

If Nothing Is Accepted — Repairing Argumentation Frameworks

Ringo Baumann, Markus Ulbricht

Department of Computer Science, Leipzig University, Germany
{baumann,mulbricht}@informatik.uni-leipzig.de

Abstract

Conflicting information in an agent’s knowledge base may lead to a semantical defect, that is, a situation where it is impossible to draw any plausible conclusion. Finding out the reasons for the observed inconsistency (so-called *diagnosis*) and/or restoring consistency in a certain minimal way (so-called *repairs*) are frequently occurring issues in knowledge representation and reasoning. In this paper we provide a series of first results for these problems in the context of abstract argumentation theory regarding the two most important reasoning modes, namely credulous as well as sceptical acceptance. Our analysis includes the following problems regarding minimal repairs/diagnosis: existence, verification, computation of one and enumeration of all solutions. The latter problem is tackled with a version of the so-called *hitting set duality* first introduced by Raymond Reiter in 1987. It turns out that grounded semantics plays an outstanding role not only in terms of complexity, but also as a useful tool to reduce the search space for diagnosis regarding other semantics.

1 Introduction

A well-known problem in knowledge representation and reasoning is the semantical collapse of an agent’s knowledge base \mathcal{K} , i.e. \mathcal{K} is inconsistent and thus does not allow any plausible conclusion. Hansson coined the term *consolidation* and defined it as an operation that withdraws parts of \mathcal{K} in such a way that, first, the resulting knowledge base \mathcal{K}' is consistent and secondly, the change is as small as possible (Hansson 1994). Even earlier, Reiter introduced the presumably first formal treatment of this problem in his seminal paper (Reiter 1987). The so-called *diagnostic problem* for a given system arises whenever we observe that the system does not behave as it should. Reiter used first-order logic as representation formalism and his definition of a diagnosis contains the concepts of consistency as well as minimality. Since then the problem of restoring consistency under the requirement of minimal change has been considered for many other formalisms like situation calculus (McIlraith 1999), logic programs (Sakama and Inoue 2003), description logic including non-monotonic versions (Lembo et al. 2011; Bienvenu 2012; Eiter, Fink, and Stepanova 2013) as well as

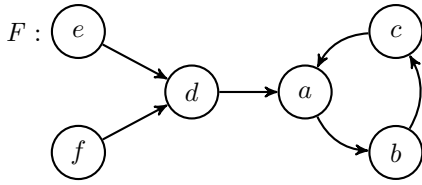
probabilistic conditional logic (Potyka and Thimm 2014) to mention a few.

In this paper we focus on the non-monotonic theory of abstract argumentation (Dung 1995). More precisely, we consider an abstract argumentation framework (AF) as an agent’s knowledge base and the associated extensions correspond to her beliefs (cf. (Coste-Marquis et al. 2014; Nouioua and Würbel 2014; Diller et al. 2018) for similar approaches). In brief, Dung-style AFs consist of arguments and attacks which are treated as primitives, i.e., the internal structure of arguments is not considered. The major focus is on resolving conflicts. To this end a variety of semantics have been defined, each of them specifying acceptable sets of arguments, so-called *extensions*, in a particular way.

The starting point of our study is a semantical defect of an agent’s AF which prevents her from drawing any plausible conclusion in the sense that nothing is accepted. Our aim is to obtain an agent which is able to act. Therefore we want to know what are minimal diagnoses of the given knowledge base, i.e., which parts are causing the semantical defect. The knowledge about these diagnoses may make it easier to decide what to do next. For instance, a certain minimal diagnosis may consist of arguments which are somehow out of date in comparison to the others. Consequently, one may tend to discard these arguments. This is why our repair approach focusses on removal of certain arguments. In general, it is easily conceivable that one may create a certain internal hierarchy over the stored arguments encoding the willingness to drop them. For instance, let us assume that dropping either the argument A_1 or the argument A_2 represent minimal repairs. Let us further assume that A_1 represents “Do not drive over the lawn, because this would destroy the lawn.” and A_2 stands for “Drive over the lawn, whenever you may save lives with this action.” In this situation it might be reasonable to drop A_1 .

Let us consider an abstract example under the most prominent semantics, namely the stable one. Informally, a set of arguments is a stable extensions if there are no conflicts between them and moreover, all other arguments are attacked by at least one argument of the set.

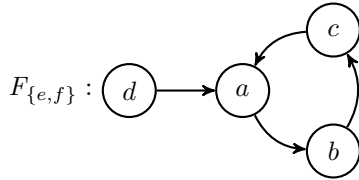
Example 1.1 (Semantical Collapse). The AF F does not possess any stable extension, i.e., $stb(F) = \emptyset$.



The main aim is to study such semantical defects with regard to the following naturally arising questions/tasks:

1. *Diagnosis* – Which sets of arguments are causing the collapse?
2. *Computation* – How to compute one/all diagnoses?
3. *Properties* – Do diagnoses always exist? Are there certain preferred diagnoses? How computationally costly is it to verify a candidate diagnosis?
4. *Repair* – How to use this information to obtain an agent which is indeed able to act?

Example 1.2 (Diagnosis and Repair). One may argue that the arguments e and f together can be seen as a diagnosis for the semantical defect of F since ignoring these arguments and their corresponding attacks result in a meaningful AF denoted by $F_{\{e,f\}}$. In particular, $stb(F_{\{e,f\}}) = \{\{d, b\}\}$.



Note that neither of both arguments can be omitted since the resulting frameworks would collapse too. In this sense, the presented diagnosis $\{e, f\}$ and the corresponding repair $F_{\{e,f\}}$ are minimal.

The main results of the paper are as follows: After discussing necessary preliminaries for abstract argumentation in Section 2 we tackle the second question in Section 3, namely: How to systematically find all minimal diagnoses of a given AF under a given semantics. In particular, we will see that the considered AF F indeed possesses further minimal repairs. These formal results are mainly due to the well-known *hitting set duality* first introduced in (Reiter 1987) and a recently generalized version of it (Brewka, Thimm, and Ulbricht 2017). In the subsequent Section 4 we consider the question of existence of minimal diagnoses and we provide relations between diagnoses w.r.t. different semantics. Grounded semantics plays a central role here since its minimal diagnoses represent bounds for diagnosis of other semantics. We also briefly discuss subclasses of AFs, namely *symmetric*, *compact* and *acyclic* frameworks. Afterwards, we study the computational complexity for the associated existence and verification problem. In the subsequent Section 6 we briefly discuss repair strategies for stable and preferred semantics. Finally, we conclude as well as discuss some related work in Section 7. In almost all cases we provide full proofs. If not available we included some short comments indicating how to prove the statement in question.

2 Background in Abstract Argumentation

In the original formulation (Dung 1995), an *abstract argumentation framework* is a directed graph $F = (A, R)$ where nodes in A represent arguments and the relation R models *attack*, i.e., for $a, b \in A$, if $(a, b) \in R$ we say that a *attacks* b or a is an *attacker* of b . We say that F is *self-controversial* if any argument attacks itself. In this paper we restrict ourselves to non-empty finite AFs (cf. (Baumann and Spanring 2015; 2017) for a treatment of unrestricted AFs). Moreover, for a set E we use E^+ for $\{b \mid (a, b) \in R, a \in E\}$ and define $E^\oplus = E \cup E^+$. A further essential notion in argumentation is *defense*. Formally, an argument b is *defended* by a set A if each attacker of b is counter-attacked by some $a \in A$. For a given set $S \subseteq A$ we use F_S as a shorthand for $F|_{A \setminus S} = (A \setminus S, R|_{A \setminus S})$.

