



HAL
open science

Stratified evidence logics

Philippe Balbiani, David Fernández-Duque, Andreas Herzig, Emiliano Lorini

► **To cite this version:**

Philippe Balbiani, David Fernández-Duque, Andreas Herzig, Emiliano Lorini. Stratified evidence logics. 28th International Joint Conference on Artificial Intelligence (IJCAI 2019), Aug 2019, Macao, China. pp.1523-1529, <10.24963/ijcai.2019/211>. <hal-02411254>

HAL Id: hal-02411254

<https://hal.science/hal-02411254v1>

Submitted on 14 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Stratified evidence logics

Philippe Balbiani¹, David Fernández-Duque², Andreas Herzig¹ and Emiliano Lorini¹

¹Institut de Recherche en Informatique de Toulouse, Toulouse University ²Ghent University
{philippe.balbani, andreas.herzig, emiliano.lorini}@irit.fr, david.fernandezduque@ugent.be

Abstract

Evidence logics model agents’ belief revision process as they incorporate and aggregate information obtained from multiple sources. This information is captured using *neighbourhood structures*, where individual neighbourhoods represent pieces of evidence. Nevertheless, these formalisms represent evidence aggregation qualitatively, and do not allow one to quantify the amount of resources (i.e., pieces of evidence) required to reach a specific belief. In this paper we propose an extended framework which allows such information to be recorded, provide a complete deductive calculus and a proof of decidability, and show how existing frameworks can be embedded into ours.

1 Introduction

Evidence logics were first proposed by van Benthem and Pacuit [van Benthem and Pacuit, 2011a; van Benthem and Pacuit, 2011b] and then further developed by van Benthem et al. [van Benthem et al., 2012; van Benthem et al., 2014] and Baltag et al. [Baltag et al., 2016a; Baltag and Occhipinti Liberman, 2017]. More recently, they were extended to the multi-agent setting by [Liu and Lorini, 2017]. They aim to model belief formation in settings where an agent has access to different pieces of evidence. Evidence is modelled as a family of sets of possible states of the world, which the agent must then combine in various ways in order to form her beliefs. This is similar to the conception of evidence held by Shafer’s theory of uncertainty [Shafer, 1976] as well as by theories of information fusion [Benferhat et al., 1993; Benferhat et al., 1995; Dubois et al., 2016] that conceive belief as the result of aggregating incomplete or uncertain information coming from various sources of evidence. It is also connected with justification logic by [Artemov, 2008; Fitting, 2005] in which evidence is expressed as a term, and possible manipulations of evidence are conceived as operations over terms.

Evidence logics can be applied both to situations where an agent receives information from possibly fallible sources such as news outlets, as well as to situations where an agent performs experiments to obtain data of varying precision about the state of the world, as in e.g. physics. Their semantics

exploit neighbourhood models that have also been used to model *explicit*, as opposed to *implicit*, belief [Balbiani et al., 2016; Balbiani et al., 2018; Velázquez-Quesada, 2013].

The logics considered so far provide a qualitative account of belief formation: either an agent has evidence for a certain fact, or she does not. Our aim is to introduce a quantitative element to these frameworks resulting in a new family of logics called *stratified evidence logics*. This may be used to measure the number of distinct sources needed to pool enough evidence to justify a certain belief, or else be interpreted as a measure of effort, say in performing precise measurements. As we want to allow for the possibility of considering arbitrary but finite amount of data (or, perhaps, countably many pieces of information), we will model this effort using cardinal numbers: $[1]\varphi$ means that the belief φ can be justified using up to one piece of evidence, $[2]\varphi$ that φ can be justified using up to two pieces of evidence, and $[\omega]\varphi$ that φ can be justified using some finite (but arbitrarily large) amount of pieces of evidence. The idea of adding a quantitative parameter to epistemic modalities goes back to [van der Hoek and Meyer, 1991] in which graded epistemic logics are studied. More recently, it was further developed by [Naumov and Tao, 2015] who introduce modal operators capturing the “cost” for an agent of justifying a certain conclusion.

We believe that stratified evidence logics provide a powerful and useful alternative to existing logics of evidence suitable for applications in artificial intelligence (AI), where an artificial agent such as a chatbot or a conversational agent is expected to gather information, to filter it and to transfer the selected information to the human user. Such agents may be endowed with the capability of forming beliefs on the basis of the collected evidences. By way of example, consider a chatbot connected to the Internet who has to provide information to the human user about the quality of a certain movie. The chatbot has access to different recommendation systems about movies (e.g., Netflix, Rotten Tomatoes, IMDb). Stratified evidence logics allow us to identify the number of recommendations that are sufficient for the chatbot to form an opinion about the quality of the movie and to be able to inform the human about this.

As we will show, existing frameworks naturally embed into our own (with some exceptions). Moreover we show that stratified evidence logic has a natural axiomatization and is decidable.

2 Preliminaries

We will need a few basic facts from cardinal arithmetic, which we briefly review. The cardinality of a set X is denoted $|X|$, and recall that cardinal addition is defined so that $|X| + |Y| = |X \cup Y|$ whenever X, Y are disjoint; $+$ will always denote cardinal addition. Cardinal addition is associative and commutative, and coincides with standard addition on natural numbers; on the other hand, if either α or β is infinite, then $\alpha + \beta = \max\{\alpha, \beta\}$. We use ω to denote the cardinality of \mathbb{N} (i.e., of the set of natural numbers), and ω_1 to denote the first uncountable cardinal. We also use ∞ as a formal symbol which is decreed to be greater than any cardinal (and hence not itself an actual cardinal).

Given a set A and a cardinal λ , we use the notation $\binom{A}{\lambda}$ for the set of subsets of A with cardinality λ (for example, $\binom{A}{0} = \{\emptyset\}$, $\binom{A}{1}$ is the set containing all singletons $\{a\}$ with $a \in A$, $\binom{A}{\omega}$ is the set of countably infinite subsets of A). A *cardinal degree set* is a set Ω whose elements are all cardinals and such that $0 \in \Omega$. We assume that Ω is countable and comes equipped with a computable notation system such that the relation $\alpha + \beta \leq \gamma$ is decidable.¹ Typically we take $\Omega = \mathbb{N}$ or $\Omega = \mathbb{N} \cup \{\omega\}$.

3 Syntax and axiomatics

Our formal language will be parametrized by a cardinal degree set, Ω . Fix a countably infinite set \mathbb{P} of propositional variables. We define a language \mathcal{L}_Ω by the following grammar:

$$\varphi, \psi ::= \top \mid p \mid (\varphi \wedge \psi) \mid \neg\varphi \mid [\alpha]\varphi,$$

where $p \in \mathbb{P}$ and $\alpha \in \Omega$. Other Boolean connectives are defined by the usual abbreviations, $\langle \alpha \rangle \varphi$ is defined by $\neg[\alpha]\neg\varphi$, and we follow the standard rules for elimination of parentheses. We read $[\alpha]\varphi$ as “ φ can be justified using α pieces of evidence” and $\langle \alpha \rangle \varphi$ as “ φ cannot be refuted using α pieces of evidence”.

