Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning
Main Track - Short paper

On the Epistemic Logic of Incomplete Argumentation Frameworks

Andreas Herzig[1] , **Antonio Yuste-Ginel**[2]
[1]IRIT, CNRS, Toulouse
[2]Departamento de Filosofía, Universidad de Málaga (UMA)
herzig@irit.fr , antonioyusteginel@gmail.com

Abstract

We study the relation between two existing formalisms: incomplete argumentation frameworks (IAFs) and epistemic logic of visibility (ELV). We show that the set of completions of a given IAF naturally corresponds to a specific equivalence class of possible worlds within the model of visibility. This connection is further strengthened in two directions. First, we show how to reduce argument acceptance problems of IAFs to ELV model-checking problems. Second, we highlight the epistemic assumptions that underlie IAFs by providing a minimal epistemic logic for IAFs.

1 Introduction

Argumentation is a well-established research area within the field of Artificial Intelligence. It has been proved useful for a variety of purposes (Bench-Capon and Dunne 2007; Atkinson et al. 2017), including (but not limited to): providing an operational semantics for diverse non-monotonic logics, establishing a bridge between human and machine reasoning, and improving artificial agents' interaction. At the abstract level, that is, without analysing the structure and nature of arguments, Dung's *argumentation frameworks* (AFs) are by far the most studied formal construct (Dung 1995). Conceptually, AFs respond to one fundamental question of (general) argumentation theory: given a set of possibly conflicting arguments, which subsets of arguments are to be accepted by a rational agent?

AFs are, however, too abstract for some purposes. Consequently, they have been extended in many different directions, including the handling of preferences (Amgoud and Vesic 2011), and more refined forms of interaction between arguments (Cayrol and Lagasquie-Schiex 2005). One of the main assumptions that AFs come equipped with is perfect knowledge about the structure of the modelled discussion (about the structure of the argumentation framework). The formal argumentation community has recently made efforts to relax this assumption. This has been done using quantitative methods, mainly probabilistic (Li, Oren, and Norman 2011), and qualitative ones. Among the second type, *incomplete argumentation frameworks* (IAFs) are prominent (Coste-Marquis et al. 2007; Baumeister, Neugebauer, and Rothe 2018; Baumeister et al. 2018b; Fazzinga, Flesca, and Furfaro 2020). An IAF partitions both the set of arguments A and the set of attacks R of an AF into two sets, distinguishing arguments (attacks) known with certainty from those perceived as uncertain by the agent.

In this paper, we pursue a simple research question: what is the epistemic logic underlying IAFs? Our answer, presented in Section 3, consists in establishing a strong connection between IAFs and the *logic of visibility* (ELV) (Herzig, Lorini, and Maffre 2018). We further exploit this connection in two different senses. First, we show how it can be used to reduce acceptance reasoning tasks of IAFs to ELV model-checking problems (Section 4). Second, we provide a minimal epistemic logic for IAFs that makes the epistemic assumptions underlying them explicit (Section 5). For space reasons, results are stated without proof; they can be found in Antonio Yuste-Ginel's forthcoming PhD dissertation.

2 Preliminaries

2.1 Abstract Argumentation Frameworks

An **abstract argumentation framework** (AF) is a directed graph, i.e., a pair (A, R) where A is a finite set whose elements stand for *arguments* and $R \subseteq A \times A$ represents some kind of conflict-based relation among arguments, typically, an *attack relation*. AFs were first studied by (Dung 1995), where the author proposed to capture the informal notion of argument acceptability through different argumentation semantics. For the sake of brevity, we limit ourselves to stable semantics, but our analysis applies to all semantics that can be captured in propositional logic (see (Besnard, Cayrol, and Lagasquie-Schiex 2020)). A **stable extension** of an AF (A, R) is a set of arguments $E \subseteq A$ s.th.: (i) there are no $x, y \in E$ s.th. $(x, y) \in R$ ('E is conflict-free'); and (ii) for every $x \in A \setminus E$, there is a $y \in E$ such that $(y, x) \in R$ ('E attacks every argument outside itself').

