τ attacks t , and thus the attacked $f_1 \cap f_2$ is not defended by LFP_τ ; therefore, $f_1 \cap f_2$ is not in LFP_τ . (\Leftarrow) From right to left, take arbitrary $f_1, f_2 \in \text{LFP}_\tau$. If there is no $t \in \tau$ such that $f_1 \cap f_2 \leftarrow t$, then we are done. Otherwise, take arbitrary $t \in \tau$ such that $f_1 \cap f_2 \leftarrow t$. It implies that there is $f' \in \text{LFP}_\tau$ such that $t \cap f' = \emptyset$. Thus either f' attacks t or else t attacks f' . The former case implies that there is $f'' \in \text{LFP}_\tau$ such that $t \leftarrow f''$ by virtue of f'' 's membership in LFP_τ ; together with the latter case, i.e. $t \leftarrow f'$, we can conclude that there is $f \in \text{LFP}_\tau$ such that $t \leftarrow f$. Hence for any $t \in \tau$ such that $f_1 \cap f_2 \leftarrow t$, there is $f \in \text{LFP}_\tau$ such that $t \leftarrow f$, which implies that $f_1 \cap f_2 \in \text{LFP}_\tau$. \square

Now, for proposition 3.2. The proof for transitivity and the relation being unambiguous, proceeds by contraposition, so take any $f_1, f_2 \in \text{LFP}_\tau$ such that $f_1 \cap f_2$ is not in LFP_τ . Then, by Lemma .1, there is an open $t \in \tau$ who attacks $f_1 \cap f_2$ (i.e., $f_1 \cap f_2 \leftarrow t$) and who is not in conflict with elements of LFP_τ (i.e., $f \in \text{LFP}_\tau$ implies $t \cap f \neq \emptyset$). The goal is to show that \leftarrow is not transitive.

Define t_1, t_2 and t_3 as

$$t_1 := f_1 \cap t, \quad t_2 := f_2 \cap t, \quad t_3 := f_1 \cap f_2,$$

and note that none of them are empty. Note also that due to the fact that t attacks $f_1 \cap f_2$ ($f_1 \cap f_2 \leftarrow t$), t must be in conflict with $f_1 \cap f_2$ ($(f_1 \cap f_2) \cap t = \emptyset$); hence, $t_1 \cap t_2 = t_2 \cap t_3 = t_3 \cap t_1 = \emptyset$.

For transitivity, consider now the following two mutually exclusive and collectively exhaustive cases: (1) $t_1 \leftarrow t_3$ or $t_2 \leftarrow t_3$, (2) $t_1 \not\leftarrow t_3$ and $t_2 \not\leftarrow t_3$. In the first, assume the leftmost disjunct $t_1 \leftarrow t_3$; that, together with $t_3 \leftarrow t$ (recall: $t_3 = f_1 \cap f_2$) and $t_1 \not\leftarrow t$ (as $t_1 \cap t = f_1 \cap t \neq \emptyset$) shows that \leftarrow is not transitive. If the rightmost disjunct $t_2 \leftarrow t_3$ is assumed, a similar reasoning yields $t_3 \leftarrow t$ and $t_2 \not\leftarrow t$, and thus lack of transitivity again.

Now the second case, which implies $t_3 \leftarrow t_1$ and $t_3 \leftarrow t_2$, as the t_i are all in conflict. Since t_1 and f_2 are in conflict, at least one of them should attack the other. If $t_1 \leftarrow f_2$, then this together with $t_3 \leftarrow t_1$ and $t_3 \not\leftarrow f_2$ (as $t_3 \cap t = f_1 \cap f_2 \neq \emptyset$, and thus they are not in conflict) shows that \leftarrow is not transitive. Otherwise, $f_2 \leftarrow t_1$ should be the case; hence, since f_2 is in LFP_τ , it should be defended from t_1 by some $f \in \text{LFP}_\tau$, that is, $t_1 \leftarrow f$. But then we have $f_2 \leftarrow t_1, t_1 \leftarrow f$ and $f_2 \not\leftarrow f$ (since LFP_τ is conflict-free). Thus, again, \leftarrow is not transitive.