We define the *stratified evidence logic* SEL_Ω by the following axioms and rules, where the nomenclature follows the modal logic literature.

(Taut) : All classical propositional tautologies

(Mon) : $[\alpha]\varphi \rightarrow [\beta]\varphi$ (if $\alpha < \beta$)

(Dist₊) : $[\alpha]\varphi \wedge [\beta]\psi \rightarrow [\gamma](\varphi \wedge \psi)$ (if $\alpha + \beta \leq \gamma$)

(T) : $[\alpha]\varphi \rightarrow \varphi$

(4) : $[\alpha]\varphi \rightarrow [\alpha][\alpha]\varphi$

(5₀) : $\langle 0 \rangle \varphi \rightarrow [0]\langle 0 \rangle \varphi$

(MP) : $\frac{\varphi \quad \varphi \rightarrow \psi}{\psi}$ (RM) : $\frac{\varphi \rightarrow \psi}{[\alpha]\varphi \rightarrow [\alpha]\psi}$

(N₀) : $\frac{\varphi}{[0]\varphi}$

¹For example, one can enumerate cardinals using a standard ordinal notation system [Miller, 1976].

As we will see, these axioms and rules are sound and complete for the class of *evidence models*, defined in the following section. First let us exhibit a useful family of derivable formulas.

Lemma 3.1. *If $\alpha + \beta \leq \gamma \in \Omega$, then $[\alpha]\varphi \wedge [\beta]\psi \rightarrow [\gamma]([\alpha]\varphi \wedge [\beta]\psi)$ is derivable. More generally, if $\alpha_1 + \dots + \alpha_n \leq \alpha \in \Omega$, then*

$$([\alpha_1]\varphi_1 \wedge \dots \wedge [\alpha_n]\varphi_n) \rightarrow [\alpha]([\alpha_1]\varphi_1 \wedge \dots \wedge [\alpha_n]\varphi_n)$$

is derivable.

Proof. By 4 we have that $\vdash [\alpha]\varphi \wedge [\beta]\psi \rightarrow [\alpha][\alpha]\varphi \wedge [\beta][\beta]\psi$, while by Dist₊ we have that $\vdash [\alpha][\alpha]\varphi \wedge [\beta][\beta]\psi \rightarrow [\gamma]([\alpha]\varphi \wedge [\beta]\psi)$. The second claim follows from similar reasoning by induction on n . \square

4 Evidence semantics

Our models are based on evidence models as in e.g. [van Benthem *et al.*, 2012]. However, since our language allows us to represent quantitative, as well as qualitative, information about evidence, it is convenient to encode such information into our models. Thus we arrive at the notion of *stratified evidence frames*.

Definition 4.1. *Fix a cardinal degree set Ω . A stratified evidence pre-frame is a tuple $\mathcal{F} = (W, \mathcal{E})$, where W is a non-empty set of worlds and $\mathcal{E} \subseteq \Omega \times (2^W \setminus \{\emptyset\})$. We call \mathcal{E} an evidence hierarchy. For $\alpha \in \Omega$ and $\mathcal{H} \subseteq \mathcal{E}$, we write*

$$\mathcal{H}_\alpha = \{X \subseteq W : (\alpha, X) \in \mathcal{H}\}.$$

We say that \mathcal{F} is a stratified evidence frame if $\mathcal{E}_0 = \{W\}$ and, whenever $(\alpha_1, X_1), \dots, (\alpha_n, X_n) \in \mathcal{E}$, $\gamma \geq \alpha_1 + \dots + \alpha_n$ and $X := X_1 \cap \dots \cap X_n \neq \emptyset$, we have $(\gamma, X) \in \mathcal{E}$.

We say that \mathcal{F} is a strict stratified evidence frame if it is a stratified evidence frame and for all $X \subseteq W$ and cardinals α , $X \in \mathcal{E}_\alpha$ if and only if $X \neq \emptyset$ and there are $\beta < \alpha + 1$ and $\mathcal{X} \in \binom{\mathcal{E}_\beta}{1}$ such that $X = \bigcap \mathcal{X}$. In this case, we will identify \mathcal{F} with (W, \mathcal{E}_1) .

As usual, (strict) stratified evidence models are frames with a valuation $V : \mathbb{P} \rightarrow 2^W$.

The inequality $\beta < \alpha + 1$ is a succinct way of writing $\beta \leq \alpha$ if α is finite, $\beta < \alpha$ otherwise. The reason for this choice is that, in a strict evidence frame, for $n < \omega$ we want $X \in \mathcal{E}_n$ to mean ‘there are n basic evidence sets whose intersection is X .’ We also want $X \in \mathcal{E}_\omega$ to mean ‘there are finitely many evidence sets whose intersection is X ’: as we will see, this coincides with the interpretation of \square of Baltag *et al.* (note that $\omega + 1 = \omega$ by the definition of cardinal addition). If we want to consider countable intersections, we write $[\omega_1]\varphi$. We remark that, by set-theoretic convention, $\bigcap \emptyset = W$, which is the reason that we stipulate $\mathcal{E}_0 = \{W\}$. Note that our definitions then imply that $W \in \mathcal{E}_\alpha$ for all $\alpha \in \Omega$.

Definition 4.2. *Let Ω be a cardinal degree set and $\mathcal{M} = (W, \mathcal{E}, V)$ a stratified evidence model. The truth set of φ is defined as follows:*

- $\llbracket p \rrbracket = V(p)$;
- $\llbracket \neg\varphi \rrbracket = W \setminus \llbracket \varphi \rrbracket$;
- $\llbracket \varphi \wedge \psi \rrbracket = \llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket$;
- $w \in \llbracket [\alpha]\varphi \rrbracket$ iff there is $X \in \mathcal{E}_\alpha$ such that $w \in X \subseteq \llbracket \varphi \rrbracket$.

We may also write $\mathcal{M}, w \models \varphi$ if $w \in \llbracket \varphi \rrbracket$. The logical notions of satisfiability and validity are defined as usual.

Theorem 4.3. *Given a degree set Ω , the logic SEL_Ω is sound for the class of stratified evidence models.*

Proof. Let $\mathcal{M} = (W, \mathcal{E}, V)$ be an arbitrary stratified evidence model; the proof proceeds by checking that axioms are valid on \mathcal{M} and rules preserve validity. We consider only a few examples.

