Fuzzy Argumentation Frameworks

Jeroen Janssen[*]
Dept. of Computer Science
Vrije Universiteit Brussel
jeroen.janssen@vub.ac.be

Martine De Cock
Dept. of Appl. Math.
and Computer Science
Universiteit Gent
martine.decock@ugent.be

Dirk Vermeir
Dept. of Computer Science
Vrije Universiteit Brussel
dvermeir@vub.ac.be

January 2, 2008

Abstract

We propose fuzzy argumentation frameworks as a conservative extension of traditional argumentation frameworks[Dun95]. The fuzzy approach enriches the expressive power of the classical argumentation model by allowing to represent the relative strength of the attack relationships between arguments, as well as the degree to which arguments are accepted. Furthermore, we explore the relationship with fuzzy answer set programming, more in particular the correspondence between the stable extensions of a fuzzy argumentation framework and the fuzzy models of an associated program.

Keywords: Fuzzy argumentation frameworks, fuzzy answer set programming, query expansion.

1 Introduction

The challenge of understanding argumentation and its role in human reasoning has been addressed by many researchers in different fields, including philosophy, logic, and AI. A formal theory of argumentation has been proposed

[*]Funded by the Fund for Scientific Research - Flanders (FWO-Vlaanderen)

in a seminal article by Dung[Dun95]. This theory is based on the idea that a statement (argument) is believable if it can be *defended* successfully against contesting arguments. Thus, the theory considers so-called *argumentation frameworks* that model the interactions between different arguments, while abstracting from the meaning or internal structure of an argument. Within a given framework of interacting arguments, there might be one, or more, sets of conclusions, called *extensions*, that are deemed collectively acceptable.

In human argumentation however, in many cases, not all attacks have equal strength. For example, in a murder trial, the argument "her fingerprints are on the murder weapon" is a stronger attack on the premise "the accused is innocent" than the argument that "she was the last person to see the victim alive". Furthermore, statements like "the accused is innocent" and "the accused did not have a good relationship with the victim" may hold to, and hence may be deemed acceptable to, a certain extent only, as opposed to being either accepted or not. In the present paper, we therefore introduce *fuzzy argumentation frameworks* as a conservative extension of [Dun95]'s notion. Advantages of the fuzzy approach include the possibility to represent the relative strength of the (attack) relationships between arguments and a more sophisticated approach to extensions: being fuzzy sets of arguments, an extension may commit only to a certain degree to the acceptance of a certain argument, thus weakening any (unanswered) attacks on it.

Already in [Dun95], a strong link between non-monotonic reasoning formalisms and argumentation frameworks was presented. E.g. the stable model semantics for logic programs naturally maps to argumentation frameworks where an argument constitutes a derivation in the program. In the present paper, we take a different approach: we show that a program can be obtained from a fuzzy argumentation framework such that the fuzzy models of the program correspond to the stable extensions of the framework.

The rest of the paper is organised as follows: to introduce the notation used in the paper, in Section 2 we recall some preliminaries re. fuzzy logic. Next we introduce fuzzy argumentation frameworks by providing the basic definitions and properties (Section 3) as well as a sample application (Section 4). The relationship with fuzzy answer set programs is established in Section 5 and Section 6 concludes.

Due to space restrictions, all proofs have been omitted. They can be obtained from `http://tinf2.vub.ac.be/~dvermeir/papers/faf-full.pdf`.

2 Preliminaries

Throughout this paper, membership values are taken from a complete lattice called a *truth lattice*, i.e. a partially ordered set $(\mathcal{L}, \leq_{\mathcal{L}})$ where every subset of $\mathcal{L}$ has a greatest lower bound (inf) and a least upper bound (sup)[Bir67]. Where the order is clear from the context, we refer to the lattice as $\mathcal{L}$ instead of $(\mathcal{L}, \leq_{\mathcal{L}})$. In addition we use $0_{\mathcal{L}}$ and $1_{\mathcal{L}}$, or 0 and 1 if $\mathcal{L}$ is understood, to denote the smallest and greatest element of $\mathcal{L}$, respectively. The traditional logical operations of negation, conjunction, disjunction, and implication are generalized to logical operators acting on truth values of $\mathcal{L}$ in the usual way (see e.g. [NPM99]):

- A *negator* on $\mathcal{L}$ is any anti-monotone $\mathcal{L} \to \mathcal{L}$ mapping $\sim$ satisfying $\sim 0_{\mathcal{L}} = 1_{\mathcal{L}}$ and $\sim 1_{\mathcal{L}} = 0_{\mathcal{L}}$. A negator $\sim$ is called *involutive* iff $\forall x \in \mathcal{L} \cdot \sim\sim x = x$.

- A *triangular norm*, t-norm for short, on $\mathcal{L}$ is any commutative and associative $\mathcal{L}^2 \to \mathcal{L}$ (infix) operator $\curlywedge$ satisfying $\forall x \in \mathcal{L} \cdot 1_{\mathcal{L}} \curlywedge x = x$. Moreover we require $\curlywedge$ to be increasing in both of its arguments, i.e.[1] for $x_1, x_2 \in \mathcal{L}$, $x_1 \leq_{\mathcal{L}} x_2$ implies $x_1 \curlywedge y \leq_{\mathcal{L}} x_2 \curlywedge y$. Intuitively, a t-norm corresponds to conjunction.