For the relation being unambiguous, consider two cases, $t_1 \leftarrow t_2$ or $t_2 \leftarrow t_1$. In the case of $t_1 \leftarrow t_2$, if $t_2 \leftarrow t_3$, then no matter $t_1 \leftarrow t_3$ or $t_3 \leftarrow t_1$, it is not unambiguous. If $t_3 \leftarrow t_2$, then no matter $t_1 \leftarrow t_3$ or $t_3 \leftarrow t_1$, \leftarrow is not unambiguous. The proof for the case of $t_2 \leftarrow t_1$ follows a similar argument. Therefore, we reach the conclusion that \leftarrow cannot be unambiguous.

For symmetry, assume that \leftarrow is symmetric. Observe that $\text{LFP}_\tau = \{t \in \tau \mid \forall e \in E : e \cap t \neq \emptyset\}$, which is closed under conjunction.

PROOF OF PROPOSITION 3.3 Take any t in $J_{\mathcal{M}}$. Since t is consistent with any $e \in E$, it is also consistent with any open $t' \in \tau$, that is, for any such t' we have $t \cap t' \neq \emptyset$. But then t is not attacked at all, and thus it is defended by the empty set, i.e. $t \in d_{\tau}(\emptyset)$. Therefore, $J_{\mathcal{M}} \subseteq d_{\tau}(\emptyset)$.

PROOF OF LEMMA 4.1 Given a topological argumentation model, it is easy to build a point-wise modally equivalent belief neighborhood model: it is enough to define the neighborhood function N_B as LFP_{τ} plus all its element's supersets ($N_B := \{b \in 2^X \mid f \subseteq b \text{ for some } f \in \text{LFP}_{\tau}\}$). For the other direction, given a belief neighborhood model $M = (X, N_B, V)$, let $\mathcal{M}_M = (X, E_0, \tau, \leftarrow_{N_B}, V)$ be the topological argumentation model that shares domain and atomic valuation with M , and in which the family of pieces of evidence is given by the singletons in X ($E_0 := \{\{x\} \in 2^X \mid x \in X\}$) and therefore the generated topology is the power set of the domain ($\tau = 2^X$). Moreover, define the attack relation \leftarrow_{N_B} as

$$t \leftarrow_{N_B} t' \quad \text{iff} \quad \begin{cases} t \cap t' = \emptyset \text{ and } t \notin N_B & \text{when } t' \neq \emptyset \\ t = \emptyset & \text{when } t' = \emptyset \end{cases}$$

for every $t, t' \in \tau$, so a non-empty t' attacks a non-empty t if and only if they are in conflict and t is not in N_B , and while the empty set does not attack non-empty sets, it is attacked by everybody (including itself). It is easy to verify that this model is a topological argumentation model.

PROOF OF LEMMA 4.2 (\supseteq) Take any $b \in N_B$. By \leftarrow_{N_B} 's definition, no element of τ who is in conflict with b attacks it (i.e., there is no $t \in \tau$ such that $b \cap t = \emptyset$ and $b \leftarrow_{N_B} t$); therefore, no one attacks b and hence $b \in \text{LFP}_{\tau}$.

(\subseteq) Take an arbitrary $t \in \tau$ such that $t \notin N_B$. If t is not in conflict with any open in τ , then t must be X , so $X \notin N_B$, contradicting the fact that N_B contains the unit. Hence, t should have a conflict with some open, that is, there is at least one $t' \in \tau$ such that $t \cap t' = \emptyset$. From this and \leftarrow_{N_B} 's definition it follows that t' attacks t , $t \leftarrow_{N_B} t'$.