For Mon, suppose that $\alpha < \beta$ and $w \in \llbracket [\alpha]\varphi \rrbracket$. Then, there is $X \in \mathcal{E}_\alpha$ such that $w \in X \subseteq \llbracket \varphi \rrbracket$, so that $\alpha =: \alpha_1 \leq \beta$ (seen as a one-element sum) and thus $X \in \mathcal{E}_\beta$ (seen as a one-element intersection). Thus $X \in \mathcal{E}_\beta$ is such that $w \in X \subseteq \llbracket \varphi \rrbracket$, hence $w \in \llbracket [\beta]\varphi \rrbracket$. The argument for Dist_+ is essentially the same, except that now we have $X \in \mathcal{E}_\alpha$ such that $w \in X \subseteq \llbracket \varphi \rrbracket$ and $Y \in \mathcal{E}_\beta$ such that $w \in Y \subseteq \llbracket \psi \rrbracket$, so that if $\alpha + \beta \leq \gamma$ then $X \cap Y \in \mathcal{E}_\gamma$ and $w \in X \cap Y \subseteq \llbracket \varphi \wedge \psi \rrbracket$. Axiom T holds because $w \in \llbracket [\alpha]\varphi \rrbracket$ means that there is $X \in \mathcal{E}_\alpha$ with $w \in X \subseteq \llbracket \varphi \rrbracket$, which yields $w \in \llbracket \varphi \rrbracket$. For 4 we note that if $w \in \llbracket [\alpha]\varphi \rrbracket$ then there is $X \in \mathcal{E}_\alpha$ such that $w \in X \subseteq \llbracket \varphi \rrbracket$. Choose $v \in X$; then $v \in X \subseteq \llbracket \varphi \rrbracket$, which means that $v \in \llbracket [\alpha]\varphi \rrbracket$. Since v was arbitrary, $X \subseteq \llbracket [\alpha]\varphi \rrbracket$, thus witnessing that $w \in \llbracket [\alpha][\alpha]\varphi \rrbracket$. Other axioms are fairly standard. \square

5 Motivating examples

Hamiltonian cycles. Suppose we are given a graph $G = (X, E)$ with $|X| = n$, and we want to know if G contains a Hamiltonian cycle. Since this problem is NP-complete, there is no known method that is substantially better than guessing a sequence of nodes and checking whether it is indeed a Hamiltonian cycle.

To model this, let W be the set of all graphs with vertices X . Let V be a valuation such that $V(h)$ is the set of all graphs of W that contain a Hamiltonian cycle, and $\text{Perm}(X)$ be the set of permutations of X . To each permutation $v = (v_1, \dots, v_n) \in \text{Perm}(X)$, let us assign two evidence sets: H_v , the set of $G \in W$ such that v is a Hamiltonian cycle of G , and N_v , the set of $G \in W$ such that v is not a Hamiltonian cycle of G . We let \mathcal{E}_1 be the family of sets of the forms H_v or N_v with $v \in \text{Perm}(V)$. Finally, we define the strict evidence model $\mathcal{M} = (W, \mathcal{E}_1, V)$.

We claim that $\mathcal{M} \models h \rightarrow [1]h$. Indeed, if $G \in \llbracket h \rrbracket$ then G has a Hamiltonian cycle given by some permutation v . We then note that $G \in H_v \subseteq \llbracket h \rrbracket$, which witnesses that $G \in \llbracket [1]p \rrbracket$. On the other hand, $\mathcal{M} \not\models \neg h \rightarrow [1]\neg h$, as if v is any permutation of X , we have that there is $G \in N_v$ which has a Hamiltonian cycle (just choose G' to have a Hamiltonian cycle different from v) and thus $N_v \not\subseteq \llbracket \neg h \rrbracket$. However, we do have that $\mathcal{M} \models \neg h \rightarrow [\omega]\neg h$, since $\llbracket \neg h \rrbracket = \bigcap_{v \in \text{Perm}(X)} N_v$.

The halting problem. Now let us consider the halting problem, where we are given a Turing machine T (where we assume that the input is hard-wired into T) and we want to know whether T halts. As before, the best that one can do in general is to simulate T and wait for it to halt.

Let W be the set of all Turing machines, and for each natural number n let $\mathcal{E}_n = \{H_n, N_n\}$, where H_n is the set of machines that halt in at most n steps and N_n its complement. Let $V(h)$ be the set of all Turing machines that halt, and define a non-strict evidence model $\mathcal{M} = (W, \mathcal{E}, V)$. Then, $\mathcal{M} \models h \rightarrow [\omega]h$, since if $T \in \llbracket h \rrbracket$ it follows that T halts after n steps for some n and thus $T \in H_n \subseteq \llbracket h \rrbracket$. On the other hand, every N_n contains a Turing machine that halts (just choose M to halt after $n + 1$ steps), so that $\mathcal{M}, T \not\models \neg h \rightarrow [\omega]\neg h$. On the other hand, $\bigcap_{n < \omega} N_n = \llbracket \neg h \rrbracket$, which is a countably infinite intersection and thus witnesses that $\mathcal{M} \models \neg h \rightarrow [\omega_1]\neg h$.

6 Comparison to existing frameworks

Our framework is meant to be an extension of existing evidence logics, so it will be convenient to review such logics. For convenience we consider the following language, \mathcal{L}^{ev} of evidence logic:

$$\varphi, \psi ::= p \mid \top \mid (\varphi \wedge \psi) \mid \neg\varphi \mid E\varphi \mid B\varphi \mid \Box_0\varphi \mid \Box\varphi \mid \forall\varphi.$$

The intention is for \mathcal{L}^{ev} to include the languages of both [Baltag *et al.*, 2016a] and [van Benthem *et al.*, 2012]. If we indicate sublanguages of \mathcal{L}^{ev} by displaying the allowed modalities as subindexes, van Benthem *et al.* consider $\mathcal{L}_{EB\forall}^{\text{ev}}$ and Baltag *et al.* consider $\mathcal{L}_{\Box_0\Box\forall}^{\text{ev}}$ (in both cases, additional definable operators are also considered, as well as some proper extensions). Models are defined over *evidence frames*, which are pairs (W, \mathcal{E}) where $\mathcal{E} \subseteq 2^W$; such frames can be seen as a special case of stratified evidence frames where $\Omega = \{0, 1\}$, and viewing \mathcal{E}_1 as the set of neighbourhoods.

Semantics are given as in Definition 4.2, with the new clauses

- $w \in \llbracket E\varphi \rrbracket$ iff there is $X \in \mathcal{E}$ such that $X \subseteq \llbracket \varphi \rrbracket$;
- $w \in \llbracket \Box_0\varphi \rrbracket$ iff there is $X \in \mathcal{E}$ such that $w \in X \subseteq \llbracket \varphi \rrbracket$;
- $w \in \llbracket \Box\varphi \rrbracket$ iff there is finite $\mathcal{X} \subseteq \mathcal{E}$ such that $w \in \bigcap \mathcal{X} \subseteq \llbracket \varphi \rrbracket$;
- $w \in \llbracket \forall\varphi \rrbracket$ iff $\llbracket \varphi \rrbracket = W$.