- A *triangular conorm*, t-conorm for short, on $\mathcal{L}$ is any commutative and associative $\mathcal{L}^2 \to \mathcal{L}$ (infix) operator $\curlyvee$ satisfying $\forall x \in \mathcal{L} \cdot 0_{\mathcal{L}} \curlyvee x = x$. Moreover we require $\curlyvee$ to be increasing in both of its arguments. A t-conorm corresponds to disjunction.

- An *implicator* $\rightsquigarrow$ on $\mathcal{L}$ is any $\mathcal{L}^2 \to \mathcal{L}$ (infix) operator $\rightsquigarrow$ satisfying $0_{\mathcal{L}} \rightsquigarrow 0_{\mathcal{L}} = 1_{\mathcal{L}}$, and $\forall x \in \mathcal{L} \cdot 1_{\mathcal{L}} \rightsquigarrow x = x$. Moreover $\rightsquigarrow$ must be decreasing in its first, and increasing in its second argument, i.e., for $x_1, x_2, y \in \mathcal{L}$, $x_1 \leq_{\mathcal{L}} x_2$ implies $x_1 \rightsquigarrow y \geq_{\mathcal{L}} x_2 \rightsquigarrow y$ as well as $y \rightsquigarrow x_1 \leq_{\mathcal{L}} y \rightsquigarrow x_2$.

The *residual implicator* of a t-norm $\curlywedge$ is defined by $x \rightsquigarrow y = \sup\{\lambda \in \mathcal{L} \mid x \curlywedge \lambda \leq_{\mathcal{L}} y\}$, while a t-conorm $\curlyvee$ and a negator $\sim$ induce an S-implicator defined by $x \rightsquigarrow y = \sim x \curlyvee y$.

In this paper, we will mostly assume that truth lattices are finite. E.g. finite subsets of $[0, 1]$ such as $\{0.0, 0.1, \ldots, 0.9, 1.0\}$ are used frequently. Well-known fuzzy logical operators include the *minimum t-norm* $x \curlywedge y = \min(x, y)$, its residual implicator $x \rightsquigarrow y = 1$ if $x \leq y$, and $x \rightsquigarrow y = y$ otherwise,

[1] Note that the monotonicity of the second component immediately follows from that of the first component due to the commutativity.

as well as the corresponding S-implicator $x \rightsquigarrow y = \max(1-x, y)$. The *Łukasiewicz t-norm* $x \curlywedge y = \max(x+y-1, 0)$ induces the residual implicator $x \rightsquigarrow y = \min(1-x+y, 1)$, which is also an S-implicator. For negation, often the *standard negator* $\sim x = 1 - x$ is used.

A *fuzzy set* A over some (ordinary) set X and a truth lattice $\mathcal{L}$ is an $X \to \mathcal{L}$ mapping. We use $\mathcal{F}(X)$, where $\mathcal{L}$ is understood, to denote the set of all fuzzy sets over X. We sometimes use the notation x^l to denote that $A(x) = l$. The *support* of a fuzzy set A is defined by $supp(A) = \{x \mid A(x) > 0_{\mathcal{L}}\}$. Fuzzy set intersection is defined by $(A \cap B)(x) = A(x) \curlywedge B(x)$. Fuzzy set inclusion is also defined as usual by $A \subseteq_{\mathcal{L}} B$, or $A \subseteq B$ if $\mathcal{L}$ is understood, iff $\forall x \in X \cdot A(x) \leq_{\mathcal{L}} B(x)$. A fuzzy relation over X is a fuzzy set over $X \times X$.

3 Fuzzy Argumentation Frameworks

An argumentation framework[Dun95] consists of a set of *arguments*, some of which *attack* each other. Intuitively, an *extension* of such a framework represents a position, i.e. a set of arguments a rational agent may subscribe to, that can be defended against attacks.

Definition 1 *An* argumentation framework *(*AF *for short) is a tuple* $\langle \mathcal{A}, \not\rightarrow \rangle$ *where* $\mathcal{A}$ *is a set of arguments and* $\not\rightarrow \subseteq \mathcal{A} \times \mathcal{A}$ *is a binary* attack *relation between arguments. When* $a_1 \not\rightarrow a_2$*, we say that the argument* a_1 attacks *the argument* a_2*. The notation is extended to sets in the obvious way: i.e. for* b *an argument, and* A *and* B *sets of arguments,* $A \not\rightarrow b$ *iff* $\exists a \in A \cdot a \not\rightarrow b$*, while* $b \not\rightarrow A$ *iff* $\exists a \in A \cdot b \not\rightarrow a$*, and* $A \not\rightarrow B$ *iff* $\exists b \in B \cdot A \not\rightarrow b$*.*