An *extension-based semantics* σ is a function which assigns to any AF $F = (A, R)$ a set of sets of arguments $\sigma(F) \subseteq 2^A$. Each one of them, so-called *extension*, is considered to be acceptable with respect to F . Besides conflict-free and admissible sets (abbr. *cf* and *ad*) we consider stage, stable, semi-stable, complete, preferred, grounded, ideal and eager semantics (abbr. *stg*, *stb*, *ss*, *co*, *pr*, *gr*, *il* and *eg* respectively). A recent overview can be found in (Baroni, Caminada, and Giacomin 2011).

Definition 2.1. Let $F = (A, R)$ be an AF and $E \subseteq A$.

1. $E \in cf(F)$ iff for no $a, b \in E$, $(a, b) \in R$,
2. $E \in ad(F)$ iff $E \in cf(F)$ and E defends all its elements,
3. $E \in stb(F)$ iff $E \in cf(F)$ and $E^\oplus = A$,
4. $E \in ss(F)$ iff $E \in ad(F)$ and for no $\mathcal{I} \in ad(F)$, $E^\oplus \subset \mathcal{I}^\oplus$,
5. $E \in co(F)$ iff $E \in ad(F)$ and for any $a \in A$ defended by E , $a \in E$,
6. $E \in pr(F)$ iff $E \in co(F)$ and for no $\mathcal{I} \in co(F)$, $E \subset \mathcal{I}$,
7. $E \in gr(F)$ iff $E \in co(F)$ and for any $\mathcal{I} \in co(F)$, $E \subseteq \mathcal{I}$,
8. $E \in il(F)$ iff $E \in co(F)$, $E \subseteq \bigcap pr(F)$ and there is no $\mathcal{I} \in co(F)$ satisfying $\mathcal{I} \subseteq \bigcap pr(F)$ s.t. $E \subset \mathcal{I}$,
9. $E \in eg(F)$ iff $E \in co(F)$, $E \subseteq \bigcap ss(F)$ and there is no $\mathcal{I} \in co(F)$ satisfying $\mathcal{I} \subseteq \bigcap ss(F)$ s.t. $E \subset \mathcal{I}$.

We say that a semantics σ is *universally defined* if $\sigma(F) \neq \emptyset$ for any $F \in \mathcal{F}$. If even $|\sigma(F)| = 1$ we say that σ is *uniquely defined*. All semantics apart from stable one are universally defined. In addition, grounded, ideal and eager semantics are examples of uniquely defined semantics. Stable semantics may *collapse*, i.e. there are AFs F , s.t. $stb(F) = \emptyset$ (cf. running example F depicted in Example 1.1). For two semantics σ and τ we write $\sigma \subseteq \tau$ if for any AF F , $\sigma(F) \subseteq \tau(F)$. For instance, it is well-known that $stb \subseteq ss \subseteq pr \subseteq co \subseteq ad \subseteq cf$.

In the present paper we are interested in situations where a given AF $F = (A, R)$ does not possess accepted arguments. To make the notion of acceptance precise, we utilize the following two alternative reasoning modes, namely *credulous* as well as *sceptical acceptance*. Note that we require $\sigma(F)$ to be non-empty for sceptical reasoning in order to avoid the (for our purpose) unintended situation that every argument is sceptically accepted due to technical reasons.

Set-theoretically the intersection over the empty family of sets would yield any argument (cf. (Baumann and Spanring 2015, Section 2) for more information). However, in our setting it makes sense to define it as the empty set since the (sceptical or credulous) acceptance of an argument should imply the existence of at least one set containing a .

Definition 2.2. Given a semantics σ , an AF $F = (A, R)$ and an argument $a \in A$. We say that a is

1. *credulously accepted w.r.t. σ* if $a \in \bigcup \sigma(F)$,
2. *sceptically accepted w.r.t. σ* if $a \in \bigcap \sigma(F)$ and $\sigma(F) \neq \emptyset$.

3 A Hitting Set Duality for AFs

In his seminal paper (Reiter 1987), Reiter establishes – for his setting of a given system description – a duality result between the set of all minimal repairs and the minimal conflicts. Recently, it was shown (Brewka, Thimm, and Ulbricht 2017) that Reiter’s duality can be generalized to arbitrary logics given that the knowledge bases in question can be modeled as finite set of formulas. In order to capture even non-monotonic logics a refinement of the notion of inconsistency was necessary. In this section we will demonstrate how to utilize this generalized duality for the non-monotonic theory of abstract argumentation.

For the result in (Brewka, Thimm, and Ulbricht 2017) a generic definition of a logic is used. However, the only property we require in order to apply the hitting set duality from (Brewka, Thimm, and Ulbricht 2017) is the ability to express a given AF as a finite set of formulas (a “knowledge base”). We do this as follows.

Definition 3.1. Given an AF $F = (A, R)$ with $A = \{a_1, \dots, a_n\}$ we call $\mathcal{K} = \{a_1, \dots, a_n\}$ the atom-based representation of F . In this case, a subset $H \subseteq \mathcal{K}$ is the representation of the framework $F|_H$.

Of course, given a set A of arguments, the corresponding framework $F = (A, R)$ is ambiguous. This is why we may only make use of the atom-based representation of an AF F whenever the framework in question is given.

The hitting set duality we aim at makes explicit use of consistent and inconsistent subsets of a given knowledge base (resp. AF in our case). Hence, the following definition is central.

Definition 3.2. Given an AF $F = (A, R)$ with atom-based representation \mathcal{K} . Let σ be a semantics. We call \mathcal{K} *inconsistent* w.r.t. credulous resp. sceptical reasoning whenever there is no argument a that is credulously resp. sceptically accepted.

Now we turn to the notion of *strong inconsistency* and how it induces a hitting set duality for our setting. The following definition naturally applies to both credulous and sceptical reasoning.

Definition 3.3. Let $F = (A, R)$ be an AF with atom-based representation \mathcal{K} . Let σ be any semantics. For $H \subseteq \mathcal{K}$, H is called *strongly \mathcal{K} -inconsistent* if $H \subseteq H' \subseteq \mathcal{K}$ implies H' is inconsistent (w.r.t. credulous resp. sceptical reasoning). H is *minimal strongly \mathcal{K} -inconsistent* if H is strongly \mathcal{K} -inconsistent and $H' \subsetneq H$ implies that H' is not strongly \mathcal{K} -inconsistent.

Let $SI_{min}(\mathcal{K})$ denote the set of all minimal strongly \mathcal{K} -inconsistent subsets of \mathcal{K} . Note that the definition of $SI_{min}(\mathcal{K})$ depends on the reasoning mode. Given an AF F , an atom-based representation \mathcal{K} might have different inconsistent subsets w.r.t. credulous reasoning than w.r.t. sceptical reasoning. The reason is that “inconsistency” in the former case means that no argument is credulously accepted while “inconsistency” the latter case means that no argument is sceptically accepted.

Example 3.4 (Strongly Inconsistent Subsets). Consider the running example F with atom-based representation $\{a, b, c, d, e, f\}$. Let us focus on credulous reasoning. We already observed that F has no stable extension, i.e., \mathcal{K} itself is strongly \mathcal{K} -inconsistent w.r.t. stable semantics. The subset $H_1 \subseteq \mathcal{K}$ with $H_1 = \{a, b, c\}$ corresponds to the odd circle contained in F . However, H_1 is not strongly \mathcal{K} -inconsistent since the superset H_2 with $H_1 \subseteq H_2 \subseteq \mathcal{K}$ given as $H_2 = \{a, b, c, d\}$ has the stable extension $\{b, d\}$ (cf. AF $F_{\{e, f\}}$ depicted in Example 1.2). One may easily verify that $SI_{min}(\mathcal{K}) = \{\{a, b, c, e\}, \{a, b, c, f\}\}$.

