

# Comparison of SAT-Based and ASP-Based Algorithms for Inconsistency Measurement

**Isabelle Kuhlmann**

*University of Hagen, Germany*

ISABELLE.KUHLMANN@FERNUNI-HAGEN.DE

**Anna Gessler**

*University of Koblenz, Germany*

ANNA.GESSLER.WORK@GMAIL.COM

**Vivien Laszlo**

**Matthias Thimm**

*University of Hagen, Germany*

VIVIEN.LASZLO@FERNUNI-HAGEN.DE

MATTHIAS.THIMM@FERNUNI-HAGEN.DE

## Abstract

We present algorithms based on satisfiability problem (SAT) solving, as well as answer set programming (ASP), for solving the problem of determining inconsistency degrees in propositional knowledge bases. We consider six different inconsistency measures whose respective decision problems lie on the first level of the polynomial hierarchy. Namely, these are the contension, forgetting-based, hitting set, max-distance, sum-distance, and hit-distance inconsistency measures. In an extensive experimental analysis, we compare the SAT-based and ASP-based approaches with each other, as well as with a set of naive baseline algorithms. Our results demonstrate that, overall, both the SAT-based and the ASP-based approaches clearly outperform the naive baseline methods in terms of runtime. The results further show that the proposed ASP-based approaches perform superior to the SAT-based ones with regard to all six inconsistency measures considered in this work. Moreover, we conduct additional experiments to explain the aforementioned results in greater detail.

## 1. Introduction

The handling of conflicting information is a substantial problem in symbolic approaches to Artificial Intelligence. For instance, different expert opinions could (partially) contradict each other, rule mining algorithms could yield conflicting rules, or data gathered from sensors could be noisy or otherwise distorted. Thus, inconsistencies can occur in virtually any area of application, and require to be detected and handled. The field of *inconsistency measurement* (Grant & Martinez, 2018; Thimm, 2019) provides an analytical perspective on this matter by facilitating the quantitative assessment of the severity of inconsistency in formal knowledge representation formalisms. Representing the degree of inconsistency as a numerical value may assist automatic reasoning mechanisms on the one hand, and human modellers who aim to identify and compare multiple alternative formalizations on the other hand. Moreover, such an analysis can be used to identify conflicts, and consequently also help to restore consistency to an inconsistent knowledge base. Inconsistency measures have been used to estimate reliability of agents in multi-agent systems (Cholvy, Perrussel, & Thevenin, 2017), to analyze inconsistencies in news reports (Hunter, 2006), to support collaborative software requirements specifications (Martinez, Arias, & Vilas, 2004), to allow for inconsistency-tolerant reasoning in probabilistic logic (Potyka & Thimm, 2017),

to handle inconsistencies in business processes (Corea, Thimm, & Delfmann, 2021; Corea, Grant, & Thimm, 2022), and to monitor and maintain quality in database settings (Decker & Misra, 2017; Bertossi, 2018a). For a general overview of the subject, see the seminal work by Grant (1978) and the edited collection by Grant & Martinez (2018).

In the literature, a multitude of different inconsistency measures have been introduced. Some approaches conceptually rely on minimal inconsistent sets or maximal consistent sets (see, e. g., (Hunter, Konieczny, et al., 2008; Jabbour, Ma, & Raddaoui, 2014; Ammoura, Raddaoui, Salhi, & Oukacha, 2015)), others rely on non-classical semantics (see, e. g., (Grant & Hunter, 2011; Ma, Qi, Xiao, Hitzler, & Lin, 2009; Knight, 2002)), and yet others utilize further properties (see, e. g., (Thimm, 2016b)). However—and despite the above list of applications and consequent need for practical working solutions—algorithmic approaches to inconsistency measurement have received only little attention so far. Ma et al. (2009) propose an algorithm that approximates the inconsistency value of a newly proposed inconsistency measure and evaluate it with respect to computational complexity. Likewise, Xiao & Ma (2012) present two new inconsistency measures and perform a complexity analysis on their decision problems, both of which are found to be on the second level of the polynomial hierarchy. They also develop and evaluate a practically feasible anytime algorithm. McAreavey et al. (2014) note that there is a lack of practical implementations for inconsistency measures that employ minimal unsatisfiable subsets, and develop and evaluate an algorithm for enumerating such subsets. Jabbour & Sais (2016) describe two algorithms for their newly introduced inconsistency measure, but do not evaluate them with regard to performance or complexity. Thimm (2016b) designs and evaluates evolutionary algorithms for two inconsistency measures. Bertossi (2018b) proposes an inconsistency measure for databases that can be computed using answer set programming and analyzes its complexity.

As the above overview shows, most algorithmic studies of inconsistency measurement focus on individual inconsistency measures and have failed to address systematic comparisons of algorithms and complexities. In response, one of the contributions of the survey by Thimm & Wallner (2019) was to determine the complexity levels of a large number of inconsistency measures. The authors concluded that the problem of inconsistency measurement is hard in general, but that there are certain measures which are more suitable candidates for practical applications due to their complexity class. Based on these findings, for a selection of three different inconsistency measures whose corresponding decision problems—i. e., deciding whether a certain value is an upper or lower bound of the inconsistency value, or whether it corresponds exactly to the inconsistency value—were found to be on the first level of the polynomial hierarchy, a set of algorithms based on Answer Set Programming (ASP) has been introduced in (Kuhlmann & Thimm, 2020) and (Kuhlmann & Thimm, 2021). Namely, these measures are the *contension* inconsistency measure (Grant & Hunter, 2011), the *forgetting-based* inconsistency measure (Besnard, 2016), and the *hitting set* inconsistency measures (Thimm, 2016b) (for their formal definitions, see Section 2). Those three measures were not only selected because of their associated complexity class, but also because they each give a different perspective on the inconsistencies in a given knowledge base, meaning that they could each provide different information that could be used to subsequently restore consistency. The contension inconsistency measure tells us which propositions are involved in a conflict, while the forgetting-based inconsistency measure

tells us which occurrences of each proposition are involved in a conflict, and can therefore also point to specific formulas. In addition, the forgetting-based measure is, to the best of our knowledge, the only inconsistency measure in the literature that is based on the notion of atom occurrences. The contension inconsistency measure serves as a representative example of inconsistency measures based on non-classical semantics (see, e. g., (Ma et al., 2009; Knight, 2002) for other examples of such measures). The hitting set inconsistency measure offers a further perspective by considering how many different interpretations are minimally needed to (individually) satisfy all formulas. This measure was originally designed for streaming-based applications, such as Linked Open Data. With regard to such applications, the development of efficient algorithms is of great interest.

In (Kuhlmann & Thimm, 2021), the three ASP-based approaches were implemented and compared to naive baseline implementations in an experimental evaluation. As anticipated, the result of the study was that the ASP-based implementations were clearly superior. Furthermore, in (Kuhlmann, Gessler, Laszlo, & Thimm, 2022), a revised version of the ASP-based approach for the contension inconsistency measure is proposed, in addition to an approach based on *satisfiability problem* (SAT) solving. The latter is widely used in applications such as hardware verification (Biere, Cimatti, Clarke, Fujita, & Zhu, 1999; Vizek, Weissenbacher, & Malik, 2015), electronic design automation (Marques-Silva & Sakallah, 2000), or cryptanalysis (Mironov & Zhang, 2006; Nejati & Ganesh, 2019). This, in addition to the fact that there exist highly optimized SAT solvers (see the results of the annual SAT competition<sup>1</sup> for an overview), makes it a natural approach for computing inconsistency measures on the first level of the polynomial hierarchy. Moreover, SAT and ASP have been compared wrt. other applications before (see, e. g., (Banbara, Gebser, Inoue, Ostrowski, Peano, Schaub, Soh, Tamura, & Weise, 2015; Eyupoglu, Fidan, Gulesen, Izci, Teber, Yilmaz, Alkan, & Erdem, 2021)). The results of the study show that both the ASP and the SAT approach clearly outperform the naive baseline method, but ultimately the ASP approach performs superior to the SAT approach.

In the work at hand, we follow up on (Kuhlmann et al., 2022) by greatly extending the scope of the considered measures and the depth of the experimental evaluation. For that, we revisit the proposed SAT-based and ASP-based approaches for the contension inconsistency measure and we present SAT-based approaches for the forgetting-based and the hitting set inconsistency measure, as well as revised versions of the corresponding ASP-based approaches introduced in (Kuhlmann & Thimm, 2021). Moreover, we propose both a SAT-based and an ASP-based approach for each of three different variations of the distance-based inconsistency measure (Grant & Hunter, 2017), which are likewise on the first level of the polynomial hierarchy. Further, the distance-based approach offers yet another perspective on the notion of an inconsistency than the three previously discussed measures—Grant & Hunter (2017) view the models of the formulas in a knowledge base as points in Euclidian space. The authors also point out how these measures can be used in applications such as the evaluation of violations of integrity constraints in databases. We consider the *max-distance*, *sum-distance*, and *hit-distance* inconsistency measures. Their formal definitions follow, along with those of the other measures, in Section 2.

---

1. <http://www.satcompetition.org/>

Hence, we examine a total of six different inconsistency measures in this work (the contension, the forgetting-based, the hitting set, the max-distance, the sum-distance, and the hit-distance inconsistency measure), and present one SAT-based and one ASP-based approach for each of them. With regard to each measure, we compare the SAT and the ASP approach to each other, as well as to a naive baseline method in an experimental evaluation. More precisely, we conduct our experiments on a total of five different data sets. Whilst two of them were used in the literature before (Kuhlmann & Thimm, 2021; Kuhlmann et al., 2022), the three other data sets are newly introduced and made publicly available. Further, we investigate the runtime composition of the SAT-based and the ASP-based approaches in more detail. In addition to the runtimes, we also record the actual inconsistency values resulting from the different inconsistency measures on the various data sets. Based on that, we conduct an experiment in which we compare different search strategies wrt. the SAT-based methods. Overall, the results of our experimental analysis confirm that both the SAT-based and the ASP-based approaches perform superior to the baseline algorithms. Nevertheless, the results also show that the ASP-based approaches altogether outperform the SAT-based ones.

In the following, we give a concise overview of the structure of this paper. We first provide the required preliminaries on inconsistency measurement in Section 2. Sections 3 and 4 comprise detailed descriptions of our SAT-based and, respectively, ASP-based algorithms. Section 5 encompasses an evaluation, in which we compare the SAT-based and the ASP-based approaches with each other, and additionally, we draw a comparison with the naive baseline methods. Section 6 concludes this work by providing a brief summary of our results, as well as an overview of possible future work.

## 2. Preliminaries

The inconsistency measures examined in this work are all designed to be applied in propositional logic knowledge bases. A knowledge base  $\mathcal{K}$  is a finite set of propositional formulas, and we define  $\mathbb{K}$  as the set of all propositional knowledge bases. Formulas are constructed by means of the usual connectives *negation* ( $\neg$ ), *disjunction* ( $\vee$ ), and *conjunction* ( $\wedge$ ). Some algorithms defined in Section 3 and 4 use the notion of *subformulas*. The set of subformulas of a formula  $\Phi$  is denoted by  $\text{sub}(\Phi)$  and is inductively defined in the following manner.

**Definition 1.** *Let  $\Phi$  be a propositional formula. If  $\Phi$  is a proposition  $X$ , the only subformula of  $\Phi$  is the proposition itself, meaning  $\text{sub}(X) = \{X\}$ . If  $\Phi$  is a negation  $\neg\Psi$ , the subformulas  $\text{sub}(\neg\Psi)$  are given by  $\{\neg\Psi\} \cup \text{sub}(\Psi)$ . The subformulas  $\text{sub}(\Psi_1 \bowtie \Psi_2)$  for a formula  $\Psi_1 \bowtie \Psi_2$ , where  $\bowtie \in \{\wedge, \vee\}$  is a binary operator, are  $\{\Psi_1 \bowtie \Psi_2\} \cup \text{sub}(\Psi_1) \cup \text{sub}(\Psi_2)$ .*

Analogously, we define the set of subformulas  $\text{sub}(\mathcal{K})$  of a knowledge base  $\mathcal{K}$ :

$$\text{sub}(\mathcal{K}) = \bigcup_{A \in \mathcal{K}} \text{sub}(A)$$

Observe that we denote arbitrary formulas and subformulas as  $\Phi$ , and formulas that are explicitly elements of a knowledge base as  $A$ . If a finer granularity is required, we may use the notation  $\Psi$  for the subformula of  $\Phi$ .

We denote the *signature* of a propositional formula or knowledge base, i. e., the (propositional) atoms appearing in it, as  $\text{At}(\cdot)$ . Formulas can evaluate to either *true* (abbreviated as  $t$ ) or *false* (abbreviated as  $f$ ). An *interpretation* is a function  $\omega : \text{At} \rightarrow \{t, f\}$  that assigns truth values to all atoms. An interpretation  $\omega$  satisfies an atom  $X \in \text{At}$  if and only if  $\omega(X) = t$ , represented by  $\omega \models X$ . For non-atomic formulas, the satisfaction relation is extended recursively according to the truth-valued functions of the connectives as usual (see (van Harmelen, Lifschitz, & Porter, 2008)).

An interpretation that satisfies a formula  $\Phi$  is also called a *model* of  $\Phi$ . For every interpretation  $\omega$ , if  $\omega$  is a model of a formula  $\Phi_1$  if and only if  $\omega$  is a model of another formula  $\Phi_2$ , then  $\Phi_1$  and  $\Phi_2$  are called *logically equivalent*. A knowledge base is satisfied by an interpretation if all of its formulas are satisfied. Throughout this paper we denote the set of models of a knowledge base  $\mathcal{K}$  by  $\text{Mod}(\mathcal{K})$  and the set of interpretations wrt. the signature  $\text{At}(\mathcal{K})$  by  $\Omega(\text{At}(\mathcal{K}))$ .

A formula  $\Phi$  is *inconsistent* if there is no interpretation that satisfies it, meaning  $\text{Mod}(\Phi) = \emptyset$ . By extension, a knowledge base  $\mathcal{K}$  is inconsistent if  $\text{Mod}(\mathcal{K}) = \emptyset$ . Let  $\mathbb{R}_{\geq 0}^{\infty}$  be the set containing all non-negative real numbers and  $\infty$ .

**Definition 2.** An inconsistency measure is a function  $\mathcal{I} : \mathbb{K} \rightarrow \mathbb{R}_{\geq 0}^{\infty}$  which satisfies  $\mathcal{I}(\mathcal{K}) = 0$  iff  $\mathcal{K}$  is consistent, for all  $\mathcal{K} \in \mathbb{K}$ .

Many different inconsistency measures and properties that characterize these measures have been proposed. To illustrate why there are many ways to define the severity of inconsistency, consider the following example.

**Example 1.** Let  $\text{At} = \{\text{sunny}, \text{cloudy}\}$ , where *sunny* represents “it is sunny” and *cloudy* represents “it is cloudy”. The following knowledge bases are inconsistent:

1.  $\mathcal{K}_1 = \{\neg\text{cloudy}, \text{cloudy}, \text{sunny}\}$
2.  $\mathcal{K}_2 = \{\neg\text{cloudy} \vee \neg\text{sunny}, \text{cloudy} \vee \text{sunny}, \text{sunny} \leftrightarrow \text{cloudy}\}$

If our aim was to decide which of these knowledge bases is more severely inconsistent, there are different aspects to consider. The first knowledge base has a more obvious and easily fixable conflict (remove either  $\neg\text{cloudy}$  or  $\text{cloudy}$ ), but it also includes a formula that is not involved in any conflict (*sunny*), meaning it contains non-zero information even if the conflict cannot be repaired. The second knowledge base has a more hidden conflict (all three formulas are required to produce the inconsistency); nevertheless, since all of its formulas are involved in the conflict, there is no “safe” formula that does not participate in any conflict, and the whole knowledge base needs to be discarded if the conflict cannot be repaired.

## 2.1 The Contension Inconsistency Measure

The contension inconsistency measure  $\mathcal{I}_c$  (Grant & Hunter, 2011) is based on Priest’s three-valued logic (Priest, 1979). The latter extends the truth values of propositional logic—*true* ( $t$ ) and *false* ( $f$ )—by a third value *both* ( $b$ ), also referred to as the inconsistent or paradoxical truth value. The semantics for this three-valued logic are specified in Table 1. Further, a *three-valued interpretation*  $\omega^3 : \text{At}(\mathcal{K}) \rightarrow \{t, f, b\}$  assigns one of the three truth values to

$x$	$y$	$\neg x$	$x \vee y$	$x \wedge y$
$t$	$t$	$f$	$t$	$t$
$t$	$b$	$f$	$t$	$b$
$t$	$f$	$f$	$t$	$f$
$b$	$t$	$b$	$t$	$b$
$b$	$b$	$b$	$b$	$b$
$b$	$f$	$b$	$b$	$f$
$f$	$t$	$t$	$t$	$f$
$f$	$b$	$t$	$b$	$f$
$f$	$f$	$t$	$f$	$f$

Table 1: Truth table for Priest’s three-valued logic. The rows that differ from classical propositional logic are marked in gray.

each atom in the signature  $\text{At}(\mathcal{K})$  of the knowledge base  $\mathcal{K}$ . Further, a *three-valued model* of a knowledge base is an interpretation  $\omega^3$  that does not assign  $f$  to any formula. Therefore, a formula is satisfiable if it evaluates to either  $t$  or  $b$  in some interpretation. The set of three-valued models of a knowledge base  $\mathcal{K}$  is denoted by  $\text{Mod}^3(\mathcal{K})$ . For a given three-valued interpretation  $\omega^3$ , we denote the set of atoms that are assigned  $b$  as the  $\text{Conflictbase}(\omega^3)$ .

**Example 2.** *The knowledge base  $\mathcal{K}_3 = \{x, \neg x\}$  is unsatisfiable in propositional logic but satisfiable in Priest’s three-valued logic. The interpretation  $\omega_1^3$  on  $x$  with  $\omega_1^3(x) = b$  is a three-valued model for  $\mathcal{K}_3$  with  $\text{Conflictbase}(\omega_1^3) = \{x\}$ .*

**Definition 3** (Grant & Hunter (2011)). *The contension inconsistency measure  $\mathcal{I}_c(\mathcal{K}) : \mathbb{K} \rightarrow \mathbb{N}_0$  is defined as*

$$\mathcal{I}_c(\mathcal{K}) = \min\{|\text{Conflictbase}(\omega^3)| \mid \omega^3 \in \text{Mod}^3(\mathcal{K})\}.$$

In other words, the contension inconsistency value is the minimum number of atoms in a knowledge base that need to be assigned  $b$  to produce a three-valued model of the knowledge base. The maximum inconsistency value is  $|\text{At}(\mathcal{K})|$ , because there can only be as many  $b$  assignments as there are atoms in the signature.

**Example 3.** *Let  $\mathcal{K}_4 = \{x \wedge y, \neg y\}$ . The interpretation  $\omega_2^3 : \{x, y\} \rightarrow \{t, b, f\}$  with  $\omega_2^3(x) = t$  and  $\omega_2^3(y) = b$  is a three-valued model for  $\mathcal{K}_4$  with  $\text{Conflictbase}(\omega_2^3) = \{y\}$ . There is no model that assigns  $b$  to 0 atoms, therefore  $\min\{|\text{Conflictbase}(\omega_2^3)|\} = |\{y\}| = 1$ , and hence  $\mathcal{I}_c(\mathcal{K}_4) = 1$ .*

## 2.2 The Forgetting-Based Inconsistency Measure

In order to define the forgetting-based inconsistency measure, we first require the definition of *atom occurrences* in formulas. Let  $\Phi$  be a propositional logic formula and  $\text{At}(\Phi)$  its signature. An atom  $X \in \text{At}(\Phi)$  can appear multiple times in  $\Phi$ . Each of these appearances is called an *occurrence* of  $X$ .

**Definition 4.** Let  $\#occ(X, \Phi)$  denote the number of times  $X$  occurs in  $\Phi$ .

**Example 4.** Consider  $\phi_1 = x \wedge y \vee \neg x \wedge z$  with  $At(\phi_1) = \{x, y, z\}$ . We can see that  $x$  appears twice in  $\phi_1$ , and  $y$  and  $z$  each appear once. Consequently,  $\#occ(x, \phi_1) = 2$  and  $\#occ(y, \phi_1) = \#occ(z, \phi_1) = 1$ .

We extend Definition 4 to knowledge bases such that  $\#occ(X, \mathcal{K})$  represents the number of times  $X$  occurs in the formulas of the knowledge base  $\mathcal{K}$ :

$$\#occ(X, \mathcal{K}) = \sum_{A \in \mathcal{K}} \#occ(X, A)$$

We use superscripts to distinguish between different occurrences of the same atom. We assume that  $X^l$  refers to the  $l$ -th occurrence of  $X$  in the considered formula  $\Phi$ , or the considered knowledge base  $\mathcal{K}$ , with  $l \in \{1, \dots, \#occ(X, \Phi)\}$ , or  $l \in \{1, \dots, \#occ(X, \mathcal{K})\}$ , respectively. We also refer to  $l$  as the *label* of  $X$ .

**Definition 5.** We define the set of all atom occurrences in a formula  $\Phi$  as

$$Occ(\Phi) = \{X_i^l \mid X_i \in At(\Phi), i \in \{1, \dots, |At(\Phi)|\}, l \in \{1, \dots, \#occ(X_i, \Phi)\}\}$$

**Example 5.** Consider again the formula  $\phi_1 = x \wedge y \vee \neg x \wedge z$  from Example 4. We assign each atom  $X$  a label  $l \in \{1, \dots, \#occ(X, \phi_1)\}$  and get  $\phi_1^l = x^1 \wedge y^1 \vee \neg x^2 \wedge z^1$ .  $Occ(\phi_1)$  contains all (labeled) atom occurrences in  $\phi_1$ , i. e.,  $Occ(\phi_1) = \{x^1, x^2, y^1, z^1\}$ .

We analogously extend the definition of  $Occ(\cdot)$  to knowledge bases.

The forgetting-based inconsistency measure  $\mathcal{I}_f$  as defined by Besnard (2016) utilizes a special operation on knowledge bases (*forgetting*) to measure the level of inconsistency. The idea behind this measure is similar to the one behind the contension inconsistency measure. However, while the contension measure refers to the atoms in the signature of a knowledge base, the forgetting-based measure refers to individual atom occurrences. An atom occurrence can be “forgotten” by replacing it with  $\top$  or  $\perp$ . We denote the replacement of the  $l$ -th occurrence of atom  $X$  in formula  $\Phi$  by  $\top$  as  $\Phi \mid X^l \xrightarrow{f} \top$  (analogously with  $\perp$  instead of  $\top$ ). Formally, the forgetting operation is defined as follows.

**Definition 6** (Besnard (2016)). Let  $\Phi$  be a propositional logic formula containing  $\#occ(X, \Phi)$  occurrences of  $X \in At(\Phi)$ . The following substitution of the  $l$ -th occurrence of  $X$  is then called forgetting  $X^l$  in  $\Phi$ :

$$\Phi \mid X^l \xrightarrow{f} \top, \perp = (\Phi \mid X^l \xrightarrow{f} \top) \vee (\Phi \mid X^l \xrightarrow{f} \perp)$$

For succinct notation we express the formula obtained by forgetting the occurrence  $X^l$  in  $\Phi$  by writing  $\mathcal{F}_{l,X}.\Phi$ .

**Example 6.** Consider the formula  $\phi_2 = \neg x^1 \wedge \neg y^1 \wedge (x^2 \vee y^2) \wedge z^1$ . If the second occurrence of  $x$  (i. e.,  $x^2$ ) is forgotten, we obtain:

$$\begin{aligned} \phi_2 \mid x^2 \xrightarrow{f} \top, \perp &= (\phi_2 \mid x^2 \xrightarrow{f} \top) \vee (\phi_2 \mid x^2 \xrightarrow{f} \perp) \\ &= (\neg x^1 \wedge \neg y^1 \wedge (\top \vee y^2) \wedge z^1) \vee (\neg x^1 \wedge \neg y^1 \wedge (\perp \vee y^2) \wedge z^1) \\ &= \neg x^1 \wedge \neg y^1 \wedge z^1 \end{aligned}$$

Observe that the original formula  $\phi_2$  in Example 6 is not satisfiable wrt. classical two-valued semantics, i. e.  $\text{Mod}(\phi_2) = \emptyset$ . However, the formula obtained by forgetting  $x^2$  in  $\phi_2$ , written as  $\mathcal{F}_{2,x}.\phi_2$ , is satisfiable. In fact, any propositional logic formula  $\Phi$  can be made consistent by forgetting a sufficient number of atom occurrences. By extension, consistency can be restored for any inconsistent knowledge base  $\mathcal{K} \in \mathbb{K}$  if a sufficient number of atom occurrences are forgotten.

**Definition 7** (Besnard (2016)). *Let  $\wedge_{\mathcal{K}}$  be the conjunction of all formulas  $A \in \mathcal{K}$ . Based on the forgetting operation, the measure  $\mathcal{I}_f : \mathbb{K} \rightarrow \mathbb{N}_0$  can be defined as*

$$\mathcal{I}_f(\mathcal{K}) = \min\{n \mid \mathcal{F}_{l_1, X_1}.\mathcal{F}_{l_2, X_2} \dots \mathcal{F}_{l_n, X_n}.\wedge_{\mathcal{K}} \not\models \perp\}$$

with  $X_i \in \text{At}(\mathcal{K})$  and  $l_i$  being the respective labels with  $i \in \{1, \dots, n\}$ .

In other words, the inconsistency value  $\mathcal{I}_f(\mathcal{K})$  is the minimum number of atom occurrences which have to be forgotten in order to restore consistency in  $\mathcal{K}$ . Forgetting every atom occurrence in a knowledge base yields a trivially consistent knowledge base. It is thus guaranteed that the inconsistency value  $\mathcal{I}_f(\mathcal{K})$  will not exceed  $|\text{Occ}(\mathcal{K})|$ .

**Example 7.** *Consider the knowledge base  $\mathcal{K}_5 = \{x \wedge y, x \vee y, z, \neg x\}$ . After assigning labels to all atom occurrences, we get  $\mathcal{K}_5^l = \{x^1 \wedge y^1, x^2 \vee y^2, z^1, \neg x^3\}$ . Clearly, there is a conflict between  $x \wedge y$  and  $\neg x$ . However, if we forget either the first or the third occurrence of  $x$  (i. e.,  $x^1$  or  $x^3$ ),  $\mathcal{K}_5$  becomes consistent. Since there is no possibility of rendering  $\mathcal{K}_5$  consistent by forgetting less than one atom occurrence,  $\mathcal{I}_f(\mathcal{K}_5) = 1$ .*

### 2.3 The Hitting Set Inconsistency Measure

In general, a hitting set is a set that contains at least one element from each set of a given set of sets. In this work, we use a context-specific form of a hitting set.

**Definition 8.** *A set  $H \subseteq \Omega(\text{At}(\mathcal{K}))$  is a hitting set of a knowledge base  $\mathcal{K}$  if for every  $A \in \mathcal{K}$  there exists an  $\omega \in H$  with  $\omega \models A$ .*

We use the above notion of a hitting set to define the hitting set inconsistency measure (Thimm, 2016b). Let  $\mathbb{N}_0^\infty$  denote  $\mathbb{N}_0 \cup \{\infty\}$ .

**Definition 9** (Thimm (2016b)). *The hitting set inconsistency measure  $\mathcal{I}_h(\mathcal{K}) : \mathbb{K} \rightarrow \mathbb{N}_0^\infty$  is defined as*

$$\mathcal{I}_h(\mathcal{K}) = \min\{|H| \mid H \text{ is a hitting set of } \mathcal{K}\} - 1$$

with  $\min\{\emptyset\} = \infty \forall \mathcal{K} \in \mathbb{K} \setminus \{\emptyset\}$ , and with  $\mathcal{I}_h(\emptyset) = 0$ .

Intuitively, the hitting set inconsistency measure is based on the idea of searching for the minimum number of interpretations that are required to satisfy all formulas in a given knowledge base  $\mathcal{K}$ , subtracted by 1. If  $\mathcal{K}$  is consistent, we only need one interpretation (i. e., a model of  $\mathcal{K}$ ). Thus, if we subtract 1, we get 0—the desired inconsistency value for consistent knowledge bases. If there is a conflict in  $\mathcal{K}$ , at least 2 interpretations are needed to satisfy all formulas, i. e., the inconsistency value is  $> 0$ . The more interpretations we need to satisfy all formulas, the more severe the conflict is considered to be. Note that

$\mathcal{I}_h(\mathcal{K}) = \infty$  if and only if one of the formulas in  $\mathcal{K}$  is contradictory. In this case there exists no hitting set as there exists no interpretation which can satisfy such a formula, and consequently, no number of interpretations can satisfy all formulas in  $\mathcal{K}$ . If  $\mathcal{I}_h(\mathcal{K}) \neq \infty$  then the maximum inconsistency value is  $|\mathcal{K}| - 1$ . This value occurs when there is a model for every formula in  $\mathcal{K}$ , but no interpretation is a model for more than one formula.

**Example 8.** Consider again  $\mathcal{K}_4 = \{x \wedge y, \neg y\}$ . Let  $\omega_1$  and  $\omega_2$  be interpretations with  $\omega_1(x) = \omega_1(y) = t$ , and  $\omega_2(x) = t$  and  $\omega_2(y) = f$ . Clearly,  $\omega_1$  is a model for  $x \wedge y$ , and  $\omega_2$  is a model for  $\neg y$ . There is no single interpretation satisfying both formulas, making  $\{\omega_1, \omega_2\}$  a minimal hitting set of  $\mathcal{K}_4$ . Therefore,  $\mathcal{I}_h(\mathcal{K}_4) = |\{\omega_1, \omega_2\}| - 1 = 2 - 1 = 1$ .

Consider another knowledge base  $\mathcal{K}_6 = \{x \wedge \neg x, y, z\}$ . The formula  $x \wedge \neg x$  is contradictory, consequently there is no hitting set of  $\mathcal{K}_6$ , and  $\mathcal{I}_h(\mathcal{K}_6) = \infty$ .

## 2.4 Distance-Based Inconsistency Measures

Grant & Hunter (2013) proposed several inconsistency measures based on calculating distances between interpretations. Different distance measures could be used, but we assume that the distance measure is always the Dalal distance  $d(\omega, \omega')$  between two interpretations  $\omega$  and  $\omega'$ . The Dalal distance, also known as Hamming distance, measures distances based on the number of differing digits. When talking about interpretations, we define the Dalal distance as the number of differing truth value assignments for atoms  $X \in \text{At}$  between two interpretations  $\omega, \omega'$ .

**Definition 10.** The Dalal distance  $d_{dalal}$  between two interpretations  $\omega, \omega'$  is defined as

$$d_{dalal}(\omega, \omega') = |\{X \in \text{At} \mid \omega(X) \neq \omega'(X)\}|$$

Additionally, we define the Dalal distance between a set of interpretations  $I \subseteq \Omega(\text{At})$  and a single interpretation  $\omega$  as the smallest distance between some interpretation in  $I$  and the interpretation  $\omega$ .

**Definition 11.** The Dalal distance  $d_{dalal}$  between a set of interpretations  $I$  and an individual interpretation  $\omega$  is defined as

$$d_{dalal}(I, \omega) = \min_{\omega' \in I} d(\omega', \omega), \text{ with } d_{dalal}(\emptyset, \omega) = \infty$$

We now define three distance-based inconsistency measures. The general idea behind all three measures is to find an interpretation that has an “optimal” distance to the models of the individual formulas in a given knowledge base  $\mathcal{K}$ . The definition of “optimal” differs for each measure.

The *max-distance* measure calculates the interpretation with the smallest maximum distance to the models of each formula in  $\mathcal{K}$ . The inconsistency value is equal to the aforementioned smallest maximum distance. In case that  $\mathcal{K}$  contains a contradictory formula  $A_\perp$ , i. e.,  $\text{Mod}(A_\perp) = \emptyset$ , the maximum distance will always be  $\infty$ , due to  $d_{dalal}(\emptyset, \omega) = \infty$  for any  $\omega \in \Omega(\text{At})$ . Hence, the minimal maximum distance is  $\infty$ , and  $\mathcal{I}_d^{\max}(\mathcal{K}) = \infty$ .

**Definition 12** (Grant & Hunter (2013)). The max-distance inconsistency measure  $\mathcal{I}_d^{\max}(\mathcal{K}) : \mathbb{K} \rightarrow \mathbb{N}_0^\infty$  is defined as

$$\mathcal{I}_d^{\max}(\mathcal{K}) = \min \left\{ \max_{A \in \mathcal{K}} d_{dalal}(\text{Mod}(A), \omega) \mid \omega \in \Omega(\text{At}(\mathcal{K})) \right\}$$

The *sum-distance* measure calculates an interpretation  $\omega$  such that the sum of  $\omega$ 's distances to the models of each formula in  $\mathcal{K}$  is minimal. The inconsistency value is equal to the aforementioned smallest sum of distances. Note that in the case that  $\mathcal{K}$  contains a contradictory formula  $A_{\perp}$ , each sum will add up to  $\infty$ , and thus  $\mathcal{I}_d^{\max}(\mathcal{K}) = \infty$ .

**Definition 13** (Grant & Hunter (2013)). *The sum-distance inconsistency measure  $\mathcal{I}_d^{\Sigma}(\mathcal{K}) : \mathbb{K} \rightarrow \mathbb{N}_0^{\infty}$  is defined as*

$$\mathcal{I}_d^{\Sigma}(\mathcal{K}) = \min\{\sum_{A \in \mathcal{K}} d_{dalal}(\text{Mod}(A), \omega) \mid \omega \in \Omega(\text{At}(\mathcal{K}))\}$$

The *hit-distance* measure calculates an interpretation such that the number of distances greater than 0 to the models of each formula is minimal. A different, simpler characterization of this measure (Grant & Hunter, 2017), shows that  $\mathcal{I}_d^{\text{hit}}$  is equal to the minimum number of formulas that need to be removed from a knowledge base  $\mathcal{K}$  in order to make it consistent.

**Definition 14** (Grant & Hunter (2017)). *The hit-distance inconsistency measure  $\mathcal{I}_d^{\text{hit}}(\mathcal{K}) : \mathbb{K} \rightarrow \mathbb{N}_0$  is defined as*

$$\mathcal{I}_d^{\text{hit}}(\mathcal{K}) = \min\{|\{A \in \mathcal{K} \mid d_{dalal}(\text{Mod}(A), \omega) > 0\}| \mid \omega \in \Omega(\text{At}(\mathcal{K}))\}$$

$\mathcal{I}_d^{\max}(\mathcal{K})$  and  $\mathcal{I}_d^{\Sigma}(\mathcal{K})$  take the value  $\infty$  if and only if one of the formulas in  $\mathcal{K}$  is inconsistent individually. For the non-infinity case, the maximum value of  $\mathcal{I}_d^{\max}(\mathcal{K})$  is  $|\text{At}(\mathcal{K})|$ , because the distance is based on the number of differing atom assignments which cannot exceed the number of atoms. For the non-infinity case of  $\mathcal{I}_d^{\Sigma}(\mathcal{K})$ , the maximum value is  $|\mathcal{K}| \cdot |\text{At}(\mathcal{K})|$ , because in the worst case, each distance  $d_{dalal}(\text{Mod}(A), \omega)$  for all formulas  $A \in \mathcal{K}$  is maximal.  $\mathcal{I}_d^{\text{hit}}$  cannot take the value  $\infty$ , because the value of  $\mathcal{I}_d^{\text{hit}}$  is the number of distances greater than 0 rather than a distance itself. The maximum value of  $\mathcal{I}_d^{\text{hit}}$  is therefore  $|\mathcal{K}|$ .

**Example 9.** *Consider again  $\mathcal{K}_4 = \{x \wedge y, \neg y\}$ . For simplified readability, we name the formulas:  $x \wedge y =: \alpha_1$  and  $\neg y =: \alpha_2$ . The possible interpretations for  $\text{At}(\mathcal{K}_4)$  are  $\Omega(\text{At}(\mathcal{K}_4)) = \{\omega_0, \omega_x, \omega_y, \omega_{xy}\}$  with:*

$$\begin{aligned} \omega_0(x) &= \omega_0(y) = f \\ \omega_x(x) &= t, \omega_x(y) = f \\ \omega_y(x) &= f, \omega_y(y) = t \\ \omega_{xy}(x) &= \omega_{xy}(y) = t \end{aligned}$$

*The models for the formulas in  $\mathcal{K}_4$  are  $\text{Mod}(\alpha_1) = \{\omega_{xy}\}$  and  $\text{Mod}(\alpha_2) = \{\omega_0, \omega_x\}$ . We now calculate  $d_{dalal}$  for all formulas in  $\mathcal{K}_4$  and all interpretations in  $\Omega(\text{At}(\mathcal{K}_4))$ :*

$$\begin{array}{ll} d_{dalal}(\text{Mod}(\alpha_1), \omega_0) = 2 & d_{dalal}(\text{Mod}(\alpha_2), \omega_0) = 0 \\ d_{dalal}(\text{Mod}(\alpha_1), \omega_x) = 1 & d_{dalal}(\text{Mod}(\alpha_2), \omega_x) = 0 \\ d_{dalal}(\text{Mod}(\alpha_1), \omega_y) = 1 & d_{dalal}(\text{Mod}(\alpha_2), \omega_y) = 1 \\ d_{dalal}(\text{Mod}(\alpha_1), \omega_{xy}) = 0 & d_{dalal}(\text{Mod}(\alpha_2), \omega_{xy}) = 1 \end{array}$$

The max-distance measure  $\mathcal{I}_d^{\max}$  looks at the maximum distances per interpretation, i. e.,  $\max_{\alpha \in \mathcal{K}} d_{dalal}(\text{Mod}(\alpha), \omega)$ . Here, the set of maximal distances consists of:

$$\begin{aligned} d_{dalal}(\text{Mod}(\alpha_1), \omega_0) &= 2 \\ d_{dalal}(\text{Mod}(\alpha_1), \omega_x) &= 1 \\ d_{dalal}(\text{Mod}(\alpha_1), \omega_y) &= 1 \\ d_{dalal}(\text{Mod}(\alpha_2), \omega_{xy}) &= 1 \end{aligned}$$

The value  $\mathcal{I}_d^{\max}(\mathcal{K}_4)$  is the minimum of those maxima, meaning

$$d_{dalal}(\text{Mod}(\alpha_1), \omega_x) = 1,$$

or one of the other maxima with value 1. Therefore,  $\mathcal{I}_d^{\max}(\mathcal{K}_4) = 1$ .

For the sum-distance inconsistency measure  $\mathcal{I}_d^{\Sigma}$ , we sum up the distances for each interpretation:

$$\begin{aligned} \sum_{\alpha \in \mathcal{K}_4} d_{dalal}(\text{Mod}(\alpha), \omega_0) &= 2 \\ \sum_{\alpha \in \mathcal{K}_4} d_{dalal}(\text{Mod}(\alpha), \omega_x) &= 1 \\ \sum_{\alpha \in \mathcal{K}_4} d_{dalal}(\text{Mod}(\alpha), \omega_y) &= 2 \\ \sum_{\alpha \in \mathcal{K}_4} d_{dalal}(\text{Mod}(\alpha), \omega_{xy}) &= 1. \end{aligned}$$

The value of  $\mathcal{I}_d^{\Sigma}(\mathcal{K}_4)$  is the minimum of the above sums. Consequently,  $\mathcal{I}_d^{\Sigma}(\mathcal{K}_4) = 1$ .

For the hit-distance inconsistency measure  $\mathcal{I}_d^{\text{hit}}$ , we count the number of distances greater than 0 for each interpretation:

$$\begin{aligned} |\{\alpha \in \mathcal{K}_4 \mid d_{dalal}(\text{Mod}(\alpha), \omega_0) > 0\}| &= 1 \\ |\{\alpha \in \mathcal{K}_4 \mid d_{dalal}(\text{Mod}(\alpha), \omega_x) > 0\}| &= 1 \\ |\{\alpha \in \mathcal{K}_4 \mid d_{dalal}(\text{Mod}(\alpha), \omega_y) > 0\}| &= 2 \\ |\{\alpha \in \mathcal{K}_4 \mid d_{dalal}(\text{Mod}(\alpha), \omega_{xy}) > 0\}| &= 1 \end{aligned}$$

The inconsistency value is the minimum of those counts:  $\mathcal{I}_d^{\text{hit}}(\mathcal{K}_4) = 1$ .

### 3. SAT-Based Algorithms for Selected Inconsistency Measures

By now we have established a fundamental overview of inconsistency measurement in general, and we defined the six particular inconsistency measures considered in this work. Building on that, in the section at hand, we describe how to use SAT encodings to determine inconsistency values.

#### 3.1 Satisfiability Solving

A major problem in the realm of propositional logic is the *Boolean satisfiability problem*, which is one of the most-studied problems of computer science, and which is NP-complete (Biere, Heule, Maaren, & Walsh, 2009). The input of the Boolean satisfiability problem is a formula in *conjunctive normal form* (CNF), which is a conjunction of clauses. Note that we do not restrict our input knowledge bases for inconsistency measurement to CNF,

however, every propositional formula can be converted into an equivalent formula in CNF (see (van Harmelen et al., 2008) for a proof). Naively, this conversion can be done using Boolean transformation rules, but the resulting formulas are sometimes exponentially larger than the original formula. Using the Tseitin method (Tseitin, 1968), every formula can be converted into an equisatisfiable CNF formula with only a linear increase in size.

**Definition 15.** *The Boolean satisfiability problem (SAT) is the problem of deciding if there exists an interpretation that satisfies a given propositional formula.*

- **Input:** formula  $\Phi$  in CNF
- **Output:** true iff  $\text{Mod}(\Phi) \neq \emptyset$ , false otherwise

A SAT solver is a program that solves SAT for a given formula. There exist many high-performance SAT solvers (for an overview, see the results of the recurring SAT competition<sup>2</sup>). A consequence of the NP-completeness of SAT is that SAT solvers can also be used to solve other NP-complete problems if we transform them into SAT problems.

In order to use SAT solving for the problem at hand, we need a way to encode the concept that some value is an upper limit. *Cardinality constraints* represent that at least, at most, or exactly some number  $k$  out of a set of propositional atoms are allowed to be true. Using the formal definition of Abio et al. (2013), we define that a *cardinality constraint* is a constraint of the form

$$X_1 + \dots + X_n \bowtie k$$

where  $X_1, \dots, X_n$  are propositional atoms,  $\bowtie \in \{<, \leq, =, \geq, >\}$  is an operator, and  $k$  is a natural number. The meaning of the  $+$  operator in this context is that for every true atom the number 1 is added and for every false atom the number 0 is added, thereby counting the number of true atoms. Constraints of the form  $X_1 + \dots + X_n \leq k$  are informally referred to as *at-most- $k$  constraints*. Such a constraint is true if and only if  $\leq k$  atoms out of the set  $\{X_1, \dots, X_n\}$  are true. The direct approach to encode cardinality constraints is to enumerate all possible atom assignments that satisfy the constraints. For small knowledge bases the resulting encodings can be compact; however, in general, there are  $\binom{n}{k+1}$  subsets of size  $k+1$ , meaning we generate  $\binom{n}{k+1}$  clauses. It is inefficient for larger inputs because the size of the constraint grows exponentially with  $n$ . Several more efficient methods to encode cardinality constraints more efficiently have been developed. In the implementations used for our experimental evaluation (Section 5), we use the *sequential counter encoding* method (Sinz, 2005) which only generates  $n \cdot k$  clauses.

In the following sections, we use the notation  $\text{at\_most}(k, \mathbb{X})$  where  $\mathbb{X}$  is a set of atoms  $\{X_1, \dots, X_n\}$  to denote the constraint  $X_1 + \dots + X_n \leq k$ .

### 3.2 Scheme for SAT-Based Algorithms for Inconsistency Measures

This section proposes a binary search algorithm for the computation of the inconsistency value of a given knowledge base  $\mathcal{K}$  wrt. one of the inconsistency measures considered in this work, i. e., the computation of  $\mathcal{I}(\mathcal{K})$  with  $\mathcal{I} \in \{\mathcal{I}_c, \mathcal{I}_f, \mathcal{I}_h, \mathcal{I}_d^{\max}, \mathcal{I}_d^{\Sigma}, \mathcal{I}_d^{\text{hit}}\}$ . Our approach follows similar procedures that have been proposed in the literature (see, e.g., (Giunchiglia

---

2. <http://www.satcompetition.org/>

$\mathcal{I}$	Maximum non- $\infty$ value	$\infty$ possible?
$\mathcal{I}_c$	$ \text{At}(\mathcal{K}) $	No
$\mathcal{I}_f$	$ \text{Occ}(\mathcal{K}) $	No
$\mathcal{I}_h$	$ \mathcal{K}  - 1$	Yes
$\mathcal{I}_d^{\max}$	$ \text{At}(\mathcal{K}) $	Yes
$\mathcal{I}_d^{\Sigma}$	$ \text{At}(\mathcal{K})  \cdot  \mathcal{K} $	Yes
$\mathcal{I}_d^{\text{hit}}$	$ \mathcal{K} $	No

Table 2: Ranges of possible values of each inconsistency measure considered in this work. The minimum value is 0 in all cases.

& Maratea, 2006)). So the objective of our algorithm is to solve the function problem  $\text{VALUE}_{\mathcal{I}}$  as formalized in (Thimm & Wallner, 2019).

$\text{VALUE}_{\mathcal{I}}$     Input:     $\mathcal{K} \in \mathbb{K}$   
                   Output:    value of  $\mathcal{I}(\mathcal{K})$

All inconsistency measures considered in this paper have a clearly defined range as established in their descriptions in Section 2 (see also Table 2). This clearly defined search space prompts the use of a binary search procedure. The range is searched by the algorithm and for each iteration of the search, a call to a SAT-solver is made, to decide which half of the range needs to be searched in the following iteration. The problem that is solved at each iteration of the search procedure is the decision problem  $\text{UPPER}_{\mathcal{I}}$  (Thimm & Wallner, 2019), meaning the problem of deciding whether a given value  $u$  is an upper bound of the inconsistency value of a given knowledge base.

$\text{UPPER}_{\mathcal{I}}$     Input:     $\mathcal{K} \in \mathbb{K},$   
                                    $u \in \mathbb{N}_0$   
                   Output:    *true* iff  $\mathcal{I}(\mathcal{K}) \leq u$

For all inconsistency measures presented in Section 2, the problem  $\text{UPPER}_{\mathcal{I}}$  is NP-complete as shown in (Thimm & Wallner, 2019), and is thus reducible to SAT. Therefore, the idea is to find SAT encodings  $S : \mathbb{K} \times \mathbb{N}_0 \rightarrow \mathbb{K}$  for  $\text{UPPER}_{\mathcal{I}}$  for all inconsistency measures, that satisfy the requirement that  $(\mathcal{K}, u)$  is a positive instance of  $\text{UPPER}_{\mathcal{I}}$  if and only if  $S(\mathcal{K}, u)$  is a positive instance of SAT. At each search step, we compute the encoding for the current potential upper bound  $u$  and make a call to a SAT solver to check if  $u$  is in fact an upper limit of the inconsistency value, until the exact value is found. If no value is found,  $\infty$  is assumed to be the inconsistency value for the measures  $\mathcal{I}_h$ ,  $\mathcal{I}_d^{\max}$ , and  $\mathcal{I}_d^{\Sigma}$ . With regard to the other measures ( $\mathcal{I}_c$ ,  $\mathcal{I}_f$ , and  $\mathcal{I}_d^{\text{hit}}$ ), the binary search procedure is guaranteed to return a finite value after  $\log_2(n + 1)$  steps, with  $n$  being the maximally possible finite value.

Note that the iterative SAT-based approach described in this section is essentially a naive MaxSAT approach. Although the literature offers many different dedicated MaxSAT solvers (see, e. g., (Berg, Jarvisalo, Martins, Niskanen, & Paxian, 2024) for a recent overview), which are based on several different optimization algorithms, we decided to rely on classical SAT solvers for this work. In particular, the focus of this work is on the presented encodings;

a detailed investigation of the various optimization algorithms would exceed the scope of this paper (note, however, that we briefly touch upon the topic of optimization wrt. SAT in Section 5.3.4 and wrt. ASP in Section 5.3.5).

### 3.3 The Contension Inconsistency Measure

Beginning with this section, we describe how to construct the SAT encodings that are used in the binary search procedure described in the previous section for each specific inconsistency measure. Note that the encoding for the contension inconsistency measure covered in this section was already proposed in (Kuhlmann et al., 2022). However, for the sake of completeness, we describe it once more.

Let  $\mathcal{K}$  be a propositional knowledge base and  $u$  an integer representing a possible upper bound for the contension inconsistency value  $\mathcal{I}_c(\mathcal{K})$ . We present a SAT encoding for  $\text{UPPER}_{\mathcal{I}_c}$ , denoted  $S_c(\mathcal{K}, u)$ , which is defined by the components (SC1)–(SC17).

Recall that  $\mathcal{I}_c(\mathcal{K})$  is defined by the minimal number of atoms in  $\text{At}(\mathcal{K})$  that need to be set to  $b$  in order for  $\mathcal{K}$  to become consistent in Priest’s three-valued logic. To encode the three-valued logic in propositional logic, we use additional variables. First of all, for every atom  $X$  in the original signature  $\text{At}(\mathcal{K})$ , we use three new atoms  $X_t, X_b, X_f$  (SC1) to represent the three truth values  $t, b, f$ . We need to assure that out of each triple of new atoms that represent the three possible truth values of the original atom  $X$ , exactly one is true. We represent this by adding the following formula for every  $X \in \text{At}(\mathcal{K})$  to the encoding  $S_c(\mathcal{K}, u)$ :

$$(X_t \vee X_f \vee X_b) \wedge (\neg X_t \vee \neg X_f) \wedge (\neg X_t \vee \neg X_b) \wedge (\neg X_b \vee \neg X_f) \quad (\text{SC3})$$

To model three-valued satisfiability, we recursively represent the evaluation of formulas in three-valued logic (see the semantics described in Section 2.1). To achieve this, three variables  $v_\Phi^t, v_\Phi^f, v_\Phi^b$  (SC2) are used for every subformula  $\Phi$  in every formula  $A \in \mathcal{K}$  to represent each of the three possible valuations of  $\Phi$ . For each of those valuation atoms we add an equivalence that defines the evaluations based on the operator of the subformula. To represent all possible formulas, we need to encode the operators  $\wedge, \vee$ , and  $\neg$ .

We encode conjunctions  $\Phi_c = \Psi_{c,1} \wedge \Psi_{c,2}$  by adding the following formulas to  $S_c(\mathcal{K}, u)$ :

$$v_{\Phi_c}^t \leftrightarrow v_{\Psi_{c,1}}^t \wedge v_{\Psi_{c,2}}^t \quad (\text{SC4})$$

$$v_{\Phi_c}^f \leftrightarrow v_{\Psi_{c,1}}^f \vee v_{\Psi_{c,2}}^f \quad (\text{SC5})$$

$$v_{\Phi_c}^b \leftrightarrow (\neg v_{\Psi_{c,1}}^t \vee \neg v_{\Psi_{c,2}}^t) \wedge \neg v_{\Psi_{c,1}}^f \wedge \neg v_{\Psi_{c,2}}^f \quad (\text{SC6})$$

Analogously, disjunctions  $\Phi_d = \Psi_{d,1} \vee \Psi_{d,2}$  are encoded as follows:

$$v_{\Phi_d}^t \leftrightarrow v_{\Psi_{d,1}}^t \vee v_{\Psi_{d,2}}^t \quad (\text{SC7})$$

$$v_{\Phi_d}^f \leftrightarrow v_{\Psi_{d,1}}^f \wedge v_{\Psi_{d,2}}^f \quad (\text{SC8})$$

$$v_{\Phi_d}^b \leftrightarrow (\neg v_{\Psi_{d,1}}^f \vee \neg v_{\Psi_{d,2}}^f) \wedge \neg v_{\Psi_{d,1}}^t \wedge \neg v_{\Psi_{d,2}}^t \quad (\text{SC9})$$

To represent a negated formula  $\Phi_n = \neg \Psi_n$ , we create new variables representing the three evaluations of  $\Phi_n$ , meaning  $v_{\Phi_n}^t, v_{\Phi_n}^f$  and  $v_{\Phi_n}^b$ , and encode the evaluation of those cases by

adding the following formulas to  $S_c(\mathcal{K}, u)$ :

$$v_{\Phi_n}^t \leftrightarrow v_{\Psi_n}^f \quad (\text{SC10})$$

$$v_{\Phi_n}^f \leftrightarrow v_{\Psi_n}^t \quad (\text{SC11})$$

$$v_{\Phi_n}^b \leftrightarrow v_{\Psi_n}^b \quad (\text{SC12})$$

Thus, (SC10) encodes that the formula  $\Phi_n = \neg\Psi_n$  evaluates to  $t$  if  $\Psi_n$  evaluates to  $f$ . In the same fashion,  $\Phi_n$  evaluates to  $f$  if  $\Psi_n$  is  $t$  (SC11), and  $\Phi_n$  evaluates to  $b$  if  $\Psi_n$  is also  $b$  (SC12). Further, we add the following formulas for each subformula  $\Phi_a$  which represents an individual atom  $X$ :

$$v_{\Phi_a}^t \leftrightarrow X_t \quad (\text{SC13})$$

$$v_{\Phi_a}^f \leftrightarrow X_f \quad (\text{SC14})$$

$$v_{\Phi_a}^b \leftrightarrow X_b \quad (\text{SC15})$$

We also add a formula to  $S_c(\mathcal{K}, u)$  that represents when a formula  $A \in \mathcal{K}$  becomes true. This is the case when the subformula that contains the whole formula evaluates to  $t$  or  $b$ :

$$v_A^t \vee v_A^b \quad (\text{SC16})$$

Finally, we add a cardinality constraint representing that at most  $u$  of the  $b$ -atoms are allowed to be true. Let  $\text{At}_b$  be the set of  $b$ -atoms of  $\mathcal{K}$ , i. e.,  $\text{At}_b^{\mathcal{K}} = \{X_b \mid X \in \text{At}(\mathcal{K})\}$ . We add the following cardinality constraint to  $S_c(\mathcal{K}, u)$ :

$$\text{at\_most}(u, \text{At}_b^{\mathcal{K}}) \quad (\text{SC17})$$

$S_c(\mathcal{K}, u)$  is finally comprised of the signature and the formulas presented in (SC1)–(SC17) (see Encoding 1 in Appendix A.1.1 for an overview). The following result establishes that this encoding faithfully implements  $\mathcal{I}_c$ . The proof of the theorem below is provided in Appendix A.1.1.

