Ranking Functions over Labelings

Tjitze RIENSTRA, Matthias THIMM
University of Koblenz-Landau, Germany

Abstract. We study rankings over labelings as a generalization of traditional labeling-based semantics in abstract argumentation. Our approach is an alternative to recent developments on rankings over arguments. The formal basis is a qualitative abstraction of probability theory called ranking theory. As a guiding principle in determining rankings over labelings, we interpret argumentation frameworks similarly to ranking networks (the ranking-theoretic equivalent of Bayesian networks).

1. Introduction

Motivated by the fact that the usual distinction between acceptance and rejection of arguments is too coarse grained for many applications, various approaches to ranking-based semantics have been developed in recent years (see [12] for a survey and [10,1,2,7] for some concrete approaches). The solution proposed in these approaches is to rank arguments from most acceptable to least acceptable.

In this paper we investigate an alternative approach. Instead of ranking arguments, we rank labelings. As a formal basis we rely on *ranking theory* [20], a qualitative abstraction of probability based on *ranking functions*. These are functions that associate possible worlds with non-negative integers or ∞. These values, called *ranks*, represent degrees of surprise: 0 for not surprising, 1 for surprising, 2 very surprising, and so on. Like conditional probabilities, ranks give rise to conditional ranks, which can be used to model dynamics of belief.

We consider the notion of a *ranking-theoretic semantics*, which associates each argumentation framework with a ranking function over labelings. As a guiding principle in determining such rankings, we interpret argumentation frameworks similarly to ranking networks. These are the ranking-theoretic equivalent of Bayesian networks, i.e. directed acyclic graphs in which nodes and edges represent variables and dependencies. While straightforward in the acyclic case, the presence of cycles in argumentation frameworks complicates matters. For this case we introduce the notion of SCC stratification, which is closely related to the SCC decomposability property studied in argumentation [3,5].

Although our approach is based on rankings over *labelings*, a ranking over labelings can be translated into a ranking over *arguments* by considering marginal or absolute ranks of arguments. To explore this connection we recall a number of properties that have been proposed in the literature on rankings over arguments, and we check whether they are satisfied by our approach.

Overview of this paper: After presenting preliminaries in Section 2 we turn in Section 3 to the idea of interpreting argumentation frameworks like ranking networks, leading to the definition of SCC stratification. We then present in Section 4 a general scheme to define a ranking-theoretic semantics on the basis of a labeling-based semantics. In Section 5 we determine the conditions under which this scheme satisfies SCC stratification. In Section 6 we make a comparison with rankings over arguments. We discuss related work in Section 7 and conclude in Section 8.

2. Preliminaries

In this section we present the necessary basics of abstract argumentation and ranking theory. We start with a some graph-theoretical notions that we use.

Definition 1. A *directed graph* is an ordered pair $G = (V, \rightarrow)$ where V is a set of vertices and $\rightarrow \subseteq V \times V$ a set of edges. We write $x \rightarrow y$ whenever $(x, y) \in \rightarrow$.

Definition 2. Let $G = (V, \rightarrow)$ be a directed graph, let $x \in V$ and $B \subseteq V$.

- x is a *parent* of y iff $x \rightarrow y$. We denote by $Pa_G(x)$ the set of parents of x.
- x is a *source vertex* iff it has no parent. B is a *source set* iff all its members are source vertices.
- x is a *descendant* of y iff $x = y$ or there is a directed path from y to x.
- x is a *descendant* of B iff x is a descendant of some $y \in B$. The descendants of B are denoted by $D_G(B)$. We also write $D_G(x)$ instead of $D_G(\{x\})$.
- x is *non-descendant* of B iff x is not a descendant of B. The non-descendants of B are denoted by $ND_G(B)$. We also write $ND_G(x)$ instead of $ND_G(\{x\})$.
- x is an *outparent* of B iff x is a parent of some $y \in B$ and $x \notin B$. The set of outparents of B is denoted by $OP_G(B)$.
- The *context* of B is the set $B \cup OP_G(B)$ and is denoted by $C_G(B)$.
- The *restriction of G to B* is a new directed graph $(B, \rightarrow \cap B \times B)$ and is denoted by $G{\downarrow}B$.
- The set of *SCCs* (*strongly connected components*) of G, denoted $\mathrm{SCC}(G)$, contains all equivalence classes induced by the *path equivalence relation* $\sim_G$ defined by $x \sim_G y$ iff x and y are each others descendants.

2.1. Abstract Argumentation

The basic notion in abstract argumentation is an *argumentation framework* [14].

Definition 3. An argumentation framework (AF) is a directed graph $F = (A, \rightsquigarrow)$ whose vertices are called *arguments* and whose edges are called *attacks*.

We restrict our attention in this paper to AFs with finite sets of arguments. A *semantics* determines, given an AF, rational points of view on which arguments can be accepted. In this paper we rely on the three-valued labeling-based definition of a semantics [9], where a *labeling* maps each argument to a label **I** (*in* or *accepted*) **O** (*out* or *rejected*) or **U** (undecided).

Definition 4. A *labeling* of a set A is a function $\mathbf{L} : A \to \{\mathbf{I}, \mathbf{O}, \mathbf{U}\}$. We denote by $\mathcal{L}(A)$—or, if $F = (A, \rightsquigarrow)$, by $\mathcal{L}(F)$—the set of labelings of A. Given a labeling $\mathbf{L}$ of A and set $B \subseteq A$ we denote by $\mathbf{L}{\downarrow}B$ the *restriction of* $\mathbf{L}$ *to* B. We denote by $\mathbf{I}(\mathbf{L})$ (resp. $\mathbf{O}(\mathbf{L})$, $\mathbf{U}(\mathbf{L})$) the set of arguments labeled $\mathbf{I}$ (resp. $\mathbf{O}$, $\mathbf{U}$) by $\mathbf{L}$.