We proceed with the well-known concepts of (minimal) hitting sets and (maximal) consistent subsets.

Definition 3.5. Let \mathcal{M} be a set of sets. We call S a *hitting set* of \mathcal{M} if $S \cap M \neq \emptyset$ for each $M \in \mathcal{M}$. A hitting set S of \mathcal{M} is a *minimal hitting set* of \mathcal{M} if $S' \subsetneq S$ implies S' is not a hitting set of \mathcal{M} .

Definition 3.6. Let $F = (A, R)$ be an AF with atom-based representation \mathcal{K} . Let σ be any semantics. We say $H \subseteq \mathcal{K}$ is a *maximal consistent* subset of \mathcal{K} if H is consistent and $H \subsetneq H' \subseteq \mathcal{K}$ implies H' is inconsistent. We denote the set of all maximal consistent subsets of \mathcal{K} by $C_{max}(\mathcal{K})$.

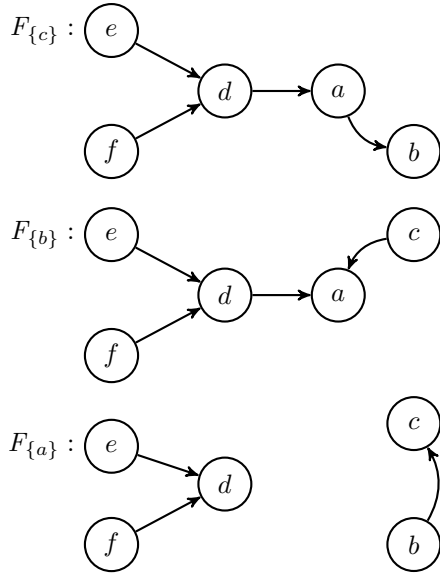
In Example 1.2 we informally discussed the notions of diagnoses and repairs. The following central definition captures these intuitions in a formally precise way. Please note that the need for a diagnosis is given whenever we are faced with a semantical situation which does not resolve any conflict with regard to the considered reasoning mode (*semantical defect*). This means, in case of credulous reasoning, either no extensions or the empty extension is considered to be acceptable only and in case of sceptical reasoning, either there are no extensions or no argument is supported by all extensions.

Definition 3.7. Given a semantics σ an AF F and a representation \mathcal{K} of it. We call $S \subseteq \mathcal{K}$ a σ -*cred-diagnosis* of F iff $\mathcal{K} \setminus S$ is consistent w.r.t. credulous reasoning. Moreover, we call the AF F_S as well as $\mathcal{K} \setminus S$ a σ -*cred-repair* of F . We use the terms *minimal* and *least* for \subseteq -minimal or \subseteq -least σ -diagnosis as well as the associated σ -repairs. We define (minimal, least) σ -*scep-diagnosis* and σ -*scep-repair* analogously.

We omit the terms “*cred*” resp. “*scep*” whenever the type of reasoning is implicitly clear or unimportant. We will thus in most cases simply speak of σ -*diagnoses* and σ -*repairs*. We are now ready to phrase the duality theorem from (Brewka, Thimm, and Ulbricht 2017) within our setting. Note that it applies to both credulous and sceptical reasoning.

Theorem 3.8. (Brewka, Thimm, and Ulbricht 2017) Let $F = (A, R)$ be an AF with atom-based representation \mathcal{K} . Let σ be any semantics. Then, S is a minimal hitting set of $SI_{min}(\mathcal{K})$ if and only if S is a minimal σ -diagnosis of F .

Example 3.9 (Maximal Consistent Subsets via Hitting Set Duality). Consider again the running example F with stable semantics and credulous reasoning. We already checked that $SI_{min}(\mathcal{K}) = \{\{a, b, c, e\}, \{a, b, c, f\}\}$. According to Definition 3.5 we obtain four minimal hitting sets, namely $\{a\}$, $\{b\}$, $\{c\}$ and $\{e, f\}$. Observe that $\{e, f\}$ is the already found *stb*-diagnosis presented in Example 1.2. The minimal hitting sets for F can be interpreted as follows: Either one argument from the odd circle needs to be removed or both e and f to facilitate d . In fact, the maximal consistent subsets of F are $C_{max}(\mathcal{K}) = \{\{a, b, c, d\}, \{a, b, d, e, f\}, \{a, c, d, e, f\}, \{b, c, d, e, f\}\}$. These sets correspond to the *stb*-repairs $F_{\{e, f\}}$ (considered in Example 1.2) as well as $F_{\{c\}}$, $F_{\{b\}}$ and $F_{\{a\}}$ depicted below. We obtain, $stb(F_{\{c\}}) = \{\{e, f, a\}\}$, $stb(F_{\{b\}}) = \{\{e, f, c\}\}$ and $stb(F_{\{a\}}) = \{\{e, f, b\}\}$.



4 On the Existence of Repairs

In the last section we dealt with a characterization of all minimal diagnoses of a given framework. Clearly, before computing potential repairs one may wonder whether there are minimal diagnoses at all. In this section we provide the formal results w.r.t. this problem and in particular, we give an affirmative answer for all considered semantics given that the framework in question is not self-controversial. Unfortunately, the existence of a least repair is not guaranteed which leads to follow-up question of *how to repair?* which will be considered in the subsequent section.

Relations between Credulous and Sceptical Reasoning Mode

We start with some general relations between credulous and sceptical diagnoses. The following theorem applies to any

semantics. It states that minimal credulous diagnoses can be found as subsets of sceptical diagnoses.

Theorem 4.1. Given an AF F and a semantics σ . If S is a *scep*- σ -diagnosis of F , then there is a minimal *cred*- σ -diagnosis S' of F , s.t. $S' \subseteq S$.

Proof. Let S be a *scep*- σ -diagnosis of F . This means, $\bigcap \sigma(F_S) \neq \emptyset$. Consequently, $\sigma(F_S) \neq \emptyset$ and therefore $\bigcup \sigma(F_S) \neq \emptyset$. Thus, S is a *cred*- σ -diagnosis of F . Moreover, by finiteness of S we deduce the existence of a minimal *cred*- σ -diagnosis S' of F with $S' \subseteq S$ concluding the proof. \square

Vice Versa, sceptical diagnosis can be found as supersets of credulous ones. We want to mention two issues. First, in contrast to the theorem before, the proof of Theorem 4.2 requires semantics specific properties and thus, does not hold for any argumentation semantics. Secondly, it is an open question whether even minimality can be claimed for the considered universally defined semantics. We conjecture that minimality can be shown but we did not find a proof so far (cf. Conjecture 4.3).

Theorem 4.2. Given an AF F and a semantics $\sigma \in \{stg, stb, ss, co, pr, gr, ul, eg\}$. If S is a *cred*- σ -diagnosis of F , then there is a *scep*- σ -diagnosis S' of F , s.t. $S \subseteq S'$.

Proof. Let S be a *cred*- σ -diagnosis of $F = (A, R)$. This means, $\bigcup \sigma(F_S) \neq \emptyset$. Consequently, there is an argument $a \in \bigcup \sigma(F_S) \neq \emptyset$. Define $S' = A \setminus \{a\}$. Obviously, $S \subseteq S'$ and $F_{S'} = (\{a\}, \emptyset)$ since σ -extensions are conflict-free. Moreover, $\sigma(F_{S'}) = \{\{a\}\}$ for any considered semantics σ . This implies $\bigcap \sigma(F_{S'}) \neq \emptyset$ justifying the assertion that S' is a *scep*- σ -diagnosis of F . \square

Conjecture 4.3. Given an AF F and a semantics $\sigma \in \{stg, stb, ss, co, pr\}$. If S is a minimal *cred*- σ -diagnosis of F , then there is a minimal *scep*- σ -diagnosis S' of F , s.t. $S \subseteq S'$.

