

Preference-Based Inconsistency Management in Multi-Context Systems (Extended Abstract)*

Thomas Eiter¹ and Antonius Weinzierl²

¹ Institut für Logic and Computation, Technische Universität Wien, Vienna, Austria

² Department of Computer Science, Aalto University, Espoo, Finland
 eiter@kr.tuwien.ac.at, antonius.weinzierl@aalto.fi

Abstract


Establishing information exchange between existing knowledge-based systems can lead to devastating inconsistency. Automatic resolution of inconsistency often is unsatisfactory, because any modification of the information flow may lead to bad or even dangerous conclusions. Methods to identify and select preferred repairs of inconsistency are thus needed. In this work, we leverage the expressive power and generality of Multi-Context Systems (MCS), a formalism for information exchange, to select most preferred repairs, by use of a meta-reasoning transformation. As for computational complexity, finding preferred repairs is not higher than the base case; finding most-preferred repairs is higher, yet worst-case optimal.

1 Introduction

Multi-Context Systems (MCS) are an approach to share information between individual knowledge bases, called *emphcontexts*, which exchange information using special rules, called *bridge rules*. The semantics of the whole system then emerges from the local semantics of the constituent knowledge bases. Practical applications of MCS are defeasible reasoning in ambient intelligence [Bikakis and Antoniou, 2010], cooperation in distributed information systems [Caire and Bikakis, 2011], and the METIS system for maritime situation awareness support [Velikova *et al.*, 2014].

MCS are rooted in the work of McCarthy [1993] and enjoyed a rich development (cf. [Giunchiglia and Serafini, 1994; Ghidini and Giunchiglia, 2001; Roelofsen and Serafini, 2005; Brewka *et al.*, 2007; Serafini and Homola, 2012; Bikakis and Antoniou, 2010; Brewka *et al.*, 2011]). For generality, this work is based on MCS with heterogeneous contexts and non-monotonic information exchange [Brewka and Eiter, 2007].

As the contexts of an MCS are typically autonomous and host knowledge bases that are inherited legacy systems, information exchange may lead to inconsistency; which means



- $r_1: (2: \textit{hyperglycemia}) \leftarrow (1: \textit{hyperglycemia}).$
- $r_2: (2: \textit{allow_anim_insulin}) \leftarrow \textit{not} (1: \textit{allergic_anim_insulin}).$
- $r_3: (3: \textit{bill_anim_insulin}) \leftarrow (2: \textit{give_anim_insulin}).$
- $r_4: (3: \textit{bill_human_insulin}) \leftarrow (2: \textit{give_human_insulin}).$
- $r_5: (3: \textit{insurance_B}) \leftarrow (1: \textit{insurance_B}).$

Figure 1: Hospital MCS $M = (C_1, C_2, C_3)$ and its bridge rules.

that the MCS has no global model (called *equilibrium*); thus, the whole MCS becomes useless. The notion of *diagnosis* can be used to repair an inconsistent MCS [Eiter *et al.*, 2010; 2014] and it indicates bridge rules that must be modified in order to restore consistency. But the modified information exchange may have serious consequences.

Example 1. Consider an MCS M employed in a hospital, which interconnects three systems: (1) a patient database with information about illnesses, insurance companies, and potential allergies of patients; (2) an expert system suggesting proper medications for illnesses; and (3) a system billing the insurance company of patients for the administered treatments. The system is shown in Figure 1 together with its bridge rules r_1 to r_5 ; for illustration we assume that there is only one patient. Bridge rule r_1 states that the medication context C_2 knows that the patient has *hyperglycemia* if the patient database C_1 believes so; bridge rule r_2 states that C_2 is allowed to administer animal insulin if an allergy to animal insulin is not present in the patient database. The other bridge rules inform the billing system C_3 about the patient’s insurer (r_3) and what insulin was administered (r_4 and r_5). The expert system only recommends treatments to which patients are not allergic to and the billing system only allows administered treatments that are covered by the insurance companies. Suppose now that the patient’s illness requires insulin and she/he is allergic to animal insulin. If the insurer of the patient only pays for animal insulin, the only allowed treat-

*This paper is an extended abstract of an article in the Journal of Artificial Intelligence Research [Eiter and Weinzierl, 2017]. This extended abstract has been supported by the Academy of Finland, project 251170 and the Austrian Science Fund (FWF) project P27730. Initial research was funded by WWTF, grant ICT08-020.

ment is rejected by the billing system and M is inconsistent.