The semantics for $B\varphi$ are somewhat more elaborate. Say that $\mathcal{S} \subseteq \mathcal{E}$ is *consistent* if any finite intersection of elements of \mathcal{S} is non-empty, and *maximal consistent* if no proper extension of \mathcal{S} is consistent. Then, $w \in \llbracket B\varphi \rrbracket$ if and only if for every maximal consistent family \mathcal{S} we have that $\bigcap \mathcal{S} \subseteq \llbracket \varphi \rrbracket$.

We do not claim that $B\varphi$ is definable in our language in general, but as pointed out in [Baltag *et al.*, 2016a], E is definable in $\mathcal{L}_{\Box_0\Box\forall}^{\text{ev}}$ by $E\varphi \leftrightarrow \exists\Box\varphi$ and B is also definable over the class of *finite* structures by $B\varphi \leftrightarrow \forall\Box\varphi$. Thus we will focus on comparing \mathcal{L}_Ω and $\mathcal{L}_{\Box_0\Box\forall}^{\text{ev}}$.

Theorem 6.1. *Over the class of strict stratified evidence models, $\mathcal{L}_{\{0,1,\omega\}}$ is expressively equivalent to $\mathcal{L}_{\Box_0\Box\forall}^{\text{ev}}$.*

Proof. By induction we show that for every $\varphi \in \mathcal{L}_{\Box_0\Box\forall}^{\text{ev}}$ there is $\varphi' \in \mathcal{L}_{\{0,1,\omega\}}$ equivalent to φ . Let $\mathcal{M} = (W, \mathcal{E}, V)$ be a

strict evidence model. If $\varphi = \forall\psi$, by induction hypothesis there is $\psi' \in \mathcal{L}_{\{0,1,\omega\}}$ equivalent to ψ . Then $[0]\psi'$ holds if and only if $\llbracket\psi'\rrbracket = W$, i.e. if and only if $\forall\psi$ holds, and we may thus set $\varphi' = [0]\psi'$. Similarly, $[1]\psi'$ holds if and only if $\Box_0\psi$ does (as the semantics are identical), and inspection on the semantics of $[\omega]\psi'$ shows that it is equivalent to $\Box\psi$, as both require that $x \in \bigcap \mathcal{X} \subseteq \llbracket\psi\rrbracket \stackrel{\text{IH}}{=} \llbracket\psi'\rrbracket$ for some finite $\mathcal{X} \subseteq \mathcal{E}_1$. Note that the translation $\varphi \mapsto \varphi'$ is invertible, so the two languages are expressively equivalent. \square

On the other hand, $\mathcal{L}_{\Box_0\Box\forall}^{\text{ev}}$ cannot express $[2]\varphi$, although we leave a proof of this fact for future work. Thus our framework naturally embeds the framework of Baltag et al., and in the case of finite structures also that of van Benthem et al.

There are other frameworks which consider quantitative notions of knowledge or belief, although these typically model *trust* rather than *effort* [Baltag et al., 2016b; Falappa et al., 2013]. Topological semantics for epistemic logic [Baltag et al., 2017; Moss and Parikh, 1992] also takes the notion of “effort” to form a new belief into consideration, although the approach is qualitative rather than quantitative.

7 Completeness for stratified evidence models

The main result of this paper is that SEL_Ω is complete for the class of strict stratified evidence models. As an intermediate step, we first show that it is complete for the class of all stratified evidence models. This intermediate result is interesting since SEL_Ω has the finite model property for this class, even when Ω contains uncountable cardinals. Thus we immediately obtain the decidability of SEL_Ω .

Theorem 7.1. *Given a cardinal degree set Ω , the logic SEL_Ω is sound and strongly complete for the class of stratified evidence models and sound and (weakly) complete for the class of finite, stratified evidence models.*

We prove this by a standard canonical model construction. Let Ω be a cardinal degree set. First let us define the set of worlds of our models. Say that a set Σ of formulas is *admissible* if it is closed under subformulas, single negations and $[0]\top \in \Sigma$.

Definition 7.2. *Let Ω be a cardinal degree set, $\Sigma \subseteq \mathcal{L}_\Omega$ be admissible, and let MC be the set of all maximal SEL_Ω -consistent subsets of Σ . For $w \in \text{MC}$, let MC_w be the set of all $v \in \text{MC}$ such that, for all φ , $[0]\varphi \in v$ if and only if $[0]\varphi \in w$.*

In order to interpret $[0]$ correctly, we need all of the worlds in a model to satisfy the same formulas of the form $[0]\varphi$. Thus we will not work with all of MC , but rather with subsets of the form MC_w , and these will be the possible worlds in a model.

Next we define the evidence sets in our canonical model.

Definition 7.3. *Fix a degree set Ω , an admissible set $\Sigma \subseteq \mathcal{L}_\Omega$ and $w \in \text{MC}$. For $\alpha \in \Omega$ and $\varphi \in \Sigma$, define*

$$X_\alpha^\varphi = \{v \in \text{MC}_w : [\alpha]\varphi \in v\}.$$

Then, let \mathcal{E}_α be the set of all non-empty sets of the form

$$X_{\alpha_1}^{\varphi_1} \cap \dots \cap X_{\alpha_n}^{\varphi_n}$$

such that $\alpha_1 + \dots + \alpha_n \leq \alpha$ and $\varphi_j \in \Sigma$ for each $j \leq n$.

With this we define the structure $\mathcal{M}_w = (\text{MC}_w, \mathcal{E}, V)$ where $u \in V(p)$ if and only if $p \in u$.

Lemma 7.4. *Given $w \in \text{MC}$, the structure \mathcal{M}_w is a stratified evidence model.*

Proof. First we check that $\mathcal{E}_0 = \{\text{MC}_w\}$; that $\text{MC}_w \in \mathcal{E}_0$ follows from the fact that $[0]\top \in v$ for all $v \in \text{MC}_w$, hence $X_0^\top = \text{MC}_w$. On the other hand, if $X \in \mathcal{E}_0$, it follows that $X = X_0^{\varphi_1} \cap \dots \cap X_0^{\varphi_n}$ for some formulas $\varphi_1, \dots, \varphi_n$. For each $i \leq n$, either $[0]\varphi_i \in w$ and hence $[0]\varphi_i \in v$ for all $v \in \text{MC}_w$, or else $[0]\varphi_i \notin w$ and hence $[0]\varphi_i \notin v$ for all $v \in \text{MC}_w$. In other words, $X_0^{\varphi_i} \in \{\emptyset, \text{MC}_w\}$. But the empty set is excluded from \mathcal{E}_0 by definition, which means that $X = \text{MC}_w$, as needed.

