

Complexity-Sensitive Decision Procedures for Abstract Argumentation*

Wolfgang Dvořák¹ and Matti Järvisalo² and Johannes Peter Wallner¹ and Stefan Woltran¹

¹Institute of Information Systems, Vienna University of Technology, Austria
{dvorak, wallner, woltran}@dbai.tuwien.ac.at

²HIIT & Department of Computer Science, University of Helsinki, Finland
matti.jarvisalo@cs.helsinki.fi

Abstract

Abstract argumentation frameworks (AFs) provide the basis for various reasoning problems in the areas of Knowledge Representation and Artificial Intelligence. Efficient evaluation of AFs has thus been identified as an important research challenge. So far, implemented systems for evaluating AFs have either followed a straight-forward reduction-based approach or been limited to certain tractable classes of AFs. In this work, we present a generic approach for reasoning over AFs, based on the novel concept of complexity-sensitivity. Establishing the theoretical foundations of this approach, we derive several new complexity results for preferred, semi-stable and stage semantics which complement the current complexity landscape for abstract argumentation, providing further understanding on the sources of intractability of AF reasoning problems. The introduced generic framework exploits decision procedures for problems of lower complexity whenever possible. This allows, in particular, instantiations of the generic framework via harnessing in an iterative way current sophisticated Boolean satisfiability (SAT) solver technology for solving the considered AF reasoning problems. First experimental results show that the SAT-based instantiation of our novel approach outperforms existing systems.

Introduction

Formal argumentation has evolved as an important field in knowledge representation and reasoning with abstract argumentation frameworks (AFs for short), as introduced by Dung (1995), being its central formalization, providing a simple yet powerful formalism to reason about conflicts between arguments. The power of the formalism, however, comes at a price. In particular, many important reasoning problems for AFs are located on the second level of the polynomial hierarchy, including reasoning in the preferred semantics (Dunne and Bench-Capon 2002), the semi-stable and the stage semantics (Dvořák and Woltran 2010). This naturally raises the question about the origin of this high complexity and, in particular, calls for research on lower complexity fragments of the reasoning tasks. The focus

of this paper is both on the identification of such lower-complexity fragments of second-level reasoning problems arising from abstract argumentation, and on exploiting this knowledge in developing efficient *complexity-sensitive* decision procedures for the generic second-level problems.

Tractable (i.e., polynomial-time decidable) fragments have been quite thoroughly studied in the literature (see, e.g., (Coste-Marquis, Devred, and Marquis 2005; Dunne 2007; Dvořák, Szeider, and Woltran 2010; Dvořák, Pichler, and Woltran 2011; Ordyniak and Szeider 2011)). However, there is only little work on the identification of fragments which are located on the first level (NP-coNP layer), that is, *inbetween* tractability and full second-level complexity.

Identification of first-level fragments of second-level reasoning tasks is important due to several reasons. First, from a theoretical point of view, such fragments show particular (but not all) sources of complexity of the considered problems and pave the way towards “trichotomy”-like results (e.g. (Truszczynski 2011) in the context of answer-set programming). Second, NP fragments can be efficiently reduced to the problem of satisfiability in classical propositional logic (SAT). This allows for realizations of argumentation procedures by employing highly sophisticated SAT solver technology in reasoning on argumentation problems.

Going even further, we aim at designing decision procedures for *larger* fragments based on decision procedures developed for an NP-fragment, using the NP decision procedures as an NP oracle in an iterative fashion. Such procedures fall under the general *counter-example guided abstraction refinement* (CEGAR) approach originating from the field of model checking (Clarke et al. 2003; Clarke, Gupta, and Strichman 2004). For problems complete for the second level of the polynomial hierarchy, this leads to a general procedure which, in the worst case, requires an exponential number of calls to the NP oracle, which is indeed unavoidable under the assumption that the polynomial hierarchy does not collapse. Nevertheless, such procedures can be designed to behave adequately on input instances that fall into the considered NP fragment and on instances for which a relatively low number of oracle calls is sufficient. As a generic notion, we say that such a procedure is *complexity-sensitive* w.r.t. the NP fragment at hand. For instance, for the second level problem of answer-set existence for disjunctive logic programs, the successful loop-formula

*This work has been funded by Vienna Science and Technology Fund (WWTF) through project ICT08-028 and by Academy of Finland (grant 132812).
Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

approach (see, e.g. (Lierler 2005)) yields a polynomial reduction to SAT for the fragment of tight programs, although in general the resulting SAT instance is of exponential size. This approach gives thus a practical decision procedure for (the second level problems of) answer-set programming that is at the same time complexity-sensitive w.r.t. the NP fragment of tight disjunctive programs.

In this paper we identify various lower-complexity fragments of second-level reasoning problems arising from abstract argumentation, and show how some of the fragments can be exploited in complexity-sensitive CEGAR-style decision procedures for the generic second-level problems. The fragments identified and exploited are based on notions of “distance” to particular NP fragments. This leads to the intuition that, the higher the distance, the more iterative calls to the NP oracle are needed. We also employ the concept of distance to *generalize* known classes of NP fragments.

In more detail, we focus on three important semantics for abstract argumentation, namely preferred, semi-stable and stage semantics. Our complexity analysis is based on six different classes of argumentation frameworks which are known to yield milder complexity results for at least one of these semantics. We first present complexity results for these classes in cases where the exact complexity for a particular semantics has not been established yet. Moreover, we categorize the classes into *syntactical* and *semantical* families. For the former family, we consider the known concepts of acyclic and odd-cycle free AFs as well as a new class (so-called weak cyclic AFs). As semantical subclasses we consider the prominent class of coherent AFs (Dunne and Bench-Capon 2002); the class of AFs which possess at least one stable extension (stable-consistent AFs); and the class of AFs which possess a unique preferred extension.

In a second step, we consider different notions of distance in order to capture AFs which are “close” to one of the aforementioned classes. We consider the following realizations of distance: (i) *graph-based distance measures*, where the parameter is given by the number of arguments to be deleted from a given AF in order to fall into a specified class; (ii) *extension-based distance measures* which apply to the semantical subclasses. For instance (among others), starting from the class of coherent AFs (where the preferred and stable extensions coincide), we consider as parameter the number of additional preferred extensions.

The main contributions of the paper are the following.

- We show new complexity results for acceptance problems in argumentation on certain fragments. In particular, for the class of frameworks which possess a unique preferred extension, semi-stable semantics yields milder complexity than stage semantics. To the best of our knowledge, this is the first result that indicates a difference between the complexities of these two semantics.
- We show that graph-based distance measures are in most cases tight: already a small distance from the subclass at hand leads to the full second-level complexity. This reveals that syntactic fragments based on such distance measures do not hint towards complexity-sensitive decision procedures.
- Towards the design of complexity-sensitive decision pro-

cedures, we also identify extension-based distance measures and show that certain problems can be solved by a bounded number (in terms of the distance) of NP-oracle calls.

- Exploiting the suitable extension-based distance measures, we develop a generic framework of complexity-sensitive decision procedures for the different second-level reasoning problems within abstract argumentation. We present our procedures in terms of (first-level) argumentation problems, i.e., we give novel characterizations of preferred, semi-stable, and stage semantics in terms of simpler semantics (such as stable and complete). The actual computation of the simpler semantics can be instantiated in various ways.

- We show in detail how the generic framework can be instantiated using a SAT-based CEGAR-style approach. For this, we develop novel SAT-encodings for the oracle calls, differing from previously suggested SAT-encodings of first-level AF reasoning problems (Besnard and Doutre 2004). Notably, we exploit possibilities of learning from counterexamples both on the level of the original argumentation framework as well as the SAT oracle during computation. Importantly, while monolithic SAT-encodings of second-level argumentation problems are deemed to be of exponential size, our procedures are truly complexity-sensitive in that the exponential space requirements may be circumvented in cases where it suffices to consider a small part of a monolithic encoding to decide the actual query.

- We have implemented a prototype of the SAT-based instantiation of our approach, exploiting a state-of-the-art *conflict-driven clause learning* (CDCL) SAT solver as the underlying NP-oracle. First experiments show the high potential of the proposed approach compared to other state-of-the-art implementations for abstract argumentation, in particular the logic-programming approach based on monolithic encodings of second-level argumentation problems (Egly, Gaggl, and Woltran 2010; Dvořák et al. 2011a).

Preliminaries

In this section we review (abstract) argumentation frameworks (Dung 1995) and the semantics studied in this paper (see also (Baroni and Giacomin 2009)).