To repair M , consider its set of minimal diagnoses:

$$D_m^\pm(M) = \{ (\{r_1\}, \emptyset), (\{r_4\}, \emptyset), (\{r_5\}, \emptyset), (\emptyset, \{r_2\}) \},$$

where $(\{r_1\}, \emptyset)$ indicates to remove r_1 , which means the illness of the patient is ignored; $(\{r_4\}, \emptyset)$ removes r_4 , hence the (correct) medication is not billed; $(\{r_5\}, \emptyset)$ ignores the insurer, likely billing it for some medication it will not pay; $(\emptyset, \{r_2\})$ indicates to make the bridge rule r_2 *condition-free*, i.e., remove all conditions from r_2 and by that add the allowance to administer animal insulin although the patient is allergic to it; this may threaten the health of the patient. No diagnosis is clearly the best, and depending on one's preferences one diagnosis may be better than another.

The goal of this work is to develop a machinery for automatic identification of preferred diagnoses and pruning unwanted ones, in order to only require from a human operator to select from a reduced set (of most preferred diagnoses) the best diagnosis manually. What constitutes a preferred or best diagnosis cannot be decided in general since it will be different for each MCS and depends on the environment into which an MCS is embedded. In the above example, the health of patients may be paramount, alternatively billing may be of highest importance from an economic perspective.

Automatic selection of preferred diagnoses according to some preference requires in turn a formalism for expressing and evaluating preferences, like e.g. *ceteris paribus* preferences [Doyle *et al.*, 1991], CP-nets [Boutilier *et al.*, 2004; Domshlak *et al.*, 2001; Goldsmith *et al.*, 2008], or utility functions [Von Neumann and Morgenstern, 1944]. As there is no one-fits-all preference formalism our approach is based on the idea that a user-customized preference on diagnoses, specified in a formalism chosen by the user, can be seen as a knowledge-base or context of an MCS. We enable such a context to “see” the diagnoses of the MCS and provide extended notions to select the most preferred diagnoses based on the observations and preferences expressed by this context.

This paper gives a quick overview of our main contributions: 1. two basic methods for preference-based selection of diagnoses, 2. a meta-reasoning transformation to enable that a context observes diagnoses applied to an MCS, 3. extended notions of diagnosis for the selection of preferred diagnoses, and 4. computational complexity results on these notions.

2 Preliminaries

To capture all kinds of knowledge-representation formalisms, MCS use the concept of an *abstract “logic”*, which is a triple $L = (\mathbf{KB}, \mathbf{BS}, \mathbf{ACC})$ where \mathbf{KB} is the set of knowledge bases (“inputs”), \mathbf{BS} is the set of possible belief sets (“outputs”), and $\mathbf{ACC} : \mathbf{KB} \rightarrow 2^{\mathbf{BS}}$ is a function assigning each knowledge base a set of acceptable belief sets (“semantics”). This captures classical logic, ontologies, logic programs, etc.

Information exchange is specified using *bridge rules*, which are similar to rules in logic programming, of form:

$$(k : s) \leftarrow (c_1 : p_1), \dots, (c_i : p_i), \\ \mathbf{not} (c_{i+1} : p_{i+1}), \dots, \mathbf{not} (c_m : p_m).$$

Intuitively such a bridge rule reads: if all beliefs p_1, \dots, p_i are present at contexts c_1, \dots, c_i , respectively, and beliefs

p_{i+1}, \dots, p_m are not present at c_{i+1}, \dots, c_m , respectively, then the knowledge-base formula s is added to context k .