If at least one of such t' is in N_B , then it is also in LFP_{τ} (the (\supseteq) case above), and thus the fact that t' attacks t implies that $t \notin \text{LFP}_{\tau}$, as this set is conflict-free. Then we are done. Otherwise no such t' is in N_B , that is, every open in conflict with t is outside N_B ; in other words, no element of N_B is in conflict with t ($b \in N_B$ implies $t \cap b \neq \emptyset$).

Define now the set s as the union of the non- t parts of every element of N_B , that is, $s := \bigcup_{b \in N_B} (b \setminus t)$. Since τ is the full power set, s is in τ . This s is in conflict with t ($s \cap t = \emptyset$, by construction), but it is consistent with all $b \in N_B$ ($b \cap s \neq \emptyset$ for all such b , as otherwise the $b' \in N_B$ in conflict with s would have been a subset of t , and thus closure under supersets of N_B would imply $t \in N_B$, contradicting our initial assumption). These two facts tells us that we have found an open $s \in \tau$ that attacks t ($t \leftarrow_{N_B} s$, from the first fact, $t \notin N_B$ and \leftarrow_{N_B} 's definition) but it is not attacked by anyone in N_B ($b \in N_B$ implies $s \not\leftarrow_{N_B} b$, from the second fact and \leftarrow_{N_B} 's definition). Moreover, it attacks every open not in N_B with whom it has conflict ($x \notin N_B$ implies $x \cap s = \emptyset$, again from \leftarrow_{N_B} 's definition). Thus, s (1) attacks everybody in τ that is not in N_B , and (2) is not attacked by anybody in N_B ; hence no one outside N_B can be protected by N_B against s , which means that $d_{\tau}(N_B) = N_B$. Together with the fact that $N_B \subseteq \text{LFP}_{\tau}$, it implies that $N_B = \text{LFP}_{\tau}$. In particular, since $t \notin N_B$, we get the required $t \notin \text{LFP}_{\tau}$.

References

- [1] Sergei N. Artemov (2008): *The Logic of Justification*. *The Review of Symbolic Logic* 1(4), pp. 477–513, doi:10.1017/S1755020308090060.