Finally, we show two helpful, but not unexpected relations between different semantics and their reasoning modes.

Theorem 4.4. Given two semantics σ and τ , s.t. $\sigma \subseteq \tau$ and σ is universally defined. For any AF F we have:

1. If S is a *cred*- σ -diagnosis of F , then there is a minimal *cred*- τ -diagnosis S' of F , s.t. $S' \subseteq S$.
2. If S is a *scep*- τ -diagnosis of F , then there is a minimal *scep*- σ -diagnosis S' of F , s.t. $S' \subseteq S$.

Proof. We prove the second item only. Let S be a *scep*- τ -diagnosis of F . This means, $\bigcap \tau(F_S) \neq \emptyset$. Since $\sigma \subseteq \tau$ is assumed we deduce $\emptyset \neq \bigcap \tau(F_S) \subseteq \bigcap \sigma(F_S)$. Since σ is universally defined we have $\sigma(F_S) \neq \emptyset$ which implies that S is a *scep*- σ -diagnosis of F . Moreover, by finiteness of S we deduce the existence of a minimal *scep*- σ -diagnosis S' of F with $S' \subseteq S$ concluding the proof. \square

Uniquely Defined Semantics

Please note that in case of uniquely defined semantics we have that any (minimal) sceptical diagnosis is a (minimal) credulous one and vice versa.¹ Thus, Conjecture 4.3 is fulfilled by any uniquely defined semantics.

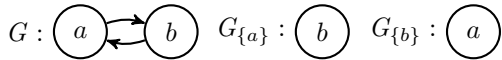
We proceed with grounded semantics since these results will play a central role for all other semantics considered in this paper. Dung originally defined the grounded extension of an AF $F = (A, R)$ as the \subseteq -least fixpoint of the so-called *characteristic function* $\Gamma_F : 2^A \rightarrow 2^A$ with $E \mapsto \{a \in A \mid a \text{ is defended by } E\}$. Moreover, he showed that this definition coincides with the \subseteq -least complete extension (Dung 1995, Theorem 25) as introduced in Definition 2.1. Since Γ_F is shown to be \subseteq -monotonic we may compute the unique grounded extension G stepwise, i.e., applying Γ_F iteratively starting from the empty set. More precisely, $G = \bigcup_{i=1}^{|A|} \Gamma_F^i(\emptyset)$ (cf. (Baumann and Spanring 2017, Section 3.2)). For instance, the unique grounded extensions of $F_{\{c\}}$ and $F_{\{a\}}$ are $\{e, f, a\} = \Gamma_{F_{\{c\}}}^2(\emptyset)$ and $\{e, f, b\} = \Gamma_{F_{\{a\}}}^1(\emptyset)$, respectively. Consequently, an AF possesses a non-empty grounded extension if and only if there exists at least one unattacked argument.

Fact 4.5. *For any AF F which is not self-controversial there exists a gr -repair.*

The subsequent main theorem claims the existence of minimal σ -diagnoses for the considered uniquely defined semantics (recall that we do not need to distinguish between credulous and sceptical reasoning for those semantics). The facts that ideal semantics accepts more arguments than grounded semantics and eager semantics is even more credulous than ideal semantics combined with the above fact yield the existence of diagnoses. Moreover, the restriction to finite AFs even gives us the existence of minimal ones. As an aside, the main theorem does not carry over to infinite AFs.²

Theorem 4.6. *For any semantics $\sigma \in \{gr, eg, il\}$ and any not self-controversial AF F there exists a minimal σ -repair.*

Example 4.7. The following simple framework G demonstrates that least σ -repairs does not necessarily exist. For $\sigma \in \{gr, eg, il\}$ we have $\sigma(G) = \{\emptyset\}$, i.e. nothing is credulously/sceptically accepted.



The \subseteq -incomparable sets $S_1 = \{a\}$ and $S_2 = \{b\}$ are minimal σ -diagnosis which yield the σ -repairs $G_{\{a\}}$ and $G_{\{b\}}$. Clearly, both possess non-empty σ -extensions, namely $\{b\}$ or $\{a\}$, respectively.

Finally, diagnoses for ideal and eager semantics can be found as subsets of a grounded diagnosis.

Lemma 4.8. *Let $\sigma \in \{il, eg\}$. If S' is a gr -diagnosis of an AF F , then there is a minimal σ -diagnosis S of F , s.t. $S \subseteq S'$.*

¹Formally, if $|\sigma(F)| = 1$, then $\bigcap \sigma(F) = \bigcup \sigma(F)$.

²It can be checked that the AF $L = (\mathbb{N}, \{(i, j) \mid i < j\})$ possesses gr -repairs but no minimal ones.

Universally Defined Semantics

Let us consider now semantics which provide us with at least one acceptable position. The following lemma shows that for these semantics minimal credulous as well as sceptical diagnoses are guaranteed, whenever there is a grounded diagnosis.

Lemma 4.9. *Let $\sigma \in \{ss, pr, co\}$. For any AF F there exists a minimal σ -diagnosis S , whenever there exists a gr -diagnosis S' of F . Moreover, even $S \subseteq S'$ can be guaranteed.*

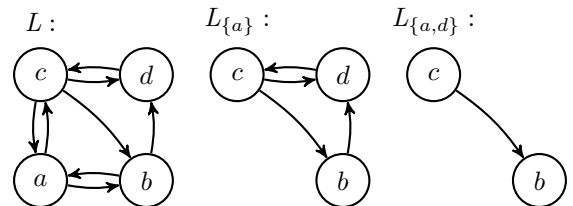
Proof. Let $\sigma \in \{ss, pr, co\}$ and S' a gr -diagnosis of F . Hence, $gr(F_{S'}) = \{G\}$ with $G \neq \emptyset$. Since G is the \subseteq -least fixpoint of $\Gamma_{F_{S'}}$, we deduce $G \subseteq C$ for any $C \in co(F_{S'})$. Due to $ss \subseteq pr \subseteq co$ and the universal definedness of σ we have $\emptyset \neq G \subseteq \bigcap \sigma(F_{S'})$ as well as $\emptyset \neq G \subseteq \bigcup \sigma(F_{S'})$. Hence, S' is a sceptical as well as credulous σ -diagnosis of F . Due to finiteness of F , there exists a minimal σ -diagnosis $S \subseteq S'$ concluding the proof. \square

Combining Theorem 4.6 and Lemma 4.9 yields the subsequent main theorem for the considered universally defined semantics.

Theorem 4.10. *For any semantics $\sigma \in \{ss, pr, co\}$ any not self-controversial AF F possesses a minimal σ -repair.*

The following example shows, as promised in Lemma 4.9, that already computed grounded diagnoses can be used to find minimal preferred diagnoses. Moreover, in contrast to uniquely defined semantics we observe that minimal sceptical and minimal credulous diagnoses do not necessarily coincide.

Example 4.11. Consider the following AF L . Since we have no unattacked arguments we deduce $gr(L) = \{\emptyset\}$, i.e., nothing is accepted. Moreover, the same applies to $L_{\{a\}}$ and $L_{\{d\}}$. Consequently, $L_{\{a,d\}}$ is a minimal gr -repair since $gr(L_{\{a,d\}}) = \{\{c\}\}$. Note that $\{a, d\}$ is even a sceptical as well credulous preferred diagnosis of L . These diagnoses are not minimal for preferred semantics since $pr(L) = \{\{a, d\}, \{c\}\}$ implies $\bigcup pr(L) \neq \emptyset$ as well as $pr(L_{\{a\}}) = \{\{c\}\}$ entails $\bigcap pr(L_a) \neq \emptyset$. Altogether, we have strict subset relation ($\emptyset \subsetneq \{a\} \subsetneq \{a, d\}$) between minimal credulous preferred, minimal sceptical preferred and minimal grounded diagnoses.