A *Multi-Context System* $M = (C_1, \dots, C_n)$ consists of contexts $C_i = (L_i, kb_i, br_i)$, $1 \leq i \leq n$, where (i) $L_i = (\mathbf{KB}_i, \mathbf{BS}_i, \mathbf{ACC}_i)$ is an abstract logic, (ii) $kb_i \in \mathbf{KB}_i$ is a knowledge base, and (iii) br_i is a set of bridge rules. By $br(M) = \bigcup_{i=1}^n br_i$ we denote the set of all bridge rules of M . The semantics of an MCS $M = (C_1, \dots, C_n)$ is defined over sequences $S = (S_1, \dots, S_n)$ of belief sets $S_i \in \mathbf{BS}_i$, $1 \leq i \leq n$ called *belief states*. A bridge rule r is applicable in S , if for each positive body literal $(j : p)$ holds $p \in S_j$, while for each (negative) **not** $(j : p)$ holds $p \notin S_j$. We denote by $(k : s) \leftarrow \perp$ and $(k : s) \leftarrow \top$ bridge rules that are never and always applicable, respectively. A belief state S of M is an *equilibrium* if every belief set is locally accepted given applicable bridge rules. $\text{EQ}(M)$ denotes the set of all equilibria; M is *inconsistent* if M has no equilibrium, i.e., $\text{EQ}(M) = \emptyset$.

For repairing (i.e., diagnosing) an inconsistent MCS, pairs (D_1, D_2) of sets of bridge rules are considered, such that if we deactivate the rules in D_1 and add the rules in D_2 in condition-free form, then the MCS admits an equilibrium. Formally, a *diagnosis* of M is a pair (D_1, D_2) , $D_1, D_2 \subseteq br(M)$, such that $\text{EQ}(M[br(M) \setminus D_1 \cup cf(D_2)]) \neq \emptyset$; $D^\pm(M)$ denotes the set of all diagnoses. A diagnosis indicates which bridge rules are assumed to require modification in order to obtain a consistent MCS, i.e., a diagnosis constitutes a way to *repair* an MCS. We call a pair $D = (D_1, D_2) \in 2^{br(M)} \times 2^{br(M)}$ a *candidate diagnosis*. By Occam's razor, one may consider preferable diagnoses to be the ones that require the least modifications, i.e., (subset-)minimal diagnoses. For a pairs $A = (A_1, A_2)$ and $B = (B_1, B_2)$ we say $A \subset B$ iff $A_1 \subseteq B_1$, $A_2 \subseteq B_2$ and $A \neq B$. Then, a diagnosis $D \in D^\pm(M)$ is (pointwise) subset *minimal*, if no $D' \subset D$ is in $D^\pm(M)$. Furthermore, $D_m^\pm(M)$ denotes all such D .

3 Preferences

In the literature two forms of preference prevail: (1) filters, which consider a single outcome (i.e., a diagnosis) and discard it if it fails some condition; (2) (binary) preferences, which compare two outcomes and decide which is better. In the spirit of MCS we also want this approach to be open to any kind of formalism for specifying filters or preferences.

A filter can be seen as hard constraints on diagnosis candidates; for generality, a filter basically is a Boolean function.

Definition 1. A *diagnosis filter for an MCS M with bridge rules $br(M)$* is a function $f : 2^{br(M)} \times 2^{br(M)} \rightarrow \{0, 1\}$, and the set of filtered diagnoses is:

$$D_f^\pm(M) = \{(D_1, D_2) \in D^\pm(M) \mid f(D_1, D_2) = 1\}.$$

$D_{m,f}^\pm(M)$ is the set of subset-minimal diagnoses passing f .

In general, a preference on diagnoses is just a binary order relation; to avoid counter-intuitive results, we require preferences to be transitive. Since virtually every other preference formalism yields an order relation, we introduce it as such.

Definition 2. A *preference order over diagnoses for an MCS M* is a transitive and reflexive binary relation \preceq on $2^{br(M)} \times 2^{br(M)}$; for $D, D' \in 2^{br(M)} \times 2^{br(M)}$ we say that D is *preferred to D'* if $D \preceq D'$.

