

The Consistency of the CADIAG-2 Knowledge Base: A Probabilistic Approach

Pavel Klinov, Bijan Parsia and David Picado-Muiño

University of Manchester
Oxford Road, Manchester M13 9PL, UK
{pklinov|bparsia}@cs.man.ac.uk

Institut für Diskrete Mathematik und Geometrie
Wiedner Hauptstrasse 8 / 104. A-1040 Vienna, Austria
picado@logic.at

Abstract. The paper presents the methodology and the results of checking consistency of the knowledge base of CADIAG-2, a large-scale medical expert system. Such knowledge base consists of a large collection of rules representing knowledge about various medical entities (symptoms, signs, diseases...) and relationships between them. The major portion of the rules are uncertain, i.e., they specify *to what degree* a medical entity is *confirmed* by another medical entity or a combination of them. Given the size of the system and the uncertainty it has been challenging to validate its consistency. Recent attempts to partially formalise CADIAG-2's knowledge base into decidable Gödel logics have shown that, on that formalisation, CADIAG-2 is inconsistent. In this paper we verify this result with an alternative, more expressive formalisation of CADIAG-2 as a set of probabilistic conditional statements and apply a state-of-the-art probabilistic logic solver to determine *satisfiability* of the knowledge base and to extract conflicting sets of rules. As CADIAG-2 is too large to be handled out of the box we describe an approach to split the knowledge base into fragments that can be tested independently and prove that such methodology is complete (i.e., is guaranteed to find all conflicts). With this approach we are able to determine that CADIAG-2 contains numerous sets of conflicting rules and compute all of them for a slightly relaxed version of the knowledge base.

1 Introduction

CADIAG-2 (Computer Assisted DIAGnosis) is a well-known rule-based expert system aimed at providing support in diagnostic decision making in the field of internal medicine. Its design and construction was initiated in the early 80's at the Medical University of Vienna by K.P. Adlassnig – see [1], [2], [3] or [4] for more on the origins and design of CADIAG-2.

CADIAG-2 consists of two fundamental pieces: the inference engine and the knowledge base. The inference engine — see [5] or [6] for alternative formalisations and analyses of CADIAG-2's inference — is based on methods of approximate reasoning in fuzzy set theory, in the sense of [7] and [8]. In fact CADIAG-2

is presented in some monographs as an example of a fuzzy expert system — see [9], [10]. The knowledge base consists of a set of *IF-THEN* rules — known in the literature as *production* rules — intended to represent relationships between distinct medical entities: symptoms, findings, signs and test results on the one hand and diseases and therapies on the other. The vast majority of them are binary (i.e., they relate single medical entities) and only such rules are considered in this paper. The one that follows is an example of a binary rule of CADIAG-2 (taken from [3]):

```
IF suspicion of liver metastases by liver palpation
THEN pancreatic cancer
with degree of confirmation 0.3
```

The *degree of confirmation* refers, intuitively, to the degree to which the antecedent (i.e., '*suspicion of liver metastases by liver palpation*' in the example above) confirms the consequent (i.e., '*pancreatic cancer*' above).

In this paper we present a formalisation of a coded version of the binary fragment of CADIAG-2's knowledge base (i.e., that contains only codes for the identification of the distinct medical entities) as a probabilistic logic theory. We then check the satisfiability of that formalisation with *Pronto*, our probabilistic Description Logic (DL) reasoner, which we briefly introduce. We find that CADIAG-2 is highly unsatisfiable (confirming the results of an alternative, weaker formalisation, [11]) and analyse the sources of unsatisfiability.

To our knowledge, the probabilistic version of CADIAG-2 is the largest PSAT (Probabilistic SATisfiability) problem to be solved by an automated reasoner and is certainly the largest non-artificial one. This is, perhaps, a bit misleading as it is comparatively easy to detect unsatisfiability by first heuristically detecting small but likely unsatisfiable fragments, and then performing a satisfiability check on each fragment. While this might suffice to validate that the knowledge base is unsatisfiable it is not sufficient, without further qualification, to detect all conflicting sets of rules, nor can it ensure that a satisfiable fragment is so in the context of the entire knowledge base. As CADIAG-2 is too large (the number of rules in the binary fragment we are concerned with is over 18000) we describe an approach to split the knowledge base into comparatively large fragments that can be tested independently and prove that such methodology is complete (i.e., is guaranteed to find all conflict sets). With this methodology we are able to determine that CADIAG-2 contains numerous sets of conflicting rules and compute all of them for a slightly relaxed interpretation of the knowledge base.

2 Notation and preliminary definitions

Throughout we will be working with a finite set $\mathcal{L} = S \cup D = \{P_1, \dots, P_n\}$ of unary predicates in a first-order language, for some $n \in \mathbb{N}$. \mathcal{L} is intended to represent the set of medical entities occurring in the inference rules of the expert system CADIAG-2, with S the set of symptoms, findings, signs and test results

(to which we will commonly refer as *symptoms*) and D the set of therapies and diseases (to which we will commonly refer as *diseases*).

Definition 1. An interpretation \mathcal{I} of \mathcal{L} is a pair $(D^{\mathcal{I}}, V^{\mathcal{I}})$, where $D^{\mathcal{I}}$ is a finite, non-empty set and $V^{\mathcal{I}}$ is a map from $\mathcal{L} \times D^{\mathcal{I}}$ to $[0, 1]$.

An interpretation \mathcal{I} is said to be *classical* if $V^{\mathcal{I}}(P, a) \in \{0, 1\}$ for all $(P, a) \in \mathcal{L} \times D^{\mathcal{I}}$. It is said to be *rational* if $V^{\mathcal{I}}(P, a) \in [0, 1] \cap \mathbb{Q}$ for all $(P, a) \in \mathcal{L} \times D^{\mathcal{I}}$.