- [2] Sergei N. Artemov & Elena Nogina (2005): *Introducing Justification into Epistemic Logic*. *Journal of Logic and Computation* 15(6), pp. 1059–1073, doi:10.1093/logcom/exi053.
- [3] Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün & Sonja Smets (2013): *The Topology of Belief, Belief Revision and Defeasible Knowledge*. In Davide Grossi, Olivier Roy & Huaxin Huang, editors: *Logic, Rationality, and Interaction: 4th International Workshop, LORI 2013, Hangzhou, China, October 9-12, 2013, Proceedings*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 27–40, doi:10.1007/978-3-642-40948-6_3.
- [4] Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün & Sonja Smets (2015): *The topological theory of belief*. Technical Report PP-2015-18, ILLC, University of Amsterdam. Available at <https://www.illc.uva.nl/Research/Publications/Reports/PP-2015-18.text.pdf>.
- [5] Alexandru Baltag, Nick Bezhanishvili, Aybüke Özgün & Sonja Smets (2016): *Justified Belief and the Topology of Evidence*. In Jouko Väänänen, Åsa Hirvonen & Ruy de Queiroz, editors: *Logic, Language, Information, and Computation: 23rd International Workshop, WoLLIC 2016, Puebla, Mexico, August 16-19th, 2016. Proceedings*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 83–103, doi:10.1007/978-3-662-52921-8_6.
- [6] Alexandru Baltag, Virginie Fiutek & Sonja Smets (2016): *Beliefs and Evidence in Justification Models*. In S. Demri L. Beklemishev & A. Mate, editors: *Proceedings of Advances in Modal Logic*, College Publications, Oxford, UK, pp. 156–176. Available at <http://www.aiml.net/volumes/volume11/Baltag-Fiutek-Smets.pdf>.
- [7] Alexandru Baltag, Bryan Renne & Sonja Smets (2012): *The Logic of Justified Belief Change, Soft Evidence and Defeasible Knowledge*. In C.-H. Luke Ong & Ruy J. G. B. de Queiroz, editors: *Logic, Language, Information and Computation - 19th International Workshop, WoLLIC 2012, Buenos Aires, Argentina, September 3-6, 2012. Proceedings, Lecture Notes in Computer Science 7456*, Springer, pp. 168–190, doi:10.1007/978-3-642-32621-9_13.
- [8] Alexandru Baltag, Bryan Renne & Sonja Smets (2014): *The Logic of Justified Belief, Explicit Knowledge and Conclusive Evidence*. *Annals of Pure and Applied Logic* 165(1), pp. 49–81, doi:10.1016/j.apal.2013.07.005.
- [9] Alexandru Baltag & Sonja Smets (2008): *Probabilistic dynamic belief revision*. *Synthese* 165(2), p. 179, doi:10.1007/s11229-008-9369-8.
- [10] Alexandru Baltag & Sonja Smets (2008): *A qualitative theory of dynamic interactive belief revision*. *Texts in logic and games* 3, pp. 9–58. Available at <http://www.vub.ac.be/CLWF/SS/chapter.pdf>.
- [11] Johan van Benthem, David Fernández Duque & Eric Pacuit (2014): *Evidence and plausibility in neighborhood structures*. *Ann. Pure Appl. Logic* 165(1), pp. 106–133, doi:10.1016/j.apal.2013.07.007.
- [12] Johan van Benthem & Eric Pacuit (2011): *Dynamic Logics of Evidence-Based Beliefs*. *Studia Logica* 99(1-3), pp. 61–92, doi:10.1007/s11225-011-9347-x.
- [13] Oliver Board (2004): *Dynamic interactive epistemology*. *Games and Economic Behaviour* 49(1), pp. 49–80, doi:10.1016/j.geb.2003.10.006.
- [14] Martin W. A. Caminada & Dov M. Gabbay (2009): *A Logical Account of Formal Argumentation*. *Studia Logica* 93(2-3), p. 109, doi:10.1007/s11225-009-9218-x.
- [15] Brian F. Chellas (1980): *Modal Logic: An Introduction*. Cambridge University Press, Cambridge Eng. ; New York, doi:10.1017/CBO9780511621192.
- [16] Sylvie Coste-Marquis, Caroline Devred, Sébastien Konieczny, Marie-Christine Lagasquie-Schiex & Pierre Marquis (2007): *On the Merging of Dung’s Argumentation Systems*. *Artificial Intelligence* 171(10), pp. 730–753, doi:10.1016/j.artint.2007.04.012.
- [17] Jon Doyle (1979): *A Truth Maintenance System*. *Artificial Intelligence* 12(3), pp. 231–272, doi:10.1016/0004-3702(79)90008-0.
- [18] Phan Minh Dung (1995): *On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games*. *Artificial intelligence* 77, pp. 321–357, doi:10.1016/0004-3702(94)00041-X.