Using \preceq we define what constitutes a most preferred diagnosis. Intuitively, such a diagnosis incurs a minimal set of modifications and no other diagnosis is strictly more preferred. We first introduce \preceq -preferred diagnoses, which are those diagnoses such that no other diagnosis is strictly more preferred. The most preferred diagnoses then are the subset-minimal ones from the set of \preceq -preferred diagnoses.

Definition 3. A diagnosis $D \in D^\pm(M)$ of an MCS M is called \preceq -preferred if for all $D' \in 2^{br(M)} \times 2^{br(M)}$ with $D' \prec D$ it holds that $D' \notin D^\pm(M)$. D is minimal \preceq -preferred if D is subset-minimal among all \preceq -preferred diagnoses. The set of all \preceq -preferred diagnoses is denoted by $D_{\preceq}^\pm(M)$ and the set of all minimal \preceq -preferred by $D_{m,\preceq}^\pm(M)$.

We do not require \preceq to be acyclic and by relying on \prec (the irreflexive and anti-symmetric version of \preceq) we consider all diagnoses in a cycle to be equally preferred.

4 Meta-Reasoning for Diagnosis

In order to realize filters and preference orders inside an MCS, some context must be able to reason on the diagnoses of the MCS. We achieve this based on the following insights.

(A) If the rules of a context, say context C_{n+1} , are of certain form, it knows exactly whether some rules have been removed or made condition-free by a diagnosis and it can then exhibit this knowledge as beliefs.

$$(n+1 : not_removed_r) \leftarrow \top. \quad (1)$$

$$(n+1 : uncond_r) \leftarrow \perp. \quad (2)$$

Rule (1) is always firing, so if a diagnosis removes it, the knowledge *not_removed_r* is no longer present (provided it is not obtained in any other way). If a diagnosis makes this rule condition-free, it has no effect since it already is, i.e., its body is \top . For rule (2) it is exactly the other way round: it never fires so removing it has no consequences but if it is made condition-free, then the knowledge *uncond_r* will be present.

(B) We can simulate with beliefs of an equilibrium the effects to a bridge rule of the modification by a diagnosis, i.e., if certain beliefs are present, the rule either stops firing or fires unconditionally. For example, let r be the bridge rule $(3 : c) \leftarrow (1 : a), \mathbf{not} (2 : b)$, and consider these bridge rules:

$$(3 : c) \leftarrow (1 : a), \mathbf{not} (2 : b), \mathbf{not} (n+1 : removed_r). \quad (3)$$

$$(3 : c) \leftarrow (n+1 : uncond_r). \quad (4)$$

Rule (3) behaves like r as long as context $n+1$ does not believe *removed_r* and it behaves as removed otherwise (i.e., it does not derive c for context 3). Similarly, rule (4) behaves like r being condition-free if $n+1$ believes *uncond_r*. Replacing r by (3) and (4) thus enables simulation of a diagnosis.

By (A), a context can observe (parts) of a diagnosis and by (B), it can exhibit this observation as beliefs to other contexts whose rules then behave as being modified by the diagnosis. We can use this to establish a 1-1 correspondence between (effects of) diagnoses on an MCS M and (effects of) diagnoses on an extended MCS M' where diagnoses are observed by an *observation context* C_{n+1} . To also ensure that bridge rules like (3) and (4) are not modified in a diagnosis, we introduce the notion of a diagnosis excluding protected rules.

Definition 4. Given an MCS M with protected rules $br_P \subseteq br(M)$, a diagnosis excluding protected rules is a diagnosis $(D_1, D_2) \in D^\pm(M)$, where $D_1, D_2 \subseteq br(M) \setminus br_P$. We denote the set of all such diagnoses by $D^\pm(M, br_P)$ and by $D_m^\pm(M, br_P)$ the set of all minimal such diagnoses.

Consider a diagnosis filter f on diagnosis candidates of M , if C_{n+1} observes all bridge rules of the original MCS M , then C_{n+1} sees in M' whether a diagnosis candidate (D_1, D_2) of M is applied (assuming that only bridge rules of the form (3) and (4) are not protected). So, in principle, the logic inside C_{n+1} could evaluate $f(D_1, D_2) = t$ and in case of $t = 0$ context C_{n+1} could become inconsistent. By that, all diagnoses excluding protected rules of M' are diagnoses passing the filter f . That is, we found a way to obtain filtered diagnoses and since we did not specify what logic is used inside C_{n+1} , any formalism deemed appropriate may be used.