2.2 Incomplete AFs

An IAF is a tuple $\mathsf{IAF} = (A, A^?, R, R^?)$ s.th. $A \cap A^? = \emptyset$; $R, R^? \subseteq (A \cup A^?) \times (A \cup A^?)$; and $R \cap R^? = \emptyset$. (A, R) is the *definite part* of IAF while $(A^?, R^?)$ is the *uncertain part* of IAF. We assume that $A \cup A^?$ is finite. A **completion** of IAF is any pair (A^*, R^*) s.th.:

- $A \subseteq A^* \subseteq A \cup A^?$; and
- $R_{|A^*} \subseteq R^* \subseteq (R \cup R^?)_{|A^*}$

where we note $R_{|Y} = R \cap (Y \times Y)$ the restriction of a relation $R \subseteq X \times X$ to $Y \subseteq X$.

All the classic reasoning problems regarding AFs were generalized to IAFs and their complexity was characterised (Baumeister et al. 2018b; Baumeister et al. 2021). We here focus on argument acceptance problems. Classic acceptance problems for standard AFs can be generalized to IAFs as follows. The **set of all stable extensions** of (A, R) is noted $\mathsf{st}(A, R)$.

stable-Possible-Sceptical-Acceptance (st-PSA)
Given: An incomplete AF $\mathsf{IAF} = (A, A^?, R, R^?)$ and an argument $a \in A$. **Question:** Does there exist a completion (A^*, R^*) of IAF s.th. for all $E \in \mathsf{st}(A^*, R^*), a \in E$?

stable-Necessary-Credulous-Acceptance (st-NCA)
Given: An incomplete AF $\mathsf{IAF} = (A, A^?, R, R^?)$ and an argument $a \in A$. **Question:** Is it true that for all completions (A^*, R^*) of IAF there is an $E \in \mathsf{st}(A^*, R^*)$ s.th. $a \in E$?

The other two variants of the problem, stable-Possible-Credulous-Acceptance (st-PCA) and stable-Necessary-Sceptical-Acceptance (st-NSA), are obtained by changing the quantifiers of the definitions above in the obvious way; see (Baumeister et al. 2021) for details.

2.3 Epistemic Logic of Visibility

Epistemic logic of visibility or observability was born as a lightweight alternative to standard epistemic logic (van der Hoek, Troquard, and Wooldridge 2011; Herzig, Lorini, and Maffre 2018; Cooper et al. 2021). We consider its single-agent version and adapt it to our purposes. Let Σ be a finite set of arguments, called the **signature**. We define the **set of atoms for signature** Σ in two steps:

$$\mathsf{Prop}^\Sigma = \{\mathsf{aw}_x \mid x \in \Sigma\} \cup \{\mathsf{r}_{x,y} \mid (x,y) \in \Sigma \times \Sigma\} \cup \{\mathsf{in}_x \mid x \in \Sigma\};$$

$$\mathsf{ATM}^\Sigma = \mathsf{Prop}^\Sigma \cup \{\mathbf{S}p \mid p \in \mathsf{Prop}^\Sigma\}.$$

Intuitively, aw_x means that the agent is aware of argument x; $\mathsf{r}_{x,y}$ means that x attacks y; in_x means that argument x is in the extension; and $\mathbf{S}p$ means that the agent sees the value of p (equivalently, she knows whether p). We use the term *awareness* in a rather loose sense here. Among different possible interpretations, it can be seen as the arguments the agent's *opponent* is aware of. This explains why an agent can be aware of an argument without knowing whether she is aware of it. Moreover, note that Σ does not need to coincide with $A \cup A^?$: there can be arguments whose awareness the agent knows not to hold.

The **epistemic language** for signature Σ is generated by the following BNF:

$$\varphi ::= \alpha \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \mathbf{K}\varphi \qquad (\alpha \in \mathsf{ATM}^\Sigma)$$

where $\mathbf{K}\varphi$ reads "the agent knows that φ".[1] We also consider the dual $\hat{\mathbf{K}} = \neg\mathbf{K}\neg$, read "the agent considers possible that...".

The **single-agent visibility model** for ATM^Σ is the pair $M_{vis} = (2^{\mathsf{ATM}^\Sigma}, \sim)$, where the *accessibility relation* $\sim \subseteq 2^{\mathsf{ATM}^\Sigma} \times 2^{\mathsf{ATM}^\Sigma}$ is such that $v \sim v'$ iff for every $p \in \mathsf{Prop}^\Sigma$:

1. if $\mathbf{S}p \in v$ then ($p \in v$ iff $p \in v'$);
2. $\mathbf{S}p \in v$ iff $\mathbf{S}p \in v'$.