Definition 2 *Let* $\langle \mathcal{A}, \not\rightarrow \rangle$ *be an* AF*. A set* $A \subseteq \mathcal{A}$ *is* conflict-free *if no argument in* A *attacks an argument in* A*, i.e.* $\neg(A \not\rightarrow A)$*. A conflict-free set* A *is an* admissible extension *if it defends itself against all attacks, i.e.* $\forall b \not\rightarrow A \cdot A \not\rightarrow b$*. A* preferred extension *is a maximal (w.r.t.* $\subseteq$*) admissible extension. A* stable extension *is a preferred extension* A *that attacks all external arguments, i.e.* $\forall b \notin A \cdot A \not\rightarrow b$*.*

Example 1 *[GK96] The* AF $\langle \{porsche, volvo, safe, sporty\}, R \rangle$*, where* R *contains* $porsche \not\rightarrow volvo$*,* $volvo \not\rightarrow porsche$*,* $safe \not\rightarrow sporty$*,* $safe \not\rightarrow porsche$*,* $sporty \not\rightarrow safe$*, and* $sporty \not\rightarrow volvo$*, represents a discussion between a wife and her husband about buying a car. There are two stable extensions:* $\{sporty, porsche\}$ *and* $\{safe, volvo\}$*. Note that, on its own, e.g.* $\{safe\}$ *is already admissible (but not preferred).*

Example 2 *The* AF $\langle\{red, pink, blue\}, R\rangle$*, where* R *contains* $red \not\rightharpoonup pink$*,* $pink \not\rightharpoonup red$*,* $red \not\rightharpoonup blue$*,* $blue \not\rightharpoonup red$*,* $pink \not\rightharpoonup blue$*, and* $blue \not\rightharpoonup pink$*, represents a discussion about the colour of a certain sweater in the context of a school uniform policy. There are three stable extensions:* $\{red\}$*,* $\{pink\}$*, and* $\{blue\}$*.*

All three arguments in Example 2 are conflicting. However, the argument that the sweater is red (e.g. as raised by a teacher) and the argument that the sweater is pink (e.g. as claimed by the school principal) are not conflicting to a very high degree, because for a light shade of red one person might call it pink, while another might still prefer to call it red. Suppose that the school policy demands for a blue uniform. Both the teacher's and the school principal's argument strongly attack the argument that the sweater conforms to the school policy (e.g. as claimed by the student). To harvest the best arguments to attack the student's argument, it is therefore desirable to be able to include both the teacher's and the school principal's argument in the extension to a high degree. Indeed the latter two attack each other only to a small degree, hence, intuitively, committing to them both to a relatively high degree does not significantly violate the conflict-freeness of the extension.

To allow this kind of expressivity, we extend the classical argumentation model in two ways: (1) by allowing the attack relation to be a fuzzy relation over the set of arguments, we can represent the degree to which arguments attack each other, and (2) by allowing an extension to be a fuzzy set over the set of arguments, we can model that some arguments are accepted to a higher degree than others. In the following definition, we formalize these intuitions by extending $\not\rightharpoonup$ to cover fuzzy sets of arguments. A fuzzy argumentation framework consists of a set of arguments, some of which *attack each other to a certain degree.*

Definition 3 *A* fuzzy argumentation framework *(*FAF *for short) is a tuple* $\langle\mathcal{A}, \not\rightharpoonup\rangle$ *where* $\mathcal{A}$ *is a set of arguments and* $\not\rightharpoonup$ *is a fuzzy relation over* $\mathcal{A}$*. For* b *an argument, and* A *and* B *fuzzy sets of arguments, we define the degree to which* A *attacks* b *as* $A \not\rightharpoonup b = \sup_{a\in\mathcal{A}}(A(a) \curlywedge (a \not\rightharpoonup b))$*, and the degree to which* b *attacks* A *as* $b \not\rightharpoonup A = \sup_{a\in\mathcal{A}}(A(a) \curlywedge (b \not\rightharpoonup a))$*. Furthermore, the degree to which* A *attacks* B *is given by* $A \not\rightharpoonup B = \sup_{b\in\mathcal{A}}(B(b) \curlywedge (A \not\rightharpoonup b))$*.*

According to Definition 3, the strength of an attack $A \not\rightharpoonup b$ does not only depend on the strength of an attack $a \not\rightharpoonup b$, where a is an argument supported by A, but also on the degree $A(a)$ to which A supports a: the stronger the presence of a in A, the stronger the attacks from A through a. On the other hand, if an argument b is only present to a marginal degree in B, it should be clear that attacking b does not greatly contribute to the "global" attack on B.