For the meta-reasoning to serve as uniform foundation for filters and preferences, we introduce a property θ that describes the additional behavior of the context C_{n+1} . For preferences a set \mathcal{K}_p of additional bridge rules at C_{n+1} whose head beliefs are all distinct and whose bodies are \perp , is required. Intuitively, \mathcal{K}_p are additional bridge rules which may be forced by θ to become condition-free. By that, θ can (arbitrarily) map diagnosis candidates of M into diagnoses modifying \mathcal{K}_p . So θ specifies how a diagnosis candidate maps into \mathcal{K}_p . An (informal) definition of $M^{mr(\theta, \mathcal{K}_p)}$ is as follows.

Definition 5. Given an MCS $M = (C_1, \dots, C_n)$, \mathcal{K}_p , and $\theta : 2^{br(M)} \times 2^{br(M)} \times 2^{\mathcal{K}_p}$, the MCS $M^{mr(\theta, \mathcal{K}_p)} = (C'_1, \dots, C'_n, C_{n+1})$ is a meta-reasoning encoding if:

(i) C'_i equals C_i except that all its bridge rules are replaced with ones like in (3) and (4).

(ii) $C_{n+1} = (L_{n+1}, kb_{n+1}, br_{n+1})$ is a context such that:

- the logic $L_{n+1} = (\mathbf{KB}_{n+1}, \mathbf{BS}_{n+1}, \mathbf{ACC}_{n+1})$ is compatible with bridge rules in \mathcal{K}_p and of form (1) and (2);
- the bridge rules br_{n+1} are those in \mathcal{K}_p and for each $r \in br(M)$ a bridge rule of form (1) and of form (2); and
- the knowledge base $kb_{n+1} \in \mathbf{KB}_{n+1}$ fulfills: for every set H of possible inputs by bridge rules, there is an acceptable belief set $S_{n+1} \in \mathbf{ACC}_{n+1}(kb_{n+1} \cup H)$ iff S_{n+1} exhibits the bridge rule modifications through beliefs *removed_r* and *uncond_r*, and θ holds for H .

The context C_{n+1} is similar to an interface in programming, which defines certain properties but is open to an arbitrary implementation. In the next section we illustrate how θ and \mathcal{K}_p can be used for preference realization.

5 Preference Realization

Since a preference order compares two diagnosis candidates, but meta-reasoning can only observe one diagnosis at a time, our solution is to clone the original MCS M such that we obtain an MCS $2M$ that contains two independent copies of M . Every diagnosis of $2M$ then is comprised of two diagnoses of M . Since the latter are independent of each other, applying the meta-reasoning encoding yields the MCS $(2M)^{mr(\theta, \mathcal{K}_p)}$ where the observation context sees two (independent) diagnoses of M and it may compare them, mapping the result

of this comparison to \mathcal{K}_p . The minimal \preceq -preferred diagnoses can then be selected from $M^{\preceq} = (2M)^{mr(\theta, \mathcal{K}_p)}$ using a slightly more involved diagnosis notion with prioritized bridge rules. Instead of defining the clone encoding here, we illustrate M^{\preceq} in Figure 2 for Example 1.

There are three types of bridge rules at the observation context C_7 : (i) bridge rules of the form $di(r_j)$, $i \in \{1, 2\}$, that encode two diagnoses for M , one per each clone, (ii) prioritized bridge rules of the forms $in_i(r_j)$ and $\bar{in}_i(r_j)$, $i \in \{1, 2\}$, that encode the diagnosis of the first clone again, and (iii) the bridge rule t_{max} . The latter two types comprise the set \mathcal{K}_p of the meta-reasoning transformation.