Definition 1. An argumentation framework (AF) is a pair $F = (A, R)$ where A is a finite set of arguments and $R \subseteq A \times A$ is the attack relation. For a given AF $F = (A, R)$ we use A_F to denote the set A of its arguments and R_F to denote its attack relation R . We sometimes use the notation $a \succ^R b$ instead of $(a, b) \in R$. For $S \subseteq A$ and $a \in A$, we also write $S \succ^R a$ (resp. $a \succ^R S$) in case there is a $b \in S$, such that $b \succ^R a$ (resp. $a \succ^R b$). In case no ambiguity arises, we use \succ instead of \succ^R .

Semantics for argumentation frameworks assign to each AF $F = (A, R)$ a set $\sigma(F) \subseteq 2^A$ of extensions. We consider here for σ the functions *stb*, *adm*, *prf*, *com*, *stg*, and *sem* which stand for stable, admissible, preferred, complete, stage, and respectively, semi-stable semantics. Before giving the actual definitions for these semantics, we need to define a few more formal concepts.

Definition 2. Given an AF $F = (A, R)$, an argument $a \in A$ is defended (in F) by a set $S \subseteq A$ if for each $b \in A$, such that $b \rightarrow a$, also $S \rightarrow b$ holds. Moreover, for $S \subseteq A$, we denote by S_R^+ the set $S \cup \{b \mid S \rightarrow b\}$.

Definition 3. Let $F = (A, R)$ be an AF. A set $S \subseteq A$ is conflict-free (in F), denoted $S \in cf(F)$, iff there are no $a, b \in S$, such that $(a, b) \in R$. For $S \in cf(F)$, it holds that

- $S \in stb(F)$, if for each $a \in A \setminus S$, $S \rightarrow a$, i.e. $S_R^+ = A$;
- $S \in adm(F)$, if each $a \in S$ is defended by S ;
- $S \in prf(F)$, if $S \in adm(F)$ and there is no $T \in adm(F)$ with $T \supset S$;
- $S \in com(F)$, if $S \in adm(F)$ and for each $a \in A$ defended by S , $a \in S$ holds;
- $S \in stg(F)$, if there is no $T \in cf(F)$, with $T_R^+ \supset S_R^+$;
- $S \in sem(F)$, if $S \in adm(F)$ and there is no $T \in adm(F)$ with $T_R^+ \supset S_R^+$.

We recall that for each AF F , $stb(F) \subseteq sem(F) \subseteq prf(F) \subseteq com(F) \subseteq adm(F)$ holds, and that for each of the considered semantics σ (except stable) $\sigma(F) \neq \emptyset$ holds. Moreover, in case an AF has at least one stable extension, its stable, semi-stable, and respectively, stage extensions coincide.

Example 1. Consider the AF $F = (A, R)$, with $A = \{a, b, c, d, e\}$ and $R = \{(a, b), (c, b), (c, d), (d, c), (d, e), (e, e)\}$. The graph representation of F is given as follows.



Here $stb(F) = stg(F) = sem(F) = \{\{a, d\}\}$. The admissible sets of F are $\emptyset, \{a\}, \{c\}, \{d\}, \{a, c\}, \{a, d\}$, and $prf(F) = \{\{a, c\}, \{a, d\}\}$. The complete extensions are $\{a\}, \{a, c\}$ and $\{a, d\}$. \diamond

We now recall the complexity of reasoning in AFs for the following decision problems under different semantics σ .

- **Credulous Acceptance** ($Cred_\sigma$): Given an AF F and argument $a \in A_F$, is a contained in some $S \in \sigma(F)$?
- **Skeptical Acceptance** ($Skept_\sigma$): Given an AF F and argument $a \in A_F$, is a contained in each $S \in \sigma(F)$?
- **Verification** (Ver_σ): Given an AF F and a set $S \subseteq A_F$, is $S \in \sigma(F)$?
- **Existence** ($Exists_\sigma$): Given an AF F , is $\sigma(F) \neq \emptyset$?
- **Non-emptiness** ($Exists_\sigma^{-\emptyset}$): Given an AF F , is there a set $S \subseteq A_F, S \neq \emptyset$ such that $S \in \sigma(F)$?

In accordance with the above problems we say that an argument is credulously (resp. skeptically) accepted iff it is contained in at least one extension (resp. in all extensions).

Table 1 summarizes the computational complexity of these problems (Coste-Marquis, Devred, and Marquis 2005; Dimopoulos and Torres 1996; Dung 1995; Dunne and Bench-Capon 2002; Dunne and Caminada 2008; Dvořák and Woltran 2010; Dvořák and Woltran 2011).

We will focus on $Skept_{prf}$, $Cred_{sem}$, $Skept_{sem}$, $Cred_{stg}$ and $Skept_{stg}$: as Table 1 indicates, these problems are the ones on the second-level of the polynomial hierarchy.

Table 1: Complexity of decision problems for AFs.

σ	$Cred_\sigma$	$Skept_\sigma$	Ver_σ	$Exists_\sigma$	$Exists_\sigma^{-\emptyset}$
<i>stb</i>	NP-c	coNP-c	in L	NP-c	NP-c
<i>adm</i>	NP-c	trivial	in L	trivial	NP-c
<i>com</i>	NP-c	P-c	in L	trivial	NP-c
<i>prf</i>	NP-c	Π_2^P -c	coNP-c	trivial	NP-c
<i>sem</i>	Σ_2^P -c	Π_2^P -c	coNP-c	trivial	NP-c
<i>stg</i>	Σ_2^P -c	Π_2^P -c	coNP-c	trivial	in L

Subclasses of Argumentation Frameworks

In this section we review several classes of AFs where reasoning with preferred, stage or semi-stable semantics becomes easier compared to the results for the general case. Both earlier and new results are discussed. First, we consider the classes of acyclic and weakly cyclic AFs.

Definition 4. An AF F is acyclic if there is no directed cycle of attacks in F ; F is weakly cyclic if F can be made acyclic by deleting one argument (and its incident attacks) from each strongly connected component (SCC) of F . We denote these classes of AFs by *acyc* and *wcyc*.

One can easily show that deciding whether a given AF falls into one of these classes can be done efficiently.

It is well known that the problems we are interested here become tractable when restricted to acyclic AFs. For weakly cyclic AFs (these are the AFs where the graph parameter cycle-rank is at most 1 (Dvořák, Pichler, and Woltran 2011)), we can make direct use of the following complexity result.

Proposition 1 ((Dvořák, Pichler, and Woltran 2011)). *For weakly cyclic AFs, the problem $Skept_{prf}$ is coNP-complete.*

The reasoning problems for stage and semi-stable semantics still maintain their full complexity when restricted to weak cyclic AFs (Dvořák, Pichler, and Woltran 2011).

We now turn to semantical subclasses.

Definition 5. An AF F for which $stb(F) \neq \emptyset$ is stable-consistent. We denote the class of such AFs by *stablecons*.

We recall that testing for the existence of a stable extension is NP-complete (Dimopoulos and Torres 1996).

The following result is immediate from the fact that, in case an AF has at least one stable extension, its stable, semi-stable, and respectively, stage extensions coincide.

Proposition 2. *For stable-consistent AFs, $Cred_\sigma$ is NP-complete and $Skept_\sigma$ is coNP-complete, $\sigma \in \{sem, stg\}$.*

However, in case of the preferred semantics, stable consistency is of no help for deciding skeptical acceptance.

Proposition 3. *For stable-consistent AFs, the problem $Skept_{prf}$ is Π_2^P -complete,*

Proof. Hardness follows from a reduction in (Dunne and Bench-Capon 2002) that maps the Π_2^P -hard problem of deciding whether given an a QBF $\Phi = \forall Y \exists Z \varphi(Y, Z)$, where φ is a CNF formula $\bigwedge_{c \in C} c$ with each clause $c \in C$ a disjunction of literals from $X = Y \cup Z$, is true to $Skept_{prf}$.




Figure 1: F_Φ for the QBF $\Phi = \forall y_1 y_2 \exists z_3 z_4 (y_1 \vee y_2 \vee z_3) \wedge (y_2 \vee \neg z_3 \vee \neg z_4) \wedge (y_2 \vee z_3 \vee z_4)$.