Given an interpretation \mathcal{I} of \mathcal{L} , we will refer to the elements in $D^{\mathcal{I}}$ by latin characters a, b, c, \dots and to the elements in \mathcal{L} by uppercase latin characters P, Q, \dots (possibly with suffices).

Let $L = \{p_1, \dots, p_n\}$ be a finite propositional language, for $n \in \mathbb{N}$, and SL its closure under boolean connectives.

Definition 2. Let $w : SL \rightarrow [0, 1]$. We say that w is a probability function on L if the following two conditions hold, for all $\theta, \phi \in SL$:

- If $\models \theta$ then $w(\theta) = 1$.
- If $\models \neg(\theta \wedge \phi)$ then $w(\theta \vee \phi) = w(\theta) + w(\phi)$.¹

We can restrict probability functions on L to values in $[0, 1] \cap \mathbb{Q}$. We will call such probability functions *rational*.

A probability distribution w on L can be characterized by the values it assigns to the expressions of the form $\pm p_1 \wedge \dots \wedge \pm p_n$, which we call *states* or *worlds*, where $+p$ and $-p$ stand for p and $\neg p$ respectively. We denote the set of states in L by W and define, for $\phi \in SL$, W_ϕ as follows:

$$W_\phi = \{s \in W \mid s \models \phi\}.$$

We define conditional probability from the notion of unconditional probability on conditional statements in L in the conventional way.

3 The knowledge base in CADIAG-2

We can classify CADIAG-2's binary rules (Φ_{CB}) into three different types: rules in which both antecedent and consequent are medical entities in S (*symptom-symptom*, $\Phi_{S|S}$), rules in which both antecedent and consequent are medical entities in D (*disease-disease*, $\Phi_{D|D}$) and those in which the antecedent is a medical entity in S and the consequent an entity in D (*symptom-disease*, $\Phi_{D|S}$).² The degree of confirmation in a rule of the first two types is a value in the set $\{0, 1\}$ and it is in this sense that we say that rules of these types are *classical*.

¹ Here (and throughout) \models is classical entailment.

² CADIAG-2's knowledge base formally contains values for conditional relations with a medical entity in D as the antecedent and a medical entity in S as the consequent. However, such rules are not used by CADIAG-2's inference mechanism and therefore are not taken into account in this paper.

Let $\langle P, Q, \eta \rangle \in \Phi_{CB}$ be a rule in CADIAG-2's binary knowledge base, with $P, Q \in \mathcal{L}$ and $\eta \in [0, 1] \cap \mathbb{Q}$. The value η is intended to quantify the *degree* to which P (the antecedent) *confirms* Q (the consequent) and claimed in most of the literature on CADIAG-2 (see, for example, [2] or [3]) to have been calculated from a certain database or interpretation \mathcal{I} as follows:³

$$\frac{\sum_{a \in D^{\mathcal{I}}} \min\{V^{\mathcal{I}}(P, a), V^{\mathcal{I}}(Q, a)\}}{\sum_{a \in D^{\mathcal{I}}} V^{\mathcal{I}}(P, a)} = \eta \quad (1)$$

Notice that this expression generalises the concept of *conditional probability* or *frequency* that one gets when restricting the model to classical interpretations. Under such a restriction η becomes a probability (in the sense of frequency) and its meaning is intuitive and formally well understood. However, when allowing in (1) valuations other than classical the intuitive meaning of η is lost. Certainly, (1) would benefit from a serious attempt to clarify its meaning when allowing valuations other than classical but it is not the purpose of this paper to do so. As we will see later, whether the interpretation of (1) is assumed in terms of any valuations or in terms of only classical valuations (i.e., a probabilistic interpretation) will be indifferent to our satisfiability-checking purposes.

Throughout we will use the expression $Q/\mathcal{I}P = \eta$ to abbreviate (1). Sometimes, in order to generalise results, we will be considering an interval, say $\Omega \subseteq [0, 1]$, instead of a single value (i.e., η in (1)) and we will be using the expression $Q/\mathcal{I}P \in \Omega$ to abbreviate the corresponding modification of (1). Such modification is motivated by the possibility of alternative, suitable interpretations of the rules in Φ_{CB} that one could consider interesting in the view of some theoretical or practical aspects. Among these alternative interpretations we consider replacing η in equation (1) by the interval $[\eta, 1]$ (i.e., consider η a lower bound for the degrees of confirmation instead of a precise one) or replacing η whenever $\eta \in (0, 1)$ by an interval of the form $[\eta - \epsilon, \eta + \epsilon]$, for ϵ small (i.e., a slightly relaxed interpretation of Φ_{CB}).

We will denote the collection of real intervals in $[0, 1]$ by \mathfrak{J} . We will normally refer to intervals of the form $[\eta, \eta] \in \mathfrak{J}$ by η itself.

For the next definition and proposition let \mathcal{I} be an interpretation of \mathcal{L} and $\Phi \subseteq \mathcal{R}_{\mathcal{L}}$, for

$$\mathcal{R}_{\mathcal{L}} = \{\langle P, Q, \Omega \rangle \mid P, Q \in \mathcal{L}, \Omega \in \mathfrak{J}\}.$$

Definition 3. We say that \mathcal{I} is a model of Φ (denoted $\models_{\mathcal{I}} \Phi$) if $Q/\mathcal{I}P \in \Omega$ for all $\langle P, Q, \Omega \rangle \in \Phi$.

Proposition 1. Φ has a classical model if and only if it has a rational model.

Proof. The right implication follows trivially from the fact that every classical interpretation is also rational. In order to prove the left implication let us assume that $\mathcal{I} = (D^{\mathcal{I}}, V^{\mathcal{I}})$ is a rational interpretation such that $\models_{\mathcal{I}} \Phi$.