The idea of the clone encoding M^{\preceq} is that t_{max} is used as a flag whether the diagnosis applied to clone 2 is preferred over the diagnosis applied to clone 1. Specifically, if the diagnosis of clone 2 is preferred over the diagnosis of clone 1, then t_{max} does not lead to an inconsistency, but in all other cases it does. This means, that every diagnosis of M^{\preceq} where the diagnosis of clone 1 is minimal \preceq -preferred must contain t_{max} . So θ can be readily given by: θ holds for H iff H corresponds to a candidate diagnosis (D_1, D_2) of clone 1 and a candidate diagnosis (D'_1, D'_2) of clone 2, and one of the following holds, where $\mathcal{K}(D_1, D_2)$ is the representation of (D_1, D_2) by prioritized bridge rules of the form $in_i(r_j)$ and $\bar{in}_i(r_j)$, $i \in \{1, 2\}$, and R_p are all such rules in H :

- $(D_1, D_2) = (D'_1, D'_2)$ and $R_p = \mathcal{K}(D_1, D_2) \cup \{t_{max}\}$, or
- $(D'_1, D'_2) \prec (D_1, D_2)$ and $R_p = \mathcal{K}(D_1, D_2)$.

If for a given diagnosis of clone 1, there exists a more preferred diagnosis of clone 2, then there exists a diagnosis without t_{max} . Selecting a diagnosis modifying a minimal set of prioritized bridge rules thus selects a \preceq -preferred diagnosis.

Definition 6. Let M be an MCS with protected rules br_P , and prioritized rules br_H . A diagnosis $D \in D_m^{\pm}(M, br_P)$ is prioritized-minimal if no $D' \in D_m^{\pm}(M, br_P)$ modifies a smaller set of prioritized bridge rules; $D^{\pm}(M, br_P, br_H)$ is the set of all prioritized-minimal diagnoses.

Below we denote by $t(D_1, D_2)$ the candidate diagnosis in M^{\preceq} that directly corresponds to (D_1, D_2) of M . Then:

Theorem 1. Let M be an MCS and let \preceq be a preference order on the diagnoses of M . Then $(D_1, D_2) \in D^{\pm}(M)$ is \preceq -preferred iff $t(D_1, D_2) \in D^{\pm}(M^{\preceq}, br_P, br_H)$ holds.

Finally, selecting a diagnosis that modifies a minimal set of prioritized bridge rules and contains t_{max} corresponds to selecting a minimal \preceq -preferred diagnosis.

Definition 7. Let M be an MCS with protected rules br_P and prioritized rules br_H . A subset-minimal prioritized-minimal (mpm) diagnosis wrt. t_{max} is a diagnosis that is prioritized-minimal, contains t_{max} , and is subset-minimal (w.r.t. non-prioritized bridge rules) among all prioritized-minimal diagnoses that contain t_{max} . The set of all such diagnoses is denoted by $D_{m, t_{max}}^{\pm}(M, br_P, br_H)$.

We show that the clone encoding is sound and complete.

Theorem 2. Let M be an MCS and let \preceq be a preference order on diagnoses of M . Then $(D_1, D_2) \in D_{m, \preceq}^{\pm}(M)$ holds iff $t(D_1, D_2) \in D_{m, t_{max}}^{\pm}(M^{\preceq}, br_P, br_H)$ holds.