Now, if $\alpha + \beta \leq \gamma$ and $X \in \mathcal{E}_\alpha$, $Y \in \mathcal{E}_\beta$, we have that $X = X_{\alpha_1}^{\varphi_1} \cap \dots \cap X_{\alpha_n}^{\varphi_n}$ and $Y = X_{\alpha_{n+1}}^{\varphi_{n+1}} \cap \dots \cap X_{\alpha_{n+m}}^{\varphi_{n+m}}$ for some sequence of formulas $\varphi_1, \dots, \varphi_{n+m} \in \Sigma$ and $\alpha_1, \dots, \alpha_{n+m} \in \Omega$ such that $\alpha_1 + \dots + \alpha_n \leq \alpha$ and $\alpha_{n+1} + \dots + \alpha_{n+m} \leq \beta$. But then, $\alpha_1 + \dots + \alpha_{n+m} \leq \alpha + \beta \leq \gamma$, hence $X \cap Y = X_{\alpha_1}^{\varphi_1} \cap \dots \cap X_{\alpha_{n+m}}^{\varphi_{n+m}} \in \mathcal{E}_\gamma$. \square

Lemma 7.5. *Fix $w_0 \in \text{MC}$. Then, for all $w \in \text{MC}_{w_0}$ and all formulas φ , $\varphi \in w$ if and only if $\mathcal{M}_{w_0}, w \models \varphi$.*

Proof. A standard induction on formulas, of which only the case $[\alpha]\varphi$ is interesting. Suppose $[\alpha]\varphi \in w$; then, $X_\alpha^\varphi \neq \emptyset$, hence $w \in X_\alpha^\varphi \in \mathcal{E}_\alpha$. By the truth axiom, for all $v \in X_\alpha^\varphi$ we have that $\varphi \in v$, hence by the induction hypothesis, $X_\alpha^\varphi \subseteq \llbracket\varphi\rrbracket$, which means that $\mathcal{M}_{w_0}, w \models [\alpha]\varphi$.

Conversely, if $\mathcal{M}_{w_0}, w \models [\alpha]\varphi$, then there are formulas $[\beta_1]\psi_1, \dots, [\beta_n]\psi_n \in \Sigma$ such that $\beta_1 + \dots + \beta_n \leq \alpha$, $X_{\beta_i}^{\psi_i} \neq \emptyset$, and

$$X := X_{\beta_1}^{\psi_1} \cap \dots \cap X_{\beta_n}^{\psi_n} \subseteq \llbracket\varphi\rrbracket. \quad (1)$$

Define $\psi = [\beta_1]\psi_1 \wedge \dots \wedge [\beta_n]\psi_n$.

Let us use the notation w^α to denote the set of $\theta \in w^\alpha$ of the forms $[\alpha]\theta'$ or $\neg[\alpha]\theta'$. We claim that there exists a finite $\Theta \subseteq (w_0)^0$ such that $\Theta \vdash \psi \rightarrow \varphi$; otherwise $(w_0)^0 \cup \{\psi, \neg\varphi\}$ would be consistent, hence extendible to a maximal consistent $v \subseteq \Sigma$, and then we would have that $v \in X \setminus \llbracket\varphi\rrbracket$, contrary to (1).

Let $\theta = \bigwedge \Theta$; observe that

$$\vdash \theta \leftrightarrow [0]\theta \quad (2)$$

by axioms T 4 and N₀, and thus $\vdash ([0]\theta \wedge \psi) \rightarrow \varphi$. By the monotonicity rule, it follows that

$$\vdash [\alpha]([0]\theta \wedge \psi) \rightarrow [\alpha]\varphi. \quad (3)$$

On the other hand, since $\beta_1 + \dots + \beta_n \leq \alpha$, by Lemma 3.1 we obtain

$$\vdash [0]\theta \wedge \psi \rightarrow [\alpha]([0]\theta \wedge \psi),$$

and by (2) and (3), we see that $\vdash (\theta \wedge \psi) \rightarrow [\alpha]\varphi$.

But,

$$\Theta \cup \{[\beta_1]\psi_1, \dots, [\beta_n]\psi_n\} \subseteq w,$$

and hence $w \vdash [\alpha]\varphi$; since $[\alpha]\varphi \in \Sigma$ by assumption, it follows that $[\alpha]\varphi \in w$, as needed. \square

From this, Theorem 7.1 immediately follows.

Proof of Theorem 7.1. For strong completeness, let Φ be a consistent set of formulas and $\Sigma = \mathcal{L}_\Omega$. Then, Φ can be extended to a maximal consistent set w_* . The model \mathcal{M}_{w_*} is a stratified evidence model and for any $\varphi \in \Phi$, from $\varphi \in w_*$ and Lemma 7.5 we obtain $\mathcal{M}_{w_*}, w_* \models \varphi$, as needed.

For weak completeness, let φ be a consistent formula and let Σ be the least admissible set with $\varphi \in \Sigma$. Clearly MC is finite and since φ is consistent, $\varphi \in w_*$ for some $w_* \in \text{MC}$. Thus \mathcal{M}_{w_*} is a finite stratified evidence model satisfying φ . \square

8 Ranks and decidability

Note that Ω may be infinite, and as such, in a stratified evidence frame $\mathcal{F} = (W, \mathcal{E})$ the hierarchy \mathcal{E} is technically an infinite object, even when W is finite. However, \mathcal{E} may admit a finite representation using ranks.

Definition 8.1. Given a stratified evidence frame $\mathcal{F} = (W, \mathcal{E})$ and $X \in 2^W$, we define $\text{rk}(X) = \alpha$ if $\alpha \in \Omega$ is least such that $X \in \mathcal{E}_\alpha$, $\text{rk}(X) = \infty$ if there is no such α .

Recall that ∞ is a symbol that is decreed to be larger than all cardinals. Some properties of the rank function are as follows:

Lemma 8.2. Let $\mathcal{F} = (W, \mathcal{E})$ be a stratified evidence frame. Then:

1. If $X \in \mathcal{E}_\alpha$ for some $\alpha \in \Omega$ then $\text{rk}(X) < \infty$.
2. If $\alpha \geq \text{rk}(X)$ then $X \in \mathcal{E}_\alpha$.

Thus the frame (W, \mathcal{E}) gives rise to a structure (W, rk) . Moreover $\text{rk}(X) = 0$ iff $X = W$ and $\text{rk}(X \cap Y) \leq \text{rk}(X) + \text{rk}(Y)$. Similarly, we can reconstruct (W, \mathcal{E}) from (W, rk) .