³ We say in *most* of the literature. There are some references in which the interpretation suggested for η in $\langle P, Q, \eta \rangle$ is different. For example in [1] it is claimed that η can be interpreted as a frequency and thus $\langle P, Q, \eta \rangle$ as a probabilistic conditional statement.

Let C be the set given by the values $V^{\mathcal{I}}(P, a)$, for $(P, a) \in \mathcal{L} \times D^{\mathcal{I}}$. It is assumed that all the values in C are rational. Let us consider the minimum common multiple of the denominators of all the elements of C , say $q \in \mathbb{N}$. We next construct a new interpretation \mathcal{J} from \mathcal{I} such that $\models_{\mathcal{J}} \Phi$.⁴

We first define $D^{\mathcal{J}}$ from $D^{\mathcal{I}}$. For each element $a \in D^{\mathcal{I}}$ we set q elements in the domain $D^{\mathcal{J}}$, labelled as follows: $\{a_1, \dots, a_q\}$. Let us consider now $P \in \mathcal{L}$ and $a \in D^{\mathcal{I}}$ and assume that $V^{\mathcal{I}}(P, a) = \frac{p}{q}$. We define $V^{\mathcal{J}}$ on $\mathcal{L} \times D^{\mathcal{J}}$ from $V^{\mathcal{I}}$ as follows, for $i \in \{1, \dots, q\}$:

$$V^{\mathcal{J}}(P, a_i) = \begin{cases} 1 & \text{if } i \leq p \\ 0 & \text{otherwise} \end{cases}$$

It is easy to see that \mathcal{J} thus defined is such that $\models_{\mathcal{J}} \Phi$.

For what follows we will be considering the collection of intervals \mathfrak{J}^* in the set $[0, 1]$. \mathfrak{J}^* differs from \mathfrak{J} in that an interval $\Omega \in \mathfrak{J}^*$ needs to have its maximum and/or minimum in \mathbb{Q} , provided it has a maximum and/or a minimum.

We define

$$R_{\mathcal{L}}^* = \{\langle P, Q, \Omega \rangle \mid P, Q \in \mathcal{L}, \Omega \in \mathfrak{J}^*\}$$

and consider $\Phi \subseteq R_{\mathcal{L}}^*$ for the next proposition.

Proposition 2. *If Φ has a model then it has a rational model.*

Proof. Let \mathcal{I} be an interpretation such that $\models_{\mathcal{I}} \Phi$. We then have that, for all $\langle P, Q, \Omega \rangle$ in Φ , $Q/\mathcal{I}P \in \Omega$.

For each $Q/\mathcal{I}P \in \Omega$ we consider the inequalities

$$\sum_{a \in D^{\mathcal{I}}} \eta_1 V^{\mathcal{I}}(P, a) < \sum_{a \in D^{\mathcal{I}}} \min\{V^{\mathcal{I}}(P, a), V^{\mathcal{I}}(Q, a)\} < \sum_{a \in D^{\mathcal{I}}} \eta_2 V^{\mathcal{I}}(P, a),$$

where Ω is assumed to be of the form $(\eta_1, \eta_2) \in \mathfrak{J}^*$ (for $\Omega \in \mathfrak{J}^*$ of any other form we replace ' $<$ ' in the inequalities above by ' \leq ' as required) and with

$$\min\{V^{\mathcal{I}}(P, a), V^{\mathcal{I}}(Q, a)\}$$

replaced by $V^{\mathcal{I}}(P, a)$ or $V^{\mathcal{I}}(Q, a)$ accordingly.

Let us also consider, for each $Q/\mathcal{I}P \in \Omega$, the inequalities

$$0 \leq V^{\mathcal{I}}(P, a), V^{\mathcal{I}}(Q, a) \leq 1,$$

$$\sum_{a \in D^{\mathcal{I}}} V^{\mathcal{I}}(P, a) > 0$$

and, for $V^{\mathcal{I}}(P, a)$ greater than $V^{\mathcal{I}}(Q, a)$, the inequality $V^{\mathcal{I}}(Q, a) \leq V^{\mathcal{I}}(P, a)$ (the inequality $V^{\mathcal{I}}(Q, a) \geq V^{\mathcal{I}}(P, a)$ otherwise).

The solution set of the linear system above with unknown values $V^{\mathcal{I}}(P, a)$, $V^{\mathcal{I}}(Q, a)$ is not empty (\mathcal{I} is assumed to be a solution of the system) and needs to contain rational solutions (due to the form of the intervals in \mathfrak{J}^*). Therefore, there has to exist a rational interpretation of \mathcal{L} that models Φ .

⁴ It is worth stressing here again that the predicates in \mathcal{L} , and thus the predicates involved in (1), are atomic predicates.

Corollary 1. Φ has a model if and only if it has a classical model.

By Corollary 1 we have that Φ_{CB} will have a model if and only if it has a classical model.

For the next definition let w be a probability function on \mathcal{L} , where \mathcal{L} here is regarded as a propositional language, and $\langle P, Q, \Omega \rangle$ a triple in $\mathcal{R}_{\mathcal{L}}$, with $P, Q \in \mathcal{L}$ and $\Omega \in \mathcal{J}$.

Definition 4. We say that the probability function w on \mathcal{L} satisfies $\langle P, Q, \Omega \rangle$ (denoted $\models_w \langle P, Q, \Omega \rangle$) if $w(P) > 0$ and

$$\frac{w(P \wedge Q)}{w(P)} \in \Omega.$$

In that sense we say that $\langle P, Q, \Omega \rangle$ is *satisfiable* (or *consistent*)⁵ if there exists a probability function w on L that satisfies $\langle P, Q, \Omega \rangle$. We extend the notion of satisfiability for sets of rules in $\mathcal{R}_{\mathcal{L}}$ in the trivial way.

For the next proposition let $\Phi \subseteq \mathcal{R}_{\mathcal{L}}^*$.

Proposition 3. Φ is satisfiable if and only if it has a model.