The semantics we consider in this paper are based on *complete* labelings [9].

Definition 5. Let $F = (A, \rightsquigarrow)$ be an AF. A labeling $\mathbf{L} \in \mathcal{L}(F)$ is a *complete* labeling of F iff, for all $x \in A$:

1. If $\forall y \in Pa_F(x)$, $\mathbf{L}(y) = \mathbf{O}$, then $\mathbf{L}(x) = \mathbf{I}$.
2. If $\exists y \in Pa_F(x)$ s.t. $\mathbf{L}(y) = \mathbf{I}$, then $\mathbf{L}(x) = \mathbf{O}$.
3. If $\forall y \in Pa_F(x)$, $\mathbf{L}(y) \neq \mathbf{I}$ and $\exists y \in Pa_F(x)$ s.t. $\mathbf{L}(y) = \mathbf{U}$, then $\mathbf{L}(x) = \mathbf{U}$.

A semantics σ maps each AF F to a set $\mathcal{L}_\sigma(F) \subseteq \mathcal{L}(F)$. Besides the complete semantics we consider the *preferred*, *grounded* and *semi-stable* semantics, which yield, respectively, the $\mathbf{I}$-maximal, $\mathbf{I}$-minimal and $\mathbf{U}$-minimal complete labelings.

Definition 6. A *semantics* σ associates each AF $F = (A, \rightsquigarrow)$ with a set $\mathcal{L}_\sigma(F) \subseteq \mathcal{L}(A)$ of labelings. The **co** (complete), **pr** (preferred), **gr** (grounded) and **ss** (semi-stable) semantics are defined by

$$\begin{aligned}
\mathcal{L}_{\mathbf{co}}(F) &= \{\mathbf{L} \in \mathcal{L}(F) \mid \mathbf{L} \text{ is a complete labeling of } F\},\\
\mathcal{L}_{\mathbf{pr}}(F) &= \{\mathbf{L} \in \mathcal{L}_{\mathbf{co}}(F) \mid \nexists \mathbf{L}' \in \mathcal{L}_{\mathbf{co}}(F) \text{ s.t. } \mathbf{I}(\mathbf{L}) \subset \mathbf{I}(\mathbf{L}')\},\\
\mathcal{L}_{\mathbf{gr}}(F) &= \{\mathbf{L} \in \mathcal{L}_{\mathbf{co}}(F) \mid \nexists \mathbf{L}' \in \mathcal{L}_{\mathbf{co}}(F) \text{ s.t. } \mathbf{I}(\mathbf{L}) \supset \mathbf{I}(\mathbf{L}')\},\\
\mathcal{L}_{\mathbf{ss}}(F) &= \{\mathbf{L} \in \mathcal{L}_{\mathbf{co}}(F) \mid \nexists \mathbf{L}' \in \mathcal{L}_{\mathbf{co}}(F) \text{ s.t. } \mathbf{U}(\mathbf{L}) \supset \mathbf{U}(\mathbf{L}')\}.
\end{aligned}$$

Under all these semantics, the existence of at least one labeling is guaranteed (in the finite case). We omit the stable semantics to avoid technical difficulties due to the possible non-existence of a labeling. The grounded labeling is unique.

2.2. Ranking Theory

Ranking theory is a qualitative abstraction of probability theory in which events receive *ranks* [16,20]. A rank is a non-negative integer or ∞ and can be understood as a degree of surprise: 0 for not surprising, 1 for surprising, 2 for very surprising, and so on, and ∞ for impossible. The central notion in ranking theory is a *ranking function* (also known as an *ordinal conditional function* or *kappa function*).

Definition 7. A ranking function over a set Ω is a function $\kappa : \Omega \to \mathbb{N} \cup \{\infty\}$ such that $\kappa(w) = 0$ for at least one $w \in \Omega$. A ranking function κ is extended to a function over propositions or events (i. e., subsets of Ω) by defining $\kappa(X) = \infty$ if $X = \emptyset$, and $\kappa(X) = min(\{\kappa(w) \mid w \in X\})$, otherwise.

Like probabilities ranks give rise to *conditional ranks*. These are ranks of events given that we learn that some other event occurred. Like [16] we define any rank conditional on an impossible event to be ∞.

Definition 8. Let κ be a ranking function over Ω and let $X, Y \subseteq \Omega$. We define the *rank of X conditional on Y*, denoted $\kappa(X \mid Y)$, by $\kappa(X \mid Y) = \kappa(X \cap Y) - \kappa(Y)$ if $\kappa(Y) < \infty$, and $\kappa(X \mid Y) = \infty$, otherwise.

A ranking function κ induces beliefs using the principle that X is believed iff the complement $\overline{X} = \Omega \setminus X$ is surprising (i.e. $\kappa(\overline{X}) > 0$). Similarly, X is believed *conditional on Y* iff $\kappa(\overline{X} \mid Y) > 0$.

2.3. Ranking Networks

A *ranking network* (also known as an *OCF-network*) is the ranking-theoretic analogue of a Bayesian network [16]. Like a Bayesian network, a ranking network is a directed acyclic graph (DAG) that encodes conditional independence relationships among variables and can be used to compactly represent a ranking function.