Lemma 8.3. Let W be a set and $r: 2^W \rightarrow \Omega$ a function such that $r(X) = 0$ iff $X \in \{\emptyset, W\}$ and for all $X, Y \in 2^W$, $r(X \cap Y) \leq r(X) + r(Y)$. Define \mathcal{E}^r by $(\alpha, X) \in \mathcal{E}^r$ if and only if $\text{rk}(X) \leq \alpha$. Then, (W, \mathcal{E}^r) is a stratified evidence frame with rank function r .

Moreover, if (W, \mathcal{E}) is any evidence frame with rank function rk then $\mathcal{E} = \mathcal{E}^{\text{rk}}$.

The upshot is that (W, rk) is a finite object (since 2^W is finite, as is each of its elements) and we can computably check if rk satisfies the required properties, provided that the relevant operations on Ω are computable.

Theorem 8.4. Given a cardinal degree set Ω , the set of theorems of SEL_Ω is decidable.

Proof. This is an immediate consequence of Theorems 4.3 and 7.1, since any formula φ is either derivable or refutable in a finite model (W, \mathcal{E}) . Note that \mathcal{E} may be infinite even if W is finite; however in view of Lemma 8.3 we may represent it in the form (W, rk) , which is finite since rk has finite domain. \square

9 The representation theorem

Now we show that SEL_Ω is also complete for the class of strict stratified models. In this case we will rely on finite models given by Theorem 7.1, so we will no longer obtain strong completeness. Moreover, the finite model property will inevitably be lost if Ω contains any uncountable cardinals; note however that this is not an issue in the ‘intended’ case where $\Omega = \mathbb{N} \cup \{\omega\}$.

The idea is to start with a finite stratified model and use that to construct a new strict stratified model. The relation between the two models will be witnessed by an honest map, as defined below.

Definition 9.1. Let $\mathcal{A} = (W^{\mathcal{A}}, \mathcal{E}^{\mathcal{A}}, V^{\mathcal{A}})$ and $\mathcal{B} = (W^{\mathcal{B}}, \mathcal{E}^{\mathcal{B}}, V^{\mathcal{B}})$ be stratified evidence models. Let Ξ be a set of cardinals. A map $\pi: W^{\mathcal{A}} \rightarrow W^{\mathcal{B}}$ is reliable (for Ξ) if, whenever $\xi \in \Xi$ and $X \in \mathcal{E}_\xi^{\mathcal{A}}$, it follows that $\pi[X] \in \mathcal{E}_\xi^{\mathcal{B}}$. It is forthright (for Ξ) if, whenever $\xi \in \Xi$, $Y \in \mathcal{E}_\xi^{\mathcal{B}}$, $\pi^{-1}[Y] \in \mathcal{E}_\xi^{\mathcal{A}}$. A map that is reliable and forthright is honest (for Ξ).

Note that honest maps are automatically surjective provided $0 \in \Xi$, since from $\pi[W^{\mathcal{A}}] \in \mathcal{E}_0^{\mathcal{B}}$ we obtain $\pi[W^{\mathcal{A}}] = W^{\mathcal{B}}$. Honest maps are useful because they preserve the truth of formulas. The following is easily verified by a standard induction on formulas.

Lemma 9.2. Let $\mathcal{A} = (W^{\mathcal{A}}, \mathcal{E}^{\mathcal{A}}, V^{\mathcal{A}})$ and $\mathcal{B} = (W^{\mathcal{B}}, \mathcal{E}^{\mathcal{B}}, V^{\mathcal{B}})$ be stratified evidence models and $\pi: W^{\mathcal{A}} \rightarrow W^{\mathcal{B}}$ be an honest map such that for all $w \in W^{\mathcal{A}}$ and $p \in \mathbb{P}$, $\pi^{-1}[V^{\mathcal{B}}(p)] = V^{\mathcal{A}}(p)$. Let φ be any formula and Ξ a set containing all cardinals that appear in φ . Then for any $w \in W^{\mathcal{A}}$ we have that $\mathcal{A}, w \models \varphi$ if and only if $\mathcal{B}, \pi(w) \models \varphi$.

Thus if given a finite stratified evidence model \mathcal{M} we can obtain a strict stratified evidence model \mathcal{A} and an honest map $\pi: W^{\mathcal{A}} \rightarrow W^{\mathcal{M}}$, we will immediately obtain a completeness result for the class of strict stratified models. The following proposition shows that this is indeed the case.

Proposition 9.3. Let $\mathcal{M} = (W^{\mathcal{M}}, \mathcal{E}^{\mathcal{M}})$ be a finite stratified evidence frame and Ξ be a finite set of cardinals. Then, there exists a stratified evidence model $\mathcal{A} = (W^{\mathcal{A}}, \mathcal{E}_1^{\mathcal{A}}, V^{\mathcal{A}})$ and an honest map $\pi: W^{\mathcal{A}} \rightarrow W^{\mathcal{M}}$.

Proof. Define $\mathcal{E}_\Xi^{\mathcal{M}} := \mathcal{E}^{\mathcal{M}} \cap (\Xi \times 2^{W^{\mathcal{M}}})$ and let $\mathcal{H} \subseteq \mathcal{E}_\Xi^{\mathcal{M}}$ be maximal with the property that there exists a stratified evidence frame $\mathcal{A} = (W^{\mathcal{A}}, \mathcal{E}_1^{\mathcal{A}})$ and a reliable map $\pi: W^{\mathcal{A}} \rightarrow W^{\mathcal{M}}$ with $\pi^{-1}[X] \in \mathcal{E}_\xi^{\mathcal{A}}$ whenever $X \in \mathcal{H}_\xi$.

We claim that $\mathcal{H} = \mathcal{E}_\Xi^{\mathcal{M}}$. If not, let $(\alpha, N) \in \mathcal{E}_\Xi^{\mathcal{M}} \setminus \mathcal{H}$. Since Ξ is finite, choose $\beta < \alpha + 1$ such that if $\beta < \alpha' + 1$ and $\alpha' \in \Xi$ it follows that $\alpha \leq \alpha'$. Define a new model \mathcal{B} and a map $\sigma: W^{\mathcal{B}} \rightarrow W^{\mathcal{A}}$ as follows. Let I_β be a set with cardinality β and set $\mathcal{B} = W^{\mathcal{A}} \times I_\beta$, and $\sigma(w, i) = w$. Then, set $\rho = \pi \sigma: W^{\mathcal{B}} \rightarrow W^{\mathcal{M}}$. For elements of \mathcal{B} , we will write X_i instead of $X \times \{i\}$.