Proof. That Φ has a model if and only if it is satisfied by a rational probability function is clear. On the other hand, one can prove that if a probability function satisfies Φ then there exists a rational probability function that satisfies Φ by an argument similar to that of the proof of Proposition 2 by considering the corresponding system of linear inequalities with variables the *worlds* in \mathcal{L} .

Proposition 3 implies that Φ_{CB} can be regarded, for consistency-checking purposes, as a knowledge base formalised in propositional probabilistic logic (or *PPL*).

For the last definition of this section let $\Phi \subseteq \mathcal{R}_{\mathcal{L}}$.

Definition 5. We say that Φ is a *minimal unsatisfiable set* (or *minimal inconsistent set*) if it is not satisfiable and, for all $\Phi^* \subset \Phi$, Φ^* is satisfiable.

4 Checking satisfiability of CADIAG-2

Our main goals are to solve the PSAT problem (Probabilistic SATisfiability) for Φ_{CB} and, provided Φ_{CB} is unsatisfiable (which turns out to be the case), to figure out the minimal sets of conflicting rules (see Definition 5). Despite the fact that Φ_{CB} is formalised as a propositional knowledge base and that there exist several propositional PSAT solvers, we use *Pronto* (see [12]), our probabilistic Description Logic reasoner, for several reasons.

The first reason is that, unlike propositional solvers, it treats classical (i.e., certain) and probabilistic knowledge separately and scales perfectly with respect

⁵ We use both terms indistinctively throughout.

to the amount of the former. Φ_{CB} contains many classical formulas (for example, the number of rules in $\Phi_{D|D}$ is over 200) and so, given the scalability limits of PPL solvers (about 1000 formulas), they are likely to be unable to handle a sufficient number of uncertain *symptom-disease* rules in addition to $\Phi_{D|D}$ and (a fragment of) $\Phi_{S|S}$.

The second reason is that *Pronto* has *pinpointing* capabilities for finding all minimal sets of conflicting formulas in an unsatisfiable knowledge base. This feature is critical in the context of this work given the size of CADIAG-2 and, as we will see shortly, the number of potentially overlapping inconsistencies. It must be noted that finding all minimal inconsistencies is by no means a trivial extension of the PSAT algorithm (for example, its naive implementation using a PSAT solver as a black-box reasoner is not practical due to the hardness of PSAT). Even with our implementation, the number of conflicts in most fragments results in a significant slowdown of reasoning.

Finally, we are interested in evaluating our algorithms (see the next section) on a large and naturally occurring knowledge base such as Φ_{CB} .

4.1 Algorithms

In this section we briefly sketch the PSAT and conflict-finding algorithms implemented in *Pronto* within the frame of classical propositional logic (since for Φ_{CB} a formalisation in terms of the probabilistic DL language is not necessary).

Probabilistic satisfiability algorithm. For the sake of clarity and brevity we will consider the case of PSAT for sets of probabilistic conditional statements of the form $\langle \top, \phi, \eta \rangle$ on L (a finite propositional language), with $\phi \in SL$, where ' \top ' abbreviates 'always true' (a tautology in SL) and η represents the probability assigned to it (i.e., all probabilistic statements considered are unconditional and assigned point-valued probabilities. It is straightforward, but technically awkward and space consuming, to generalise the procedure to handle conditional interval statements, see [12]).

We say that a collection of probabilistic conditional statements of such form, say Φ , is satisfiable if and only if the objective value of the following linear program is equal to 1:

$$\begin{aligned} \max \quad & \sum_{s \in W} x_s \\ \text{s.t.} \quad & \sum_{s \in W_\phi} x_s = \eta \times \sum_{s \in W} x_s, \text{ for each } \langle \top, \phi, \eta \rangle \in \Phi \\ & \sum_{s \in W} x_s \leq 1 \text{ and all } x_s \geq 0 \end{aligned} \quad (2)$$

where x_s is the assignment to the possible world $s \in W$.

Let A denote the matrix of linear coefficients in (2). At every step of the simplex algorithm, A is represented as a combination (B, N) where B and N

are the submatrices of the *basic* and *non-basic* variables, respectively. Values of non-basic variables are fixed to zero and the solver proceeds by replacing one basic variable (i.e., column in A) by a non-basic one until the optimal solution is found. As the size of A is exponential in the size of our language L , one should determine the entering column without representing A explicitly. This is done using the column generation technique in which entering columns are computed by optimizing a subproblem (sometimes referred to as the *pricing-out problem*). Observe that the above system of linear inequalities always admits a solution, e.g. all $x_s = 0$ even if Φ is unsatisfiable, which facilitates the column generation process. Note, however, that the actual linear programs solved in Pronto are considerably more involved, in particular, they include stabilization variables to improve convergence.

The critical step is to formulate linear constraints for the pricing-out problem such that every solution (a column) corresponds to a possible world in W . In the propositional case this can be done by employing a well-known formulation of SAT as a mixed-integer linear program [13]. In the case of an expressive language, such as Description Logic [14], there appears to be no easy way of determining a set of constraints H for the pricing-out problem such that its set of solutions is in one-to-one correspondence with W . *Pronto* implements a novel *hybrid* procedure to compute H iteratively via interaction with a DL reasoner.

The main idea of the algorithm is that every column produced as a solution to the pricing-out problem is converted to a DL concept expression which is then checked for satisfiability by the DL reasoner. If the expression is satisfiable, it means that the column corresponds to a possible world (W in our context) and can be added to (2). Otherwise, the justifications of unsatisfiability (see [15]) are converted into linear constraints and added to the pricing-out problem, which is then re-optimized. Finally, either an entering column is found or the pricing-out problem becomes infeasible, which implies that the system (2) is optimal.

A detailed description of the PSAT algorithm is beyond the scope of this paper and is left as the core of a future paper.