To define it, we assume that Ω is determined by a set of variables. This requires some notation. Let $\mathcal{X}$ be a finite set of variables. For simplicity we assume each variable to have the same domain $Dom(\mathcal{X})$. A *valuation* of a set $B \subseteq \mathcal{X}$ is a function $\mathbf{V}$ from B to $\mathrm{Dom}(\mathcal{X})$. We say that Ω is determined by $\mathcal{X}$ iff Ω consists of all valuations of $\mathcal{X}$. If $\mathbf{V}$ is a valuation of B and $B' \subseteq B$, we denote by $\mathbf{V}{\downarrow}B'$ the restriction of $\mathbf{V}$ to B'. If $x \in \mathcal{X}$ and $v \in \mathrm{Dom}(\mathcal{X})$, we denote by $x =_\Omega v$ the set $\{\mathbf{V} \in \Omega \mid \mathbf{V}(x) = v\}$ (i.e., the event that x equals v), omitting the subscript Ω if clear from context. Similarly, $x \neq_\Omega v$ denotes the event that x does not equal v. If $\mathbf{V}$ is a valuation of B we also denote (abusing notation) by $\mathbf{V}$ the set $\cap_{x \in B} x = \mathbf{V}(x)$. A ranking network over $\mathcal{X}$ is a DAG $G = (\mathcal{X}, \rightarrow)$. A ranking function is said to be *stratified* w.r.t. G if it can be decomposed into the sum of conditional ranks corresponding to each variable given its parents [16]:

Definition 9. If Ω is determined by the set $\mathcal{X}$ of variables then a ranking function κ over Ω is *stratified* w.r.t. a ranking network $G = (\mathcal{X}, \rightarrow)$ iff for all $\mathbf{V} \in \Omega$,

$$\kappa(\mathbf{V}) = \sum_{x \in \mathcal{X}} \kappa(x = \mathbf{V}(x) \mid \mathbf{V}{\downarrow}Pa_G(x)). \tag{1}$$

Stratification w.r.t. a ranking network implies that each variable depends only on its parents or, more precisely, is independent of its non-descendants given its parents [16]. This independence condition resembles the probabilistic one [18].

Proposition 1. *Let Ω be determined by the set $\mathcal{X}$ of variables and let κ be a ranking function over Ω. If κ is stratified w.r.t. G then, for each $x \in \mathcal{X}$ and $v \in Dom(\mathcal{X})$, and each valuation $\boldsymbol{Pa}_x$ of $Pa_G(x)$ and $\boldsymbol{ND}_x$ of $ND_G(x)$ that agree on the values assigned to $Pa_G(x)$*[1]*, we have*

$$\kappa(x = v \mid \boldsymbol{ND}_x) = \kappa(x = v \mid \boldsymbol{Pa}_x) \text{ whenever } \kappa(\boldsymbol{ND}_x) < \infty.$$

A *conditional ranking table* (CRT) plays the same role as a conditional probability table in the probabilistic setting [18]. Given a ranking function κ over a

[1] Note that, because G is acyclic, we have $Pa_G(x) \subseteq ND_G(x)$.

set Ω determined by variables $\mathcal{X}$ with domain $\{v_1, \ldots, v_n\}$, a CRT for a variable $x \in \mathcal{X}$ specifies, for each valuation $\boldsymbol{Pa}_x$ of $Pa_G(x)$, the conditional ranks

$$\kappa(x = v_1 \mid \boldsymbol{Pa}_x), \ldots, \kappa(x = v_n \mid \boldsymbol{Pa}_x),$$

such that at least one conditional rank per valuation $\boldsymbol{Pa}_x$ is zero. If κ is stratified w.r.t. G and if we know the CRTs of all variables, then κ is uniquely determined, because for each $\mathbf{V} \in \Omega$ we can fill in the terms on the right-hand side of Eq. (1).

3. SCC stratification

Our aim is to generalize labeling-based semantics to ranking-theoretic semantics, which associate each AF with a ranking function over labelings. The idea is that the rank of a labeling represents its degree of surprise, and that external information about the status of an argument is processed via conditionalization.

In the next section we present a constructive definition of a family of ranking-theoretic semantics. Apart from a constructive definition, however, one may ask: which properties should such a semantics satisfy? Such properties (like principles studied in abstract argumentation [4]) can then be used to check whether a particular semantics behaves as desired.

In this section we propose one such property, called *SCC stratification*. We will argue that any reasonable ranking function over labelings of an AF should be SCC stratified and, hence, that any ranking-theoretic semantics should associate each AF F with a ranking function that is SCC stratified.

3.1. Stratified Rankings over Labelings

We first limit our attention to acyclic AFs. A plausible requirement for a ranking function κ over the labelings of an acyclic AF F is that the status of each argument x is independent of the status of its nondescendants given the status of its parents. In other words, the status of an argument x depends only on the status of its attackers. We already know how to express this requirement formally: κ should be stratified w.r.t. F. That is, for each $\mathbf{L} \in \mathcal{L}(F)$,

$$\kappa(\mathbf{L}) = \sum_{x \in A} \kappa(x = \mathbf{L}(x) \mid \mathbf{L}{\downarrow}Pa_F(x)).$$

This implies, via proposition 1, that the status of an argument is independent of the status of its nondescendants given the status of its parents.

We can now determine a ranking over labelings of F by filling in the CRTs for each argument. These CRTs determine *how* the status of each argument depends on the status of its attackers. Here we present a possible scheme to fill in these CRTs. Apart from serving as an example, we will later show that this particular scheme also makes sense formally.