Formulas are interpreted at states of M_{vis}, i.e., at propositional valuations. We only state the clause for the epistemic operator:

$$M_{vis}, v \models \mathbf{K}\varphi \quad \text{iff} \quad v \sim v' \text{ implies } M_{vis}, v' \models \varphi.$$

Note that the second constraint in the definition of the visibility model guarantees that $\mathbf{K}$ is a fully introspective operator ($\mathbf{K}\varphi \to \mathbf{K}\mathbf{K}\varphi$ and $\neg\mathbf{K}\varphi \to \mathbf{K}\neg\mathbf{K}\varphi$ are both valid in the visibility model). Moreover, $\mathbf{K}p \leftrightarrow (\mathbf{S}p \wedge p)$ is valid too.

3 Incomplete AFs and Visibility Models

Relations between IAFs and possible worlds semantics were spotted in the literature (Baumeister et al. 2018a; Baumeister, Neugebauer, and Rothe 2018; Baumeister et al. 2021), and just recently explored in more depth by reducing IAFs to general epistemic logic (Proietti and Yuste-Ginel 2021). Nevertheless, a basic research question about the relation between these two formalisms remains still open: what is *the* epistemic logic underlying IAFs? We provide an answer by connecting IAFs to the single-agent logic of visibility. Formally, we show a one-to-one correspondence between completions of a given IAF and a specific equivalence class of the visibility model. For the sake of simplicity, we dispense with acceptance variables in this section, that is, we assume that $\mathsf{Prop}^\Sigma = \{\mathsf{aw}_x \mid x \in \Sigma\} \cup \{\mathsf{r}_{x,y} \mid (x,y) \in \Sigma \times \Sigma\}$.

Let $\mathsf{IAF} = (A, A^?, R, R^?)$ be an IAF with $A \cup A^? \subseteq \Sigma$. We associate to IAF **its single-agent propositional valuation**

$$v_{\mathsf{IAF}} = \{\mathbf{S}\mathsf{aw}_x \mid x \in \Sigma \setminus A^?\} \cup \{\mathsf{aw}_x \mid x \in A\} \cup \{\mathbf{S}\mathsf{r}_{x,y} \mid (x,y) \in (\Sigma \times \Sigma) \setminus R^?\} \cup \{\mathsf{r}_{x,y} \mid (x,y) \in R\}.$$

The other way round, for every propositional valuation $v \subseteq \mathsf{ATM}^\Sigma$ we define its **induced AF** (A_v, R_v) by:

$$A_v = \{x \in \Sigma \mid \mathsf{aw}_x \in v\};$$

$$R_v = \{(x,y) \in A_v \times A_v \mid \mathsf{r}_{x,y} \in v\}.$$

3.1 From IAFs to the Visibility Model

Using the notions we have just defined, we can provide the first clear connection between IAFs and the visibility model.

Proposition 1. *Let* $\mathsf{IAF} = (A, A^?, R, R^?)$ *be an IAF s.th.* $A \cup A^? \subseteq \Sigma$*, let* v_{IAF} *be its propositional valuation, and let* $M_{vis} = (2^{\mathsf{ATM}^\Sigma}, \sim)$ *be the single-agent visibility model for the set of atoms* ATM^Σ*, then:*

- *For each completion* (A^*, R^*) *of* IAF *there is a* $u \in \sim[v_{\mathsf{IAF}}]$ *s.th.* $(A^*, R^*) = (A_u, R_u)$.
- *For each* $u \in \sim[v_{\mathsf{IAF}}]$ *there is a completion* (A^*, R^*) *of* IAF *s.th.* $(A^*, R^*) = (A_u, R_u)$.

Hence completions can be understood as possible worlds of the visibility model (i.e., as propositional valuations over ATM^Σ) that are indistinguishable for the formalized agent.

[1] As shown in (Herzig, Lorini, and Maffre 2018), in ELV the operator $\mathbf{K}$ can be recursively eliminated from any visibility formula. However, we prefer to include it for presentational purposes.