Collapsing Semantics

Stable semantics is the only prominent semantics which may collapse even for finite AFs (cf. running example F). However, in terms of existence of repairs we do not observe any differences to all other considered semantics since the

conflict-freeness of at least one non-empty set ensures the existence of repairs.

Fact 4.12. *Any not self-controversial AF F possesses a minimal stb -repair.*

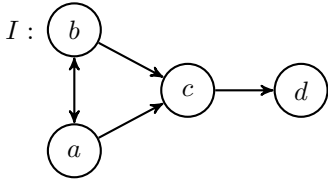
In contrast to all other considered semantics we have that stable diagnoses can not be necessarily found as subsets of an already computed grounded one (Lemmata 4.8, 4.9). For instance, the running example F possesses the unique grounded extension $\{e, f\}$. Consequently, we have the trivial (least) gr -diagnosis, namely the empty set. As we have already seen in Example 3.9 all minimal stb -diagnoses are non-empty. Nevertheless, credulous as well as sceptical diagnoses for stable semantics can be found as supersets of grounded ones.

Lemma 4.13. *If S' is a gr -diagnosis of an AF F , then there is a stb -diagnosis S of F , s.t. $S' \subseteq S$.*

Proof. Given S' as gr -diagnosis of $F = (A, R)$, i.e. $gr(F_{S'}) = \{E\}$ with $E \neq \emptyset$. Consider now E^\oplus w.r.t. the attack-relation of $F_{S'}$. Obviously, $S' \subseteq A \setminus E^\oplus$ and moreover, $gr(F_{A \setminus E^\oplus}) = \{E\}$. Obviously, by construction $E \in stb(F_{A \setminus E^\oplus})$. Furthermore, since E is non-empty we deduce that there is at least one unattacked argument $a \in E$. Hence, for any $E' \in stb(F_{A \setminus E^\oplus})$ we have $a \in E'$. Consequently, $A \setminus E^\oplus$ serves as a credulous as well as sceptical diagnosis for stable semantics. \square

Please note that Lemma 4.13 does not claim minimality of the stb -diagnosis S . Indeed, the following example illustrates that minimality is not obtained in general.

Example 4.14. Consider the following AF I



Since every argument is attacked we infer that empty set cannot be a gr -diagnosis. A possible one is the set $\{a\}$. Indeed, this is also a stb -diagnosis (w.r.t. credulous as well as sceptical reasoning), but not minimal since I itself possesses the sceptically accepted argument d .

Subclasses of AFs

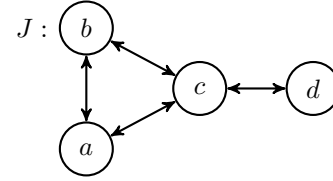
We may obtain more advanced results by restricting the AF under consideration. Let us start with so-called *symmetric* AFs. According to (Coste-Marquis, Devred, and Marquis 2005) an AF $F = (A, R)$ is symmetric if R is symmetric, nonempty and irreflexive. In case of stable, semi-stable and preferred semantics we obtain a very useful property, namely any argument $a \in A$ belongs to at least one extension (Coste-Marquis, Devred, and Marquis 2005, Proposition 6). Consequently, we may show the following properties.

Proposition 4.15. *Given a semantics $\sigma \in \{stb, ss, pr\}$, a symmetric AF $F = (A, R)$ and a set $S \subseteq A$.*

1. \emptyset is the least $cred$ - σ -diagnosis and
2. S is a (minimal) $scep$ - σ -diagnosis iff S is a (minimal) gr -diagnosis.

Let us assume that our current knowledge base underlies further external revision processes (cf. (Coste-Marquis et al. 2014; Baumann and Brewka 2015; Diller et al. 2018) for belief revision in abstract argumentation). Both items can be gainfully used if we know that certain types of revision do not affect the symmetry of an AF. More precisely, the items 1 and 2 ensure that we have either nothing to do (if interested in credulous reasoning) or we may act according to grounded semantics instead of σ (if sceptical reasoning is chosen). Regarding our setup we note that removing arguments does not affect the symmetry status of an argument.

Example 4.16. Let us consider the following symmetric version of I presented in Example 4.14.



We have $stb(J) = \{\{a, d\}, \{b, d\}, \{c\}\}$. This means, no argument is sceptically accepted. In order to repair regarding grounded semantics we have to ensure the existence of at least one unattacked argument. Consequently, the least $scep$ - gr -repair is given as $J_{\{c\}}$. As promised by Item 2 in Proposition 4.15 this indeed coincides with the least stb -repair.

Let us briefly consider two further classes of frameworks, namely so-called *compact* and *acyclic* ones. The first one is semantically defined and characterized by the feature that each argument of the AF occurs in at least one extension of the AF (Baumann et al. 2014). For instance, the AF J depicted in Example 4.16 is compact w.r.t. stable semantics. Compact frameworks obviously fulfill item 1 of Proposition 4.15 and thus build an interesting subclass of AFs if interested in credulous reasoning. The second class is syntactically defined and as expected an AF is acyclic if it does not contain any cycles. Such frameworks are known to be *well-founded* (Dung 1995) which means, they possess exactly one complete extension which is grounded, preferred and stable (Coste-Marquis, Devred, and Marquis 2005, Propositions 1 and 2). This means, the agent is able to act (in both reasoning modes) whenever we are faced with an acyclic AF.

Due to space restrictions, we leave a further intensive study of this issue for future work.

5 Computational Complexity

In this section we discuss the computational complexity of two decision problems that naturally arise, namely the existence problem as well as the verification problem regarding minimal repairs. We assume the reader to be familiar with the polynomial hierarchy. Furthermore, we consider the differences class $D_1^p = \{L_1 \cap L_2 \mid L_1 \in NP, L_2 \in coNP\}$ as defined in (Papadimitriou 1994). In the following let σ be a semantics and \diamond one of the considered reasoning modes, i.e., $\diamond \in \{cred, scept\}$.

EX-MIN-REPAIR $_{\sigma, \diamond}$

Input: A knowledge base \mathcal{K}

Output: TRUE iff there is a minimal σ - \diamond -diagnosis for \mathcal{K}

VER-MIN-REPAIR $_{\sigma, \diamond}$

Input: A knowledge base \mathcal{K} , $\mathcal{S} \subseteq \mathcal{K}$

Output: TRUE iff \mathcal{S} minimal σ - \diamond -diagnosis for \mathcal{K}

To keep this section varied and within a reasonable space, we mainly focus on *gr*, *ad*, *co*, *pr* and *stb* semantics.

Credulous Reasoning

We start with the problem of deciding whether a minimal repair exists. As we know from Theorems 4.6, 4.10 and Fact 4.12 it suffices to perform a simple syntactical check, which can be done in linear time. We thus find:

Proposition 5.1. *For $\sigma \in \{ad, gr, eg, il, ss, pr, co, stb\}$, the problem EX-MIN-REPAIR $_{\sigma, cred}$ can be solved in linear time.*

The problem of verifying a potential minimal σ -diagnosis for a given AF F is a bit more challenging. Considering the computational complexity of different reasoning problems in AFs, it is quite unsurprising that VER-MIN-REPAIR $_{A, \sigma}$ is intractable for most semantics σ as it requires checking whether a *non-empty* extension exists. For the semantics *ad*, *stb*, *pr* and *co* this is NP-hard as it requires guessing and verifying a non-empty extension (see e. g. (Dimopoulos and Torres 1996)). Due to the additional minimality check we require, our problem turns out to be in the corresponding difference class.

Theorem 5.2. VER-MIN-REPAIR $_{\sigma, cred}$ is Π_1^P -complete for semantics $\sigma \in \{ad, stb, pr, co\}$.