We build $F_\Phi = (A_\Phi, R_\Phi)$ as follows (see also Figure 1): $A_\Phi = \{\varphi, b\} \cup C \cup X \cup \bar{X}$; and $R_\Phi = \{(c, \varphi) \mid c \in C\} \cup \{(x, \bar{x}), (\bar{x}, x) \mid x \in X\} \cup \{(x, c) \mid x \text{ occurs in } c\} \cup \{(\bar{x}, c) \mid \neg x \text{ occurs in } c\} \cup \{(\varphi, b), (b, b)\} \cup \{(b, z), (b, \bar{z}) \mid z \in Z\}$. We have that Φ is valid iff the argument φ is skeptically accepted in F_Φ (Dunne and Bench-Capon 2002). Moreover each model M of φ corresponds to a stable extension of F_Φ . Since the QBF-problem remains hard for instances where φ has at least one model, the result follows. \square

The next class relates preferred and stable extensions and has been discussed in (Dunne and Bench-Capon 2002).

Definition 6. An AF F is coherent if $\text{prf}(F) = \text{stb}(F)$. We denote the class of such AFs by coherent.

Proposition 4. For coherent AFs and $\sigma \in \{\text{prf}, \text{sem}, \text{stg}\}$, Cred_σ is NP-complete and Skept_σ is coNP-complete.

Proof. The result for $\sigma = \text{prf}$ is clear by definition when taking the complexity of stable semantics (see Table 1) into account. Since each AF possesses at least one preferred extension, each coherent AF is also stable-consistent. The remaining results thus follow immediately as well. \square

Testing coherence is in general even worse (Π_2^P -complete (Dunne and Bench-Capon 2002)) than testing stable-consistency. At first glance this restricts the practical value of this fragment, but we can identify a class of easy detectable coherent AFs, namely the AFs without odd-length cycles. Thus, we introduce one further syntactical subclass.

Definition 7. An AF F is odd-cycle free if there is no directed cycle consisting of an odd number of attacks in F . We denote the class of odd-cycle free AFs by ocf.

In fact, testing for odd-length cycles in digraphs can be done in polynomial time (see e.g. (Bang-Jensen and Gutin 2010)).

Proposition 5. For AFs $F \in \text{ocf}$, and $\sigma \in \{\text{prf}, \text{sem}, \text{stg}\}$, Cred_σ is NP-complete and Skept_σ is coNP-complete.

Proof. Membership is immediate from the well-known result by Dung (1995), that every AF without odd-length cycles is coherent. Hardness results from a standard reduction for stable semantics (Dimopoulos and Torres 1996). \square

The final fragment we introduce is another semantical one. It makes use of the complexity gap between credulous and skeptical acceptance for preferred semantics.

Table 2: Complexity when the AF belongs to a sub-class \mathcal{G} .

\mathcal{G}	$\text{Skept}_{\text{prf}}$	Cred_{sem}	$\text{Skept}_{\text{sem}}$	Cred_{stg}	$\text{Skept}_{\text{stg}}$
acyc	P-c	P-c	P-c	P-c	P-c
wcyc	coNP-c	Σ_2^P -c	Π_2^P -c	Σ_2^P -c	Π_2^P -c
ocf	coNP-c	NP-c	coNP-c	NP-c	coNP-c
stablecons	Π_2^P -c	NP-c	coNP-c	NP-c	coNP-c
coherent	coNP-c	NP-c	coNP-c	NP-c	coNP-c
uniqpref	in NP	in NP	in NP	Σ_2^P -c	Π_2^P -c

Definition 8. We denote the class of AFs F satisfying $|\text{prf}(F)| = 1$ by uniqpref.

One can easily show that testing whether there is at most one preferred extensions is coNP-complete.

Proposition 6. For AFs $F \in \text{uniqpref}$ and $\sigma \in \{\text{prf}, \text{sem}\}$, problems Cred_σ and Skept_σ are NP-easy.

Proof. On AFs with a unique extension credulous and skeptical acceptance coincide. Moreover, for each $F \in \text{uniqpref}$, $\text{sem}(F) = \text{prf}(F)$ holds since the existence of a semi-stable extension is guaranteed for finite AFs, and each semi-stable extension is also preferred. \square

It is open whether these problems are also NP-hard. However, we can show NP-hardness under so-called *randomized reductions* via a version of the SAT problem where it is guaranteed that there is at most one model (Valiant and Vazirani 1986) and a standard reduction from SAT to Cred_{prf} . For the stage semantics, the complexity of acceptance problems remains as high as in the general case, however.

Proposition 7. For AFs $F \in \text{uniqpref}$, Cred_{stg} is Σ_2^P -complete and $\text{Skept}_{\text{stg}}$ is Π_2^P -complete.

Proof. Consider an arbitrary AF $F = (A, R)$ and let $t \notin A$ be a fresh argument. We construct $F' = (A \cup \{t\}, R \cup \{t, t\} \cup \{(t, a) \mid a \in A\})$. It is easy to see that $\text{cf}(F) = \text{cf}(F')$ and thus also $\text{stg}(F) = \text{stg}(F')$. Further we have that $\text{prf}(F') = \{\emptyset\}$. Hence, any decision problem for stg can be directly expressed in AFs from uniqpref. \square

To summarize, we have introduced several kinds of AF-subclasses. They can be grouped into syntactical (acyc, wcyc, ocf), and semantical classes (stablecons, coherent, uniqpref). The complexity results are summarized in Table 2. In the next two sections, we study possibilities of extending the “good” complexity behavior of these classes. To this end, we introduce certain distance measures with the aim of maintaining lower complexity as long as the distance to such a class is bound.

Graph-Based Distance Measures

A natural way to generalize the introduced subclasses is to consider the minimal number of arguments we have to delete from an AF such that the modified AF falls into the respective class (see also (Ordyniak and Szeider 2011)). This gives rise to the following distance measure.




Figure 2: F_{Φ}^{sem} for the QBF $\Phi = \forall y_1 y_2 \exists z_3 z_4 (y_1 \vee y_2 \vee z_3) \wedge (y_2 \vee \neg z_3 \vee \neg z_4) \wedge (y_2 \vee z_3 \vee z_4)$.

Definition 9. Let \mathcal{G} be a graph class and $F = (A, R)$ an AF. We define $dist_{\mathcal{G}}(F)$ as the minimal number k such that there is a set $S \subseteq A$ with $|S| = k$ and $(A \setminus S, R \cap (A \setminus S \times A \setminus S)) \in \mathcal{G}$. If there is no such set S we define $dist_{\mathcal{G}}(F) = \infty$.

We start with the class *wcyc* of weakly cyclic AFs. Recall that this class decreases complexity only in the case of skeptical acceptance with respect to preferred semantics.

Proposition 8. $Skept_{prf}$ is Π_2^P -hard for AFs F with $dist_{wcyc}(F) = 1$.

Proof. Π_2^P -hardness is established by re-using the reduction from the proof of Proposition 3 and observing that deletion of the argument b in F_{Φ} results in a framework with SCCs of size at most 2 which is easily seen to be weakly cyclic. \square

In words, the subclass *wcyc* is tight w.r.t. the introduced distance in the sense that a single argument violating membership in *wcyc* is sufficient for the general Π_2^P -hardness. An analogous result can be shown for the class *ocf*.

Proposition 9. $Skept_{prf}$ is Π_2^P -hard for AFs F with $dist_{ocf}(F) = 1$.

Proof. Similar as for Proposition 8, deleting argument b yields an AF that is free of odd-length cycles. \square

The same effect can be shown for semi-stable acceptance.

Proposition 10. $Cred_{sem}$ is Σ_2^P -hard and $Skept_{sem}$ is Π_2^P -hard, for AFs F with $dist_{ocf}(F) = 1$.

Proof. For a QBF $\Phi = \forall Y \exists Z \varphi(Y, Z)$ with $\varphi = \bigwedge_{c \in C} c$ in CNF build the AF $F_{\Phi}^{sem} = (A_{\Phi}, R_{\Phi})$ with $X = Y \cup Z$ as follows (see also Figure 2): $A_{\Phi} = \{\varphi, \bar{\varphi}, b, g\} \cup C \cup X \cup \bar{X} \cup Y' \cup \bar{Y}'$; and $R_{\Phi} = \{(c, \varphi) \mid c \in C\} \cup \{(x, \bar{x}), (\bar{x}, x) \mid x \in X\} \cup \{(x, c) \mid x \text{ occurs in } c\} \cup \{(\bar{x}, c) \mid \neg x \text{ occurs in } c\} \cup \{(y, y'), (\bar{y}, \bar{y}') \mid y \in Y\} \cup \{(\varphi, b), (g, g), (g, b)\} \cup \{(g, y'), (g, \bar{y}') \mid y \in Y\} \cup \{(\varphi, \bar{\varphi}), (\bar{\varphi}, \varphi)\}$.