3.2 From the Visibility Model to IAFs

The visibility model for a signature Σ contains *all* IAFs that can be built from Σ. Formally, we say that $(A, A^?, R, R^?)$ is **built over** Σ whenever $A \cup A^? \subseteq \Sigma$.

Proposition 2. *Let Σ be a signature, let $M_{vis} = (2^{\mathsf{ATM}^\Sigma}, \sim)$ be the single-agent visibility model for ATM^Σ. For every $\mathsf{IAF} = (A, A^?, R, R^?)$ built over Σ there is a $u \in 2^{\mathsf{ATM}^\Sigma}$ s.th. $u = v_{\mathsf{IAF}}$.*

In the other direction things get more subtle. Given a valuation $v \subseteq \mathsf{ATM}^\Sigma$, we define the **IAF represented by** v as $\mathcal{IAF}(v) = (A(v), A^?(v), R(v), R^?(v))$, where:

$$A(v) = \{x \in \Sigma \mid v \models \mathbf{K}\mathsf{aw}_x\};$$

$$A^?(v) = \{x \in \Sigma \mid v \models \neg\mathbf{S}\mathsf{aw}_x\};$$

$$R(v) = \{(x, y) \in \Sigma\times\Sigma \mid v \models \mathbf{K}\mathsf{r}_{x,y} \land \hat{\mathbf{K}}\mathsf{aw}_x \land \hat{\mathbf{K}}\mathsf{aw}_y\};$$

$$R^?(v) = \{(x, y) \in \Sigma\times\Sigma \mid v \models \neg\mathbf{S}\mathsf{r}_{x,y} \land \hat{\mathbf{K}}\mathsf{aw}_x \land \hat{\mathbf{K}}\mathsf{aw}_y\}.$$

Informally, $\mathcal{IAF}(v)$ is the IAF where certain arguments are those that the agent knows that she is aware of; uncertain arguments are those whose awareness the agent is not sure about; and analogously for attacks (conditionally on the awareness of the involved arguments). It is easy to check that, for each valuation v, (i) $\mathcal{IAF}(v)$ is indeed an IAF; and (ii) the set of all completions of $\mathcal{IAF}(v)$ equals the set of AFs (A_u, R_u) such that $u \in \sim[v]$.

Moreover, given IAF, it holds that

$$\mathcal{IAF}(v_{\mathsf{IAF}}) = \mathsf{IAF}.$$

In words, we can go from IAFs to states of the visibility model and back without loosing any information. However, this is not true if we start from the set of all valuations: it does not hold that for any u, $v_{\mathcal{IAF}(u)} = u$. The reason is that there are some valuations containing "defective information": there can be attacks that are true in all indistinguishable valuations, and hence known, even if the agent knows that she is not aware of one of the involved arguments. These pieces of defective information are cleaned out by applying the function $\mathcal{IAF}$. Alternatively we could impose that valuations should satisfy the property $\hat{\mathbf{K}}\mathsf{r}_{x,y} \to (\hat{\mathbf{K}}\mathsf{aw}_x \land \hat{\mathbf{K}}\mathsf{aw}_y)$.

4 Checking Acceptance via Model-Checking

As an application of the correspondence between IAFs and the logic of visibility, we can reduce acceptance problems regarding the former to model-checking problems of the latter. This can be done modulo some propositional logic encoding of standard argumentation semantics. We illustrate this idea for the case of stable semantics. Let us first recall how stable extensions of (awareness-relativised) AFs can be captured in propositional logic. We make use of the shorthand $\mathsf{r}^{\mathsf{aw}}_{x,y} = \mathsf{r}_{x,y} \land \mathsf{aw}_x \land \mathsf{aw}_y$ and define the formula

$$\mathsf{Stable} = \bigwedge_{x\in\Sigma} \Big((\mathsf{in}_x \to \mathsf{aw}_x) \land \big(\mathsf{aw}_x \to (\mathsf{in}_x \leftrightarrow \neg \bigvee_{y\in\Sigma} (\mathsf{in}_y \land \mathsf{r}^{\mathsf{aw}}_{y,x}))\big)\Big).$$