Example 3 *Let $\langle \mathcal{A}, \not\rightarrow \rangle$ be a* FAF *where $\mathcal{A}$ contains all arguments that appear in two position papers A and B. Since a position paper does not support each of its arguments in equal measure, the papers are best represented by fuzzy sets $A, B \in \mathcal{F}(\mathcal{A})$. Now suppose that there are two arguments $a, b \in \mathcal{A}$ such that $a \not\rightarrow^{0.9} b$, i.e. a strongly attacks b. On the other hand, a is supported by A but not by B while b is not supported by A and only weakly by B, i.e. b represents only a minor aspect of B. E.g. $A(a) = 0.9$, $B(a) = 0$, while $A(b) = 0$ and $B(b) = 0.1$.*

Then, although in supporting a, A attacks b strongly, i.e. (using the minimum as t-norm in Definition 3) $A \not\rightarrow^{0.9} b$, this attack does not noticeably affect B's position as a whole since it contributes only 0.1 *to $A \not\rightarrow B$.*

Intuitively, an extension of a fuzzy argumentation framework is a fuzzy set A such that for each a in $\mathcal{A}$, $A(a)$ represents the degree to which a rational agent accepts argument a. Definition 4 provides more flexibility than Definition 2 in three aspects: (1) while in classical argumentation frameworks extensions are required to be entirely conflict-free, in the fuzzy approach they are allowed to contain *minor internal attacks*, (2) an admissible extension of a fuzzy argumentation framework only needs to *defend itself well enough against all attacks*, and (3), likewise, a stable extension only needs to *sufficiently attack all external arguments*, in other words, any argument that is to a high degree outside of the extension should be strongly attacked by it.

Definition 4 *Let $\langle \mathcal{A}, \not\rightarrow \rangle$ be a* FAF*. A fuzzy set A over $\mathcal{A}$ is x-*conflict-free*, $x \in \mathcal{L}$, iff $(\sim (A \not\rightarrow A)) \geq x$. A fuzzy set A is a y-*admissible extension* if it defends itself well enough against all attacks, i.e. $\inf_{b \in \mathcal{A}}((b \not\rightarrow A) \rightsquigarrow (A \not\rightarrow b)) \geq y$. A y-*preferred extension*, $y \in \mathcal{L}$, is a maximal (w.r.t. $\subseteq$ over fuzzy sets) y-admissible extension. A z-*stable extension*, $z \in \mathcal{L}$, is a fuzzy set A that sufficiently attacks all external arguments, i.e. $\inf_{b \in \mathcal{A}}(\sim A(b) \rightsquigarrow (A \not\rightarrow b)) \geq z$.*

Example 4 *Assume that we replace the* AF *in Example 2 by a* FAF *where the fuzzy attack relation is given by red $\not\rightarrow^{0.1}$ pink, pink $\not\rightarrow^{0.1}$ red, red $\not\rightarrow^{1}$ blue, blue $\not\rightarrow^{1}$ red, pink $\not\rightarrow^{1}$ blue, and blue $\not\rightarrow^{1}$ pink. We consider the fuzzy set $A = \{pink^{0.8}, red^{0.7}\}$, in other words we accept that the sweater has a colour between pink and red.*

To evaluate the conflict-freeness of A, we determine the degree to which there are internal attacks in A. Using the minimum t-norm we obtain that (cfr. Definition 3) $A \not\rightarrow pink = 0.1$ and $A \not\rightarrow red = 0.1$, hence $A \not\rightarrow A = 0.1$, *which indicates a minor level of conflict among the accepted arguments. Using the standard negator we obtain* $\sim (A \not\rightarrow A) = 0.9$*, hence A is* 0.9*-conflict-free. Note that when using the Łukasiewicz t-norm, which is inherently more tolerant to minor inconsistencies, we would obtain that* $A \not\rightarrow pink = 0$ *and* $A \not\rightarrow red = 0$*, hence that A is 1-conflict-free. Because of limited space, in the rest of this example we only consider the Łukasiewicz t-norm and its residual implicator.*

Next we verify how well A defends itself against all attacks. Since $pink \not\rightarrow A = 0$ *and* $red \not\rightarrow A = 0$*, the only real attack on A comes from the argument that the sweater is blue, namely* $blue \not\rightarrow A = 0.8$*. However, from* $A \not\rightarrow blue = 0.8$ *and the previously mentioned attack values, it is clear that A strikes back with equal force on all attacking arguments. Since for any residual implicator* $a \not\rightarrow A \rightsquigarrow A \not\rightarrow a$ *equals 1 as soon as* $a \not\rightarrow A \leq A \not\rightarrow a$*, we obtain that A is 1-admissible. It is however not a* 1*-preferred set, since* $\{pink^{0.8}, red^{0.8}\}$ *is a superset of A which is also* 1*-admissible.*

As for the stability, $\sim A(pink) \rightsquigarrow A \not\rightarrow pink = 0.2 \rightsquigarrow 0 = 0.8$, $\sim A(red) \rightsquigarrow A \not\rightarrow red = 0.3 \rightsquigarrow 0 = 0.7$ *and* $\sim A(blue) \rightsquigarrow A \not\rightarrow blue = 1 \rightsquigarrow 0.8 = 0.8$*. By which we can conclude that A is a* 0.7*-stable extension. Note that A does not have a perfect score for stability because, among other things, A neither fully supports nor opposes the argument that the sweater is red.*

It is straightforward to verify that FAFs represent a conservative extension of the classical notion.