This reduction is a variation of the reduction presented in (Dvořák and Woltran 2010) and can be easily shown to be equivalent. Hence φ is skeptically accepted iff $\bar{\varphi}$ is not credulously accepted iff Φ is true. Finally deleting the argument g results in an AF that is free of odd-length cycles. \square




Figure 3: Illustration of the AF F_{φ, x_1} for $\varphi = (x_1 \vee x_2 \vee x_3) \wedge (\neg x_2 \vee \neg x_3 \vee \neg x_4) \wedge (\neg x_1 \vee \neg x_2 \vee x_4)$.

For the stage semantics, we can give an even stronger result, namely in terms of acyclic frameworks.

Proposition 11. $Cred_{stg}$ is Σ_2^P -hard and $Skept_{stg}$ is Π_2^P -hard, for AFs F with $dist_{acyc}(F) = 1$.

Proof. Hardness results from a reduction from the MIN-SAT problem, i.e. deciding whether a variable z is in a \subseteq -minimal model of a propositional formula $\varphi = \bigwedge_{c \in C} c$ in CNF over atoms X . For the reduction we additionally assume an arbitrary order $<$ on the clauses of φ . We build the AF $F_{\varphi, z} = (A, R)$ where $z \in X$ as follows (see Figure 3 for illustration): $A = \{\varphi, b, q\} \cup C \cup C' \cup X \cup \bar{X}$; and $R = \{(c, \varphi) \mid c \in C\} \cup \{(t, b), (b, b), (q, z)\} \cup \{(\bar{x}, x) \mid x \in X\} \cup \{(x, c) \mid x \text{ occurs in } c\} \cup \{(\bar{x}, c) \mid \neg x \text{ occurs in } c\} \cup \{(c', a) \mid c \in C, a \in (C \setminus \{c\}) \cup \{c' : c' < c\}\}$.

Now the following statements are equivalent: (1) z is in a minimal model of φ ; (2) z is credulously accepted (w.r.t. *stg*) in $F_{\varphi, z}$; (3) q is not skeptically accepted (w.r.t. *stg*) in $F_{\varphi, z}$. Finally, $dist_{acyc}(F_{\varphi, z}) = 1$ as deleting the argument b would result in an acyclic AF. \square

For each $F \in acyc$ we have $F \in \mathcal{G}$, where $\mathcal{G} \in \{wcyc, ocf, coherent, uniqpref\}$, and thus the previous result generalizes to all these classes when $dist_{\mathcal{G}}(F) = 1$.

Recall that Propositions 9 and 10 directly yield corresponding hardness results for the classes *coherent* and *stablecons*. It thus remains to consider the class *uniqpref*.

Proposition 12. $Cred_{sem}$ is Σ_2^P -hard and $Skept_{prf}, Skept_{sem}$ are Π_2^P -hard for AFs F with $dist_{uniqpref}(F) = 1$.

Proof. Consider an arbitrary AF $F = (A, R)$, and consider the modified AF $F' = (A \cup \{r, g\}, R \cup \{(r, g)\} \cup \{(g, a) \mid a \in A\})$. Then, $adm(F') = \{E \cup \{r\} \mid E \in adm(F)\}$, and thus $prf(F') = \{E \cup \{r\} \mid E \in prf(F)\}$ and $sem(F') = \{E \cup \{r\} \mid E \in sem(F)\}$. Moreover, $dist_{uniqpref}(F') = 1$ since deleting argument r would result in an AF where $\{g\}$ is the only admissible set. Hence credulous and skeptical acceptance (under the considered semantics) on arbitrary AFs reduces to AFs F with $dist_{uniqpref}(F) = 1$. \square

Table 3 summarizes our results which are all negative in the sense that full second-level complexity is reached when

Table 3: Complexity when parameterized by the distance to a sub-class \mathcal{G} (hardness holds for AFs F of $dist_{\mathcal{G}}(F) = 1$).

\mathcal{G}	Skept _{prf}	Cred _{sem}	Skept _{sem}	Cred _{stg}	Skept _{stg}
acyc	FPT	FPT	FPT	Σ_2^P -c	Π_2^P -c
wcyc	Π_2^P -c	Σ_2^P -c	Π_2^P -c	Σ_2^P -c	Π_2^P -c
stablecons	Π_2^P -c	Σ_2^P -c	Π_2^P -c	Σ_2^P -c	Π_2^P -c
ocf	Π_2^P -c	Σ_2^P -c	Π_2^P -c	Σ_2^P -c	Π_2^P -c
coherent	Π_2^P -c	Σ_2^P -c	Π_2^P -c	Σ_2^P -c	Π_2^P -c
uniqpref	Π_2^P -c	Σ_2^P -c	Π_2^P -c	Σ_2^P -c	Π_2^P -c

fragments are parameterized in a “syntactic” way; only acyc yields some positive results (due to (Ordyniak and Szeider 2011)). Here FPT (fixed-parameter tractability) means that for a fixed distance, a problem can be solved in polynomial time and the order of the polynomial time bound does not depend on the distance. However, we showed here that under the stage semantics, full complexity is obtained even for acyc for AFs F with $dist_{acyc}(F) \geq 1$.

Extension-Based Distance Measures

Here we consider different distance measures which take the number of extensions into account and thus naturally apply only to the semantical subclasses of AFs, i.e. stablecons, coherent, and uniqpref.

We start by generalizing the class stablecons for semantics sem and stg . In fact, for stable-consistent AFs we have that each semi-stable (resp. stage) extension has a range that covers the whole set of arguments. Hence a natural approach to relax this definition is to bound the number of arguments which are not in the range of the extensions.

Definition 10. For a semantics σ and $k \geq 0$, we call an AF $F = (A, R)$ k -stable-consistent under σ if for each $E \in \sigma(F)$, $|A \setminus E_R^+| \leq k$ holds. We use $stablecons_{\sigma}^k$ to denote the respective classes of AFs for given k and σ .

Theorem 1. For AFs in $stablecons_{\sigma}^k$ ($\sigma \in \{sem, stg\}$), $Cred_{\sigma}$ and $Skept_{\sigma}$ are in P^{NP} .

Proof. Consider AF $F = (A, R)$. For $S \subseteq A$, let $F_S^{stg} = (A \cap S, R \cap (S \times S) \cup \{(b, b) \mid b \mapsto (A \setminus S)\})$ and $F_S^{sem} = F_S^{stg} \cup \{(b, b) \mid (A \setminus S) \mapsto b\}$). By construction we have that (i) $E \in stb(F_S^{stg})$ iff $E \in cf(F)$ and $E_R^+ = S$; and (ii) $E \in stb(F_S^{sem})$ iff $E \in adm(F)$ and $E_R^+ = S$. The following yields a P^{NP} procedure deciding $Cred_{\sigma}$ (resp. $coSkept_{\sigma}$) for the class $stablecons_{\sigma}^k$ of AFs, based on NP oracles for $Cred_{stb}$, $coSkept_{stb}$ and $Exists_{stb}$.

1. Initialize \mathcal{S} with all sets $S \subseteq A$ such that $|A \setminus S| \leq k$.
2. While $\mathcal{S} \neq \emptyset$ take an $S \in \mathcal{S}$ with maximum cardinality:
 - (a) If $Cred_{stb}(F_S^{\sigma})$ (resp. $coSkept_{stb}(F_S^{\sigma})$) holds, then *accept*.
 - (b) If $Exists_{stb}(F_S^{\sigma})$, remove all subsets of S from \mathcal{S} .
 - (c) remove S from \mathcal{S} and continue with the loop.
3. *reject*

Since the cardinality of \mathcal{S} is polynomial for fixed k , the procedure runs in polynomial time using NP-oracles. \square

Next, we parameterize coherence.

Definition 11. An AF F is k -coherent, where $k \geq 0$, if $|prf(F) \setminus stb(F)| \leq k$. We use $coherent^k$ to denote the respective class of AFs for given k .

We do not consider the corresponding definitions for sem and stg as in Definition 10: the reason is that either a stable extension exists (hence, the AF is stable-consistent), or the parameter k , as used in Definition 11, would simply mention the number of semi-stable, resp. stage, extensions. In fact, we will consider these classes of AFs of bounded solution cardinality at the end of this section.

Here the potential exponential number of stable extensions appears to cause additional difficulties. While we are unable to provide a hardness result using standard reductions at the moment, we provide a slightly weaker result (using randomized reductions) which still suggests that parameterized coherence does not allow for more efficient algorithms.