Conflict finding algorithm. A satisfiability algorithm is not sufficient for a comprehensive analysis of an inconsistent knowledge base. Typically users need to identify those fragments of the knowledge base which cause the inconsistency in order to repair them. Such fragments are also required to be minimal so that the user can choose a repair strategy with minimal impact on the rest of the knowledge base.

We are interested in determining the minimal unsatisfiable subsets of a certain collection of probabilistic conditional statements in L , say Φ .

We apply the classical approach to finding minimal unsatisfiable sets based on hitting sets which dates back to Reiter (see [16]). Reiter's hitting set tree (HST) algorithm requires, as a subroutine, a satisfiability procedure which can extract one minimal unsatisfiable set from the knowledge base. It then systematically removes each axiom from that minimal unsatisfiable set and applies the satisfiability procedure again to generate a new minimal unsatisfiable set. By

being systematic in the “repairs”, the procedure finds all minimal unsatisfiable sets in the knowledge base. We reduce the problem of finding a single minimal unsatisfiable subset of Φ to the problem of finding a minimal infeasible subset of inequalities in the corresponding linear system of the form (2) above. Such subsets are known as *irreducible infeasible systems* (IIS) in the LP literature [17]. However, given that the system (2) is never represented in its full version the application of the Ryan and Parker’s algorithm is far from straightforward.⁶ If the optimal value of the system (2) is less than 1 then some inequalities have non-zero dual values. Such inequalities correspond to conflicting constraints in Φ but are not guaranteed to be minimal (though is typically quite small and close to the minimal set). We then do a brute force trial and error search to remove all superfluous constraints.

4.2 Decomposition of CADIAG-2

To our knowledge, none of the existing probabilistic solvers can solve PSAT for Φ_{CB} taken as a whole within reasonable amount of time (see Table 1 for a precise account of the size of Φ_{CB}). However, Φ_{CB} has a certain structure that allows splitting it into fragments that can be examined independently. A crucial property of our probabilistic formalisation of CADIAG-2 is that Φ_{CB} is satisfiable if and only if all of the fragments are individually satisfiable, as we show below.

Table 1. Characteristics of CADIAG-2

Number of distinct symptoms	1761
Number of distinct disease	341
Number of symptom-symptom rules (size of $\Phi_{S S}$)	720
Number of disease-disease rules (size of $\Phi_{D D}$)	218
Number of symptom-disease rules (size of $\Phi_{D S}$)	17573

We can regard Φ_{CB} as a directed graph where the nodes are the medical entities in \mathcal{L} and the edges are given by the rules in Φ_{CB} (i.e., a rule of the form $\langle P, Q, \eta \rangle$ in Φ_{CB} would correspond to an edge directed from P to Q).

Let $P \in \mathcal{L}$. We denote by $\Phi_P \subseteq \Phi_{CB}$ the set of rules that yield a directed edge in a path from P to any other medical entity in \mathcal{L} or a directed edge in a path from any medical entity in \mathcal{L} to P .

⁶ We leave it for future research to investigate how multiple irreducible infeasible systems can be generated at once when the linear system is constructed through column generation.

For the next two results let us consider two medical entities $P_1, P_2 \in S$ and assume that there is no path from P_1 to P_2 in Φ_{CB} or vice versa and that there is no medical entity $P \in \mathcal{L}$ from which there exists a path both to P_1 and P_2 .

Proposition 4. *If Φ_{P_1} and Φ_{P_2} are satisfiable then $\Phi_{P_1} \cup \Phi_{P_2}$ is satisfiable.*

Proof. Let \mathcal{I}_1 and \mathcal{I}_2 be interpretations that satisfy Φ_{P_1} and Φ_{P_2} respectively. We can assume without loss of generality that $D^{\mathcal{I}_1} \cap D^{\mathcal{I}_2} = \emptyset$. We can construct an interpretation \mathcal{I} from \mathcal{I}_1 and \mathcal{I}_2 that satisfies $\Phi_{P_1} \cup \Phi_{P_2}$ in a pretty trivial way by setting $D^{\mathcal{I}} = D^{\mathcal{I}_1} \cup D^{\mathcal{I}_2}$ and $V^{\mathcal{I}}(P, a) = 1$ if and only if $V^{\mathcal{I}_1}(P, a) = 1$ or $V^{\mathcal{I}_2}(P, a) = 1$ and 0 otherwise, for all $(P, a) \in \mathcal{L} \times D^{\mathcal{I}}$. It can be easily seen that \mathcal{I} thus defined satisfies $\Phi_{P_1} \cup \Phi_{P_2}$.

Corollary 2. *If Φ is a minimal unsatisfiable set of rules in $\Phi_{P_1} \cup \Phi_{P_2}$ then Φ needs to be contained in Φ_{P_1} or Φ_{P_2} .*

Proof. It follows trivially from Proposition 4. Notice that if Φ were a minimal inconsistent set in $\Phi_{P_1} \cup \Phi_{P_2}$ and that it were neither contained in Φ_{P_1} nor in Φ_{P_2} then we could define satisfiable subsets from Φ_{P_1} and Φ_{P_2} of the same structure (possibly by removing rules from other minimal inconsistent subsets of Φ_{P_1} and Φ_{P_2} , but none from Φ).

Although trivial, it is worth mentioning that the previous propositions also hold for any alternative interpretation of the rules in terms of probabilistic intervals (i.e., by taking certain real intervals in place of precise probabilities).

Φ_{CB} has the following properties which will enable us to decompose it into a set of fragments:

- P1** All formulas contain only atomic medical entities (i.e., entities in \mathcal{L}).
- P2** All probabilistic formulas in $\Phi_{D|S}$ condition only on symptoms (uncertain rules are unidirectional).
- P3** The graph of $\Phi_{S|S}$ contains many disconnected subgraphs.