Definition 10. Let $F = (A, \rightsquigarrow)$ be an acyclic AF and let κ be a ranking function over $\mathcal{L}(F)$. We say that κ is *rank-complete* w.r.t. F iff, for each $x \in A$ and each labeling $\boldsymbol{Pa}_x$ of $Pa_F(x)$,

(1) If for all $y \in Pa_F(x)$, $\boldsymbol{Pa}_x(y) = \mathbf{O}$, then

$$\kappa(x = \mathbf{O} \mid \boldsymbol{Pa}_x) = 1, \qquad \kappa(x = \mathbf{U} \mid \boldsymbol{Pa}_x) = \infty, \qquad \kappa(x = \mathbf{I} \mid \boldsymbol{Pa}_x) = 0.$$

(2) If for some $y \in Pa_F(x)$, $\boldsymbol{Pa}_x(y) = \mathbf{I}$, then

$$\kappa(x = \mathbf{O} \mid \boldsymbol{Pa}_x) = 0, \qquad \kappa(x = \mathbf{U} \mid \boldsymbol{Pa}_x) = \infty, \qquad \kappa(x = \mathbf{I} \mid \boldsymbol{Pa}_x) = \infty.$$

(3) If for all $y \in Pa_F(x)$, $\boldsymbol{Pa}_x(y) \neq \mathbf{I}$ and for some $y \in Pa_F(x)$, $\boldsymbol{Pa}_x(y) = \mathbf{U}$, then

$$\kappa(x = \mathbf{O} \mid \boldsymbol{Pa}_x) = 1, \qquad \kappa(x = \mathbf{U} \mid \boldsymbol{Pa}_x) = 0, \qquad \kappa(x = \mathbf{I} \mid \boldsymbol{Pa}_x) = \infty.$$

This definition consists of three parts: (1) If we learn that all attackers of x are **O**, then x is normally **I**, surprisingly **O**, but never **U**; (2) If we learn that an attacker of x is **I**, then x is always **O**; (3) If we learn that no attacker of x is **I** but at least one is **U**, then x is normally **U**, surprisingly **O**, but never **I**. These conditions mirror those of completeness (Def. 5) except that they leave open the surprising possibility that an argument is **O** while no attacker is **I**.

Stratification plus rank-completeness uniquely determines a ranking over labelings. Thus, together, they form a ranking-theoretic semantics, although *only for acyclic AFs.* This semantics generalizes the **gr** semantics in the following sense.

Proposition 2. *Let $F = (A, \rightsquigarrow)$ be an AF and let κ be a ranking function over $\mathcal{L}(F)$. If κ is stratified and rank-complete w.r.t. F then $\kappa(\mathbf{L}) = 0$ iff $\mathbf{L}$ is the grounded labeling of F.*

Ranking functions that are stratified and rank-complete encode a general principle that the rank of a labeling equals the number of violations that it contains (i. e., arguments being rejected while no attacker is accepted). Thus, (conditional) beliefs are based on the principle of minimizing such violations.

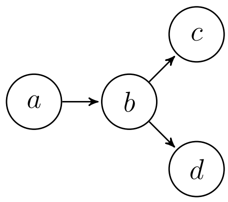

Figure 1.: An acyclic AF F

	a	b	c	d	Rank
$\mathbf{L}_1$	**I**	**O**	**I**	**I**	0
$\mathbf{L}_2$	**O**	**I**	**O**	**O**	1
$\mathbf{L}_4$	**I**	**O**	**O**	**I**	1
$\mathbf{L}_5$	**I**	**O**	**I**	**O**	1
$\mathbf{L}_3$	**O**	**O**	**I**	**I**	2
$\mathbf{L}_6$	**I**	**O**	**O**	**O**	2
$\mathbf{L}_7$	**O**	**O**	**O**	**O**	4

Table 1.: A ranking function over $\mathcal{L}(F)$

Example 1. Let F be the AF shown in Figure 1. Table 1 shows the unique ranking κ over $\mathcal{L}(F)$ that is stratified and rank-complete w.r.t. F (labelings with rank ∞ are omitted). Note that only the grounded labeling $\mathbf{L}_1$ is ranked zero. We have, for example, that rejection of c and d leads to belief in rejection of a: $\kappa(a \neq \mathbf{O} \mid c = \mathbf{O} \cap d = \mathbf{O}) = 1$ (rejection of a explains rejection of c and d with the least number of violations). Because κ is stratified we have, for example, that if we know that b is rejected then learning that c or a is also rejected does not affect our belief about d:

$$\kappa(d = \mathbf{I} \mid b = \mathbf{O} \cap c = \mathbf{O}) = \kappa(d = \mathbf{I} \mid b = \mathbf{O} \cap a = \mathbf{O}) = \kappa(d = \mathbf{I} \mid b = \mathbf{O}).$$

3.2. SCC Stratified Rankings over Labelings

We now introduce the notion of SCC stratification as a generalisation of stratification that enables us to deal with general (i.e., possibly cyclic) AFs. To motivate it we first discuss the *SCC decomposability* property known from abstract argumentation. This property is what Baroni et al. [3] call *full decomposability w.r.t. SCC partitioning.* Here we simplify their definition somewhat. A semantics σ is said to be SCC decomposable if the σ labelings of an AF F can be computed seperately for *each SCC S given the labelings of the outparents of S*:

Definition 11. A *local function* $\mathbb{L}$ is a function that assigns to each AF $F = (A, \rightsquigarrow)$, each source set I of F, and each labeling $\mathbf{L}_I \in \mathcal{L}(I)$, a set $\mathbb{L}(F, I, \mathbf{L}_I) \subseteq \mathcal{L}(A \setminus I)$. A semantics σ is *SCC decomposable* iff there exists a local function $\mathbb{L}$ such that, for each AF $F = (A, \rightsquigarrow)$ and each $\mathbf{L} \in \mathcal{L}(A)$,

$$\mathbf{L} \in \mathcal{L}_\sigma(F) \text{ iff } \forall S \in \mathrm{SCC}(F), \mathbf{L}{\downarrow}S \in \mathbb{L}(F{\downarrow}C_F(S), OP_F(S), \mathbf{L}{\downarrow}OP_F(S)). \quad (2)$$

Of the semantics we consider, only semi-stable fails SCC decomposability [3].