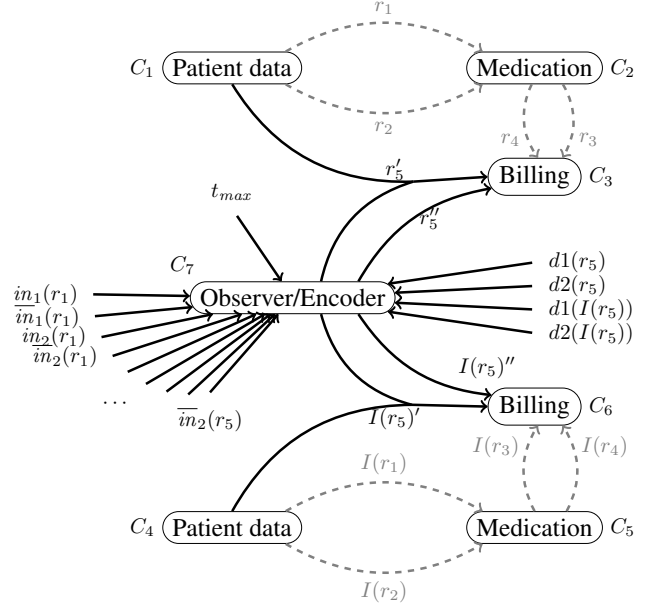


Figure 2: The clone encoding MCS $M^{\preceq} = (C_1, C_2, \dots, C_7)$ with two clones (top and bottom) of the MCS M from Example 1 and the observation context C_7 . Only some bridge rules of C_7 are shown and the bridge rules stemming from r_5 . Dashed and gray lines indicate the other bridge rules of the clones whose resulting bridge rules in M^{\preceq} are omitted. The prioritized bridge rules of M^{\preceq} are t_{max} and all bridge rules $in_i(r_j)$ and $\bar{in}_i(r_j)$.

MCS D_m	MCS DP_m	MCS DPH	MCS $DPH_{m, t_{max}}$
D_1^P	D_1^P	D_1^P	Π_2^P
D_i^P	D_i^P	D_i^P	Π_{i+1}^P

Table 1: Complexity results for extended diagnoses (selection).

6 Computational Complexity

We analyzed the computational complexity of each of the following decision problems for a given diagnosis candidate D of an MCS M with bridge rules $br(M)$, protected rules br_P and prioritized rules br_H : MCS DP_m , decide whether $D \in D_m^{\pm}(M, br_P)$ holds; MCS DPH , decide whether $D \in D^{\pm}(M, br_P, br_H)$ holds; MCS $DPH_{m, t_{max}}$, decide whether $D \in D_{m, t_{max}}^{\pm}(M, br_P, br_H)$ holds; We also analyzed MCS D_{MPREF} : given a preference order \preceq decide whether $D \in D_{m, \preceq}^{\pm}(M)$, i.e., if D is minimal \preceq -preferred. The results are given in Table 1, where MCS D_m is the complexity of deciding whether $D \in D_m^{\pm}(M)$ holds (cf. [Eiter et al., 2014]). Notably, the problems for our extended notions of diagnosis do not have higher complexity than the basic notion, except for MCS $DPH_{m, t_{max}}$. But we show that the underlying problem MCS D_{MPREF} is hard for the same complexity class, thus our solution is worst-case optimal.

7 Conclusion

The full article [Eiter and Weinzierl, 2017] contains formal notions, much more details, examples, and many results omitted for this abstract, including another encoding for total preference orders, and an appendix with full proofs.