Theorem 2. $Skept_{prf}$ for AFs in $coherent^k$ is Π_2^P -hard under randomized reductions; hardness holds even for $k = 1$.

Proof. Consider the following promise problem: given a QBF $\Phi = \forall Y \exists Z \varphi(Y, Z)$ together with the fact that $\exists^{\leq 1} Y \forall Z \neg \varphi(Y, Z)$ is true, decide whether Φ is true. By results in (Valiant and Vazirani 1986) and (Dunne 2009) (extension to QBFs), this problem is Π_2^P -hard under randomized reductions. Now one can apply the standard reduction from QBFs to $Skept_{prf}$. The sets Y with $\forall Z \neg \varphi(Y, Z)$ are in one-to-one correspondence to the preferred extensions which are not stable. Hence using the promise problem results in an AF from the class $coherent^1$. \square

The final class we discuss is $uniqpref$. The natural distance here is to consider frameworks which possess at most k preferred extensions. We will also apply the same idea for semi-stable and stage semantics.

Definition 12. Let $\sigma \in \{prf, sem, stg\}$. We denote by sol_{σ}^k the class of all AFs F such that $|\sigma(F)| \leq k$.

Theorem 3. For AFs $F \in sol_{prf}^k$, $Skept_{prf}$ is in P^{NP} .

Proof. We provide an algorithm which iteratively constructs the set $prf(F)$. The set \mathcal{E} serves this purpose and is initialized by $\mathcal{E} := \emptyset$. At each stage of the algorithm we construct a new preferred extension E as follows: start with $E = \emptyset$, and iterate over all arguments a asking an NP-oracle whether there is a complete extension C of the given framework such that $E \cup \{a\} \subseteq C$ and there is no $E' \in \mathcal{E}$ such that $C \subseteq E'$. If the oracle returns *yes*, add a to E and proceed with the next arguments. In the end, E is either a preferred extension (i.e. we can not add further arguments) or the empty set. If $E \neq \emptyset$, simply add E to the set \mathcal{E} and proceed with constructing the next preferred extension. If $E = \emptyset$, there is no non-empty complete extension that is not already contained in one of the extensions in \mathcal{E} , and hence the algorithm terminates. If $\mathcal{E} = \emptyset$, add \emptyset to \mathcal{E} .

Table 4: Complexity when the AFs belong to a sub-class \mathcal{G} .

\mathcal{G}	Skept _{prf}	Cred _{sem}	Skept _{sem}	Cred _{stg}	Skept _{stg}
stablecons ^k _σ	–	in P ^{NP}	in P ^{NP}	in P ^{NP}	in P ^{NP}
coherent ^k	Π ₂ ^P -c	in P ^{NP}	in P ^{NP}	in P ^{NP}	in P ^{NP}
sol ^k _σ	in P ^{NP}	in P ^{NP}	in P ^{NP}	in P ^{NP}	in P ^{NP}

By the assumption $F \in \text{sol}_{prf}^k$, $prf(F)$ is of polynomial size and for each extension only a linear number of steps (iterating over all arguments) is needed. Hence the overall run-time is polynomial (using an NP oracle). \square

Theorem 4. *For AFs in sol_σ^k ($\sigma \in \{sem, stg\}$), Cred_σ and Skept_σ are in P^{NP}.*

Proof. Let $F = (A, R)$. Instead of computing $\sigma(F)$, we construct here $\sigma^+(F) := \{E_R^+ \mid E \in \sigma(F)\}$. Again, below we use \mathcal{E} as the set of the currently computed elements from $\sigma^+(F)$, initialized as $\mathcal{E} := \emptyset$ and iteratively extended. Moreover, let $b_{sem} = com$ and $b_{stg} = cf$.

Within each iteration, start with $S = \emptyset$ and iterate over all arguments $a \in A$ asking an NP-oracle whether there is a set $C \in b_\sigma(F)$ such that $S \cup \{a\} \subseteq C_R^+$ and there is no $T \in \mathcal{E}$ such that $C_R^+ \subseteq T$. If the oracle returns *yes*, add a to S and proceed with the next arguments. In the end, S is either the range of a σ -extension or the empty set. If $S \neq \emptyset$, simply add S to the set \mathcal{E} and continue with constructing the next set. If $S = \emptyset$, there is no non-empty set in $b_\sigma(F)$ such that S_R^+ is not already contained in a set of \mathcal{E} , and hence the algorithm terminates. If $\mathcal{E} = \emptyset$, add \emptyset to \mathcal{E} .

By assumption, the number of σ -extensions is bounded and so is the size of $\sigma^+(F)$. To decide Cred_σ (resp. co-Skept_σ) for an argument a , iterate over all sets $S \in \sigma^+(F)$ and ask an NP oracle whether there is an $E \in b_\sigma(F)$ such that $E_R^+ = S$ and $a \in E$ (resp. $a \notin E$). \square

Our results are summarized in Table 4. Recall that the Π_2^P -hardness of Skept_{prf} for class $\text{stablecons}_\sigma^k$ was shown for randomized reductions. Furthermore, recall that coherent^k for $\sigma \in \{sem, stg\}$ reduces either to $\text{stablecons}_\sigma^k$ or to sol_σ^k . We did not consider k -stable-consistent AFs under prf since the full complexity was already reached for the class of stable-consistent AFs (recall Table 2).

SAT-Based Complexity-Sensitive Procedures

In this section we develop a generic framework instantiations of which provide complexity-sensitive decision procedures for second-level skeptical and credulous acceptance problems. The framework builds on Theorems 1, 3, and 4. The generic idea is to use the oracle calls to check candidate extensions of the input instance. At the beginning of the algorithm the candidate extensions are the admissible sets, in case of preferred and semi-stable semantics, or conflict-free sets, in case of stage semantics.

The framework exploits SAT-oracles, and hence the candidate extension checks are encoded as propositional formulas. In addition to reducing the set of remaining candi-

date solutions, we also exploit the results of the SAT-oracle calls to *strengthen* the base formula. This can be seen as a form of *no-good learning* which further prunes the possible IN/OUT-labellings (Caminada and Gabbay 2009) of the arguments on the level of the propositional encoding.

Furthermore, by applying SAT as the NP oracle, the highly optimized and efficient conflict-driven clause learning SAT solvers available today can be directly exploited in practical implementations of the framework. This reveals the implementer from the non-straightforward task of implementing the actual NP search procedure.

A SAT-Based Complexity-Sensitive Framework

In the following we develop a generic SAT-based complexity-sensitive framework for solving AF reasoning problems. The framework builds on the observation that sol_σ^k applies well to all semantics under our consideration.

In what follows we consider semantics $\sigma \in \{prf, sem, stg\}$, a reasoning mode $M \in \{\text{Skept}, \text{Cred}\}$, an AF $F = (A, R)$, and an argument $\alpha \in A$. The generic structure of our SAT-based framework is the following.

1. $q \leftarrow \begin{cases} \neg x_\alpha & \text{if } M = \text{Skept} \\ x_\alpha & \text{if } M = \text{Cred} \end{cases}$
2. $\varphi \leftarrow \varphi_{\text{BASE-SEM}(\sigma)}(F) \wedge q \wedge \text{SHORTCUTS}_\sigma(F, \alpha, M)$
3. **while** (φ is satisfiable)
 - (a) find model I of φ
 - (b) **while** (there is a model I' of $\psi_\sigma^I \wedge q$)
 $I \leftarrow I'$
 - (c) **if** ($\psi_\sigma^I(F)$ is unsatisfiable) **then accept**
 - (d) **else** $\varphi \leftarrow \varphi \wedge \gamma_\sigma^I$
4. *reject*

We will provide details for the generic concepts $\varphi_{\text{BASE-SEM}(\sigma)}(F)$, $\text{SHORTCUTS}_\sigma(F, \alpha, M)$, $\psi_\sigma^I(F)$, γ_σ^I below. Let us note at this point already, that the procedure might terminate in $\text{SHORTCUTS}_\sigma(F, \alpha, M)$ in certain cases. Then, we also terminate above procedure with the status returned by $\text{SHORTCUTS}_\sigma(F, \alpha, M)$, i.e. *accept* or *reject*.