As shown in (Doutre, Maffre, and McBurney 2017, Proposition 1), we have that $v \models \mathsf{Stable}$ iff $E_v = \{x \in \Sigma \mid \mathsf{in}_x \in v\}$ is a stable extension of (A_v, R_v). Using this encoding and the correspondence between completions and accessible states in the visibility model of Proposition 1 we obtain:

Proposition 3. *Let $\mathsf{IAF} = (A, A^?, R, R^?)$ be an IAF built over Σ, let v_{IAF} be its propositional valuation, let $a \in A$ and let $M_{vis} = (2^{\mathsf{ATM}^\Sigma}, \sim)$ be the single-agent visibility model for the set of atoms ATM^Σ, then:*

- *The answer to the* st-NSA *(stable-Necessary-Sceptical-Acceptance) problem with input* IAF *and $a \in A$ is yes iff $M_{vis}, v_{\mathsf{IAF}} \models \mathbf{K}(\mathsf{Stable} \to \mathsf{in}_a)$.*
- *The answer to the* st-PCA *(stable-Possible-Credulous-Acceptance) problem with input* IAF *and $a \in A$ is yes iff $M_{vis}, v_{\mathsf{IAF}} \models \hat{\mathbf{K}}(\mathsf{Stable} \land \mathsf{in}_a)$.*

Curiously enough, capturing st-PSA and st-NCA in the logic of visibility requires augmenting the language, jumping from the single-agent version to the two-agent version. The underlying reason for this is that while the definitions of st-PCA and st-NSA use a $\exists\exists$ (resp. $\forall\forall$) pattern for quantifying over completions and extensions, the definitions of the other two reasoning tasks alternate quantifiers. Our solution is to modularize the formalized agent: the agent consists of a part for reasoning about completions (named 1 or the *epistemic part*) and another one for reasoning about extensions (named 2 or the *argumentative part*).

First of all, we generalize the set of logical atoms so as to include a finite, non-empty set of agents $\mathtt{Agt}$:

$$\mathsf{ATM}^\Sigma_{\mathtt{Agt}} = \mathsf{Prop}^\Sigma \cup \{\mathbf{S}_i p \mid i \in \mathtt{Agt}, p \in \mathsf{Prop}^\Sigma\},$$

where Prop^Σ is the set of propositional variables as defined in Section 2.3.

Note that we do not need to parametrize awareness variables for each agent. As mentioned, we are going to modularize the agent regarding his epistemic and argumentative reasoning, but not her awareness. In the remaining of this section, we shall restrict our attention to $\mathtt{Agt} = \{1, 2\}$. The multi-agent epistemic language extends the epistemic language used before by including a different epistemic operator $\mathbf{K}_i$ for each agent in $\mathtt{Agt}$. Nevertheless, since agent 2 represents the "argumentative part" of the formalize agent, we use $\Box_2$ instead of $\mathbf{K}_2$, for conceptual clarity. We also use $\Diamond_2$ as a shorthand for $\neg\Box_2\neg$.

The multi-agent visibility model is the pair $M_{vis} = (2^{\mathsf{ATM}^\Sigma_{\mathtt{Agt}}}, \{\sim_i\}_{i\in\mathtt{Agt}})$ where each $\sim_i \subseteq 2^{\mathsf{ATM}^\Sigma_{\mathtt{Agt}}} \times 2^{\mathsf{ATM}^\Sigma_{\mathtt{Agt}}}$ is such that $v \sim_i v'$ iff for every $p \in \mathsf{Prop}^\Sigma$,

1. if $\mathbf{S}_i p \in v$ then ($p \in v$ iff $p \in v'$);
2. $\mathbf{S}_j p \in v$ iff $\mathbf{S}_j p \in v'$ for every $j \in \mathtt{Agt}$.

We adapt the notion of **propositional valuation associated to an IAF** to the two-agent language as follows. Let $\mathsf{IAF} = (A, A^?, R, R^?)$ be an IAF s.th. $A \cup A^? \subseteq \Sigma$. Define

$$v_{\mathsf{IAF}} = \{\mathbf{S}_1\mathsf{aw}_x \mid x \in \Sigma \setminus A^?\} \cup \{\mathbf{S}_1\mathsf{r}_{x,y} \mid (x, y) \in (\Sigma\times\Sigma) \setminus R^?\} \cup \{\mathbf{S}_2\mathsf{aw}_x \mid x \in \Sigma\} \cup \{\mathbf{S}_2\mathsf{r}_{x,y} \mid (x, y) \in \Sigma\times\Sigma\} \cup \{\mathsf{aw}_x \mid x \in A\} \cup \{\mathsf{r}_{x,y} \mid (x, y) \in R\}.$$

This extends the associated valuation of the single-agent visibility language so as to include visibility atoms indexed with agent 2 for *all* awareness and attack variables. Conceptually, 2 is a part of the agent that always reasons under the assumption that the underlying AF is completed.