**Theorem 1.** *For a given value  $u$ , the encoding  $S_c(\mathcal{K}, u)$  is satisfiable if and only if  $\mathcal{I}_c(\mathcal{K}) \leq u$ .*

**Example 10.** *We illustrate the construction of  $S_c(\mathcal{K}, u)$  for the following knowledge base:*

$$\mathcal{K}_7 = \left\{ \underbrace{\left\{ \begin{array}{cc} \phi_{1,1} & \phi_{1,2} \\ x & \wedge & y \end{array} \right\}}_{\alpha_1}, \underbrace{\left\{ \begin{array}{cc} \phi_{2,1} & \phi_{2,2} \\ x & \vee & y \end{array} \right\}}_{\alpha_2}, \underbrace{\left\{ \begin{array}{c} \phi_3 \\ \neg & x \end{array} \right\}}_{\alpha_3} \right\}$$

*In this example, we aim to construct  $S_c(\mathcal{K}_7, 1)$ , i. e., a SAT encoding which returns true if and only if 1 is an upper bound of the contention inconsistency value of  $\mathcal{K}_7$ .*

*As  $\text{At}(\mathcal{K}_7) = \{x, y\}$ , we require  $2 \cdot 3 = 6$  new atoms (SC1):*

$$x_t, x_b, x_f, \quad y_t, y_b, y_f$$

We also have a total of 8 (sub)formulas, so we need  $8 \cdot 3 = 24$  corresponding atoms (SC2):

$$\begin{aligned} v_{\alpha_1}^t, v_{\alpha_1}^b, v_{\alpha_1}^f, & \quad v_{\alpha_2}^t, v_{\alpha_2}^b, v_{\alpha_2}^f, & \quad v_{\alpha_3}^t, v_{\alpha_3}^b, v_{\alpha_3}^f, \\ v_{\phi_{1,1}}^t, v_{\phi_{1,1}}^b, v_{\phi_{1,1}}^f, & \quad v_{\phi_{1,2}}^t, v_{\phi_{1,2}}^b, v_{\phi_{1,2}}^f, \\ v_{\phi_{2,1}}^t, v_{\phi_{2,1}}^b, v_{\phi_{2,1}}^f, & \quad v_{\phi_{2,2}}^t, v_{\phi_{2,2}}^b, v_{\phi_{2,2}}^f, \\ & \quad v_{\phi_3}^t, v_{\phi_3}^b, v_{\phi_3}^f \end{aligned}$$

We add the following constraints to the encoding  $S_c(\mathcal{K}_7, u)$  via (SC3):

$$\begin{aligned} (x_t \vee x_f \vee x_b) \wedge (\neg x_t \vee \neg x_f) \wedge (\neg x_t \vee \neg x_b) \wedge (\neg x_b \vee \neg x_f) \\ (y_t \vee y_f \vee y_b) \wedge (\neg y_t \vee \neg y_f) \wedge (\neg y_t \vee \neg y_b) \wedge (\neg y_b \vee \neg y_f) \end{aligned}$$

We encode the first formula  $\alpha_1 = \phi_{1,1} \wedge \phi_{1,2}$ , which is a conjunction, by adding the following formulas corresponding to (SC4)–(SC6):

$$\begin{aligned} v_{\alpha_1}^t &\leftrightarrow v_{\phi_{1,1}}^t \wedge v_{\phi_{1,2}}^t \\ v_{\alpha_1}^f &\leftrightarrow v_{\phi_{1,1}}^f \vee v_{\phi_{1,2}}^f \\ v_{\alpha_1}^b &\leftrightarrow (\neg v_{\phi_{1,1}}^t \vee \neg v_{\phi_{1,2}}^t) \wedge \neg v_{\phi_{1,1}}^f \wedge \neg v_{\phi_{1,2}}^f \end{aligned}$$

We encode the second formula  $\alpha_2 = \phi_{2,1} \vee \phi_{2,2}$ , which is a disjunction, by adding the formulas corresponding to (SC7)–(SC9):

$$\begin{aligned} v_{\alpha_2}^t &\leftrightarrow v_{\phi_{2,1}}^t \vee v_{\phi_{2,2}}^t \\ v_{\alpha_2}^f &\leftrightarrow v_{\phi_{2,1}}^f \wedge v_{\phi_{2,2}}^f \\ v_{\alpha_2}^b &\leftrightarrow (\neg v_{\phi_{2,1}}^f \vee \neg v_{\phi_{2,2}}^f) \wedge \neg v_{\phi_{2,1}}^t \wedge \neg v_{\phi_{2,2}}^t \end{aligned}$$

The last formula ( $\alpha_3 = \neg \phi_3$ ), which is a negation, is modeled by adding the following formulas via (SC10)–(SC12):

$$\begin{aligned} v_{\alpha_3}^t &\leftrightarrow v_{\phi_3}^f \\ v_{\alpha_3}^f &\leftrightarrow v_{\phi_3}^t \\ v_{\alpha_3}^b &\leftrightarrow v_{\phi_3}^b \end{aligned}$$

We add formulas to encode that  $\phi_{1,1}$ ,  $\phi_{1,2}$ ,  $\phi_{2,1}$ ,  $\phi_{2,2}$ , and  $\phi_3$  are subformulas consisting of individual atoms, following (SC13)–(SC15):

$$\begin{array}{ccccc} \phi_{1,1}^t \leftrightarrow x_t & \phi_{1,2}^t \leftrightarrow y_t & \phi_{2,1}^t \leftrightarrow x_t & \phi_{2,2}^t \leftrightarrow y_t & \phi_3^t \leftrightarrow x_t \\ \phi_{1,1}^f \leftrightarrow x_f & \phi_{1,2}^f \leftrightarrow y_f & \phi_{2,1}^f \leftrightarrow x_f & \phi_{2,2}^f \leftrightarrow y_f & \phi_3^f \leftrightarrow x_f \\ \phi_{1,1}^b \leftrightarrow x_b & \phi_{1,2}^b \leftrightarrow y_b & \phi_{2,1}^b \leftrightarrow x_b & \phi_{2,2}^b \leftrightarrow y_b & \phi_3^b \leftrightarrow x_b \end{array}$$

For each formula in  $\mathcal{K}$  we additionally add a constraint wrt. (SC16):

$$\begin{aligned} v_{\alpha_1}^t \vee v_{\alpha_1}^b \\ v_{\alpha_2}^t \vee v_{\alpha_2}^b \\ v_{\alpha_3}^t \vee v_{\alpha_3}^b \end{aligned}$$

Finally, via (SC17), we add the cardinality constraint  $\text{at\_most}(1, \{x_b, y_b\})$ .

### 3.4 The Forgetting-Based Inconsistency Measure

Let  $\mathcal{K}$  again be a propositional knowledge base and  $u$  an integer representing a candidate for an upper bound of  $\mathcal{I}_f(\mathcal{K})$ . Recall that  $\mathcal{I}_f(\mathcal{K})$  is defined as the minimal number of atom occurrences which need to be forgotten in order to resolve all inconsistencies in  $\mathcal{K}$ . We construct an encoding  $S_f(\mathcal{K}, u)$  that represents  $\text{UPPER}_{\mathcal{I}_f}$  for the instance  $(\mathcal{K}, u)$ .

To begin with, we label the given knowledge base  $\mathcal{K}$  to retrieve  $\text{Occ}(\mathcal{K})$ . For each labeled atom occurrence  $X^l \in \text{Occ}(\mathcal{K})$  we also add two atoms  $t_{X,l}$  and  $f_{X,l}$  to the signature (SF1). Further, for each formula  $A \in \mathcal{K}$  a corresponding constraint  $A'$  is created and added to  $S_f(\mathcal{K}, u)$ .  $A'$  is an extension of  $A$ , gained by substituting every atom occurrence  $X^l$  in  $A$  with the subformula  $(t_{X,l} \vee X) \wedge \neg f_{X,l}$ . Assigning truth values to the new atoms  $t_{X,l}$ ,  $f_{X,l}$  models the forgetting operation on  $X^l$  in  $A$ . More precisely, setting  $t_{X,l}$  to true represents forgetting  $X^l$  and replacing it with  $\top$ , while setting  $f_{X,l}$  to true represents forgetting  $X^l$  and replacing it with  $\perp$ . We define  $\Phi_{X,l}$  to be the subformula  $X^l$  is substituted with:

$$\Phi_{X,l} := (t_{X,l} \vee X) \wedge \neg f_{X,l} \quad (\text{SF2})$$

- If  $t_{X,l}$  is set to true and  $f_{X,l}$  is set to false, then  $\Phi_{X,l}$  becomes true, regardless of the truth value assigned to  $X$ . This represents the case in which  $X^l$  has been forgotten and was replaced by  $\top$ .
- If  $t_{X,l}$  is set to false and  $f_{X,l}$  is set to true, then  $\Phi_{X,l}$  becomes false, regardless of the truth value assigned to  $X$ . This represents the case in which  $X^l$  has been forgotten and was replaced by  $\perp$ .
- If both  $t_{X,l}$  and  $f_{X,l}$  are set to false, then the truth value of  $\Phi_{X,l}$  is equal to the truth value assigned to  $X$ . This represents the case in which  $X^l$  has not been forgotten.

The case in which both  $t_{X,l}$  and  $f_{X,l}$  are set to true holds no meaning in the chosen representation, as this would imply that  $X^l$  had somehow been replaced by  $\top$  and  $\perp$  simultaneously. This invalid state needs to be avoided, which motivates the addition of constraints that prevent  $t_{X,l}$  and  $f_{X,l}$  from both being set to true at the same time:

$$\neg t_{X,l} \vee \neg f_{X,l} \quad (\text{SF3})$$

Finally, an at-most- $u$  constraint is added to limit the number of occurrences that can be forgotten. Let  $\text{At}_{\text{forget}}$  denote the set of all atoms introduced by substitutions of atom occurrences in formulas  $A \in \mathcal{K}$ , i. e., the set of all  $t_{X,l}, f_{X,l}$  wrt.  $X^l \in \text{Occ}(\mathcal{K})$ . The constraint

$$\text{at\_most}(u, \text{At}_{\text{forget}}) \quad (\text{SF4})$$

is added to  $S_f(\mathcal{K}, u)$  in order to keep the number of forgotten occurrences from exceeding  $u$ . In total,  $S_f(\mathcal{K}, u)$  is comprised of (SF1)–(SF4) (see Encoding 2 for a complete overview). The proof for the following theorem, along with Encoding 2, is provided in Appendix A.1.2

**Theorem 2.** *For a given value  $u$ , the encoding  $S_f(\mathcal{K}, u)$  is satisfiable if and only if  $\mathcal{I}_f(\mathcal{K}) \leq u$ .*

**Example 11.** Consider again the knowledge base  $\mathcal{K}_7$  from Example 10:

$$\mathcal{K}_7 = \left\{ \underbrace{\begin{array}{c} \phi_{1,1} \quad \phi_{1,2} \\ x \wedge y \\ \alpha_1 \end{array}} \right., \left. \underbrace{\begin{array}{c} \phi_{2,1} \quad \phi_{2,2} \\ x \vee y \\ \alpha_2 \end{array}} \right., \left. \underbrace{\begin{array}{c} \phi_3 \\ \neg x \\ \alpha_3 \end{array}} \right\}$$

In this example, we aim to construct  $S_f(\mathcal{K}_7, 2)$ , i. e., a SAT encoding which returns true if and only if 2 is an upper bound of the forgetting-based inconsistency value of  $\mathcal{K}_7$ .

To begin with, we need  $|\text{Occ}(\mathcal{K}_7)| \cdot 2 = 5 \cdot 2 = 10$  atoms that represent the forgetting operation wrt. each atom occurrence (SF1):

$$t_{x,1}, f_{x,1}, t_{y,1}, f_{y,1}, \quad t_{x,2}, f_{x,2}, t_{y,2}, f_{y,2}, \quad t_{x,3}, f_{x,3}$$

We substitute each occurrence in every formula in  $\mathcal{K}_7$  as described by (SF2). This produces the following 3 formulas, which are all added to  $S_f(\mathcal{K}_7, 2)$ :

$$\begin{aligned} & ((t_{x,1} \vee x) \wedge \neg f_{x,1}) \wedge ((t_{y,1} \vee y) \wedge \neg f_{y,1}) \\ & ((t_{x,2} \vee x) \wedge \neg f_{x,2}) \vee ((t_{y,2} \vee y) \wedge \neg f_{y,2}) \\ & \neg((t_{x,3} \vee x) \wedge \neg f_{x,3}) \end{aligned}$$

Next, we apply (SF3):

$$\begin{aligned} & \neg t_{x,1} \vee \neg f_{x,1} & \neg t_{y,1} \vee \neg f_{y,1} \\ & \neg t_{x,2} \vee \neg f_{x,2} & \neg t_{y,2} \vee \neg f_{y,2} \\ & \neg t_{x,3} \vee \neg f_{x,3} \end{aligned}$$

Finally, we add the at-most- $u$  constraint (SF4) to  $S_f(\mathcal{K}_7, 2)$ :

$$\text{at\_most}(2, \{t_{x,1}, f_{x,1}, t_{y,1}, f_{y,1}, t_{x,2}, f_{x,2}, t_{y,2}, f_{y,2}, t_{x,3}, f_{x,3}\})$$

### 3.5 The Hitting Set Inconsistency Measure

In the following, we propose a SAT encoding, denoted  $S_h(\mathcal{K}, u)$ , for the problem  $\mathcal{I}_h(\mathcal{K})$ , with  $\mathcal{K}$  being a knowledge base and  $u$  an integer representing a candidate for an upper bound. We reiterate that  $\mathcal{I}_h(\mathcal{K})$  is the minimum cardinality of a set of interpretations such that there is at least one model for every formula in the knowledge base  $\mathcal{K}$ , subtracted by 1. This is equivalent to finding the minimal number of blocks of a partitioning of  $\mathcal{K}$  such that each block is satisfiable (again, subtracted by 1).

The formulas in  $\mathcal{K}$  are partitioned into  $u$  blocks, and every formula belongs to at least one block. For every formula  $A \in \mathcal{K}$ , we create copies  $A_i$  with  $i = 1, \dots, u$ . The meaning of formula  $A_i$  is “the  $i$ -th copy of formula  $A$ ”. Then, we create  $|\mathcal{K}| \cdot u$  new atoms  $p_{A,i}$  (SH2). The meaning of the atom  $p_{A,i}$  is “formula  $A$  is a member of block  $i$ ”. For every  $i$ , all atoms  $X$  that appear in a formula  $A_i$  are replaced with new atoms  $X_i$  (SH1), meaning each block has its own copy of  $\text{At}(\mathcal{K})$ , allowing us to check for satisfiability block-wise.

For every block, the set of formulas inside it is satisfiable. We add  $|\mathcal{K}| \cdot u$  constraints of the form

$$p_{A,i} \rightarrow A_i. \tag{SH3}$$

This models that if a formula  $A$  belongs to block  $i$ , then the corresponding formula  $A_i$  must be satisfied (within this block). Every formula must be sorted into at least one block, meaning for all  $A \in \mathcal{K}$ , at least one  $p_{A,i}$  is true. A simple encoding for the “at-least-one” constraint is to add the following clauses for every  $A \in \mathcal{K}$ :

$$\bigvee_{1 \leq i \leq u} p_{A,i} \quad (\text{SH4})$$

When a value for the minimal number of satisfiable partition blocks is found, we subtract 1 to generate the value  $\mathcal{I}_h(\mathcal{K})$ . If no value is found in the range, the result is  $\infty$ .

The following result establishes that the above encoding  $S_h(\mathcal{K}, u)$ , defined by (SH1)–(SH4) (for an overview, see Encoding 3 in Appendix A.1.3), faithfully implements  $\mathcal{I}_h$ . The proof of the theorem below is given in Appendix A.1.3 as well.

**Theorem 3.** *For a given value  $u$ , the encoding  $S_h(\mathcal{K}, u)$  is satisfiable if and only if  $\mathcal{I}_h(\mathcal{K}) \leq u - 1$ .  $S_h(\mathcal{K}, u)$  is unsatisfiable for all  $u = 1, \dots, |\mathcal{K}|$  if and only if  $\mathcal{I}_h(\mathcal{K}) = \infty$ .*

**Example 12.** *Consider again the knowledge base  $\mathcal{K}_7$  from Example 10:*

$$\mathcal{K}_7 = \left\{ \underbrace{\begin{array}{cc} \phi_{1,1} & \phi_{1,2} \\ x & \wedge & y \end{array}}_{\alpha_1}, \quad \underbrace{\begin{array}{cc} \phi_{2,1} & \phi_{2,2} \\ x & \vee & y \end{array}}_{\alpha_2}, \quad \underbrace{\begin{array}{c} \phi_3 \\ \neg & x \end{array}}_{\alpha_3} \right\}$$

In this example, we aim to construct  $S_h(\mathcal{K}_7, 2)$ , i. e., a SAT encoding which returns true if and only if  $3 - 1 = 2$  is an upper bound of the hitting set inconsistency value of  $\mathcal{K}_7$ .

First, we create variables for each atom  $x_i \in \text{At}(\mathcal{K}_7) = \{x, y\}$  with  $i \in \{1, 2\}$  (SH1):

$$x_1, x_2, \quad y_1, y_2$$

In addition, for each formula in  $\mathcal{K}_7$  we create the following variables (SH2):

$$p_{\alpha_1,1}, p_{\alpha_1,2}, \quad p_{\alpha_2,1}, p_{\alpha_2,2}, \quad p_{\alpha_3,1}, p_{\alpha_3,2},$$

Now we create indexed copies of the formulas in  $\mathcal{K}_7$ :

$$\begin{array}{ccc} x_1 \wedge y_1, & x_1 \vee y_1, & \neg x_1, \\ x_2 \wedge y_2, & x_2 \vee y_2, & \neg x_2 \end{array}$$

Following (SH3), we add the subsequent formulas to  $S_h(\mathcal{K}_7, 2)$ :

$$\begin{array}{ccc} p_{\alpha_1,1} \rightarrow (x_1 \wedge y_1), & p_{\alpha_2,1} \rightarrow (x_1 \vee y_1), & p_{\alpha_3,1} \rightarrow (\neg x_1), \\ p_{\alpha_1,2} \rightarrow (x_2 \wedge y_2), & p_{\alpha_2,2} \rightarrow (x_2 \vee y_2), & p_{\alpha_3,2} \rightarrow (\neg x_2) \end{array}$$

At last, we add the following constraints via (SH4):

$$p_{\alpha_1,1} \vee p_{\alpha_1,2}, \quad p_{\alpha_2,1} \vee p_{\alpha_2,2}, \quad p_{\alpha_3,1} \vee p_{\alpha_3,2}$$

### 3.6 The Max-Distance Inconsistency Measure

Let, once again,  $\mathcal{K}$  be a propositional knowledge base, and let  $u$  be an integer representing a potential upper bound for  $\mathcal{I}_d^{\max}(\mathcal{K})$ . In the following, we construct a SAT encoding  $S_d^{\max}(\mathcal{K}, u)$  for the problem  $\text{UPPER}_{\mathcal{I}_d^{\max}}$ .

As a reminder, the distance-based inconsistency measures look for an interpretation with an “optimal” distance. Regarding  $\mathcal{I}_d^{\max}$ , the “optimal” interpretation has the minimal maximum distance to the models of all formulas in  $\mathcal{K}$ . In case one or more formulas in  $\mathcal{K}$  are unsatisfiable, i. e., they possess no models,  $\mathcal{I}_d^{\max}(\mathcal{K}) = \infty$ .

First we need a representation of the “optimal” interpretation that will be computed by the measure. For each atom  $X \in \text{At}(\mathcal{K})$ , we add a new atom  $X_o$  that represents this atom’s value assignment in the “optimal interpretation” (SDM1). Let  $\text{At}_O$  be the set of atoms of the optimal interpretation, i. e.,  $\text{At}_O = \{X_o \mid X \in \text{At}(\mathcal{K})\}$ .

With the exception of the case that  $\mathcal{I}_d^{\max}(\mathcal{K}) = \infty$  (i. e., the case in which at least one of the formulas is contradictory), all formulas in  $\mathcal{K}$  must be satisfied by some interpretation. For every formula  $A \in \mathcal{K}$ , we create a clone  $A_i$  with  $i \in \{1, \dots, |\mathcal{K}|\}$ . For every  $i$ , all atoms  $X$  that appear in a formula  $A_i$  are replaced with new atoms  $X_i$  (SDM2). We add the modified clones  $A_i$  to the encoding (SDM4). This models that the formulas are all satisfiable individually and allows us to put their models and the optimal interpretation in relation in the next step.

For each formula, the distance between at least one model of the formula and the optimal interpretation consists of at most  $u$  differing atom assignments. We now aim to model that the minimum distance between some model of each formula and the optimal world is at most  $u$ . This can also be formulated as “for each formula, there is some model that can be converted into the optimal world by inverting at most  $u$  value assignments” (excluding the  $\infty$  case). To represent this, we create new atoms  $\text{inv}_{X,i}$  with  $i \in \{1, \dots, |\mathcal{K}|\}$  for all  $X_i \in \text{At}(S_d^{\max}(\mathcal{K}, u)) \setminus \text{At}_O$  that model the inverted value assignments (SDM3). Then, we add the following formulas for every  $X_i$ , using the optimal interpretation atoms  $X_o \in \text{At}_O$ :

$$X_i \rightarrow X_o \vee \text{inv}_{X,i} \quad (\text{SDM5})$$

$$\neg X_i \rightarrow \neg X_o \vee \text{inv}_{X,i} \quad (\text{SDM6})$$

The above formulas represent that for every atom  $X_i$  of a model of a formula, the corresponding atom in the optimal interpretation has the same value unless it is inverted by setting  $\text{inv}_{X,i}$  to true. We then represent the upper bound by adding at-most- $u$  constraints for all  $\text{inv}_{X,i}$ . Let  $\text{INV}_i$  denote the set containing all  $\text{inv}_{X,i}$  for a fixed  $i$ :

$$\text{at\_most}(u, \text{INV}_i) \quad (\text{SDM7})$$

In total, the SAT encoding  $S_d^{\max}(\mathcal{K}, u)$  comprises (SDM1)–(SDM7). The following result establishes that the above encoding faithfully implements  $\mathcal{I}_d^{\max}$ .

**Theorem 4.** *For a given value  $u$ , the encoding  $S_d^{\max}(\mathcal{K}, u)$  is satisfiable if and only if  $\mathcal{I}_d^{\max}(\mathcal{K}) \leq u$ .  $S_d^{\max}(\mathcal{K}, u)$  is unsatisfiable for all  $u \in \{0, \dots, |\text{At}(\mathcal{K})|\}$  if and only if  $\mathcal{I}_d^{\max}(\mathcal{K}) = \infty$ .*

The proof of the above theorem, as well as Encoding 4, which contains a complete overview of  $S_d^{\max}(\mathcal{K}, u)$ , are provided in Appendix A.1.4.

**Example 13.** Consider again the knowledge base  $\mathcal{K}_7$  from Example 10:

$$\mathcal{K}_7 = \left\{ \underbrace{\left\{ \begin{array}{cc} \phi_{1,1} & \phi_{1,2} \\ x & \wedge & y \end{array} \right\}}_{\alpha_1}, \underbrace{\left\{ \begin{array}{cc} \phi_{2,1} & \phi_{2,2} \\ x & \vee & y \end{array} \right\}}_{\alpha_2}, \underbrace{\left\{ \begin{array}{c} \phi_3 \\ \neg & x \end{array} \right\}}_{\alpha_3} \right\}$$

In this example, we aim to construct  $S_d^{\max}(\mathcal{K}_7, 1)$ , i. e., a SAT encoding which returns true if and only if 1 is an upper bound of the max-distance inconsistency value of  $\mathcal{K}_7$ .

At first, we create the variables  $x_o, y_o$  for the atoms in  $\text{At}(\mathcal{K}_7) = \{x, y\}$  (SDM1). We create additional variables for each atom in  $\text{At}(\mathcal{K}_7)$  wrt. indices  $i \in \{1, \dots, |\mathcal{K}_7|\}$  with  $|\mathcal{K}_7| = 3$  via (SDM2) and (SDM3):

$$\begin{array}{ll} x_1, x_2, x_3, & y_1, y_2, y_3 \\ \text{inv}_{x,1}, \text{inv}_{x,2}, \text{inv}_{x,3}, & \text{inv}_{y,1}, \text{inv}_{y,2}, \text{inv}_{y,3} \end{array}$$

We add indexed copies of the formulas in  $\mathcal{K}_7$  to  $S_d^{\max}(\mathcal{K}_7, 1)$  (SDM4):

$$x_1 \wedge y_1, \quad x_2 \vee y_2, \quad \neg x_3$$

We now add the following constraints conforming to (SDM5) and (SDM6) to  $S_d^{\max}(\mathcal{K}_7, 1)$ :

$$\begin{array}{ll} x_1 \rightarrow x_o \vee \text{inv}_{x,1}, & \neg x_1 \rightarrow \neg x_o \vee \text{inv}_{x,1}, \\ y_1 \rightarrow y_o \vee \text{inv}_{y,1}, & \neg y_1 \rightarrow \neg y_o \vee \text{inv}_{y,1}, \\ x_2 \rightarrow x_o \vee \text{inv}_{x,2}, & \neg x_2 \rightarrow \neg x_o \vee \text{inv}_{x,2}, \\ y_2 \rightarrow y_o \vee \text{inv}_{y,2}, & \neg y_2 \rightarrow \neg y_o \vee \text{inv}_{y,2}, \\ x_3 \rightarrow x_o \vee \text{inv}_{x,3}, & \neg x_3 \rightarrow \neg x_o \vee \text{inv}_{x,3}, \\ y_3 \rightarrow y_o \vee \text{inv}_{y,3}, & \neg y_3 \rightarrow \neg y_o \vee \text{inv}_{y,3} \end{array}$$

Lastly, the following at-most- $u$  constraints are added:

$$\text{at\_most}(1, \{\text{inv}_{x,1}, \text{inv}_{y,1}\}), \quad \text{at\_most}(1, \{\text{inv}_{x,2}, \text{inv}_{y,2}\}), \quad \text{at\_most}(1, \{\text{inv}_{x,3}, \text{inv}_{y,3}\})$$

### 3.7 The Sum-Distance Inconsistency Measure

The definition of the sum-distance measure  $\mathcal{I}_d^\Sigma$  is quite similar to the definition of the previously addressed max-distance measure  $\mathcal{I}_d^{\max}$ . The only difference is the notion of optimality—while  $\mathcal{I}_d^{\max}$  seeks the interpretation with the smallest maximum distance to the models of the formulas in a given knowledge base  $\mathcal{K}$ ,  $\mathcal{I}_d^\Sigma$  seeks an interpretation such that the sum of its distances to the models of the formulas in  $\mathcal{K}$  is minimal.

Let once again  $u$  be an integer representing a candidate for an upper bound of  $\mathcal{I}_d^\Sigma(\mathcal{K})$ . In the following, we construct a SAT encoding  $S_d^\Sigma(\mathcal{K}, u)$  for the problem  $\text{UPPER}_{\mathcal{I}_d^\Sigma}$ .

The constraints (SDS1)–(SDS6) of the encoding  $S_d^\Sigma(\mathcal{K}, u)$  are identical to constraints (SDM1)–(SDM6) of  $S_d^{\max}(\mathcal{K}, u)$ . Only the use of the at-most- $u$  constraints is changed—we represent the upper bound for the sum of the distances by adding a single at-most- $u$  constraint for all  $\text{inv}_{X,i}$  with  $i \in \{1, \dots, |\mathcal{K}|\}$ . Let  $\text{INV}$  be the set containing all  $\text{inv}_{X,i}$ :

$$\text{at\_most}(u, \text{INV}) \tag{SDS7}$$

Altogether,  $S_d^{\max}(\mathcal{K}, u)$  is comprised of (SDS1)–(SDS7) (for an overview, see Encoding 5 in Appendix A.1.5). The following result establishes that the above encoding faithfully implements  $\mathcal{I}_d^\Sigma$ . The corresponding proof can likewise be found in Appendix A.1.5.

**Theorem 5.** *For a given value  $u$ , the encoding  $S_d^\Sigma(\mathcal{K}, u)$  is satisfiable if and only if  $\mathcal{I}_d^\Sigma(\mathcal{K}) \leq u$ .  $S_d^\Sigma(\mathcal{K}, u)$  is unsatisfiable for all  $u \in \{0, \dots, |\text{At}(\mathcal{K})| \cdot |\mathcal{K}|\}$  if and only if  $\mathcal{I}_d^\Sigma(\mathcal{K}) = \infty$ .*

**Example 14.** *Consider again the knowledge base  $\mathcal{K}_7$  from Example 10:*

$$\mathcal{K}_7 = \left\{ \underbrace{\begin{array}{cc} \phi_{1,1} & \phi_{1,2} \\ x & \wedge & y \end{array}}_{\alpha_1}, \underbrace{\begin{array}{cc} \phi_{2,1} & \phi_{2,2} \\ x & \vee & y \end{array}}_{\alpha_2}, \underbrace{\begin{array}{c} \phi_3 \\ \neg & x \end{array}}_{\alpha_3} \right\}$$

In this example, we aim to construct  $S_d^\Sigma(\mathcal{K}_7, 2)$ , i. e., a SAT encoding which returns true if and only if 2 is an upper bound of the sum-distance inconsistency value of  $\mathcal{K}_7$ .

The variables we require (defined by (SDS1)–(SDS3)) are the same as those in Example 13. Furthermore, the constraints we need to add to  $S_d^\Sigma(\mathcal{K}_7, 2)$  via (SDS4)–(SDS6) correspond exactly to those defined by (SDM4)–(SDM6) in Example 13. However, we use a different at-most- $u$  constraint following (SDS7):

$$\text{at\_most}(2, \{\text{inv}_{x,1}, \text{inv}_{y,1}, \text{inv}_{x,2}, \text{inv}_{y,2}, \text{inv}_{x,3}, \text{inv}_{y,3}\})$$

### 3.8 The Hit-Distance Inconsistency Measure

Let  $\mathcal{K}$  once again be a knowledge base and  $u$  an integer representing a possible upper bound of  $\mathcal{I}_d^{\text{hit}}(\mathcal{K})$ . In the following, we construct a SAT encoding  $S_d^{\text{hit}}(\mathcal{K}, u)$  for  $\text{UPPER}_{\mathcal{I}_d^{\text{hit}}}$ .

In contrast to the other distance-based measures, the value of  $\mathcal{I}_d^{\text{hit}}(\mathcal{K})$  is not a distance value, but a count of distances greater than 0. In other words,  $\mathcal{I}_d^{\text{hit}}(\mathcal{K})$  is the maximum number of formulas in  $\mathcal{K}$  that can be satisfied by one interpretation, or the minimum number of formulas that need to be removed in order to make  $\mathcal{K}$  consistent. To represent this, we add one new variable  $\text{hit}_A$  for every formula  $A \in \mathcal{K}$  (SDH1). Let  $\text{HIT}_{\mathcal{K}} = \{\text{hit}_A \mid A \in \mathcal{K}\}$ . For each  $A \in \mathcal{K}$ , we also add a disjunction that uses one of the new variables:

$$A \vee \text{hit}_A \tag{SDH3}$$

This formula represents that if  $A$  is not true, then the new variable  $\text{hit}_A$  is set to true. Note that all atoms of the input knowledge base  $\mathcal{K}$  still appear in the encoding as part of  $A$  (SDH2). We then represent the upper bound for the number of models with distances  $> 0$  by adding one at-most- $u$  constraint that restricts the number of  $\text{hit}_A \in \text{HIT}_{\mathcal{K}}$  that are allowed to be set to true:

$$\text{at\_most}(u, \text{HIT}_{\mathcal{K}}) \tag{SDH4}$$

In total,  $S_d^{\text{hit}}(\mathcal{K}, u)$  is defined by (SDH1)–(SDH4). Appendix A.1.6 provides an overview of  $S_d^{\text{hit}}(\mathcal{K}, u)$  (i. e., Encoding 6), as well as the correctness proof of the theorem below.

**Theorem 6.** *For a given value  $u$ , the encoding  $S_d^{\text{hit}}(\mathcal{K}, u)$  is satisfiable if and only if  $\mathcal{I}_d^{\text{hit}}(\mathcal{K}) \leq u$ .*

**Example 15.** Consider again the knowledge base  $\mathcal{K}_7$  from Example 10:

$$\mathcal{K}_7 = \left\{ \underbrace{\begin{array}{c} \phi_{1,1} \quad \phi_{1,2} \\ x \quad \wedge \quad y \\ \alpha_1 \end{array}} \right\}, \quad \left\{ \underbrace{\begin{array}{c} \phi_{2,1} \quad \phi_{2,2} \\ x \quad \vee \quad y \\ \alpha_2 \end{array}} \right\}, \quad \left\{ \underbrace{\begin{array}{c} \phi_3 \\ \neg x \\ \alpha_3 \end{array}} \right\}$$

In this example, we aim to construct  $S_d^{\text{hit}}(\mathcal{K}_7, 1)$ , i. e., a SAT encoding which returns true if and only if 1 is an upper bound of the hit-distance inconsistency value of  $\mathcal{K}_7$ .

First, following (SDH1), we define the variables  $\text{hit}_{\alpha_1}, \text{hit}_{\alpha_2}, \text{hit}_{\alpha_3}$  wrt. the formulas in  $\mathcal{K}_7$ . Moreover, we require the variables  $\text{At}(\mathcal{K}_7) = \{x, y\}$  (SDH2). Following (SDH3), we add the subsequent constraints to  $S_d^{\text{hit}}(\mathcal{K}_7, 1)$ :

$$(x \wedge y) \vee \text{hit}_{\alpha_1}, \quad (x \vee y) \vee \text{hit}_{\alpha_2}, \quad (\neg x) \vee \text{hit}_{\alpha_3}$$

At last, we add  $\text{at\_most}(1, \{\text{hit}_{\alpha_1}, \text{hit}_{\alpha_2}, \text{hit}_{\alpha_3}\})$  via (SDH4).

## 4. ASP-Based Algorithms for Selected Inconsistency Measures

After the introduction of a general SAT-based approach for computing inconsistency degrees, as well as specific encodings for the problem  $\text{UPPER}_{\mathcal{I}}$  wrt. all six inconsistency measures addressed in this work, we now introduce answer set programming as an alternative method to compute inconsistency values.

### 4.1 Answer Set Programming

*Answer Set Programming* (ASP) (Lifschitz, 2019; Gebser, Kaminski, Kaufmann, & Schaub, 2012; Lifschitz, 2008; Brewka, Eiter, & Truszczynski, 2011) is a declarative problem solving approach targeted at difficult search problems. Thus, rather than modeling instructions on how to solve a problem, a representation of the problem itself is modeled. The goal is to represent a problem in a logical format (a *logic program*) such that the models of this representation describe solutions of the original problem. These models are called *answer sets*.

In short, a logic program is a finite set of *rules* of the form

$$r = a_0 \text{ :- } a_1, \dots, a_n, \text{ not } a_{n+1}, \dots, \text{ not } a_m.$$

with each  $a_i$  ( $0 \leq i \leq n \leq m$ ) being atoms, and “not” denoting the default negation in the sense of Reiter (1980). In ASP, an atom has the form  $p(\mathbf{t}_1, \dots, \mathbf{t}_n)$ , with  $p$  being a predicate symbol, and  $\mathbf{t}_1, \dots, \mathbf{t}_n$  being terms. Terms are either constants, variables, arithmetic terms (i. e.,  $-\mathbf{t}_1$  or  $\mathbf{t}_1 \circ \mathbf{t}_2$  with  $\circ \in \{+, -, *, /\}$  wrt. some terms  $\mathbf{t}_1, \mathbf{t}_2$ ), or functional terms (i. e.,  $\varphi(\mathbf{t}_1, \dots, \mathbf{t}_m)$  with  $\varphi$  being a functor,  $\mathbf{t}_1, \dots, \mathbf{t}_m$  being terms, and  $m > 0$ ) (Calimeri, Faber, Gebser, Ianni, Kaminski, Krennwallner, Leone, Maratea, Ricca, & Schaub, 2020). Moreover, we express the arity  $n$  of a predicate or function  $\varphi$  as  $\varphi/n$ , and an ASP literal  $\ell$  is either an atom or a default-negated atom.

The *head* of a rule  $r$  (as shown above) is  $\text{head}(r) = \{a_0\}$  and the *body* is  $\text{body}(r) = \{a_1, \dots, a_n, \text{not } a_{n+1}, \dots, \text{not } a_m\}$ . If  $\text{body}(r) = \emptyset$ , then  $r$  is a *fact*, and if  $\text{head}(r) = \emptyset$ , then  $r$  is a *constraint*. We further divide the body elements of  $r$  into  $\text{body}^+(r) = \{a_1, \dots, a_n\}$  and  $\text{body}^-(r) = \{a_{n+1}, \dots, a_m\}$ .

An atom/rule/program is *ground* if it does not contain any variables.<sup>3</sup> Let  $X$  be a set of ground atoms. We define  $X$  to be a *model* of a ground logic program  $P$  if for all  $r \in P$ ,  $\text{head}(r) \in X$  whenever  $\text{body}^+(r) \subseteq X$  and  $\text{body}^-(r) \cap X = \emptyset$ . The *reduct* (Gelfond & Lifschitz, 1988) of  $P$  wrt.  $X$  is defined as

$$P^X = \{\text{head}(r) \text{ :- body}^+(r) \mid \text{body}^-(r) \cap X = \emptyset \text{ with } r \in P\}$$

Moreover,  $X$  is an *answer set* of  $P$  if it is a subset-minimal model of  $P^X$ .

Another language concept we make use of is the *conditional literal*, which is of the form  $\ell_0 : \ell_1, \dots, \ell_n$ , with  $\ell_0, \dots, \ell_n$  being literals. The idea behind a conditional literal is to regulate the instantiation of  $\ell_0$  by means of  $\ell_1, \dots, \ell_n$ . In other words, we can view a conditional literal as the set  $\{\ell_0 \mid \ell_1, \dots, \ell_n\}$ . Moreover, we use *cardinality constraints*. A cardinality constraint is of the form  $l\{c_1; \dots; c_m\}u$ , with  $c_1, \dots, c_m$  being conditional literals,  $l$  constituting an optional lower bound, and  $u$  an optional upper bound. Intuitively, this can be read as “at least  $l$ , and at most  $u$  of the literals specified by  $c_1, \dots, c_m$  must be satisfied”.

It is possible to use cardinality constraints not only in rule bodies, but also in heads. A rule with a cardinality constraint as the head is referred to as an (*extended*) *choice rule*. Formally, a choice rule has the form

$$r_c = l\{a_1, \dots, a_m\}u \text{ :- } a_{m+1}, \dots, a_n, \text{ not } a_{n+1}, \dots, \text{ not } a_o$$

with  $0 \leq m \leq n \leq o$ ,  $a_1, \dots, a_o$  being atoms,  $0 \leq l \leq u$ , and  $l, u$  being (optional) lower and upper bounds, respectively. The intuition behind such a rule is that any subset of the head atoms (which complies with the upper and lower bound, if specified) can be included in the answer set. Note that choice rules can be transformed into sets of normal rules (Gebser et al., 2012).

Furthermore, we use *aggregates* and *optimization statements*. The former are used to reason about minima, maxima, sums, and counts over sets of literals. Let an *aggregate element*  $g$  be defined as  $g = t_1, \dots, t_m : \ell_1, \dots, \ell_n$  with  $t_1, \dots, t_m$  being terms, and  $\ell_1, \dots, \ell_n$  being literals. An *aggregate* is then defined as  $\#\text{agg}\{g_1, \dots, g_n\} \prec t$ , with  $\#\text{agg} \in \{\#\text{count}, \#\text{min}, \#\text{max}, \#\text{sum}\}$  being an aggregate function name,  $g_1, \dots, g_n$  being aggregate elements,  $\prec \in \{<, >, \leq, \geq, =, \neq\}$  being an aggregate relation, and  $t$  being a term (Calimeri et al., 2020). Optimization statements serve the purpose of expressing cost functions that are subject to minimization or maximization. In this work, we only use a specific form of *minimize statements*, which are of the form  $\#\text{minimize}\{t_1, \dots, t_m : \ell_1, \dots, \ell_n\}$ , with  $t_1, \dots, t_m$  being terms and  $\ell_1, \dots, \ell_n$  being literals. We refer to a set that complies with the minimization as an *optimal* answer set.

We further make use of the interval (“..”) and pooling (“;”) operators to abbreviate notation. Intervals let us create multiple instances of a predicate determined by an interval of numerical values, and pooling lets us create multiple instances of a predicate by separating elements by “;”. For more details on ASP we refer the reader to the ASP-Core-2 standard (Calimeri et al., 2020).

3. Note that, following the Clingo syntax, all variable names we use start with an uppercase letter, and all constant names start with a lowercase letter. We also make use of *anonymous variables*. Those are variables that do not recur within a rule, and that are denoted by “\_”.

## 4.2 The Contension Inconsistency Measure

Given a knowledge base  $\mathcal{K}$ , we construct an extended logic program  $P_c(\mathcal{K})$  to determine the contension inconsistency value  $\mathcal{I}_c(\mathcal{K})$  as described in the following. Note that, like the SAT-based approach for  $\mathcal{I}_c$ , the ASP-based method was already proposed in (Kuhlmann et al., 2022), but is covered again for the sake of completeness.

At first, we declare some facts describing the overall composition of the knowledge base, i. e., we describe which formulas are elements of the knowledge base, which atoms are in the knowledge base’s signature, and whether a (sub)formula is a conjunction, disjunction, negation, or atom. We define for every formula  $A \in \mathcal{K}$

$$\text{kbMember}(A). \tag{AC1}$$

and for every atom  $X \in \text{At}(\mathcal{K})$

$$\text{atom}(X). \tag{AC2}$$

Note that the  $A$  in  $\mathcal{K}$  and the  $A$  in `kbMember/1` are formally not the same—in  $\mathcal{K}$  it is a propositional formula, while in `kbMember/1` it is a string representation of that formula. More precisely, a formula  $A$  is “translated” to a uniquely defined ASP constant, i. e., a string starting with a lowercase letter.<sup>4</sup> The same applies to the  $X$  in `At( $\mathcal{K}$ )` and the  $X$  in `atom/1`. In the following, whenever a propositional logic symbol appears in an ASP rule, it is to be interpreted as a uniquely defined ASP constant.

Every conjunction  $\Phi_c = \Psi_{c,1} \wedge \Psi_{c,2}$  is encoded as

$$\text{conjunction}(\Phi_c, \Psi_{c,1}, \Psi_{c,2}). \tag{AC3}$$

In the same fashion, we define disjunctions  $\Phi_d = \Psi_{d,1} \vee \Psi_{d,2}$  as

$$\text{disjunction}(\Phi_d, \Psi_{d,1}, \Psi_{d,2}). \tag{AC4}$$

For each negated formula, i. e., for each  $\Phi_n = \neg\Psi_n$ , we define

$$\text{negation}(\Phi_n, \Psi_n). \tag{AC5}$$

Further, we need to encode subformulas which consist of individual atoms. Hence, for each formula  $\Phi_a$  which is equal to an atom  $X$  we define

$$\text{formulaIsAtom}(\Phi_a, X). \tag{AC6}$$

Additionally, we need to declare the truth values of Priest’s three-valued logic ( $t$ ,  $f$ ,  $b$ ) as facts:

$$\text{tv}(t; f; b). \tag{AC7}$$

To encode how the contension inconsistency measure actually works, we first need to ensure that an atom is not assigned multiple truth values at the same time, i. e., we need to ensure that each atom gets a distinct evaluation. To this end, we introduce the predicate

---

4. For example, a given formula  $\alpha_1$  could be represented as `alpha.1`.

`truthValue/2` and define that an atom has exactly one truth value by making use of a cardinality constraint:

$$1\{\text{truthValue}(A,T) : \text{tv}(T)\}1 \text{ :- atom}(A). \quad (\text{AC8})$$

In order to encode the evaluation of formulas, we need to model the role of the connectives  $\wedge$ ,  $\vee$ , and  $\neg$  in Priest's three-valued logic (see Table 1). For a conjunction to be true in Priest's three-valued logic, both of its conjuncts have to be true:

$$\begin{aligned} \text{truthValue}(F,t) \text{ :- conjunction}(F,G,H), \\ \text{truthValue}(G,t), \\ \text{truthValue}(H,t). \end{aligned} \quad (\text{AC9})$$

For a conjunction to be false, it is sufficient if at least one of its conjuncts is false:

$$\begin{aligned} \text{truthValue}(F,f) \text{ :- conjunction}(F,G,H), \\ 1\{\text{truthValue}(G,f); \text{truthValue}(H,f)\}. \end{aligned} \quad (\text{AC10})$$

A conjunction is  $b$  in three-valued logic if it is neither  $t$  nor  $f$ :

$$\begin{aligned} \text{truthValue}(F,b) \text{ :- conjunction}(F,-,-), \\ \text{not truthValue}(F,t), \\ \text{not truthValue}(F,f). \end{aligned} \quad (\text{AC11})$$

Analogously, we define that a disjunction is false if both its disjuncts are false, and true if at least one of its disjuncts is true. Again, if a disjunction is neither  $t$  nor  $f$ , it is  $b$ .

$$\begin{aligned} \text{truthValue}(F,f) \text{ :- disjunction}(F,G,H), \\ \text{truthValue}(G,f), \\ \text{truthValue}(H,f). \end{aligned} \quad (\text{AC12})$$

$$\begin{aligned} \text{truthValue}(F,t) \text{ :- disjunction}(F,G,H), \\ 1\{\text{truthValue}(G,t); \text{truthValue}(H,t)\}. \end{aligned} \quad (\text{AC13})$$

$$\begin{aligned} \text{truthValue}(F,b) \text{ :- disjunction}(F,-,-), \\ \text{not truthValue}(F,t), \\ \text{not truthValue}(F,f). \end{aligned} \quad (\text{AC14})$$

A negation is  $t$  in three-valued logic if its base formula is  $f$  and vice versa. A negation is  $b$  if its base formula is  $b$  as well. Thus, we can formulate the following rules to model three-valued negation in ASP:

$$\text{truthValue}(F,t) \text{ :- negation}(F,G), \text{truthValue}(G,f). \quad (\text{AC15})$$

$$\text{truthValue}(F,f) \text{ :- negation}(F,G), \text{truthValue}(G,t). \quad (\text{AC16})$$

$$\text{truthValue}(F,b) \text{ :- negation}(F,G), \text{truthValue}(G,b). \quad (\text{AC17})$$

Moreover, if a (sub)formula only consists of a single atom, it must have the same truth value:

$$\begin{aligned} \text{truthValue}(F, T) & :- \text{formulaIsAtom}(F, G), \\ & \quad \text{truthValue}(G, T), \\ & \quad \text{tv}(T). \end{aligned} \tag{AC18}$$

Further, all formulas  $A \in \mathcal{K}$  need to be either  $t$  or  $b$  wrt. Priest's three-valued logic. Thus, no formula is allowed to be evaluated to  $f$ . We accomplish this by introducing the following integrity constraint:

$$:- \text{truthValue}(F, f), \text{kbMember}(F). \tag{AC19}$$

At last, to actually compute  $\mathcal{I}_c(\mathcal{K})$ , we require the minimal number of atoms in  $\mathcal{K}$  evaluated to  $b$ . This is achieved by means of a minimize statement:

$$\#\text{minimize}\{1, A: \text{truthValue}(A, b), \text{atom}(A)\}. \tag{AC20}$$

Now  $P_c(\mathcal{K})$  is the union of all rules defined in (AC1)–(AC20) (for an overview, see Encoding 7 in Appendix A.2.1). Further, let  $\omega_M^3$  be the three-valued interpretation represented by an answer set  $M$  of  $P_c(\mathcal{K})$ . To be precise, if  $M$  is an answer set of  $P_c(\mathcal{K})$  then the three-valued interpretation  $\omega_M^3$  defined as

$$\omega_M^3(X) = \begin{cases} t & \text{truthValue}(X, t) \in M \\ f & \text{truthValue}(X, f) \in M \\ b & \text{truthValue}(X, b) \in M \end{cases}$$

is a model of  $\mathcal{K}$  with  $X \in \text{At}(\mathcal{K})$ . The proofs regarding the well-definedness of the previous definition as well as the following theorem are provided in Appendix A.2.1.

**Theorem 7.** *Let  $M_o$  be an optimal answer set of  $P_c(\mathcal{K})$ . Then  $|(\omega_{M_o}^3)^{-1}(b)| = \mathcal{I}_c(\mathcal{K})$ .<sup>5</sup>*

**Example 16.** *Consider again the knowledge base  $\mathcal{K}_7$  from Example 10:*

$$\mathcal{K}_7 = \left\{ \underbrace{\begin{matrix} \phi_{1,1} & \phi_{1,2} \\ x & \wedge & y \end{matrix}}_{\alpha_1}, \underbrace{\begin{matrix} \phi_{2,1} & \phi_{2,2} \\ x & \vee & y \end{matrix}}_{\alpha_2}, \underbrace{\begin{matrix} \phi_3 \\ \neg & x \end{matrix}}_{\alpha_3} \right\}$$

We aim to construct  $P_c(\mathcal{K}_7)$ , i. e., an ASP encoding which allows us to retrieve  $\mathcal{I}_c(\mathcal{K}_7)$ . First, we represent the formulas and atoms in  $\mathcal{K}_7$  by applying (AC1) and (AC2):

$$\begin{aligned} \text{kbMember}(\alpha_1). & \quad \text{atom}(x). \\ \text{kbMember}(\alpha_2). & \quad \text{atom}(y). \\ \text{kbMember}(\alpha_3). & \end{aligned}$$

We use (AC3) to represent the only conjunction in  $\mathcal{K}_7$  ( $\alpha_1$ ) as  $\text{conjunction}(\alpha_1, \phi_{1,1}, \phi_{1,2})$ . In the same fashion, we apply (AC4) to  $\alpha_2$ , the only disjunction occurring in  $\mathcal{K}_7$ , and get

5. For any function  $\varphi : X \mapsto Y$  and  $y \in Y$  we define  $\varphi^{-1}(y) = \{x \in X \mid A(x) = y\}$

$\text{disjunction}(\alpha_2, \phi_{2,1}, \phi_{2,2})$ . Now, the only formula that is not represented in ASP yet is  $\alpha_3$ , so we represent it as  $\text{negation}(\alpha_3, \phi_3)$  via (AC5). To represent the subformulas which consist of individual atoms, we use (AC6):

$$\begin{aligned} \text{formulaIsAtom}(\phi_{1,1}, x). & \quad \text{formulaIsAtom}(\phi_{1,2}, y). \\ \text{formulaIsAtom}(\phi_{2,1}, x). & \quad \text{formulaIsAtom}(\phi_{2,2}, y). \\ \text{formulaIsAtom}(\phi_3, x). & \end{aligned}$$

The remainder of the logic program, i. e., (AC7)–(AC20), is static.

### 4.3 The Forgetting-Based Inconsistency Measure

The forgetting-based inconsistency measure  $\mathcal{I}_f$  is determined by the number of atom occurrences that need to be “forgotten” in order to make the knowledge base  $\mathcal{K}$  consistent. In the following, we construct an extended logic program  $P_f(\mathcal{K})$  which computes  $\mathcal{I}_f(\mathcal{K})$ . Just as described in the preceding section, we begin with the definition of facts. Again, every formula  $A \in \mathcal{K}$  is represented as  $\text{kbMember}(A)$  (AF1). Conjunctions (AF3), disjunctions (AF4), and negations (AF5) are also represented in exactly the same form as introduced in Section 4.2. On the other hand, formulas which consist of individual atoms have to be handled differently. Since the forgetting-based inconsistency measure does not only consider each atom, but each atom *occurrence*, we have to represent this in our ASP encoding. Therefore, we define the predicate  $\text{formulaIsAtomOcc}/3$  which contains the formula  $\Phi$ , the atom  $X$  it consists of, and the atom’s label  $l$ :

$$\text{formulaIsAtomOcc}(\Phi, X, l). \tag{AF2}$$

The representation of an atom occurrence can be extracted from the above rule as follows:

$$\text{atomOcc}(A, L) \text{ :- formulaIsAtomOcc}(\_, A, L). \tag{AF7}$$

Moreover, we can gain the representation of an atom by applying the following rule:

$$\text{atom}(A) \text{ :- atomOcc}(A, \_). \tag{AF8}$$

We model the truth values in the usual manner:

$$\text{tv}(t; f). \tag{AF6}$$

Next, we include a rule that ensures unique atom evaluation. With regard to  $\mathcal{I}_f$ , this means that an atom as a whole must be evaluated to either  $t$  or  $f$ , even though individual occurrences of that atom might be replaced by  $\top$  or  $\perp$ :

$$1\{\text{truthValue}(A, T) : \text{tv}(T)\}1 \text{ :- atom}(A). \tag{AF9}$$

The connective encodings simply model propositional entailment. Thus, the evaluation of a conjunction  $\Phi_c = \Psi_{c,1} \wedge \Psi_{c,2}$  is modeled as follows:

$$\begin{aligned} \text{truthValue}(F, t) \text{ :- conjunction}(F, G, H), \\ \text{truthValue}(G, t), \\ \text{truthValue}(H, t). \end{aligned} \tag{AF10}$$

$$\begin{aligned} \text{truthValue}(F, f) \text{ :- conjunction}(F, \_, \_), \\ \text{not truthValue}(F, t). \end{aligned} \tag{AF11}$$

In the same fashion, we define the evaluation of a disjunction  $\Phi_d = \Psi_{d,1} \vee \Psi_{d,2}$ :

$$\begin{aligned} \text{truthValue}(F, f) & :- \text{disjunction}(F, G, H), \\ & \quad \text{truthValue}(G, f), \\ & \quad \text{truthValue}(H, f). \end{aligned} \tag{AF12}$$

$$\begin{aligned} \text{truthValue}(F, t) & :- \text{disjunction}(F, -, -), \\ & \quad \text{not truthValue}(F, f). \end{aligned} \tag{AF13}$$

At last we define the evaluation of a negation  $\Phi_n = \neg\Psi_n$ :

$$\begin{aligned} \text{truthValue}(F, t) & :- \text{negation}(F, G), \\ & \quad \text{truthValue}(G, f). \end{aligned} \tag{AF14}$$

$$\begin{aligned} \text{truthValue}(F, f) & :- \text{negation}(F, G), \\ & \quad \text{truthValue}(G, t). \end{aligned} \tag{AF15}$$

If a formula  $\Phi_a$  consists of a single atom occurrence  $X^l$ , we need to consider that the forgetting operation might be applied. To model this, we first guess for each atom occurrence whether it is being forgotten:

$$\{\text{atomOccForgotten}(A, L)\} :- \text{atomOcc}(A, L). \tag{AF16}$$

If  $X^l$  is not forgotten, it needs to evaluate to the truth value of  $X$  itself:

$$\begin{aligned} \text{truthValue}(F, T) & :- \text{formulaIsAtomOcc}(F, A, L), \\ & \quad \text{truthValue}(A, T), \\ & \quad \text{not atomOccForgotten}(A, L). \end{aligned} \tag{AF17}$$

If, on the other hand,  $X^l$  needs to be forgotten, it means that the (sub)formula consisting of  $X^l$  evaluates to the opposite truth value:

$$\begin{aligned} \text{truthValue}(F, t) & :- \text{formulaIsAtomOcc}(F, A, L), \\ & \quad \text{truthValue}(A, f), \\ & \quad \text{atomOccForgotten}(A, L). \end{aligned} \tag{AF18}$$

$$\begin{aligned} \text{truthValue}(F, f) & :- \text{formulaIsAtomOcc}(F, A, L), \\ & \quad \text{truthValue}(A, t), \\ & \quad \text{atomOccForgotten}(A, L). \end{aligned} \tag{AF19}$$

All formulas  $A \in \mathcal{K}$  must evaluate to  $t$  after the forgetting operation is applied. Hence, we add the following integrity constraint which corresponds exactly to the one we used for the contension inconsistency measure (see (AC19) in Section 4.2):

$$:- \text{truthValue}(F, f), \text{kbMember}(F). \tag{AF20}$$

Lastly, we need to minimize the number of atom occurrences which are forgotten:

$$\#\text{minimize}\{1, A, L : \text{atomOccForgotten}(A, L)\}. \tag{AF21}$$

The union of the rules (AF1)–(AF21) defined above (see Encoding 8 in Appendix A.2.2 for an overview) constitute the extended logic program  $P_f(\mathcal{K})$ . With  $M$  being an answer set of  $P_f(\mathcal{K})$ , we denote the set of atom occurrences that are forgotten as

$$F_M = \{X^l \mid X \in \text{At}(\mathcal{K}), \text{atomOccForgotten}(X, l) \in M\}$$

**Theorem 8.** *Let  $M_o$  be an optimal answer set of  $P_f(\mathcal{K})$ . Then  $|F_{M_o}| = \mathcal{I}_f(\mathcal{K})$ .*

The proof of the above theorem is provided in Appendix A.2.2.

**Example 17.** *Consider again the knowledge base  $\mathcal{K}_7$  from Example 10:*

$$\mathcal{K}_7 = \left\{ \underbrace{\begin{array}{cc} \phi_{1,1} & \phi_{1,2} \\ x & \wedge & y \end{array}}_{\alpha_1}, \underbrace{\begin{array}{cc} \phi_{2,1} & \phi_{2,2} \\ x & \vee & y \end{array}}_{\alpha_2}, \underbrace{\begin{array}{c} \phi_3 \\ \neg & x \end{array}}_{\alpha_3} \right\}$$

We aim to construct  $P_f(\mathcal{K}_7)$ , i. e., an ASP encoding which allows us to retrieve  $\mathcal{I}_f(\mathcal{K}_7)$ . To begin with, the formulas  $\alpha_1$ ,  $\alpha_2$ , and  $\alpha_3$  are represented using `kbMember/1` (AF1), exactly as in Example 16. Likewise, the representation of the conjunction (AF3), disjunction (AF4), and negation (AF5) is the same as in Example 16. We label each atom occurrence to get  $\mathcal{K}_7^l = \{x^1 \wedge y^1, x^2 \vee y^2, \neg x^3\}$ , and represent each atom occurrence via (AF2):

$$\begin{aligned} \text{formulaIsAtomOcc}(\phi_{1,1}, x, 1). & \quad \text{formulaIsAtomOcc}(\phi_{1,2}, y, 1). \\ \text{formulaIsAtomOcc}(\phi_{2,1}, x, 2). & \quad \text{formulaIsAtomOcc}(\phi_{2,2}, y, 2). \\ \text{formulaIsAtomOcc}(\phi_3, x, 3). & \end{aligned}$$

The remainder of the logic program, i. e., (AF6)–(AF21), is static.

#### 4.4 The Hitting Set Inconsistency Measure

The hitting set inconsistency measure  $\mathcal{I}_h(\mathcal{K})$  is defined by the size of the minimal hitting set wrt. a knowledge base  $\mathcal{K}$ , subtracted by 1. The maximal size of such a hitting set is determined by the number of formulas in  $\mathcal{K}$ . Further,  $\omega_i$  refers to the  $i$ -th interpretation out of the  $|\mathcal{K}|$  possible interpretations we need to consider, assuming that the interpretations have an arbitrary, but fixed order. We construct an extended logic program  $P_h(\mathcal{K})$  which computes  $\mathcal{I}_h(\mathcal{K})$  as follows.