Overall, the procedure works as follows. Depending on the reasoning mode, we test whether there is an extension containing α , or whether there is an extension not containing α . This is encoded via the query-atom q . In step 2, a formula is built to encode extensions of the base semantics, i.e. not taking maximality into account together with the query q as well as semantics-specific *shortcuts* that can be applied for pruning the search space via learning inferred information; this will be discussed later in more detail. The SHORTCUTS function allows for refining the base encoding using the inferred information it outputs. The loop in step 3 follows the ideas of Theorems 3 and 4: starting with a model that corresponds to an argument set satisfying the base semantics and q , each iteration extends the set to a larger one satisfying q until we have a maximal set satisfying q . In step 3c, the condition q is dropped for testing whether the set is maximal among all sets, i.e. whether it is an extension. If so, the algorithm accepts. If not, we learn that none of the smaller sets can be an extension (step 3d). Finally, after excluding all

sets satisfying the base semantics and q from being a valid extension, the algorithm rejects the query (step 4).

Instantiating the Framework

A key aspect of instantiating the SAT-based framework is how the AF reasoning tasks are encoded as propositional formulas over the variables $X = \{x_\alpha \mid \alpha \in A\}$ and $X_r = \{x'_\alpha \mid \alpha \in A\}$ such that the models of the formulas are in the correspondence with certain sets of arguments. The intuition behind the variables is that x_α is true iff α is in the set, and x'_α is true iff α is in the range of the set.

Encoding the Base Semantics To ensure the relation between the set and its range, we apply the formula

$$x'_a \leftrightarrow x_a \vee \bigvee_{(b,a) \in R} x_b$$

together with a propositional encoding that restrict models to correspond to particular types of AF extensions. To this end, we use the following propositional formulas $\varphi_{cf}(F)$, $\varphi_{adm}(F)$, and $\varphi_{com}(F)$, respectively, to encode the conflict-free sets, admissible sets, and resp. complete extensions, for a given AF $F = (A, R)$:

$$\begin{aligned} \varphi_{cf}(F) &= \bigwedge_{(a,b) \in R} (\neg x_a \vee \neg x_b) \\ \varphi_{adm}(F) &= \varphi_{cf}(F) \wedge \bigwedge_{(b,c) \in R} (\neg x_c \vee \bigvee_{(a,b) \in R} x_a) \\ \varphi_{com}(F) &= \varphi_{adm}(F) \wedge \bigwedge_{a \in A} (x'_a \vee \bigvee_{(b,a) \in R} \neg x'_b) \end{aligned}$$

Preferred Semantics For preferred semantics we are only interested in skeptical reasoning, since credulous reasoning for preferred semantics has lower complexity.

Given an AF $F = (A, R)$ and an argument $\alpha \in A$. For preferred semantics we can choose $\text{BASE-SEM}(prf)$ from $\{adm, com\}$. The **SHORTCUTS** function for preferred semantics is in pseudocode as follows.

SHORTCUTS_{prf}(F, α, Skept):

1. **if** ($\varphi_{\text{BASE-SEM}(prf)}(F) \wedge (\bigvee_{(\beta, \alpha) \in R} x_\beta)$ is satisfiable) **then reject**
2. **else return** $\bigwedge_{(\beta, \alpha) \in R} \neg x_\beta$

Here we simply check whether there is a counter-example for skeptical acceptance of α under $\text{BASE-SEM}(prf)$ witnessed by a set attacking α , and if not, learn that this is the case. Furthermore, in the case of preferred semantics the formulas $\psi_\sigma^I(F)$ and γ_σ^I are:

$$\begin{aligned} \psi_{prf}^I(F) &= \varphi_{\text{BASE-SEM}(prf)}(F) \wedge \bigwedge_{x \in I \cap X} x \wedge (\bigvee_{x \in X \setminus I} x) \\ \gamma_{prf}^I &= \bigvee_{x \in X \setminus I} x. \end{aligned}$$

Here, given a set I satisfying $\text{BASE-SEM}(prf)$, ψ_{prf}^I encodes the supersets $I' \supset I$ satisfying $\text{BASE-SEM}(prf)$, and γ_{prf}^I that we are no longer interested in subsets of I .

Semi-Stable and Stage Semantics Let us next consider $\sigma \in \{sem, stg\}$. For sem , we have the choice of using either adm or com as $\text{BASE-SEM}(sem)$. For stg , we use $\text{BASE-SEM}(stg) = cf$. The **SHORTCUTS** function for semi-stable and stage semantics, based on k -coherence, is in pseudocode as follows.

SHORTCUTS _{σ} (F, α, M):

1. $\psi \leftarrow \top; U \leftarrow \{S \subseteq A\}$
2. **while** ($U \neq \emptyset \wedge |A \setminus S| \leq d$)
 - (a) $S \leftarrow \text{MAXIMAL}(U)$
 - (b) $f_S \leftarrow \bigwedge_{s \in S} x'_s \wedge \bigwedge_{s \in A \setminus S} \neg x'_s$
 - (c) **if** ($\varphi_{\text{BASE-SEM}(\sigma)}(F) \wedge q \wedge f_S(X)$ is satisfiable) **then accept**
 - (d) **else if** ($\varphi_{\text{BASE-SEM}(\sigma)}(F) \wedge f_S(X)$ is satisfiable) **then**
 - i. learn $\psi \leftarrow \psi \wedge (\bigvee_{s \in A \setminus S} x'_s)$
 - ii. $U \leftarrow U \setminus \{E \mid E \subseteq S\}$
 - (e) **else**
 - i. learn $\psi \leftarrow \psi \wedge (\bigvee_{s \in S} \neg x'_s \vee \bigvee_{s \in A \setminus S} x'_s)$
 - ii. $U \leftarrow U \setminus S$
3. **if** ($U = \emptyset$) **then reject**
4. **else return** ψ

The procedure follows the ideas in the proof of Theorem 1, and thus in principle can decide the reasoning problems without any further investigations, but would be expensive for high values of k . To profit in cases where k is small from the complexity-sensitive perspective, and avoid unnecessary computational cost otherwise, we include the parameter d to bound the search-depth to sets of arguments to those whose range S satisfies $|A \setminus S| \leq d$.

The loop in step 2 iterates over sets $S \subseteq A$. In step 2a, **MAXIMAL** returns a maximal-cardinality set S from U , (starting with $S = A$). Step 2b builds the formula f_S encoding that sets of interest have range S . Step 2c tests whether there is a set with range S satisfying q under the base semantics. If so, we have found an extension satisfying q and accept. Step 2d tests whether there is an extension under base semantics which has range S . If so, we learn that we are no longer interested in sets $S' \subseteq S$. Otherwise (step 2e) we learn that we are not interested in extensions with range S . Finally, if all sets have been excluded from being the range of an extension satisfying q , the procedure rejects. In case the bound d is exceeded before this, we return the formula ψ as learnt information.

The helper formulas ψ_σ^I and γ_σ^I ($\sigma \in \{sem, stg\}$) are

$$\psi_\sigma^I(F) = \varphi_{\text{BASE-SEM}(\sigma)}(F) \wedge \bigwedge_{x' \in I \cap X_r} x' \wedge (\bigvee_{x' \in X_r \setminus I} x')$$

$$\gamma_\sigma^I = \bigvee_{x' \in X_r \setminus I} x'.$$

where $F = (A, R)$ is an AF. For a set I satisfying $\text{BASE-SEM}(\sigma)$ the formula ψ_σ^I encodes sets I' satisfying $\text{BASE-SEM}(\sigma)$ and $I_R^+ \subset (I')_R^+$, while formula γ_σ^I encodes that we are no longer interested in sets I' with $(I')_R^+ \subseteq I_R^+$.

Experiments

In order to study the practical relevance of our SAT-based complexity-sensitive framework, we implemented a prototype instantiation called **CEGARTIX**.¹ In this section

¹Available at <http://www.dbai.tuwien.ac.at/research/project/argumentation/cegartix/>.

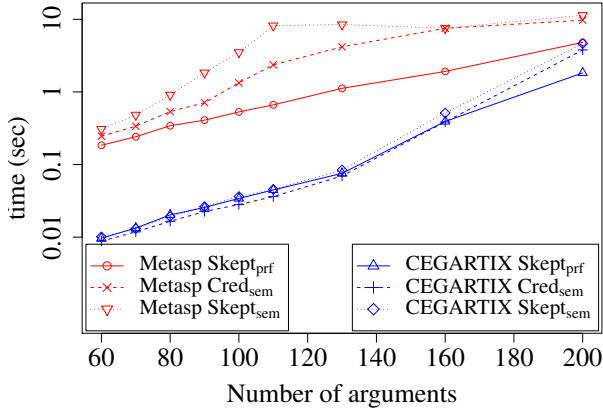


Figure 4: Average runtimes of Skept_{prf} , Cred_{sem} , Skept_{sem} on logarithmic scale.

we present preliminary experimental results on the efficiency of CEGARTIX, comparing CEGARTIX to a recently proposed, state-of-the-art argumentation reasoning system (Dvořák et al. 2011a) that exploits advances in answer set programming (ASP) via the so-called `metasp` approach². This system is a further improvement of the ASPARTIX system (Egly, Gaggl, and Woltran 2010). It turns out that, although CEGARTIX is only a prototype implementation of our SAT-based framework, it improves quite significantly over the state-of-the-art ASP-based system.