Theorem 1 *Let* $F = \langle \mathcal{A}, \not\rightarrow \rangle$ *be a* AF *and let* $F_f = \langle \mathcal{A}, \not\rightarrow_f \rangle$ *be the fuzzy version over* $\{0,1\}$ *of F with* $a \not\rightarrow_f^1 b$ *iff* $a \not\rightarrow b$*. Then the set of stable extensions of F coincides with the set of supports of the* 1*-conflict-free* 1*-stable extensions of* F_f*.*

The following theorems describe some of the effects of accepting more or less arguments (or accepting them to a higher or a lower degree). The extended attack relation is monotone w.r.t. fuzzy set inclusion, i.e. accepting more arguments does not decrease the power to attack.

Theorem 2 *Let* $\langle \mathcal{A}, \not\rightarrow \rangle$ *be a* FAF *and* $B \subseteq A \in \mathcal{F}(\mathcal{A})$*. Then* $\forall a \in \mathcal{A} \cdot B \not\rightarrow a \leq A \not\rightarrow a$*.*

Proof

Let $A, B \in \mathcal{F}(\mathcal{A})$ such that $B \subseteq A$ and let $a \in \mathcal{A}$, then if $B \not\rightarrow a$ this means, by definition, that $\sup_{b \in \mathcal{A}}(B(b) \curlywedge (b \not\rightarrow a))$. Now, due to the fact that $\forall a \in \mathcal{A} \cdot B(a) \leq A(a)$ and the monotonicity of $\curlyvee$ and $\curlywedge$ it follows that $\sup_{b \in \mathcal{A}}(B(b) \curlywedge (b \not\rightarrow a)) \leq \sup_{b \in \mathcal{A}}(A(b) \curlywedge (b \not\rightarrow a))$, thus $B \not\rightarrow a \leq A \not\rightarrow a$ $\square$

Furthermore, z-stability is monotone with respect to fuzzy set inclusion. This might come as a surprise since it is not true in the boolean case, but it is due to the fact that z-stability is independent of the degree of conflict-freeness.

Theorem 3 *Let $\langle \mathcal{A}, \not\rightharpoonup \rangle$ be a* FAF *and $B \subseteq A \in \mathcal{F}(\mathcal{A})$. Then if B is a z-stable extension, so is A.*

Proof
Let $B \in \mathcal{F}(\mathcal{A})$ be z-stable and $B \subseteq A \in \mathcal{F}(\mathcal{A})$. Due to the definition of z-stable, this means that $\inf_{a \in \mathcal{A}}(\sim B(a) \rightsquigarrow B \not\rightharpoonup a) \geq z$. ¿From this it follows that $\inf_{a \in \mathcal{A}}(\sim B(a) \rightsquigarrow A \not\rightharpoonup a) \geq z$ because of Theorem 2 and the monotonicity of $\rightsquigarrow$ in its second argument. Due to the fact that $\sim$ is monotonically decreasing, that $\forall a \in \mathcal{A} \cdot B(a) \leq A(a)$ and that $\rightsquigarrow$ is anti-monotone in its first argument, we obtain $\inf_{a \in \mathcal{A}}(\sim A(a) \rightsquigarrow A \not\rightharpoonup a) \geq z$ and thus A is also z-stable. □

For x-conflict-freeness, we obtain anti-monotonicity: adding more arguments, results in more conflicts, so conflict-freeness decreases.

Theorem 4 *Let $\langle \mathcal{A}, \not\rightharpoonup \rangle$ be a* FAF *and let $B \subseteq A \in \mathcal{F}(\mathcal{A})$. Then if A is x-conflict-free, so is B.*

Proof
Let $\langle \mathcal{A}, \not\rightharpoonup \rangle$ be a FAF and let $B \subseteq A \in \mathcal{F}(\mathcal{A})$ with A x-conflict-free. Then,
$(B \not\rightharpoonup B) \leq (A \not\rightharpoonup A)$ follows from $B \subseteq A$,
$B \not\rightharpoonup B = \sup_{a \in \mathcal{A}}(B(a) \curlywedge \sup_{b \in \mathcal{A}}(B(b) \curlywedge b \not\rightharpoonup a))$,
$A \not\rightharpoonup A = \sup_{a \in \mathcal{A}}(A(a) \curlywedge \sup_{b \in \mathcal{A}}(A(b) \curlywedge b \not\rightharpoonup a))$, and the monotonicity of $\curlywedge$ and sup. □

4 Query Expansion

Query expansion is the process of expanding keyword based queries with more terms, related to the intended meaning of the query, to refine the search results. One option is to use an available thesaurus such as WordNet, expanding the query by adding synonyms[Voo94]. Related terms can also be automatically discovered from the searchable documents though, taking into account statistical information such as co-occurrences of words in documents: the more frequently two terms co-occur, the more they are assumed to be related[XC96].