We begin by representing each element  $A \in \mathcal{K}$  (AH1) as well as each conjunction (AH4), disjunction (AH5), and negation (AH6) in the same manner described in the two preceding sections. Atoms are represented as `atom/1` (AH2), and formulas consisting of individual atoms are represented as `formulaIsAtom/2` (AH7). The two classical truth values are encoded in the same fashion as shown in Section 4.3 (AH9). In addition, the hitting set inconsistency measure requires the use of interpretations. Therefore, we add  $|\mathcal{K}|$  interpretations as follows:

$$\text{interpretation}(1..|\mathcal{K}|). \tag{AH3}$$

Again, we need to model that each atom takes exactly one (classical) truth value. However, as opposed to the two previously discussed measures, wrt.  $\mathcal{I}_h$  we need to take into

account the formulas' interpretations. Consequently, each atom requires a unique evaluation wrt. each interpretation. We model this by introducing the predicate `truthValueInt/3` which represents the truth value of an atom wrt. a specific interpretation:

$$1\{\text{truthValueInt}(A,I,T) : \text{tv}(T)\}1 \text{ :- atom}(A), \\ \text{interpretation}(I). \quad (\text{AH10})$$

The connective encodings for each (sub)formula follow classical propositional entailment. Hence, the rules which model conjunction, disjunction, and negation wrt.  $\mathcal{I}_h$  are essentially the same as those wrt.  $\mathcal{I}_f$ , but with an additional reference to an interpretation:

$$\text{truthValueInt}(F,I,t) \text{ :- conjunction}(F,G,H), \\ \text{interpretation}(I), \\ \text{truthValueInt}(G,I,t), \\ \text{truthValueInt}(H,I,t). \quad (\text{AH11})$$

$$\text{truthValueInt}(F,I,f) \text{ :- conjunction}(F,-,-), \\ \text{interpretation}(I), \\ \text{not truthValueInt}(F,I,t). \quad (\text{AH12})$$

$$\text{truthValueInt}(F,I,f) \text{ :- disjunction}(F,G,H), \\ \text{interpretation}(I), \\ \text{truthValueInt}(G,I,f), \\ \text{truthValueInt}(H,I,f). \quad (\text{AH13})$$

$$\text{truthValueInt}(F,I,t) \text{ :- disjunction}(F,-,-), \\ \text{interpretation}(I), \\ \text{not truthValueInt}(F,I,f). \quad (\text{AH14})$$

$$\text{truthValueInt}(F,I,t) \text{ :- negation}(F,G), \\ \text{truthValueInt}(G,I,f). \quad (\text{AH15})$$

$$\text{truthValueInt}(F,I,f) \text{ :- negation}(F,G), \\ \text{truthValueInt}(G,I,t). \quad (\text{AH16})$$

If a formula is composed of an individual atom, it needs to be assigned the same truth value as the atom itself wrt. an interpretation:

$$\text{truthValueInt}(F,I,T) \text{ :- formulaIsAtom}(F,G), \\ \text{truthValueInt}(G,I,T), \\ \text{interpretation}(I), \\ \text{tv}(T). \quad (\text{AH17})$$

In order to meet the definition of the hitting set measure, each formula  $A \in \mathcal{K}$  must evaluate to  $t$  wrt. at least one interpretation included in the hitting set. We model this by

using the predicate `truthValue/2` as follows:

$$\begin{aligned} \text{truthValue}(F,t) \text{ :- } & \text{truthValueInt}(F,I,t), \\ & \text{kbMember}(F), \\ & \text{interpretation}(I), \\ & \text{interpretationActive}(I). \end{aligned} \quad (\text{AH18})$$

The predicate `interpretationActive/1` serves the purpose of marking which interpretations are included in the final hitting set. Later on, we minimize the number of those “active interpretations”. However, to ensure that at least one (and at most  $|\mathcal{K}|$ ) interpretations are included in the hitting set, we require the following cardinality constraint:

$$1\{\text{interpretationActive}(I) : \text{interpretation}(I)\}|\mathcal{K}|. \quad (\text{AH8})$$

To avoid symmetries introduced by (AH8), we add another integrity constraint:

$$\text{:- interpretationActive}(I), 1 < I, \text{not interpretationActive}(I-1). \quad (\text{AH20})$$

If a formula  $A \in \mathcal{K}$  is not  $t$  wrt. any interpretation, it is  $f$ . In other words, if `truthValue(A,t)` is not included in the answer set, then  $A$  is false with regard to all interpretations:

$$\text{truthValue}(F,f) \text{ :- } \text{kbMember}(F), \text{not truthValue}(F,t). \quad (\text{AH19})$$

However, since every formula in  $\mathcal{K}$  must be satisfied by at least one interpretation, we need to include the following integrity constraint:

$$\text{:- truthValue}(F,f), \text{kbMember}(F). \quad (\text{AH21})$$

At last, we minimize the number of interpretations that are required to be “active” in order to satisfy all formulas in the given knowledge base:

$$\#\text{minimize}\{1, I : \text{interpretationActive}(I)\}. \quad (\text{AH22})$$

We define  $P_h(\mathcal{K})$  to be the extended logic program specified by the union of all rules defined in (AH1)–(AH22) (see Encoding 9 in Appendix A.2.3 for a complete list of all rules). Let  $M$  be an answer set of  $P_h(\mathcal{K})$ . We define the set of interpretations  $\omega_i$  represented in  $M$  as  $\Omega(M) = \{\omega_i \mid \text{interpretationActive}(\omega_i^{\text{asp}}) \in M\}$ , with  $\omega_i^{\text{asp}}$  being an ASP representation of  $\omega_i$ . The proof of the following theorem can be found in Appendix A.2.3.

**Theorem 9.** *Let  $M_o$  be an optimal answer set of  $P_h(\mathcal{K})$ . Then  $|\Omega(M_o)| - 1 = \mathcal{I}_h(\mathcal{K})$ . If no answer set of  $P_h(\mathcal{K})$  exists,  $\mathcal{I}_h(\mathcal{K}) = \infty$ .*

**Example 18.** *Consider again the knowledge base  $\mathcal{K}_7$  from Example 10:*

$$\mathcal{K}_7 = \left\{ \underbrace{\begin{array}{cc} \phi_{1,1} & \phi_{1,2} \\ x & \wedge & y \end{array}}_{\alpha_1}, \underbrace{\begin{array}{cc} \phi_{2,1} & \phi_{2,2} \\ x & \vee & y \end{array}}_{\alpha_2}, \underbrace{\begin{array}{c} \phi_3 \\ \neg & x \end{array}}_{\alpha_3} \right\}$$

Our aim is to construct  $P_h(\mathcal{K}_7)$ , i. e., an ASP encoding which allows us to retrieve the  $\mathcal{I}_h(\mathcal{K}_7)$ . The formulas in  $\mathcal{K}_7$  are represented the same way as in Example 16 and 17, using `kbMember/1` (AH1). Atoms are represented like in Example 16 as well (by using `atom/1`). Likewise, subformulas consisting of individual atoms (AH7) are modeled as in Example 16. Moreover, the conjunction (AH4), disjunction (AH5), and negation (AH6) are also defined as before. However, we now need to define  $|\mathcal{K}_7| = 3$  interpretations via (AH3):

$$\text{interpretation}(1..3).$$

Also, we ensure that at least one and at most  $|\mathcal{K}_7| = 3$  interpretations are included in the hitting set (AH8):

$$1\{\text{interpretationActive}(X) : \text{interpretation}(X)\}3.$$

The remainder of the logic program, i. e., (AH9)–(AH22), is static.

#### 4.5 The Max-Distance Inconsistency Measure

Recall that, given a knowledge base  $\mathcal{K}$ , the max-distance inconsistency measure  $\mathcal{I}_d^{\max}$  aims to find an interpretation that has a minimal maximum distance to the models of the formulas  $A \in \mathcal{K}$ . We construct an extended logic program  $P_d^{\max}(\mathcal{K})$  that calculates  $\mathcal{I}_d^{\max}(\mathcal{K})$  as described in the following.

Atoms are represented as `atom/1`, as described previously in the context of  $\mathcal{I}_c$ ,  $\mathcal{I}_f$ , and  $\mathcal{I}_h$  (see also (ADM2) in Encoding 10). Conjunctions, disjunctions, negations, and formulas consisting of single atoms are handled exactly as presented in the previous section (see (ADM4)–(ADM7) and (ADM11)–(ADM17) in Encoding 10). Truth values are again defined by `tv/1` (ADM9), and unique atom evaluation is also expressed as for  $\mathcal{I}_h$  (ADM10). Moreover, we use the predicate `interpretation/1` to represent  $|\mathcal{K}| + 1$  interpretations:

$$\text{interpretation}(0..|\mathcal{K}|). \tag{ADM3}$$

As opposed to the encoding of  $\mathcal{I}_h$ , which only requires  $|\mathcal{K}|$  interpretations, the encoding of  $\mathcal{I}_d^{\max}$  requires  $|\mathcal{K}| + 1$  interpretations. This is due to the fact that for  $\mathcal{I}_d^{\max}$  we first need to provide  $|\mathcal{K}|$  interpretations in order to ensure that each formula is satisfied by at least one interpretation (i. e., that we have at least one model for each formula). However, the interpretation with the smallest maximum distance to the formulas' models is not necessarily a model of one of the formulas itself. Thus, we need to provide one additional interpretation and end up with a total of  $|\mathcal{K}| + 1$  interpretations.

As stated previously, each  $A \in \mathcal{K}$  must be satisfied by at least one interpretation. To achieve this, we demand that the  $i$ -th formula in  $\mathcal{K}$  must be satisfied by the  $i$ -th interpretation. In order to model this in ASP, we first add a fact for every  $A \in \mathcal{K}$  which additionally incorporates an index  $i$ :

$$\text{kbMember}(A, i). \tag{ADM1}$$

Further, we need to add an integrity constraint which ultimately ensures that the formula with index  $i$  cannot be set to  $f$  (i. e., must evaluate to  $t$ ) under the  $i$ -th interpretation:

$$\begin{aligned} & :- \text{truthValueInt}(F, I, f), \\ & \text{kbMember}(F, I). \end{aligned} \tag{ADM18}$$

We know that the  $i$ -th formula in  $\mathcal{K}$  must be satisfied by the  $i$ -th interpretation. Thus, each interpretation  $\omega_i$  with  $i \in \{0, \dots, |\mathcal{K}| - 1\}$  is a model of at least one formula in  $\mathcal{K}$ . We are now looking for an interpretation  $\omega_{|\mathcal{K}|}$  with the smallest maximum distance to the models of each formula. In the worst case,  $\mathcal{I}_d^{\max}(\mathcal{K}) = |\text{At}(\mathcal{K})|$ . We represent each possible value as follows.

$$\text{diff}(1..X) \text{ :- } X = \#\text{count}\{\text{A: atom}(\text{A})\}. \quad (\text{ADM19})$$

By means of the `#count` aggregate above, we count the number of (ground) `atom/1` instances. We can now calculate the distance between  $\omega_{|\mathcal{K}|}$  and the models of each formula using another `#count` aggregate:

$$\begin{aligned} \text{d}(X) \text{ :- } & \text{diff}(X), \\ & \text{interpretation}(\text{I}), \\ & X \leq \#\text{count}\{\text{A: atom}(\text{A}), \\ & \quad \text{truthValueInt}(\text{A}, \text{I}, \text{T}), \\ & \quad \text{not truthValueInt}(\text{A}, |\mathcal{K}|, \text{T})\}. \end{aligned} \quad (\text{ADM8})$$

Essentially, `d/1` checks how many different values the distance between  $\omega_{|\mathcal{K}|}$  and each model corresponding to a formula can take. Since `diff/1` ensures that `d/1` is only derived once for a given value, we can count how many different values (i. e., different distances) are calculated. By minimizing this value, we indirectly get the minimal maximum distance:

$$\#\text{minimize}\{1, X : \text{d}(X)\}. \quad (\text{ADM20})$$

Let  $P_d^{\max}$  be the extended logic program specified by the union of rules (ADM1)–(ADM20) (see Encoding 10 in Appendix A.2.4 for a complete overview). Further, let  $D_M^{\max} = \{n \mid \text{d}(n) \in M, n \in \mathbb{N}\}$  wrt. an answer set  $M$ . The proof of the theorem below is given in Appendix A.2.4 as well.

**Theorem 10.** *Let  $M_o$  be an optimal answer set of  $P_d^{\max}(\mathcal{K})$ . Then  $|D_{M_o}^{\max}| = \mathcal{I}_d^{\max}(\mathcal{K})$ . If no answer set of  $P_d^{\max}(\mathcal{K})$  exists,  $\mathcal{I}_d^{\max}(\mathcal{K}) = \infty$ .*

**Example 19.** *Consider again the knowledge base  $\mathcal{K}_7$  from Example 10:*

$$\mathcal{K}_7 = \left\{ \underbrace{\left\{ \begin{array}{cc} \phi_{1,1} & \phi_{1,2} \\ x & \wedge & y \end{array} \right\}}_{\alpha_1}, \underbrace{\left\{ \begin{array}{cc} \phi_{2,1} & \phi_{2,2} \\ x & \vee & y \end{array} \right\}}_{\alpha_2}, \underbrace{\left\{ \begin{array}{c} \phi_3 \\ \neg & x \end{array} \right\}}_{\alpha_3} \right\}$$

*In this example, we aim to construct  $P_d^{\max}(\mathcal{K}_7)$ , i. e., an ASP encoding which allows us to retrieve  $\mathcal{I}_d^{\max}(\mathcal{K}_7)$ . Following (ADM1), we assign each formula in  $\mathcal{K}_7$  an index  $i \in \{0, \dots, |\mathcal{K}_7| - 1\}$ , i. e.,  $i \in \{0, \dots, 2\}$ :*

$$\text{kbMember}(\alpha_1, 0). \quad \text{kbMember}(\alpha_2, 1). \quad \text{kbMember}(\alpha_3, 2).$$

*Moreover, we need to define  $|\mathcal{K}_7| + 1$  (i. e., 4) interpretations (ADM3):*

$$\text{interpretation}(0..3).$$

*Atoms are represented using `atom/1` (ADM2), as we saw in Example 16 and 18. The same applies to the conjunction (ADM4), disjunction (ADM5), and negation (ADM6), as well as all (sub)formulas consisting of individual atoms (ADM7).*

*The remainder of the program, i. e., (ADM8)–(ADM20), is static.*

#### 4.6 The Sum-Distance Inconsistency Measure

The difference between the sum-distance inconsistency measure  $\mathcal{I}_d^\Sigma$  and the max-distance inconsistency measure  $\mathcal{I}_d^{\max}$  is the definition of the “optimal” distance. While  $\mathcal{I}_d^{\max}$  aims to find the smallest maximum distance,  $\mathcal{I}_d^\Sigma$  seeks to find the smallest sum of all distances between an interpretation  $\omega_{|\mathcal{K}|}$  and the models of all  $A \in \mathcal{K}$  (i. e.,  $\omega_i$  with  $0 \leq i < |\mathcal{K}|$ ). As opposed to  $\mathcal{I}_d^{\max}$ , we do not need to model the different distances explicitly—we can model the summation directly by means of the following minimize statement:

$$\begin{aligned} \# \text{minimize} \{ & 1, \text{A}, \text{I}: \text{atom}(\text{A}), \\ & \text{truthValueInt}(\text{A}, \text{I}, \text{T}), \\ & \text{not truthValueInt}(\text{A}, |\mathcal{K}|, \text{T}) \}. \end{aligned} \quad (\text{ADS8})$$

Intuitively, we consider all models  $\omega_i$  and sum up the number of differing atom valuations wrt.  $\omega_{|\mathcal{K}|}$ . (Each differing atom valuation per  $\omega_i$  is counted as 1.) Due to the minimization, we get the minimal sum of distances.

Let  $P_d^\Sigma$  be the extended logic program specified by the union of rules (ADS1)–(ADS18) (see Encoding 11 in Appendix A.2.5 for a complete overview). Further, let  $M$  be an answer set of  $P_d^\Sigma$ , and let  $\theta \in \{t, f\}$ . We define

$$D_M^\Sigma = \{ \text{diff}_M(X, i) \mid \text{truthValueInt}(X, i, \theta) \in M, \text{truthValueInt}(X, |\mathcal{K}|, \theta) \notin M \}.$$

The proof of the following theorem is provided in Appendix A.2.5.

**Theorem 11.** *Let  $M_o$  be an optimal answer set of  $P_d^\Sigma(\mathcal{K})$ . Then  $|D_{M_o}^\Sigma| = \mathcal{I}_d^\Sigma(\mathcal{K})$ . If no answer set of  $P_d^\Sigma(\mathcal{K})$  exists,  $\mathcal{I}_d^\Sigma(\mathcal{K}) = \infty$ .*

**Example 20.** *Consider again the knowledge base  $\mathcal{K}_7$  from Example 10:*

$$\mathcal{K}_7 = \left\{ \underbrace{\left\{ \begin{array}{cc} \phi_{1,1} & \phi_{1,2} \\ x & y \end{array} \right\}}_{\alpha_1}, \underbrace{\left\{ \begin{array}{cc} \phi_{2,1} & \phi_{2,2} \\ x & y \end{array} \right\}}_{\alpha_2}, \underbrace{\left\{ \begin{array}{c} \phi_3 \\ \neg x \end{array} \right\}}_{\alpha_3} \right\}$$

In this example, we aim to construct  $P_d^\Sigma(\mathcal{K}_7)$ , i. e., an ASP encoding which allows us to retrieve the sum-distance inconsistency value of  $\mathcal{K}_7$ . The formulas in  $\mathcal{K}_7$ , the atoms in  $\text{At}(\mathcal{K}_7)$ , the interpretations, the conjunction, the disjunction, the negation, and the subformulas consisting of individual atoms are encoded as described in (ADS1)–(ADS7) (see Encoding 11), which corresponds exactly to (ADM1)–(ADM7) wrt.  $\mathcal{I}_d^{\max}$  (see Example 19). Since  $|\mathcal{K}| = 3$ , the minimize statement (ADM8) is expressed as follows.

$$\begin{aligned} \# \text{minimize} \{ & 1, \text{A}, \text{I}: \text{atom}(\text{A}), \\ & \text{truthValueInt}(\text{A}, \text{I}, \text{T}), \\ & \text{not truthValueInt}(\text{A}, 3, \text{T}) \}. \end{aligned}$$

The remaining rules, i. e., (ADS9)–(ADS18), are static.

#### 4.7 The Hit-Distance Inconsistency Measure

As for the SAT-based approach, in order to model the hit-distance inconsistency measure  $\mathcal{I}_d^{\text{hit}}$  in ASP, we use the characterization that  $\mathcal{I}_d^{\text{hit}}(\mathcal{K})$  is equal to the minimal number of formulas that need to be removed from the given knowledge base  $\mathcal{K}$  in order to render it consistent. Utilizing the latter characterization, we construct an extended logic program  $P_d^{\text{hit}}$  which calculates  $\mathcal{I}_d^{\text{hit}}(\mathcal{K})$  as described below.

To start with, atoms, conjunctions, disjunctions, negations, and formulas consisting of single atoms are represented exactly as presented before wrt.  $\mathcal{I}_d^{\text{max}}$  and  $\mathcal{I}_d^{\Sigma}$  (see also (ADH2)–(ADH6) in Encoding 12). Likewise, we define the two truth values via  $\text{tv}/1$  (ADH7). Moreover, we need to ensure once again that each atom is assigned a distinct truth value. This is realized in the same manner as for  $\mathcal{I}_c$  and  $\mathcal{I}_f$  (ADH8).

As we do not need to take into account multiple interpretations at the same time, the assignment of truth values as well as the functionality of the connectives does not need to include a representation of the concept of interpretations. Thus, conjunction, disjunction, and negation are therefore modeled in the same manner as presented for  $\mathcal{I}_f$  (ADH9)–(ADH14). If a formula consists of an individual atom, its evaluation is modeled in the same way as presented for  $\mathcal{I}_c$  (ADH15).

As opposed to the encodings of the other two distance-based measures, elements of the given knowledge base  $\mathcal{K}$  do not require an index in the encoding of  $\mathcal{I}_d^{\text{hit}}$ . Thus, instead of representing  $A \in \mathcal{K}$  in combination with an index  $i$  as  $\text{kbMember}(A, i)$ , we simply use  $\text{kbMember}(A)$  (ADH1), like we saw previously in the encodings of  $\mathcal{I}_c$ ,  $\mathcal{I}_f$ , and  $\mathcal{I}_h$ .

The objective is to minimize the number of formulas in  $\mathcal{K}$  which are evaluated to  $f$ . To achieve this, we first define a rule that extracts the truth values of those formulas which are elements of  $\mathcal{K}$ .

$$\begin{aligned} \text{truthValueKbMember}(F, T) & :- \text{kbMember}(F), \\ & \text{tv}(T), \\ & \text{truthValue}(F, T). \end{aligned} \quad (\text{ADH16})$$

Now we can minimize the number of formulas which are evaluated to  $f$ —i.e., we minimize the number of formulas which need to be removed in order to make  $\mathcal{K}$  consistent.

$$\#\text{minimize}\{1, F : \text{truthValueKbMember}(F, f)\}. \quad (\text{ADH17})$$

Let  $P_d^{\text{hit}}$  be the extended logic program specified by the union of rules (ADH1)–(ADH17) (see Encoding 12 in Appendix A.2.6 for a complete overview). Further, with  $M$  being an answer set, we define  $K_M = \{A^{\text{asp}} \mid A \in \mathcal{K}, \text{truthValueKbMember}(A^{\text{asp}}, f) \in M\}$  with  $A^{\text{asp}}$  being an ASP representation of  $A$ . The proof of the theorem below can also be found in Appendix A.2.6.

**Theorem 12.** *Let  $M_o$  be an optimal answer set of  $P_d^{\text{hit}}(\mathcal{K})$ . Then  $|K_{M_o}| = \mathcal{I}_d^{\text{hit}}(\mathcal{K})$ .*

**Example 21.** *Consider again the knowledge base  $\mathcal{K}_7$  from Example 10:*

$$\mathcal{K}_7 = \left\{ \underbrace{\begin{array}{cc} \phi_{1,1} & \phi_{1,2} \\ x & \wedge & y \end{array}}_{\alpha_1}, \underbrace{\begin{array}{cc} \phi_{2,1} & \phi_{2,2} \\ x & \vee & y \end{array}}_{\alpha_2}, \underbrace{\neg x}_{\alpha_3} \right\}$$

In this example, we aim to construct  $P_d^{\text{hit}}(\mathcal{K}_7)$ , i. e., an ASP encoding which allows us to retrieve  $\mathcal{I}_d^{\text{hit}}(\mathcal{K}_7)$ . At first, we use `kbMember/1` (ADH1) to represent the formulas  $\alpha_1$ ,  $\alpha_2$ , and  $\alpha_3$  in the same fashion as wrt.  $\mathcal{I}_c$ ,  $\mathcal{I}_f$ , and  $\mathcal{I}_h$  (see Example 16). Further, the atoms, the conjunction, the disjunction, the negation, and the subformulas consisting of individual atoms (see (ADH2)–(ADH6)) are represented exactly as wrt.  $\mathcal{I}_c$ ,  $\mathcal{I}_h$ ,  $\mathcal{I}_d^{\text{max}}$  and  $\mathcal{I}_d^\Sigma$  (see again Example 16). The remaining rules, i. e., (ADS7)–(ADS17), are static.

## 5. Evaluation

The aim of our experimental evaluation is to compare the proposed SAT-based and ASP-based approaches to each other wrt. all six inconsistency measures considered in this work. We additionally compare our methods to naive baseline algorithms (see Section 5.2) in terms of runtime. Although we can expect both the SAT-based and the ASP-based approaches to be superior to the baseline methods (regarding ASP this has already been demonstrated for  $\mathcal{I}_c$ ,  $\mathcal{I}_f$ , and  $\mathcal{I}_h$  (Kuhlmann & Thimm, 2021), and regarding SAT this has been shown for  $\mathcal{I}_c$  in (Kuhlmann et al., 2022)), we still draw this comparison in order to concretely quantify this assumption. Besides, the baseline algorithms are, to the best of our knowledge, the only existing implementations for the inconsistency measures in question. However, as the result of comparing the SAT and ASP methods is far less predictable (both SAT and ASP are established formalisms for dealing with problems on the first level of the polynomial hierarchy), we examine the two approaches more closely.

### 5.1 Experimental Setup

Due to the unavailability of standard benchmarking sets for inconsistency measurement, we compiled a series of different data sets, either by artificially generating knowledge bases or by translating benchmark data from other fields. For all these data sets the goal was to have benchmark instances that are (mostly) inconsistent and feature a structure that can be expected in real-world applications. We compiled five data sets in total that we briefly describe in the following. Note that the SRS data set and the ML data set were already used in (Kuhlmann et al., 2022); the remaining three data sets are novel.<sup>6</sup>

**SRS** Data set SRS consists of knowledge bases randomly generated using the *SyntacticRandomSampler* provided by *TweetyProject*<sup>7</sup>. This generator randomly generates a single propositional formula  $\Phi$  as follows. Given input parameters  $p_d, p_c, p_n \in [0, 1]$  with  $p_d + p_c + p_n \leq 1$ , it is randomly chosen whether  $\Phi$  is a disjunction  $\Phi = \Psi_1 \vee \Psi_2$  (with probability  $p_d$ ), a conjunction  $\Phi = \Psi_1 \wedge \Psi_2$  (with probability  $p_c$ ), a negation  $\Phi = \neg\Psi_1$  (with probability  $p_n$ ), or a proposition  $\Phi = X$  (with probability  $1 - p_d - p_c - p_n$ ; the exact proposition is chosen uniformly at random from some given set of propositions). If  $\Phi$  is not a proposition then this process is repeated for the subformula  $\Psi_1$  (and additionally  $\Psi_2$  in the first two cases), where the probabilities for the first three cases are multiplied by some discount parameter  $d \in (0, 1)$  (in order to ensure that the process will, in practice, terminate at some point).

6. Download (all data sets): <https://fernuni-hagen.sciebo.de/s/sML2faFBiCib1nm>

7. <http://tweetyproject.org/api/1.19/org/tweetyproject/logics/pl/util/SyntacticRandomSampler.html>

Signature size	Formulas / KB	Mean sig. size / formula	Max. sig. size / formula	Mean #atom occ. / formula	Max #atom occ. / formula
3	5–15	1.69	3	2.23	6
5	15–25	2.32	5	3.11	11
10	15–25	2.70	8	3.13	10
15	15–25	2.80	9	3.11	11
15	25–50	2.81	9	3.11	11
20	25–50	2.88	11	3.11	11
25	25–50	2.91	12	3.09	13
25	50–100	2.92	10	3.11	11
30	50–100	2.94	11	3.10	13

Table 3: Overview of the sets of knowledge bases making up data set SRS. Each row represents a set of 200 knowledge bases. The columns describe (from left to right) the signature size, the number of formulas per knowledge base, the mean (resp. maximum) signature size per formula, and the mean (resp. maximum) number of atom occurrences per formula.

We created a total of 1800 knowledge bases using the *SyntacticRandomSampler*. In order to include instances of varying complexity, we used different parameter settings to create nine sets which each consist of 200 knowledge bases. To be specific, we used different signature sizes and different numbers of formulas per knowledge base. The parameters  $p_d$ ,  $p_c$ , and  $p_n$  were always set to 0.3. A more detailed overview of the composition of the SRS data set is given in Table 3.

**ML** The ML data set consists of knowledge bases containing formulas learnt from machine learning data (Thimm & Rienstra, 2020). More precisely, the data set “Animals with attributes”<sup>8</sup> describes 50 animals, e.g. ox, mouse, or dolphin, using 85 binary attributes such as “swims”, “black”, and “arctic”. We used the Apriori algorithm (Agrawal & Srikant, 1994) to mine association rules from this data set for a given minimal confidence value  $c$  and minimal support value  $s$ . All these rules were then interpreted as propositional logic implications. Finally, we selected one animal at random and added all its attributes as facts (thus making the knowledge base inconsistent as even rules with low confidence values were interpreted as strict implications). We set

$$c \in \{0.6, 0.65, 0.70, 0.75, 0.8, 0.85, 0.90, 0.95\},$$

$$s \in \{0.6, 0.65, 0.70, 0.75, 0.8, 0.85, 0.90, 0.95\},$$

and allowed maximally 4 literals per rule. The final data set contains 1920 instances.

**ARG** Data set ARG consists of a total of 326 knowledge bases extracted from benchmark data of the International Competition on Computational Models of Argumentation

8. <http://attributes.kyb.tuebingen.mpg.de>

Name	#KBs	#Formulas		Sig. size		#Con./formula	
		Mean	Std.	Mean	Std.	Mean	Std.
SRS	1800	36	24	16	9	3.1	2.0
ML	1920	7506	17984	76	0	5.5	1.4
ARG	326	989	2133	827	1781	198.3	455.1
SC	100	14254	12788	2559	2616	4.5	3.9
LP	45	9150	4683	10097	4834	11.0	905.1

Table 4: Overview of some properties of the data sets used in the evaluation, which include (from left to right) the total number of knowledge bases, the number of formulas per knowledge base (mean and standard deviation), the signature size per KB (mean and standard deviation), and the number of connectives per formula (mean and standard deviation).

2019 (ICCMA'19)<sup>9</sup>. An abstract argumentation framework (Dung, 1995) is a directed graph  $F = (A, R)$  where  $A$  is a set of arguments, and  $R$  models a conflict relation between such arguments. A computational task here is to find a *stable extension*, i. e., a set  $E \subseteq A$  with  $(a, b) \notin R$  for all  $a, b \in E$  and  $(a, c) \in R$  for all  $c \in A \setminus E$  and some  $a \in E$ . For each instance from ICCMA'19, we encoded the instance and the problem of finding such a stable extension via the approach from (Besnard, Doutre, & Herzig, 2014) and, additionally, added constraints to ensure that 20% of randomly selected arguments have to be contained in  $E$ . Note that the latter constraints usually make the knowledge base inconsistent.

**SC** Data set SC consists of the 100 smallest (in terms of file size) instances of the benchmark data set from the main track of the SAT competition 2020.<sup>10</sup>

**LP** Data set LP is generated from benchmark data for answer set programming. More specifically, we used the problems *15-Puzzle* (11 instances), *Hamiltonian Cycle* (20 instances), and *Labyrinth* (14 instances). available from the Asparagus website<sup>11</sup>. Each problem specification and a single instance was first grounded and then converted to a propositional logic knowledge base using Clark completion (Clark, 1978), i. e., a procedure that, basically, translates rules  $L_0 :- L_1, \dots, L_n$ . with identical head to a single formula  $L_0 \leftrightarrow \bigwedge L_1 \vee \dots \vee \bigwedge L_n$  (where each  $L_i, i \in \{1, \dots, n\}$ , may be a set of literals). Both classical negation and default negation in the rules are converted to logical negation in the resulting propositional formula.<sup>12</sup>

An overview of some basic statistics (including the number of instances, the mean number of formulas, the mean signature size, and the mean number of connectives per formula) of all five data sets is provided in Table 4.

9. <http://argumentationcompetition.org/2019/>

10. <https://satcompetition.github.io/2020/downloads.html>

11. <https://asparagus.cs.uni-potsdam.de>

12. Note that this procedure does not maintain the semantics of the original logic program. However, since we are not interested in computing solutions to the original problems but only require benchmarks that feature a similar structure as application scenarios, we believe that this does not pose an issue.

Both the SAT-based and the ASP-based approaches are implemented in C++. The SAT solver we use is CaDiCal sc2021<sup>13</sup> (Biere, Fazekas, Fleury, & Heisinger, 2020), and the ASP solver we use is Clingo 5.5.1<sup>14</sup> (Gebser, Kaminski, Kaufmann, & Schaub, 2019). For the computation of cardinality constraints in SAT we use sequential counter encoding (Sinz, 2005), and for transforming formulas to CNF we use Tseitin’s method (Tseitin, 1968).

Due to the significant amount of time required to complete the evaluation, two different servers were used to facilitate the evaluation of multiple data sets in parallel. The experiments regarding the SRS data set were run on a computer with 128 GB RAM and an Intel Xeon E5-2690 CPU which has a basic clock frequency of 2.90 GHz. The other data sets (ML, ARG, SC, and LP) were evaluated on server instances with 32 GB RAM and an Intel Xeon Platinum 8260M CPU with a basic clock frequency of 2.40 GHz.

## 5.2 Baseline Approaches

We will compare our two families of approaches against existing baseline implementations of the respective measures from *TweetyProject*. These are quite simple computational approaches mainly relying on brute force methods which are written in Java. We briefly explain these baseline approaches as follows:

- *The contension inconsistency measure.*<sup>15</sup> On input  $\mathcal{K}$  this approach first converts  $\mathcal{K}$  into conjunctive normal form  $\mathcal{K}'$  (exploiting the fact that the contension inconsistency value does not change when applying syntactic transformations). Then a SAT solver is used to check whether  $\mathcal{K}'$  is satisfiable. If it is, the inconsistency value is 0 and the algorithm terminates. Otherwise, for some atom  $X$  appearing in  $\mathcal{K}'$ , all clauses containing  $X$  or  $\neg X$  are removed (this is equivalent to setting the truth value of  $X$  to  $b$ ). If the resulting knowledge base is satisfiable, the inconsistency value 1 is returned. Otherwise, every other atom appearing in  $\mathcal{K}'$  is tested in the same way. If none of the resulting knowledge bases is satisfiable, the procedure is continued with all pairs of atoms appearing in  $\mathcal{K}'$ , and so forth.
- *The forgetting-based inconsistency measure.*<sup>16</sup> This approach works similar as the approach for the contension inconsistency measure. First, satisfiability of the input  $\mathcal{K}$  is tested using transformation into conjunctive normal form and application of a SAT solver. Then, all possible substitutions of forgetting a single atom occurrence are applied on the input and it is checked whether the resulting knowledge base is consistent (again by using a SAT solver). Then, all pairs of possible substitutions are tested, and so on.
- *The hitting set inconsistency measure.*<sup>17</sup> This approach first checks whether there is a single interpretation that satisfies the input  $\mathcal{K}$  (by exhaustive enumeration) and returns the inconsistency value 0 in that case. Then, we check whether there is any pair of interpretations (by exhaustively enumerating all possible combinations) such

13. <https://github.com/arminbiere/cadical>

14. <https://potassco.org/clingo/>

15. [http://tweetyproject.org/r/?r=base\\_contension](http://tweetyproject.org/r/?r=base_contension)

16. [http://tweetyproject.org/r/?r=base\\_forget](http://tweetyproject.org/r/?r=base_forget)

17. [http://tweetyproject.org/r/?r=base\\_hs](http://tweetyproject.org/r/?r=base_hs)

that every formula is satisfied by at least one of them. If such a pair can be found, the inconsistency value 1 is returned. Otherwise, the process is repeated with three interpretations, and so on.

- *The max-distance inconsistency measure.*<sup>18</sup> This approach simply iterates over all possible interpretations and calculates the distances to each formula (i. e., the minimal distance to all models of the formula). The maximum value found is returned.
- *The sum-distance inconsistency measure.*<sup>19</sup> This approach is analogous to the max-distance inconsistency measure, but sums are calculated rather than maxima.
- *The hit-distance inconsistency measure.*<sup>20</sup> This approach is analogous to the max-distance inconsistency measure, but only the number of unsatisfied formulas is checked.

It is apparent that these approaches solve the corresponding problems in quite a naive fashion, but they are still able to produce results for small (toy) examples.

### 5.3 Results

In this section, we present the results of our experimental analysis. To begin with, we compare the overall runtime of the SAT-based, ASP-based, and baseline approaches. Afterwards, we analyze more specific aspects, namely the runtime composition of the SAT- and ASP-based methods, the choice of search strategy used in the SAT-based approaches, the use of MaxSAT, and a comparison of previous versions of ASP encodings for  $\mathcal{I}_c$  with the current version.

#### 5.3.1 OVERALL RUNTIME COMPARISON

Our first goal of the evaluation is to obtain an overview of the runtimes of the different approaches. To get a most thorough overview, we consider all three approaches for each of the six measures, and use all five previously described data sets (see Section 5.1).

A timeout was set to 600 seconds (i. e., 10 minutes) for instances from the SRS, ML, and ARG data sets. For instances from the SC and LP data sets, which are overall more challenging compared to those instances from the previously mentioned data sets, we increased the timeout to 5000 seconds (i. e., 83.3 minutes). This corresponds to the time limit used in the SAT competition 2020, from which the instances in the SC data set were selected.

**SRS** First we consider the SRS data set, which has been used previously in (Kuhlmann & Thimm, 2021; Kuhlmann et al., 2022). Table 5 provides the number of solved instances as well as the cumulative runtime of each approach wrt. each measure and each data set. (For a more detailed overview, see also the plots in B.1 and B.3.) Looking at the SRS part of Table 5, a first noticeable observation is that, overall, the ASP-based approaches perform best, while the naive approaches perform, as expected, poorest. More precisely, the ASP-based approaches only time out in 175 out of 10,800 cases (1.62%). To set this into perspective: the SAT approaches time out in 1552 cases (14.37%), and the naive ones in

18. [http://tweetyproject.org/r/?r=base\\_dmax](http://tweetyproject.org/r/?r=base_dmax)

19. [http://tweetyproject.org/r/?r=base\\_dsum](http://tweetyproject.org/r/?r=base_dsum)

20. [http://tweetyproject.org/r/?r=base\\_dhit](http://tweetyproject.org/r/?r=base_dhit)

5703 cases (52.81 %). The only exception to this pattern is the sum-distance inconsistency measure ( $\mathcal{I}_d^\Sigma$ )—here, the naive approach outperforms the SAT-based one. This might be due to the fact that the search space for  $\mathcal{I}_d^\Sigma$  is quite large compared to most other measures<sup>21</sup> (see Table 2). As the maximum non- $\infty$  value for  $\mathcal{I}_d^\Sigma$  is  $|\text{At}(\mathcal{K})| \times |\mathcal{K}|$ , the search procedure consists of  $\log(|\text{At}(\mathcal{K})| \times |\mathcal{K}| + 1)$  steps, and we also need to calculate a new cardinality constraint exactly as many times. This results in the SAT method for  $\mathcal{I}_d^\Sigma$  being slower than those for the other two distance-based measures, while the respective naive methods are about equally fast for all three measures.

From concrete comparisons between each pair of approaches (see Appendix B.3 for some visualizations) we learn that all ASP-based approaches outperform the SAT-based ones in almost all cases (see Figure 17). The only exception worth mentioning are a few knowledge bases for which the SAT approach for  $\mathcal{I}_h$  was faster. With regard to the comparison of the ASP-based and naive methods we observe a similar pattern—again, the ASP methods are faster than the naive ones in most cases (see Figure 18). There are again a few instances for which  $\mathcal{I}_h$  is computed faster by the naive method. Overall, these results reinforce the perception conveyed by Table 5 that ASP performs strongest across this data set.

The comparison between the SAT-based and the naive approaches shows that the naive approach to solving  $\mathcal{I}_d^\Sigma$  is faster than its SAT-based counterpart in the majority of cases (see Figure 19), and in addition, there are numerous instances for which the SAT-based method could not return a result within the time limit, while the naive method could. Wrt. all other measures (with the exception of  $\mathcal{I}_c$ ), there exist some instances for which the naive approach was faster as well. However, in total, the SAT-based approaches perform superior to the naive ones.

**ML** A first glance at the ML part of Table 5 immediately reveals that this data set appears to be more challenging than the SRS data set, as there are significantly more timeouts across the three approaches wrt. all six inconsistency measures. To be precise, the ASP approaches time out in 3666 out of 11,520 cases (31.82 %), the SAT approaches time out in 7760 cases (67.36 %), and the naive approaches in 9831 cases (85.34 %). Thus, overall, the relationships between the approaches are again similar to those for the SRS data set—the ASP approaches achieve the strongest results, the naive ones, as expected, the weakest.

Particularly noticeable are the results regarding the three distance-based measures. The naive methods did not solve a single instance for any one of the three measures within the time limit. The SAT approaches did not yield any results for  $\mathcal{I}_d^{\max}$  and  $\mathcal{I}_d^\Sigma$  either; only for  $\mathcal{I}_d^{\text{hit}}$ , a total of 770 instances could be solved. The ASP approach was the only one that could actually solve a number of instances for both  $\mathcal{I}_d^{\max}$  and  $\mathcal{I}_d^\Sigma$ .

Another noticeable observation is that the naive approach for  $\mathcal{I}_f$  slightly outperforms its SAT-based equivalent—which was clearly not the case wrt. the SRS data set. This is most likely due to the fact that the SRS data set contains significantly fewer instances that are consistent, i.e., have an inconsistency degree of 0, than the ML data set (19 vs. 548 instances; see Figures 12 and 13 for more information). Since the naive approach first checks whether

21. The only other measure that could possibly result in a higher value wrt. a given knowledge base  $\mathcal{K}$  is  $\mathcal{I}_f$ . However, since  $\mathcal{I}_f$  depends on the number of atom occurrences in  $\mathcal{K}$ , and  $\mathcal{I}_d^\Sigma$  depends on the number of formulas in  $\mathcal{K}$  as well as the signature size  $\text{At}(\mathcal{K})$ , a general assessment of which measure has a larger search space is not possible.

		Naive		SAT		ASP	
		#solved	CRT (s)	#solved	CRT (s)	#solved	CRT (s)
SRS (1800)	$\mathcal{I}_c$	1162	49237	<b>1800</b>	793	<b>1800</b>	48
	$\mathcal{I}_f$	271	22230	1580	103172	<b>1667</b>	16482
	$\mathcal{I}_h$	647	16176	<b>1792</b>	57644	<b>1792</b>	1612
	$\mathcal{I}_d^{\max}$	1006	27156	1618	97589	<b>1800</b>	3919
	$\mathcal{I}_d^{\Sigma}$	1006	27188	658	67065	<b>1766</b>	11785
	$\mathcal{I}_d^{\text{hit}}$	1005	26557	<b>1800</b>	4585	<b>1800</b>	63
ML (1920)	$\mathcal{I}_c$	1134	19075	1763	99260	<b>1920</b>	9013
	$\mathcal{I}_f$	555	3373	523	33994	<b>1467</b>	67220
	$\mathcal{I}_h$	0	–	704	76745	<b>1163</b>	58203
	$\mathcal{I}_d^{\max}$	0	–	0	–	<b>733</b>	89206
	$\mathcal{I}_d^{\Sigma}$	0	–	0	–	<b>651</b>	18818
	$\mathcal{I}_d^{\text{hit}}$	0	–	770	38730	<b>1920</b>	7803
ARG (326)	$\mathcal{I}_c$	103	3045	196	10252	<b>246</b>	9071
	$\mathcal{I}_f$	70	1175	68	6395	<b>143</b>	1908
	$\mathcal{I}_h$	47	1412	157	8539	<b>218</b>	12975
	$\mathcal{I}_d^{\max}$	23	762	85	6222	<b>145</b>	8408
	$\mathcal{I}_d^{\Sigma}$	23	773	30	3694	<b>133</b>	2644
	$\mathcal{I}_d^{\text{hit}}$	23	772	<b>164</b>	3340	157	1382
SC (100)	$\mathcal{I}_c$	11	15481	<b>13</b>	21244	7	8619
	$\mathcal{I}_f$	<b>11</b>	11859	1	111	4	285
	$\mathcal{I}_h$	0	–	2	153	<b>3</b>	1461
	$\mathcal{I}_d^{\max}$	0	–	1	313	<b>2</b>	1172
	$\mathcal{I}_d^{\Sigma}$	0	–	0	–	<b>3</b>	322
	$\mathcal{I}_d^{\text{hit}}$	0	–	3	2513	<b>5</b>	2347
LP (45)	$\mathcal{I}_c$	0	–	0	–	<b>45</b>	15869
	$\mathcal{I}_f$	0	–	0	–	<b>11</b>	36315
	$\mathcal{I}_h$	0	–	0	–	0	–
	$\mathcal{I}_d^{\max}$	0	–	0	–	0	–
	$\mathcal{I}_d^{\Sigma}$	0	–	0	–	0	–
	$\mathcal{I}_d^{\text{hit}}$	0	–	0	–	<b>31</b>	7983

Table 5: Overview of the number of solved instances and cumulative runtime (CRT) per data set, measure, and approach. The total number of instances contained in each data set is provided in braces after the data set name. Timeout: 600 seconds (10 minutes) for SRS, ML, and ARG, and 5000 seconds for SC and LP.

the given knowledge base is consistent (by means of a SAT solver), consistent instances can be solved rather fast, in particular when compared to the SAT-based approach, where  $\log(|\text{Occ}(\mathcal{K})| + 1)$  calls to a SAT solver are required. Besides, the SAT approach requires the computation of a SAT encoding for  $\text{UPPER}_{\mathcal{I}_f}$  (including the computation of an at-most- $k$  constraint in each step), which leads to an additional overhead. The same applies to the naive approach for  $\mathcal{I}_c$ . Furthermore, the naive approaches for  $\mathcal{I}_c$  and  $\mathcal{I}_f$  might also yield a result rather quickly if the corresponding inconsistency value is not 0, but still very low, as they both search for the correct value from low to high in a linear manner. Deeper analysis revealed that the naive approach for  $\mathcal{I}_c$  was faster than the corresponding SAT-based approach in a total of 895 cases. In 507 out of these 895 cases, the contention inconsistency degree of the respective instances was in fact 0. Moreover, the inconsistency degree was 1 in 340 cases, and 2 in 48 cases. Whenever the inconsistency degree was  $> 2$ , the naive approach for  $\mathcal{I}_c$  was slower than the SAT-based approach. Regarding  $\mathcal{I}_f$ , the naive approach was also faster than the SAT-based one in 507 cases in which the inconsistency value is 0. Further, there are 19 instances  $\mathcal{K}$  with  $\mathcal{I}_f(\mathcal{K}) = 1$  which could be solved faster by the naive approach; there are no instances with  $\mathcal{I}_f(\mathcal{K}) > 1$  where this is the case.

Wrt. ASP, the picture looks a bit different: the naive method was faster than the ASP-based method in only 65 cases for  $\mathcal{I}_c$ , and 57 cases for  $\mathcal{I}_f$ . In 45 of the cases regarding  $\mathcal{I}_c$ , the inconsistency value is 0, and in the remaining 20 cases it is 1. With regard to  $\mathcal{I}_f$ , in 56 out of the 57 cases in which the naive approach was faster, the corresponding inconsistency value is 0, and in the single remaining case it is 1.

A comparison between the runtimes of the ASP-based and the SAT-based approaches (see Figure 20) shows that the former are faster wrt. all instances. When comparing the ASP approaches with the naive ones (see also Figure 21), we can see that there is a number of instances which could be solved by the naive method for  $\mathcal{I}_f$ , while the ASP-based one resulted in a timeout. Yet overall, the ASP methods perform significantly superior compared to the naive ones. The comparison between the SAT approaches and the naive ones looks similar, however the naive approaches overall perform a bit stronger than when compared to the ASP-based versions. However, although the SAT-based methods are outperformed by the naive ones in a total of 1465 cases, the former still solve 1799 instances more than the latter (across all inconsistency measures).

**ARG** From the ARG part of Table 5 we can observe that all three approaches wrt. all six inconsistency measures can solve some instances, but none can solve all instances from the data set. In total, the ASP-based approaches time out in 914 out of 1956 cases (46.73%), the SAT-based approaches time out in 1256 cases (64.21%), and the naive ones in 1667 cases (85.22%). Once again, the ASP approaches overall clearly perform strongest, and the naive ones weakest. However, similarly to what we observed for the ML data set, the naive approach for  $\mathcal{I}_f$  performs a bit stronger than the SAT-based one. Another exception to the overall pattern is that the SAT-based and the ASP-based approaches for  $\mathcal{I}_d^{\text{hit}}$  solve a quite similar number of instances; in fact, the ASP-based approach produces a few more timeouts than the SAT-based one. This has not been the case for any measure with regard to the SRS or ML data set. However, it should also be mentioned that the ASP approach is clearly faster in the majority of cases (see also the scatter plot in Figure 23f). Further, wrt.  $\mathcal{I}_d^{\Sigma}$ , the naive method and the SAT-based method perform similarly. Nevertheless,

although the SAT-based approach is a bit slower on average, the naive one still produces 7 more timeouts.

The results of a comparison regarding the individual runtimes per knowledge base between the ASP-based, SAT-based, and naive approaches underline the proposition that the ASP-based approaches perform superior to the SAT-based and naive ones on the ARG data set. The comparison to the SAT-based approaches (see Figure 23) shows that ASP is faster in all cases, except for a few instances wrt.  $\mathcal{I}_c$  and  $\mathcal{I}_d^{\text{hit}}$ . Compared to the naive approaches (see Figure 24), the ASP-based ones were faster in all cases, except for one instance in which the naive version was slightly faster wrt.  $\mathcal{I}_f$ .

The comparison between the SAT-based and the naive approaches (see Figure 25) shows that the naive ones outperformed the SAT-based ones in multiple cases. More precisely, wrt. all inconsistency measures (except  $\mathcal{I}_d^{\text{hit}}$ ) there are instances for which the naive approach was faster. In particular, there are some instances for which the SAT-based versions of  $\mathcal{I}_f$  or  $\mathcal{I}_d^\Sigma$  time out, but their naive counterparts do not. At large, the SAT approaches still perform superior to the naive ones. This becomes particularly clear on inspection of the number of instances that the SAT-based methods can solve within the time limit, but the naive versions cannot.

**SC** To begin with, we would like to recall the fact that we increased the timeout for the data sets SC and LP to 5000 seconds (83.3 minutes). However, even with this increased timeout, the data set still poses a great challenge for all three approaches wrt. all six inconsistency measures. On the other hand, this is not entirely surprising—in a SAT competition, the problem at hand is to decide whether a given knowledge base is satisfiable or not (which is NP-complete), while the problem we are dealing with in this work is to determine an inconsistency value (which is in  $\text{FP}^{\text{NP}[\log n]}$  wrt.  $\mathcal{I}_h$  and  $\mathcal{I}_d^{\text{max}}$ , and proven to be  $\text{FP}^{\text{NP}[\log n]}$ -complete wrt.  $\mathcal{I}_c$ ,  $\mathcal{I}_f$ ,  $\mathcal{I}_d^\Sigma$ , and  $\mathcal{I}_d^{\text{hit}}$  (Thimm & Wallner, 2019)).

In summary, all three approaches time out in most cases—the ASP approaches in 576 out of 600 cases, the SAT approaches in 580, and the naive approaches in 578 cases (see Table 5 for more details). It should be noted that wrt. the naive methods, only the those for  $\mathcal{I}_c$  and  $\mathcal{I}_f$  could solve any instances whatsoever. The reason for this behavior is most likely the previously mentioned fact that the naive approaches for both  $\mathcal{I}_c$  and  $\mathcal{I}_f$  use a SAT solver to check whether the given knowledge base is consistent as a first step, which allows for retrieving the inconsistency value fast in case it is 0. If the inconsistency value is 1, both approaches can still yield a result rather fast if the “correct” atom (in the case of  $\mathcal{I}_c$ ) or atom occurrence (in the case of  $\mathcal{I}_f$ ) is removed at an early stage. Our analysis (see Figure 15) shows that, in fact, all instances from the SC data set that could be solved have either inconsistency value 0 or 1 (wrt. all measures).<sup>22</sup> Furthermore, all knowledge bases in the data set at hand are already provided in CNF—hence, as opposed to all other data sets considered in this work, the formulas do not require any additional transformation steps for the SAT approaches.

**LP** At last we inspect the runtime results regarding the LP data set. Again, the timeout is set to 5000s. The results (see Table 5) show that neither one of the SAT-based or naive approaches could solve any instance for any of the six measures within the time limit. As

22. This is a quite reasonable result, as the knowledge bases in this data set were designed for a SAT competition. Hence, it is supposed to be challenging to decide whether they are unsatisfiable at all.

for the ASP-based approaches, some instances could be solved wrt.  $\mathcal{I}_c$ ,  $\mathcal{I}_f$ , and  $\mathcal{I}_d^{\text{hit}}$ , but wrt.  $\mathcal{I}_h$ ,  $\mathcal{I}_d^{\text{max}}$ , and  $\mathcal{I}_d^{\Sigma}$ , again, not a single inconsistency value could be determined within the time limit. To be exact, we can observe that the ASP-based method for  $\mathcal{I}_c$  could, in fact, solve all 45 instances, while the one for  $\mathcal{I}_f$  could solve 11 instances, but timed out 34 times. The ASP method for  $\mathcal{I}_d^{\text{hit}}$  solved 31 instances, and consequently timed out in 14 cases. A closer look at the runtimes (see also Figure 7) reveals that, although the ASP-based approaches for  $\mathcal{I}_c$ ,  $\mathcal{I}_f$ , and  $\mathcal{I}_d^{\text{hit}}$  could solve a number of instances, they did require a rather large amount of time in a lot of cases, in particular wrt.  $\mathcal{I}_f$ . More precisely, the ASP-based method for  $\mathcal{I}_c$  still solved 16 instances between 10 and 100s, however, the remaining 29 instances took between 100 and 1000s, with 12 of them exceeding the 500s mark. With regard to the ASP method for  $\mathcal{I}_f$ , all instances (except one) even require  $> 1700$ s. These comparatively long runtimes might be an indicator of why none of the SAT-based or naive approaches could solve even a single instance.

### 5.3.2 RUNTIME COMPOSITION

The experiments illustrated in the preceding section revealed that the ASP-based approaches overall outperform their SAT-based and naive counterparts. Whereas this is not surprising wrt. the naive methods, it is not immediately clear why the SAT methods are outperformed as well. The objective of this section is to explore the reasons for this in more detail by investigating how the runtimes of both the SAT and the ASP methods are composed on average.

Because both the ASP and the SAT approaches solved a significant number of instances from the SRS data set, and, respectively, the ARG data set, wrt. all six inconsistency measures, we selected those two data sets for further analysis. More specifically, we examined the average runtime *composition*, meaning we defined a number of categories representing the main elements of the approaches and measured how much of the overall runtime is spent on each of them. The categories are “encoding generation”, “solving”, and for the SAT approaches additionally “transformation to CNF”. Moreover, we defined the category “other” for any remaining tasks, such as loading the given knowledge base. Note that “solving” does not only include the plain solving time, but also the solver initialization for both ASP and SAT methods.

Furthermore, we compare the *average* runtimes of the SAT and ASP approaches wrt. each measure and regarded data set. For the SAT methods, this includes all required search steps, i.e., multiple encoding generation steps (this particularly affects cardinality constraints), multiple solver calls, and multiple transformations to CNF.

To begin with, we consider the results wrt. the SRS data set. In total, the ASP approaches take between 0.03s ( $\mathcal{I}_c$  and  $\mathcal{I}_d^{\text{hit}}$ ) and 9.89s ( $\mathcal{I}_f$ ), and the SAT approaches between 0.44s ( $\mathcal{I}_c$ ) and 101.92s ( $\mathcal{I}_d^{\Sigma}$ ) on average. Figure 1 visualizes the runtime composition of the ASP-based and the SAT-based approach for  $\mathcal{I}_c$ . We see that the encoding generation covers the largest fraction of the overall average runtime of the SAT-based approach. Wrt. the ASP-based approach on the other hand, encoding generation merely takes a very small fraction of a second (and is thus not even visible in the figure). Solving makes up the second-largest share of the overall runtime of the SAT-based method; the other two categories (i.e., the transformation to CNF, and “other”) only have a relatively minor impact.

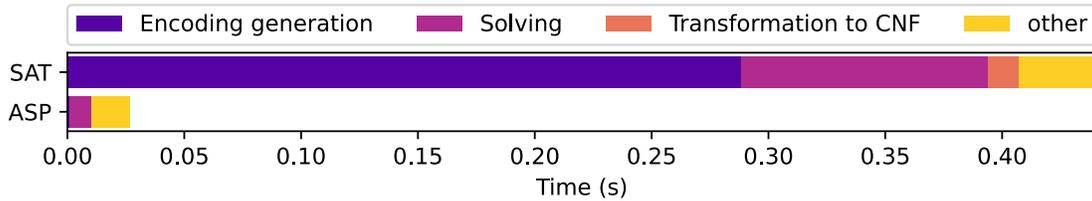


Figure 1: Overview of the average runtime composition of the ASP-based and SAT-based approaches for  $\mathcal{I}_c$  wrt. the SRS data set.



Figure 2: Overview of the average runtime composition of the ASP-based and SAT-based approaches for  $\mathcal{I}_f$  wrt. the SRS data set.

Regarding the ASP-based method, both solving and “other” contribute significantly to the overall runtime. A probable reason for the relatively large proportion of “other” is the overall very short average runtime of the ASP-based approach.

The previously described findings do not, however, reflect the results regarding the remaining five inconsistency measures on the SRS data set. More precisely, wrt. the SAT-based approaches for the latter measures (i.e.,  $\mathcal{I}_f$ ,  $\mathcal{I}_h$ ,  $\mathcal{I}_d^{\max}$ ,  $\mathcal{I}_d^{\Sigma}$ , and  $\mathcal{I}_d^{\text{hit}}$ ), the encoding part does not make up the largest fraction of the overall average runtime—in fact, it makes up a significantly smaller proportion than solving or CNF transformation in all cases (see also Figure 33). The transformation to CNF, which represents the smallest runtime component for  $\mathcal{I}_c$ , plays a more prominent role for the other measures. In the case of  $\mathcal{I}_h$ , it even represents the largest fraction. Regarding the forgetting-based, as well as all three distance-based measures, the largest runtime component is comprised by solving. Figure 2 shows the runtime composition with regard to  $\mathcal{I}_f$ , which represents the previously described pattern that can also be applied to  $\mathcal{I}_d^{\max}$ ,  $\mathcal{I}_d^{\Sigma}$ ,  $\mathcal{I}_d^{\text{hit}}$ , and (under the restriction that the solving-to-CNF transformation ratio does not quite fit) also to  $\mathcal{I}_h$ . Moreover, this figure shows that the only relevant runtime component of the ASP-based approach for  $\mathcal{I}_f$  is the solving part. The same applies to the four remaining measures ( $\mathcal{I}_h$ ,  $\mathcal{I}_d^{\max}$ ,  $\mathcal{I}_d^{\Sigma}$ ,  $\mathcal{I}_d^{\text{hit}}$ ) as well.