CEGARTIX employs the CDCL SAT solver Minisat (Eén and Sörensson 2004) (v2.2.0) as the SAT oracle in an incremental mode, which allows us to retain from starting search from scratch when adding new learnt information to a current satisfiable working formula. Within CEGARTIX, we employ the base semantics $\text{BASE-SEM}(\sigma) = com$ for $\sigma \in \{prf, sem\}$, and $\text{BASE-SEM}(stg) = cf$. Within $\text{SHORTCUTS}_\sigma(F, \alpha, M)$ for $\sigma \in \{sem, stg\}$ we used $d = 2$.

For comparing CEGARTIX with the ASP-based approach, we used the state-of-the-art no-good learning disjunctive ASP solver `claspD` (Drescher et al. 2008) (v1.1.1) combined with the grounder `gringo` (Gebser et al. 2011) (v3.0.3). In the experiments, we considered all five reasoning tasks we have focused on in the paper. The experiments were executed under OpenSUSE with Intel Xeon processors (2.33 GHz) and 49 GB memory. A timeout of 5 minutes was enforced on each individual run.

As benchmarks, we randomly generated 2948 AFs over 60–200 arguments, using two parameterized methods for generating the attack relation. The first generates random AFs and inserts for any pair of arguments (a, b) the attack from a to b with a given probability p . The other method generates AFs of an $n \times m$ grid-structure. We consider two different neighborhoods, one connecting arguments ver-

²Skeptical and credulous reasoning in `metasp` is done by introducing constraints in the meta answer-set programs.

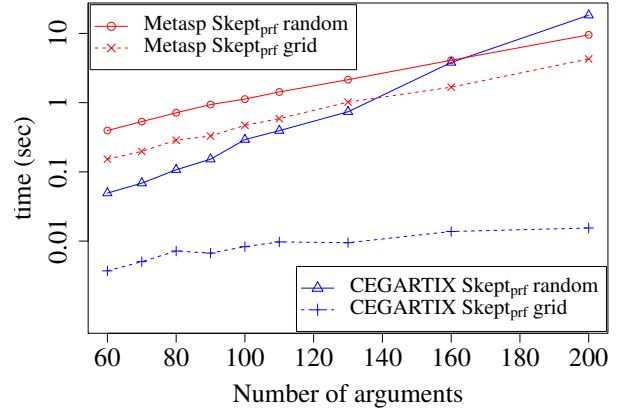


Figure 5: Average runtimes of Skept_{prf} for random and grid structured instances on logarithmic scale.

tically and horizontally and one that additionally connects the arguments diagonally. Such a connection is a mutual attack with a given probability p and in only one direction otherwise. Values used are $p \in \{0.1, 0.2, 0.3, 0.4\}$. This resulted in a total of 46200 benchmark instances for the semantics prf , sem and stg (using skeptical reasoning with the first, and additionally credulous reasoning with the latter two). For the random instances three different arguments were queried, while for the grid instances we used five arguments. For the largest instances of size 130 and more we only tested preferred and semi-stable semantics, a total of 9636 instances. For the AFs of size 200 we queried over at most three different arguments. We would like to stress that the generated AFs are by no means tailored to the fragments our approach is based on.

The average computation times, using logarithmic scale, are shown in Figure 4 for the preferred skeptical acceptance and semi-stable semantics with skeptical and credulous acceptance. The averages *do not include timed out runs*, which actually favors `metasp`; the plot excludes a total of 510 timeouts encountered with `metasp`: 23 with preferred semantics on AFs from size 130 to 200, and the rest with semi-stable semantics (119 for credulous and 368 for skeptical reasoning on AFs from size 110 to 200). In contrast only one timeout was encountered with CEGARTIX, which was for a preferred acceptance instance with 200 arguments. Overall, CEGARTIX outperforms the `metasp`-based approach: for stage semantics the results (details excluded due to page limit) are similar to Figure 4. In Figure 5 we show the results separated for random and grid instances for the preferred skeptical acceptance, the other semantics having a similar result. CEGARTIX performs better on grid structured instances, while on the random instances `metasp` has a lower average runtime on instances with 200 arguments. Note also that the number of attacks scales linearly with the number of arguments for grid instances, while it scales quadratically with the number of arguments for random instances.

To further test the scalability of CEGARTIX, we also tested its performance on larger instances with 300 and 500 arguments (without grid-structure). Of the instances of size 300 CEGARTIX solved 90 percent of the queries within the time limit, whereas 80 percent of the queries on instances of size 500 timed out.

Conclusion

In this work, we developed a novel method for solving hard problems in the area of argumentation in a “complexity-sensitive” way. From one perspective, the approach can be seen as an argumentation-customized incarnation of the counter-example guided abstraction refinement approach originating from the field of model checking. Our prototype implementation CEGARTIX employs SAT solvers as underlying inference engines. First experiments show that CEGARTIX significantly outperforms existing systems developed for hard argumentation problems (i.e. problems under the preferred, semi-stable, or stage semantics). The fundamental aspects of our approach are generic, allowing in principle to exploit as the underlying NP-oracle systems developed for other reasoning problems such as CSP or ASP, or even native argumentation systems for “easier” semantics such as the stable or complete semantics. Building necessary ground for the complexity-sensitive approach, we also presented an extensive complexity theoretic analysis, providing new results for certain fragments of argumentation frameworks, as well as distance-based complexity analysis, complementing recent results by Ordyniak and Szeider (2011).

As for future directions, a more in-depth study on which types of instances our approach is well suited for is needed; we are currently running larger test series. The promising first experimental results suggest to apply our approach also to formalisms extending the Dung-style frameworks which we focused on here. In particular, abstract dialectical frameworks (Brewka and Woltran 2010) would be an appealing target formalism since it generalizes other proposals such as bipolar (Amgoud et al. 2008) and extended argumentation frameworks (Modgil 2009). In the opposite direction, it is also on our agenda to consider further fragments of Dung-style frameworks. For preferred semantics, an interesting class are AFs having a bound number of odd cycles; the complexity of evaluating such AFs is currently open. We also aim to find complexity-sensitive approaches for this fragment.

Related Work To the best of our knowledge, our complexity-sensitive approach to developing decision procedures for second-level abstract argumentation problems is novel. The CEGAR approach has been harnessed for solving various other intrinsically hard reasoning problems (de Moura, Ruess, and Sorea 2002; Wintersteiger, Hamadi, and de Moura 2010; Janota, Grigore, and Marques-Silva 2010; Janota and Marques-Silva 2011) using SAT solvers iteratively as the underlying NP-oracle. However, we are not aware of earlier work on developing complexity-sensitive CEGAR-based procedures. Systems developed for solving second-level AF reasoning problems typically rely on monolithic encodings in other reasoning problems of sim-

ilar complexity, see e.g. (Egly and Woltran 2006) for encodings in terms of quantified Boolean formulas or (Egly, Gaggl, and Woltran 2010) for an answer-set programming based approach. Also “easier”, i.e. first-level, AF problems approaches and systems are mostly monolithic; this includes SAT-encodings (Besnard and Doutre 2004) as well as CSP-encodings (Amgoud and Devred 2011; Bistarelli, Campli, and Santini 2011). A noteworthy exception is the family of labelling-based algorithms (Verheij 2007; Modgil and Caminada 2009; Podlaszewski, Caminada, and Pigozzi 2011). Other current branches in abstract argumentation are certain preprocessing techniques (Baumann 2011; Baumann, Brewka, and Wong 2011) to divide the problem into small pieces as well as implementations of fixed-parameter tractable algorithms (Dvořák, Pichler, and Woltran 2010; Dvořák, Szeider, and Woltran 2010; Dvořák et al. 2011b). While the concept of complexity-sensitivity as used in our approach bears some resemblance to fixed-parameter tractability and parameterized complexity, there are fundamental differences. Especially, while parameterized complexity usually studies parameterizations with respect to polynomial-time tractability, the CEGAR approach in general is typically based on exploiting “NP-tractability” via iterative calls to efficient NP decision procedures. Hence our work takes further steps towards a full classification for argumentation problems of high complexity. This also distinguishes our complexity results to the ones obtained in (Ordyniak and Szeider 2011). Finally there is recent work on average-case algorithms for value-based argumentation (Nofal, Dunne, and Atkinson 2012), using an elaborate handling of the possible value orderings.