In either case, the thesaurus can be thought of as being a (fuzzy) relation R over the set X of terms. Note that R should be reflexive (i.e. every term is obviously related to itself) and symmetric (i.e. if we say that term a is related to term b to a degree of p, obviously we want that term b is related to term a in the same degree). Transitivity is not necessarily required, for reasons explained in [CCK07]. From X and R, a FAF F can be generated such that the conflict-free and stable extensions of F correspond to optimal queries.

More in particular, we generate the FAF $F = \langle X, \not\rightarrow \rangle$ where $a \not\rightarrow b = \sim (a\, R\, b)$. In other words, the arguments of F correspond to terms, and a term a attacks a term b to the extent to which these terms are not related. Note that, due to the symmetry of R, every argument in X attacks each one of its attackers to the same degree.

A set $A \subseteq X$ is then highly conflict-free if every term in A is related to all other terms in A to a high degree, in other words A is a coherent query. Furthermore, A is highly stable if it strongly attacks all external arguments, i.e. a term can only be left out of the query to the extent that it is not related to at least one of the terms in the query. Together these requirements fit our intuition about what an optimally expanded query should look like, i.e. a query that has a high stability (refined with as many terms as possible) and is highly conflict-free (without losing the coherence). The following example illustrates the interplay between the conflict-freeness and the stability requirements.

Table 1: Fuzzy Thesaurus from [CCK07]

	mac	recipe	computer	apple	fruit	pie
mac	1	0	0.89	0.89	0	0.01
recipe		1	0.56	0.83	0.66	1
computer			1	0.94	0.44	0.44
apple				1	0.83	0.99
fruit					1	0.44
pie						1

Example 5 *Suppose we have a set of terms* $X = \{apple, fruit, mac, pie, recipe, computer\}$ *and a fuzzy relation* R *over* X*, defined as in Table 1. The corresponding attack relation, using the standard negator, is depicted in Table 2. The user query that we aim to expand is* $Q_1 = \{apple, pie\}$*. Independently of the choice of t-norm in Definition 3, it holds that* $Q_1 \not\rightarrow mac = 0.99$*,* $Q_1 \not\rightarrow recipe = 0.17$*,* $Q_1 \not\rightarrow computer = 0.66$*,* $Q_1 \not\rightarrow apple = 0.01$*,* $Q_1 \not\rightarrow fruit = 0.66$ *and* $Q_1 \not\rightarrow pie = 0.01$*.*

Table 2: Attack Relation

	mac	recipe	computer	apple	fruit	pie
mac	0	1	0.11	0.11	1	0.99
recipe		0	0.44	0.17	0.44	0
computer			0	0.06	0.66	0.66
apple				0	0.17	0.01
fruit					0	0.66
pie						0

The query Q_1 is highly conflict-free, namely to degree 0.99 when using the standard negator, since the strength of the attack apple $\not\rightharpoonup$ pie is almost neglectable. On the other hand, Q_1 is only a 0.17*-stable extension (independently of the choice of negator and of the choice of implicator in Definition 4). This unstability is due to the term recipe being excluded from the query without there being a good reason to. Indeed, recipe is not strongly attacked by any of the keywords from the query so, in terms of stability, there is no reason not to include it.*

However, it is not possible to add more terms to the query while keeping the original high level of conflict-freeness. We therefore lower our standards and look for expanded queries that are 0.8*-conflict-free. The only expanded query that satisfies this condition is $Q_2 = \{apple, pie, recipe\}$, which is a* 0.66*-stable extension.*

Example 6 *Adding more terms to query Q_2 in Example 5 results in a significant drop in conflict-freeness. Indeed, the remaining terms mac, computer, and fruit are under an attack with a strength of at least 0.66 by at least one of the terms already in Q_2, so including yet another term in Q_2 would leave the conflict-freeness at most at 0.44. However, in the setting described in this section, it is also possible to consider weighted queries. For instance, one can verify that, using the minimum t-norm, for $Q_3 = \{apple^1, pie^1, recipe^1, fruit^{0.3}\}$ we obtain that $Q_3 \not\rightharpoonup mac = 1$, $Q_3 \not\rightharpoonup recipe = 0.3$, $Q_3 \not\rightharpoonup computer = 0.66$, $Q_3 \not\rightharpoonup apple = 0.17$, $Q_3 \not\rightharpoonup fruit = 0.66$ and $Q_3 \not\rightharpoonup pie = 0.3$. Hence, Q_3 is* 0.7*-conflict-free, and, using the residual or the S-implicator associated with the minimum t-norm, Q_3 is still a* 0.66*-stable extension.*

5 FAF and Fuzzy Answer Set Programming

The answer set programming (ASP) paradigm[Bar03] has gained a lot of popularity in the last years, due to its truly declarative non-monotonic semantics. ASP and fuzzy logic can be be combined into the single framework of fuzzy answer set programming to increase the flexibility and hence the application potential of ASP. A *fuzzy answer set program*[NCV07b], FASP for short, is a finite set of rules[2] of the form $a \leftarrow \alpha$ with a an *atom* and α a set of *literals*, each of which is either an atom or of the form *not* a, a an atom, representing the "negation as failure" of a.