We split Φ_{CB} into a set of fragments of the form Φ_P , where $P \in S$ is a symptom such that there is no rule in Φ_{CB} of the form $\langle Q, P, \eta \rangle$. For simplicity we include the entire $\Phi_{D|D}$ in each fragment since it is decomposable to a much less extent than $\Phi_{S|S}$. The largest fragments have around 200 probabilistic formulas that normally relate two or three connected symptoms to diseases.

Corollary 2 guarantees that all minimal unsatisfiable sets of formulas in Φ_{CB} can be found by computing such sets for each fragment. Thus, our methodology is simply a systematic analysis of the fragments of the form Φ_P , which involves a PSAT test and, if the fragment is unsatisfiable, the computation of all minimal unsatisfiable sets in it (see the algorithms in the two previous subsections).

5 Results

We present here results concerning the consistency check of Φ_{CB} when considering a slightly relaxed interpretation of Φ_{CB} by replacing each rule of type *symptom-disease* of the form $\langle P, Q, \eta \rangle \in \Phi_{CB}$, for some $P, Q \in \mathcal{L}$ and $\eta \in (0, 1) \cap \mathbb{Q}$, by $\langle P, Q, \Omega_\eta \rangle$, with

$$\Omega_\eta = [\eta - 0.01, \eta + 0.01] = [\eta^-, \eta^+].^7$$

We have opted for checking consistency of this slightly relaxed interpretation of the rules in Φ_{CB} against a precise interpretation (i.e., the standard interpretation with precise values) because of time constraints. The implementation of our algorithms for the relaxed interpretation of Φ_{CB} completes the task of finding all minimal unsatisfiable subsets in a reasonable amount of time (around one hour). It is a well-known fact in model-diagnosis theory that computing *all* minimal unsatisfiable subsets of a certain knowledge base requires a number of satisfiability tests (in our case, PSAT tests) that is exponential in the number of unsatisfiable subsets. Our relaxed interpretation of Φ_{CB} already contains a high number of unsatisfiable sets (as we will just see) and a precise interpretation adds up more. Furthermore, some of the unsatisfiable sets that are present in the precise interpretation and not in our relaxed one are relatively large (some contain 7 rules) and do not overlap with other unsatisfiable sets. Such facts bring the algorithm's running time closer to its worst case.

An example of a type of minimal unsatisfiable set detected under a precise interpretation of the rules but not under our relaxed version is the one that follows:

$$\begin{aligned} &\langle P_1, Q_1, \eta_1 \rangle, \langle P_1, Q_2, \eta_2 \rangle \langle P_2, Q_1, \eta_3 \rangle \langle P_2, Q_3, \eta_4 \rangle, \\ &\langle Q_1, Q_3, 1 \rangle \langle Q_2, Q_3, 1 \rangle \langle P_1, P_2, 1 \rangle, \end{aligned}$$

for $P_1, P_2 \in S$, $Q_1, Q_2, Q_3 \in D$, $\eta_1, \eta_2, \eta_3, \eta_4 \in [0, 1]$, with $\eta_3 = \eta_4$ and $\eta_1 < \eta_2$. Notice that the rules $\langle P_2, Q_1, \eta_3 \rangle$ and $\langle P_2, Q_3, \eta_4 \rangle$ along with $\langle Q_1, Q_3, 1 \rangle$ intuitively claim that the set of patients with symptom P_2 and disease Q_2 coincides with the set of patients with symptom P_2 and disease Q_3 when assuming $\eta_3 = \eta_4$. Under such an assumption the rules $\langle P_1, Q_1, \eta_1 \rangle$ and $\langle P_1, Q_2, \eta_2 \rangle$ along with the remaining classical rules generate an inconsistency whenever $\eta_1 < \eta_2$. Notice also that, for example, for $\eta_3 < \eta_4$ the set would not be unsatisfiable and thus our relaxed interval interpretation would yield this set consistent (assuming $\eta_3, \eta_4 < 1$).

For the sake of simplicity we will adopt the same notation for the rules of type *symptom-disease* of the form $\langle P, Q, \eta \rangle$, with $\eta \in \{0, 1\}$. We will write $\langle P, Q, \Omega_\eta \rangle$, with $\Omega_\eta = [\eta, \eta] = [\eta^-, \eta^+]$.

We list the different types of minimal unsatisfiable sets encountered in Φ_{CB} under this relaxed interpretation of the rules:

⁷ The degrees of confirmation of the rules in Φ_{CB} are all of the form $\frac{k}{100}$, for some $k \in \{0, 1, \dots, 100\} \subset \mathbb{Z}$. Thus Ω_η is well defined.

Type 1. Our first type of minimal unsatisfiable set in Φ_{CB} is given by a collection of rules of the form

$$\langle P, Q_1, \Omega_\eta \rangle, \langle P, Q_2, \Omega_\zeta \rangle, \langle Q_1, Q_2, 1 \rangle,$$

for $P \in S$, $Q_1, Q_2 \in D$, $\eta, \zeta \in [0, 1]$ and $\zeta^+ < \eta^-$.

By $\zeta^+ < \eta^-$ we are intuitively assuming that the number of patients that have both symptom P and disease Q_1 is greater than the number of patients with both symptom P and disease Q_2 , which contradicts $\langle Q_1, Q_2, 1 \rangle$ (i.e., the assumption that all patients that have disease Q_1 have also disease Q_2).

Type 2. Our second type of minimal unsatisfiable set in Φ_{CB} is given by a set of rules of the form

$$\langle P, Q_1, \Omega_\eta \rangle, \langle P, Q_2, \Omega_\zeta \rangle, \langle Q_1, Q_2, 0 \rangle,$$

for $P \in S$, $Q_1, Q_2 \in D$, $\eta, \zeta \in [0, 1]$ and $\eta^- + \zeta^- > 1$.