References

- Amgoud, L., and Devred, C. 2011. Argumentation frameworks as constraint satisfaction problems. In *Proc. SUM*, volume 6929 of *LNCS*, 110–122. Springer.
- Amgoud, L.; Cayrol, C.; Lagasque, M.-C.; and Livet, P. 2008. On bipolarity in argumentation frameworks. *International Journal of Intelligent Systems* 23:1–32.
- Bang-Jensen, J., and Gutin, G. 2010. *Digraphs: Theory, Algorithms and Applications*. Springer Monographs in Mathematics. Springer.
- Baroni, P., and Giacomin, M. 2009. Semantics of abstract argument systems. In Rahwan, I., and Simari, G., eds., *Argumentation in Artificial Intelligence*. Springer. 25–44.
- Baumann, R.; Brewka, G.; and Wong, R. 2011. Splitting argumentation frameworks: An empirical evaluation. In *Proc. TAAFA*.
- Baumann, R. 2011. Splitting an argumentation framework. In *Proc. LPNMR*, volume 6645 of *LNCS*, 40–53. Springer.
- Besnard, P., and Doutre, S. 2004. Checking the acceptability of a set of arguments. In *Proc. NMR*, 59–64.
- Bistarelli, S.; Campli, P.; and Santini, F. 2011. Finding partitions of arguments with Dung’s properties via SCSPs. In *Proc. SAC*, 913–919. ACM.
- Brewka, G., and Woltran, S. 2010. Abstract dialectical frameworks. In *Proc. KR*, 102–111. AAAI Press.

- Caminada, M., and Gabbay, D. 2009. A logical account of formal argumentation. *Studia Logica* 93(2–3):109–145.
- Clarke, E. M.; Grumberg, O.; Jha, S.; Lu, Y.; and Veith, H. 2003. Counterexample-guided abstraction refinement for symbolic model checking. *J. ACM* 50(5):752–794.
- Clarke, E. M.; Gupta, A.; and Strichman, O. 2004. SAT-based counterexample-guided abstraction refinement. *IEEE T-CAD* 23(7):1113–1123.
- Coste-Marquis, S.; Devred, C.; and Marquis, P. 2005. Symmetric argumentation frameworks. In *Proc. ECSQARU*, volume 3571 of *LNCS*, 317–328. Springer.
- de Moura, L.; Ruess, H.; and Sorea, M. 2002. Lazy theorem proving for bounded model checking over infinite domains. In *Proc. CADE*, volume 2392 of *LNCS*, 438–455. Springer.
- Dimopoulos, Y., and Torres, A. 1996. Graph theoretical structures in logic programs and default theories. *Theor. Comput. Sci.* 170(1-2):209–244.
- Drescher, C.; Gebser, M.; Grote, T.; Kaufmann, B.; König, A.; Ostrowski, M.; and Schaub, T. 2008. Conflict-driven disjunctive answer set solving. In *Proc. KR*, 422–432. AAAI Press.
- Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif. Intell.* 77(2):321–358.
- Dunne, P. E., and Bench-Capon, T. J. M. 2002. Coherence in finite argument systems. *Artif. Intell.* 141(1/2):187–203.
- Dunne, P. E., and Caminada, M. 2008. Computational complexity of semi-stable semantics in abstract argumentation frameworks. In *Proc. JELIA*, volume 5293 of *LNCS*, 153–165. Springer.
- Dunne, P. E. 2007. Computational properties of argument systems satisfying graph-theoretic constraints. *Artif. Intell.* 171(10-15):701–729.
- Dunne, P. E. 2009. The computational complexity of ideal semantics. *Artif. Intell.* 173(18):1559–1591.
- Dvořák, W., and Woltran, S. 2010. Complexity of semi-stable and stage semantics in argumentation frameworks. *Inf. Process. Lett.* 110(11):425–430.
- Dvořák, W.; Szeider, S.; and Woltran, S. 2010. Reasoning in argumentation frameworks of bounded clique-width. In *Proc. COMMA*, volume 216 of *FAIA*, 219–230. IOS Press.
- Dvořák, W., and Woltran, S. 2011. On the intertranslatability of argumentation semantics. *J. Artif. Intell. Res.* 41:445–475.
- Dvořák, W.; Pichler, R.; and Woltran, S. 2010. Towards fixed-parameter tractable algorithms for argumentation. In *Proc. KR*, 112–122. AAAI Press.
- Dvořák, W.; Gaggl, S. A.; Wallner, J. P.; and Woltran, S. 2011a. Making use of advances in answer-set programming for abstract argumentation systems. In *Proc. INAP*. <http://arxiv.org/abs/1108.4942>.
- Dvořák, W.; Morak, M.; Nopp, C.; and Woltran, S. 2011b. A dynamic programming reasoner for abstract argumentation. In *Proc. INAP*. <http://arxiv.org/abs/1108.4804>.
- Dvořák, W.; Pichler, R.; and Woltran, S. 2011. Towards fixed-parameter tractable algorithms for abstract argumentation. Technical Report DBAI-TR-2011-74, Technische Universität Wien. <http://www.dbai.tuwien.ac.at/research/project/argumentation/papers/dbai-tr-2011-74.pdf>.
- Eén, N., and Sörensson, N. 2004. An extensible SAT-solver. In *Proc. SAT 2003*, volume 2919 of *LNCS*, 502–518. Springer.
- Egly, U., and Woltran, S. 2006. Reasoning in argumentation frameworks using quantified boolean formulas. In *Proc. COMMA*, volume 144 of *FAIA*, 133–144. IOS Press.
- Egly, U.; Gaggl, S. A.; and Woltran, S. 2010. Answer-set programming encodings for argumentation frameworks. *Argument and Computation* 1(2):147–177.
- Gebser, M.; Kaminski, R.; König, A.; and Schaub, T. 2011. Advances in *gringo* series 3. In *Proc. LPNMR*, volume 6645 of *LNCS*, 345–351. Springer.
- Janota, M., and Marques-Silva, J. P. 2011. Abstraction-based algorithm for 2QBF. In *Proc. SAT*, volume 6695 of *LNCS*, 230–244. Springer.
- Janota, M.; Grigore, R.; and Marques-Silva, J. 2010. Counterexample guided abstraction refinement algorithm for propositional circumscription. In *Proc. JELIA*, volume 6341 of *LNCS*, 195–207. Springer.
- Lierler, Y. 2005. cmodels - SAT-based disjunctive answer set solver. In *Proc. LPNMR*, volume 3662 of *LNCS*, 447–451. Springer.
- Modgil, S., and Caminada, M. 2009. Proof theories and algorithms for abstract argumentation frameworks. In *Argumentation in Artificial Intelligence*. Springer. 105–132.
- Modgil, S. 2009. Reasoning about preferences in argumentation frameworks. *Artif. Intell.* 173(9-10):901–934.
- Nofal, S.; Dunne, P. E.; and Atkinson, K. 2012. Towards average-case algorithms for abstract argumentation. In *Proc. ICAART 2012*.
- Ordyniak, S., and Szeider, S. 2011. Augmenting tractable fragments of abstract argumentation. In *Proc. IJCAI*, 1033–1038. IJCAI/AAAI.
- Podlaszewski, M.; Caminada, M.; and Pigozzi, G. 2011. An implementation of basic argumentation components. In *Proc. AAMAS*, 1307–1308. IFAAMAS.
- Truszczynski, M. 2011. Trichotomy and dichotomy results on the complexity of reasoning with disjunctive logic programs. *Theory and Practice of Logic Programming* 11(6):881–904.
- Valiant, L. G., and Vazirani, V. V. 1986. NP is as easy as detecting unique solutions. *Theor. Comput. Sci.* 47(3):85–93.
- Verheij, B. 2007. A labeling approach to the computation of credulous acceptance in argumentation. In *Proc. IJCAI*, 623–628.
- Wintersteiger, C. M.; Hamadi, Y.; and de Moura, L. 2010. Efficiently solving quantified bit-vector formulas. In *Proc. FMCAD*, 239–246. IEEE.