A *fuzzy interpretation* of a FASP P is a mapping $I : \mathcal{B}_P \rightarrow \mathcal{L}$ assigning a truth value from the lattice $\mathcal{L}$ to each of the atoms appearing in P. I is extended to literals by defining $I(not\ a) = \sim I(a)$ and to rules r: $a \leftarrow \alpha$ using $I(\alpha) = \curlywedge_{l\in\alpha} I(l)$ and $I(r) = I(\alpha) \rightsquigarrow I(a)$.

A fuzzy y-*model*, $y \in \mathcal{L}$, of P is a fuzzy interpretation I that satisfies $\mathcal{A}_p(P, I) \geq y$ where $\mathcal{A}_p$ is a function that takes as input a program and an interpretation, yielding a value denoting the degree in which I is a model of P. Naturally, $\mathcal{A}_p$ should be increasing whenever the degrees of satisfaction $I(r)$ of the rules in P are increasing.

We show that, under certain conditions, the y-stable extensions of a FAF $F = \langle \mathcal{A}, \not\rightarrow \rangle$ correspond to the fuzzy y-models of a FASP Π_F that can be constructed from F as follows: intuitively, for each argument $a \in \mathcal{A}$, Π_F will contain exactly one rule r_a introducing a. The body of r_a contains one literal *not* b_a for each $b \in \mathcal{A}$ such that $(b \not\rightarrow a) > 0_\mathcal{L}$. Each literal b_a is itself defined through a single rule[3] $b_a \leftarrow b, (b \not\rightarrow a)$.

Definition 5 *For a* FAF $F = \langle \mathcal{A}, \not\rightarrow \rangle$*, the associated* FASP Π_F *is defined by* $\Pi_F = R_a \cup R_b$ *where*

$$\begin{aligned} R_a &= \{a \leftarrow \{not\ b_a \mid (b \not\rightarrow a > 0_\mathcal{L})\} \mid a \in \mathcal{A}\} \\ R_b &= \{b_a \leftarrow b, (b \not\rightarrow a) \mid (b \not\rightarrow a) > 0_\mathcal{L}\} \end{aligned}$$

Where we take min as the t-norm for aggregating the body in R_a-rules and the t-norm of the argumentation framework for the R_b-rules.
The aggregator $\mathcal{A}_{\Pi_F}$ is such that all rules from R_b must evaluate to $1_\mathcal{L}$, while for other rules, the minimal degree of satisfaction is taken, i.e.

$$\mathcal{A}_{\Pi_F}(\Pi_F, I) = \begin{cases} \inf_{r_a \in R_a}(I(r_a)) & \textit{iff } (\alpha) \\ 0_\mathcal{L} & \textit{otherwise} \end{cases}$$

where $(\alpha) \equiv \forall r \in R_b \cdot I(r) = 1_\mathcal{L}$.

[2] In the present paper we do not consider negative literals of the form $\neg a$, a an atom. In terms of [NCV07b], this makes every interpretation 1-consistent. Also, we do not consider constraints (rules with empty head).

[3] Although program rules in [NCV07b] syntactically may not contain constants from $\mathcal{L}$, the semantics in [NCV07b] easily supports such an extension using the "fuzzy input literals" mechanism (Section 4 in [NCV07b]).

Note that any argument a that is not attacked will have a corresponding fact rule $a \leftarrow$ in R_a.

We show that y-stable extensions of F correspond to certain fuzzy y-models of Π_F.

Theorem 5 *Let $F = \langle \mathcal{A}, \not\rightarrow \rangle$ be a* FAF. *If $\rightsquigarrow$ is a contrapositive implicator, $x = y \Rightarrow x \rightsquigarrow y = 1$, $\sim$ is an involutive negator and $\sim \sup_{x \in X}(F(x)) = \inf_{x \in X}(\sim F(x))$, then for any $y \in \mathcal{L}$, X is a y-stable extension of F iff $X' = X \cup \{b_a^q \mid q = (X(b) \curlywedge (b \not\rightarrow a))\}$ is a y-model of Π_F.*

Proof

Suppose X is a y-stable extension of F. By the definition of y-stable extensions, this is equivalent to

$$\inf_{a \in \mathcal{A}} (\sim X(a) \rightsquigarrow (X \not\rightarrow a)) \geq y$$

Using contraposition this is equivalent to

$$\inf_{a \in \mathcal{A}} (\sim (X \not\rightarrow a) \rightsquigarrow X(a)) \geq y$$

Then, by the definition of $X \not\rightarrow a$

$$\inf_{a \in \mathcal{A}} (\sim \sup_{b \in \mathcal{A}} (X(b) \curlywedge (b \not\rightarrow a)) \rightsquigarrow X(a)) \geq y$$

Which is equivalent to

$$\inf_{a \in \mathcal{A}} (\sim \sup_{b \in \mathcal{A} \wedge (b \not\rightarrow a) > 0_{\mathcal{L}}} (X(b) \curlywedge (b \not\rightarrow a)) \rightsquigarrow X(a)) \geq y$$