However, given an AF (and a potential diagnosis), we know that the grounded extension is non-empty if and only if there is an argument which is not attacked. Thus, verifying that \mathcal{S} is a *gr*-diagnosis is quite easy. It turns out that minimality is tractable as well.

Proposition 5.3. VER-MIN-REPAIR $_{gr, cred}$ is in P .

Let us briefly discuss the subclasses of AFs we mentioned in Section 4. If the AF in question is *symmetric*, EX-MIN-REPAIR $_{\sigma, cred}$ trivializes for any σ we considered in this paper since irreflexivity ensures that the AF is not self-controversial. The problem VER-MIN-REPAIR $_{\sigma, cred}$ is trivial for $\sigma \in \{ad, stb, pr, co\}$ since any symmetric AF possesses a credulousl accepted argument. Hence, (H, \mathcal{K}) is a positive instance of VER-MIN-REPAIR $_{\sigma, cred}$ (with $\sigma \in \{ad, stb, pr, co\}$) iff $H = \mathcal{K}$. It is easy to see that VER-MIN-REPAIR $_{gr, cred}$ can be solved in linear time. If the framework is *compact* or *acyclic*, there is inherently nothing to solve w.r.t. credulous reasoning.

Sceptical Reasoning

We now turn to sceptical reasoning. It is not hard to see that EX-MIN-REPAIR $_{\sigma, scep}$ is basically the same since there is a repair if and only if at least one argument does not attack itself.

Proposition 5.4. *For $\sigma \in \{ad, gr, eg, il, ss, pr, co, stb\}$, the problem EX-MIN-REPAIR $_{\sigma, scep}$ can be solved in linear time.*

We thus turn to VER-MIN-REPAIR $_{\sigma, scep}$. This problem is more involved (and presumably more interesting) than in the credulous case. One reason is the following: The decision problem VER-MIN-REPAIR $_{\sigma, scep}$ involves checking all supersets of H within \mathcal{K} . In Theorem 5.2 we had a situation where this resulted in moving to the corresponding difference class. However, the reason was that the underlying decision problem, i. e., “is the framework consistent w.r.t. credulous reasoning?” is in NP. Thus, verifying inconsistency is in coNP. Now, verifying inconsistency for all supersets H' with $H \subseteq H' \subseteq \mathcal{K}$ does not induce a quantifier alternation. That is why the upper bounds in Theorem 5.2 are (under standard assumptions) below Π_2^P . This is not always the case anymore since for the decision problem “is the framework consistent w.r.t. sceptically reasoning?” we naturally face coNP resp. Π_2^P lower bounds.

First observe that for *ad* semantics, the problem is trivial since any framework possess an empty admissible extension. For *gr* semantics the problem is equivalent to the corresponding one for credulous reasoning since the extension is unique, anyway.

Proposition 5.5. VER-MIN-REPAIR $_{gr, scep}$ is in P .

Recall that the unique grounded extension of an AF F is complete as well. Moreover we have $gr(F) \subseteq \bigcap co(F)$. Hence, any framework F possesses a sceptically accepted argument w.r.t. grounded semantics if and only if this is the case for complete semantics. This yields:

Proposition 5.6. VER-MIN-REPAIR $_{co, scep}$ is in P .

Now we consider *pr* semantics. Recall that deciding whether an argument is sceptically accepted is Π_2^P -complete (Dvorák and Dunne 2018). Thus, given a framework \mathcal{K} and a subset $H \subseteq \mathcal{K}$, the decision problem VER-MIN-REPAIR $_{pr, scep}$ involves checking whether for *all* H' with $H \subseteq H' \subseteq \mathcal{K}$ the framework H' does not possess a sceptically accepted argument. Since the latter check is in Σ_2^P for each H' , we immediately see a Π_3^P upper bound for VER-MIN-REPAIR $_{pr, scep}$. Indeed, hardness holds as well.

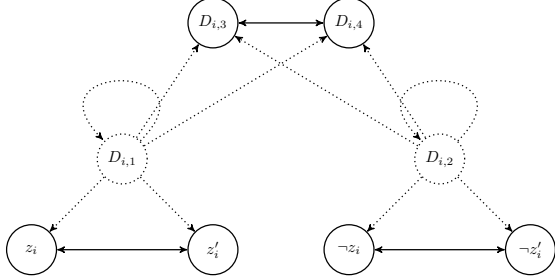
Theorem 5.7. VER-MIN-REPAIR $_{pr, scep}$ is Π_3^P -complete.

Proof. (Sketch.)

Membership was sketched above. Recall the construction from (Dvorák and Dunne 2018) with the property that the AF accepts an argument sceptically w.r.t. preferred semantics if and only if a formula $\Phi = \forall Y \exists X : \phi(X, Y)$ in CNF evaluates to true.

In order to prove hardness in Π_3^P for our problem we somehow need to simulate an additional quantifier. This, however, comes natural since the decision problem VER-MIN-REPAIR $_{pr, scep}$ involves consideration of *all* supersets H' of a given subset $H \subseteq \mathcal{K}$. Given a formula $\Psi = \exists Z \forall Y \exists X : \psi(X, Y, Z)$ in CNF we augment the construction from (Dvorák and Dunne 2018) with the intention that Ψ evaluates to true if and only if (\mathcal{K}, H) is a “no” instance of VER-MIN-REPAIR $_{pr, scep}$ (that is, there exists a

superset H' with $H \subseteq H' \subseteq \mathcal{K}$ such that H' possesses a sceptically accepted argument w.r.t. preferred semantics). The choice of the superset H' with $H \subseteq H' \subseteq \mathcal{K}$ corresponds to the choice of an assignment $\omega : Z \rightarrow \{0, 1\}$ which allows us to recycle the construction from (Dvorák and Dunne 2018) in a rather smooth way. The main idea now is to include the following gadget for any Z variable:



The dummy arguments $D_{i,1}$ and $D_{i,2}$ as well as their attacks are depicted with dotted lines to illustrate that they do *not* occur in H , but in $\mathcal{K} \setminus H$. Including $D_{i,1}$, for example, ensures that z_i and z'_i are never defended and thus occur in no preferred extension. Hence, this choice corresponds to letting z_i be *false*. The $D_{i,3}$ and $D_{i,4}$ are auxiliary arguments to ensure that at least one of $D_{i,1}$ and $D_{i,2}$ needs to be chosen. Moreover, the arguments corresponding to the clauses within ψ in the natural way. Now, there is a superset H' of H (corresponding to an assignment $\omega : Z \rightarrow \{0, 1\}$) where an argument is sceptically accepted if and only if Ψ evaluates to true. The details are omitted due to space restrictions. \square

We turn to stable semantics. Since finding a stable extension is NP-complete it is not hard to see that there is a coNP lower bound for sceptical reasoning. However, as the framework in question might collapse, we also need to verify that there is *at least one* stable extension of a given framework. The result is a D_1^P lower bound (see (Rahwan and Simari 2009)). Interestingly, however, this observation does not change anything in our case. The coNP lower bound is already responsible for $\text{VER-MIN-REPAIR}_{\text{stb}, \text{scep}}$ to have a Π_2^P lower bound. Given $H \subseteq \mathcal{K}$ the decision problem $\text{VER-MIN-REPAIR}_{\text{stb}, \text{scep}}$ involves checking whether *all* sets H' with $H \subseteq H' \subseteq \mathcal{K}$ do *not* possess any sceptically accepted argument. Since the latter test has a NP lower bound, we have a Π_2^P lower bound for $\text{VER-MIN-REPAIR}_{\text{stb}, \text{scep}}$. More precisely:

Theorem 5.8. $\text{VER-MIN-REPAIR}_{\text{stb}, \text{scep}}$ is Π_2^P -complete.