As mentioned in the beginning of this section, the total runtime of both ASP and SAT approaches differs from measure to measure. Moreover, the ratio of the average ASP runtime to the average SAT runtime differs wrt. the inconsistency measures as well. For example, wrt.  $\mathcal{I}_d^{\text{hit}}$ , ASP is roughly 73 times as fast as SAT (0.03 s vs. 2.55 s), however, wrt.  $\mathcal{I}_f$ , ASP is only about 7 times as fast (9.89 s vs. 65.30 s).

Some observations we made wrt. the SRS data set also apply to the ARG data set:

- For the ASP-based approaches, the solving part is essentially the only relevant component regarding the overall runtime.

- For the SAT-based approaches, solving overall makes up the most relevant part of the runtime composition.
- Wrt. the SAT approach for  $\mathcal{I}_c$ , the encoding generation part takes up a larger share of the overall runtime than wrt. the SAT approaches for the other five measures.
- Wrt. the SAT approach for  $\mathcal{I}_h$ , the transformation to CNF takes up a larger fraction of the overall runtime than wrt. the SAT approaches for the other measures.

On the other hand, the runtime composition wrt. the ARG data set also provides some new insights (see Figure 34 for an overview of all six measures). To begin with, the ratio between the average ASP runtime and the average SAT runtime is smaller than wrt. the SRS data set. More precisely, an ASP approach is at most roughly 7 times as fast as its corresponding SAT version ( $\mathcal{I}_f$ ). Moreover, in the case of  $\mathcal{I}_h$ , the SAT-based method is on average even a bit faster than the ASP-based one. However, this observation has to be taken with a grain of salt—the number of instances the ASP-based method could solve is higher than the number of instances the SAT-based approach could solve, and when considering individual knowledge bases, the ASP method was faster in each case (see Figure 23c).

### 5.3.3 LINEAR SEARCH IN SAT APPROACHES

In addition to the runtimes of the different approaches considered in this work we logged the inconsistency values wrt. the different measures and data sets (in Appendix B.2, we include histograms over all inconsistency values that could be retrieved for each data set and measure). From the results we can see that the range of inconsistency values varies greatly, in particular when considering the size of the search space. For instance, the contention inconsistency values range from 0 to 25 on the SRS data set, with a maximum possible value of  $|\text{At}(\mathcal{K})| = 30$  (for the largest fraction of instances). In contrast, the max-distance inconsistency measure has the same search space as  $\mathcal{I}_c$ , but the values we actually measured in our experiments are only ranging from 0 to 2, with roughly 1000 instances resulting in the value 1. Another example is the sum-distance measure—here, the values are ranging from 0 to 42. However, the search space is much larger, with a maximum possible value of  $|\text{At}(\mathcal{K})| \cdot |\mathcal{K}| = 30 \cdot 100 = 3000$  (again, for the largest instances). These observations raise the question of whether the binary search approach used in the SAT-based methods is the most practical solution in all cases. We therefore conduct an experiment to compare the previously used SAT approach which includes a binary search procedure to an adapted version which uses linear search instead.

In our experiment, we focus on the three previously mentioned inconsistency measures ( $\mathcal{I}_c$ ,  $\mathcal{I}_d^{\max}$ , and  $\mathcal{I}_d^{\Sigma}$ ) wrt. the SRS data set. Table 6 displays the results (see also Figures 8, 29, and 30 for further visualizations). We can see that with regard to  $\mathcal{I}_c$ , the linear search variant of the SAT approach is on average a bit slower than the binary search variant. Nevertheless, the two methods perform quite similarly. Concerning  $\mathcal{I}_d^{\max}$ , the two approaches perform again similarly. However, despite being a bit faster on average, the linear search version times out in 33 cases more than its binary search counterpart. Regarding  $\mathcal{I}_d^{\Sigma}$ , the results look quite different than those described before—the linear search variant performs clearly superior, exhibiting 410 fewer timeouts than the binary search

	Binary search		Linear search	
	#solved	CRT (s)	#solved	CRT (s)
$\mathcal{I}_c$	<b>1800</b>	793	<b>1800</b>	1670
$\mathcal{I}_d^{\max}$	<b>1618</b>	97589	1585	43506
$\mathcal{I}_d^{\Sigma}$	658	67065	<b>1068</b>	77712

Table 6: Comparison between binary search and linear search in the SAT approach. Displayed are the number of solved instances and the cumulative runtime (CRT) per measure on the SRS data set (1800 knowledge bases). Timeout: 10 minutes.

version. Thus, although binary search is in theory a more efficient search strategy than linear search, it is not necessarily the superior option in practice.

#### 5.3.4 MAXSAT APPROACHES

Another aspect we aim to investigate regarding SAT is the use of a *MaxSAT* solver. The *maximum satisfiability problem* (MaxSAT) is the problem of finding an assignment of truth values that satisfies a maximum number of clauses (for an overview on this topic, see, e.g., (Bacchus, Jarvisalo, & Martins, 2021; Li & Manyà, 2021)). Hence, our proposed technique of using iterative SAT checks is essentially a naive method of solving a MaxSAT problem. In the following, we use the contension measure as an example to show how our SAT encodings can be modified to become MaxSAT encodings.

A (partial) MaxSAT encoding consists of *hard clauses* and *soft clauses*. While the hard clauses must be satisfied, the soft clauses do not. Nevertheless, the goal is to satisfy as many soft clauses as possible. Recall the SAT encoding for  $\mathcal{I}_c$  being determined by (SC1)–(SC17) in Section 3.3. The signature, defined by (SC1)–(SC2), remains the same in the MaxSAT case, and we use the constraints defined via (SC3)–(SC16) as hard clauses. Only (SC17), the at-most- $u$  constraint, is not required anymore. Instead, we define a soft clause  $\neg X_b$  for each  $X \in \text{At}(\mathcal{K})$ . The intuition behind this is that as many atoms as possible are supposed to be *not* set to  $b$ . Note that the SAT-based approaches for the other inconsistency measures can be adapted to MaxSAT in a similar fashion.

We perform an experiment to compare the MaxSAT approach with the previously discussed approaches for  $\mathcal{I}_c$ . To achieve this, we use the SRS data set, and as a MaxSAT solver, we use *EvalMaxSAT*<sup>23</sup> (Avellaneda, 2020). The runtime results are presented in Table 7 (see also Figures 10 and 31). We can observe that the MaxSAT approach is indeed overall faster than the iterative SAT approach. However, it is, in total, still a bit inferior to the ASP-based method. On the other hand, it should be noted that the runtimes for both the MaxSAT and the ASP approach are quite short (< 1 s for each instance). Altogether, the results of this experiments show that the use of MaxSAT has potential and is worth being investigated in future work.

23. <https://github.com/FlorentAvellaneda/EvalMaxSAT>

SAT binary search		SAT linear search		MaxSAT	
#solved	CRT (s)	#solved	CRT (s)	#solved	CRT (s)
1800	793	1800	1670	1800	<b>221</b>

Table 7: Comparison between the two variants of the iterative SAT approach for  $\mathcal{I}_c$  and their MaxSAT counterpart. Displayed are the number of solved instances and the cumulative runtime (CRT) per measure on the SRS data set (1800 knowledge bases). Timeout: 10 minutes.

### 5.3.5 CORE-GUIDED OPTIMIZATION IN ASP

Modern ASP systems, such as Clingo, allow for the use of different optimization strategies. Clingo’s default option (which is used in all experiments described above) is a *branch-and-bound* approach (Gebser, Kaufmann, & Schaub, 2012). In this approach, we first search for an initial solution, i. e., an initial stable model, and (in case a solution exists) obtain its objective value. The next step consists of a loop in which we try to find a solution with a strictly “better” objective value—i. e., in the case of minimization, a strictly lower value, and in the case of maximization, a strictly higher value. This is achieved by adding constraints which ensure that any new solution must have a lower (resp. higher) objective value than the solution derived before. If no “better” solution can be found, we know that the previous solution corresponds to an optimum (Kaminski, Schaub, & Wanko, 2017).

Although the branch-and-bound optimization approach led to promising results in our previous experiments, Clingo has further optimization techniques to offer, which we have yet to explore. As an example, we perform an experiment in which we set Clingo’s optimization strategy to “usc” (Gebser, Kaminski, Kaufmann, Romero, & Schaub, 2015), employs a *core-guided* approach (Andres, Kaufmann, Matheis, & Schaub, 2012). The latter emerged in the area of MaxSAT solving (Fu & Malik, 2006; Marques-Silva & Manquinho, 2008) and is based on the following procedure. We first check if all soft clauses can be satisfied, i. e., we implicitly introduce a lower bound of 0 soft clauses that are allowed to be falsified. If this is not the case, we allow for at most one soft clause to be falsified (i. e., we increase the lower bound to 1), and check for satisfiability again. This process is repeated until we find a satisfiable solution. In order to guide this process, at each step an *unsatisfiable core* (an unsatisfiable set of clauses) is extracted and only soft clauses from that a core are allowed to be falsified (Fu & Malik, 2006). This approach has been transferred to ASP (Andres et al., 2012), with the literals of ASP optimization statements being interpreted as the soft clauses of a MaxSAT problem.

Since, in the case of minimization, core-guided optimization is useful when the target values are close to 0, we selected inconsistency measures and data sets for which this applies (see the histograms in Appendix B.2). To be precise, we performed an experiment with the ASP approach for  $\mathcal{I}_h$  on the ML data set; here, the inconsistency values tend to be very low (either 0 or 1; see Figure 13), while the search space (i. e., the number of formulas) is still quite large (the mean number of formulas is 7506; see Table 4). In addition, we considered the ARG data set in combination with  $\mathcal{I}_h$  and  $\mathcal{I}_d^{\max}$ , as the inconsistency values are likewise always 0 or 1.

		ASP core-guided		ASP	
		#solved	CRT (s)	#solved	CRT (s)
ML	$\mathcal{I}_h$	<b>1220</b>	51441	1163	58203
ARG	$\mathcal{I}_h$	<b>221</b>	11146	218	12975
ARG	$\mathcal{I}_d^{\max}$	<b>176</b>	7266	145	8408

Table 8: Comparison between a core-guided optimization strategy for the ASP approach and the default optimization strategy (branch-and-bound). Displayed are the number of solved instances and the cumulative runtime (CRT) for  $\mathcal{I}_h$  on the ML data set (1920 knowledge bases), as well as  $\mathcal{I}_h$  and  $\mathcal{I}_d^{\max}$  on the ARG data set (326 knowledge bases). Timeout: 10 minutes.

The results in Table 8 show that the core-guided optimization strategy can indeed lead to improved runtimes (see also Figure 9 in Appendix B.1 for the cactus plots and Figure 32 in Appendix B.3 for the corresponding scatter plots). The core-guided variant solves more instances while also maintaining a lower cumulative runtime in all three cases. More specifically, the core-guided approach solves between 1.36 % (ARG/ $\mathcal{I}_h$ ) and 17.61 % (ARG/ $\mathcal{I}_d^{\max}$ ) more instances.

### 5.3.6 PREVIOUS ASP APPROACHES

There already exist two ASP-based approaches for  $\mathcal{I}_c$ , and one for  $\mathcal{I}_f$ , and  $\mathcal{I}_h$ , respectively, in the literature. To be precise, in (Kuhlmann & Thimm, 2020), the authors propose a method similar to our overall SAT approach, which uses ASP encodings for the problem  $\text{UPPER}_{\mathcal{I}_c}$  in order to find  $\text{VALUE}_{\mathcal{I}_c}$  via binary search. A revised version of this approach, which calculates  $\text{VALUE}_{\mathcal{I}_c}$  directly within ASP by means of a minimize statement, is introduced in (Kuhlmann & Thimm, 2021). Note that in the latter version, only propositional language concepts are used, which leads to a program that is already ground. The authors also propose ASP encodings for  $\mathcal{I}_f$  and  $\mathcal{I}_h$  in the same manner. In contrast, the ASP approach presented in (Kuhlmann et al., 2022) (which is also used in the work at hand) makes use of first-order predicates and variables, which enables an automated, and internally optimized, grounding procedure. In order to determine whether this actually has a positive effect on the runtime, we compare the different versions with each other using the contension measure as an example. Note that we apply exactly those implementations which were used in the two corresponding papers (Kuhlmann & Thimm, 2020, 2021).

The results of this experiment, presented in Table 9 (see Figure 11 for an additional visualization), confirm that the newest version of the ASP-based method in fact outperforms its predecessors. The first ASP approach (Kuhlmann & Thimm, 2020), which is based on a binary search procedure, clearly performs the poorest, and hits the timeout of 10 minutes in precisely 600 cases. The second version of the approach (Kuhlmann & Thimm, 2021) solves all 1800 instances within the time limit, nevertheless it performs on average roughly 10 times slower than the current version (0.266 vs. 0.027 seconds). Although the new version might have an advantage by being implemented in C++, both rely on the same ASP solver (Clingo 5.5.1). In fact, the solving time itself is around 3 times shorter wrt. the new ASP version compared to the previous one (0.009 vs. 0.028 seconds on average).

ASP binary search		ASP minimize v1		ASP	
#solved	CRT (s)	#solved	CRT (s)	#solved	CRT (s)
1200	59618	<b>1800</b>	479	<b>1800</b>	48

Table 9: Runtime comparison of the different versions of the ASP approach for  $\mathcal{I}_c$  on the SRS data set. “ASP binary search” refers to the version from (Kuhlmann & Thimm, 2020), “ASP minimize v1” to the version from (Kuhlmann & Thimm, 2021), and “ASP” to the current one. Displayed are the number of solved instances and the cumulative runtime (CRT). Timeout: 10 minutes.

## 5.4 Discussion

In the experiments presented above, we investigated three approaches for the calculation of six different inconsistency measures on five different data sets. Overall, we focused our analysis on the comparison of runtimes; however, we took additional aspects into account, such as the distribution of measured inconsistency values with regard to the different measures and data sets. As a first step, we compared the runtimes of all three approaches per inconsistency measure and data set. Our results confirmed that both the SAT-based and the ASP-based approaches perform, as expected, superior to the naive baseline algorithms. Moreover, the results showed that altogether, the ASP methods are faster than the SAT variants. Only in very isolated cases can the SAT-based or naive approaches retrieve inconsistency values faster than the ASP-based ones or result in fewer timeouts. Furthermore, the ASP-based methods are not only faster than the other two approaches on average, but also wrt. the vast majority of individual instances.

Two factors that presumably play a role in the performance differences between the SAT approaches and the ASP approaches are the following. First, in each SAT-based approach, we need to generate a new at-most- $k$  constraint in each iteration (note that this has been addressed in Section 5.3.2 as well). Second, the SAT solver does not reuse learned clauses from previous calls, i. e., each instance has to be constructed from scratch. The ASP encodings, on the other hand, only have to be constructed and solved once for each knowledge base.

We could further observe that the naive baseline methods for the contension and the forgetting-based measure outperformed the ASP-based and SAT-based approaches in some cases. The most likely reason for this is the fact that the particular instances for which this effect was observed had a very low inconsistency value—in fact, most of these instances were consistent. The baseline methods for both the contension measure and the forgetting-based measure include a satisfiability check (via a SAT solver) as a first step. Hence, consistent knowledge bases can be identified rather quickly. Very small inconsistency values might also be retrieved rather fast by the two naive methods, as they both search for the correct value in a linear manner, starting from the lowest possible value. Another factor that can further decrease the runtime of the two methods is if the given knowledge base is already in CNF, so that no additional transformation steps are required in order to guarantee a valid input for the SAT solver.

A general observation from the evaluation results is that none of the three approaches for either one of the considered inconsistency measures could solve all instances from all data

sets. Hence, even the ASP approaches, which altogether performed superior, clearly reached their limits. In the following, we aim to identify factors that make the computation of inconsistency values for a given knowledge base “difficult” (either for individual approaches or in general).

- A rather obvious factor is the size of the given knowledge base in terms of the number of formulas, number of atoms in the signature, and number of connectives per formula (see Table 4 for an overview regarding the data sets used in our experiments). A larger knowledge base consequently leads to larger encodings, as well as a larger range of possible inconsistency values, which in turn leads to a more difficult problem that the SAT or ASP solver has to solve, as well as to more search steps wrt. the SAT-based approaches. The SAT/ASP solver having to deal with a more difficult problem is directly connected to an increase in runtime, as the results in Section 5.3.2 indicated that the solving process overall makes up the largest fraction of the runtimes of both the SAT-based and ASP-based methods.
- Although all six measures considered in this paper lie on the same level of the polynomial hierarchy (Thimm, 2016a), the ranges of their possible values vary greatly. A larger search space particularly influences the SAT-based approaches. This is reflected, e.g., in the fact that for the sum-distance measure, fewer instances could be solved by the SAT-based approach than for any of the other measures.
- The range of the resulting inconsistency values may also be a relevant factor to take into account. If the inconsistency values are close to 0, a SAT check as a preprocessing step could prove useful—we observed this effect with the naive approaches for the contension measure and the forgetting-based measure, in particular if the given knowledge base is already in CNF. Moreover, the range of the inconsistency values is also relevant for the choice of search strategy regarding the SAT-based approaches (see Section 5.3.3).

Although the previously discussed points hint at which measures might be more suitable in practice in terms of runtime, the choice of an appropriate inconsistency measure for a specific practical application still depends on other factors as well. For instance, in some scenarios it could be useful to look for conflicting formulas (which could be done by using the hitting set measure or the hit-distance measure), while in other scenarios it could be of greater interest to look for the atoms involved in a conflict (which could be done by using the contension or forgetting-based measure). Moreover, the granularity of the measured inconsistency could be of interest (e. g., the forgetting-based measure has a finer granularity than the contension measure). Yet another aspect that should be taken into consideration when dealing with a practical application is the usefulness of the inconsistency measure for resolving the conflicts in the given knowledge base.

## 6. Conclusion

In the course of this work, we proposed a SAT-based and an ASP-based approach for each of six different inconsistency measures (i. e., the contension, the forgetting-based, the hitting set, the max-distance, the sum-distance, and the hit-distance inconsistency measure). With

the SAT-based approaches, we encode the problem of whether a given value is an upper bound of the inconsistency degree, and retrieve the actual inconsistency value by means of iterative SAT solver calls in a (binary) search procedure. In ASP, we can encode the problem of finding an inconsistency value directly by utilizing optimization statements.

In an extensive experimental evaluation we compared the ASP-based to the SAT-based approaches, focusing on runtime. To achieve this, we used a total of five different data sets, two of which were used in preceding works (Kuhlmann & Thimm, 2021; Kuhlmann et al., 2022), while the remaining three are novel. All data sets are publicly available for use in future work. The results of our experiments first of all demonstrated that both the SAT-based and the ASP-based methods overall clearly outperform the naive implementations we used as a baseline—which, to the best of our knowledge, are the only previously existing implementations for the inconsistency measures considered. Moreover, the ASP approaches performed superior to the SAT approaches (with only very few exceptions).

Nonetheless, the results also showed that even the ASP methods are by far not able to solve all instances from all data sets wrt. all considered inconsistency measures. Some data sets turned out to be generally challenging; an example of this is the SC data set, which contains benchmark data from the SAT competition 2020. A reason for this lies in the sheer size of the individual knowledge bases, i. e., the number of formulas contained in a knowledge base, its signature size, and the number of connectives per formula. A large knowledge base requires a large SAT/ASP encoding (in terms of variables and clauses/rules), which in turn leads to an increased solving time and therefore an increased runtime—our results also showed that the solving process altogether makes up the largest fraction of the runtime composition of both the SAT and the ASP approaches.

In order to use approaches to inconsistency measurement in a practical application in a sensible fashion, an analysis of the given data is vital. Information about signature sizes or numbers of formulas per knowledge base can be an indicator for which inconsistency measure will likely be a reasonable choice in terms of runtime. Moreover, if there is any information available about what degrees of inconsistency can roughly be expected, an adjusted search strategy can be used. Such information can further be used to develop preprocessing techniques. For instance, if there is a certain probability that the knowledge bases at hand might be consistent, an initial satisfiability check by means of a SAT solver (as we saw with the naive implementations of the contension and the forgetting-based measure) could prove useful. Certainly there are many other options regarding the topic of preprocessing; however, this is subject to future work.

Another issue to contemplate in terms of future work is a more in-depth examination of MaxSAT approaches. We have already observed in our experiments that a MaxSAT approach for the contension inconsistency measure performs superior to its iterative SAT counterpart. Moreover, there are numerous problem solving paradigms that have not been explored yet in the field of inconsistency measurement. Examples include *Integer Linear Programming*, *Quantified Boolean Formulas*, and *Counterexample-Guided Abstraction Refinement* approaches. The choice of problem solving paradigm depends on the complexity of the measure at hand—which is related to yet another possible topic of future work: the investigation of algorithms for inconsistency measures beyond the first level of the polynomial hierarchy. Furthermore, we merely covered inconsistency measures designed for propositional knowledge bases. The development of approaches for computing inconsistencies in

other domains, such as different logics, data bases, or business process models, constitutes another open problem. In addition, our evaluation showed that all approaches introduced in this work hit their limit at some point during the experiments. To counteract this effect, i. e., to develop more scalable approaches, we already mentioned the possibility of applying preprocessing methods. A different strategy to improve scalability could be to apply approximation techniques.

## Acknowledgments

The research reported here was supported by the Deutsche Forschungsgemeinschaft under grant 506604007.

## Appendix A. Proofs

This appendix contains the correctness proofs for the SAT-based approaches introduced in Section 3, and the ASP-based approaches introduced in Section 4.

### A.1 SAT-based Algorithms

Each of the following sections comprises a correctness proof corresponding to the SAT encoding for one of the six inconsistency measures considered in this work.

#### A.1.1 THE CONTENSION INCONSISTENCY MEASURE

<b>Signature</b>	
For every atom $X \in \text{At}(\mathcal{K})$ :	
Create atoms $X_t, X_f, X_b$	(SC1)
For every formula $\Phi \in \text{sub}(\mathcal{K})$ :	
Create atoms $v_{\Phi}^t, v_{\Phi}^f, v_{\Phi}^b$	(SC2)
<b>Constraints</b>	
For every atom $X \in \text{At}(\mathcal{K})$ :	
$(X_t \vee X_f \vee X_b) \wedge (\neg X_t \vee \neg X_f) \wedge (\neg X_t \vee \neg X_b) \wedge (\neg X_b \vee \neg X_f)$	(SC3)
For every conjunction $\Phi_c = \Psi_{c,1} \wedge \Psi_{c,2}$ appearing in some formula:	
$v_{\Phi_c}^t \leftrightarrow v_{\Psi_{c,1}}^t \wedge v_{\Psi_{c,2}}^t$	(SC4)
$v_{\Phi_c}^f \leftrightarrow v_{\Psi_{c,1}}^f \vee v_{\Psi_{c,2}}^f$	(SC5)
$v_{\Phi_c}^b \leftrightarrow (\neg v_{\Psi_{c,1}}^t \vee \neg v_{\Psi_{c,2}}^t) \wedge \neg v_{\Psi_{c,1}}^f \wedge \neg v_{\Psi_{c,2}}^f$	(SC6)
For every disjunction $\Phi_d = \Psi_{d,1} \vee \Psi_{d,2}$ appearing in some formula:	
$v_{\Phi_d}^t \leftrightarrow v_{\Psi_{d,1}}^t \vee v_{\Psi_{d,2}}^t$	(SC7)
$v_{\Phi_d}^f \leftrightarrow v_{\Psi_{d,1}}^f \wedge v_{\Psi_{d,2}}^f$	(SC8)
$v_{\Phi_d}^b \leftrightarrow (\neg v_{\Psi_{d,1}}^f \vee \neg v_{\Psi_{d,2}}^f) \wedge \neg v_{\Psi_{d,1}}^t \wedge \neg v_{\Psi_{d,2}}^t$	(SC9)
For each negation $\Phi_n = \neg \Psi_n$ appearing in some formula:	
$v_{\Phi_n}^t \leftrightarrow v_{\Psi_n}^f$	(SC10)
$v_{\Phi_n}^f \leftrightarrow v_{\Psi_n}^t$	(SC11)
$v_{\Phi_n}^b \leftrightarrow v_{\Psi_n}^b$	(SC12)

For every formula $\Phi_a$ which consists of a single atom $X$ :	
$v_{\Phi_a}^t \leftrightarrow X_t$	(SC13)
$v_{\Phi_a}^f \leftrightarrow X_f$	(SC14)
$v_{\Phi_a}^b \leftrightarrow X_b$	(SC15)
For every formula $A \in \mathcal{K}$ :	
$v_A^t \vee v_A^b$	(SC16)
Let $\text{At}_b^{\mathcal{K}} = \{X_b \mid X \in \text{At}(\mathcal{K})\}$ . Cardinality constraint:	
$\text{at\_most}(u, \text{At}_b^{\mathcal{K}})$	(SC17)

Encoding 1: Overview of the SAT encoding  $S_c(\mathcal{K}, u)$  of the contension inconsistency measure  $\mathcal{I}_c$ .  $\mathcal{K}$  is a given knowledge base and  $\text{At}(\mathcal{K})$  its signature.  $u$  is a candidate for an upper limit of  $\mathcal{I}_c(\mathcal{K})$ . Let  $\text{sub}(\mathcal{K})$  be the set of subformulas of  $\mathcal{K}$  as defined in Section 3.1. The signature size of the encoding is  $|\text{At}(\mathcal{K})| \cdot 3 + |\text{sub}(\mathcal{K})| \cdot 3$ .

**Theorem 1.** *For a given value  $u$ , the encoding  $S_c(\mathcal{K}, u)$  is satisfiable if and only if  $\mathcal{I}_c(\mathcal{K}) \leq u$ .*

*Proof.* Assume  $S_c(\mathcal{K}, u)$  is satisfiable and let  $\omega$  be a model of  $S_c(\mathcal{K}, u)$ . Define the three-valued interpretation  $\omega^3$  via

$$\omega^3(X) = \begin{cases} t & \text{if } \omega(X_t) = t \\ f & \text{if } \omega(X_f) = t \\ b & \text{if } \omega(X_b) = t \end{cases}$$

Due to (SC3) (see Encoding 1) it is clear that  $\omega^3$  is well-defined. We now show, via structural induction, that  $\omega^3(\Phi) = \theta$  iff  $\omega(v_{\Phi}^{\theta}) = t$  for all  $\Phi \in \text{sub}(\mathcal{K})$ :

- $\Phi_a = X$  for  $X \in \text{At}$ : this follows already from the definition of  $\omega^3$  above and (SC13)–(SC15).
- $\Phi = \neg\Psi$ : by induction hypothesis,  $\omega^3(\Psi) = \theta$  iff  $\omega(v_{\Psi}^{\theta}) = t$ . This implies that exactly one of  $v_{\Psi}^t, v_{\Psi}^f, v_{\Psi}^b$  is true. Due to (SC10)–(SC12), it follows that exactly one of  $v_{\neg\Psi}^t, v_{\neg\Psi}^f, v_{\neg\Psi}^b$  is true. Observe that (SC10)–(SC12) models exactly the semantics of negation in three-valued logic. It follows  $\omega^3(\neg\Psi) = \theta$  iff  $\omega(v_{\neg\Psi}^{\theta}) = t$ .
- $\Phi = \Psi_1 \wedge \Psi_2$ : by induction hypothesis,  $\omega^3(\Psi_1) = \theta$  iff  $\omega(v_{\Psi_1}^{\theta}) = t$  and  $\omega^3(\Psi_2) = \theta$  iff  $\omega(v_{\Psi_2}^{\theta}) = t$ . This implies that exactly one of  $v_{\Psi_1}^t, v_{\Psi_1}^f, v_{\Psi_1}^b$  and exactly one of  $v_{\Psi_2}^t, v_{\Psi_2}^f, v_{\Psi_2}^b$  is true. Due to (SC4)–(SC6), it follows that exactly one of  $v_{\Psi_1 \wedge \Psi_2}^t, v_{\Psi_1 \wedge \Psi_2}^f, v_{\Psi_1 \wedge \Psi_2}^b$  is true. Observe that (SC4)–(SC6) models exactly the semantics of conjunction in three-valued logic: (SC4) models the case that both  $\Psi_1$  and  $\Psi_2$  are *true* (then  $\Psi_1 \wedge \Psi_2$  is also *true*), (SC5) models the case that one of  $\Psi_1$  and  $\Psi_2$  is *false* (then  $\Psi_1 \wedge \Psi_2$  is also *false*), and (SC6) covers all remaining cases (then  $\Psi_1 \wedge \Psi_2$  is *both*).
- $\Phi = \Psi_1 \vee \Psi_2$ : by induction hypothesis,  $\omega^3(\Psi_1) = \theta$  iff  $\omega(v_{\Psi_1}^{\theta}) = t$  and  $\omega^3(\Psi_2) = \theta$  iff  $\omega(v_{\Psi_2}^{\theta}) = t$ . This implies that exactly one of  $v_{\Psi_1}^t, v_{\Psi_1}^f, v_{\Psi_1}^b$  and exactly one of  $v_{\Psi_2}^t, v_{\Psi_2}^f, v_{\Psi_2}^b$

$v_{\Psi_2}^f, v_{\Psi_2}^b$  is true. Due to (SC7)–(SC9), it follows that exactly one of  $v_{\Psi_1 \vee \Psi_2}^t, v_{\Psi_1 \vee \Psi_2}^f, v_{\Psi_1 \vee \Psi_2}^b$  is true. Observe that (SC7)–(SC9) models exactly the semantics of disjunction in three-valued logic: (SC7) models the case that at least one of  $\Psi_1$  and  $\Psi_2$  is *true* (then  $\Psi_1 \vee \Psi_2$  is also *true*), (SC8) models the case that both  $\Psi_1$  and  $\Psi_2$  are *false* (then  $\Psi_1 \vee \Psi_2$  is also *false*), and (SC9) covers all remaining cases (then  $\Psi_1 \vee \Psi_2$  is *both*).

So it follows that  $\omega^3(\Phi) = \theta$  iff  $\omega(v_\Phi^\theta) = t$  for all  $\Phi \in \text{sub}(\mathcal{K})$ . Note that this includes all  $\Phi \in \mathcal{K}$  as well. Due to (SC16), for all  $\Phi \in \mathcal{K}$  we have  $\omega^3(\Phi) \in \{t, b\}$  and therefore  $\omega^3 \in \text{Mod}^3(\mathcal{K})$ . Due to (SC17), we also have  $|\text{Conflictbase}(\omega^3)| \leq u$  and it follows  $\mathcal{I}_c(\mathcal{K}) \leq u$ .

For the other direction, assume  $\mathcal{I}_c(\mathcal{K}) \leq u$ . Then there exists a three-valued interpretation  $\omega^3$  with  $\omega^3 \in \text{Mod}^3(\mathcal{K})$  and  $|\text{Conflictbase}(\omega^3)| \leq u$ . Define  $\omega$  on the signature of  $S_c(\mathcal{K}, u)$  via

$$\begin{aligned}
 \omega(x_t) &= \begin{cases} t & \text{if } \omega^3(X) = t \\ f & \text{otherwise} \end{cases} \\
 \omega(x_f) &= \begin{cases} t & \text{if } \omega^3(X) = f \\ f & \text{otherwise} \end{cases} \\
 \omega(x_b) &= \begin{cases} t & \text{if } \omega^3(X) = b \\ f & \text{otherwise} \end{cases} \\
 \omega(v_\Phi^t) &= \begin{cases} t & \text{if } \omega^3(\Phi) = t \\ f & \text{otherwise} \end{cases} \\
 \omega(v_\Phi^f) &= \begin{cases} t & \text{if } \omega^3(\Phi) = f \\ f & \text{otherwise} \end{cases} \\
 \omega(v_\Phi^b) &= \begin{cases} t & \text{if } \omega^3(\Phi) = b \\ f & \text{otherwise} \end{cases}
 \end{aligned}$$

for all  $X \in \text{At}(\mathcal{K})$  and  $\Phi \in \text{sub}(\mathcal{K})$ . Using the same argumentation as above, it can be shown that  $\omega$  satisfies  $S_c(\mathcal{K}, u)$ .  $\square$

### A.1.2 THE FORGETTING-BASED INCONSISTENCY MEASURE

<b>Signature</b>	
We make use of the original signature $\text{At}(\mathcal{K})$ .	
For every occurrence $X^l \in \text{Occ}(\mathcal{K})$ :	
Create atoms $t_{X,l}, f_{X,l}$	(SF1)
<b>Constraints</b>	
For every formula $A \in \mathcal{K}$ , replace every atom occurrence $X^l$ with:	
$(t_{X,l} \vee X) \wedge \neg f_{X,l}$	(SF2)
Let $\text{At}_{\text{forget}}$ be the set of all atoms defined by (SF3).	
For each pair $t_{X,l}, f_{X,l} \in \text{At}_{\text{forget}}$ :	
$\neg t_{X,l} \vee \neg f_{X,l}$	(SF4)
Cardinality constraint:	

at_most( $u, \text{At}_{\text{forget}}$ )	(SF4)
---	-------

Encoding 2: Overview of the SAT encoding  $S_f(\mathcal{K}, u)$  of the forgetting-based inconsistency measure  $\mathcal{I}_f$ .  $\mathcal{K}$  is a knowledge base and  $\text{At}(\mathcal{K})$  its signature.  $u$  is a candidate for an upper limit for  $\mathcal{I}_f(\mathcal{K})$ . Let  $\text{Occ}(\mathcal{K})$  be the set of all atom occurrences in  $\mathcal{K}$  as defined in Section 3.1. The signature size of the encoding is  $|\text{At}(\mathcal{K})| + 2 \cdot |\text{Occ}(\mathcal{K})|$ .

Let  $\mathcal{K}$  be an arbitrary knowledge base and  $S_f(\mathcal{K}, u)$  the corresponding SAT encoding defined via (SF1)–(SF4). Assume that  $S_f(\mathcal{K}, u)$  is satisfiable, and let  $\omega$  be a model of  $S_f(\mathcal{K}, u)$ . We define

$$\begin{aligned} T_{S_f(\mathcal{K}, u)}(\omega) &= \{t_{X,l} \mid \omega(t_{X,l}) = t\} \\ F_{S_f(\mathcal{K}, u)}(\omega) &= \{f_{X,l} \mid \omega(f_{X,l}) = t\} \end{aligned}$$

Due to (SF3) (i. e.,  $\neg t_{X,l} \vee \neg f_{X,l}$ ),  $T_{S_f(\mathcal{K}, u)}(\omega) \cap F_{S_f(\mathcal{K}, u)}(\omega) = \emptyset$  for every model  $\omega$  of  $S_f(\mathcal{K}, u)$ .

**Lemma 1.** *If  $\omega$  is a model of  $S_f(\mathcal{K}, u)$  then  $(\bigwedge \mathcal{K})[X_1^1 \rightarrow \top; \dots; X_p^n \rightarrow \top; Y_1^1 \rightarrow \perp; \dots; Y_q^m \rightarrow \perp]$  is consistent with  $T_{S_f(\mathcal{K}, u)}(\omega) = \{t_{X_1,1}, \dots, t_{X_p,n}\}$  and  $F_{S_f(\mathcal{K}, u)}(\omega) = \{f_{Y_1,1}, \dots, f_{Y_q,m}\}$ .*

*Proof.* To begin with, each  $X^l$  represents exactly the  $l$ -th occurrence of an atom  $X$ . Moreover, each  $X^l$  gets a distinct pair of atoms  $t_{X,l}$  and  $f_{X,l}$  (SF1). Thus,  $t_{X,l}$  and  $f_{X,l}$  represent exactly the  $l$ -th occurrence of  $X$ .

Following (SF2), each atom occurrence  $X^l$  is replaced by the term  $(t_{X,l} \vee X) \wedge \neg f_{X,l}$ . If  $\omega(t_{X,l}) = t$  (and  $\omega(f_{X,l}) = f$  due to (SF3)), the term becomes true, regardless of whether  $\omega(X) = t$  or  $\omega(X) = f$ . Hence, this corresponds exactly to replacing the  $l$ -th occurrence of  $X$  in the original formula by  $\top$ . On the other hand, if  $\omega(f_{X,l}) = t$  (and  $\omega(t_{X,l}) = f$  due to (SF3)), the term becomes false, regardless of the truth value of  $X$ . Consequently, this corresponds to replacing the  $l$ -th occurrence of  $X$  in the original formula by  $\perp$ . If  $\omega(t_{X,l}) = f$  and  $\omega(f_{X,l}) = f$ , the right part of the conjunction is true, and the left part is true if  $\omega(X) = t$  and false if  $\omega(X) = f$ . Thus, the term evaluates to the truth value of  $X^l$ , i. e., the truth value of the  $l$ -th occurrence of  $X$  in the original formula.  $\square$

**Theorem 2.** *For a given value  $u$ , the encoding  $S_f(\mathcal{K}, u)$  is satisfiable if and only if  $\mathcal{I}_f(\mathcal{K}) \leq u$ .*

*Proof.* From Lemma 1 we know that  $T_{S_f(\mathcal{K}, u)}(\omega)$  and  $F_{S_f(\mathcal{K}, u)}(\omega)$  contain exactly those  $t_{X,l}$  and  $f_{X,l}$  which correspond to those occurrences of  $X$  that are being forgotten (i. e., replaced by  $\top$ , and  $\perp$ , respectively). Moreover, the cardinality constraint (SF4) ensures that  $|T_{S_f(\mathcal{K}, u)}(\omega) \cup F_{S_f(\mathcal{K}, u)}(\omega)| \leq u$ . Thus, if more than  $u$  atom occurrences must be forgotten in order for  $\mathcal{K}$  to be satisfiable, then  $S_f(\mathcal{K}, u)$  is unsatisfiable. Otherwise, i. e., if a minimum of  $\leq u$  atom occurrences must be forgotten,  $S_f(\mathcal{K}, u)$  is satisfiable.  $\square$

## A.1.3 THE HITTING-SET INCONSISTENCY MEASURE

<b>Signature</b>	
For every atom $X \in \text{At}(\mathcal{K})$ :	
Create variables $X_i$ with $i \in \{1, \dots, u\}$	(SH1)
For every formula $A \in \mathcal{K}$ :	
Create variables $p_{A,i}$ with $i \in \{1, \dots, u\}$	(SH2)
<b>Constraints</b>	
For every formula $A \in \mathcal{K}$ and $i \in \{1, \dots, u\}$ , create $A_i$ , which is a copy of $A$ where each instance of atom $X$ is replaced by $X_i$ , and add:	
$p_{A,i} \rightarrow A_i$	(SH3)
For every formula $A \in \mathcal{K}$ :	
$\bigvee_{1 \leq i \leq u} p_{A,i}$	(SH4)

Encoding 3: Overview of the SAT encoding  $S_h(\mathcal{K}, u)$  of the hitting set inconsistency measure  $\mathcal{I}_h$ .  $\mathcal{K}$  is a knowledge base and  $\text{At}(\mathcal{K})$  its signature.  $u$  is a candidate for an upper limit for  $\mathcal{I}_h(\mathcal{K})$ . The signature size of the encoding is  $u \cdot |\text{At}(\mathcal{K})| + u \cdot |\mathcal{K}|$ .

Let  $\mathcal{K}$  be an arbitrary knowledge base and  $S_h(\mathcal{K}, u)$  the corresponding SAT encoding defined via (SH1)–(SH4).

**Lemma 2.** *If  $\mathcal{K}$  contains at least one contradictory formula,  $S_h(\mathcal{K}, u)$  is unsatisfiable for all  $u \in \{1, \dots, |\mathcal{K}|\}$ .*

*Proof.* Let  $u \in \{1, \dots, |\mathcal{K}|\}$  be arbitrary, and let some  $A^\perp \in \mathcal{K}$  be contradictory. Then there exists no interpretation which could satisfy a copy  $A_i^\perp$  ( $i \in \{1, \dots, u\}$ ) of  $A^\perp$ . Hence, in order to satisfy the constraints defined by (SH3) (i. e.,  $p_{A^\perp,i} \rightarrow A_i^\perp$ ), each  $p_{A^\perp,i}$  would need to evaluate to false. However, due to (SH4), at least one  $p_{A^\perp,i}$  must be true. Thus,  $S_h(\mathcal{K}, u)$  is unsatisfiable, regardless of the value of  $u$ .  $\square$

Assume that  $\mathcal{K}$  does not contain any contradictory formulas and that  $S_f(\mathcal{K}, u)$  is satisfiable, and let  $\omega$  be a model of  $S_f(\mathcal{K}, u)$ . We define

$$I_{S_h(\mathcal{K}, u)}(\omega) = \{i_a \mid \exists \omega(p_{A, i_a}) = t\}$$

with  $A \in \mathcal{K}$ , and  $i_a \in \{1, \dots, u\}$ .

**Lemma 3.**  *$I_{S_h(\mathcal{K}, u)}(\omega)$  corresponds to a hitting set of  $\mathcal{K}$ .*

*Proof.* For each formula  $A \in \mathcal{K}$  we have at least one  $p_{A, i_a}$  that evaluates to true, because of (SH4). Consequently, wrt. that  $p_{A, i_a}$ ,  $A_{i_a}$  must be true as well (SH3). Hence, at least one copy  $A_{i_a}$  of each  $A \in \mathcal{K}$  is satisfied.

Further, each  $i \in \{1, \dots, u\}$  corresponds to an interpretation, since each  $A_i$  uses the  $i$ -th copy of  $\text{At}(\mathcal{K})$ . Moreover, each  $i_a \in I_{S_h(\mathcal{K}, u)}$  corresponds to an “active” interpretation, i. e., one that is actually used to satisfy one or more formulas. Thus,  $I_{S_h(\mathcal{K}, u)}(\omega)$  corresponds to a hitting set of  $\mathcal{K}$ .  $\square$

**Theorem 3.** *For a given value  $u$ , the encoding  $S_h(\mathcal{K}, u)$  is satisfiable if and only if  $\mathcal{I}_h(\mathcal{K}) \leq u - 1$ .  $S_h(\mathcal{K}, u)$  is unsatisfiable for all  $u = 1, \dots, |\mathcal{K}|$  if and only if  $\mathcal{I}_h(\mathcal{K}) = \infty$ .*

*Proof.* Lemma 2 shows that  $S_h(\mathcal{K}, u)$  is unsatisfiable for all  $u \in \{1, \dots, |\mathcal{K}|\}$  if  $\mathcal{K}$  contains a contradictory formula, i. e., if  $\mathcal{I}_h(\mathcal{K}) = \infty$ . Lemma 3 shows that, if  $\mathcal{K}$  does not contain any contradictory formula,  $I_{S_h(\mathcal{K}, u)}(\omega) = \{i_a \mid \exists \omega(p_{A, i_a}) = t\}$  corresponds to a hitting set of  $\mathcal{K}$ . Since the maximum cardinality of  $I_{S_h(\mathcal{K}, u)}(\omega)$  is restricted by  $u$ ,  $S_h(\mathcal{K}, u)$  is unsatisfiable if more than  $u$  copies of each formula are required, i. e., if more than  $u$  different interpretations are required to satisfy each formula in  $\mathcal{K}$ . Otherwise,  $S_h(\mathcal{K}, u)$  is satisfiable.

Thus, if a hitting set of size  $\leq u$  exists,  $S_h(\mathcal{K}, u)$  is satisfiable. Otherwise, i. e., if no hitting set of size  $\leq u$  exists,  $S_h(\mathcal{K}, u)$  is unsatisfiable. Since we have to subtract 1 from the cardinality of the hitting set in order to get  $\mathcal{I}_h(\mathcal{K})$ , it follows that if  $\mathcal{I}_h(\mathcal{K}) \leq u - 1$ , then  $S_h(\mathcal{K}, u)$  is satisfiable, and if  $\mathcal{I}_h(\mathcal{K}) > u - 1$ , then  $S_h(\mathcal{K}, u)$  is unsatisfiable.  $\square$

#### A.1.4 THE MAX-DISTANCE INCONSISTENCY MEASURE

<b>Signature</b>	
For every atom $X \in \text{At}(\mathcal{K})$ :	
Create variable $X_o$	(SDM1)
For every atom $X \in \text{At}(\mathcal{K})$ and $i \in \{1, \dots,  \mathcal{K} \}$ :	
Create variables $X_i$	(SDM2)
Create variables $\text{inv}_{X,i}$	(SDM3)
<b>Constraints</b>	
For every formula $A \in \mathcal{K}$ and $i \in \{1, \dots,  \mathcal{K} \}$ :	
Create $A_i$ , which is a copy of $A$ where each instance of atom $X$ is replaced by $X_i$	(SDM4)
For every $X_i$ with $i \in \{1, \dots,  \mathcal{K} \}$ :	
$X_i \rightarrow X_o \vee \text{inv}_{X,i}$	(SDM5)
$\neg X_i \rightarrow \neg X_o \vee \text{inv}_{X,i}$	(SDM6)
Let $\text{INV}_i$ be the set containing all $\text{inv}_{X,i}$ for a fixed $i$ .	
For every $i \in \{1, \dots,  \mathcal{K} \}$ :	
$\text{at\_most}(u, \text{INV}_i)$	(SDM7)

Encoding 4: Overview of the SAT encoding  $S_d^{\max}(\mathcal{K}, u)$  of the max-distance inconsistency measure  $\mathcal{I}_d^{\max}$ .  $\mathcal{K}$  is a knowledge base and  $\text{At}(\mathcal{K})$  its signature.  $u$  is a candidate for an upper limit for  $\mathcal{I}_d^{\max}(\mathcal{K})$ . The signature size of the encoding is  $|\text{At}(\mathcal{K})| + 2 \cdot |\mathcal{K}| \cdot |\text{At}(\mathcal{K})|$ .

Let  $\mathcal{K}$  be an arbitrary knowledge base and  $S_d^{\max}(\mathcal{K}, u)$  the corresponding SAT encoding defined via (SDM1)–(SDM7).

**Lemma 4.** *If  $\mathcal{K}$  contains at least one contradictory formula,  $S_d^{\max}(\mathcal{K}, u)$  is unsatisfiable for all  $u \in \{0, \dots, |\text{At}(\mathcal{K})|\}$ .*

*Proof.* Let some  $A^\perp \in \mathcal{K}$  be contradictory. Following (SDM4), we add indexed copies of all formulas  $A \in \mathcal{K}$  to  $S_d^{\max}(\mathcal{K}, u)$ . Thus, we also have an indexed copy  $A_i^\perp$  of  $A^\perp$ . Since adding an index  $i$  to each atom in  $A$  does not resolve the conflict within the formula,  $A_i^\perp$  is still

contradictory. Therefore,  $S_d^{\max}(\mathcal{K}, u)$  contains an unsatisfiable formula and is consequently itself unsatisfiable, regardless of the value of  $u$ .  $\square$

Assume that  $\mathcal{K}$  does not contain any contradictory formulas and that  $S_d^{\max}(\mathcal{K}, u)$  is satisfiable, and let  $\omega$  be a model of  $S_d^{\max}(\mathcal{K}, u)$ . For each  $i \in \{1, \dots, |\mathcal{K}|\}$  we create variables  $X_i$  (SDM2) and variables  $\text{inv}_{X,i}$  (SDM3) for each  $X \in \text{At}(\mathcal{K})$ . We define

$$I_{\max}(\omega, i) = \{\text{inv}_{X,i} \mid \omega(\text{inv}_{X,i}) = t\}$$

Recall that we represent an “optimal” interpretation by introducing a variable  $X_o$  for each  $X \in \text{At}(\mathcal{K})$  (SDM1). The truth values assigned to each  $X_o$  then correspond to the “optimal” interpretation.

**Lemma 5.** *If  $\omega(X_i) \neq \omega(X_o)$  then  $\text{inv}_{X,i} \in I_{\max}(\omega, i)$ .*

*Proof.*  $X_i$  must be either true or false. If  $\omega(X_i) = t$ , then the left part of the implication  $X_i \rightarrow X_o \vee \text{inv}_{X,i}$  (SDM5) is true. Thus, in order to satisfy the implication, the right part must be true as well. If  $X_o$  is false, i. e.,  $\omega(X_i) \neq \omega(X_o)$ , then  $\text{inv}_{X,i}$  must be true. Consequently,  $\text{inv}_{X,i} \in I_{\max}(\omega, i)$ . Moreover, the implication  $\neg X_i \rightarrow \neg X_o \vee \text{inv}_{X,i}$ , defined via (SDM6), evaluates to true as well, because the left part ( $\neg X_i$ ) is false.

If  $\omega(X_i) = f$ , the left part of the implication defined by (SDM6), i. e.,  $\neg X_i \rightarrow \neg X_o \vee \text{inv}_{X,i}$ , is true. Thus, to satisfy this implication,  $\neg X_o \vee \text{inv}_{X,i}$  must evaluate to true. If  $\omega(X_o) = t$ , meaning if  $\omega(X_i) \neq \omega(X_o)$ , then  $\neg X_o$  is false, and  $\text{inv}_{X,i}$  must be true. Hence,  $\text{inv}_{X,i} \in I_{\max}(\omega, i)$ . In addition, the implication  $X_i \rightarrow X_o \vee \text{inv}_{X,i}$ , defined by (SDM5), is also satisfied, because the left part ( $X_i$ ) is false.  $\square$

**Theorem 4.** *For a given value  $u$ , the encoding  $S_d^{\max}(\mathcal{K}, u)$  is satisfiable if and only if  $\mathcal{I}_d^{\max}(\mathcal{K}) \leq u$ .  $S_d^{\max}(\mathcal{K}, u)$  is unsatisfiable for all  $u \in \{0, \dots, |\text{At}(\mathcal{K})|\}$  if and only if  $\mathcal{I}_d^{\max}(\mathcal{K}) = \infty$ .*

*Proof.* Lemma 4 shows that  $S_d^{\max}(\mathcal{K}, u)$  is unsatisfiable for all  $u \in \{1, \dots, |\text{At}(\mathcal{K})|\}$  if  $\mathcal{K}$  contains a contradictory formula, i. e., if  $\mathcal{I}_d^{\max}(\mathcal{K}) = \infty$ .

Let us now assume that  $\mathcal{K}$  does not contain any contradictory formulas, i. e.,  $\mathcal{I}_d^{\max}(\mathcal{K}) \neq \infty$ . From Lemma 5 we know that for each  $i \in |\mathcal{K}|$ , if  $\omega(X_i) \neq \omega(X_o)$ , then the corresponding  $\text{inv}_{X,i}$  must be included in  $I_{\max}(\omega, i)$  (i. e.,  $\omega(\text{inv}_{X,i}) = t$ ). This means that if a formula  $A \in \mathcal{K}$  with index  $i$  requires an interpretation that differs from the “optimal” interpretation, then for each  $X_i \in \text{At}(A_i)$  with  $\omega(X_i) \neq \omega(X_o)$ ,  $\text{inv}_{X,i} \in I_{\max}(\omega, i)$ . Further, (SDM7) restricts  $|I_{\max}(\omega, i)|$  to  $u$  (for each  $i$ ). If  $\mathcal{I}_d^{\max}(\mathcal{K}) \leq u$ , there exists a solution for which  $|I_{\max}(\omega, i)| \leq u$  for each  $i \in \{1, \dots, |\mathcal{K}|\}$ .

Hence, the constraints defined by (SDM7) are satisfied, and  $S_d^{\max}(\mathcal{K}, u)$  is satisfiable. If  $\mathcal{I}_d^{\max}(\mathcal{K}) > u$ , then for at least one  $i$  we have  $|I_{\max}(\omega, i)| > u$ . Consequently, at least one constraint defined by (SDM7) cannot be satisfied, which makes  $S_d^{\max}(\mathcal{K}, u)$  unsatisfiable.  $\square$

## A.1.5 THE SUM-DISTANCE INCONSISTENCY MEASURE

<b>Signature</b>	
For every atom $X \in \text{At}(\mathcal{K})$ :	
Create variable $X_o$	(SDS1)
For every atom $X \in \text{At}(\mathcal{K})$ and $i \in \{1, \dots,  \mathcal{K} \}$ :	
Create variables $X_i$	(SDS2)
Create variables $\text{inv}_{X,i}$	(SDS3)
<b>Constraints</b>	
For every formula $A \in \mathcal{K}$ and $i \in \{1, \dots,  \mathcal{K} \}$ :	
Create $A_i$ , which is a copy of $A$ where each instance of atom $X$ is replaced by $X_i$	(SDS4)
For every $X_i$ with $i \in \{1, \dots,  \mathcal{K} \}$ :	
$X_i \rightarrow X_o \vee \text{inv}_{X,i}$	(SDS5)
$\neg X_i \rightarrow \neg X_o \vee \text{inv}_{X,i}$	(SDS6)
Let $\text{INV}$ be the set containing all $\text{inv}_{X,i}$ for all $i$ . Cardinality constraint:	
$\text{at\_most}(u, \text{INV})$	(SDS7)

Encoding 5: Overview of the SAT encoding  $S_{\text{d}}^{\Sigma}(\mathcal{K}, u)$  of the sum-distance inconsistency measure  $\mathcal{I}_{\text{d}}^{\Sigma}$ .  $\mathcal{K}$  is a knowledge base and  $\text{At}(\mathcal{K})$  its signature.  $u$  is a candidate for an upper limit for  $\mathcal{I}_{\text{d}}^{\Sigma}(\mathcal{K})$ . The signature size of the encoding is  $|\text{At}(\mathcal{K})| + 2 \cdot |\mathcal{K}| \cdot |\text{At}(\mathcal{K})|$ .

Let  $\mathcal{K}$  be an arbitrary knowledge base and  $S_{\text{d}}^{\Sigma}(\mathcal{K}, u)$  the corresponding SAT encoding defined via (SDS1)–(SDS7).

**Lemma 6.** *If  $\mathcal{K}$  contains at least one contradictory formula,  $S_{\text{d}}^{\Sigma}(\mathcal{K}, u)$  is unsatisfiable for all  $u \in \{0, \dots, |\text{At}(\mathcal{K})| \cdot |\mathcal{K}|\}$ .*

*Proof.* Analogous to the proof of Lemma 4. □

Assume that  $\mathcal{K}$  does not contain any contradictory formulas and that  $S_{\text{d}}^{\Sigma}(\mathcal{K}, u)$  is satisfiable, and let  $\omega$  be a model of  $S_{\text{d}}^{\Sigma}(\mathcal{K}, u)$ . For each  $i \in \{1, \dots, |\mathcal{K}|\}$  we create variables  $X_i$  (SDS2) and variables  $\text{inv}_{X,i}$  (SDS3) for each  $X \in \text{At}(\mathcal{K})$ . We define

$$I_{\Sigma}(\omega) = \{\text{inv}_{X,i} \mid \omega(\text{inv}_{X,i}) = t\}$$

Recall that we represent an “optimal” interpretation by introducing a variable  $X_o$  for each  $X \in \text{At}(\mathcal{K})$  (SDS1). The truth values assigned to each  $X_o$  then correspond to the “optimal” interpretation.

**Lemma 7.** *If  $\omega(X_i) \neq \omega(X_o)$  then  $\text{inv}_{X,i} \in I_{\Sigma}(\omega)$ .*

*Proof.* Analogous to the proof of Lemma 5, since (SDM1)–(SDM6) is equivalent to (SDS1)–(SDS6). □

**Theorem 5.** *For a given value  $u$ , the encoding  $S_{\text{d}}^{\Sigma}(\mathcal{K}, u)$  is satisfiable if and only if  $\mathcal{I}_{\text{d}}^{\Sigma}(\mathcal{K}) \leq u$ .  $S_{\text{d}}^{\Sigma}(\mathcal{K}, u)$  is unsatisfiable for all  $u \in \{0, \dots, |\text{At}(\mathcal{K})| \cdot |\mathcal{K}|\}$  if and only if  $\mathcal{I}_{\text{d}}^{\Sigma}(\mathcal{K}) = \infty$ .*

*Proof.* Lemma 6 shows that  $S_d^\Sigma(\mathcal{K}, u)$  is unsatisfiable for all  $u \in \{0, \dots, |\text{At}(\mathcal{K})| \cdot |\mathcal{K}|\}$  if  $\mathcal{K}$  contains a contradictory formula, i. e., if  $\mathcal{I}_d^\Sigma(\mathcal{K}) = \infty$ .