Notice that the rule $\langle Q_1, Q_2, 0 \rangle$ assumes *disjointness* between Q_1 and Q_2 (intuitively, there cannot be a patient with both disease Q_1 and Q_2), which rules out the possibility of consistency whenever $\eta^- + \zeta^- > 1$.

Type 3. The third type of minimal conflict set in Φ_{CB} is given by a set of the form

$$\langle P_1, Q, \Omega_\eta \rangle, \langle P_2, Q, \Omega_1 \rangle, \langle P_1, P_2, 1 \rangle,$$

for $P_1, P_2 \in S$, $Q \in D$, $\eta \in [0, 1]$ and $\eta^+ < 1$.

Intuitively, the rule $\langle P_1, P_2, 1 \rangle$ says that all patients with symptom P_1 also have symptom P_2 . The rule $\langle P_2, Q, \Omega_1 \rangle$ intuitively says that all patients with symptom P_2 have disease Q . These two facts together imply that patients with symptom P_1 should all have disease Q (i.e., $\eta^+ = 1$).

Type 4 The fourth and last type of minimal unsatisfiable set is given by a collection of rules of the form

$$\langle P, Q_1, \Omega_\eta \rangle, \langle P, Q_2, \Omega_\zeta \rangle, \langle P, Q_3, \Omega_\lambda \rangle, \langle Q_1, Q_3, 1 \rangle, \langle Q_2, Q_3, 1 \rangle, \langle Q_1, Q_2, 0 \rangle,$$

with $P \in S$, $Q_1, Q_2, Q_3 \in D$, $\eta, \zeta, \lambda \in [0, 1]$, $\lambda^+ < \eta^- + \zeta^- \leq 1$ and $\zeta^-, \eta^- \leq \lambda^+$ (to guarantee minimality).

Intuitively, assuming $\langle P, Q_1, \Omega_\eta \rangle$, $\langle P, Q_2, \Omega_\zeta \rangle$ and $\langle Q_1, Q_2, 0 \rangle$, the proportion of patients that, having symptom P , have either disease Q_1 or Q_2 is at least $\eta^- + \zeta^-$. On the other hand, assuming $\langle Q_1, Q_3, 1 \rangle$ and $\langle Q_2, Q_3, 1 \rangle$, we have that all patients with disease either Q_1 or Q_2 have also disease Q_3 . Thus, under such assumptions, satisfiability requires that $\lambda^+ \geq \eta^- + \zeta^-$.

A thorough analysis of these types of inconsistencies in connection with the whole knowledge base and with possible repair strategies and in relation to other sets of inconsistencies obtained under alternative interpretations of Φ_{CB} (as briefly pointed above, under the standard interpretation of the rules or when regarding η in $\langle P, Q, \eta \rangle \in \Phi_{CB}$ as a lower-bound threshold) is being done at present.

Table 2. Number of minimal unsatisfiable sets in Φ_{CB} and size (relaxed interpretation)

Type of minimal unsatisfiable set	Amount	Number of rules involved
<i>Type 1</i>	420	3
<i>Type 2</i>	5	3
<i>Type 3</i>	1	3
<i>Type 4</i>	269	6

6 Related work

Consistency-checking methods and algorithms for large-scale databases have long been of relevance in scientific computational research. In relation to expert systems and in particular to CADIAG-2 it is worth referring to [18] as an example of research of this nature. In this paper, a classical first-order logic theorem prover was used to analyze the predecessor of CADIAG-2 (CADIAG-1), which did not contain any uncertain rules, and that helped to detect some inconsistent sets of rules. Consistency-checking in CADIAG-2 by means of formal methods is harder mostly because one has to use an appropriate formalism for representing degrees of confirmation in rules (in particular in *symptom-disease* rules). Very recently the first such attempt (see [11]) was made using a specific fragment of monadic infinite-valued Gödel logic G (denoted by G^\sim) extended with classical, involutive negation.

In [11] a sentence in G^\sim is associated to each rule of the form $\langle P, Q, \eta \rangle$ in Φ_{CB} . It is proved that, for $\langle P, Q, \eta \rangle \in \Phi_{CB}$ and θ the sentence in G^\sim associated to it, the following holds:

- If $\eta \in (0, 1)$ then θ is satisfied by a certain interpretation \mathcal{I} of \mathcal{L} if and only if $Q/\mathcal{I}P \in (0, 1)$.
- If $\eta \in \{0, 1\}$ then θ is satisfied by a certain interpretation \mathcal{I} of \mathcal{L} if and only if $Q/\mathcal{I}P = \eta$.

In [11] the problem of checking satisfiability of the set of sentences in G^\sim associated to the rules in Φ_{CB} is proved to be equivalent to the problem of satisfiability in classical first-order logic for such sentences (i.e., equivalent to determining whether there is a classical interpretation of \mathcal{L} that satisfies the sentences associated to the rules).

The relation between our approach and that in [11] is clear in the light of the results stated in Section 3. We will have that a certain collection of rules $\Phi \subseteq \Phi_{CB}$ will be found to be inconsistent according to the approach defined in [11] if and only if there is no probability function w on \mathcal{L} such that $\models_w \Phi^*$, where Φ^* is defined from Φ by replacing η in each rule $\langle P, Q, \eta \rangle \in \Phi$ by the interval $(0, 1)$.

Unlike the approach in [11], our probabilistic formalisation is equisatisfiable with Φ_{CB} (see Proposition 3) and ensures finding all minimal unsatisfiable sets

of rules. In [11] the fragment of compound rules in CADIAG-2 is also considered in addition to the binary fragment (at the expense of a further weakening in the information expressed by the formulas in G^\sim representing these rules). As mentioned earlier, we do not consider CADIAG-2's compound rules in this paper. Certainly, we would require additional efforts to ensure (if possible at all) equisatisfiability and completeness of our decomposition procedure (as given by Proposition 4 for Φ_{CB}) if we consider CADIAG-2's compound rules, although the former would not be necessary if we assumed a probabilistic interpretation of the rules from the outset.