Again let us take a look at *symmetric* AFs. The observation that $\text{EX-MIN-REPAIR}_{\sigma, \text{cred}}$ is trivial for any σ we considered in this paper also holds for sceptical reasoning. For $\sigma \in \{\text{stb}, \text{ss}, \text{pr}\}$ the problem $\text{VER-MIN-REPAIR}_{\sigma, \text{scep}}$ is equivalent to $\text{VER-MIN-REPAIR}_{\text{gr}, \text{scep}}$ (see Proposition 4.15) and can thus be solved in linear time since this is the case for the latter problem as mentioned above. If the framework is *acyclic*, the problem is trivial for $\sigma \in \{\text{stb}, \text{pr}, \text{co}, \text{ad}\}$ since the AF possesses exactly one (nonempty) extension, except for $\sigma = \text{ad}$ which is as usual.

This concludes our discussion on the computational complexity of the two decision problems we named above. As a final remark we want to mention that one can compute *all* *gr*-diagnoses in P . We believe this observation is relevant since the grounded repairs play an essential role as the results from Section 4 suggest. Assume we are given the AF $F = (A, R)$ with $\text{gr}(F) = \{\emptyset\}$. Since a grounded diagnosis needs to ensure that at least one argument $a \in A$ is not attacked anymore, we can successively look at any $a \in A$ and consider $S = \{b \in A \mid (b, a) \in R\}$. If S is a minimal *gr*-diagnosis (verification is in P due to Proposition 5.3), we add S to our list, otherwise we delete it. Since there are at most $|A|$ *gr*-diagnoses, this procedure is in P . So:

Proposition 5.9. Computing all *gr*-diagnoses of a given AF F can be done in P .

Now, even though finding a σ -diagnosis may become rather hard depending on σ , we can efficiently compute all *gr*-diagnoses and then utilize Lemmata 4.8, 4.9 and 4.13 in order to reduce the search space. This approach explains the central role of grounded semantics. In a nutshell, the *gr*-repairs can be seen as a (polynomial time computable) starting point in order to find minimal repairs for other semantics. A thorough investigation of this approach is part of future work. Moreover, the investigation of further subclasses of AFs seems rather promising considering computational complexity. For example, the ones we considered trivialize credulous diagnoses in almost any case. Other restrictions might ensure tractability of certain problems we considered here while being less restrictive.

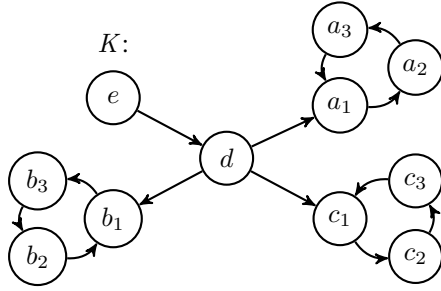
6 How to Repair? — A Short Case Study

As mentioned before, due to Lemmata 4.8, 4.9 and 4.13 we may reduce the search space for diagnoses as long as we are equipped with an already computed grounded one. If one is interested in *all* diagnoses, the notion of strong inconsistency in order to use the hitting set duality might be useful. We discuss both credulous and sceptical reasoning.

First let us consider an example with stable semantics. Let us start with credulous reasoning. It is well-known that in case of finite AFs the non-existence of acceptable positions implies the existence of odd-cycles. This means, by contraposition, one possible strategy for repairing AFs in case of stable semantics is to break odd-cycles. This approach corresponds to the minimal *stb*-repairs $F_{\{a\}}$, $F_{\{b\}}$ and $F_{\{c\}}$ from our running example F . Since possessing odd-cycles is not sufficient for the collapse of stable semantics further considerations are required. Indeed, in case of our running example, we have seen that eliminating the arguments e and f results in a minimal *stb*-repair, namely $F_{\{e, f\}}$, too. Regarding the principle of minimal change one may argue that breaking the odd-cycle in F has to be preferred over the latter strategy since less arguments are involved. The following slightly modified version of our running example shows that this observation is not true in general. A further intensive study of this issue will be part of future work.

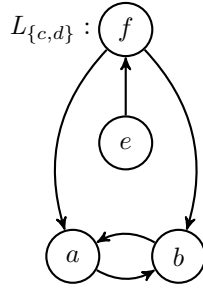
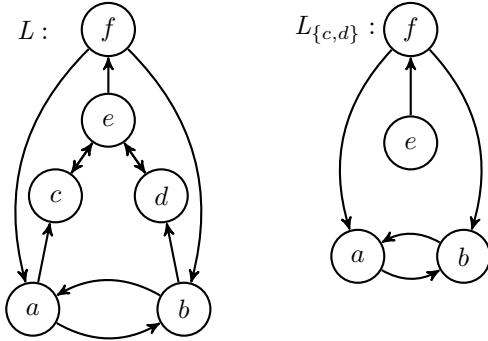
Example 6.1. Consider the following AF K . One may easily confirm that there are 9 minimal *cred-stb*-diagnoses, namely $\{a_i, b_j, c_k\}$ with $i, j, k \in \{1, 2, 3\}$. They comply

with the idea to break all odd loops of the given AF. However, $\{e\}$ is a minimal *cred-stb*-diagnosis as well, and arguably the most immediate one.

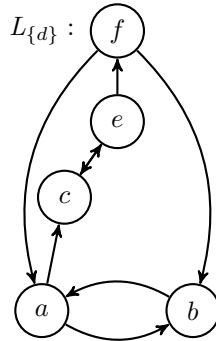
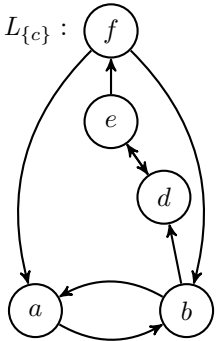


The subsequent example considers a semantical defect w.r.t. preferred semantics which is tackled via grounded repairs.

Example 6.2. The AF L exemplifies a situation where preferred semantics do not possess any sceptically accepted argument. More precisely, $\bigcap pr(L) = \emptyset$ since $pr(L) = \{\{a, e\}, \{b, e\}, \{c, d, f\}\}$.



Our goal is to find a minimal *scep-pr*-diagnosis S , i.e., a set S such that $\bigcap pr(L_S) \neq \emptyset$ and $pr(L_S) \neq \emptyset$. Lemma 4.9 suggests that looking for *gr*-repairs is a reasonable starting point. In order to guarantee at least one unattacked argument one finds $\{c, d\}$ as minimal *gr*-diagnosis. Let $L_{\{c,d\}}$ denote the associated minimal repair. We have $gr(L_{\{c,d\}}) = \{\{e\}\}$. Hence, $\bigcap pr(L_{\{c,d\}}) \neq \emptyset$ is implied. This means, $\{c, d\}$ is a *scep-pr*-diagnosis. Moreover, $\{c, d\}$ is even minimal proven by the following two AFs $L_{\{c\}}$ and $L_{\{d\}}$.



Indeed, we have $\bigcap pr(L_{\{c\}}) = \bigcap pr(L_{\{d\}}) = \emptyset$ since $\{a, e\}, \{d, f\} \in pr(L_{\{c\}})$ and $\{a, e\}, \{c, f\} \in pr(L_{\{d\}})$.

7 Related Work and Future Directions

We studied the notion of repairing in the context of abstract argumentation. This topic as introduced in (Reiter 1987) is less developed in this area. The closest one to our work is (Nouioua and Würbel 2014). The authors define an operator and provide an algorithm, s.t. the resulting framework does not collapse. The mentioned work considers a semantical defect as the absence of any extension. Consequently, only stable semantics can be considered in contrast to our setup which additionally includes a treatment of semantics which may provide the empty set as unique extension. All semantics known from the literature do so. Moreover, restoring consistency is achieved via dropping a minimal set of attacks. All arguments survive the revision process. We mention that this approach can be modeled using a slightly different setup as used in this paper. The key idea is to change the canonical representation as presented in Definition 3.1. A precise implementation of this idea will be part of future work.