As \mathcal{B} is a strict stratified evidence model, we only need to define $\mathcal{E}_1^{\mathcal{B}}$. We will define $\mathcal{E}_1^{\mathcal{B}} = \mathcal{N} \cup \mathcal{O}$ (for ‘new’ and ‘old’) as follows. We set $\mathcal{N} = \{N_i : i \in I_\beta\}$, where

$$N_i = \rho^{-1}[N] \cup \bigcup_{j \in I_\beta \setminus \{i\}} W_j^{\mathcal{A}},$$

and we set $\mathcal{O} = \{\sigma^{-1}[A] : A \in \mathcal{E}_1^A\}$.

We have that $\rho^{-1}[Y] \in \mathcal{E}_\gamma^B$ whenever $Y \in \mathcal{H}_\gamma$, since $\pi^{-1}[Y] \in \mathcal{E}_\gamma^A$ and thus there are $\zeta < \gamma + 1$, a set I_ζ with cardinality ζ and a collection $\{A_i : i \in I_\zeta\} \subseteq \mathcal{E}_1^A$ so that $\pi^{-1}[Y] = \bigcap_{i \in I_\zeta} A_i$, and therefore $\rho^{-1}[Y] = \sigma^{-1}\pi^{-1}[Y] = \sigma^{-1}[\bigcap_{i \in I_\zeta} A_i] = \bigcap_{i \in I_\zeta} \sigma^{-1}[A_i] \in \mathcal{E}_\gamma^B$, where the last step holds because $\sigma^{-1}[A_i] \in \mathcal{O}$ and \mathcal{B} is a strict stratified evidence frame and $|I_\zeta| = \zeta < \gamma + 1$. Moreover,

$$\rho^{-1}[N] = \bigcap_{i \in I_\beta} N_i \in \mathcal{E}_\alpha^B.$$

We conclude that $\rho^{-1}[Y] \in \mathcal{H}_\gamma$ whenever $(\gamma, Y) \in \mathcal{H}_\gamma \cup \{(\alpha, N)\}$.

Next we check that π is reliable. Let $\mathcal{X} \in (\mathcal{E}_\gamma^B)$. We need to check that $\rho[\bigcap \mathcal{X}] \in \mathcal{E}_\delta^M$ whenever $\gamma < \delta + 1$. Let $J = \{j \in I_\beta : N_j \in \mathcal{X}\}$ and $\mathcal{Y} = \mathcal{X} \cap \mathcal{O}$. By definition of \mathcal{O} , there is $\mathcal{Z} \subseteq \mathcal{E}_1^A$ such that $|\mathcal{Z}| = |\mathcal{Y}|$ and $\mathcal{Y} = \{\sigma^{-1}[Z] : Z \in \mathcal{Z}\}$.

Consider two cases. If $J = I_\beta$, then

$$\bigcap_{i \in J} N_i = N \times I_\beta = \rho^{-1}[N]$$

and hence

$$\begin{aligned} \rho\left[\bigcap \mathcal{X}\right] &= \rho\left[\bigcap_{i \in I_\beta} N_i \cap \bigcap \mathcal{Y}\right] = \rho\left[\rho^{-1}[N] \cap \bigcap \mathcal{Y}\right] \\ &= N \cap \rho\left[\bigcap \mathcal{Y}\right] = N \cap \pi\sigma\left[\bigcap \mathcal{Y}\right] \\ &= N \cap \pi\sigma\left[\bigcap_{Z \in \mathcal{Z}} \sigma^{-1}[Z]\right] \\ &= N \cap \pi\sigma\sigma^{-1}\left[\bigcap \mathcal{Z}\right] = N \cap \pi\left[\bigcap \mathcal{Z}\right], \end{aligned}$$

the last equality being because σ is surjective. Now, note that $\pi[\bigcap \mathcal{Z}] \in \mathcal{E}_{|\mathcal{Y}|}^M$ since π is reliable, and $\beta + |\mathcal{Y}| = \gamma < \delta + 1$, hence $\rho[\bigcap \mathcal{X}] \in \mathcal{E}_\delta^M$, as needed.

Otherwise, there is some $j < \beta$ such that $N_j \notin \mathcal{X}$. Here, we claim that

$$\rho\left[\bigcap \mathcal{X}\right] = \rho\left[W_j^A \cap \bigcap \mathcal{Y}\right]. \quad (4)$$

To see this, note that clearly $\rho[\bigcap \mathcal{X}] \supseteq \rho[W_j^A \cap \bigcap \mathcal{Y}]$, since $W_j^A \subseteq \bigcap_{i \in J} N_i$. For the other direction, take $w \in \rho[\bigcap \mathcal{X}]$. Then, $w = \rho(w, i)$ for some $(w, i) \in \bigcap \mathcal{X}$. But, for any $i \in J$, $(w, i) \in N_j$, and for any $Y = \sigma^{-1}[Z] \in \mathcal{Y}$, since $(w, i) \in Y$ it follows that $w \in Z$ and hence $(w, j) \in \sigma^{-1}[Z]$. Thus $w = \rho(w, j) \in \rho[W_j^A \cap \bigcap \mathcal{Y}]$, and (4) holds as claimed. From this, we see that

$$\begin{aligned} \rho\left[\bigcap \mathcal{X}\right] &= \rho\left[W_j^A \cap \bigcap \mathcal{Y}\right] = \pi\sigma\left[W_j^A \cap \bigcap \mathcal{Y}\right] \\ &= \pi\sigma\left[W_j^A \cap \bigcap_{Z \in \mathcal{Z}} \sigma^{-1}[Z]\right] \\ &= \pi\sigma\left[W_j^A \cap \sigma^{-1}[\bigcap \mathcal{Z}]\right] \\ &= \pi\left[\sigma[W_j^A] \cap \bigcap \mathcal{Z}\right] = \pi\left[W^A \cap \bigcap \mathcal{Z}\right] \\ &= \pi\left[\bigcap \mathcal{Z}\right], \end{aligned}$$

which is an element of \mathcal{E}_δ^M by our assumption on \mathcal{A} and the fact that $|\mathcal{Z}| \leq |\mathcal{X}| < \delta + 1$.

Thus \mathcal{B} is a stratified evidence model and $\rho^{-1}[Y] \in \mathcal{E}^B$ whenever $Y \in \mathcal{H}_\gamma \cup \{(\alpha, N)\}$, contradicting the maximality of \mathcal{H} . Hence we conclude that $\mathcal{H} = \mathcal{E}_\Xi^A$, and π is honest, as needed. \square

In view of Theorem 7.1, we obtain:

Theorem 9.4. *The logic SEL_Ω is sound and (weakly) complete for the class of strict stratified evidence models.*

Proof. If φ is consistent then it is satisfiable on some finite stratified evidence model \mathcal{M} on some world w_* . Let Ξ be a finite set containing 0 and all cardinals appearing in φ . Then by Proposition 9.3 there is a strict stratified evidence model \mathcal{A} and an honest map $\pi: W^A \rightarrow W^M$. Since $0 \in \Xi$, π is surjective, so for v_* with $\pi(v_*) = w_*$ we have that $\mathcal{M}, v_* \models \varphi$, as needed. \square

10 Concluding remarks

We have presented a new family of logics called stratified evidence logics. In contrast to traditional evidence logics that only account for the qualitative aspects of the connection between evidence and belief, our logics focus on the quantitative aspects. In particular, they allow us to represent the amount of pieces of evidence that are sufficient for an agent to form a certain belief. We have studied the mathematical and computational properties of our framework by providing a complete deductive calculus and proving that its satisfiability problem is decidable. Moreover, we have shown that existing evidence logics can be embedded in our framework.