7 Conclusion

While CADIAG-2's knowledge base is, when formalised as a probabilistic logic theory, *highly* unsatisfiable it is unclear what action this calls for. Inconsistency in a knowledge base may capture critical information and maintaining it may be critical to the integrity of the represented knowledge (see [19] for a more comprehensive discussion). An ongoing research challenge is to study and *measure* the inconsistencies in CADIAG-2 in order to understand them better and to determine suitable repair strategies.

Regardless of one's preferred strategy for resolving the conflicts, it is clear that detecting them is critical to a complete understanding of the knowledge base which is challenging when we reach CADIAG-2's size. Even with our (fortuitous) decomposition, the extraction of all conflicts under the standard interpretation of the rules in CADIAG-2 is unfeasible for everyday knowledge base development (we estimate that it will take weeks to extract all conflicts. Of course, if the modellers decide that producing and maintaining a satisfiable version is the right course of action then even several weeks would not be unreasonable as a one-time cost). Subsequent satisfiability checks would go much faster, especially as one can check only the relevant fragment a modeller is working on. This is similar to various proposals from the description logic community for modular ontology development (see [20–22]). As part of our future work we intend to integrate more general modular analysis into our reasoner as an optimization. We intend to investigate whether it is necessary to do this decomposition outside of the solver (that is, by decomposing the input knowledge base before even starting to solve PSAT) in the rather crude manner we currently do, or whether modular analysis can be more tightly integrated with the reasoning process.

We have, as yet, to attempt entailment from CADIAG-2 or any of its fragments. It is not clear yet the extent to which one could generate interesting queries for CADIAG-2's knowledge base (or, more generally, for CADIAG-2 like knowledge bases) once it has been repaired and possibly modified for inferential purposes (see [6]).

We hope that CADIAG-2, or CADIAG-2 like problems, will be taken up by the PSAT solving community. CADIAG-2 is interestingly different in kind, not only in size, from traditional generated problems while its size sets a new base line for scalable PSAT.

References

1. Adlassnig, K., Kolarz, G., Scheithauer, W., Grabner, H.: Approach to a hospital-based application of a medical expert system. *Informatics for Health and Social Care* **11**(3) (1986) 205–223
2. Adlassnig, K., Kolarz, G., Effenberger, W., Grabner, H.: Cadiag: Approaches to computer-assisted medical diagnosis. *Computers in Biology and Medicine* **15** (1985) 315–335
3. Adlassnig, K.: Fuzzy set theory in medical diagnosis. *IEEE Transactions on Systems, Man and Cybernetics* **16**(2) (1986) 260–265
4. Leitich, H., Adlassnig, K., Kolarz, G.: Evaluation of two different models of semiautomatic knowledge acquisition for the medical consultant system CADIAG-2/RHEUMA. *Artificial Intelligence in Medicine* **25** (2002) 215–225
5. Ciabattoni, A., Vetterlein, T.: On the fuzzy (logical) content of Cadiag2. *Fuzzy Sets and Systems* (to appear shortly) (2009)
6. Picado Muño, D.: The (probabilistic) logical content of cadiag2. In: *Proceedings of ICAART 2010*. (2010) 28–35
7. Zadeh, L.: Fuzzy sets. *Information and Control* **8** (1965) 338–353
8. Zadeh, L.: Fuzzy logic and approximate reasoning. *Synthese* **30** (1975) 407–428
9. Klir, G., Folger, T.: *Fuzzy Sets, Uncertainty and Information*. Prentice-Hall International (1988)
10. Zimmermann, H.: *Fuzzy Set Theory and its Applications*. Kluwer Academic Publisher, Boston - Dordrecht - London (1991)
11. Ciabattoni, A., Rusnok, P.: On the classical content of monadic G^{\sim} and its applications to a fuzzy medical expert system. In: *Proceedings of the 12th International Conference on the Principles of Knowledge Representation and Reasoning*. (2010)
12. Klinov, P., Parsia, B.: Pronto: A practical probabilistic description logic reasoner. In: *International Workshop on Uncertainty in Description Logics*. (2010)
13. Hooker, J.N.: Quantitative approach to logical reasoning. *Decision Support Systems* **4** (1988) 45–69
14. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P.F.: *Description Logic Handbook*. Cambridge University Press (2003)
15. Horridge, M., Parsia, B., Sattler, U.: Laconic and precise justifications in OWL. In: *International Semantic Web Conference*. (2008) 323–338
16. Reiter, R.: A theory of diagnosis from first principles. *Artificial Intelligence* **32** (1987) 57–95
17. Parker, M., Ryan, J.: Finding the minimum weight IIS cover of an infeasible system of linear inequalities. *Ann. Math. Artif. Intell.* **17**(1-2) (1996) 107–126
18. Moser, W., Adlassnig, K.: Consistency checking of binary categorical relationships in a medical knowledge bases. *Artificial Intelligence in Medicine* **8** (1992) 389–407
19. Gabbay, D.M., Hunter, A.: Making inconsistency respectable: a logical framework for inconsistency in reasoning. In: *FAIR*. (1991) 19–32
20. Sattler, U., Schneider, T., Zakharyashev, M.: Which kind of module should I extract? In: *Grau, B.C., Horrocks, I., Motik, B., Sattler, U., eds.: Description Logics. Volume 477 of CEUR Workshop Proceedings., CEUR-WS.org* (2009)
21. Cuenca Grau, B., Horrocks, I., Kazakov, Y., Sattler, U.: Extracting modules from ontologies: A logic-based approach. In: *Modular Ontologies*. (2009) 159–186
22. Cuenca Grau, B., Parsia, B., Sirin, E.: Ontology integration using ϵ -connections. In: *Modular Ontologies*. (2009) 293–320