Proposition 4. *Let* $\mathsf{IAF} = (A, A^?, R, R^?)$ *be an IAF built over* Σ*, let* v_{IAF} *be its propositional valuation, let* M_{vis} *be the visibility model for* $\mathsf{ATM}^{\Sigma}_{\{1,2\}}$*, and let* $a \in A$*, then:*

- *The answer to the* st-*NCA (stable-Necessary-Credulous-Acceptance) problem with input* IAF *and* $a \in A$ *is yes iff* $M_{vis}, v_{\mathsf{IAF}} \models \mathbf{K}_1 \Diamond_2(\mathsf{Stable} \wedge \mathsf{in}_a)$.
- *The answer to the* st-*PCA (stable-Possible-Sceptical-Acceptance) problem with input* IAF *and* $a \in A$ *is yes iff* $M_{vis}, v_{\mathsf{IAF}} \models \hat{\mathbf{K}}_1 \Box_2(\mathsf{Stable} \rightarrow \mathsf{in}_a)$.

5 A Minimal Epistemic Logic for IAFs

As it is well-known, modal operators interpreted over equivalence relations (as the one we have just used in the logic of visibility) validate S5 axioms, which informally correspond to a factive, fully introspective kind of knowledge; see e.g. (Fagin et al. 2004, Chapter 2). However, one may want to use epistemic logic to reason about AFs without embracing such strong principles. Interestingly, the connection we have established in Section 3 can be used to unravel a *minimal epistemic logic for IAFs*. By doing so, we clarify the main epistemic assumptions underlying IAFs.

We start defining the logic semantically. Let Σ be an argument signature, Prop^{Σ} its associated set of propositions, and $\mathtt{Agt}$ a finite, non-empty set of agents. A **Kripke model** is a tuple $(W, \mathcal{R}, V)$ where $W \neq \emptyset$ is a set of *possible worlds*, $\mathcal{R} : \mathtt{Agt} \rightarrow \wp(W \times W)$ assigns to each agent an *epistemic accessibility relation* $\mathcal{R}_i$, and $V : \mathsf{Prop}^{\Sigma} \rightarrow \wp(W)$ is a *valuation function*. Since the current discussion is orthogonal to acceptance problems, we dispense again with acceptance variables and assume $\mathsf{Prop}^{\Sigma} = \{\mathsf{aw}_x \mid x \in \Sigma\} \cup \{\mathsf{r}_{x,y} \mid (x,y) \in \Sigma \times \Sigma\}$. Given a Kripke model $(W, \mathcal{R}, V)$ and a world $u \in W$, we define the **AF associated to** u as the tuple (A_u, R_u), where $A_u = \{x \in \Sigma \mid u \in V(\mathsf{aw}_x)\}$ and $R_u = \{(x,y) \in \Sigma \times \Sigma \mid u \in V(\mathsf{r}_{x,y})\} \cap (A_u \times A_u)$. Given the epistemic multi-agent language, the concepts of *truth* $(M, w \models \varphi)$, *validity in a model* $(M \models \varphi)$, and *validity* $(\models \varphi)$ are the usual ones in epistemic modal logic, see e.g. (Fagin et al. 2004, Chapter 2). Moreover, an **IAF-friendly Kripke model** is a Kripke model where for every $w \in W$, every $x, y \in \Sigma$, and every $i \in \mathtt{Agt}$:

(AWAR) if $\mathcal{R}_i[w] \cap V(\mathsf{r}_{x,y}) \neq \emptyset$, then $\mathcal{R}_i[w] \cap V(\mathsf{aw}_x) \neq \emptyset$ and $\mathcal{R}_i[w] \cap V(\mathsf{aw}_y) \neq \emptyset$; and

(COMP) there is an incomplete AF, IAF, built over Σ s.th.: the set of completions of IAF is equal to $\{(A_v, R_v) \mid v \in \mathcal{R}_i[w]\}$.