- [19] Ulle Endriss & Umberto Grandi (2017): *Graph Aggregation*. *Artificial Intelligence* 245, pp. 86–114, doi:10.1016/j.artint.2017.01.001.
- [20] Richard Foley (2009): *Beliefs, Degrees of Belief, and the Lockean Thesis*. In Franz Huber & Christoph Schmidt-Petri, editors: *Degrees of Belief, Synthese Library* 342, Springer Netherlands, pp. 37–47, doi:10.1007/978-1-4020-9198-8_2.
- [21] Davide Grossi (2013): *Abstract argument games via modal logic*. *Synthese* 190, pp. 5–29, doi:10.1007/s11229-012-0237-1.
- [22] Davide Grossi & Wiebe van der Hoek (2014): *Justified Beliefs by Justified Arguments*. In: *Proceedings of the Fourteenth International Conference on Principles of Knowledge Representation and Reasoning*, AAAI Press, pp. 131–140. Available at <http://www.aaai.org/ocs/index.php/KR/KR14/paper/view/7997>.
- [23] Jaakko Hintikka (1962): *Knowledge and Belief*. Cornell University Press.
- [24] Kevin T. Kelly (1996): *The logic of reliable inquiry*. Oxford University Press.
- [25] Thomas Kelly (2016): *Evidence*. In Edward N. Zalta, editor: *The Stanford Encyclopedia of Philosophy*, winter 2016 edition, Metaphysics Research Lab, Stanford University. Available at <https://plato.stanford.edu/archives/win2016/entries/evidence/>.
- [26] Johan de Kleer (1986): *An Assumption-Based TMS*. *Artificial Intelligence* 28(2), pp. 127–162, doi:10.1016/0004-3702(86)90080-9.
- [27] Dominik Klein, Norbert Gratzl & Olivier Roy (2015): *Introspection, Normality and Agglomeration*. In: *Logic, Rationality, and Interaction*, Springer, Berlin, Heidelberg, pp. 195–206, doi:10.1007/978-3-662-48561-3_16.
- [28] Keith Lehrer (1990): *Theory of Knowledge*, 1st edition. Routledge.
- [29] Wolfgang Lenzen (1980): *Glauben, Wissen und Wahrscheinlichkeit: Systeme der epistemischen Logik*. Springer, doi:10.1007/978-3-7091-8594-0. Google-Books-ID: 6mwIAQAIAAJ.
- [30] John-Jules Ch. Meyer & Wiebe van Der Hoek (1995): *Epistemic Logic for AI and Computer Science*. Cambridge University Press, New York, N.Y., U.S.A., doi:10.1017/CBO9780511569852.
- [31] Richard Montague (1970): *Universal Grammar*. *Theoria* 36(3), pp. 373–398, doi:10.1111/j.1755-2567.1970.tb00434.x.
- [32] George Pappas (2014): *Internalist vs. Externalist Conceptions of Epistemic Justification*. In Edward N. Zalta, editor: *The Stanford Encyclopedia of Philosophy*, fall 2014 edition, Metaphysics Research Lab, Stanford University. Available at <https://plato.stanford.edu/archives/fall2014/entries/justep-intext/>.
- [33] François Schwarzentruber, Srdjan Vesic & Tjitze Rienstra (2012): *Building an Epistemic Logic for Argumentation*. In Luis Fariñas del Cerro, Andreas Herzig & Jérôme Mengin, editors: *Logics in Artificial Intelligence - 13th European Conference, JELIA 2012, Toulouse, France, September 26-28, 2012. Proceedings, Lecture Notes in Computer Science* 7519, Springer, pp. 359–371, doi:10.1007/978-3-642-33353-8_28.
- [34] Dana Scott (1970): *Advice on modal logic*. In Karel Lambert, editor: *Philosophical Problems in Logic*, Reidel, Dordrecht, The Netherlands, pp. 143–173, doi:10.1007/978-94-010-3272-8_7.
- [35] Wolfgang Spohn (1988): *Ordinal conditional functions: A dynamic theory of epistemic states*. In William Leonard Harper & Brian Skyrms, editors: *Causation in decision, belief change, and statistics: Proceedings of the Irvine Conference on Probability and Causation, The Western Ontario Series in Philosophy of Science* 42, Kluwer, Dordrecht, pp. 105–134, doi:10.1007/978-94-009-2865-7_6.
- [36] Jan van Eijck & Bryan Renne (2014): *Belief as Willingness to Bet*. *arXiv:1412.5090 [cs]*. Available at <https://arxiv.org/abs/1412.5090>.
- [37] Steven Vickers (1989): *Topology via Logic*. Cambridge University Press, Cambridge, UK.