Proposition 3. *The* **co**, **pr** *and* **gr** *semantics are SCC decomposable but the* **ss** *semantics is not.*

We can similarly require, of a ranking function over labelings of F, that it can be decomposed into the sum of conditional ranks corresponding to *each SCC S of F given the outparents of S*. This is what we call SCC stratification. Formally:

Definition 12. Let F be an AF. A ranking function κ over $\mathcal{L}(F)$ is *SCC stratified* w.r.t. F iff, for each $\mathbf{L} \in \mathcal{L}(A)$,

$$\kappa(\mathbf{L}) = \sum_{S \in \mathrm{SCC}(F)} \kappa(\mathbf{L}{\downarrow}S \mid \mathbf{L}{\downarrow}OP_F(S)). \quad (3)$$

SCC stratification implies that each SCC S is independent of its non-descendant given its outparents and reduces to stratification in the acyclic case.

Proposition 4. *If κ is SCC stratified w.r.t. F then, for all $S \in SCC(F)$ and all valuations $\mathbf{S}$ of S, $\boldsymbol{OP}_S$ of $OP_F(S)$, and $\boldsymbol{ND}_S$ of $ND_F(S)$ that agree on the values assigned to $OP_F(S)$*[2]*, we have*

$$\kappa(\mathbf{S} \mid \boldsymbol{ND}_S) = \kappa(\mathbf{S} \mid \boldsymbol{OP}_S) \text{ whenever } \kappa(\boldsymbol{ND}_S) < \infty. \quad (4)$$

Proposition 5. *Let F be an AF and κ a ranking function over $\mathcal{L}(F)$. If F is acyclic then κ is SCC stratified w.r.t. F iff κ is stratified w.r.t. F.*

We believe that, in the presence of cycles, SCC stratification adequately captures the (in)dependence relationships between (sets of) arguments that should hold for a ranking function over labelings. After presenting a constructive definition of a family of ranking-theoretic semantics, we will establish, in section 5, the conditions under which these semantics yield SCC stratified rankings.

[2] Note that for each SCC S we have $OP_F(S) \subseteq ND_F(S)$.

4. The σ^* Ranking-Theoretic Semantics

We now define the general notion of a *ranking-theoretic semantics*, which associates each AF F with a ranking function over $\mathcal{L}(F)$.

Definition 13. A ranking-theoretic semantics ρ maps each AF $F = (A, \rightsquigarrow)$ to a ranking function over $\mathcal{L}(A)$ denoted by $K_\rho(F)$.

Here we propose a family of ranking-theoretic semantics based on a scheme that turns any semantics σ into a ranking-theoretic semantics denoted by σ^*. First a definition. We say that a set of arguments B *σ-enforces* a labeling $\mathbf{L}$ (written $(F, B, \sigma) \rightarrow \mathbf{L}$) if adding a new argument attacking the members of B turns $\mathbf{L}$ into a σ labeling. This is a special case of enforcement [6]. We assume that there is a unique argument (denoted Q) playing the role of the added attacker.

Definition 14. Let $F = (A, \rightsquigarrow)$ be an AF. Give a set $B \subseteq A$ we denote by $F \blacktriangleright B$ the AF $(A \cup \{Q\}, \rightsquigarrow \cup \{(Q, x) \mid x \in B\})$. Given a semantics σ, we say that B *σ-enforces* a labeling $\mathbf{L}$ (written $((A, \rightsquigarrow), B, \sigma) \rightarrow \mathbf{L}$) iff $\mathbf{L} \in \mathcal{L}_\sigma(F \blacktriangleright B){\downarrow}A$.

The σ^* semantics assigns to each labeling a rank that equals the size of a cardinality-wise minimal set that σ-enforces $\mathbf{L}$. Thus, the rank of a labeling is the minimal number of arguments we need to attack to turn it into a σ labeling.

Definition 15. Let σ be a semantics. Let $F = (A, \rightsquigarrow)$ be an AF. The ranking-theoretic semantics σ^* of F is defined by

$$K_{\sigma^*}(F)(\mathbf{L}) = min(\{|B| \mid B \subseteq A, ((A, \rightsquigarrow), B, \sigma) \rightarrow \mathbf{L}\} \cup \{\infty\}).$$

The σ^* semantics generalizes the σ semantics in the sense that $K_{\sigma^*}(F)(\mathbf{L}) = 0$ iff $\mathbf{L} \in \mathcal{L}_\sigma(F)$. Furthermore, in the acyclic case, the $\mathbf{gr}^*$ semantics (as well as the $\mathbf{co}^*$, $\mathbf{pr}^*$ and $\mathbf{ss}^*$ semantics, which all coincide in this case) is completely characterized by stratification plus rank-completeness.