Informally, an IAF-friendly Kripke model is just a Kripke model where (i) if an agent considers possible that x attacks y then she considers possible being aware of both x and y; and (ii) the set of i-successors of each world represents the set of completions of an incomplete AF.

Now, consider the following two schemas:

(awar) $\hat{\mathbf{K}}_i \mathsf{r}_{x,y} \rightarrow (\hat{\mathbf{K}}_i \mathsf{aw}_x \wedge \hat{\mathbf{K}}_i \mathsf{aw}_y)$;

(comp) $(\hat{\mathbf{K}}_i l_1 \wedge \cdots \wedge \hat{\mathbf{K}}_i l_n) \rightarrow \hat{\mathbf{K}}_i (l_1 \wedge \cdots \wedge l_n)$,

where $\{l_1, ..., l_n\}$ is a consistent set of Prop^{Σ}-literals: elements of Prop^{Σ} or their negations (hence the set cannot contain both p and $\neg p$). Note that (comp) is true in every world of the single-agent visibility model and, moreover, it is the contra-positive of one of the axioms of the logic of visibility of (Herzig, Lorini, and Maffre 2018).

A sound and complete axiomatization with respect to IAF-friendly Kripke models can be obtained by simply adding all instances of (awar) and (comp) as well as axiom D: $\mathbf{K}_i \varphi \rightarrow \hat{\mathbf{K}}_i \varphi$ to the minimal modal proof system K.[2] A key for obtaining this result is the following claim:

Proposition 5. *A Kripke model* $M = (W, \mathcal{R}, V)$ *for* Prop^{Σ} *is an IAF-friendly model iff* $\mathcal{R}_i$ *is serial for every* $i \in \mathtt{Agt}$ *and all instances of (awar) and (comp) are valid in* M.

Discussion. The fact that the logic of IAF-friendly Kripke models is $KD + (comp) + (awar)$ reveals some interesting things. First, the minimal epistemic attitude underlying IAFs is a form of belief for which consistency is required (axiom D). Second, IAFs assume that all combinations of uncertain elements are considered as doxastically/epistemically possible by the agent (axiom (comp)). Equivalently, the belief/knowledge operator distributes over disjunctions of consistent literals.

6 Conclusion

To the best of our knowledge, this is the first paper studying the relations between IAFs and ELV, and it can be understood as an improvement of the reduction of IAFs to general epistemic logic provided in (Proietti and Yuste-Ginel 2021). We have briefly sketched two applications of this connection that seem conceptually valuable, as they dig into an epistemic interpretation of IAFs. As to future work, our findings can be used to provide a multi-agent generalization of IAFs. More precisely, we are now able to systematically investigate different multi-agent generalizations of IAFs, corresponding to different epistemic attitudes. A further perspective is to investigate the dynamics of IAFs in dynamic extensions of epistemic logic and ELV (cf. (Herzig and Yuste Ginel 2021)), taking inspiration from what has been done for Dung AFs in (Doutre, Maffre, and McBurney 2017; Doutre, Herzig, and Perrussel 2018).

Acknowledgements

Andreas Herzig is partially supported by the EU ICT-48 2020 project TAILOR (No. 952215). Antonio Yuste-Ginel gratefully acknowledges funding received from the PhD grant No. MECDFPU 2016/04113.

[2]We recall that axiom D is the modal schema axiomatizing seriality of each relation $\mathcal{R}_i$ and conceptually corresponding to consistency of each agents' belief. It is needed to axiomatise the logic of IAF-friendly Kripke models because, by definition, every IAF has a non-empty set of completions.

References

Amgoud, L., and Vesic, S. 2011. A new approach for preference-based argumentation frameworks. *Annals of Mathematics and Artificial Intelligence* 63(2):149–183.

Atkinson, K.; Baroni, P.; Giacomin, M.; Hunter, A.; Prakken, H.; Reed, C.; Simari, G.; Thimm, M.; and Villata, S. 2017. Towards artificial argumentation. *AI Magazine* 38(3):25–36.