Let us now assume that  $\mathcal{K}$  does not contain any contradictory formulas, i. e.,  $\mathcal{I}_d^\Sigma(\mathcal{K}) \neq \infty$ . From Lemma 7 we know that for any  $i \in |\mathcal{K}|$ , if  $\omega(X_i) \neq \omega(X_o)$ , then the corresponding  $\text{inv}_{X,i}$  must be included in  $I_\Sigma(\omega)$  (i. e.,  $\omega(\text{inv}_{X,i}) = t$ ). This means that if a formula  $A \in \mathcal{K}$  with index  $i$  requires an interpretation that differs from the “optimal” interpretation, then for each  $X_i \in \text{At}(A_i)$  with  $\omega(X_i) \neq \omega(X_o)$ ,  $\text{inv}_{X,i} \in I_\Sigma(\omega)$ . Further, (SDS7) restricts  $|I_\Sigma(\omega)|$  to  $u$ . If  $\mathcal{I}_d^\Sigma(\mathcal{K}) \leq u$ , there exists a solution for which  $|I_\Sigma(\omega)| \leq u$ . Hence, the constraint defined by (SDS7) is satisfied, and  $S_d^\Sigma(\mathcal{K}, u)$  is satisfiable. If  $\mathcal{I}_d^\Sigma(\mathcal{K}) > u$ , then we have  $|I_\Sigma(\omega)| > u$ . Consequently, the constraint defined by (SDS7) cannot be satisfied, which makes  $S_d^\Sigma(\mathcal{K}, u)$  unsatisfiable.  $\square$

#### A.1.6 THE HIT-DISTANCE INCONSISTENCY MEASURE

<b>Signature</b>	
For every formula $A \in \mathcal{K}$ :	
Create variable $\text{hit}_A$	(SDH1)
For every atom $X \in \text{At}(\mathcal{K})$ :	
Create variable $X$	(SDH2)
<b>Constraints</b>	
For every formula $A \in \mathcal{K}$ :	
$A \vee \text{hit}_A$	(SDH3)
Let $\text{HIT}_\mathcal{K}$ be the set containing all variables added in (SDH1).	
Cardinality constraint:	
$\text{at\_most}(u, \text{HIT}_\mathcal{K})$	(SDH4)

Encoding 6: Overview of the SAT encoding  $S_d^{\text{hit}}(\mathcal{K}, u)$  of the hit-distance inconsistency measure  $\mathcal{I}_d^{\text{hit}}$ .  $\mathcal{K}$  is a knowledge base and  $\text{At}(\mathcal{K})$  its signature.  $u$  is a candidate for an upper limit for  $\mathcal{I}_d^{\text{hit}}(\mathcal{K})$ . The signature size of the encoding is  $|\text{At}(\mathcal{K})| + |\mathcal{K}|$ .

Let  $\mathcal{K}$  be an arbitrary knowledge base and  $S_d^{\text{hit}}(\mathcal{K}, u)$  the corresponding SAT encoding defined via (SDH1)–(SDH4). Assume that  $S_d^{\text{hit}}(\mathcal{K}, u)$  is satisfiable, and let  $\omega$  be a model of  $S_d^{\text{hit}}(\mathcal{K}, u)$ . We define

$$H_{\text{hit}}(\omega) = \{\text{hit}_A \mid \omega(\text{hit}_A) = t\}.$$

**Lemma 8.** *Each  $\text{hit}_A \in H_{\text{hit}}(\omega)$  corresponds to a formula  $A \in \mathcal{K}$  being removed from  $\mathcal{K}$ .*

*Proof.* Due to (SDH3), every formula  $A \in \mathcal{K}$  is extended by “ $\vee \text{hit}_A$ ”. If  $\omega(\text{hit}_A) = t$ , the term  $A \vee \text{hit}_A$  will always be true, regardless of the truth value of  $A$  under  $\omega$ . Hence, this is equivalent to the removal of  $A$  from  $\mathcal{K}$ .  $\square$

**Theorem 6.** *For a given value  $u$ , the encoding  $S_d^{\text{hit}}(\mathcal{K}, u)$  is satisfiable if and only if  $\mathcal{I}_d^{\text{hit}}(\mathcal{K}) \leq u$ .*

*Proof.* The cardinality constraint (SDH4) restricts the number of atoms  $\text{hit}_A$  being allowed to be true, to  $u$ . If the minimal number of formulas  $A \in \mathcal{K}$  that need to be removed in

order to make  $\mathcal{K}$  consistent is  $\leq u$  (i. e., if  $\mathcal{I}_d^{\text{hit}}(\mathcal{K}) \leq u$ ) then  $\leq u$  atoms  $\text{hit}_A$  can be set to true (i. e.,  $|H_{\text{hit}}(\omega)| \leq u$ ). From Lemma 8 we know that those atoms correspond to the formulas that are removed. Thus, there exists a solution in which exactly those atoms  $\text{hit}_A$  are set to true which correspond to the formulas  $A \in \mathcal{K}$  that need to be removed to make  $\mathcal{K}$  consistent. Consequently, the remaining formulas in  $\mathcal{K}$  are satisfiable, and  $S_d^{\text{hit}}(\mathcal{K}, u)$  is satisfiable. If more than  $u$  formulas need to be removed (i. e., if  $\mathcal{I}_d^{\text{hit}}(\mathcal{K}) > u$ ), (SDH4) still restricts the number of  $\text{hit}_A$  atoms being set to true to  $u$ . Hence, there exists no model for  $\{\bigcup A \in \mathcal{K} \mid \omega(\text{hit}_A) = f\}$ , and  $S_d^{\text{hit}}(\mathcal{K}, u)$  is unsatisfiable.  $\square$

## A.2 ASP-based Algorithms

The following sections are each comprised of a correctness proof corresponding to the ASP encoding for one of the six inconsistency measures considered in this work.

### A.2.1 THE CONTENSION INCONSISTENCY MEASURE

For every $A \in \mathcal{K}$ :	
<code>kbMember(A).</code>	(AC1)
For every $X \in \text{At}(\mathcal{K})$ :	
<code>atom(X).</code>	(AC2)
For every conjunction $\Phi_c = \Psi_{c,1} \wedge \Psi_{c,2}$ appearing in some formula:	
<code>conjunction(<math>\Phi_c, \Psi_{c,1}, \Psi_{c,2}</math>).</code>	(AC3)
For every disjunction $\Phi_d = \Psi_{d,1} \vee \Psi_{d,2}$ appearing in some formula:	
<code>disjunction(<math>\Phi_d, \Psi_{d,1}, \Psi_{d,2}</math>).</code>	(AC4)
For each negation $\Phi_n = \neg\Psi_n$ appearing in some formula:	
<code>negation(<math>\Phi_n, \Psi_n</math>).</code>	(AC5)
For every formula $\Phi_a$ which consists of a single atom $X$ :	
<code>formulaIsAtom(<math>\Phi_a, X</math>).</code>	(AC6)
As the static part, we define:	
<code>tv(<math>t; f; b</math>).</code>	(AC7)
<code>1{truthValue(A, T) : tv(T)}1 :-     atom(A).</code>	(AC8)
<code>truthValue(F, t) :-     conjunction(F, G, H),     truthValue(G, t),     truthValue(H, t).</code>	(AC9)
<code>truthValue(F, f) :-     conjunction(F, G, H),     1{truthValue(G, f), truthValue(H, f)}.</code>	(AC10)
<code>truthValue(F, b) :-     conjunction(F, -, -),     not truthValue(F, t),     not truthValue(F, f).</code>	(AC11)

```

truthValue(F,f) :-
    disjunction(F,G,H),
    truthValue(G,f),
    truthValue(H,f),
    (AC12)

truthValue(F,t) :-
    disjunction(F,G,H),
    1{truthValue(G,t), truthValue(H,t)}.
    (AC13)

truthValue(F,b) :-
    disjunction(F,-,-),
    not truthValue(F,t),
    not truthValue(F,f).
    (AC14)

truthValue(F,t) :-
    negation(F,G),
    truthValue(G,f).
    (AC15)

truthValue(F,f) :-
    negation(F,G),
    truthValue(G,t).
    (AC16)

truthValue(F,b) :-
    negation(F,G),
    truthValue(G,b).
    (AC17)

truthValue(F,T) :-
    formulaIsAtom(F,G),
    truthValue(G,T),
    tv(T).
    (AC18)

:-
    truthValue(F,f),
    kbMember(F).
    (AC19)

#minimize{1,A : truthValue(A,b), atom(A)}.
    (AC20)

```

Encoding 7: Overview of the ASP encoding  $P_c(\mathcal{K})$  of the contension inconsistency measure  $\mathcal{I}_c$ .  $\mathcal{K}$  is the given knowledge base.

Consider an arbitrary knowledge base  $\mathcal{K}$  as well as the set of rules listed in Encoding 7 in Section 4.2, which make up the extended logic program  $P_c(\mathcal{K})$ . Let  $P'_c(\mathcal{K})$  denote the extended logic program  $P_c(\mathcal{K})$  *without* the minimize statement (AC20).

**Lemma 9.** *If  $M$  is an answer set of  $P'_c(\mathcal{K})$  then the three-valued interpretation  $\omega_M^3$  defined as*

$$\omega_M^3(X) = \begin{cases} t & \text{truthValue}(X,t) \in M \\ f & \text{truthValue}(X,f) \in M \\ b & \text{truthValue}(X,b) \in M \end{cases}$$

is a model of  $\mathcal{K}$  with  $X \in \text{At}(\mathcal{K})$ .

*Proof.*  $\omega_M^3$  is well-defined: As (AC8) states, each answer set of  $P'_c(\mathcal{K})$  will hold exactly one of  $\text{truthValue}(X, t)$ ,  $\text{truthValue}(X, f)$ ,  $\text{truthValue}(X, b)$  for each atom  $X \in \text{At}(\mathcal{K})$ :

$$1\{\text{truthValue}(X, T) : \text{tv}(T)\}1 \text{ :- atom}(X).$$

Via structural induction, we show that for all (sub)formulas  $\Phi$  in  $\mathcal{K}$ ,  $\omega_M^3(\Phi) = \theta$  (with  $\theta \in \{t, f, b\}$ ) if and only if  $\text{truthValue}(\Phi, \theta) \in M$ .

**Induction Hypothesis (I.H.):**  $\omega_M^3(\Phi) = \theta$  if and only if  $\text{truthValue}(\Phi, \theta) \in M$  is true for all subformulas of the considered formula.

1. Let  $\Phi = X$  for  $X \in \text{At}(\mathcal{K})$ .

• “ $\Rightarrow$ ”:

$$\omega_M^3(\Phi) = \theta \Rightarrow \omega_M^3(X) = \theta \xrightarrow{\text{Def.}} \text{truthValue}(X, \theta) \in M$$

• “ $\Leftarrow$ ”:

$$\text{truthValue}(X, \theta) \in M \xrightarrow{\text{Def.}} \omega_M^3(X) = \theta \Rightarrow \omega_M^3(\Phi) = \theta$$

2. Let  $\Phi = \Psi_1 \wedge \Psi_2$ .

(a)  $\theta = t$ :

• “ $\Rightarrow$ ”:

$$\begin{aligned} \omega_M^3(\Psi_1 \wedge \Psi_2) = \theta &\xrightarrow{\theta=t} \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = t \\ \omega_M^3(\Psi_2) = t \end{array} \right\} \\ &\xrightarrow{\text{I.H.}} \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, t) \in M \\ \text{truthValue}(\Psi_2, t) \in M \end{array} \right\} \\ &\xrightarrow{(\text{AC9})} \text{truthValue}(\Psi_1 \wedge \Psi_2, t) \in M \end{aligned}$$

• “ $\Leftarrow$ ”:

$$\begin{aligned} \text{truthValue}(\Psi_1 \wedge \Psi_2, t) \in M &\Rightarrow \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, t) \in M \\ \text{truthValue}(\Psi_2, t) \in M \end{array} \right\} \\ &\Rightarrow \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = t \\ \omega_M^3(\Psi_2) = t \end{array} \right\} \\ &\Rightarrow \omega_M^3(\Psi_1 \wedge \Psi_2) = t \end{aligned}$$

(b)  $\theta = b$ :

• “ $\Rightarrow$ ”:

$$\begin{aligned} \omega_M^3(\Psi_1 \wedge \Psi_2) = \theta &\xrightarrow{\theta=b} \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = t \\ \omega_M^3(\Psi_2) = b \end{array} \right\} \\ &\xrightarrow{\text{I.H.}} \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, t) \in M \\ \text{truthValue}(\Psi_2, b) \in M \end{array} \right\} \\ &\xrightarrow{(\text{AC11})} \text{truthValue}(\Psi_1 \wedge \Psi_2, b) \in M \end{aligned}$$

Analogous for the case where  $\omega_M^3(\Psi_1) = b$  and  $\omega_M^3(\Psi_2) = t$ .

$$\begin{aligned}
 \text{or } & \xrightarrow{\theta=b} \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = b \\ \omega_M^3(\Psi_2) = b \end{array} \right\} \\
 & \xrightarrow{\text{I.H.}} \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, b) \in M \\ \text{truthValue}(\Psi_2, b) \in M \end{array} \right\} \\
 & \xrightarrow{(\text{AC11})} \text{truthValue}(\Psi_1 \wedge \Psi_2, b) \in M
 \end{aligned}$$

- “ $\Leftarrow$ ”:

$$\begin{aligned}
 \text{truthValue}(\Psi_1 \wedge \Psi_2, b) \in M & \Rightarrow \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, t) \in M \\ \text{truthValue}(\Psi_2, b) \in M \end{array} \right\} \\
 & \Rightarrow \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = t \\ \omega_M^3(\Psi_2) = b \end{array} \right\} \\
 & \Rightarrow \omega_M^3(\Psi_1 \wedge \Psi_2) = b
 \end{aligned}$$

Analogous for the case where  $\omega_M^3(\Psi_1) = b$  and  $\omega_M^3(\Psi_2) = t$ .

$$\begin{aligned}
 \text{or } & \Rightarrow \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, b) \in M \\ \text{truthValue}(\Psi_2, b) \in M \end{array} \right\} \\
 & \Rightarrow \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = b \\ \omega_M^3(\Psi_2) = b \end{array} \right\} \\
 & \Rightarrow \omega_M^3(\Psi_1 \wedge \Psi_2) = b
 \end{aligned}$$

(c)  $\theta = f$ :

- “ $\Rightarrow$ ”:

$$\begin{aligned}
 \omega_M^3(\Psi_1 \wedge \Psi_2) = \theta & \xrightarrow{\theta=f} \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = t \\ \omega_M^3(\Psi_2) = f \end{array} \right\} \\
 & \xrightarrow{\text{I.H.}} \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, t) \in M \\ \text{truthValue}(\Psi_f, f) \in M \end{array} \right\} \\
 & \xrightarrow{(\text{AC10})} \text{truthValue}(\Psi_1 \wedge \Psi_2, f) \in M
 \end{aligned}$$

Analogous for the case where  $\omega_M^3(\Psi_1) = f$  and  $\omega_M^3(\Psi_2) = t$ .

$$\begin{aligned}
 \text{or } & \xrightarrow{\theta=f} \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = b \\ \omega_M^3(\Psi_2) = f \end{array} \right\} \\
 & \xrightarrow{\text{I.H.}} \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, b) \in M \\ \text{truthValue}(\Psi_2, f) \in M \end{array} \right\} \\
 & \xrightarrow{(\text{AC10})} \text{truthValue}(\Psi_1 \wedge \Psi_2, f) \in M
 \end{aligned}$$

Analogous for the case where  $\omega_M^3(\Psi_1) = f$  and  $\omega_M^3(\Psi_2) = b$ .

$$\begin{aligned}
 \text{or } & \xrightarrow{\theta=f} \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = f \\ \omega_M^3(\Psi_2) = f \end{array} \right\} \\
 & \xrightarrow{\text{I.H.}} \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, f) \in M \\ \text{truthValue}(\Psi_2, f) \in M \end{array} \right\} \\
 & \xrightarrow{(\text{AC10})} \text{truthValue}(\Psi_1 \wedge \Psi_2, f) \in M
 \end{aligned}$$

- “ $\Leftarrow$ ”:

$$\begin{aligned}
 \text{truthValue}(\Psi_1 \wedge \Psi_2, f) \in M & \Rightarrow \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, t) \in M \\ \text{truthValue}(\Psi_2, f) \in M \end{array} \right\} \\
 & \Rightarrow \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = t \\ \omega_M^3(\Psi_2) = f \end{array} \right\} \\
 & \Rightarrow \omega_M^3(\Psi_1 \wedge \Psi_2) = f
 \end{aligned}$$

Analogous for the case where  $\omega_M^3(\Psi_1) = f$  and  $\omega_M^3(\Psi_2) = t$ .

$$\begin{aligned}
 \text{or } & \Rightarrow \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, b) \in M \\ \text{truthValue}(\Psi_2, f) \in M \end{array} \right\} \\
 & \Rightarrow \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = b \\ \omega_M^3(\Psi_2) = f \end{array} \right\} \\
 & \Rightarrow \omega_M^3(\Psi_1 \wedge \Psi_2) = f
 \end{aligned}$$

Analogous for the case where  $\omega_M^3(\Psi_1) = f$  and  $\omega_M^3(\Psi_2) = b$ .

$$\begin{aligned}
 \text{or } & \Rightarrow \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, f) \in M \\ \text{truthValue}(\Psi_2, f) \in M \end{array} \right\} \\
 & \Rightarrow \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = f \\ \omega_M^3(\Psi_2) = f \end{array} \right\} \\
 & \Rightarrow \omega_M^3(\Psi_1 \wedge \Psi_2) = f
 \end{aligned}$$

3. Let  $\Phi = \Psi_1 \vee \Psi_2$ .

(a)  $\theta = t$ :

- “ $\Rightarrow$ ”:

$$\begin{aligned}
 \omega_M^3(\Psi_1 \vee \Psi_2) = \theta & \xrightarrow{\theta=t} \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = t \\ \omega_M^3(\Psi_2) = f \end{array} \right\} \\
 & \xrightarrow{\text{I.H.}} \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, t) \in M \\ \text{truthValue}(\Psi_2, f) \in M \end{array} \right\} \\
 & \xrightarrow{(\text{AC13})} \text{truthValue}(\Psi_1 \vee \Psi_2, t) \in M
 \end{aligned}$$

Analogous for the case where  $\omega_M^3(\Psi_1) = f$  and  $\omega_M^3(\Psi_2) = t$ .

$$\begin{aligned}
 \text{or } & \xrightarrow{\theta=t} \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = t \\ \omega_M^3(\Psi_2) = b \end{array} \right\} \\
 & \xrightarrow{\text{I.H.}} \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, t) \in M \\ \text{truthValue}(\Psi_2, b) \in M \end{array} \right\} \\
 & \xrightarrow{(\text{AC13})} \text{truthValue}(\Psi_1 \vee \Psi_2, t) \in M
 \end{aligned}$$

Analogous for the case where  $\omega_M^3(\Psi_1) = b$  and  $\omega_M^3(\Psi_2) = t$ .

$$\begin{aligned}
 \text{or } & \xrightarrow{\theta=t} \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = t \\ \omega_M^3(\Psi_2) = t \end{array} \right\} \\
 & \xrightarrow{\text{I.H.}} \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, t) \in M \\ \text{truthValue}(\Psi_2, t) \in M \end{array} \right\} \\
 & \xrightarrow{(\text{AC13})} \text{truthValue}(\Psi_1 \vee \Psi_2, t) \in M
 \end{aligned}$$

- “ $\Leftarrow$ ”:

$$\begin{aligned}
 \text{truthValue}(\Psi_1 \vee \Psi_2, t) \in M & \Rightarrow \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, t) \in M \\ \text{truthValue}(\Psi_2, f) \in M \end{array} \right\} \\
 & \Rightarrow \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = t \\ \omega_M^3(\Psi_2) = f \end{array} \right\} \\
 & \Rightarrow \omega_M^3(\Psi_1 \vee \Psi_2) = t
 \end{aligned}$$

Analogous for the case where  $\omega_M^3(\Psi_1) = f$  and  $\omega_M^3(\Psi_2) = t$ .

$$\begin{aligned}
 \text{or } & \Rightarrow \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, t) \in M \\ \text{truthValue}(\Psi_2, b) \in M \end{array} \right\} \\
 & \Rightarrow \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = t \\ \omega_M^3(\Psi_2) = b \end{array} \right\} \\
 & \Rightarrow \omega_M^3(\Psi_1 \vee \Psi_2) = t
 \end{aligned}$$

Analogous for the case where  $\omega_M^3(\Psi_1) = b$  and  $\omega_M^3(\Psi_2) = t$ .

$$\begin{aligned}
 \text{or } & \Rightarrow \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, t) \in M \\ \text{truthValue}(\Psi_2, t) \in M \end{array} \right\} \\
 & \Rightarrow \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = t \\ \omega_M^3(\Psi_2) = t \end{array} \right\} \\
 & \Rightarrow \omega_M^3(\Psi_1 \vee \Psi_2) = t
 \end{aligned}$$

- (b)  $\theta = b$ :

- “ $\Rightarrow$ ”:

$$\begin{aligned} \omega_M^3(\Psi_1 \vee \Psi_2) = \theta &\xrightarrow{\theta=b} \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = b \\ \omega_M^3(\Psi_2) = f \end{array} \right\} \\ &\xrightarrow{\text{I.H.}} \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, b) \in M \\ \text{truthValue}(\Psi_2, f) \in M \end{array} \right\} \\ &\xrightarrow{(\text{AC14})} \text{truthValue}(\Psi_1 \vee \Psi_2, b) \in M \end{aligned}$$

Analogous for the case where  $\omega_M^3(\Psi_1) = f$  and  $\omega_M^3(\Psi_2) = b$ .

$$\begin{aligned} \text{or } &\xrightarrow{\theta=b} \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = b \\ \omega_M^3(\Psi_2) = b \end{array} \right\} \\ &\xrightarrow{\text{I.H.}} \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, b) \in M \\ \text{truthValue}(\Psi_2, b) \in M \end{array} \right\} \\ &\xrightarrow{(\text{AC14})} \text{truthValue}(\Psi_1 \vee \Psi_2, b) \in M \end{aligned}$$

- “ $\Leftarrow$ ”:

$$\begin{aligned} \text{truthValue}(\Psi_1 \vee \Psi_2, b) \in M &\Rightarrow \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, b) \in M \\ \text{truthValue}(\Psi_2, f) \in M \end{array} \right\} \\ &\Rightarrow \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = b \\ \omega_M^3(\Psi_2) = f \end{array} \right\} \\ &\Rightarrow \omega_M^3(\Psi_1 \vee \Psi_2) = b \end{aligned}$$

Analogous for the case where  $\omega_M^3(\Psi_1) = f$  and  $\omega_M^3(\Psi_2) = b$ .

$$\begin{aligned} \text{or } &\Rightarrow \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, b) \in M \\ \text{truthValue}(\Psi_2, b) \in M \end{array} \right\} \\ &\Rightarrow \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = b \\ \omega_M^3(\Psi_2) = b \end{array} \right\} \\ &\Rightarrow \omega_M^3(\Psi_1 \vee \Psi_2) = b \end{aligned}$$

(c)  $\theta = f$ :

- “ $\Rightarrow$ ”:

$$\begin{aligned} \omega_M^3(\Psi_1 \vee \Psi_2) = \theta &\xrightarrow{\theta=f} \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = f \\ \omega_M^3(\Psi_2) = f \end{array} \right\} \\ &\xrightarrow{\text{I.H.}} \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, f) \in M \\ \text{truthValue}(\Psi_2, f) \in M \end{array} \right\} \\ &\xrightarrow{(\text{AC12})} \text{truthValue}(\Psi_1 \vee \Psi_2, f) \in M \end{aligned}$$

- “ $\Leftarrow$ ”:

$$\begin{aligned} \text{truthValue}(\Psi_1 \vee \Psi_2, f) \in M &\Rightarrow \left\{ \begin{array}{l} \text{truthValue}(\Psi_1, f) \in M \\ \text{truthValue}(\Psi_2, f) \in M \end{array} \right\} \\ &\Rightarrow \left\{ \begin{array}{l} \omega_M^3(\Psi_1) = f \\ \omega_M^3(\Psi_2) = f \end{array} \right\} \\ &\Rightarrow \omega_M^3(\Psi_1 \vee \Psi_2) = f \end{aligned}$$

4. Let  $\Phi = \neg\Psi$ .

(a)  $\theta = t$ :

- “ $\Rightarrow$ ”:

$$\begin{aligned} \omega_M^3(\neg\Psi) = \theta &\xrightarrow{\theta=t} \omega_M^3(\neg\Psi) = t \\ &\Rightarrow \omega_M^3(\Psi) = f \\ &\xrightarrow{\text{I.H.}} \text{truthValue}(\Psi, f) \in M \\ &\xrightarrow{(\text{AC15})} \text{truthValue}(\neg\Psi, t) \in M \end{aligned}$$

- “ $\Leftarrow$ ”:

$$\begin{aligned} \text{truthValue}(\neg\Psi, t) \in M &\Rightarrow \text{truthValue}(\Psi, f) \in M \\ &\Rightarrow \omega_M^3(\Psi) = f \\ &\Rightarrow \omega_M^3(\neg\Psi) = t \end{aligned}$$

(b)  $\theta = b$ :

- “ $\Rightarrow$ ”:

$$\begin{aligned} \omega_M^3(\neg\Psi) = \theta &\xrightarrow{\theta=b} \omega_M^3(\neg\Psi) = b \\ &\Rightarrow \omega_M^3(\Psi) = b \\ &\xrightarrow{\text{I.H.}} \text{truthValue}(\Psi, b) \in M \\ &\xrightarrow{(\text{AC17})} \text{truthValue}(\neg\Psi, b) \in M \end{aligned}$$

- “ $\Leftarrow$ ”:

$$\begin{aligned} \text{truthValue}(\neg\Psi, b) \in M &\Rightarrow \text{truthValue}(\Psi, b) \in M \\ &\Rightarrow \omega_M^3(\Psi) = b \\ &\Rightarrow \omega_M^3(\neg\Psi) = b \end{aligned}$$

(c)  $\theta = f$ :

- “ $\Rightarrow$ ”:

$$\begin{aligned}
 \omega_M^3(\neg\Psi) = \theta &\xrightarrow{\theta=f} \Rightarrow \omega_M^3(\neg\Psi) = f \\
 &\Rightarrow \omega_M^3(\Psi) = t \\
 &\xrightarrow{\text{I.H.}} \mathbf{truthValue}(\Psi, t) \in M \\
 &\xrightarrow{(\text{AC16})} \mathbf{truthValue}(\neg\Psi, f) \in M
 \end{aligned}$$

- “ $\Leftarrow$ ”:

$$\begin{aligned}
 \mathbf{truthValue}(\neg\Psi, f) \in M &\Rightarrow \mathbf{truthValue}(\Psi, t) \in M \\
 &\Rightarrow \omega_M^3(\Psi) = t \\
 &\Rightarrow \omega_M^3(\neg\Psi) = f
 \end{aligned}$$

It has been shown that  $\omega_M^3(\Phi) = \theta$  if and only if  $\mathbf{truthValue}(\Phi, \theta) \in M$ . The integrity constraint (AC23) ensures that  $\mathbf{truthValue}(A, f)$  is never true wrt. each formula  $A \in \mathcal{K}$ . In other words,  $\mathbf{truthValue}(A, f) \notin M$  for all  $A \in \mathcal{K}$ . Consequently,  $\omega_M^3(A) \neq f$  for all  $A \in \mathcal{K}$ . Thus,  $\omega_M^3(A)$  evaluates to either  $t$  or  $b$  in all cases. This corresponds to the definition of a model in Priest’s three-valued logic. Hence,  $\omega_M^3 \in \text{Models}(\mathcal{K})$ .  $\square$

**Lemma 10.** *If  $\omega^3 \in \text{Models}(\mathcal{K})$  then  $P'_c(\mathcal{K})$  has an answer set  $M$  such that  $\omega_M^3 = \omega^3$ .*

*Proof.* Analogous to the proof of Lemma 9 in reverse.  $\square$

**Theorem 7.** *Let  $M_o$  be an optimal answer set of  $P_c(\mathcal{K})$ . Then  $|(\omega_{M_o}^3)^{-1}(b)| = \mathcal{I}_c(\mathcal{K})$ .*

*Proof.* From Lemma 9 and Lemma 10 it follows that the answer sets of  $P'_c(\mathcal{K})$  correspond exactly to the 3-valued models of  $\mathcal{K}$ . The minimize statement included in  $P_c(\mathcal{K})$  (AC20) then ensures that only a minimal number of instances  $\mathbf{truthValue}(X, b)$  are included in the answer set. This corresponds to the minimal number of atoms  $X \in \text{At}(\mathcal{K})$  being evaluated to  $b$  which is exactly the definition of  $\mathcal{I}_c(\mathcal{K})$ . Note that, following (AC2), for all  $X \in \text{At}(\mathcal{K})$ ,  $\mathbf{atom}(X)$  must be included in any answer set of  $P_c(\mathcal{K})$ .  $\square$

### A.2.2 THE FORGETTING-BASED INCONSISTENCY MEASURE

For every $A \in \mathcal{K}$ :	
<code>kbMember(A)</code> .	(AF1)
For every formula $\Phi_a$ which consists of a single atom occurrence $X^l$ :	
<code>formulaIsAtomOcc(Φ, X, l)</code> .	(AF2)
For every conjunction $\Phi_c = \Psi_{c,1} \wedge \Psi_{c,2}$ appearing in some formula:	
<code>conjunction(Φ<sub>c</sub>, Ψ<sub>c,1</sub>, Ψ<sub>c,2</sub>)</code> .	(AF3)
For every disjunction $\Phi_d = \Psi_{d,1} \vee \Psi_{d,2}$ appearing in some formula:	
<code>disjunction(Φ<sub>d</sub>, Ψ<sub>d,1</sub>, Ψ<sub>d,2</sub>)</code> .	(AF4)
For each negation $\Phi_n = \neg\Psi_n$ appearing in some formula:	
<code>negation(Φ<sub>n</sub>, Ψ<sub>n</sub>)</code> .	(AF5)

For the static part, we define:

$tv(t;f)$ . (AF6)

$atomOcc(A,L) :-$   
 $formulaIsAtomOcc(_,A,L)$ . (AF7)

$atom(A) :-$   
 $atomOcc(A,_)$ . (AF8)

$1\{truthValue(A,T) : tv(T)\}1 :-$   
 $atom(A)$ . (AF9)

$truthValue(F,t) :-$   
 $conjunction(F,G,H),$   
 $truthValue(G,t),$   
 $truthValue(H,t)$ . (AF10)

$truthValue(F,f) :-$   
 $conjunction(F,-,-),$   
 $not\ truthValue(F,t)$ . (AF11)

$truthValue(F,f) :-$   
 $disjunction(F,G,H),$   
 $truthValue(G,f),$   
 $truthValue(H,f)$ . (AF12)

$truthValue(F,t) :-$   
 $disjunction(F,-,-),$   
 $not\ truthValue(F,f)$ . (AF13)

$truthValue(F,t) :-$   
 $negation(F,G),$   
 $truthValue(G,f)$ . (AF14)

$truthValue(F,f) :-$   
 $negation(F,G),$   
 $truthValue(G,t)$ . (AF15)

$\{atomOccForgotten(A,L)\} :-$   
 $atomOcc(A,L)$ . (AF16)

$truthValue(F,T) :-$   
 $formulaIsAtomOcc(F,A,L),$   
 $truthValue(A,T),$   
 $not\ atomOccForgotten(A,L)$ . (AF17)

$truthValue(F,t) :-$   
 $formulaIsAtomOcc(F,A,L),$   
 $truthValue(A,f),$   
 $atomOccForgotten(A,L)$ . (AF18)

$truthValue(F,f) :-$

<pre> formulaIsAtomOcc(F,A,L), truthValue(A,t), atomOccForgotten(A,L). </pre>	(AF19)
<pre> :-     truthValue(F,f),     kbMember(F). </pre>	(AF20)
<pre> #minimize{1,A,L : atomOccForgotten(A,L)}. </pre>	(AF21)

Encoding 8: Overview of the ASP encoding  $P_f(\mathcal{K})$  of the forgetting-based inconsistency measure  $\mathcal{I}_f$ .  $\mathcal{K}$  is the given knowledge base.

Let  $\mathcal{K}$  be an arbitrary knowledge base and  $P_f(\mathcal{K})$  the extended logic program consisting of rules (AF1)–(AF21) included Encoding 8. Let  $P'_f(\mathcal{K})$  denote the extended logic program  $P_f(\mathcal{K})$  *without* the minimize statement (AF21).

For an answer set  $M$  we define:

$$F_M = \{X^l \mid X \in \text{At}(\mathcal{K}), \text{atomOccForgotten}(X, l) \in M\}$$

**Lemma 11.** *If  $M$  is an answer set of  $P'_f(\mathcal{K})$  then  $(\bigwedge \mathcal{K})[X^1 \xrightarrow{f} \top, \perp; \dots; X^n \xrightarrow{f} \top, \perp]$  is consistent with  $F_M = \{X^1, \dots, X^n\}$ .*

*Proof.* The rules given in (AF10)–(AF15) model classical entailment wrt.  $\wedge$ ,  $\vee$ , and  $\neg$ . To be specific, the rules (AF10) and (AF11) model that the conjunction of two subformulas is only true if both conjuncts are true, and false if it is not true. In the same fashion, a disjunction is only false if both of its individual disjuncts are false (modeled by (AF12)). It is true, if it is not false (modeled by (AF13)). Classical negation, i. e.,  $\neg\Phi$  is true if  $\Phi$  is false, and vice versa, is represented by (AF14) and (AF15).

Furthermore, rule (AF9) represents that each atom  $X$  is evaluated to either  $t$  or  $f$ . In other words, since we represent that each atom  $X \in \text{At}(\mathcal{K})$  is assigned a unique truth value, we automatically represent an *interpretation* on  $\text{At}(\mathcal{K})$ .

Each atom occurrence  $X$  with label  $l$  (i. e.,  $X^l$ ) is represented as `atomOcc`( $X, l$ ), due to (AF2) and (AF7). Further, rule (AF16) models that each  $X^l$  could be forgotten by guessing a subset of atom occurrences by means of a cardinality constraint. More precisely, (AF16) assigns a subset of all atom occurrences  $X^l$  as `atomOccForgotten`( $X, l$ ). Thus, (AF16) determines  $F_M$ . For each atom occurrence in this set, the truth value of the subformula containing the atom occurrence is reversed by (AF18) and (AF19), respectively. Given an interpretation  $\omega$ , this results in setting an atom occurrence  $X^l$  with  $\omega(X) = f$  to  $t$  (AF18), and vice versa (AF19). If an atom occurrence  $X^l$  is not in  $F_M$  (i. e., if an atom is not forgotten), (AF17) is triggered, and  $X^l$  must conform to  $\omega$ .

Finally, all formulas in the given knowledge base must be true. This is ensured by the integrity constraint (AF20) which forbids elements  $A \in \mathcal{K}$  to be evaluated to  $f$ . This also implies that each  $X^l \in F_M$  must be assigned the opposite truth value of  $X$ , which is equivalent to replacing it with  $\top$  or  $\perp$ , corresponding on the context. Thus, since all formulas  $A \in \mathcal{K}$  must be true, the version of  $\mathcal{K}$  modified by the forgetting operation, i. e.,  $(\bigwedge \mathcal{K})[X^1 \xrightarrow{f} \top, \perp; \dots; X^n \xrightarrow{f} \top, \perp]$ , must be consistent.  $\square$

**Theorem 8.** *Let  $M_o$  be an optimal answer set of  $P_f(\mathcal{K})$ . Then  $|F_{M_o}| = \mathcal{I}_f(\mathcal{K})$ .*

*Proof.* From Lemma 11 it follows that the answer sets of  $P'_f(\mathcal{K})$  correspond exactly to  $(\bigwedge \mathcal{K})[X^1 \xrightarrow{f} \top, \perp; \dots; X^n \xrightarrow{f} \top, \perp]$  being consistent with  $F_M = \{X^1, \dots, X^n\}$ . The minimize statement (AF21) included in  $P_f(\mathcal{K})$  then ensures that only a minimal number of instances of `atomOccForgotten( $X, I$ )` are included in the optimal answer set. Overall, this corresponds to the minimal number of atom occurrences in  $\mathcal{K}$  being forgotten, which is exactly the definition of  $\mathcal{I}_f(\mathcal{K})$ .  $\square$

### A.2.3 THE HITTING SET INCONSISTENCY MEASURE

For every $A \in \mathcal{K}$ :	
<code>kbMember(<math>A</math>).</code>	(AH1)
For every $X \in \text{At}(\mathcal{K})$ :	
<code>atom(<math>X</math>).</code>	(AH2)
Define $ \mathcal{K} $ interpretations:	
<code>interpretation(1..<math> \mathcal{K} </math>).</code>	(AH3)
For every conjunction $\Phi_c = \Psi_{c,1} \wedge \Psi_{c,2}$ appearing in some formula:	
<code>conjunction(<math>\Phi_c, \Psi_{c,1}, \Psi_{c,2}</math>).</code>	(AH4)
For every disjunction $\Phi_d = \Psi_{d,1} \vee \Psi_{d,2}$ appearing in some formula:	
<code>disjunction(<math>\Phi_d, \Psi_{d,1}, \Psi_{d,2}</math>).</code>	(AH5)
For each negation $\Phi_n = \neg \Psi_n$ appearing in some formula:	
<code>negation(<math>\Phi_n, \Psi_n</math>).</code>	(AH6)
For every formula $\Phi_a$ which consists of a single atom $X$ :	
<code>formulaIsAtom(<math>\Phi_a, X</math>).</code>	(AH7)
At least one, and at most $ \mathcal{K} $ interpretations must be included in the hitting set:	
<code>1{interpretationActive(I) : interpretation(I)}<math> \mathcal{K} </math>.</code>	(AH8)
For the static part, we define:	
<code>tv(<math>t; f</math>).</code>	(AH9)
<code>1{truthValueInt(<math>A, I, T</math>) : tv(T)}1 :-     atom(<math>A</math>),     interpretation(I).</code>	(AH10)
<code>truthValueInt(<math>F, I, t</math>) :-     conjunction(<math>F, G, H</math>),     interpretation(I),     truthValueInt(<math>G, I, t</math>),     truthValueInt(<math>H, I, t</math>).</code>	(AH11)
<code>truthValueInt(<math>F, I, f</math>) :-     conjunction(<math>F, -, -</math>),     interpretation(I),     not truthValueInt(<math>F, I, t</math>).</code>	(AH12)
<code>truthValueInt(<math>F, I, f</math>) :-</code>	

```

        disjunction(F,G,H),
        interpretation(I),
        truthValueInt(G,I,f),
        truthValueInt(H,I,f).
    (AH13)

truthValueInt(F,I,t) :-
    disjunction(F,-,_),
    interpretation(I),
    not truthValueInt(F,I,f).
    (AH14)

truthValueInt(F,I,t) :-
    negation(F,G),
    truthValueInt(G,I,f).
    (AH15)

truthValueInt(F,I,f) :-
    negation(F,G),
    truthValueInt(G,I,t).
    (AH16)

truthValueInt(F,I,T) :-
    formulaIsAtom(F,G),
    truthValueInt(G,I,T),
    interpretation(I),
    tv(T).
    (AH17)

truthValue(F,t) :-
    truthValueInt(F,I,t),
    kbMember(F),
    interpretation(I),
    interpretationActive(I).
    (AH18)

truthValue(F,f) :-
    kbMember(F),
    not truthValue(F,t).
    (AH19)

:-
    interpretationActive(I),
    1 < I,
    not interpretationActive(I-1).
    (AH20)

:-
    truthValue(F,f),
    kbMember(F).
    (AH21)

#minimize{1,I : interpretationActive(I)}.
    (AH22)

```

Encoding 9: Overview of the ASP encoding  $P_h(\mathcal{K})$  of the hitting set inconsistency measure  $\mathcal{I}_h$ .  $\mathcal{K}$  is the given knowledge base.

Let  $\mathcal{K}$  be an arbitrary knowledge base and  $P_h(\mathcal{K})$  the extended logic program consisting of rules (AH1)–(AH22) listed in Encoding 9 in Section 4.4.

**Lemma 12.**  $P_h(\mathcal{K})$  does not have an answer set if and only if  $\mathcal{K}$  contains one or more contradictory formulas.

*Proof.* Let some  $A^\perp \in \mathcal{K}$  be contradictory. To begin with, the formulas  $A \in \mathcal{K}$  and the atoms  $X \in \text{At}(\mathcal{K})$  are represented in ASP via (AH1) and (AH2). Moreover, the operators of each (sub-)formula in  $\mathcal{K}$  is encoded using (AH4)–(AH7). Thus,  $P'_h(\mathcal{K})$  contains an ASP representation of  $\mathcal{K}$  itself. Additionally, (AH3) adds  $|\mathcal{K}|$  interpretation constants to  $P'_h(\mathcal{K})$ . By means of (AH10), we assign a distinct truth value (defined by (AH9)) to each `atom/1` for each `interpretation/1`. (AH11)–(AH16) model classical entailment (analogous to (AF11)–(AF16), but with an additional reference to an interpretation). Besides, (AH17) models that a formula consisting of an individual atom gets the same truth value as the atom. Since every atom representation in ASP is connected to a truth value, and the evaluation of the (sub-)formulas is accomplished via (AH11)–(AH17), every formula  $A \in \mathcal{K}$  (represented by `kbMember/1`) evaluates to either  $t$  or  $f$ . Since we cannot construct an interpretation which satisfies  $A^\perp$ , it would always evaluate to  $f$ . However, the integrity constraint (AH21) forbids this, therefore, no answer set can be constructed.  $\square$

Let  $P'_h(\mathcal{K})$  denote the extended logic program  $P_h(\mathcal{K})$  without the minimize statement (AH22). Let  $M$  be an answer set of  $P'_h(\mathcal{K})$ . Further, let  $\omega_i^{\text{asp}}$ ,  $i \in \{1, \dots, |\mathcal{K}|\}$ , denote the ASP constants representing interpretations, and let  $\omega_i$  denote the corresponding interpretations in classical logic. We define the set of interpretations  $\omega_i$  represented in  $M$  as

$$\Omega(M) = \{\omega_i \mid \text{interpretationActive}(\omega_i^{\text{asp}}) \in M\}.$$

**Lemma 13.** If  $M$  is an answer set of  $P'_h(\mathcal{K})$  then  $\Omega(M)$  is a hitting set of  $\mathcal{K}$ .

*Proof.*  $\Omega(M)$  is a hitting set if and only if every formula  $A \in \mathcal{K}$  is satisfied by at least one  $\omega_i \in \Omega(M)$ . We will first show that for any  $\omega_i \in \Omega(M)$  and  $A \in \mathcal{K}$ ,  $\omega_i \models A$  if and only if `interpretationActive`( $\omega_i^{\text{asp}}$ )  $\in M$ .

The cardinality constraint (AH8) ensures that the answer set includes at least one instance of `interpretationActive`( $\omega_i^{\text{asp}}$ ). According to (AH8), `interpretation`( $\omega_i$ ) must be included in  $M$  as well, however, this is always the case due to (AH3). Moreover, `interpretationActive`( $\omega_i^{\text{asp}}$ ) must be included in  $M$  in order for (AH19) to be satisfied. Further, for each  $A \in \mathcal{K}$  there must be at least one interpretation such that `truthValue`( $A, t$ )  $\in M$ , because otherwise, `truthValue`( $A, f$ )  $\in M$  (due to (AH19)), and this is forbidden by the integrity constraint (AH21).

As mentioned before, rules (AH11)–(AH17) represent entailment in classical propositional logic wrt. a given interpretation. Finally, an atom  $X \in \mathcal{K}$  can only be either true or false, but not both, under a given interpretation, as stated in (AH14).

Consequently, for each formula  $A \in \mathcal{K}$  there exists at least one interpretation  $\omega_i$  such that `truthValueInt`( $A, \omega_i^{\text{asp}}, t$ )  $\in M$ . It follows that for each  $A \in \mathcal{K}$ , there is also at least one `interpretationActive`( $\omega_i^{\text{asp}}$ )  $\in M$ . Hence, for each formula  $A \in \mathcal{K}$  there exists an  $\omega_i \in \Omega(M)$  such that  $\omega_i \models A$ . Thus,  $\Omega(M)$  is a hitting set of  $\mathcal{K}$ .  $\square$

**Theorem 9.** Let  $M_o$  be an optimal answer set of  $P_h(\mathcal{K})$ . Then  $|\Omega(M_o)| - 1 = \mathcal{I}_h(\mathcal{K})$ . If no answer set of  $P_h(\mathcal{K})$  exists,  $\mathcal{I}_h(\mathcal{K}) = \infty$ .

*Proof.* Lemma 12 shows that  $P_h(\mathcal{K})$  does not have an answer set if  $\mathcal{K}$  contains at least one contradictory formula, i. e., if  $\mathcal{I}_h(\mathcal{K}) = \infty$ .

Further, in the above proof of Lemma 13 we showed that  $\Omega(M)$  is a hitting set of  $\mathcal{K}$  wrt. some answer set  $M$ . The cardinality of  $\Omega(M)$  is exactly the number of instances `interpretationActive`( $\omega_i^{\text{asp}}$ ), with  $i \in \{1, \dots, |\mathcal{K}|\}$ , that are included in  $M$ . The minimize statement (AH22) selects the minimal number of such instances. This represents the selection of the smallest possible hitting set wrt.  $\mathcal{K}$ . At last, subtracting 1 from  $|\Omega(M_o)|$  corresponds exactly to the definition of  $\mathcal{I}_h(\mathcal{K})$ .  $\square$

#### A.2.4 THE MAX-DISTANCE INCONSISTENCY MEASURE

For every $A_i \in \mathcal{K}$ with $i \in \{0, \dots,  \mathcal{K}  - 1\}$ :	
<code>kbMember(A, i).</code>	(ADM1)
For every $X \in \text{At}(\mathcal{K})$ :	
<code>atom(X).</code>	(ADM2)
Define $ \mathcal{K}  + 1$ interpretations:	
<code>interpretation(0.. \mathcal{K} ).</code>	(ADM3)
For every conjunction $\Phi_c = \Psi_{c,1} \wedge \Psi_{c,2}$ appearing in some formula:	
<code>conjunction(\Phi_c, \Psi_{c,1}, \Psi_{c,2}).</code>	(ADM4)
For every disjunction $\Phi_d = \Psi_{d,1} \vee \Psi_{d,2}$ appearing in some formula:	
<code>disjunction(\Phi_d, \Psi_{d,1}, \Psi_{d,2}).</code>	(ADM5)
For each negation $\Phi_n = \neg\Psi_n$ appearing in some formula:	
<code>negation(\Phi_n, \Psi_n).</code>	(ADM6)
For every formula $\Phi_a$ which consists of a single atom $X$ :	
<code>formulaIsAtom(\Phi_a, X).</code>	(ADM7)
Distance between interpretation $ \mathcal{K} $ and the models of each formula:	
<code>d(X) :- diff(X), interpretation(I), X &lt;= #count{A: atom(A), truthValueInt(A, I, T), not truthValueInt(A,  \mathcal{K} , T)}.</code>	(ADM8)
For the static part, we define:	
<code>tv(t; f).</code>	(ADM9)
<code>1{truthValueInt(A, I, T) : tv(T)}1 :- atom(A), interpretation(I).</code>	(ADM10)
<code>truthValueInt(F, I, t) :- conjunction(F, G, H), interpretation(I), truthValueInt(G, I, t), truthValueInt(H, I, t).</code>	(ADM11)
<code>truthValueInt(F, I, f) :-</code>	

conjunction(F, -, -), interpretation(I), not truthValueInt(F, I, t).	(ADM12)
truthValueInt(F, I, f) :- disjunction(F, G, H), interpretation(I), truthValueInt(G, I, f), truthValueInt(H, I, f).	(ADM13)
truthValueInt(F, I, t) :- disjunction(F, -, -), interpretation(I), not truthValueInt(F, I, f).	(ADM14)
truthValueInt(F, I, t) :- negation(F, G), truthValueInt(G, I, f).	(ADM15)
truthValueInt(F, I, f) :- negation(F, G), truthValueInt(G, I, t).	(ADM16)
truthValueInt(F, I, T) :- formulaIsAtom(F, G), truthValueInt(G, I, T), interpretation(I), tv(T).	(ADM17)
:- truthValueInt(F, I, f), kbMember(F, I).	(ADM18)
diff(1..X) :- X = #count{A: atom(A)}.	(ADM19)
#minimize{1, X : d(X)}.	(ADM20)

Encoding 10: Overview of the ASP encoding  $P_d^{\max}(\mathcal{K})$  of the max-distance inconsistency measure  $\mathcal{I}_d^{\max}$ .  $\mathcal{K}$  is the given knowledge base.

Let  $\mathcal{K}$  be an arbitrary knowledge base and  $P_d^{\max}(\mathcal{K})$  the extended logic program consisting of rules (ADM1)–(ADM22) listed in Encoding 10 in Section 4.5.

**Lemma 14.**  $P_d^{\max}(\mathcal{K})$  does not have an answer set if and only if  $\mathcal{K}$  contains one or more contradictory formulas.

*Proof.* The knowledge base, i.e., its formulas, subformulas, and atoms are modelled by (ADM1)–(ADM2) and (ADM4)–(ADM7). Note that each  $A \in \mathcal{K}$  gets a label  $i$ , and is represented as  $\text{kbMember}(A, i)$  (ADM1). The two truth values are represented by (ADM9).

Moreover, via (ADM3) we represent  $|\mathcal{K}| + 1$  interpretations in  $P_d^{\max}(\mathcal{K})$ . We denote these interpretations as follows (note that we represent  $\omega_i$  simply as  $i$  in ASP):

$$\omega_i(X) = \begin{cases} t & \mathbf{truthValue}(X, i, t) \in M \\ f & \mathbf{truthValue}(X, i, f) \in M \end{cases}$$

Analogously to (AH10) (wrt.  $\mathcal{I}_h$ ), (ADM10) guesses all  $|\mathcal{K}| + 1$  interpretations by assigning each atom in each interpretation a distinct truth value. (ADM11)–(ADM17) correspond exactly to (AH11)–(AH17) (wrt.  $\mathcal{I}_h$ ). Thus, these rules model classical entailment.

Each interpretation  $\omega_i$  is supposed to be a model for formula  $A_i \in \mathcal{K}$ ,  $i \in \{0, \dots, |\mathcal{K}| - 1\}$ . To this end, the integrity constraint (ADM18) ensures that no formula  $A_i$  evaluates to  $f$  under  $\omega_i$ . If a contradictory formula  $A_i^\perp$  is included in  $\mathcal{K}$ , there exists no interpretation  $\omega_i$  that would satisfy  $A_i^\perp$ . Hence, the answer set would need to include  $\mathbf{truthValueInt}(A_i^\perp, i, f)$ . However, (ADM18) does not allow this, so no answer set can be derived.  $\square$

Let  $P_d^{\max'}(\mathcal{K})$  denote the extended logic program  $P_d^{\max}(\mathcal{K})$  *without* the minimize statement (ADM20). Let  $D_M^{\max} = \{n \mid \mathbf{d}(n) \in M, n \in \mathbb{N}\}$  wrt. an answer set  $M$ .

**Lemma 15.** *Let  $M$  be an answer set of  $P_d^{\max'}(\mathcal{K})$ . Each  $n \in D_M^{\max}$  corresponds to a value smaller than or equal to the distance between some interpretation  $\omega_{|\mathcal{K}|}$  and one of the models of the formulas in  $\mathcal{K}$ , i. e.,  $\omega_i$  with  $i \in \{0, \dots, |\mathcal{K}| - 1\}$ .*

*Proof.* If an atom  $X \in \mathbf{At}(\mathcal{K})$  evaluates to a truth value  $\theta \in \{t, f\}$  under an interpretation  $\omega_i$ , and not to truth value  $\theta$  under interpretation  $\omega_{|\mathcal{K}|}$ , with  $i \in \{0, \dots, |\mathcal{K}| - 1\}$ , then  $\mathbf{truthValueInt}(X, i, \theta) \in M$  and  $\mathbf{truthValueInt}(X, |\mathcal{K}|, \theta) \notin M$ . We also have  $\mathbf{atom}(X) \in M$ ,  $\mathbf{interpretation}(i) \in M$ , and  $\mathbf{interpretation}(|\mathcal{K}|) \in M$ . Consequently, if for  $n$  cases of atoms  $X$  wrt. interpretations  $\omega_i$  and  $\omega_{|\mathcal{K}|}$  we have different truth values, i. e., we have  $\mathbf{truthValueInt}(X, i, \theta) \in M$ , but  $\mathbf{truthValueInt}(X, |\mathcal{K}|, \theta) \notin M$ , this is counted by means of the `#count` aggregate in (ADM8). Note that this also represents the Dalal distance between  $\omega_i$  and  $\omega_{|\mathcal{K}|}$ . Moreover,  $\{\mathbf{diff}(0), \dots, \mathbf{diff}(|\mathbf{At}(\mathcal{K})|)\} \subset M$  (ADM19). Since the Dalal distance between  $\omega_i$  and  $\omega_{|\mathcal{K}|}$  cannot exceed  $|\mathbf{At}(\mathcal{K})|$  (except for the  $\infty$  case, see Lemma 14),  $\mathbf{diff}(n) \in M$ . Moreover, via (ADM8), we do not only derive  $\mathbf{d}(n)$ , but also  $\mathbf{d}(m)$  with  $m \in \{1, \dots, n - 1\}$ . Thus, we conclude that  $\{\mathbf{d}(1), \dots, \mathbf{d}(n)\} \subset M$  if the Dalal distance between  $\omega_i$  and  $\omega_{|\mathcal{K}|}$  is  $n$ . Further, each interpretation  $\omega_i$  represents a model of the  $i$ -th formula, due to the integrity constraint (ADM18). Hence, each element in  $D_M^{\max}$  always represents a value smaller than or equal to the Dalal distance between a model of a formula in  $\mathcal{K}$  and  $\omega_{|\mathcal{K}|}$ .  $\square$

**Theorem 10.** *Let  $M_o$  be an optimal answer set of  $P_d^{\max}(\mathcal{K})$ . Then  $|D_{M_o}^{\max}| = \mathcal{I}_d^{\max}(\mathcal{K})$ . If no answer set of  $P_d^{\max}(\mathcal{K})$  exists,  $\mathcal{I}_d^{\max}(\mathcal{K}) = \infty$ .*

*Proof.* Lemma 14 shows that  $P_d^{\max}(\mathcal{K})$  does not have an answer set if  $\mathcal{K}$  contains at least one contradictory formula, i. e., if  $\mathcal{I}_d^{\max}(\mathcal{K}) = \infty$ .

Furthermore, Lemma 15 shows that the elements in  $D_M^{\max}$  represent the Dalal distances between an interpretation  $\omega_{|\mathcal{K}|}$  and a model of one of the formulas, represented by  $\omega_i$ ,  $i \in \{0, \dots, |\mathcal{K}| - 1\}$ . Let  $M_o$  be an optimal answer set of  $P_d^{\max}(\mathcal{K})$ . The minimize statement (ADM20) finds the interpretation  $(\omega_{|\mathcal{K}|})$  with the “optimal” distance to the models of the

formulas of  $\mathcal{K}$ . More precisely, the objective of the minimize statement is to find as few different distance values as possible, since it “counts” each distance. Note that each distinct distance can only be derived once (due to (ADM19)), and for every  $d(n) \in M$  with  $n \in D_{M_o}^{\max}$ , we also have  $d(m) \in M$  with  $m \in \{1, \dots, n-1\}$ . In other words, every distance that is smaller than  $n$  (but at least 1, since  $\text{diff}(0) \notin M$  (ADM19)) is also counted, since (ADM8) derives  $d(n)$  for each value that is smaller than or equal to the Dalal distance (see Lemma 15). Hence, we have  $D_{M_o}^{\max} = \{1, \dots, n_{\max}\}$  with  $n_{\max}$  corresponding to the maximal value in  $D_{M_o}^{\max}$ . Due to the minimization in (ADM20),  $n_{\max}$  corresponds exactly to the minimal maximum distance. It is easy to see that  $|D_{M_o}^{\max}| = n_{\max}$ . Therefore,  $|D_{M_o}^{\max}| = \mathcal{I}_d^{\max}(\mathcal{K})$ .  $\square$

### A.2.5 THE SUM-DISTANCE INCONSISTENCY MEASURE

For every $A_i \in \mathcal{K}$ with $i \in \{0, \dots,  \mathcal{K}  - 1\}$ :	
<code>kbMember(A, i).</code>	(ADS1)
For every $X \in \text{At}(\mathcal{K})$ :	
<code>atom(X).</code>	(ADS2)
Define $ \mathcal{K}  + 1$ interpretations:	
<code>interpretation(0.. \mathcal{K} ).</code>	(ADS3)
For every conjunction $\Phi_c = \Psi_{c,1} \wedge \Psi_{c,2}$ appearing in some formula:	
<code>conjunction(\Phi_c, \Psi_{c,1}, \Psi_{c,2}).</code>	(ADS4)
For every disjunction $\Phi_d = \Psi_{d,1} \vee \Psi_{d,2}$ appearing in some formula:	
<code>disjunction(\Phi_d, \Psi_{d,1}, \Psi_{d,2}).</code>	(ADS5)
For each negation $\Phi_n = \neg\Psi_n$ appearing in some formula:	
<code>negation(\Phi_n, \Psi_n).</code>	(ADS6)
For every formula $\Phi_a$ which consists of a single atom $X$ :	
<code>formulaIsAtom(\Phi_a, X).</code>	(ADS7)
Minimize statement:	
<code>#minimize{1, A, I: atom(A), truthValueInt(A, I, T), not truthValueInt(A,  \mathcal{K} , T)}.</code>	(ADS8)
For the static part, we define:	
<code>tv(t; f).</code>	(ADS9)
<code>1{truthValueInt(A, I, T) : tv(T)}1 :- atom(A), interpretation(I).</code>	(ADS10)
<code>truthValueInt(F, I, t) :- conjunction(F, G, H), interpretation(I), truthValueInt(G, I, t), truthValueInt(H, I, t).</code>	(ADS11)
<code>truthValueInt(F, I, f) :- conjunction(F, -, -), interpretation(I),</code>	

<code>not truthValueInt(F,I,t).</code>	(ADS12)
<code>truthValueInt(F,I,f) :-                  disjunction(F,G,H),                  interpretation(I),                  truthValueInt(G,I,f),                  truthValueInt(H,I,f).</code>	(ADS13)
<code>truthValueInt(F,I,t) :-                  disjunction(F,-,-),                  interpretation(I),                  not truthValueInt(F,I,f).</code>	(ADS14)
<code>truthValueInt(F,I,t) :-                  negation(F,G),                  truthValueInt(G,I,f).</code>	(ADS15)
<code>truthValueInt(F,I,f) :-                  negation(F,G),                  truthValueInt(G,I,t).</code>	(ADS16)
<code>truthValueInt(F,I,T) :-                  formulaIsAtom(F,G),                  truthValueInt(G,I,T),                  interpretation(I),                  tv(T).</code>	(ADS17)
<code>:-                  truthValueInt(F,I,f),                  kbMember(F,I).</code>	(ADS18)

Encoding 11: Overview of the ASP encoding  $P_d^\Sigma(\mathcal{K})$  of the sum-distance inconsistency measure  $\mathcal{I}_d^\Sigma$ .  $\mathcal{K}$  is the given knowledge base.