Theorem 1. *Let $F = (A, \rightsquigarrow)$ be an acyclic AF and let κ be a ranking function over $\mathcal{L}(F)$. If κ is stratified and rank-complete w.r.t. F then $\kappa = K_{\mathbf{gr}}(F)$.*

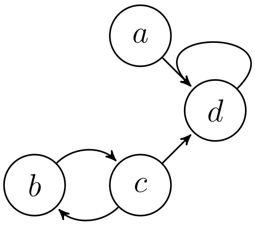

Figure 2.: The AF F

	a	b	c	d	Rank
$\mathbf{L}_1$	**I**	**I**	**O**	**O**	0
$\mathbf{L}_2$	**I**	**O**	**I**	**O**	0
$\mathbf{L}_3$	**I**	**U**	**U**	**O**	0
$\mathbf{L}_4$	**O**	**I**	**O**	**U**	1
$\mathbf{L}_5$	**O**	**O**	**I**	**O**	1
$\mathbf{L}_6$	**O**	**U**	**U**	**U**	1
$\mathbf{L}_7$	**I**	**O**	**O**	**O**	2
$\mathbf{L}_8$	**O**	**I**	**O**	**O**	2
$\mathbf{L}_9$	**O**	**U**	**U**	**O**	2
$\mathbf{L}_{10}$	**O**	**O**	**O**	**U**	3
$\mathbf{L}_{11}$	**O**	**O**	**O**	**O**	4

Table 2.: $K_{\mathbf{co}^*}(F)$

	a	b	c	d	Rank
$\mathbf{L}_1$	**I**	**O**	**I**	**O**	0
$\mathbf{L}_2$	**I**	**I**	**O**	**O**	0
$\mathbf{L}_3$	**O**	**O**	**I**	**O**	1
$\mathbf{L}_4$	**O**	**I**	**O**	**U**	2
$\mathbf{L}_5$	**I**	**O**	**O**	**O**	2
$\mathbf{L}_6$	**O**	**I**	**O**	**O**	2
$\mathbf{L}_7$	**O**	**I**	**O**	**O**	3
$\mathbf{L}_8$	**O**	**O**	**O**	**O**	4

Table 3.: $K_{\mathbf{ss}^*}(F)$

Example 2. Let F be the AF shown in Figure 2. Table 2 shows the ranking $K_{\mathbf{co}^*}(F)$ (labelings with rank ∞ are omitted). To illustrate, this ranking encodes that, if we learn that a and c are rejected, we believe that d is undecided: $K_{\mathbf{co}^*}(F)(d \neq \mathbf{U} \mid a = \mathbf{O} \cap c = \mathbf{O}) > 0$. On the other hand, learning the status of a does not affect beliefs about b and c. That is, for all $l_a, l_b, l_c \in \{\mathbf{I}, \mathbf{O}, \mathbf{U}\}$ we have

$$K_{\mathbf{co}^*}(F)(b = l_b \cap c = l_c \mid a = l_a) = K_{\mathbf{co}^*}(F)(b = l_b \cap c = l_c).$$

5. SCC Stratification and the σ^* Semantics

Does the σ^* semantics produce SCC stratified ranking functions? The answer is yes, provided that σ is SCC decomposable.

Theorem 2. *If a semantics σ satisfies SCC decomposability then for each AF F, $K_{\sigma^*}(F)$ is SCC stratified w.r.t. F.*

If σ does not satisfy SCC decomposability, then σ^* may produce ranking functions that are not SCC stratified. This is demonstrated by the following example, which is based on the non-SCC-decomposable semi-stable semantics.

Example 3. Table 3 shows the ranking $K_{\mathbf{ss}^*}(F)$ for F shown in Figure 2 (labelings with rank ∞ are omitted). The SCC $\{b, c\}$ has non-descendant a and no outparents. If $K_{\mathbf{ss}^*}(F)$ were SCC-stratified w.r.t. F then Proposition 4 would imply

$$K_{\mathbf{ss}^*}(F)(b = \mathbf{I} \cap c = \mathbf{O} \mid a = \mathbf{O}) = K_{\mathbf{ss}^*}(F)(b = \mathbf{I} \cap c = \mathbf{O}).$$

In words: b and c are independent of a. But this is false, because we have

$$K_{\mathbf{ss}^*}(F)(b = \mathbf{I} \cap c = \mathbf{O} \mid a = \mathbf{O}) = 1 \text{ and } K_{\mathbf{ss}^*}(F)(b = \mathbf{I} \cap c = \mathbf{O}) = 0.$$

Thus, $K_{\mathbf{ss}^*}(F)$ is not SCC stratified. A more direct way to verify this is to check condition (3), which implies that

$$\begin{aligned} K_{\mathbf{ss}^*}(F)(\mathbf{L}_4) = {} & K_{\mathbf{ss}^*}(F)(a = \mathbf{O}) + K_{\mathbf{ss}^*}(F)(b = \mathbf{I} \cap c = \mathbf{O}) + \\ & K_{\mathbf{ss}^*}(F)(d = \mathbf{U} \mid a = \mathbf{O} \cap c = \mathbf{O}) \\ = {} & 1 + 0 + 0 = 1. \end{aligned}$$

This is false, as shown in Table 3. Intuitively, the failure of SCC stratification here is due to the fact that, if we reject a, we must accept c to prevent d from becoming undecided, which changes our initial belief in $b = \mathbf{I} \cap c = \mathbf{O}$.