Due to the generalised De Morgan properties we get

$$\inf_{a \in \mathcal{A}} (\inf_{b \in \mathcal{A} \wedge (b \not\rightarrow a) > 0_{\mathcal{L}}} (\sim (X(b) \curlywedge (b \not\rightarrow a))) \rightsquigarrow X(a)) \geq y$$

By definition of X' this is equivalent to

$$\inf_{a \in \mathcal{A}} (\inf_{b \in \mathcal{A} \wedge (b \not\rightarrow a) > 0_{\mathcal{L}}} (\sim X'(b_a)) \rightsquigarrow X'(a)) \geq y$$

Which, due to the construction of Π_F and the interpretation of rules is equivalent to

$$\inf_{r \in R_a} (X'(r)) \geq y$$

Which is equivalent to the definition of a y-model of Π_F due to the fact that inf is the aggregator for rules in R_a if all rules in R_b are 1. The latter is established by the fact that $X'(b_a) = X'(b) \curlywedge (b \not\rightarrow a)$ and by the restriction on the chosen implicator. $\square$

Note that the above equivalence concerns FASP *models*, not answer sets where answer sets are models that are free from "assumptions"[NCV07b] and, in a sense, minimal models. However, unlike with answer set programming, minimality is not a desirable criterion for argumentation frameworks as, intuitively, one attempts to maximize the set of arguments that can be defended against attacks. The only limit on the size of an extension is the desire for conflict-freeness (see also Theorems 3 and 4) which can be imposed separately.

6 Concluding Remarks

We've motivated and introduced fuzzy argumentation frameworks as a conservative extension of the classical notion from [Dun95] which allows for more fine-grained knowledge representation in terms of doubts (some degree of conflict may be tolerated) and strength of attacks, both absolute and with respect to the degree of acceptance of the attacker and the attacked. The strong connection with Fuzzy Answer Set Programming[NCV07b] also provides a practical implementation using DLVHEX[NCV07a]. In future work, we intend to extend our approach to (applications of) bipolar argumentation frameworks[CLS05, KP01, Ver02, ACLS04]. We will also investigate dialogical semantics, see e.g. [VJ99], for fuzzy argumentation frameworks.

References

[ACLS04] Leila Amgoud, Claudette Cayrol, and Marie-Christine Lagasquie-Schiex. On the bipolarity in argumentation frameworks. In *Proceedings of the 10th International Workshop on Non-Monotonic Reasoning, NMR'2004*, pages 1–9, 2004.

[Bar03] C. Baral. *Knowledge Representation, Reasoning and Declarative Problem Solving.* Cambridge University Press, 2003.

[Bir67] Garrett Birkhoff. Lattice theory. *American Mathematical Society Colloquium Publications*, 25(3), 1967.

[CCK07] M. De Cock, C. Cornelis, and E. Kerre. Fuzzy rough sets: the forgotten step. *IEEE Transactions on Fuzzy Systems*, 15(1):121–130, 2007.

[CLS05] C. Cayrol and M. C. Lagasquie-Schiex. On the acceptability of arguments in bipolar argumentation frameworks. In *Proceedings of ECSQARU 2005*, volume 3571 of *Lecture Notes in Artificial Intelligence*, pages 378–389. Springer Verlag, 2005.

[Dun95] 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.

[GK96] T. Gordon and N. Karacapilidis. The zeno argumentation framework. In *Proceedings of the biannual International Conference on Formal and Applied Practical Reasoning (FAPR) workshop*, 1996.

[KP01] N. Karacapilidis and D. Papadias. Computer supported argumentation and collaborative decision making: the hermes system. *Information systems*, 26(4):259–277, 2001.

[NCV07a] D. Van Nieuwenborgh, M. De Cock, and D. Vermeir. Computing fuzzy answer sets using DLVHEX. In *International Conference on Logic Programming*, 2007.

[NCV07b] D. Van Nieuwenborgh, M. De Cock, and D. Vermeir. An introduction to fuzzy answer set programming. *Annals of Mathematics and Artificial Intelligence*, 50(3-4):363–388, 2007.

[NPM99] V. Novák, I. Perfilieva, and J. Močkoř. *Mathematical Principles of Fuzzy Logic*. Kluwer Academic Publishers, 1999.

[Ver02] B. Verheij. On the existence and multiplicity of extension in dialectical argumentation. In S. Benferhat and E. Giunchiglia, editors, *Proceedings of the 9th International Workshop on Non-Monotonic Reasoning (NMR2002)*, pages 416–425, 2002.

[VJ99] D. Vermeir and H. Jakobovits. Dialectic semantics for argumentation frameworks. In *Proceedings of the Seventh International Conference on Artificial Intelligence and Law*, pages 53–62. Association for Computing Machinery, june 1999.

[Voo94] E. M. Voorhees. Query expansion using lexical-semantic relations. In *Proceedings of ACM SIGIR 1994 (17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval)*, pages 61–69, 1994.

[XC96] J. Xu and W. B. Croft. Query expansion using local and global document analysis. In *Proceedings of ACM SIGIR 1996 (19th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval)*, pages 4–11, 1996.