Let  $\mathcal{K}$  be an arbitrary knowledge base and  $P_d^\Sigma(\mathcal{K})$  the extended logic program consisting of rules (ADS1)–(ADS22) listed in Encoding 11 in Section 4.6.

**Lemma 16.**  $P_d^\Sigma(\mathcal{K})$  does not have an answer set if and only if  $\mathcal{K}$  contains one or more contradictory formulas.

*Proof.* Analogous to the proof of Lemma 14. □

We define

$$D_M^\Sigma = \{\text{diff}_M(X, i) \mid \text{truthValueInt}(X, i, \theta) \in M, \text{truthValueInt}(X, |\mathcal{K}|, \theta) \notin M\}.$$

**Theorem 11.** Let  $M_o$  be an optimal answer set of  $P_d^\Sigma(\mathcal{K})$ . Then  $|D_{M_o}^\Sigma| = \mathcal{I}_d^\Sigma(\mathcal{K})$ . If no answer set of  $P_d^\Sigma(\mathcal{K})$  exists,  $\mathcal{I}_d^\Sigma(\mathcal{K}) = \infty$ .

*Proof.* Lemma 16 shows that  $P_d^\Sigma(\mathcal{K})$  does not have an answer set if  $\mathcal{K}$  contains at least one contradictory formula, i. e., if  $\mathcal{I}_d^\Sigma(\mathcal{K}) = \infty$ .

Each formula  $A \in \mathcal{K}$  is connected to one specific interpretation. More precisely, interpretation  $\omega_i$  is a model of formula  $A_i$  (ADS18) (see also the proof for Lemma 15). Further, the inner part of the minimize statement (ADS8) represents the calculation and summation of the Dalal distance between  $\omega_{|\mathcal{K}|}$  and each interpretation  $\omega_i$  with  $i \in \{0, \dots, |\mathcal{K}| - 1\}$ . To be precise, each case in which an atom  $X$  evaluates to  $\theta \in \{t, f\}$  under interpretation  $\omega_i$ , but  $\omega_{|\mathcal{K}|} \neq \theta$ , i. e., if  $\text{truthValue}(X, i, \theta) \in M$ , but  $\text{truthValue}(X, |\mathcal{K}|, \theta) \notin M$ , is counted as 1. Thus, for any answer set  $M$ ,  $|D_M^\Sigma|$  corresponds to the sum of Dalal distances between the models of the formulas  $\omega_i$  and  $\omega_{|\mathcal{K}|}$ . For an optimal answer set  $M_o$ , yielded by the *#minimize* statement in (ADS8),  $\omega_{|\mathcal{K}|}$  becomes the “optimal” interpretation, i. e., the one with the lowest sum of distances. Hence,  $|D_{M_o}^\Sigma| = \mathcal{I}_d^\Sigma(\mathcal{K})$ .  $\square$

### A.2.6 THE HIT-DISTANCE INCONSISTENCY MEASURE

For every $A \in \mathcal{K}$ :	
<code>kbMember(A).</code>	(ADH1)
For every $X \in \text{At}(\mathcal{K})$ :	
<code>atom(X).</code>	(ADH2)
For every conjunction $\Phi_c = \Psi_{c,1} \wedge \Psi_{c,2}$ appearing in some formula:	
<code>conjunction(Φ<sub>c</sub>, Ψ<sub>c,1</sub>, Ψ<sub>c,2</sub>).</code>	(ADH3)
For every disjunction $\Phi_d = \Psi_{d,1} \vee \Psi_{d,2}$ appearing in some formula:	
<code>disjunction(Φ<sub>d</sub>, Ψ<sub>d,1</sub>, Ψ<sub>d,2</sub>).</code>	(ADH4)
For each negation $\Phi_n = \neg\Psi_n$ appearing in some formula:	
<code>negation(Φ<sub>n</sub>, Ψ<sub>n</sub>).</code>	(ADH5)
For every formula $\Phi_a$ which consists of a single atom $X$ :	
<code>formulaIsAtom(Φ<sub>a</sub>, X).</code>	(ADH6)
For the static part, we define:	
<code>tv(t;f).</code>	(ADH7)
<code>1{truthValue(A,T) : tv(T)}1 :- atom(A).</code>	(ADH8)
<code>truthValue(F,t) :- conjunction(F,G,H), truthValue(G,t), truthValue(H,t).</code>	(ADH9)
<code>truthValue(F,f) :- conjunction(F,-,-), not truthValue(F,t).</code>	(ADH10)
<code>truthValue(F,f) :- disjunction(F,G,H), truthValue(G,f), truthValue(H,f).</code>	(ADH11)
<code>truthValue(F,t) :- disjunction(F,-,-),</code>	

<code>not truthValue(F,f).</code>	(ADH12)
<code>truthValue(F,t) :- negation(F,G), truthValue(G,f).</code>	(ADH13)
<code>truthValue(F,f) :- negation(F,G), truthValue(G,t).</code>	(ADH14)
<code>truthValue(F,T) :- formulaIsAtom(F,G), truthValue(G,T), tv(T).</code>	(ADH15)
<code>truthValueKbMember(F,T) :- kbMember(F), tv(T), truthValue(F,T).</code>	(ADH16)
<code>#minimize{1,F : truthValueKbMember(F,f)}.</code>	(ADH17)

Encoding 12: Overview of the ASP encoding  $P_d^{\text{hit}}(\mathcal{K})$  of the hit-distance inconsistency measure  $\mathcal{I}_d^{\text{hit}}$ .  $\mathcal{K}$  is the given knowledge base.

Let  $\mathcal{K}$  be an arbitrary knowledge base and  $P_d^{\text{hit}}(\mathcal{K})$  the extended logic program consisting of rules (ADH1)–(ADH17) listed in Encoding 12 in Section 4.7. Let  $P_d^{\text{hit}' }(\mathcal{K})$  denote the extended logic program  $P_d^{\text{hit}}(\mathcal{K})$  without the minimize statement (ADH17).

For an answer set  $M$  of  $P_d^{\text{hit}' }(\mathcal{K})$  we define:

$$K_M = \{A^{\text{asp}} \mid A \in \mathcal{K}, \text{truthValueKbMember}(A^{\text{asp}}, f) \in M\}$$

with  $A^{\text{asp}}$  being an ASP representation of  $A$ .

**Lemma 17.** *Each  $A^{\text{asp}} \in K_M$  corresponds to a formula  $A \in \mathcal{K}$  being evaluated to  $f$ .*

*Proof.* The knowledge base, i. e., its formulas, subformulas, and atoms are modeled by (ADH1)–(ADH6). (This corresponds exactly to (AC1)–(AC6) wrt.  $\mathcal{I}_c$ .) The two truth values are represented by (ADH7). Analogously to (AF10) ( $\mathcal{I}_f$ ), (ADH8) guesses an interpretation by assigning each atom a distinct truth value. We denote this interpretation as follows:

$$\omega_M(X) = \begin{cases} t & \text{truthValue}(X,t) \in M \\ f & \text{truthValue}(X,f) \in M \end{cases}$$

(ADH9)–(ADH14) model the evaluation of conjunctions, disjunctions, and negations wrt.  $\omega_M$  (this corresponds exactly to (AF11)–(AF16) wrt.  $\mathcal{I}_f$ ). In addition, (ADH15) models that a formula consisting of an individual atom gets the same truth value as that atom. Moreover, (ADH16) essentially extracts the truth values of the formulas  $A \in \mathcal{K}$ . More precisely, if the atom valuations determined by (ADH8) eventually lead to a formula being

evaluated to a truth value  $\theta \in \{t, f\}$ —i. e., if  $\text{truthValue}(A^{\text{asp}}, \theta) \in M$  and  $A$  is an element of  $\mathcal{K}$  (i. e.,  $\text{kbMember}(A^{\text{asp}}) \in M$ ), then  $\text{truthValueKbMember}(A^{\text{asp}}, \theta) \in M$ . Hence, each  $A^{\text{asp}} \in K_M$  corresponds to a formula  $A \in \mathcal{K}$  with  $\omega_M(A) = f$ .  $\square$

**Theorem 12.** *Let  $M_o$  be an optimal answer set of  $P_d^{\text{hit}}(\mathcal{K})$ . Then  $|K_{M_o}| = \mathcal{I}_d^{\text{hit}}(\mathcal{K})$ .*

*Proof.* From Lemma 17 we know that for an answer set  $M$  of  $P_d^{\text{hit}}(\mathcal{K})$ , each  $A^{\text{asp}} \in K_M$  corresponds to a formula  $A \in \mathcal{K}$  being evaluated to false (i. e., with  $\omega_M(A) = f$ ). In  $P_d^{\text{hit}}(\mathcal{K})$ , the minimization statement (ADH17) ensures that the resulting optimal answer set  $M_o$  contains a minimal number of instances of  $\text{truthValueKbMember}(A^{\text{asp}}, f)$  (with  $A \in \mathcal{K}$ ), i. e., the cardinality of  $|K_{M_o}|$  gets minimized. Hence, we are computing the minimal number of formulas in  $\mathcal{K}$  which evaluate to  $f$  (and which would need to be removed in order to render  $\mathcal{K}$  consistent), which meets the definition of  $\mathcal{I}_d^{\text{hit}}$ .  $\square$

## Appendix B. Additional Data

This section includes additional data and visualizations regarding the experimental evaluation (Section 5).

### B.1 Runtime Measurements

The following figures display cactus plots<sup>24</sup>, each corresponding to a comparison between the three approaches wrt. a specific inconsistency measure and data set.

---

24. A cactus plot is comprised of time measurements ordered from low to high. Hence, we measure the runtime of an algorithm wrt. each knowledge base of a given data set, and subsequently order the measurements from low to high and plot them. If an inconsistency value wrt. a given knowledge base could not be computed within the given time limit (i. e., if it timed out) or if a memory error occurred, the corresponding time measurement is not plotted.

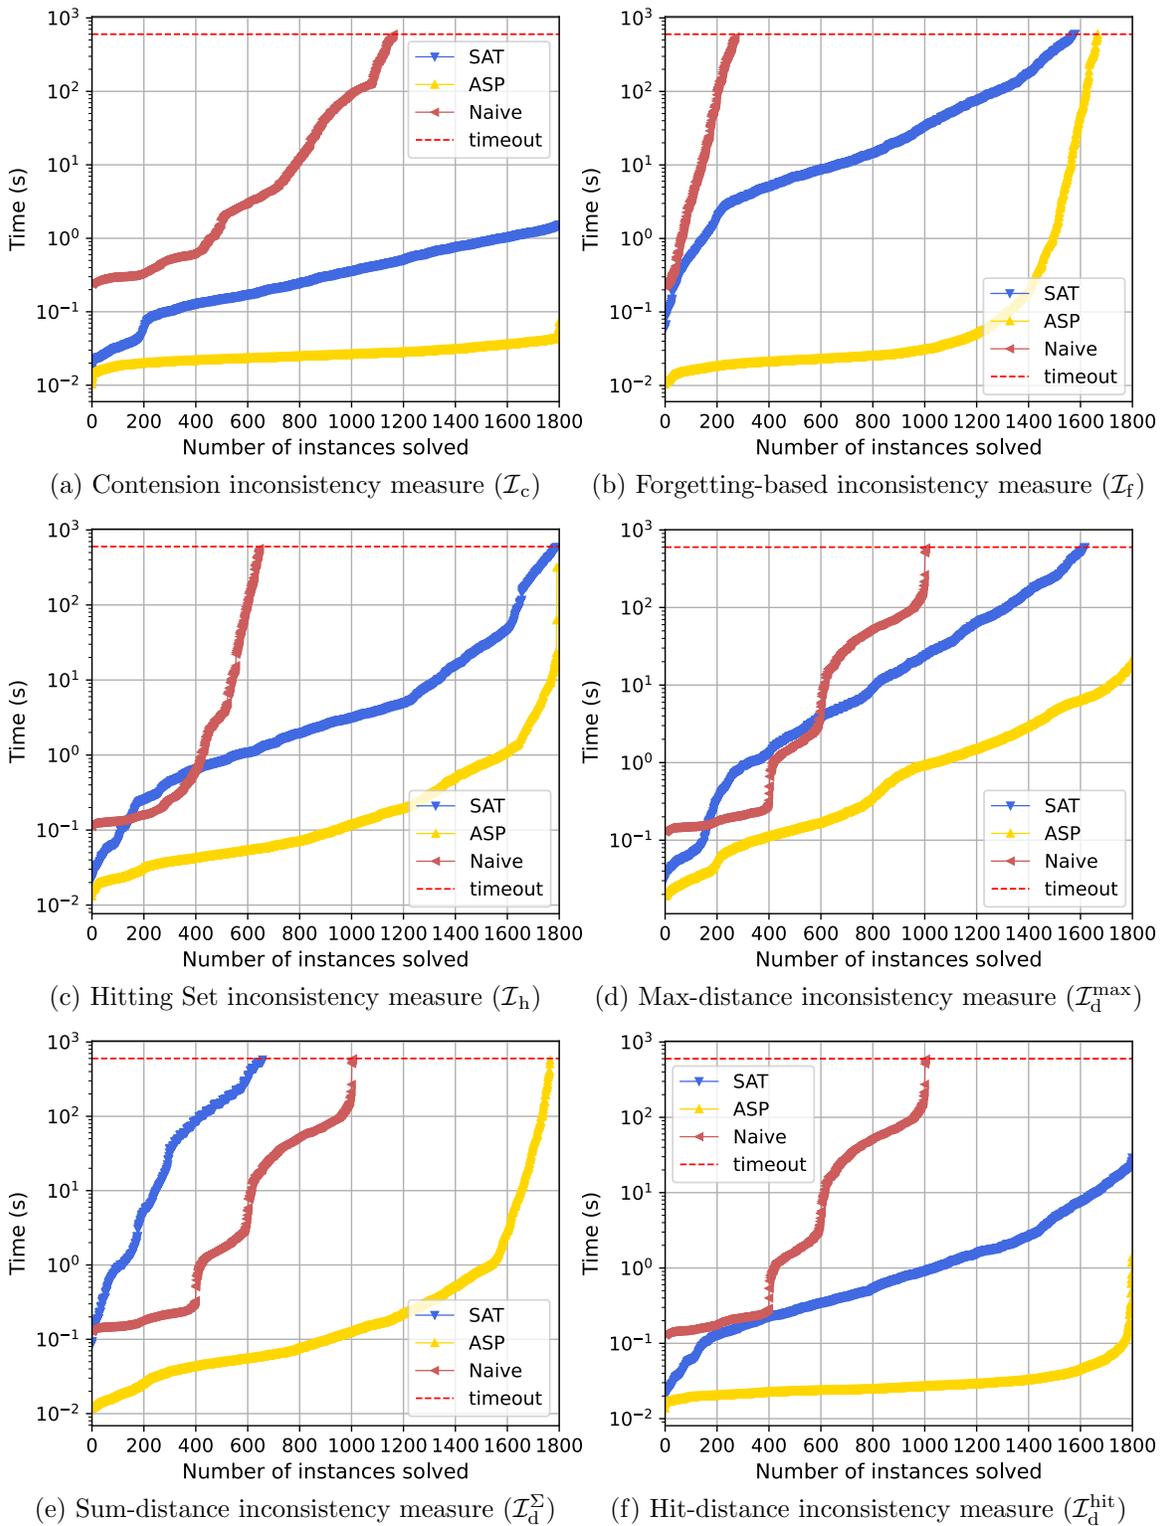


Figure 3: Runtime comparison of the ASP-based, SAT-based, and naive approaches on the SRS data set. Timeout: 10 minutes.

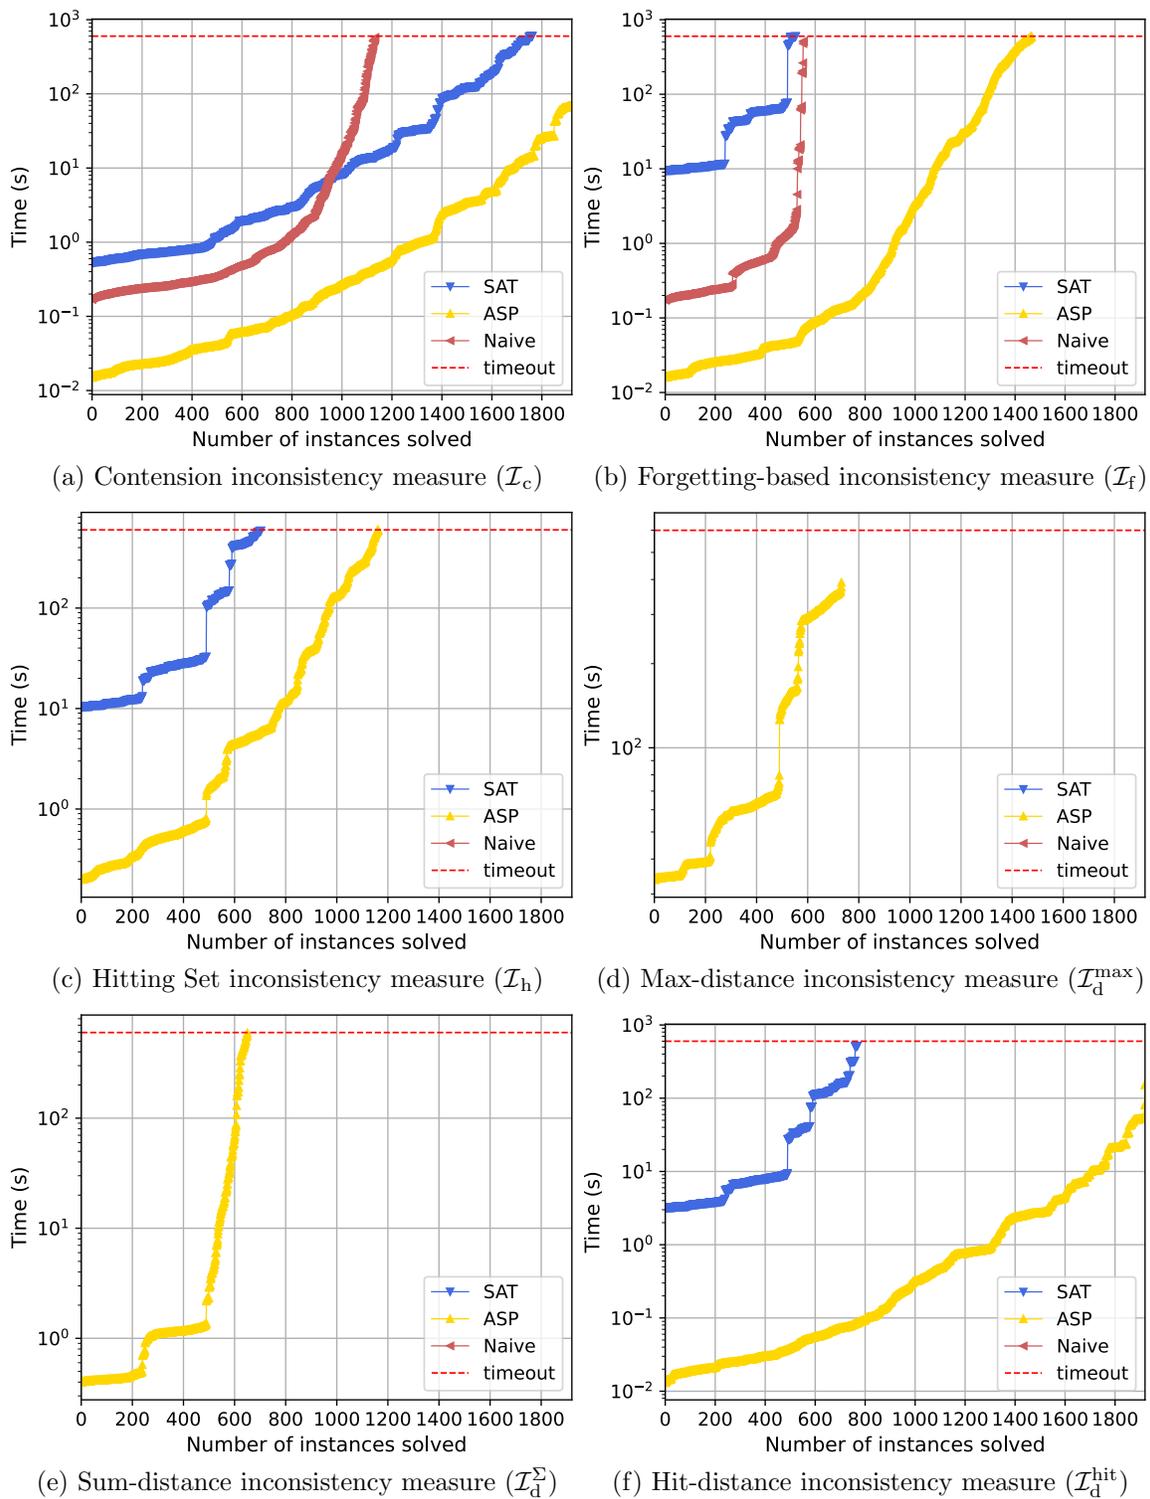


Figure 4: Runtime comparison of the ASP-based, SAT-based, and naive approaches on the ML data set. Timeout: 10 minutes.

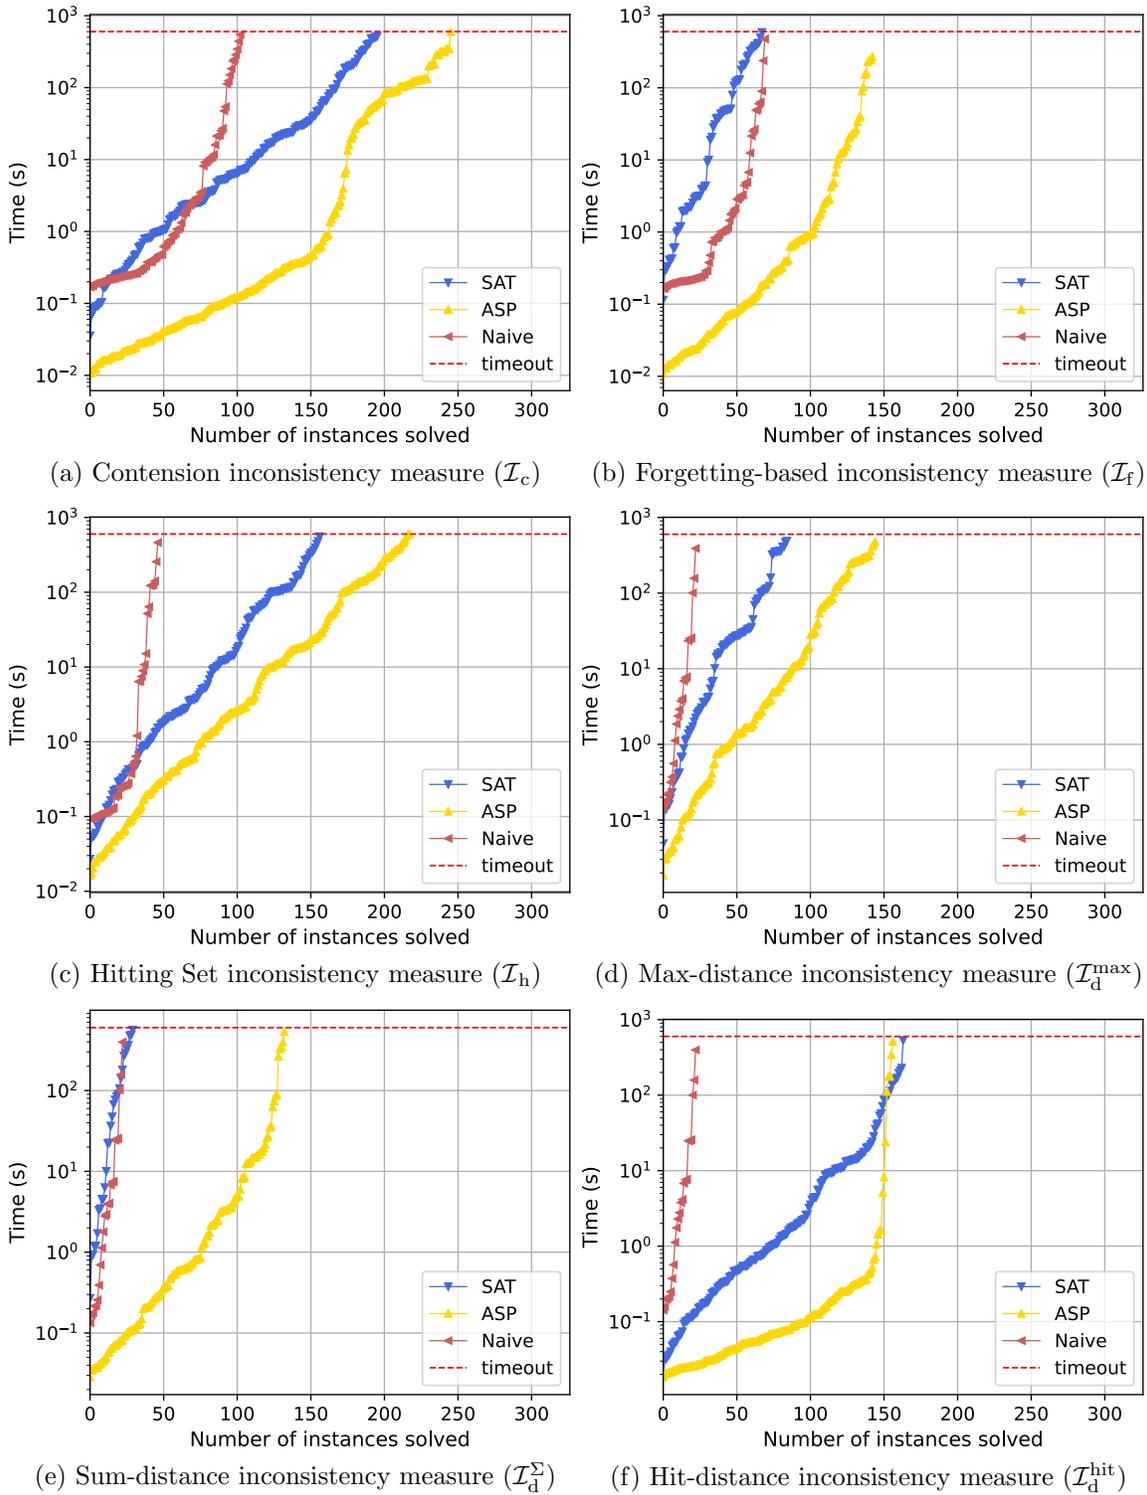


Figure 5: Runtime comparison of the ASP-based, SAT-based, and naive approaches on the ARG data set. Timeout: 10 minutes.

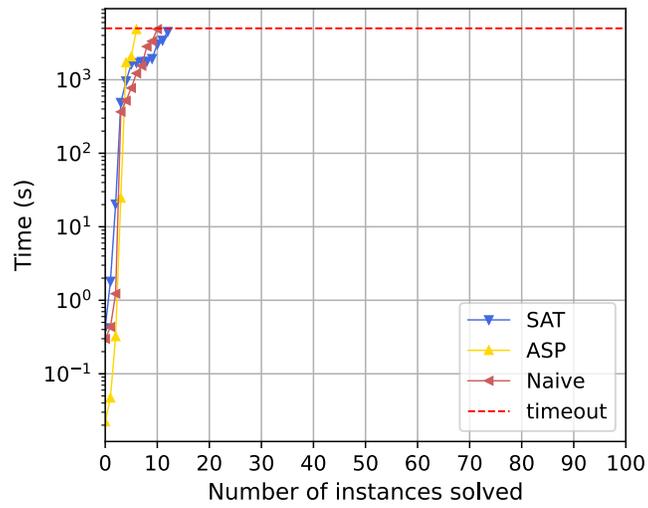
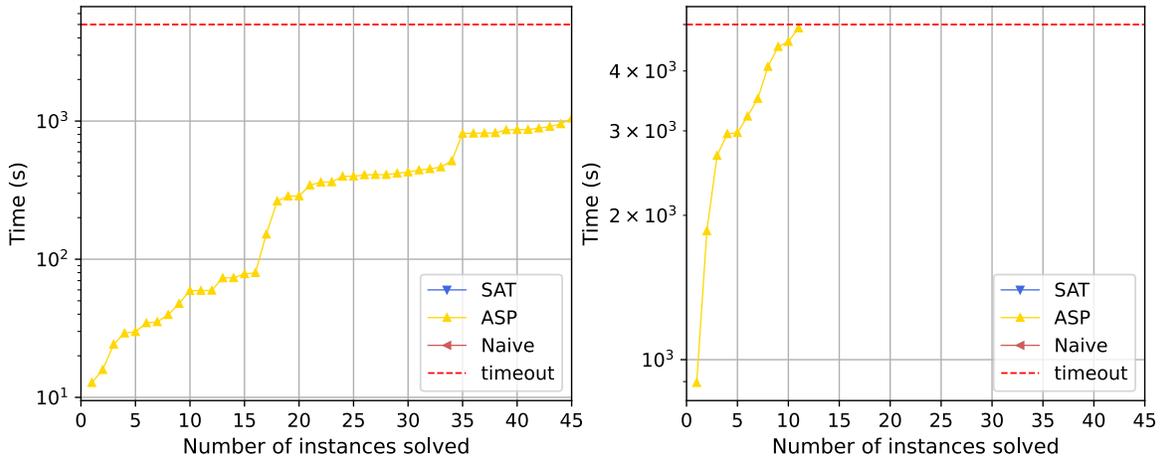
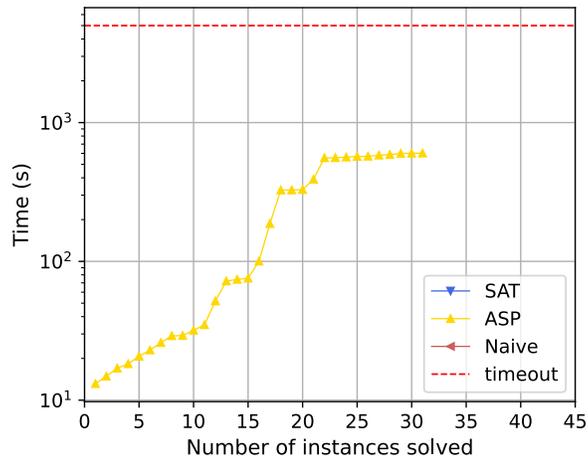


Figure 6: Runtime comparison of the ASP-based, SAT-based, and naive approaches for the contention inconsistency measure on the SC data set. Timeout: 5000 seconds (83.3 minutes).

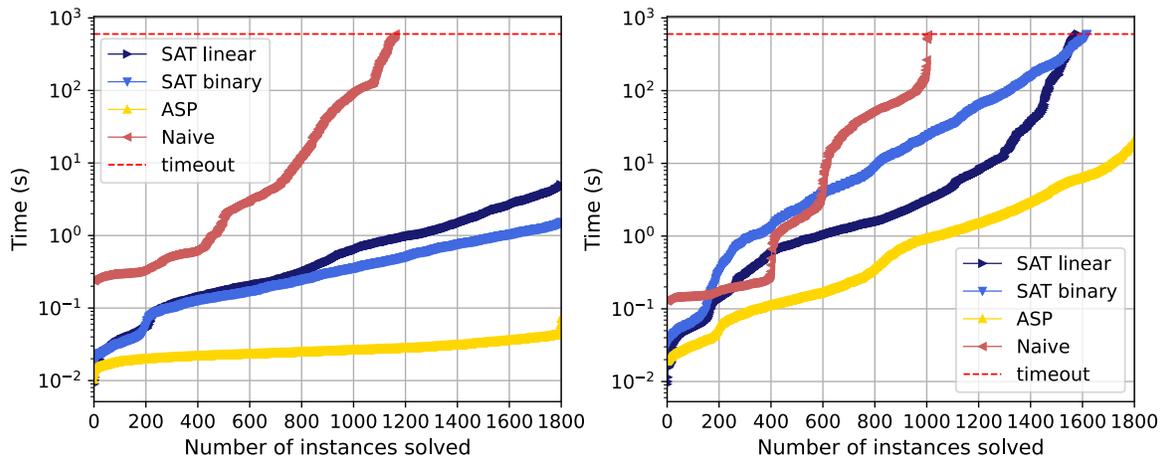


(a) Contension inconsistency measure ( $\mathcal{I}_c$ )      (b) Forgetting-based inconsistency measure ( $\mathcal{I}_f$ )

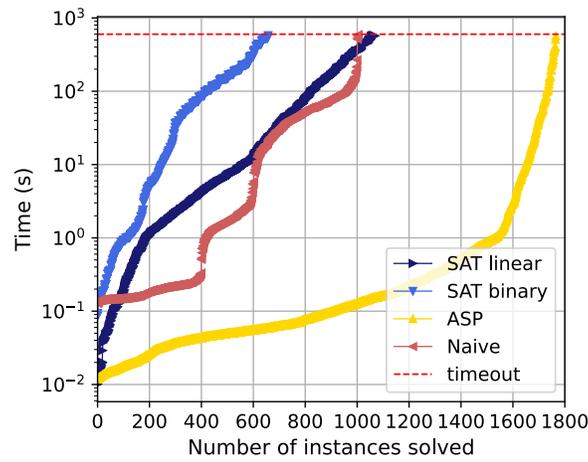


(c) Hit-distance inconsistency measure ( $\mathcal{I}_d^{\text{hit}}$ )

Figure 7: Runtime comparison of the ASP-based, SAT-based, and naive approaches on the LP data set for the contension and forgetting-based inconsistency measures. Timeout: 10 minutes.



(a) Contension inconsistency measure ( $\mathcal{I}_c$ )      (b) Max-distance inconsistency measure ( $\mathcal{I}_d^{\max}$ )



(c) Sum-distance inconsistency measure ( $\mathcal{I}_d^{\Sigma}$ )

Figure 8: Runtime comparison of the SAT-based approaches based on linear search (for  $\mathcal{I}_c$ ,  $\mathcal{I}_d^{\max}$ , and  $\mathcal{I}_d^{\Sigma}$ ) and the corresponding binary search versions, as well as the ASP-based versions, on the SRS data set. Timeout: 10 minutes.

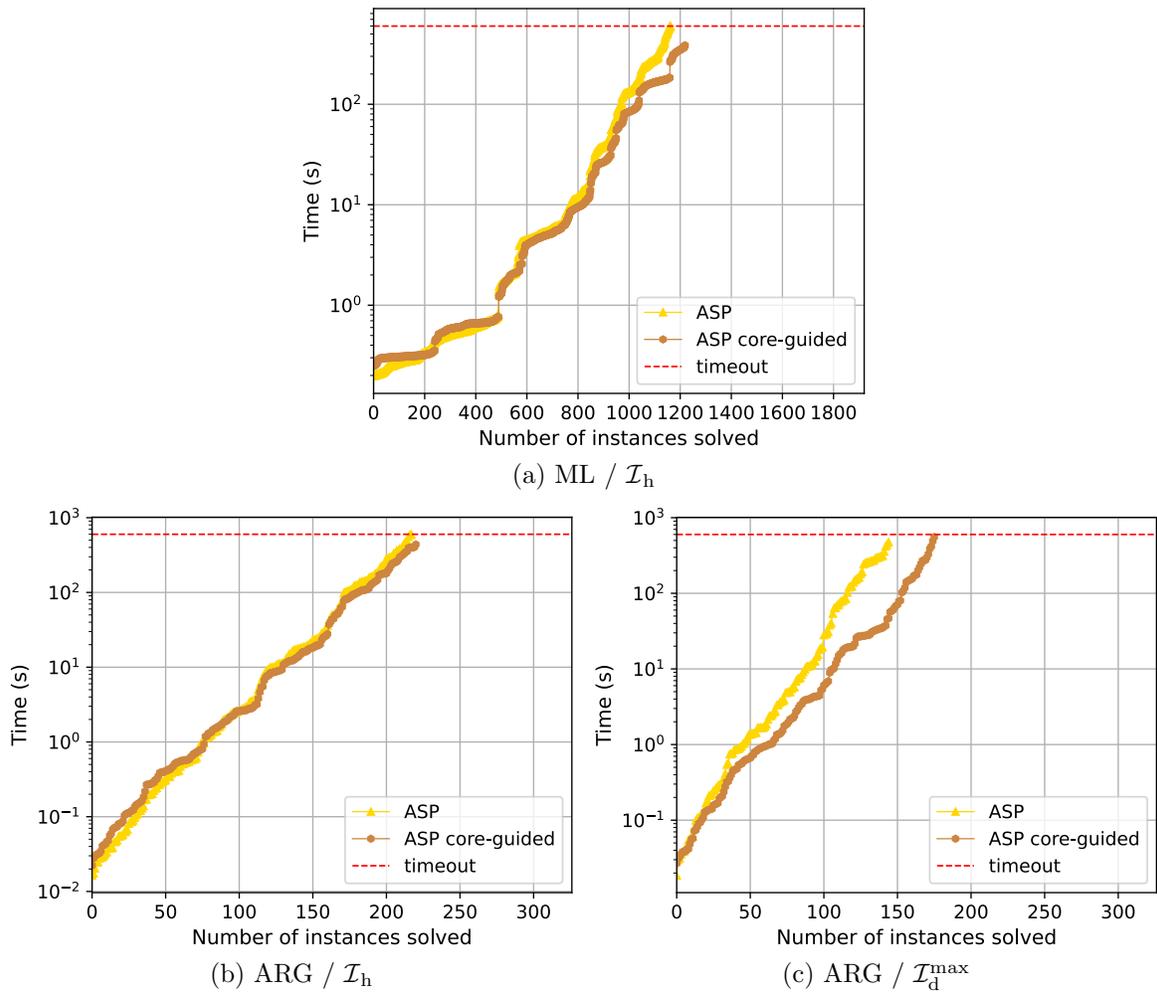


Figure 9: Runtime comparison of the ASP approach with core-guided vs. with branch-and-bound optimization (default) on the ML data set wrt. the hitting set measure ( $\mathcal{I}_h$ ) and on the ARG data set wrt.  $\mathcal{I}_h$  and the max-distance measure ( $\mathcal{I}_d^{\max}$ ). Timeout: 10 minutes.

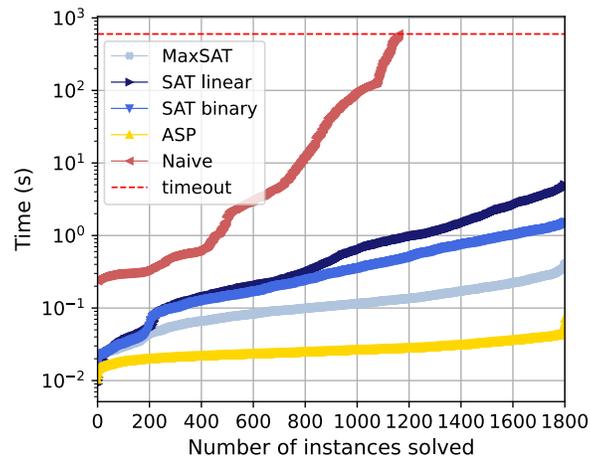


Figure 10: Runtime comparison of the MaxSAT approach for  $\mathcal{I}_c$  and the previously introduced approaches on the SRS data set. Timeout: 10 minutes.

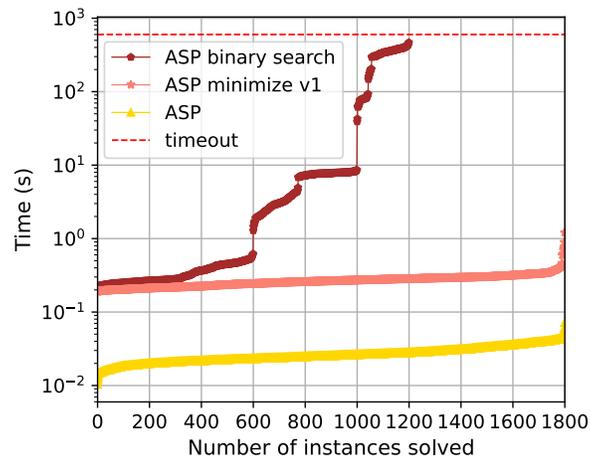


Figure 11: Runtime comparison of the different versions of the ASP approach for  $\mathcal{I}_c$  on the SRS data set. “ASP binary search” refers to the version from (Kuhlmann & Thimm, 2020), “ASP minimize v1” to the version from (Kuhlmann & Thimm, 2021), and “ASP” to the new one. Timeout: 10 minutes.

## B.2 Inconsistency values

In this section, we present histograms over the inconsistency values regarding all inconsistency measures and all data sets considered in this work. Note that we could only include inconsistency values wrt. instances for which at least one approach did not time out. To be precise, the histograms for the SRS data set are included in Figure 12, those for the ML data set in Figure 13, those for the ARG data set in Figure 14, those for the SC data set in Figure 15, and those for the LP data set in Figure 16.

Furthermore, the following plots only include non- $\infty$  values. With regard to those measures that allow for the value  $\infty$  (i.e.,  $\mathcal{I}_h$ ,  $\mathcal{I}_d^{\max}$ , and  $\mathcal{I}_d^\Sigma$ ), the number of instances resulting in that value are mentioned in the corresponding captions.

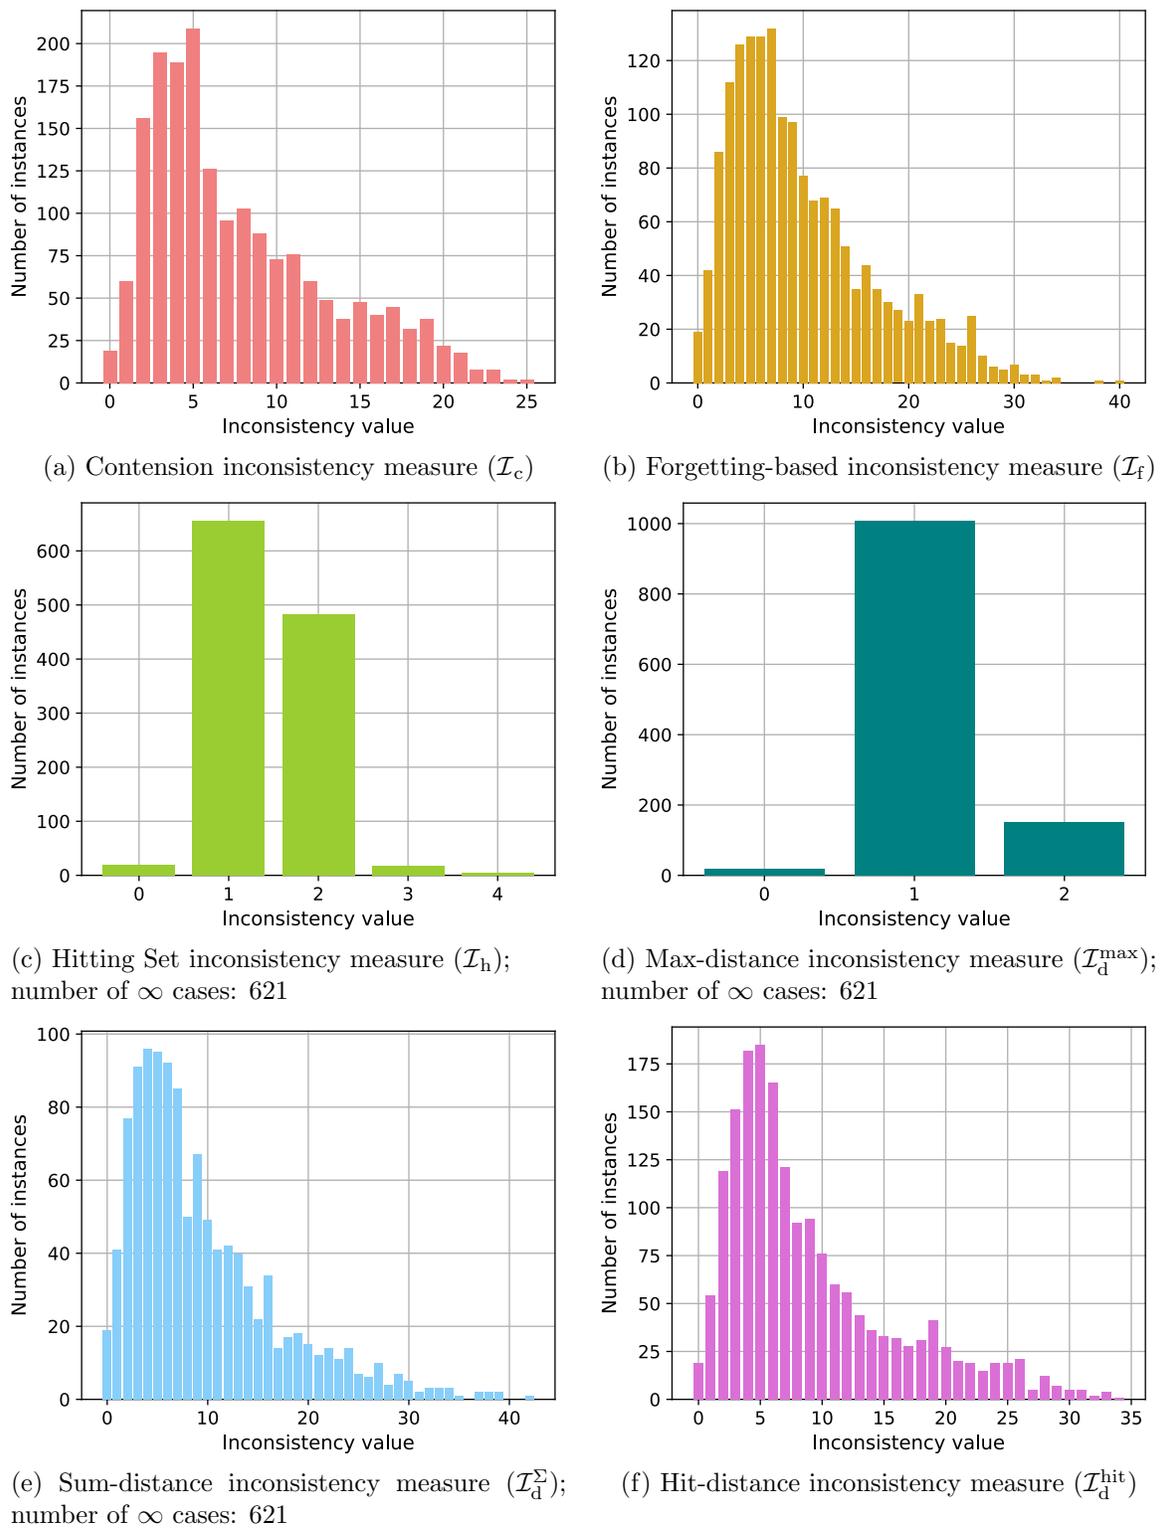
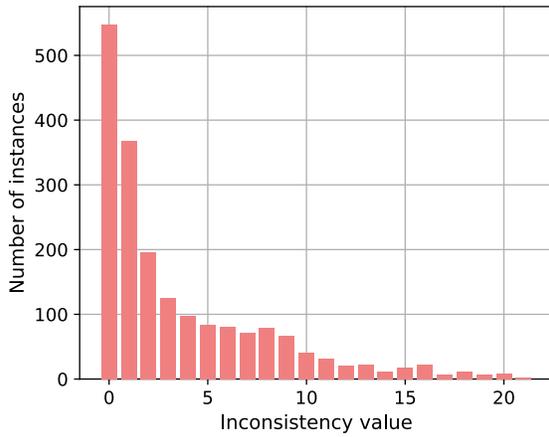
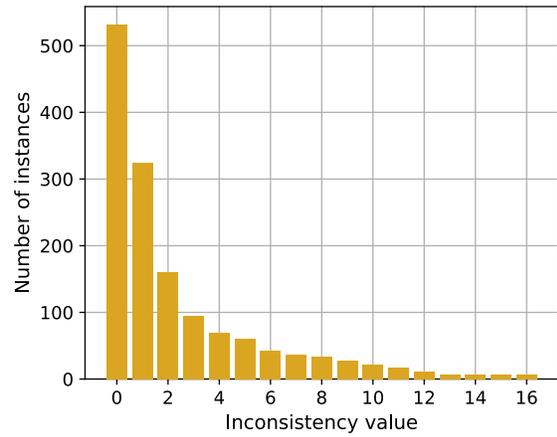


Figure 12: Histograms of the inconsistency values of the SRS data set wrt. all six inconsistency measures.



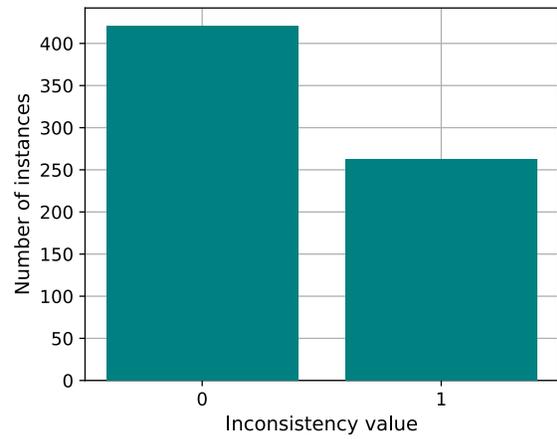
(a) Contension inconsistency measure ( $\mathcal{I}_c$ )



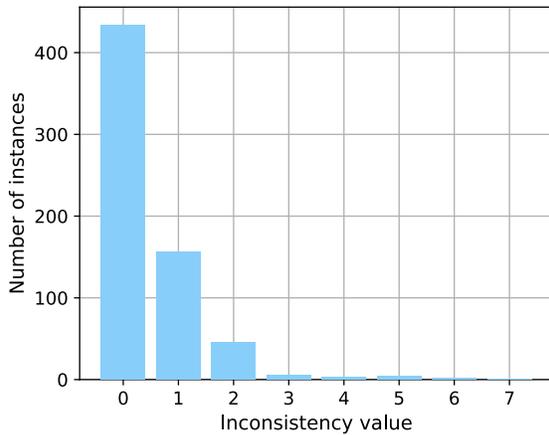
(b) Forgetting-based inconsistency measure ( $\mathcal{I}_f$ )



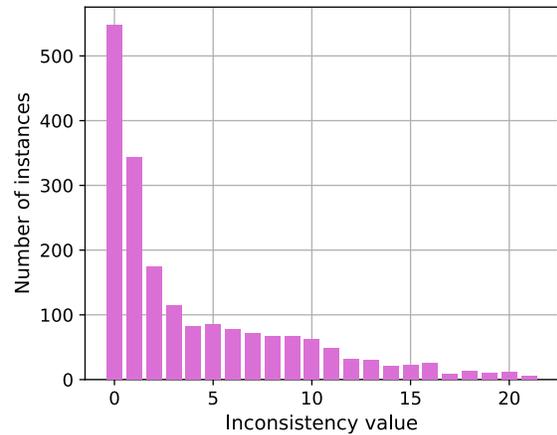
(c) Hitting Set inconsistency measure ( $\mathcal{I}_h$ );  
number of  $\infty$  cases: 0



(d) Max-distance inconsistency measure ( $\mathcal{I}_d^{\max}$ );  
number of  $\infty$  cases: 0

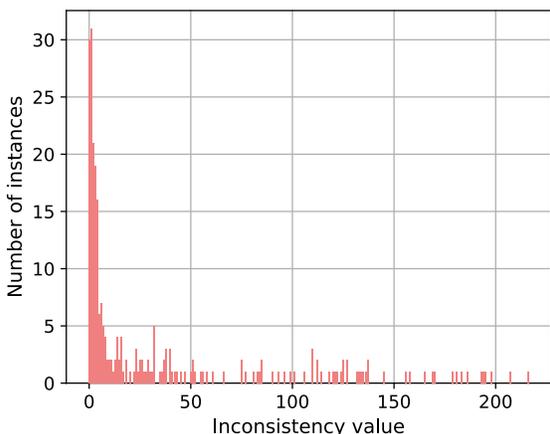


(e) Sum-distance inconsistency measure ( $\mathcal{I}_d^{\sum}$ );  
number of  $\infty$  cases: 0

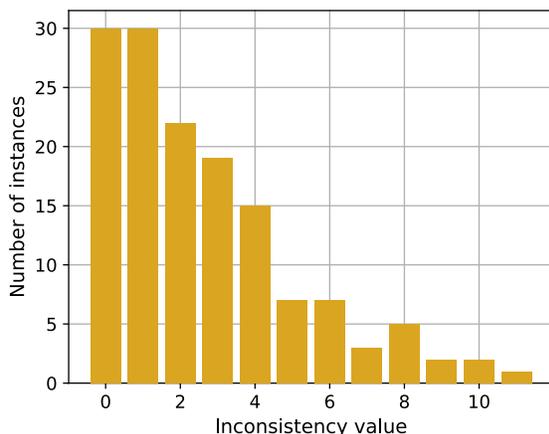


(f) Hit-distance inconsistency measure ( $\mathcal{I}_d^{\text{hit}}$ )

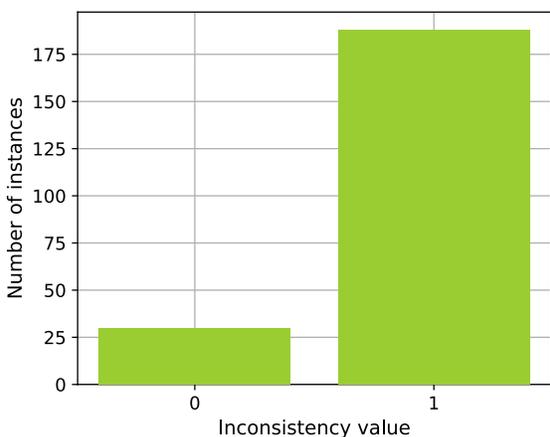
Figure 13: Histograms of the inconsistency values of the ML data set wrt. all six inconsistency measures



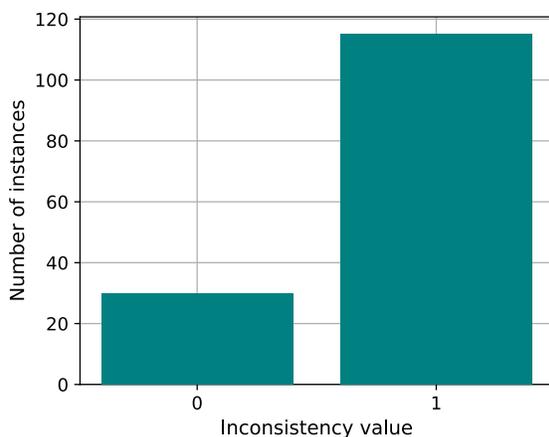
(a) Contension inconsistency measure ( $\mathcal{I}_c$ )



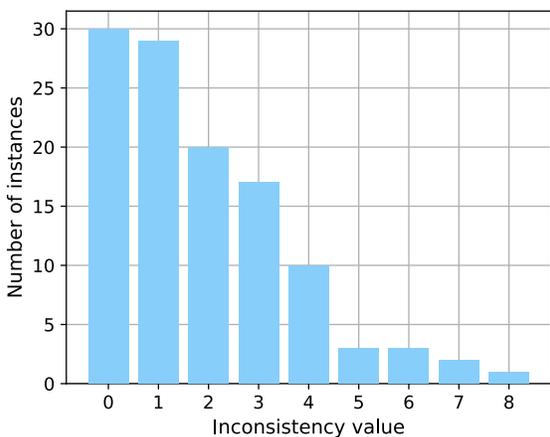
(b) Forgetting-based inconsistency measure ( $\mathcal{I}_f$ )



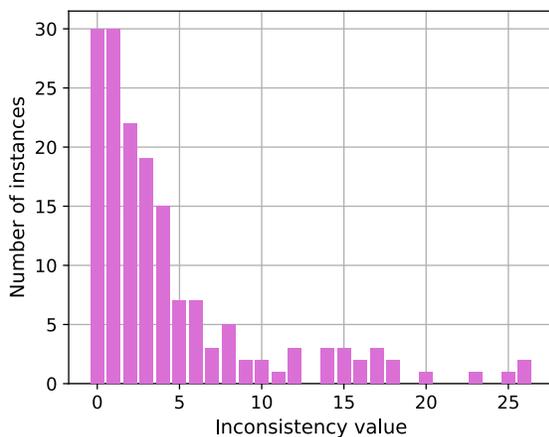
(c) Hitting Set inconsistency measure ( $\mathcal{I}_h$ );  
number of  $\infty$  cases: 0



(d) Max-distance inconsistency measure ( $\mathcal{I}_d^{\max}$ );  
number of  $\infty$  cases: 0



(e) Sum-distance inconsistency measure ( $\mathcal{I}_d^{\sum}$ );  
number of  $\infty$  cases: 0



(f) Hit-distance inconsistency measure ( $\mathcal{I}_d^{\text{hit}}$ )

Figure 14: Histograms of the inconsistency values of the ARG data set wrt. all six inconsistency measures

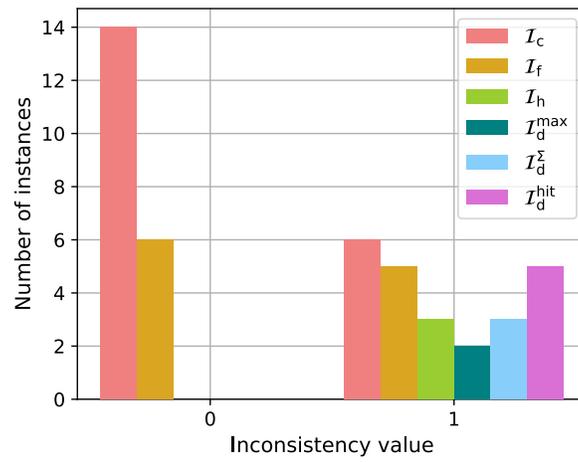


Figure 15: Histogram of the inconsistency values of the SC data set wrt. all six inconsistency measures.