6. Rankings Over Arguments

Our approach, which is based on ranking labelings, can be seen as an alternative to recent developments based on rankings over arguments [10,1,2,7,12]. In these rankings arguments are ordered from “least acceptable” to “most acceptable” (we limit our attention here to total preorders). A ranking over labelings can, however, be turned into a ranking over arguments in a straightforward way.

Definition 16. Let σ be a semantics and $F = (A, \rightsquigarrow)$ an AF. Define the preorder $\preceq_F^{\sigma^*} \subseteq A \times A$ via $a \preceq_F^{\sigma^*} b$ iff $K_{\sigma^*}(F)(a = \mathbf{I}) \leq K_{\sigma^*}(F)(b = \mathbf{I})$.

In other words, $a \preceq_F^{\sigma^*} b$ ("a is at least as acceptable as b") if the marginal rank of the event that a is accepted is less or equal than the marginal rank of the event that b is accepted. How does $\preceq_F^{\sigma^*}$ compare to other proposals? Let us recall a number of *rationality postulates* that have been proposed and that capture various intuitions behind ranked acceptability [7]. We leave a more in-depth study to future work and consider only four rather simple ones.

Let $\preceq_F$ be some ranking over the arguments of F. Two AFs $F = (A, \rightsquigarrow)$ and $F' = (A', \rightsquigarrow')$ are *isomorphic* (written $F \equiv F'$) if there is a bijective function $\gamma : A \to A'$ such that $a \rightsquigarrow b$ iff $\gamma(a) \rightsquigarrow' \gamma(b)$ for all $a, b \in A$ (γ is then called an *isomorphism*). The first postulate, called abstraction, states that names of arguments play no role in assigning ranks.

Abstraction For every pair $F = (A, \rightsquigarrow)$, $F' = (A', \rightsquigarrow')$ of isomorphic frameworks and every isomorphism $\gamma : A \to A'$, for all $a, b \in A$, $a \preceq_F b$ iff $\gamma(a) \preceq_{F'} \gamma(b)$.

A *connected component* of F is a maximal subgraph $F' = (A', \rightsquigarrow')$ of F such that every pair $a, b \in A'$ is connected through a path while *ignoring edge directions*. Let $cc(F)$ be the set of all connected components of F. *Independence* states that the evaluation of each connected component is independent.

Independence For every $F = (A, \rightsquigarrow)$ and $F' = (A', \rightsquigarrow') \in cc(F)$, for all $a, b \in A'$, $a \preceq_{F'} b$ iff $a \preceq_F b$.

Void precedence states that unattacked arguments are ranked higher than others.

Void precedence For every $F = (A, \rightsquigarrow)$, for all $a, b \in A$, if a is not attacked and b is attacked then $b \not\preceq_F a$.

Self-contradiction says that self-attacking arguments are ranked lower than others.

Self-contradiction For every $F = (A, \rightsquigarrow)$, for all $a, b \in A$, if $b \rightsquigarrow b$ and not $a \rightsquigarrow a$ then $a \preceq_F b$ and $b \not\preceq_F a$.

It turns out that our semantics satisfies three of these four postulates.

Proposition 6. *Let* $\sigma \in \{\mathbf{co}, \mathbf{gr}, \mathbf{pr}, \mathbf{ss}\}$ *and* $F = (A, \rightsquigarrow)$ *an AF. The ranking* $\preceq_F^{\sigma^*}$ *satisfies* abstraction*,* independence*, and* self-contradiction*.*

To see why $\preceq_F^{\sigma^*}$ does not satisfy void precedence, consider an AF with arguments a, b, c where a attacks b and b attacks c. We then have $a \preceq_F^{\sigma^*} c$ and $c \preceq_F^{\sigma^*} a$ for all semantics σ.

7. Related Work

Addressing belief revision in argumentation is often done, following the AGM approach, using orderings over extensions or labelings [8,11,13,19]. Our approach fits into this line of work, as it also provides an account of revision in argumentation based on orderings. The notion of SCC stratification is novel in this context, although the *conditional directionality* principle studied in [19] is related.

Other work in which ranking functions are applied in the context of argumentation includes a ranking-based semantics for defeasible logic programming [17] and an approach to structured argumentation using conditionals which induce rankings [25]. In [21], stratified labelings are introduced as a new semantics for abstract argumentation that directly correlate to ranking functions. However, this approach is based on rankings over arguments rather than labelings.

Various approaches to combining Bayesian networks and argumentation have been considered. Most of these (e.g., [22,23,24]) are based on extracting structured arguments and attacks from Bayesian network, which is quite different from what we do. An exception is [15], which deals with translating Bayesian networks into a general kind of numerical AFs, but leaves handling of cycles to future work.

8. Conclusion

We studied ranking functions over labelings as a generalization of regular labeling-based semantics. We demonstrated that, in this setting, AFs can be interpreted similarly to ranking networks. In particular, the SCC stratification property, which is related to SCC decomposability, generalizes the notion of stratification in the presence of cycles. Finally, rankings over labelings induce rankings over arguments in a natural way and we made some initial steps in comparing these induced rankings to established approaches in this direction [12].

As for future work we plan a broader study of different types of ranking-theoretic semantics, as well as a more detailed look at the relationship between rankings over labelings and over arguments. Another possibility is to apply the ideas presented here in a probabilistic setting, in which we consider probability distributions over labelings and interpret AFs like Bayesian networks.

Acknowledgements The research reported here was partially supported by the Deutsche Forschungsgemeinschaft (grant KE 1686/3-1).