References

- [Bikakis and Antoniou, 2010] Antonis Bikakis and Grigoris Antoniou. Defeasible contextual reasoning with arguments in ambient intelligence. *IEEE Trans. Knowl. Data Eng.*, 22(11):1492–1506, 2010.
- [Boutilier *et al.*, 2004] Craig Boutilier, Ronen I. Brafman, Carmel Domshlak, Holger H. Hoos, and David Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res. (JAIR)*, 21:135–191, 2004.
- [Brewka and Eiter, 2007] Gerhard Brewka and Thomas Eiter. Equilibria in heterogeneous nonmonotonic multi-context systems. In *AAAI Conference on Artificial Intelligence (AAAI)*, pages 385–390, 2007.
- [Brewka *et al.*, 2007] Gerhard Brewka, Floris Roelofsen, and Luciano Serafini. Contextual default reasoning. In Manuela M. Veloso, editor, *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007*, pages 268–273, 2007.
- [Brewka *et al.*, 2011] Gerhard Brewka, Thomas Eiter, and Michael Fink. Nonmonotonic multi-context systems: A flexible approach for integrating heterogeneous knowledge sources. In Marcello Balduccini and Tran Cao Son, editors, *Logic Programming, Knowledge Representation, and Nonmonotonic Reasoning*, volume 6565 of *Lecture Notes in Computer Science*, pages 233–258. Springer, 2011.
- [Caire and Bikakis, 2011] Patrice Caire and Antonis Bikakis. Enhancing cooperation in distributed information systems using conviviality and multi-context systems. In Chatrakul Sombatheera, Arun Agarwal, Siba K. Udgate, and Kittichai Lavangnananda, editors, *Multi-disciplinary Trends in Artificial Intelligence - 5th International Workshop, MIWAI 2011, Hyderabad, India, December 7-9, 2011. Proceedings*, volume 7080 of *Lecture Notes in Computer Science*, pages 14–25. Springer, 2011.
- [Domshlak *et al.*, 2001] Carmel Domshlak, Ronen I. Brafman, and Solomon Eyal Shimony. Preference-based configuration of web page content. In Bernhard Nebel, editor, *IJCAI*, pages 1451–1456. Morgan Kaufmann, 2001.
- [Doyle *et al.*, 1991] Jon Doyle, Yoav Shoham, and Michael P. Wellman. A logic of relative desire (preliminary report). In Zbigniew W. Ras and Maria Zemankova, editors, *Methodologies for Intelligent Systems, 6th International Symposium, ISMIS '91, Charlotte, N.C., USA, October 16-19, 1991, Proceedings*, volume 542 of *Lecture Notes in Computer Science*, pages 16–31. Springer, 1991.
- [Eiter and Weinzierl, 2017] Thomas Eiter and Antonius Weinzierl. Preference-based inconsistency management in multi-context systems. *J. Artif. Intell. Res.*, 60:347–424, 2017. <https://doi.org/10.1613/jair.5416>.
- [Eiter *et al.*, 2010] Thomas Eiter, Michael Fink, Peter Schüller, and Antonius Weinzierl. Finding explanations of inconsistency in multi-context systems. In Fangzhen Lin, Ulrike Sattler, and Mirosław Trzuszczński, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Twelfth International Conference, KR 2010, Toronto, Ontario, Canada, May 9-13, 2010*, pages 329–339. AAAI Press, 2010.
- [Eiter *et al.*, 2014] Thomas Eiter, Michael Fink, Peter Schüller, and Antonius Weinzierl. Finding explanations of inconsistency in multi-context systems. *Artificial Intelligence*, 216(0):233–274, 2014.
- [Ghidini and Giunchiglia, 2001] Chiara Ghidini and Fausto Giunchiglia. Local models semantics, or contextual reasoning=locality+compatibility. *Artif. Intell.*, 127(2):221–259, 2001.
- [Giunchiglia and Serafini, 1994] Fausto Giunchiglia and Luciano Serafini. Multilanguage hierarchical logics or: How we can do without modal logics. *Artif. Intell.*, 65(1):29–70, 1994.
- [Goldsmith *et al.*, 2008] Judy Goldsmith, Jérôme Lang, Mirosław Trzuszczński, and Nic Wilson. The computational complexity of dominance and consistency in CP-nets. *J. Artif. Intell. Res. (JAIR)*, 33:403–432, 2008.
- [McCarthy, 1993] John McCarthy. Notes on formalizing context. In Ruzena Bajcsy, editor, *Proceedings of the 13th International Joint Conference on Artificial Intelligence. Chambéry, France, August 28 - September 3, 1993*, pages 555–562. Morgan Kaufmann, 1993.
- [Roelofsen and Serafini, 2005] Floris Roelofsen and Luciano Serafini. Minimal and absent information in contexts. In Leslie Pack Kaelbling and Alessandro Saffiotti, editors, *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK, July 30-August 5, 2005*, pages 558–563. Professional Book Center, 2005.
- [Serafini and Homola, 2012] Luciano Serafini and Martin Homola. Contextualized knowledge repositories for the semantic web. *J. Web Sem.*, 12:64–87, 2012.
- [Velikova *et al.*, 2014] Marina Velikova, Peter Novák, Bas Huijbrechts, Jan Laarhuis, Jesper Hoeksma, and Steffen Michels. An integrated reconfigurable system for maritime situational awareness. In Torsten Schaub, Gerhard Friedrich, and Barry O’Sullivan, editors, *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, pages 1197–1202. IOS Press, 2014.
- [Von Neumann and Morgenstern, 1944] John Von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.