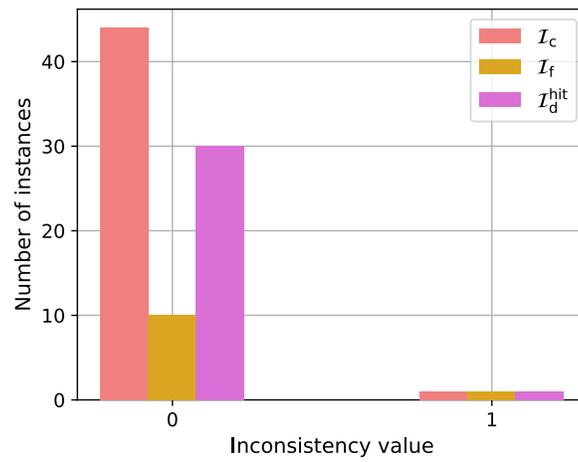


Figure 16: Histogram of the inconsistency values of the LP data set wrt.  $I_c$ ,  $I_f$ , and  $I_d^{\text{hit}}$ . Wrt. the other measures ( $I_h$ ,  $I_d^{\max}$ ,  $I_d^{\Sigma}$ ), none of the instances from the ASP data set could be solved.

### B.3 Scatter Plots

In this section we present scatter plots which compare each pair of approaches wrt. each inconsistency measure and each data set in terms of runtime per instance. Figures 17, 18, and 19 refer to the SRS data set, Figures 20, 21, and 22 refer to the ML data set, and Figures 23, 24, 25 refer to the ARG data set, and Figures 26, 27, 28 refer to the SC data set. The scatter plots regarding the LP data set are omitted, since both the naive and the SAT-based approaches exclusively produced timeouts.

Furthermore, Figures 29 and 30 depict the scatter plots regarding the comparison between a linear search variant of the SAT approach and the other approaches, as described in Section 5.3.3. Moreover, the scatter plots relating to the comparison of the MaxSAT approach for the contention inconsistency measure and the corresponding other approaches (see Section 5.3.4) are shown in Figure 31.

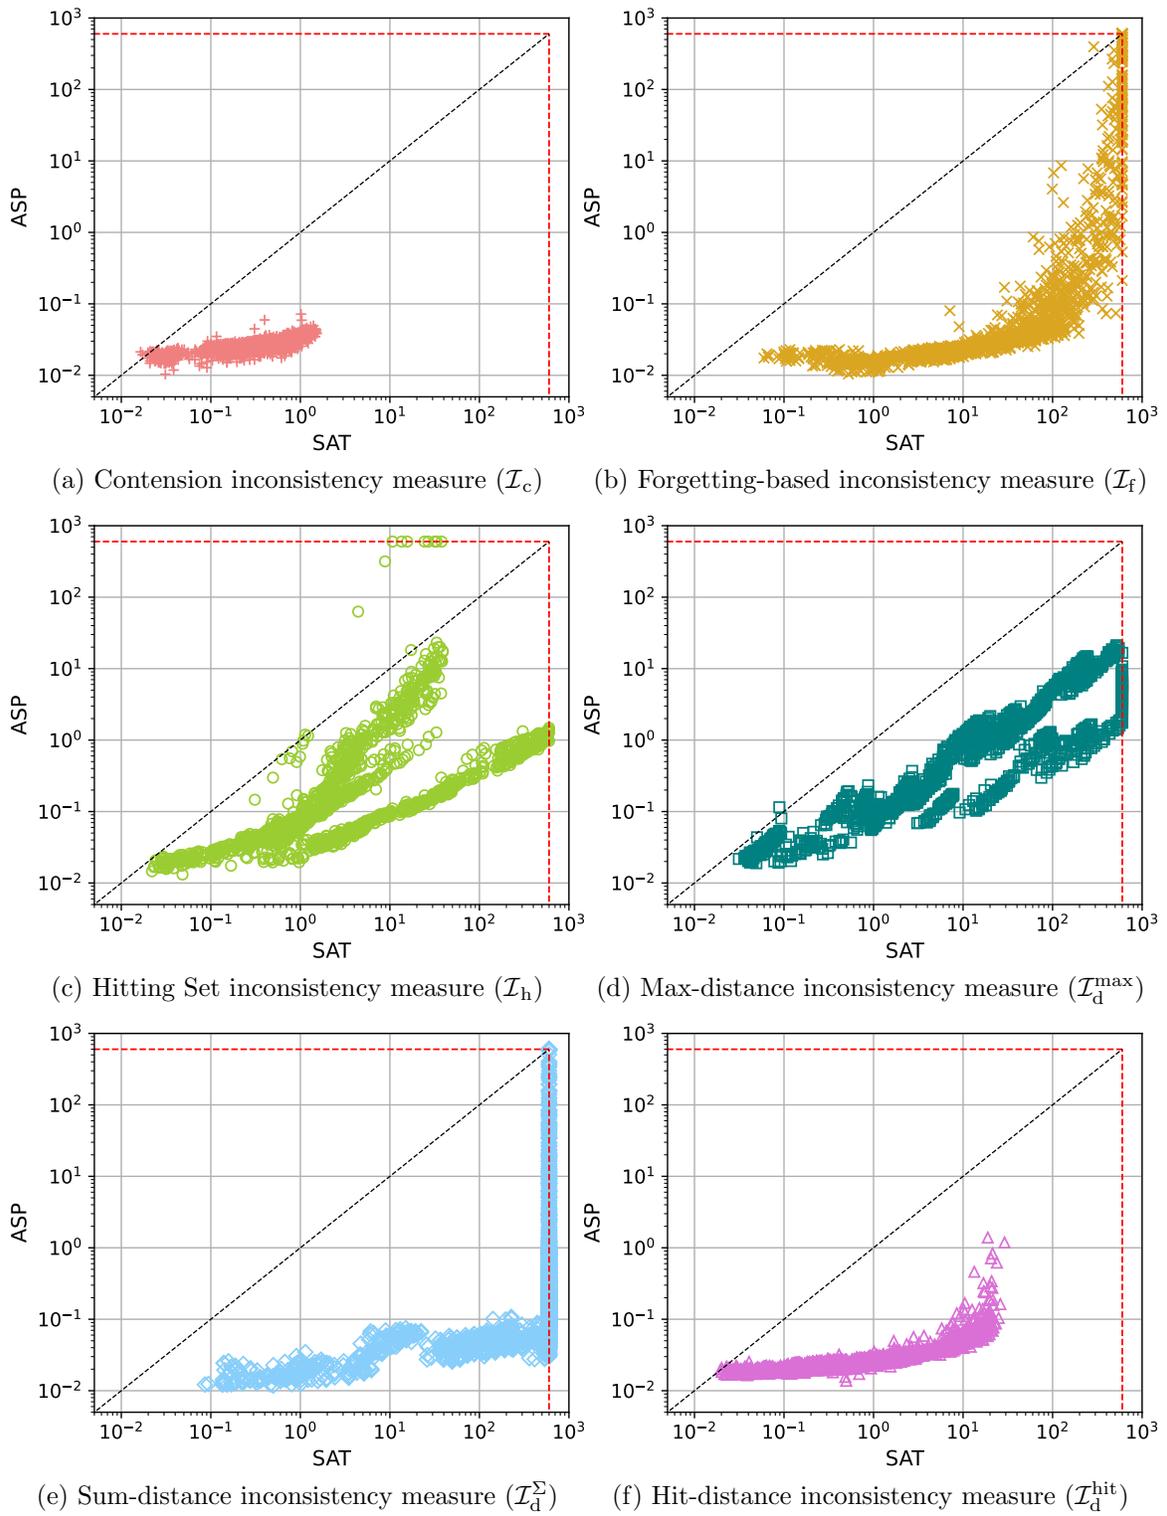


Figure 17: Runtime comparison of the ASP-based and SAT-based approaches on the SRS data set. Timeout: 10 minutes.

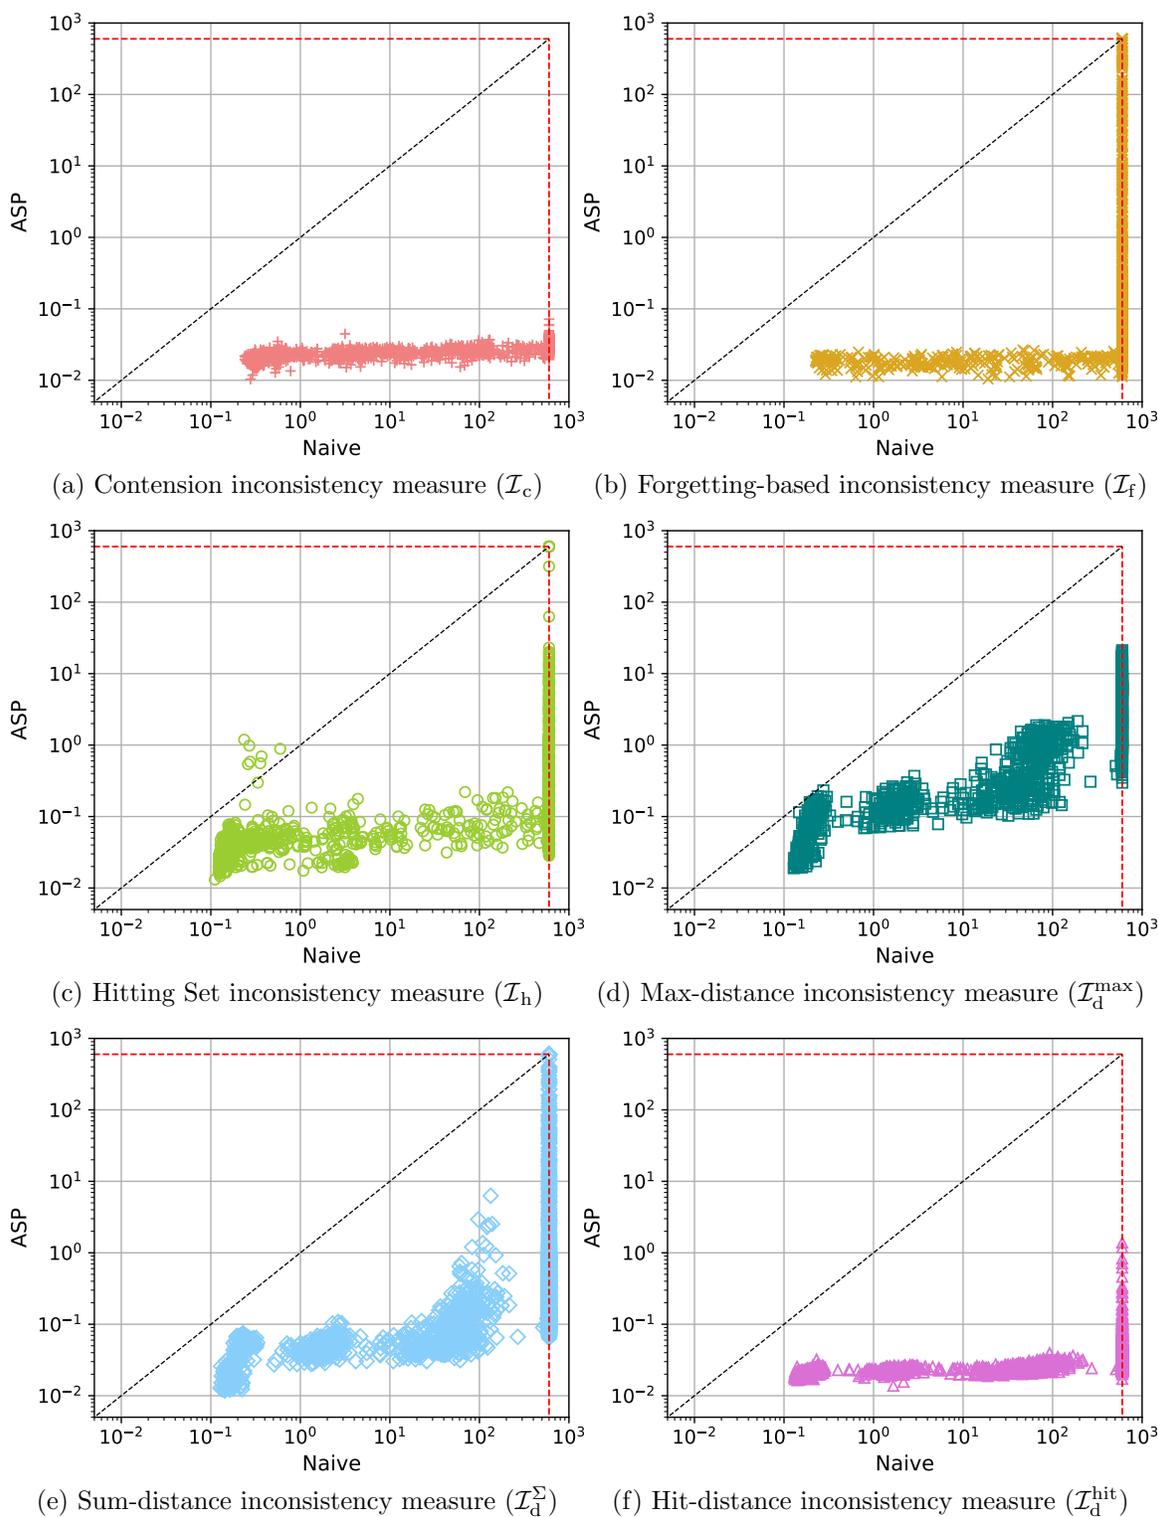


Figure 18: Runtime comparison of the ASP-based and naive approaches on the SRS data set. Timeout: 10 minutes.

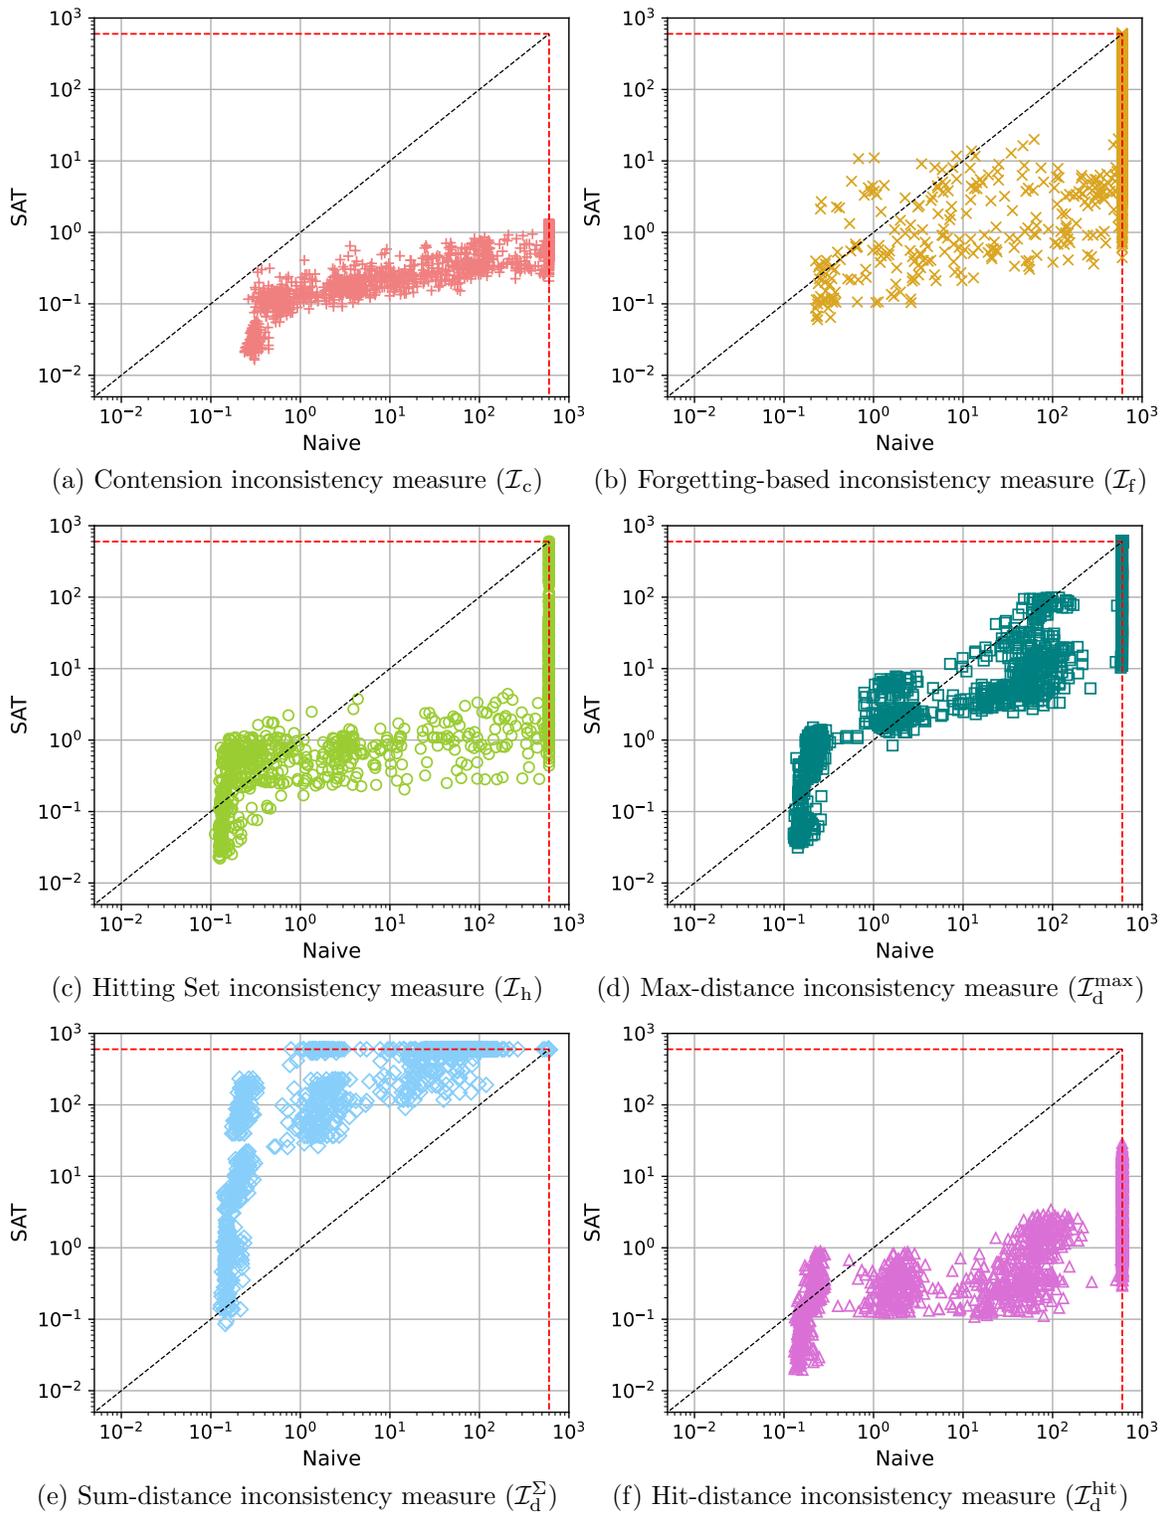


Figure 19: Runtime comparison of the SAT-based and naive approaches on the SRS data set. Timeout: 10 minutes.

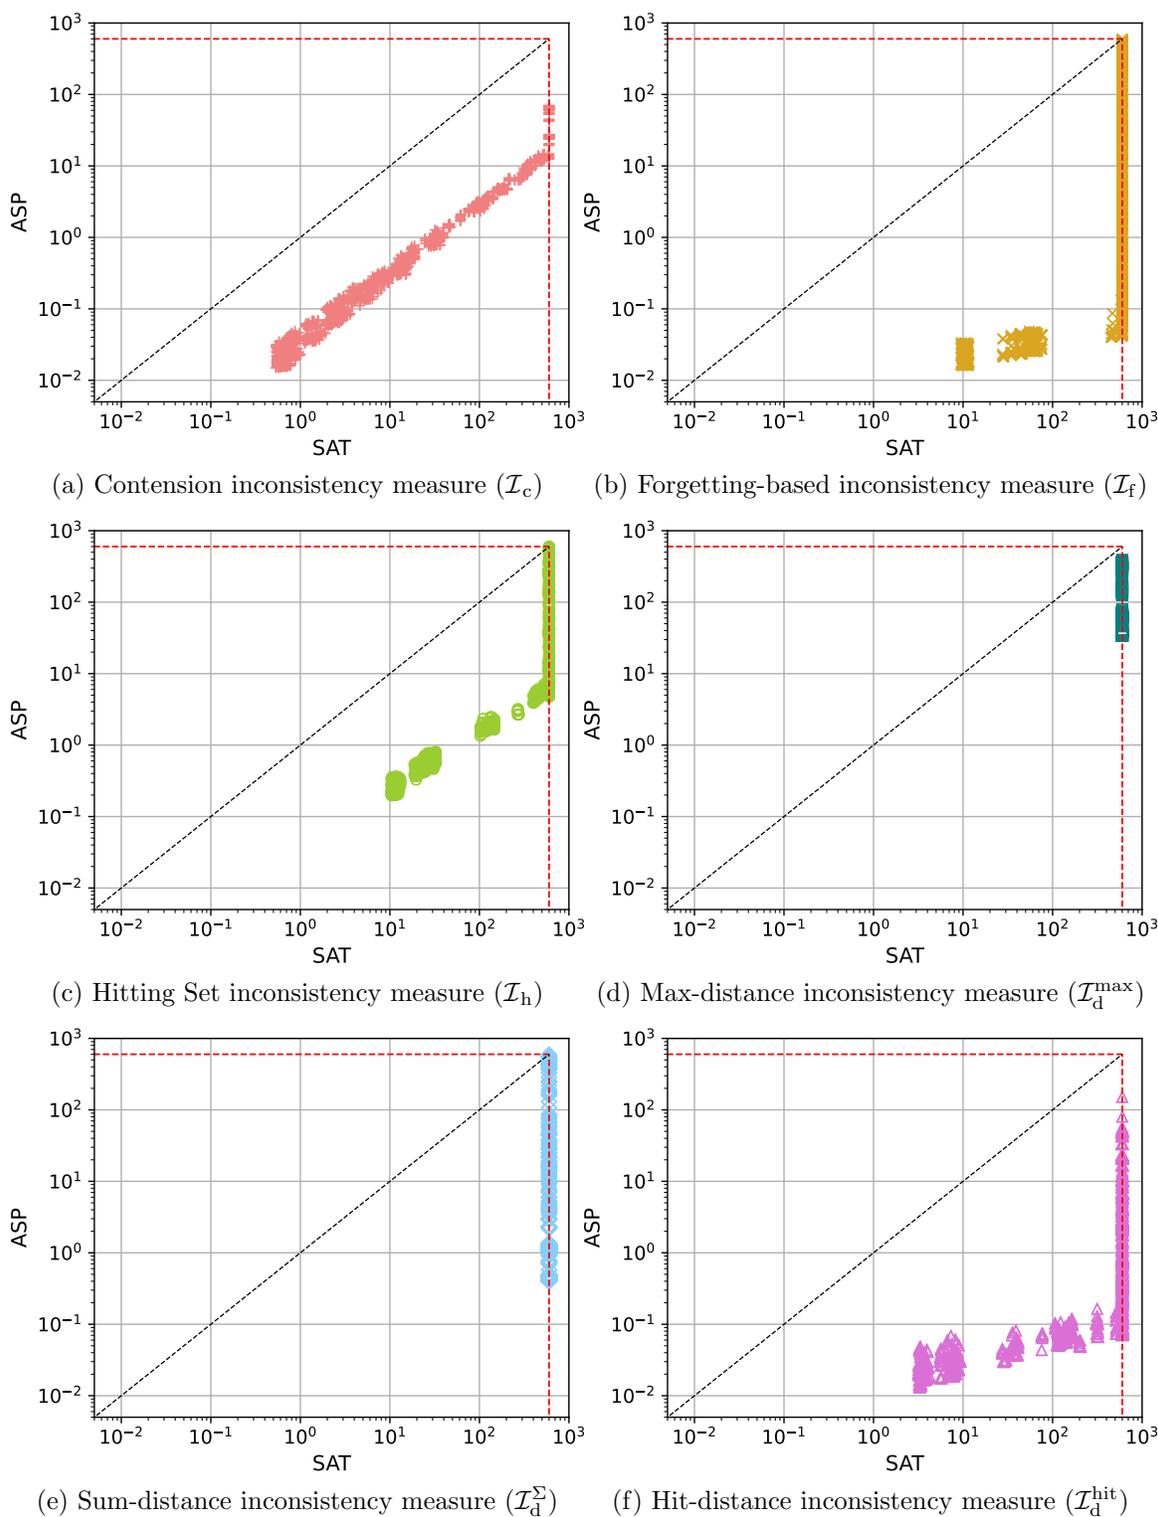


Figure 20: Runtime comparison of the ASP-based and SAT-based approaches on the ML data set. Timeout: 10 minutes.

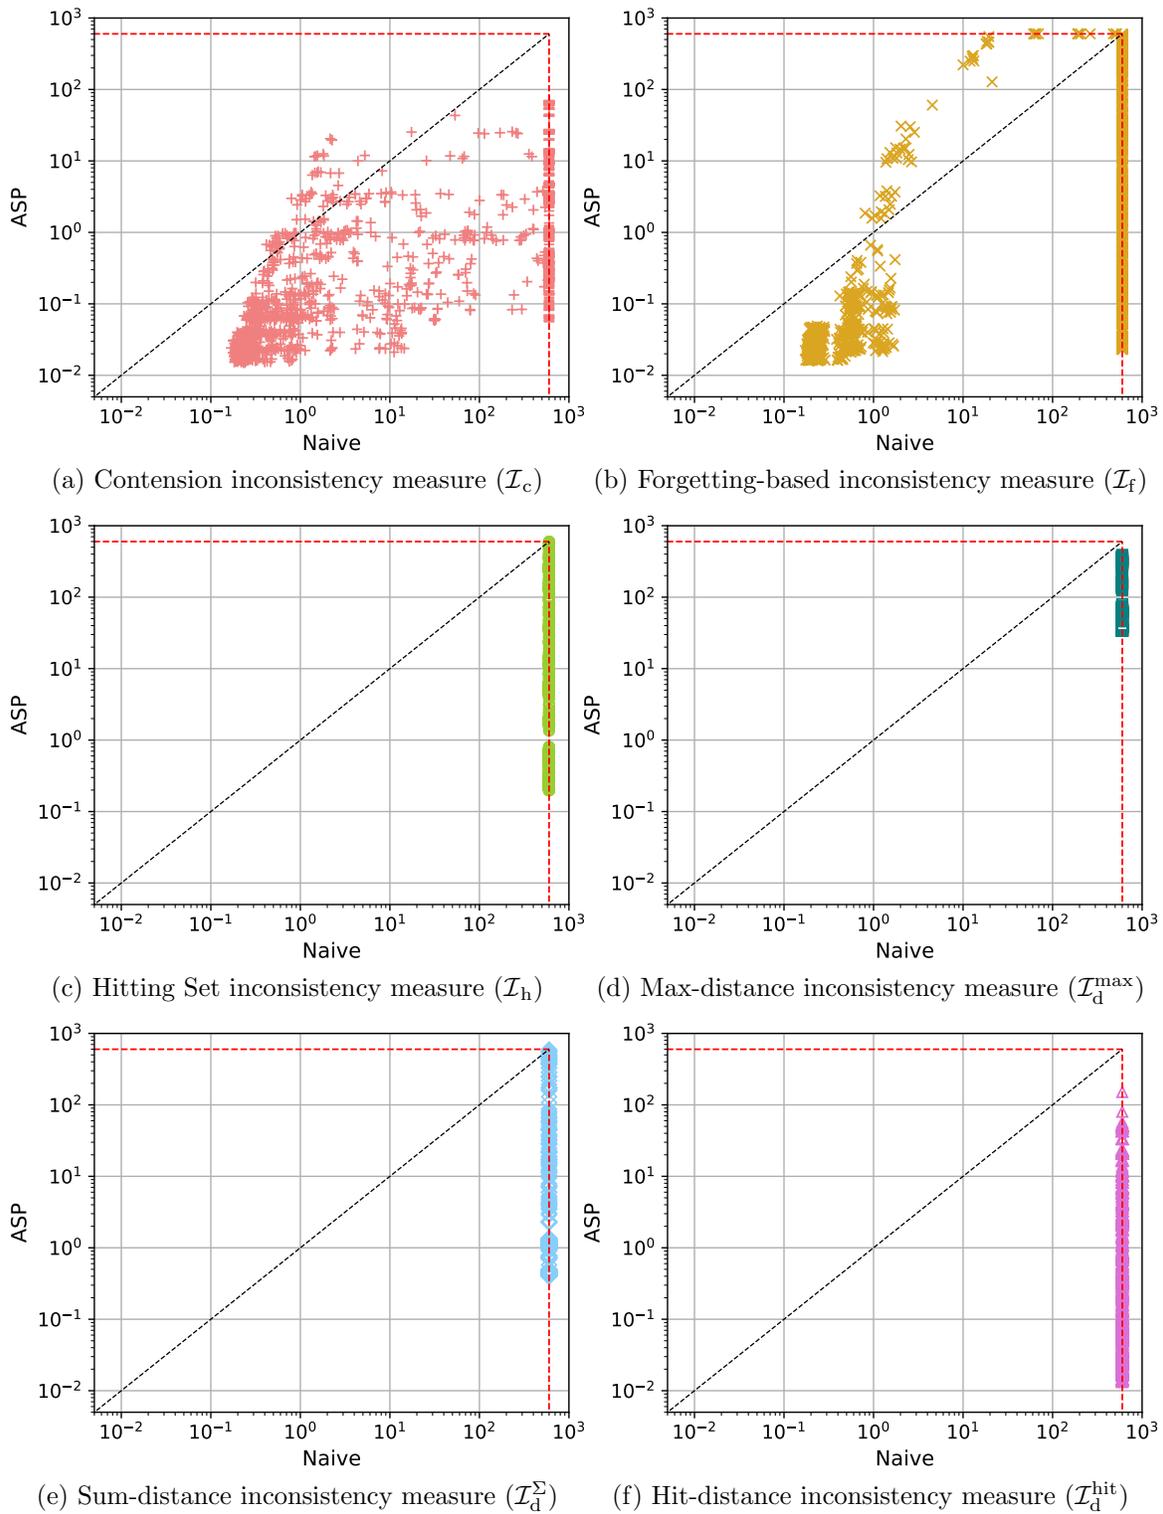


Figure 21: Runtime comparison of the ASP-based and naive approaches on the ML data set. Timeout: 10 minutes.

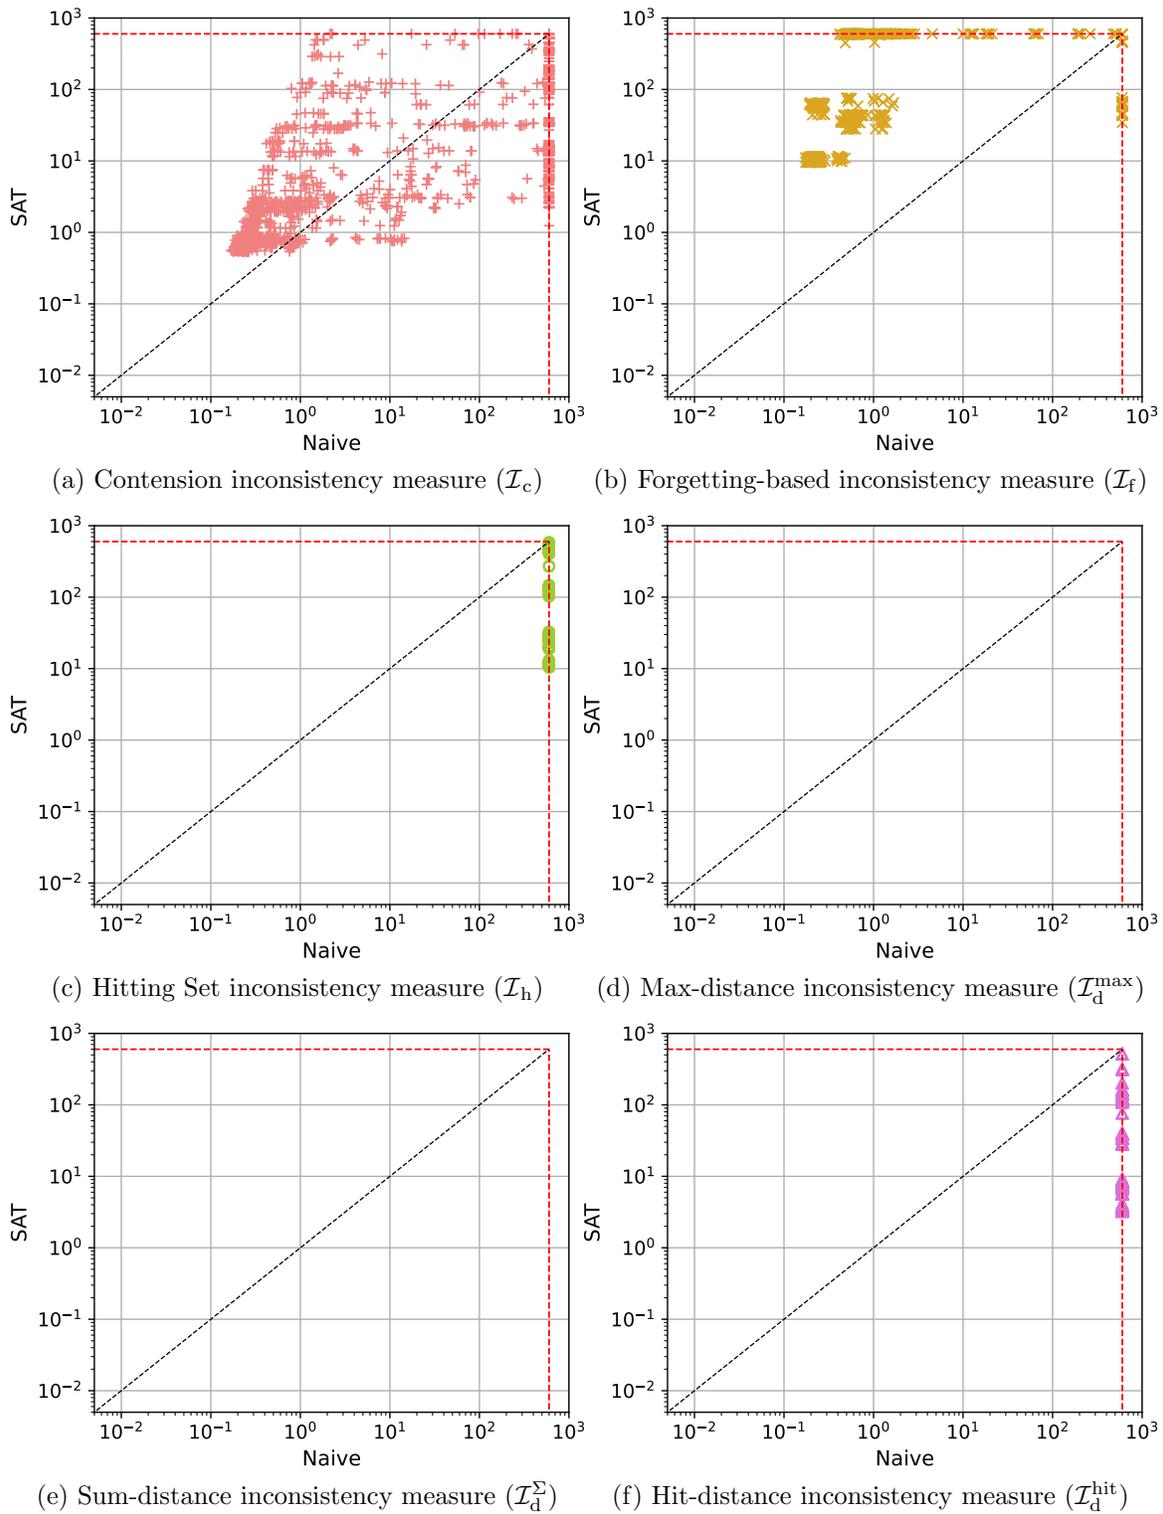


Figure 22: Runtime comparison of the SAT-based and naive approaches on the ML data set. Timeout: 10 minutes.

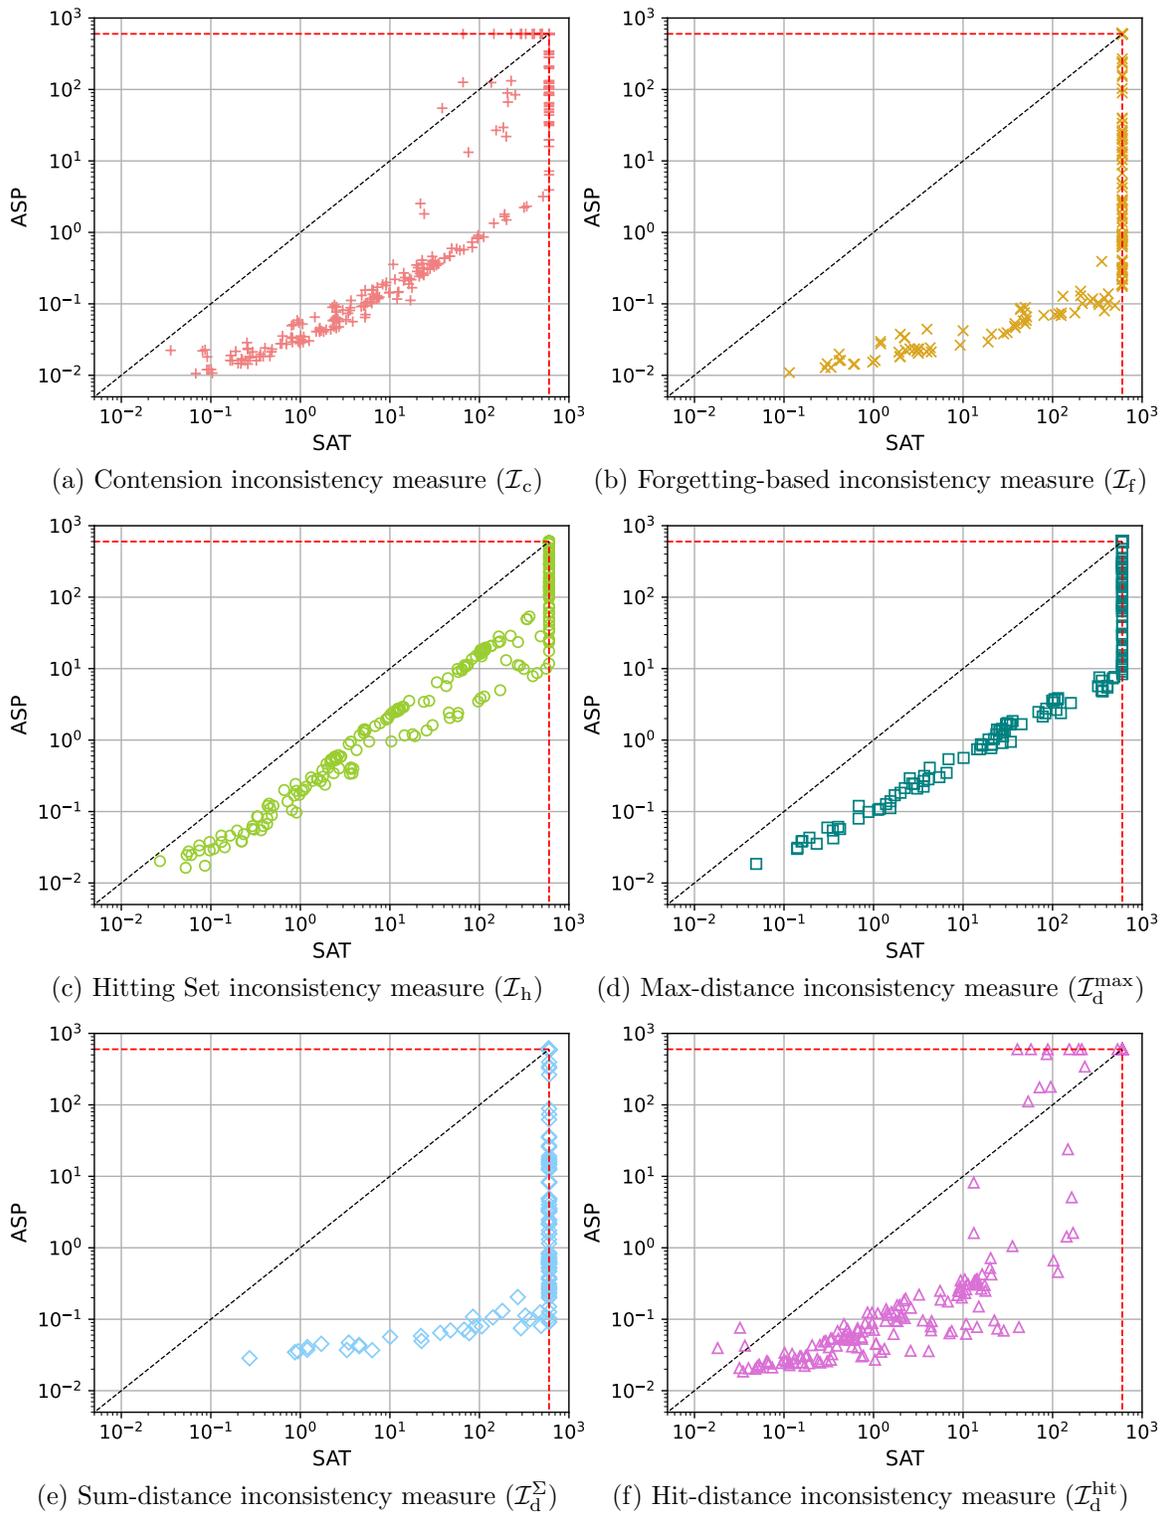


Figure 23: Runtime comparison of the ASP-based and SAT-based approaches on the ARG data set. Timeout: 10 minutes.

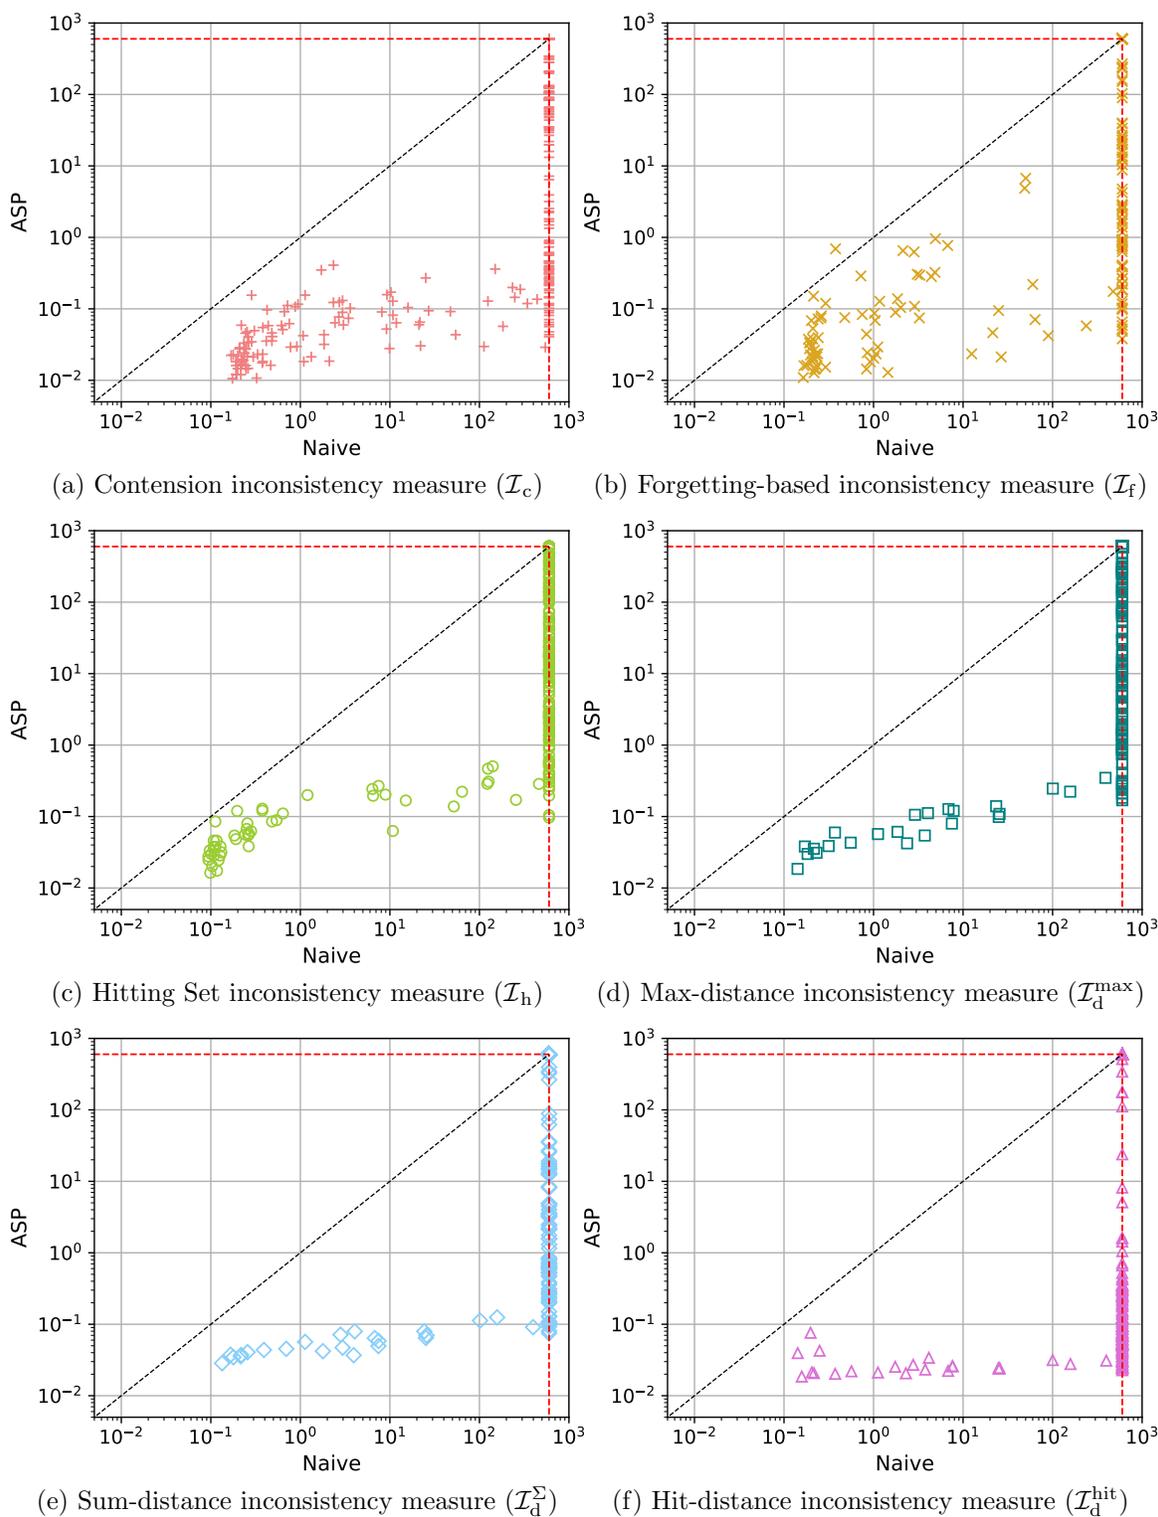


Figure 24: Runtime comparison of the ASP-based and naive approaches on the ARG data set. Timeout: 10 minutes.

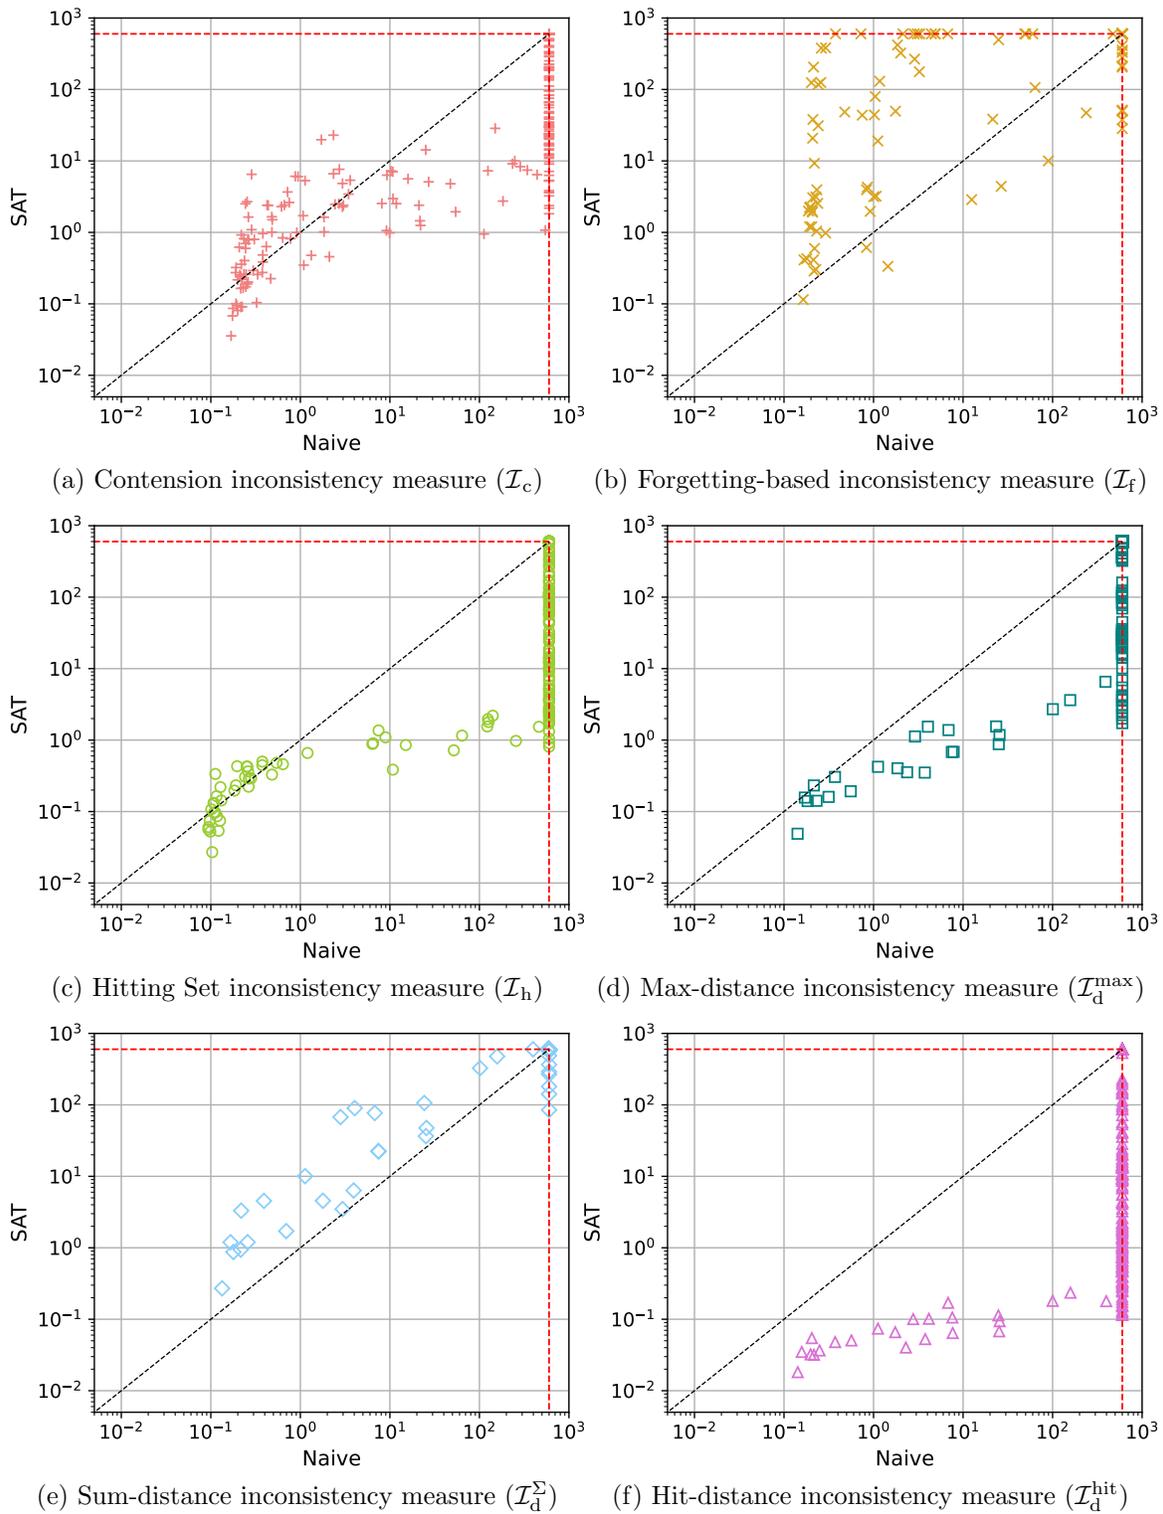


Figure 25: Runtime comparison of the SAT-based and naive approaches on the ARG data set. Timeout: 10 minutes.