Baumeister, D.; Neugebauer, D.; Rothe, J.; and Schadrack, H. 2018a. Complexity of verification in incomplete argumentation frameworks. In McIlraith, S. A., and Weinberger, K. Q., eds., *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18)*, 1753–1760. AAAI Press.

Baumeister, D.; Neugebauer, D.; Rothe, J.; and Schadrack, H. 2018b. Verification in incomplete argumentation frameworks. *Artificial Intelligence* 264:1–26.

Baumeister, D.; Järvisalo, M.; Neugebauer, D.; Niskanen, A.; and Rothe, J. 2021. Acceptance in incomplete argumentation frameworks. *Artificial Intelligence* 295:103470.

Baumeister, D.; Neugebauer, D.; and Rothe, J. 2018. Credulous and skeptical acceptance in incomplete argumentation frameworks. In *Proceedings of the COMMA 2018*, volume 305 of *Frontiers in AI and Applications*, 181–192. IOS Press.

Bench-Capon, T. J., and Dunne, P. E. 2007. Argumentation in artificial intelligence. *Artificial Intelligence* 171(10-15):619–641.

Besnard, P.; Cayrol, C.; and Lagasquie-Schiex, M.-C. 2020. Logical theories and abstract argumentation: A survey of existing works. *Argument & Computation* 11(1-2):41–102.

Cayrol, C., and Lagasquie-Schiex, M. 2005. On the acceptability of arguments in bipolar argumentation frameworks. In Godo, L., ed., *Symbolic and Quantitative Approaches to Reasoning with Uncertainty, 8th European Conference, (ECSQARU 2005)*, volume 3571 of *Lecture Notes in Computer Science*, 378–389. Springer.

Cooper, M. C.; Herzig, A.; Maffre, F.; Maris, F.; Perrotin, E.; and Régnier, P. 2021. A lightweight epistemic logic and its application to planning. *Artificial Intelligence* 298:103437.

Coste-Marquis, S.; Devred, C.; Konieczny, S.; Lagasquie-Schiex, M.-C.; and Marquis, P. 2007. On the merging of Dung's argumentation systems. *Artificial Intelligence* 171(10-15):730–753.

Doutre, S.; Herzig, A.; and Perrussel, L. 2018. Abstract argumentation in dynamic logic: Representation, reasoning and change. In Liao, B.; Ågotnes, T.; and Wáng, Y. N., eds., *Dynamics, Uncertainty and Reasoning, The Second Chinese Conference on Logic and Argumentation, (CLAR 2018)*, 153–185. Springer.

Doutre, S.; Maffre, F.; and McBurney, P. 2017. A dynamic logic framework for abstract argumentation: adding and removing arguments. In Benferhat, S.; Tabia, K.; and Ali, M., eds., *International Conference on Industrial, Engineering and Other Applications of Applied Intelligent Systems.*, volume 10351 of *LNCS*, 295–305. Springer.

Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77(2):321–357.

Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. 2004. *Reasoning about knowledge*. MIT Press.

Fazzinga, B.; Flesca, S.; and Furfaro, F. 2020. Revisiting the notion of extension over incomplete abstract argumentation frameworks. In *Proceedings of IJCAI-20*, 1712–1718. AAAI Press.

Herzig, A., and Yuste Ginel, A. 2021. Multi-agent abstract argumentation frameworks with incomplete knowledge of attacks. In *Proceedings of IJCAI-21*. AAAI Press.

Herzig, A.; Lorini, E.; and Maffre, F. 2018. Possible worlds semantics based on observation and communication. In van Ditmarsch, H., and Sandu, G., eds., *Jaakko Hintikka on Knowledge and Game-Theoretical Semantics*. Cham: Springer International Publishing. 339–362.

Li, H.; Oren, N.; and Norman, T. J. 2011. Probabilistic argumentation frameworks. In *Proceedings of TAFA*, volume 7312 of *LNCS*, 1–16. Springer.

Proietti, C., and Yuste-Ginel, A. 2021. Dynamic epistemic logics for abstract argumentation. *Synthese*. DOI:10.1007/s11229-021-03178-5.

van der Hoek, W.; Troquard, N.; and Wooldridge, M. J. 2011. Knowledge and control. In *Proceedings of AAMAS*, 719–726.