There are at least three directions that we plan to explore in future work. First of all, we plan to study complexity of the satisfiability problem of our stratified evidence logic and to provide a decision procedure for it based on tableaux. The latter will make it exploitable in the context of AI applications, such as the one briefly described in Section 5. Secondly, we plan to move from a static to a dynamic perspective by adapting the “model update” approach of dynamic epistemic logic [van Ditmarsch *et al.*, 2007] to our setting. The idea is to formalize two basic operations on evidence, conceived as update operations on strict stratified evidence models: evidence expansion and evidence contraction. The former consists in adding a new piece of evidence to the agent’s evidence set \mathcal{E}_1 , while the latter consists in removing a piece of evidence from it. Thirdly, we plan to generalize our framework to the multi-agent setting in which different agents may have different evidence sets and exchange pieces of evidence through communication.

References

- [Artemov, 2008] S. N. Artemov. The logic of justification. *The Review of Symbolic Logic*, 1(4):477–513, 2008.
- [Balbiani *et al.*, 2016] P. Balbiani, D. Fernández-Duque, and E. Lorini. A logical theory of belief dynamics for resource-bounded agents. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016*, pages 644–652, 2016.

- [Balbiani *et al.*, 2018] P. Balbiani, D. Fernández-Duque, and E. Lorini. The dynamics of epistemic attitudes in resource-bounded agents. *Studia Logica*, 2018.
- [Baltag and Occhipinti Liberman, 2017] A. Baltag and A. Occhipinti Liberman. Evidence logics with relational evidence. In *Logic, Rationality, and Interaction - 6th International Workshop, LORI 2017, Sapporo, Japan, September 11-14, 2017, Proceedings*, pages 17–32, 2017.
- [Baltag *et al.*, 2016a] A. Baltag, N. Bezhanishvili, A. Özgün, and A. Smets. Justified belief and the topology of evidence. In *Logic, Language, Information, and Computation - 23rd International Workshop, WoLLIC 2016, Puebla, Mexico, August 16-19th, 2016. Proceedings*, pages 83–103, 2016.
- [Baltag *et al.*, 2016b] A. Baltag, V. Fiutek, and S. Smets. Beliefs and evidence in justification models. In *Advances in Modal Logic*, pages 156–176, 2016.
- [Baltag *et al.*, 2017] A. Baltag, A. Özgün, and A.L. Vargas Sandoval. Topo-logic as a dynamic-epistemic logic. In *Logic, Rationality, and Interaction*, pages 330–346, 2017.
- [Benferhat *et al.*, 1993] S. Benferhat, D. Dubois, and H. Prade. Argumentative inference in uncertain and inconsistent knowledge bases. In *Uncertainty in Artificial Intelligence*, pages 411–419, 1993.
- [Benferhat *et al.*, 1995] S. Benferhat, D. Dubois, and H. Prade. How to infer from inconsistent beliefs without revising? In *International Joint Conference on Artificial Intelligence*, pages 1449–1457, 1995.
- [Dubois *et al.*, 2016] D. Dubois, W. Liub, J. Mac, and H. Prade. The basic principles of uncertain information fusion. An organised review of merging rules in different representation frameworks. *Information Fusion*, 32:12–39, 2016.
- [Falappa *et al.*, 2013] M.A. Falappa, A.J. García, G. Kern-Isberner, and G.R. Simari. Stratified belief bases revision with argumentative inference. *J. Philosophical Logic*, 42(1):161–193, 2013.
- [Fitting, 2005] M. Fitting. The logic of proofs, semantically. *Annals of Pure and Applied Logic*, 132(1):1–25, 2005.
- [Liu and Lorini, 2017] F. Liu and E. Lorini. Reasoning about belief, evidence and trust in a multi-agent setting. In *Proceedings of the 20th International Conference on Principles and Practice of Multi-Agent Systems (PRIMA 2017)*, volume 10621 of *LNCS*, pages 71–89. Springer, 2017.
- [Miller, 1976] L.W. Miller. Normal functions and constructive ordinal notations. *Journal of Symbolic Logic*, 41(2):439–459, 1976.
- [Moss and Parikh, 1992] L. S. Moss and R. Parikh. Topological reasoning and the logic of knowledge: preliminary report. In *Proceedings of the 4th conference on Theoretical Aspects of Reasoning about Knowledge (TARK)*, pages 95–105. Morgan Kaufmann Publishers Inc., 1992.
- [Naumov and Tao, 2015] P. Naumov and J. Tao. Budget-constrained knowledge in multiagent systems. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, IFAA-MAS, pages 219–226, 2015.
- [Shafer, 1976] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.
- [van Benthem and Pacuit, 2011a] J. van Benthem and E. Pacuit. Dynamic logics of evidence-based beliefs. *Studia Logica*, 99(1-3):61–92, 2011.
- [van Benthem and Pacuit, 2011b] J. van Benthem and E. Pacuit. Logical dynamics of evidence. In *Logic, Rationality, and Interaction - Third International Workshop, LORI 2011, Guangzhou, China, October 10-13, 2011. Proceedings*, pages 1–27, 2011.
- [van Benthem *et al.*, 2012] J. van Benthem, David Fernández-Duque, and E. Pacuit. Evidence logic: A new look at neighborhood structures. In *Advances in Modal Logic 9, papers from the ninth conference on "Advances in Modal Logic," held in Copenhagen, Denmark, 22-25 August 2012*, pages 97–118, 2012.
- [van Benthem *et al.*, 2014] J. van Benthem, D. Fernández-Duque, and E. Pacuit. Evidence and plausibility in neighborhood structures. *Ann. Pure Appl. Logic*, 165(1):106–133, 2014.
- [van der Hoek and Meyer, 1991] W. van der Hoek and J.-J. Ch. Meyer. Graded modalities in epistemic logic. *Logique et Analyse*, 34(133-134):251–270, 1991.
- [van Ditmarsch *et al.*, 2007] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. Berlin: Springer, 2007.
- [Velázquez-Quesada, 2013] F. R. Velázquez-Quesada. Explicit and implicit knowledge in neighbourhood models. In *Proceedings of LORI 2013*, volume 8196 of *LNCS*, pages 239–252. Springer, 2013.