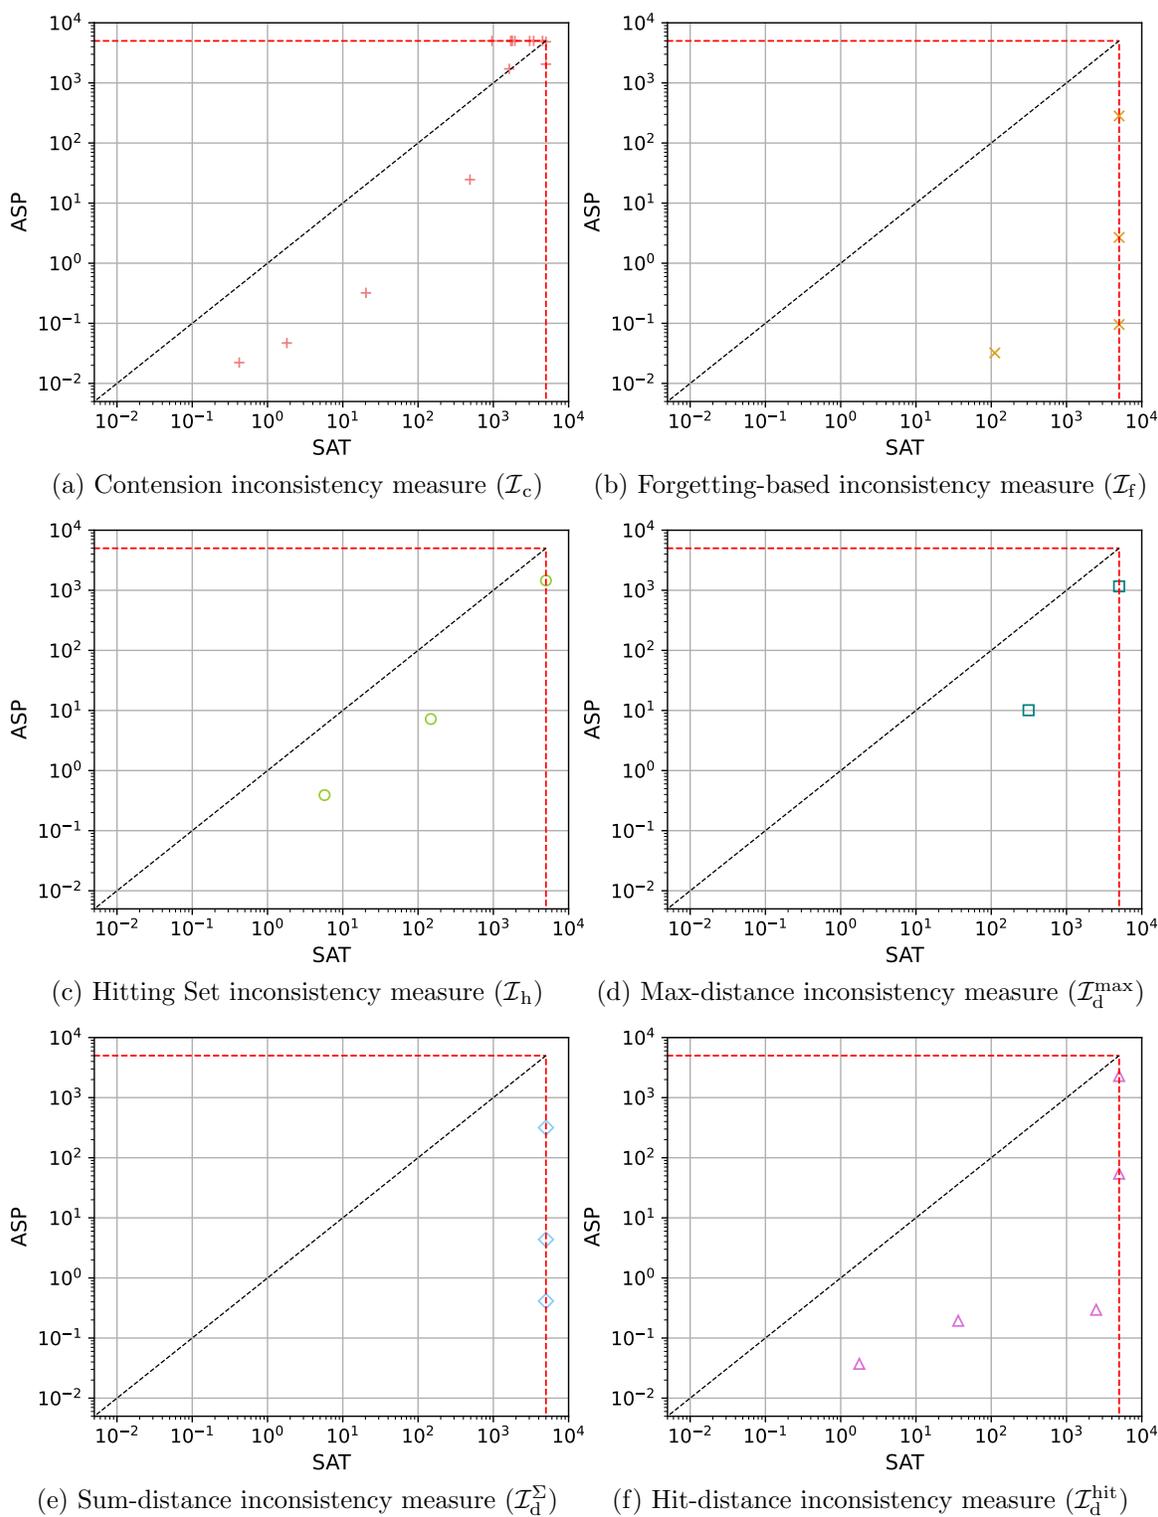


Figure 26: Runtime comparison of the ASP-based and SAT-based approaches on the SC data set. Timeout: 5000 seconds.

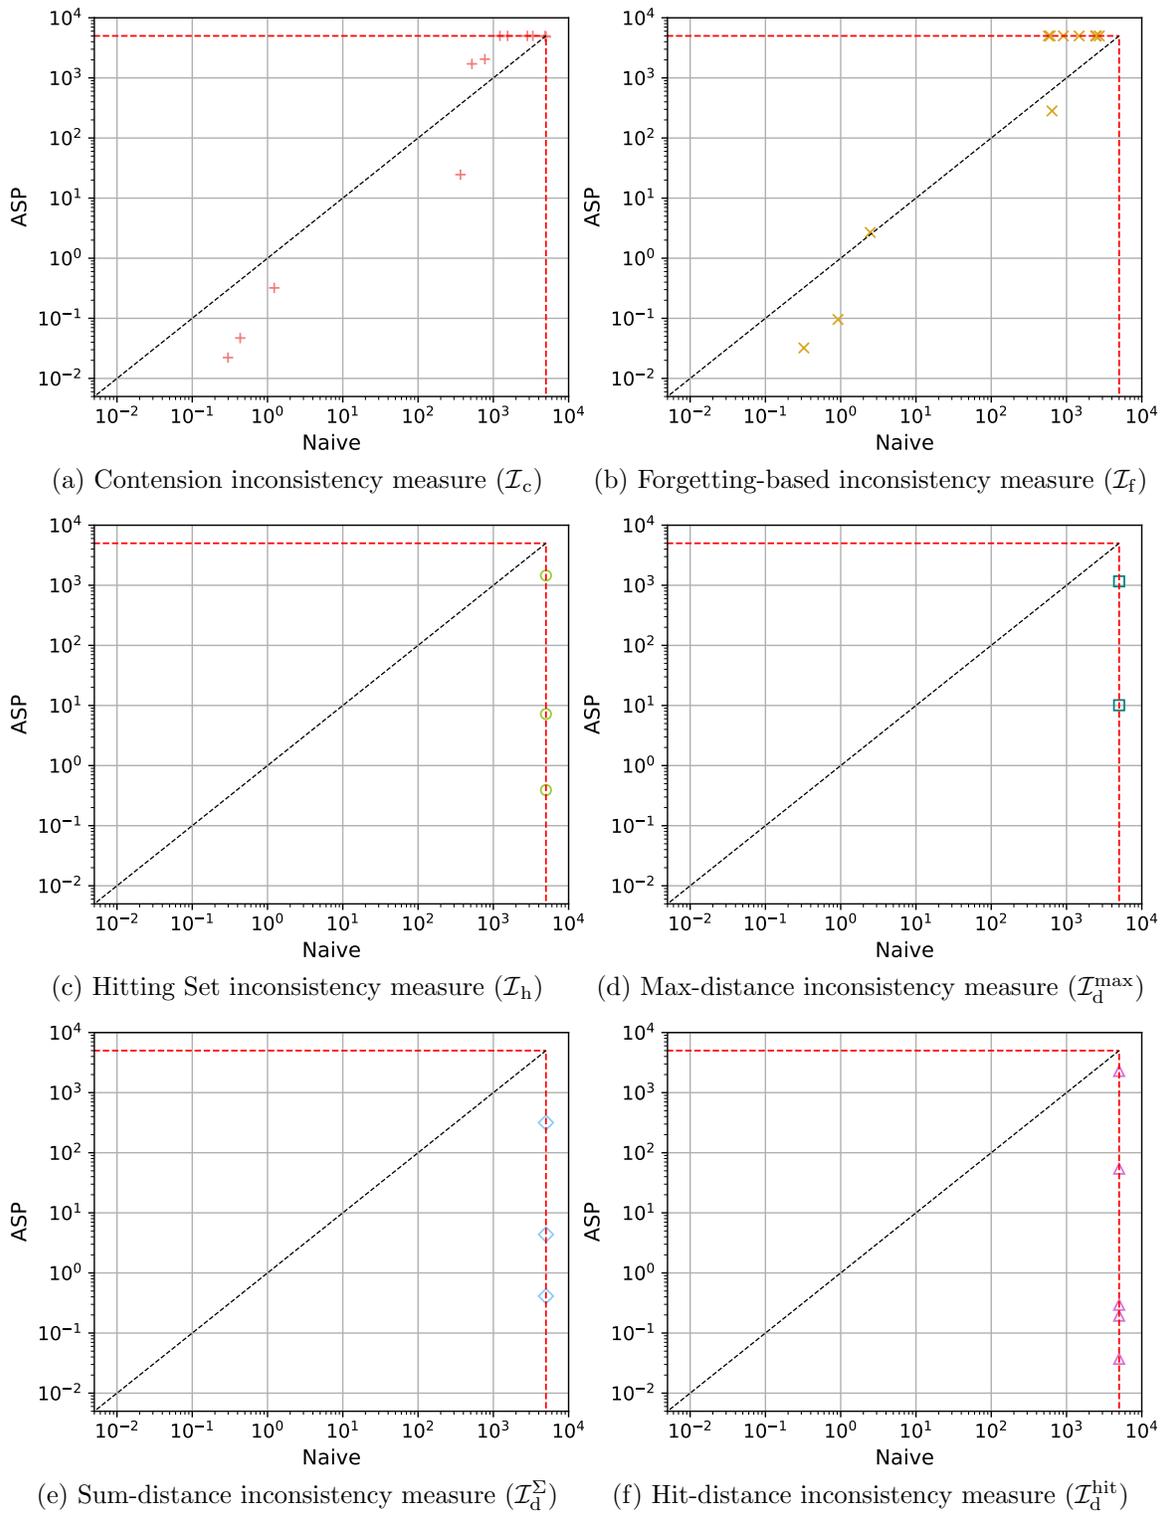


Figure 27: Runtime comparison of the ASP-based and naive approaches on the SC data set. Timeout: 5000 seconds.

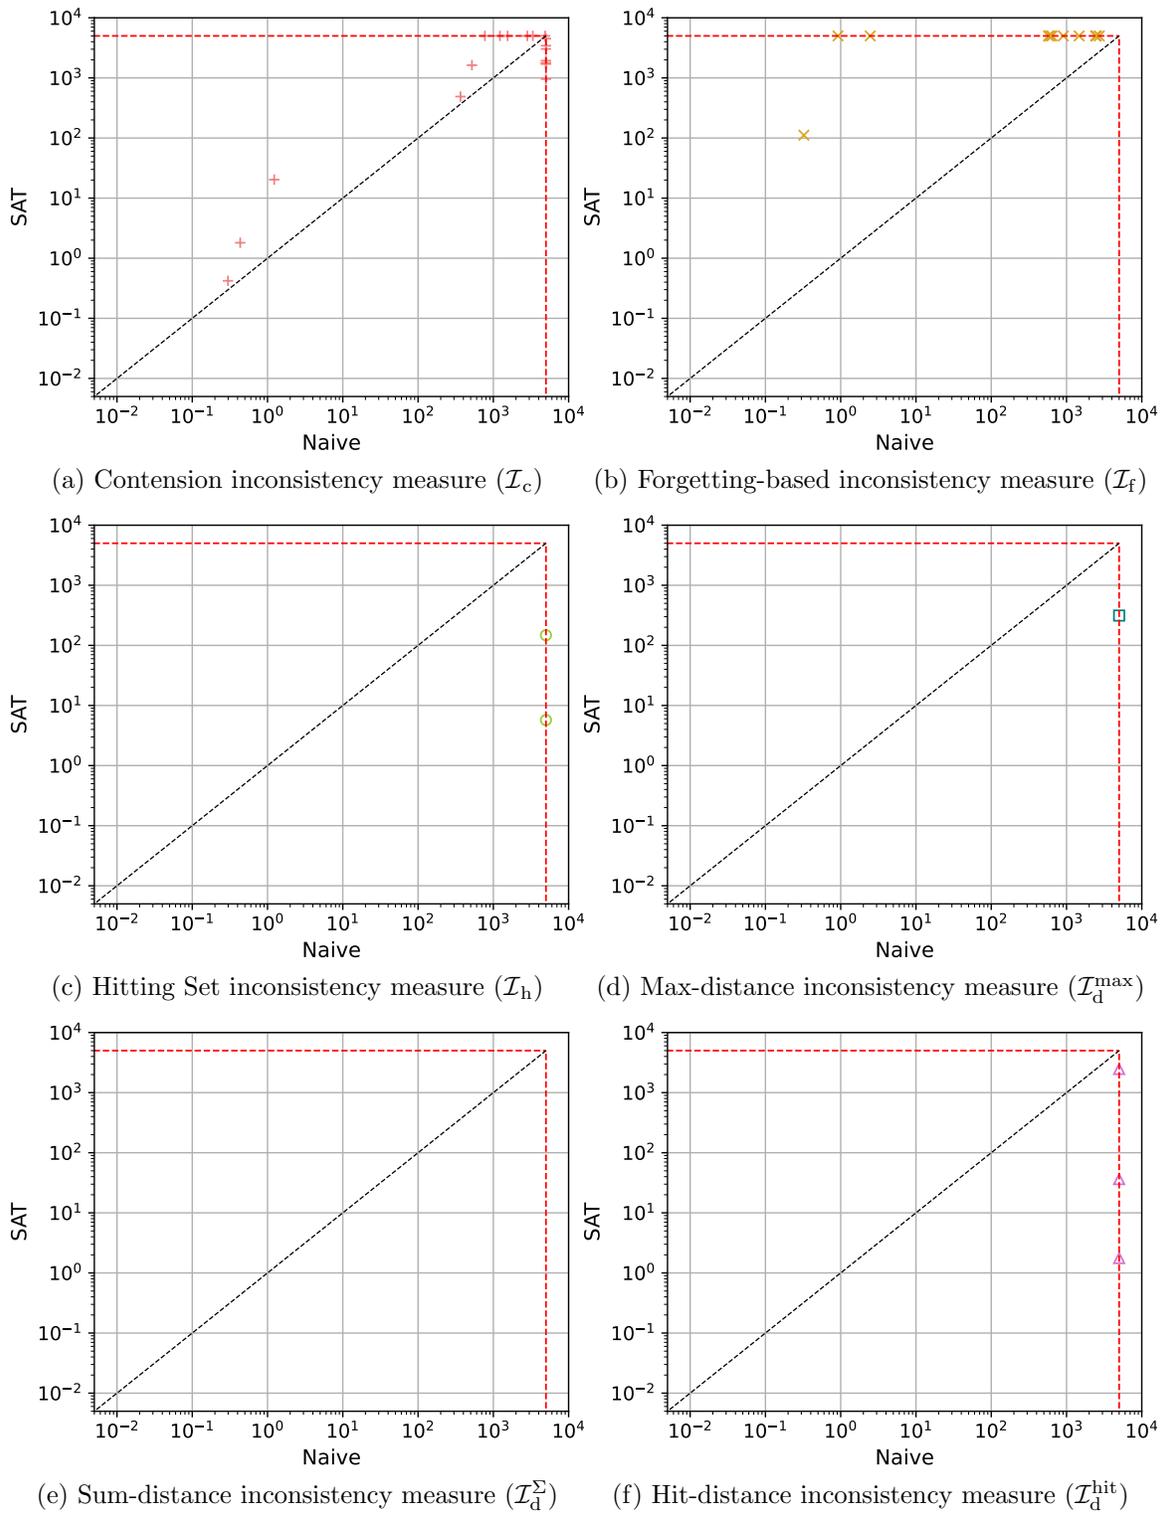


Figure 28: Runtime comparison of the SAT-based and naive approaches on the SC data set. Timeout: 5000 seconds.

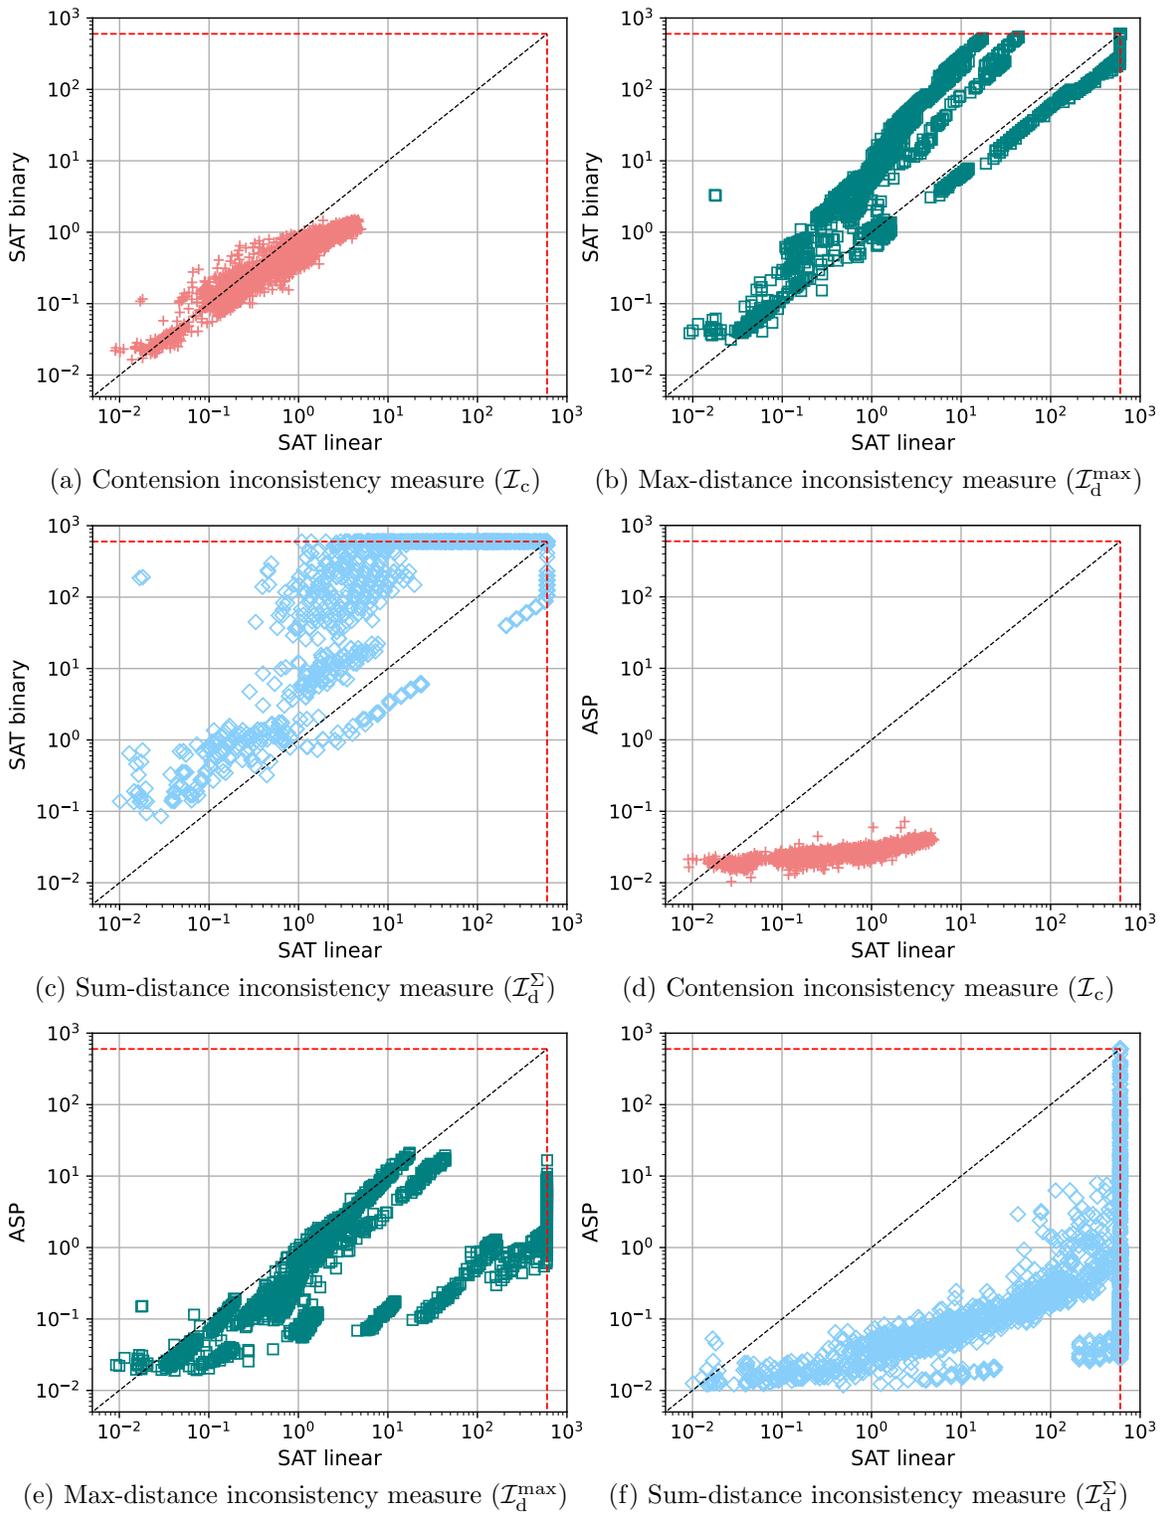


Figure 29: Runtime comparison of the SAT-based approaches based on linear search (for  $\mathcal{I}_c$ ,  $\mathcal{I}_d^{\max}$ , and  $\mathcal{I}_d^{\Sigma}$ ) and the corresponding binary search versions, as well as the ASP-based versions, on the SRS data set. Timeout: 600 seconds.

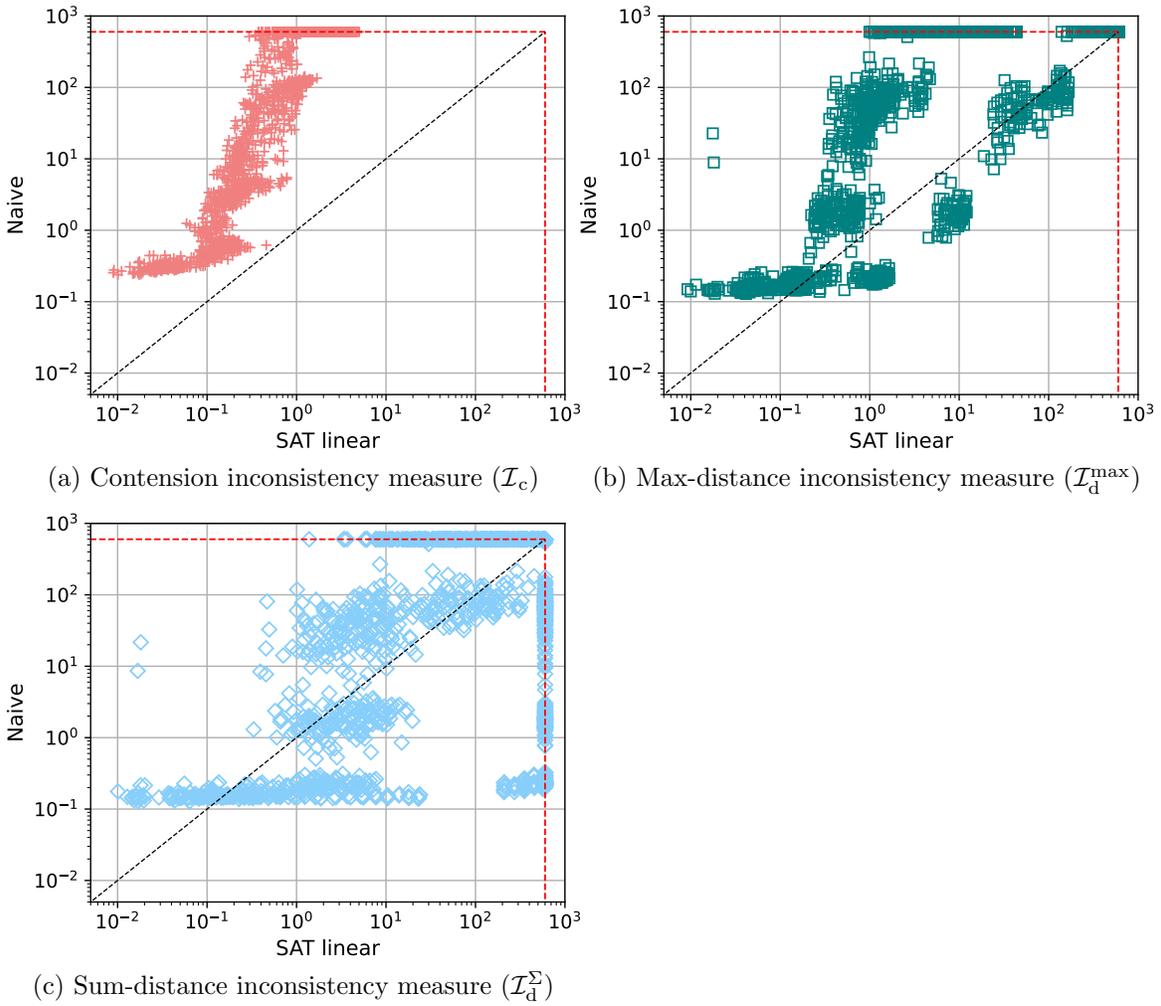


Figure 30: Runtime comparison of the SAT-based approaches based on linear search (for  $\mathcal{I}_c$ ,  $\mathcal{I}_d^{\max}$ , and  $\mathcal{I}_d^{\Sigma}$ ) and the corresponding naive methods, on the SRS data set. Timeout: 600 seconds.

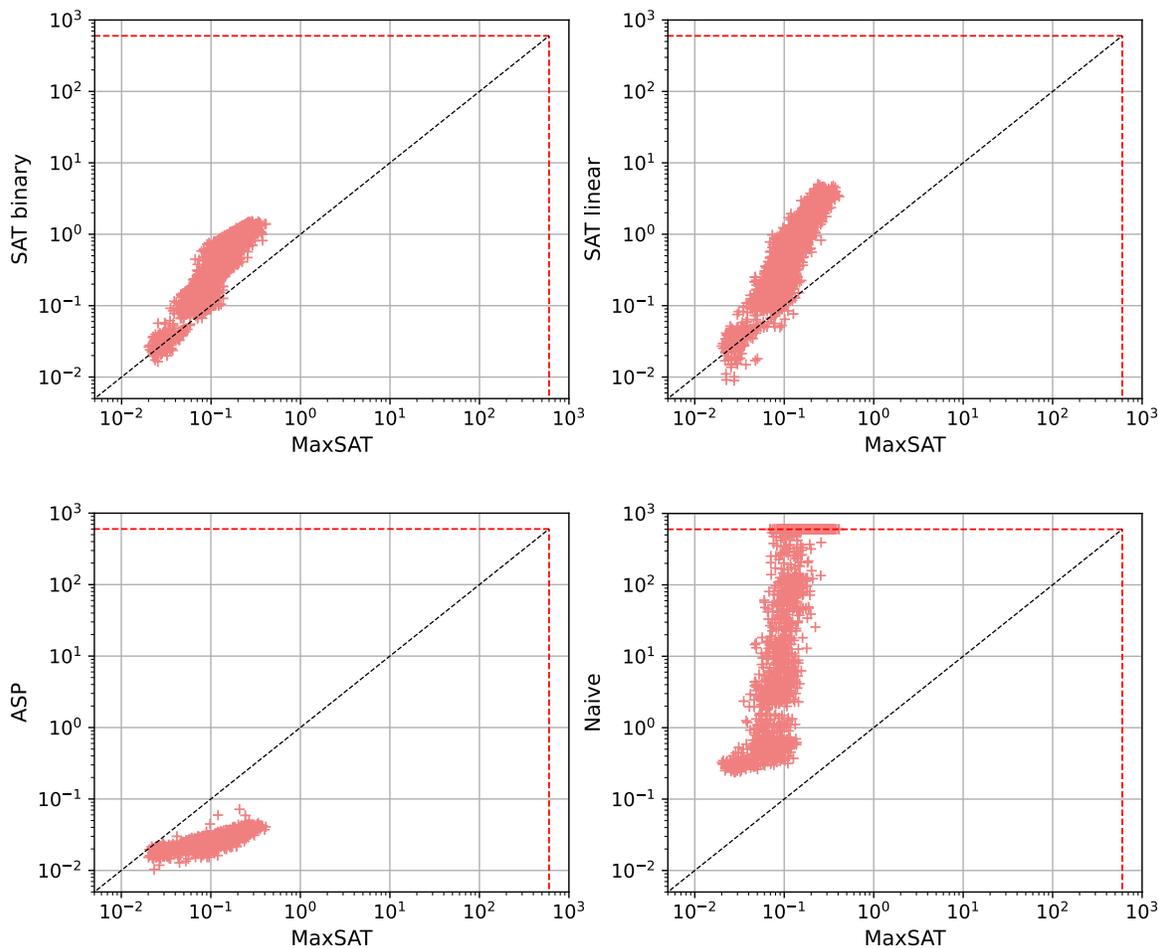


Figure 31: Runtime comparison of the MaxSAT approach for  $\mathcal{L}_c$  and the corresponding other approaches on the SRS data set. Timeout: 600 seconds.

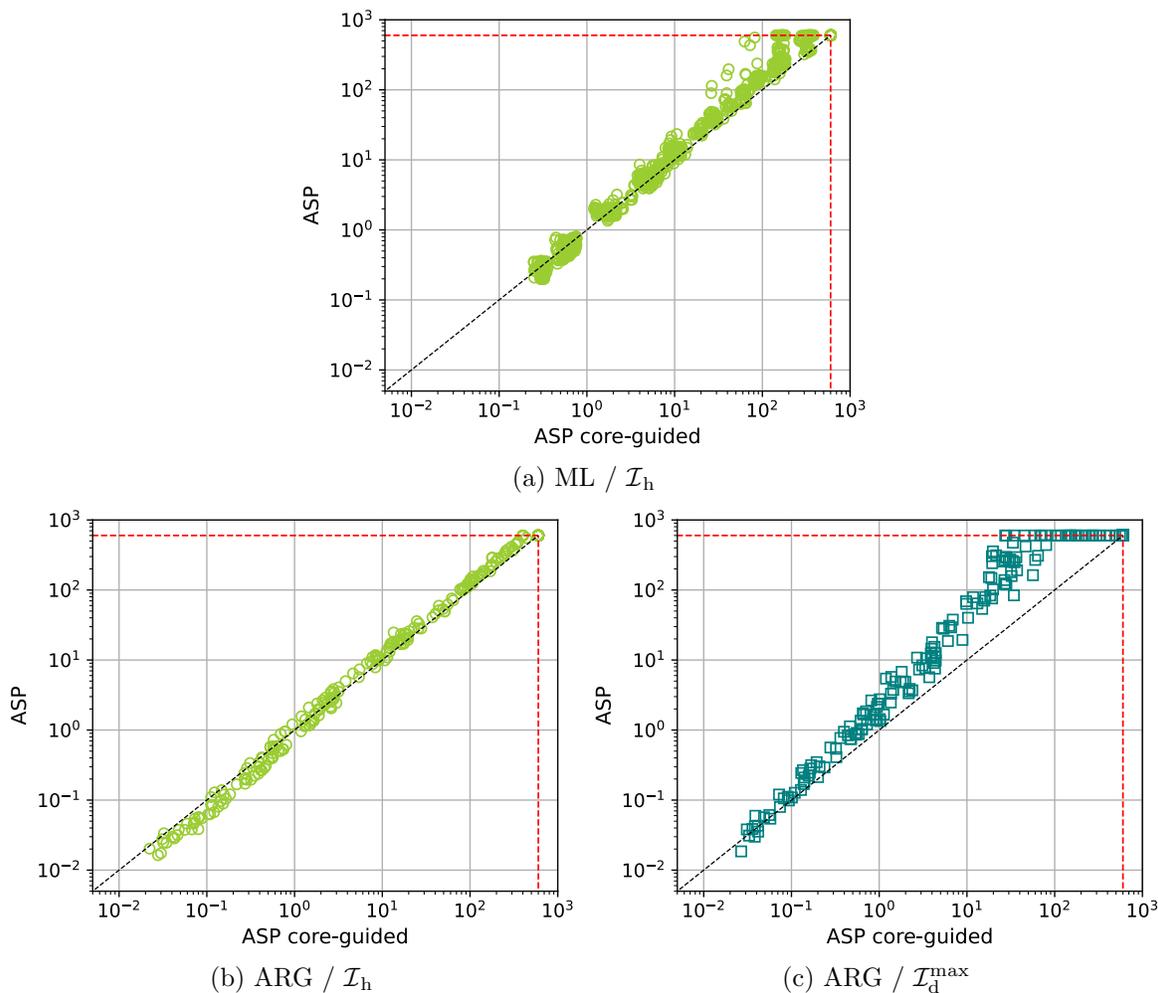
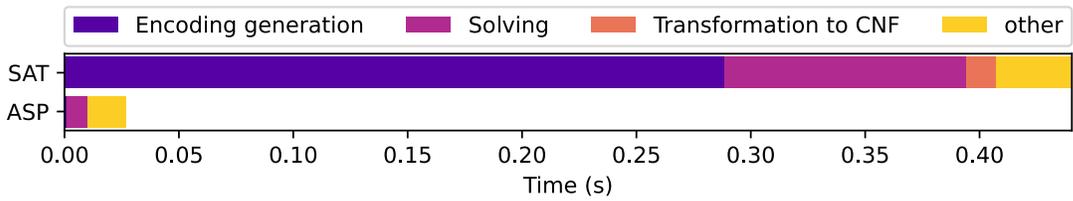


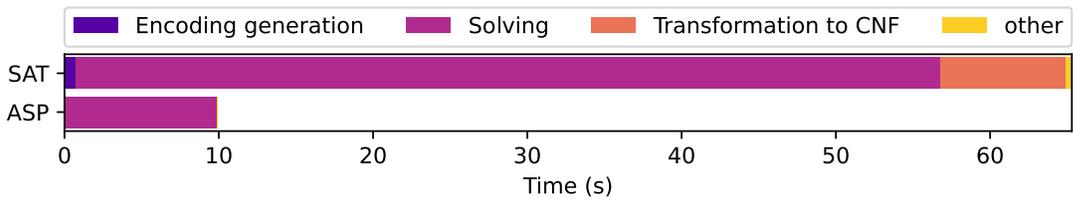
Figure 32: Runtime comparison of the ASP approach with core-guided vs. with branch-and-bound optimization (default) on the ML data set wrt. the hitting set measure ( $\mathcal{I}_h$ ) and on the ARG data set wrt.  $\mathcal{I}_h$  and the max-distance measure ( $\mathcal{I}_d^{\max}$ ). Timeout: 10 minutes.

#### **B.4 Runtime Composition**

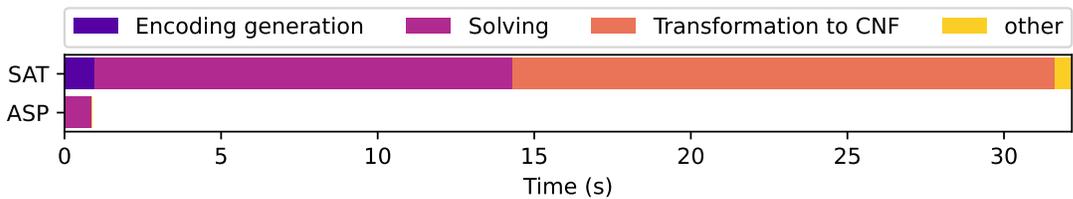
This section contains additional bar plots visualizing the average runtime composition of the SAT-based and ASP-based approaches wrt. each measure and the SRS data set (Figure 33), as well as the ARG data set (Figure 34).



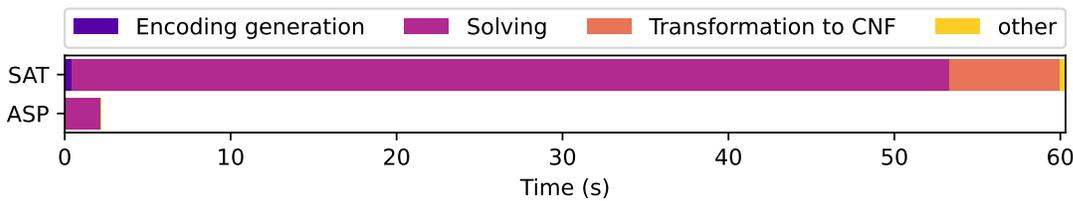
(a) Contension inconsistency measure ( $\mathcal{I}_c$ )



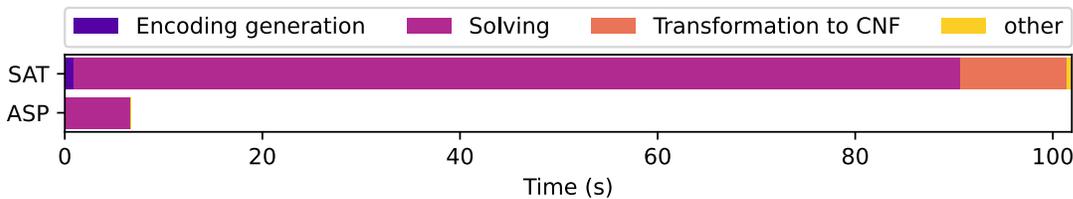
(b) Forgetting-based inconsistency measure ( $\mathcal{I}_f$ )



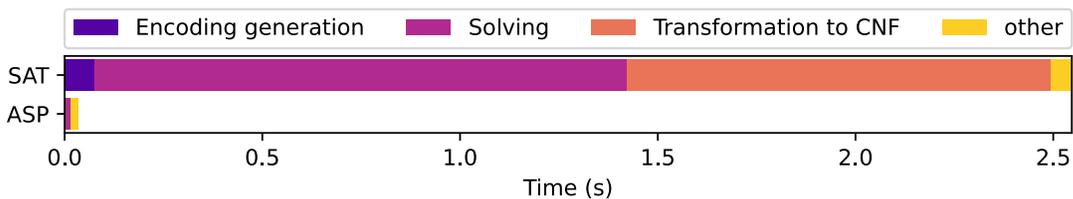
(c) Hitting Set inconsistency measure ( $\mathcal{I}_h$ )



(d) Max-distance inconsistency measure ( $\mathcal{I}_d^{\max}$ )

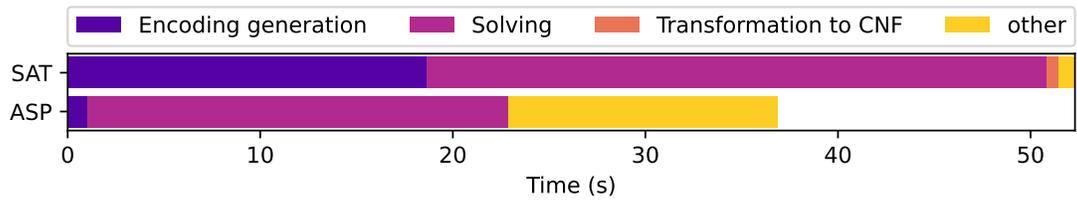


(e) Sum-distance inconsistency measure ( $\mathcal{I}_d^{\Sigma}$ )

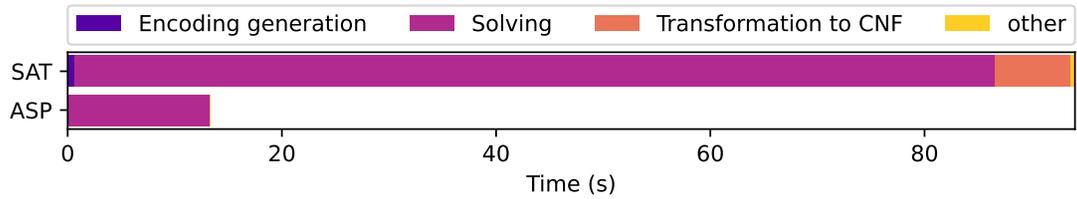


(f) Hit-distance inconsistency measure ( $\mathcal{I}_d^{\text{hit}}$ )

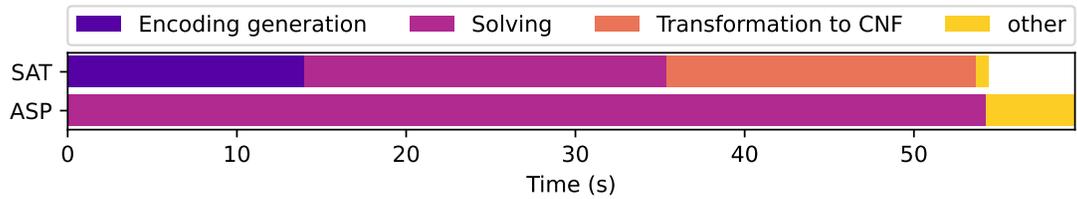
Figure 33: Comparison between the ASP-based and SAT-based approaches in terms of runtime composition wrt. the SRS data set.



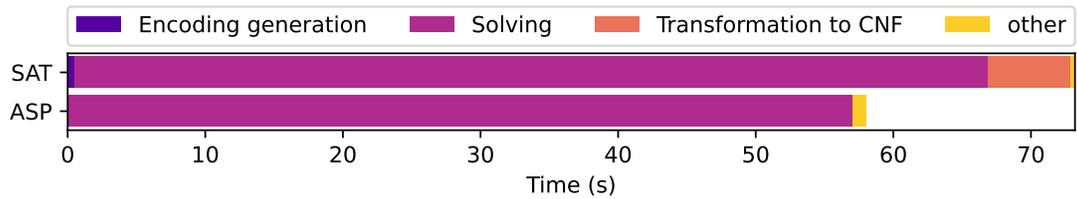
(a) Contension inconsistency measure ( $\mathcal{I}_c$ )



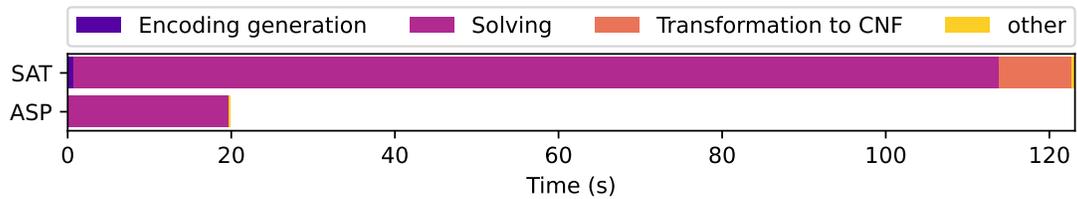
(b) Forgetting-based inconsistency measure ( $\mathcal{I}_f$ )



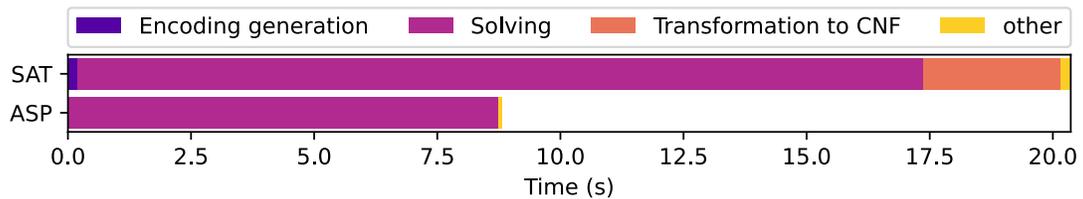
(c) Hitting Set inconsistency measure ( $\mathcal{I}_h$ )



(d) Max-distance inconsistency measure ( $\mathcal{I}_d^{\max}$ )



(e) Sum-distance inconsistency measure ( $\mathcal{I}_d^{\Sigma}$ )



(f) Hit-distance inconsistency measure ( $\mathcal{I}_d^{\text{hit}}$ )

Figure 34: Comparison between the ASP-based and SAT-based approaches in terms of runtime composition wrt. the ARG data set.

## References

- Abío, I., Nieuwenhuis, R., Oliveras, A., & Rodríguez-Carbonell, E. (2013). A parametric approach for smaller and better encodings of cardinality constraints. In *19th International Conference on Principles and Practice of Constraint Programming, CP'13*, pp. 80–96.
- Agrawal, R., & Srikant, R. (1994). Fast algorithms for mining association rules in large databases. In *Proceedings VLDB'94*, pp. 487–499.
- Ammoura, M., Raddaoui, B., Salhi, Y., & Oukacha, B. (2015). On measuring inconsistency using maximal consistent sets. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, pp. 267–276. Springer.
- Andres, B., Kaufmann, B., Matheis, O., & Schaub, T. (2012). Unsatisfiability-based optimization in clasp. In *Technical Communications of the 28th International Conference on Logic Programming (ICLP'12)(2012)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
- Avellaneda, F. (2020). A short description of the solver EvalMaxSAT. *MaxSAT Evaluation*, 8.
- Bacchus, F., Järvisalo, M., & Martins, R. (2021). Maximum satisfiability. In *Handbook of Satisfiability*, pp. 929–991. IOS Press.
- Banbara, M., Gebser, M., Inoue, K., Ostrowski, M., Peano, A., Schaub, T., Soh, T., Tamura, N., & Weise, M. (2015). aspartame: Solving constraint satisfaction problems with answer set programming. In *Logic Programming and Nonmonotonic Reasoning: 13th International Conference, LPNMR 2015, Lexington, KY, USA, September 27-30, 2015. Proceedings 13*, pp. 112–126. Springer.
- Berg, J., Järvisalo, M., Martins, R., Niskanen, A., & Paxian, T. (2024). MaxSAT evaluation 2024: Solver and benchmark descriptions..
- Bertossi, L. (2018a). Measuring and computing database inconsistency via repairs. In *Proceedings of the 12th International Conference on Scalable Uncertainty Management (SUM'18)*.
- Bertossi, L. E. (2018b). Repair-based degrees of database inconsistency: Computation and complexity. *CoRR*, abs/1809.10286.
- Besnard, P. (2016). Forgetting-based inconsistency measure. In *Scalable Uncertainty Management: 10th International Conference, SUM 2016, Nice, France, September 21-23, 2016, Proceedings 10*, pp. 331–337. Springer.
- Besnard, P., Doutre, S., & Herzig, A. (2014). Encoding argument graphs in logic. In Laurent, A., Strauss, O., Bouchon-Meunier, B., & Yager, R. R. (Eds.), *International Conference on Information Processing and Management of Uncertainty in Knowledge-based Systems - IPMU 2014*, Vol. 443 of *Communications in Computer and Information Science*, pp. 345–354. Springer.
- Biere, A., Cimatti, A., Clarke, E. M., Fujita, M., & Zhu, Y. (1999). Symbolic model checking using SAT procedures instead of BDDs. In *Proceedings of the 36th annual ACM/IEEE Design Automation Conference*, pp. 317–320.

- Biere, A., Fazekas, K., Fleury, M., & Heisinger, M. (2020). CaDiCaL, Kissat, Paracooba, Plingeling and Treengeling entering the SAT Competition 2020. In *Proc. of SAT Competition 2020 – Solver and Benchmark Descriptions*, Vol. B-2020-1 of *Department of Computer Science Report Series B*, pp. 51–53. University of Helsinki.
- Biere, A., Heule, M., Maaren, H., & Walsh, T. (2009). *Handbook of Satisfiability*. Frontiers in artificial intelligence and applications. IOS Press.
- Brewka, G., Eiter, T., & Truszczynski, M. (2011). Answer set programming at a glance. *Communications of the ACM*, 54(12), 92–103.
- Calimeri, F., Faber, W., Gebser, M., Ianni, G., Kaminski, R., Krennwallner, T., Leone, N., Maratea, M., Ricca, F., & Schaub, T. (2020). ASP-Core-2 input language format. *Theory and Practice of Logic Programming*, 20(2), 294–309.
- Cholvy, L., Perrussel, L., & Thevenin, J.-M. (2017). Using inconsistency measures for estimating reliability. *International Journal of Approximate Reasoning*, 89, 41–57.
- Clark, K. (1978). Negation as failure. In Gallaire, H., & Minker, J. (Eds.), *Logic and Data Bases*, pp. 293–322. Plenum Press, New York.
- Corea, C., Grant, J., & Thimm, M. (2022). Measuring inconsistency in declarative process specifications. In *Proceedings of the 20th International Conference on Business Process Management (BPM'22)*.
- Corea, C., Thimm, M., & Delfmann, P. (2021). Measuring inconsistency over sequences of business rule cases. In Biennu, M., Lakemeyer, G., & Erdem, E. (Eds.), *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning (KR'21)*, pp. 656–660. IJCAI Organization.
- Decker, H., & Misra, S. (2017). Database inconsistency measures and their applications. In *Proceedings of the 23rd International Conference on Information and Software Technologies (ICIST 2017)*.
- Dung, P. M. (1995). On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2), 321–358.
- Eyupoglu, S., Fidan, M., Gulesen, Y., Izci, I. B., Teber, B., Yilmaz, B., Alkan, A., & Erdem, E. (2021). Stable marriage problems with ties and incomplete preferences: An empirical comparison of ASP, SAT, ILP, CP, and local search methods..
- Fu, Z., & Malik, S. (2006). On solving the partial MAX-SAT problem. In *International Conference on Theory and Applications of Satisfiability Testing*, pp. 252–265. Springer.
- Gebser, M., Kaminski, R., Kaufmann, B., Romero, J., & Schaub, T. (2015). Progress in clasp series 3. In *Logic Programming and Nonmonotonic Reasoning: 13th International Conference, LPNMR 2015, Lexington, KY, USA, September 27-30, 2015. Proceedings 13*, pp. 368–383. Springer.
- Gebser, M., Kaminski, R., Kaufmann, B., & Schaub, T. (2012). Answer set solving in practice. *Synthesis lectures on artificial intelligence and machine learning*, 6(3), 1–238.

- Gebser, M., Kaminski, R., Kaufmann, B., & Schaub, T. (2019). Multi-shot ASP solving with clingo. *Theory and Practice of Logic Programming*, 19(1), 27–82.
- Gebser, M., Kaufmann, B., & Schaub, T. (2012). Conflict-driven answer set solving: From theory to practice. *Artificial Intelligence*, 187, 52–89.
- Gelfond, M., & Lifschitz, V. (1988). The stable model semantics for logic programming. In *Proc. ICLP/SLP*, pp. 1070–1080. MIT Press.
- Giunchiglia, E., & Maratea, M. (2006). Solving optimization problems with DLL. In *ECAI*, Vol. 141, pp. 377–381.
- Grant, J., & Hunter, A. (2011). Measuring consistency gain and information loss in stepwise inconsistency resolution. In *Proceedings ECSQARU'11*, pp. 362–373. Springer.
- Grant, J. (1978). Classifications for inconsistent theories. *Notre Dame Journal of Formal Logic*, 19(3), 435–444.
- Grant, J., & Hunter, A. (2013). Distance-based measures of inconsistency. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, pp. 230–241. Springer.
- Grant, J., & Hunter, A. (2017). Analysing inconsistent information using distance-based measures. *Int. J. Approx. Reasoning*, 89(C), 3–26.
- Grant, J., & Martinez, M. V. (Eds.). (2018). *Measuring Inconsistency in Information*, Vol. 73 of *Studies in Logic*. College Publications.
- Hunter, A. (2006). How to act on inconsistent news: Ignore, resolve, or reject. *Data & Knowledge Engineering*, 57(3), 221–239.
- Hunter, A., Konieczny, S., et al. (2008). Measuring inconsistency through minimal inconsistent sets.. *KR*, 8(358-366), 42.
- Jabbour, S., Ma, Y., & Raddaoui, B. (2014). Inconsistency measurement thanks to MUS decomposition. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pp. 877–884.
- Jabbour, S., & Sais, L. (2016). Exploiting MUS structure to measure inconsistency of knowledge bases. In *Proceedings of the Twenty-Second European Conference on Artificial Intelligence*, ECAI'16, p. 991–998, NLD. IOS Press.
- Kaminski, R., Schaub, T., & Wanko, P. (2017). A tutorial on hybrid answer set solving with Clingo. In *Reasoning Web. Semantic Interoperability on the Web: 13th International Summer School 2017*, pp. 167–203. Springer.
- Knight, K. (2002). Measuring inconsistency. *Journal of Philosophical Logic*, 31(1), 77–98.
- Kuhlmann, I., Gessler, A., Laszlo, V., & Thimm, M. (2022). A comparison of ASP-based and SAT-based algorithms for the contension inconsistency measure. In *Proceedings of the 15th International Conference on Scalable Uncertainty Management (SUM'22)*. Springer.
- Kuhlmann, I., & Thimm, M. (2020). An algorithm for the contension inconsistency measure using reductions to answer set programming. In *Proceedings of the 14th International Conference on Scalable Uncertainty Management (SUM'20)*, pp. 289–296. Springer.

- Kuhlmann, I., & Thimm, M. (2021). Algorithms for inconsistency measurement using answer set programming. In *Proceedings of the 19th International Workshop on Non-Monotonic Reasoning (NMR)*, pp. 159–168.
- Li, C. M., & Manyà, F. (2021). MaxSAT, hard and soft constraints. In *Handbook of satisfiability*, pp. 903–927. IOS Press.
- Lifschitz, V. (2008). What is answer set programming?. In *Proceedings AAAI’08*, pp. 1594–1597.
- Lifschitz, V. (2019). *Answer set programming*. Springer Berlin.
- Ma, Y., Qi, G., Xiao, G., Hitzler, P., & Lin, Z. (2009). An Anytime Algorithm for Computing Inconsistency Measurement. In *3rd International Conference on Knowledge Science, Engineering and management*, pp. 29–40, Austria.
- Marques-Silva, J., & Manquinho, V. (2008). Towards more effective unsatisfiability-based maximum satisfiability algorithms. In *International Conference on Theory and Applications of Satisfiability Testing*, pp. 225–230. Springer.
- Marques-Silva, J. P., & Sakallah, K. A. (2000). Boolean satisfiability in electronic design automation. In *Proceedings of the 37th Annual Design Automation Conference*, pp. 675–680.
- Martinez, A. B. B., Arias, J. J. P., & Vilas, A. F. (2004). On measuring levels of inconsistency in multi-perspective requirements specifications. In *Proceedings of the 1st Conference on the Principles of Software Engineering (PRISE’04)*.
- McAreavey, K., Liu, W., & Miller, P. (2014). Computational approaches to finding and measuring inconsistency in arbitrary knowledge bases. *International Journal of Approximate Reasoning*, 55.
- Mironov, I., & Zhang, L. (2006). Applications of SAT solvers to cryptanalysis of hash functions. In *International Conference on Theory and Applications of Satisfiability Testing*, pp. 102–115. Springer.
- Nejati, S., & Ganesh, V. (2019). CDCL (crypto) SAT solvers for cryptanalysis. In *Proceedings of the 29th Annual International Conference on Computer Science and Software Engineering*, CASCON ’19, p. 311–316, USA. IBM Corp.
- Potyka, N., & Thimm, M. (2017). Inconsistency-tolerant reasoning over linear probabilistic knowledge bases. *International Journal of Approximate Reasoning*, 88, 209–236.
- Priest, G. (1979). The logic of paradox. *Journal of Philosophical Logic*, 8(1), 219–241.
- Reiter, R. (1980). A logic for default reasoning. *Artificial intelligence*, 13(1-2), 81–132.
- Sinz, C. (2005). Towards an optimal CNF encoding of Boolean cardinality constraints. In van Beek, P. (Ed.), *Principles and Practice of Constraint Programming - CP 2005*, pp. 827–831, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Thimm, M. (2016a). On the expressivity of inconsistency measures. *Artificial Intelligence*, 234, 120–151.
- Thimm, M. (2016b). Stream-based inconsistency measurement. *International Journal of Approximate Reasoning*, 68, 68 – 87.

- Thimm, M. (2019). Inconsistency measurement. In Amor, N. B., Quost, B., & Theobald, M. (Eds.), *Proceedings of the 13th International Conference on Scalable Uncertainty Management (SUM'19)*, Vol. 11940 of *Lecture Notes in Artificial Intelligence*, pp. 9–23. Springer International Publishing.
- Thimm, M., & Rienstra, T. (2020). Approximate reasoning with ASPIC+ by argument sampling. In *Proceedings of the Third International Workshop on Systems and Algorithms for Formal Argumentation (SAFA'20)*.
- Thimm, M., & Wallner, J. P. (2019). On the complexity of inconsistency measurement. *Artificial Intelligence*, 275, 411–456.
- Tseitin, G. S. (1968). On the complexity of derivation in propositional calculus. In *Structures in Constructive Mathematics and Mathematical Logic, Part II*. Steklov Mathematical Institute.
- van Harmelen, F., Lifschitz, V., & Porter, B. (2008). *Handbook of Knowledge Representation*. ISSN. Elsevier Science.
- Vizel, Y., Weissenbacher, G., & Malik, S. (2015). Boolean satisfiability solvers and their applications in model checking. *Proceedings of the IEEE*, 103(11), 2021–2035.
- Xiao, G., & Ma, Y. (2012). Inconsistency measurement based on variables in minimal unsatisfiable subsets. In *Proceedings of European Conference on Artificial Intelligence (ECAI'12)*, pp. 864–869.