References

[1] Leila Amgoud and Jonathan Ben-Naim. Ranking-based semantics for argumentation frameworks. In Weiru Liu, V S Subrahmanian, and Jef Wijsen, editors, *Proceedings of the 7th International Conference on Scalable Uncertainty Management (SUM 2013)*, volume 8078 of *Lecture Notes In Artificial Intelligence*, Washington, DC, USA, September 2013.

[2] Leila Amgoud, Jonathan Ben-Naim, Dragan Doder, and Srdjan Vesic. Ranking arguments with compensation-based semantics. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference, KR 2016, Cape Town, South Africa, April 25-29, 2016.*, pages 12–21, 2016.

[3] Pietro Baroni, Guido Boella, Federico Cerutti, Massimiliano Giacomin, Leendert van der Torre, and Serena Villata. On the input/output behavior of argumentation frameworks. *Artificial Intelligence*, 217:144–197, 2014.

[4] Pietro Baroni and Massimiliano Giacomin. On principle-based evaluation of extension-based argumentation semantics. *Artificial Intelligence*, 171(10-15):675–700, 2007.

[5] Pietro Baroni, Massimiliano Giacomin, and Giovanni Guida. Scc-recursiveness: a general schema for argumentation semantics. *Artificial Intelligence*, 168(1-2):162–210, 2005.

[6] Ringo Baumann and Gerhard Brewka. Expanding argumentation frameworks: Enforcing and monotonicity results. In *COMMA 2010, Desenzano del Garda, Italy, September 8-10, 2010*, volume 216 of *Frontiers in Artificial Intelligence and Applications*, pages 75–86. IOS Press, 2010.

[7] Elise Bonzon, Jerome Delobelle, Sebastien Konieczny, and Nicolas Maudet. Argumentation ranking semantics based on propagation. In *Proceedings of the 6th International Conference on Computational Models of Argument (COMMA-2016)*, 2016.

[8] Richard Booth, Souhila Kaci, Tjitze Rienstra, and Leendert van der Torre. A logical theory about dynamics in abstract argumentation. In *Scalable Uncertainty Management - 7th International Conference, SUM 2013, Washington, DC, USA, September 16-18, 2013. Proceedings*, volume 8078 of *LNCS*, pages 148–161. Springer, 2013.

[9] Martin Caminada. On the issue of reinstatement in argumentation. In *Logics in Artificial Intelligence, 10th European Conference, JELIA 2006, Liverpool, UK, September 13-15, 2006, Proceedings*, volume 4160 of *LNCS*, pages 111–123. Springer, 2006.

[10] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. Graduality in argumentation. *Journal of Artificial Intelligence Research*, 23:245–297, 2005.

[11] Sylvie Coste-Marquis, Sébastien Konieczny, Jean-Guy Mailly, and Pierre Marquis. On the revision of argumentation systems: Minimal change of arguments statuses. In *Principles of Knowledge Representation and Reasoning: The 14th International Conference, KR 2014, Vienna, Austria, July 20-24, 2014*. AAAI Press, 2014.

[12] Jerome Delobelle. *Ranking-based Semantics for Abstract Argumentation*. PhD thesis, Centre de Recherche en Informatique de Lens, 2017.

[13] Martin Diller, Adrian Haret, Thomas Linsbichler, Stefan Rümmele, and Stefan Woltran. An extension-based approach to belief revision in abstract argumentation. *Int. J. Approx. Reasoning*, 93:395–423, 2018.

[14] Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–358, 1995.

[15] D. M. Gabbay and O. Rodrigues. Introducing bayesian argumentation networks. *The IfColog Journal of Logics and their Applications*, 3(2):241–278, 2016.

[16] Moises Goldszmidt and Judea Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artif. Intell.*, 84(1-2):57–112, 1996.

[17] Gabriele Kern-Isberner and Guillermo R. Simari. A default logical semantics for defeasible argumentation. In *Proceedings of the Twenty-Fourth International Florida Artificial Intelligence Research Society Conference (FLAIRS'11)*, 2011.

[18] Daphne Koller and Nir Friedman. *Probabilistic Graphical Models - Principles and Techniques*. MIT Press, 2009.

[19] Tjitze Rienstra. *Argumentation in Flux - Modelling Change in the Theory of Argumentation*. PhD thesis, University of Luxembourg/University of Montpellier II, 2014.

[20] Wolfgang Spohn. *The Laws of Belief - Ranking Theory and Its Philosophical Applications*. Oxford University Press, 2014.

[21] Matthias Thimm and Gabriele Kern-Isberner. Stratified labelings for abstract argumentation (preliminary report). Technical report, ArXiv, August 2013.

[22] Sjoerd T. Timmer, John-Jules Ch. Meyer, Henry Prakken, Silja Renooij, and Bart Verheij. Explaining bayesian networks using argumentation. In *ECSQARU*, volume 9161 of *Lecture Notes in Computer Science*, pages 83–92. Springer, 2015.

[23] Sjoerd T. Timmer, John-Jules Ch. Meyer, Henry Prakken, Silja Renooij, and Bart Verheij. A structure-guided approach to capturing bayesian reasoning about legal evidence in argumentation. In *ICAIL*, pages 109–118. ACM, 2015.

[24] Gerard Vreeswijk. Argumentation in bayesian belief networks. In *ArgMAS*, volume 3366 of *Lecture Notes in Computer Science*, pages 111–129. Springer, 2004.

[25] Emil Weydert. On arguments and conditionals. In *Proceedings of the ECAI-2012 Workshop on Weighted Logics for Artificial Intelligence (WL4AI)*, pages 69–76, 2012.