The very first and basic works which are dealing with dynamics in abstract argumentation are (Baumann and Brewka 2010; Baumann 2012; Baumann and Brewka 2013) as well as (Cayrol, de Saint-Cyr, and Lagasque-Schiex 2010; Bisquert et al. 2011). The first two are tackling the so-called *enforcing problem* w.r.t. possibility as well as minimal change. More precisely, they are dealing with the question whether it is possible (and if yes, as little effort as possible) to add new information in such a way that a desired set of arguments becomes an extension or at least a subset of one. In (Kim, Ordyniak, and Szeider 2013) the authors studied this problem under the name *σ -repair* and provided parametrized complexity results. Although adding information as well as desired sets are not the focus of our study there is at least one interesting similarity to our work, namely: given an AF where nothing is credulously accepted, then enforcing a certain non-empty set can be seen as a special kind of repairing. The other two works are case studies of what happens with the set of extensions if one deletes or adds one argument. The so-called *destructive change* is somehow the inverse of our notion of credulous repair since the initial framework possesses at least one credulously accepted argument whereas the result does not.

Several future directions are already mentioned in the text. For instance, a further intensive study of subclasses of AFs seems to be very promising since certain useful semantical properties are already ensured by syntactic properties. Moreover, it is already known that AFs can be seen as a restricted class of logic programs (LPs). More precisely, as shown in (Strass 2013, Theorem 4.13) there is a standard translation T from AFs to LPs, s.t. for any AF F , $\sigma(F)$ coincides with $\tau(T(F))$ for certain pairs of semantics σ and τ . This means, one interesting research question is to which extent our results can be conveyed to repairing in logic programming.

Acknowledgments

This work was partially funded by Deutsche Forschungsgemeinschaft DFG (Research Training Group 1763; project BR 1817/7-2).

References

- Baroni, P.; Caminada, M.; and Giacomin, M. 2011. An introduction to argumentation semantics. *The Knowledge Engineering Review* 26:365–410.
- Baumann, R., and Brewka, G. 2010. Expanding argumentation frameworks: Enforcing and monotonicity results. In *Proceedings of COMMA 2010*, 75–86.
- Baumann, R., and Brewka, G. 2013. Spectra in abstract argumentation: An analysis of minimal change. In *LPNMR, Proceedings of 12th International Conference in Logic Programming and Nonmonotonic Reasoning*, 174–186.
- Baumann, R., and Brewka, G. 2015. AGM meets abstract argumentation: Expansion and revision for Dung frameworks. In *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI'15*, 2734–2740. AAAI Press.
- Baumann, R., and Spanring, C. 2015. Infinite argumentation frameworks - On the existence and uniqueness of extensions. In *Advances in Knowledge Representation, Logic Programming, and Abstract Argumentation - Essays Dedicated to Gerhard Brewka on the Occasion of His 60th Birthday*, 281–295.
- Baumann, R., and Spanring, C. 2017. A study of unrestricted abstract argumentation frameworks. In *IJCAI, Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 807–813.
- Baumann, R.; Dvorák, W.; Linsbichler, T.; Strass, H.; and Woltran, S. 2014. Compact argumentation frameworks. In *21st European Conference on Artificial Intelligence*, 69–74.
- Baumann, R. 2012. What does it take to enforce an argument? minimal change in abstract argumentation. In *20th European Conference on Artificial Intelligence*, 127–132.
- Bienvenu, M. 2012. On the complexity of consistent query answering in the presence of simple ontologies. In *Proceedings of the Twenty-Sixth Conference on Artificial Intelligence*.
- Bisquert, P.; Cayrol, C.; de Saint-Cyr, F. D.; and Lagasquie-Schiex, M. 2011. Change in argumentation systems: Exploring the interest of removing an argument. In *Scalable Uncertainty Management - SUM*, 275–288.
- Brewka, G.; Thimm, M.; and Ulbricht, M. 2017. Strong inconsistency in nonmonotonic reasoning. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, 901–907.
- Cayrol, C.; de Saint-Cyr, F. D.; and Lagasquie-Schiex, M. 2010. Change in abstract argumentation frameworks: Adding an argument. *Journal of Artificial Intelligence Research* 38:49–84.
- Coste-Marquis, S.; Konieczny, S.; Mailly, J.; and Marquis, P. 2014. On the revision of argumentation systems: Minimal change of arguments statuses. In *Proceedings of Principles of Knowledge Representation and Reasoning, KR*.
- Coste-Marquis, S.; Devred, C.; and Marquis, P. 2005. Symmetric argumentation frameworks. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, 317–328. Springer.
- Diller, M.; Haret, A.; Linsbichler, T.; Rümmele, S.; and Woltran, S. 2018. An extension-based approach to belief revision in abstract argumentation. *Int. J. Approx. Reasoning* 93:395–423.
- Dimopoulos, Y., and Torres, A. 1996. Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science* 170(1-2):209–244.
- 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):321–358.
- Dvorák, W., and Dunne, P. E. 2018. Computational problems in formal argumentation and their complexity. In Baroni, P.; Gabbay, D.; Giacomin, M.; and van der Torre, L., eds., *Handbook of Formal Argumentation*. College Publications. also appears in *IfCoLog Journal of Logics and their Applications* 4(8):2623–2706.
- Eiter, T.; Fink, M.; and Stepanova, D. 2013. Data repair of inconsistent dl-programs. In *Proceedings of International Joint Conference on Artificial Intelligence, IJCAI*, 869–876.
- Hansson, S. O. 1994. *Taking Belief Bases Seriously*. Dordrecht: Springer Netherlands. 13–28.
- Kim, E. J.; Ordyniak, S.; and Szeider, S. 2013. The complexity of repairing, adjusting, and aggregating of extensions in abstract argumentation. In *Theory and Applications of Formal Argumentation - Second International Workshop, TAFE*, 158–175.
- Lembo, D.; Lenzerini, M.; Rosati, R.; Ruzzi, M.; and Savo, D. F. 2011. Inconsistency-tolerant semantics for description logic ontologies (extended abstract). In *Sistemi Evolutivi per Basi di Dati - SEBD 2011, Proceedings of the Nineteenth Italian Symposium on Advanced Database Systems*, 287–294.
- McIlraith, S. 1999. Towards a theory of diagnosis, testing and repair. In *Proc. of the 5th International Workshop on Principles of Diagnosis (DX)*.
- Nouioua, F., and Würbel, E. 2014. Removed set-based revision of abstract argumentation frameworks. In *26th IEEE International Conference on Tools with Artificial Intelligence, ICTAI 2014*, 784–791.
- Papadimitriou, C. 1994. *Computational Complexity*. Addison-Wesley.
- Potyka, N., and Thimm, M. 2014. Consolidation of probabilistic knowledge bases by inconsistency minimization. In *ECAI 2014 - 21st European Conference on Artificial Intelligence*, 729–734.
- Rahwan, I., and Simari, G. R. 2009. *Argumentation in artificial intelligence*, volume 47. Springer.
- Reiter, R. 1987. A theory of diagnosis from first principles. *Artificial Intelligence* 32(1):57–95.
- Sakama, C., and Inoue, K. 2003. An abductive framework for computing knowledge base updates. *Theory and Practice of Logic Programming* 3(6):671–713.
- Strass, H. 2013. Approximating operators and semantics for abstract dialectical frameworks. *Artificial Intelligence* 205:39–70.