Noname manuscript No.
(will be inserted by the editor)

A probabilistic interpretation of the medical expert system CADIAG-2

David Picado Muiño

Received: date / Accepted: date

Abstract CADIAG-2 is a well known expert system aimed at providing support for medical diagnose in the field of internal medicine. CADIAG-2 consists of a knowledge base in the form of a set of *IF-THEN* rules that relate distinct medical entities, in this paper interpreted as *conditional probabilistic statements*, and an inference engine constructed upon methods of *fuzzy set theory*. The aim underlying this paper is the understanding of the logical structure of the inference in CADIAG-2. To that purpose we provide a (probabilistic) logical formalisation of the inference of the system and check its adequacy with probabilistic logic.

Keywords Knowledge-based systems · rule-based expert systems · fuzzy expert systems · probabilistic inference · CADIAG-2

1 Introduction

CADIAG-2 (Computer Assisted DIAGnosis) is a well known rule-based expert system that aims 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 (Adlassnig et al., 1986), (Adlassnig et al., 1985), (Adlassnig, 1986) or (Leitich et al., 2002) 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 is based on methods of approximate reasoning in *fuzzy set theory*, in the sense of (Zadeh, 1965) and (Zadeh, 1975). In fact CADIAG-2 is presented in some monographs as an example of a *fuzzy expert system*, for example in (Klir and Folger, 1988) or (Zimmermann, 1991). The knowledge base consists of a set of *IF-THEN* rules, also known as *production* rules in the literature, 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 number of rules in the knowledge base of CADIAG-2 is approximately 40.000. The vast majority of them are binary (i.e., they relate single medical entities) and only such rules are considered in this paper. The rules in CADIAG-2 are defined along with a certain *degree of confirmation* which intuitively expresses the *degree* to which the antecedent *confirms* the consequent. For example,

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

We can identify such degrees of confirmation with probabilities and the rules themselves with conditional probabilistic statements. In (Adlassnig, 1986) it is stated that degrees of confirmation can be interpreted as *frequencies* and, more generally, as probabilities. An interpretation in terms of *degrees of belief* of the doctor (or doctors) on the truth of the consequent given that the antecedent of the rule holds is also possible. This fact motivates a probabilistic interpretation of the inference in CADIAG-2 and leads to the primary aim of this paper: formalise the inference in CADIAG-2 on probabilistic grounds and check its adequacy with probability logic –see (Halpern, 2003)– or, more generally, with probability theory. We shall not expect big surprises in

Institut für Diskrete Mathematik und Geometrie
Wiedner Hauptstrasse 8
1040 Vienna, Austria Tel.: +43-1-58801-10427
Fax: +43-1-58801-910427
E-mail: david.muino@tuwien.ac.at

this respect. The inference mechanism in CADIAG-2 proceeds in a *compositional* way and thus it is bound to be probabilistically unsound (as will be clarified later). This was soon observed in earlier studies concerning the celebrated expert system MYCIN –see (Buchanan and Shortliffe, 1984) or (Shortliffe, 1976) for a description of MYCIN and (Hajek, 1988), (Hajek, 1989), (Hajek and Valdés, 1994), (Heckerman, 1986), (Valdés, 1992) for probabilistic approaches to it–. How far the inference mechanism of CADIAG-2 is from probabilistic soundness remains to be seen though.

It is worth mentioning here that, although the interest among theoretical AI researchers in rule-based expert systems seems to be lesser today than some years ago, rule-based expert systems are very popular among AI engineers. Many CADIAG-2-like systems are in use and more are being built for future implementation. It is mainly for this reason that we believe that further analysis and understanding of CADIAG-2-like systems is of relevance.

This paper is in some way a continuation of (Ciabattoni and Vetterlein, 2009). In (Ciabattoni and Vetterlein, 2009) the inference in CADIAG-2 is formalised by means of a logical calculus, $CadL$, and compared to *t-norm-based* formalisms –see (Hajek, 1998)–. It is shown that $CadL$ does not respond to any t-norm-based (or to any fragment of a t-norm-based) logic. As far as we know (Ciabattoni and Vetterlein, 2009) constitutes the first attempt at formalising and understanding CADIAG-2 in a logical way. The present paper is the second.

This paper is structured as follows. In Section 2 we give some basic definitions and introduce most of the notation used later in the other sections. In Section 3 the inference process in CADIAG-2 is described when restricted to the binary rules of its knowledge base. In Section 4 the formal system $CadPL$ is defined and analysed in the light of probability logic. $CadPL$ is a formalisation of the inference mechanism of CADIAG-2 based on a probabilistic interpretation of it.

2 PRELIMINARY DEFINITIONS AND NOTATION

Throughout we will be working with a finite propositional language $L = \{p_1, ..., p_n\}$, for some $n \in \mathbb{N}$. We will denote by SL its closure under classical connectives. Within the context of CADIAG-2 the language L will represent the set of medical entities occurring in the system. By $S \subset L$ we will denote the set of symptoms, findings, signs and test results (*symptoms* for short) and by $D \subset L$ the set of diseases and therapies (*diseases* for short).

Let us set $L_{Lit} = \{p\ , \neg p \mid p \in L\} \subset SL$, the set of literals of the language L.

Consider $\Delta = \{\phi_1, ..., \phi_k\} \subseteq SL$, for some $k \in \mathbb{N}$. We will denote by $\bigwedge \Delta$ the sentence $\phi_1 \wedge ... \wedge \phi_k$.

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

– If $\models \theta$ then $\omega(\theta) = 1$.
– If $\models \neg(\theta \wedge \phi)$ then $\omega(\theta \vee \phi) = \omega(\theta) + \omega(\phi)$.[1]

We define probability on conditional events in the conventional way. For ω a probability function on L and $\phi, \theta \in SL$,

$$\omega(\phi|\theta) = \frac{\omega(\phi \wedge \theta)}{\omega(\theta)}.$$

Notice that if $\omega(\theta) = 0$ then $\omega(\phi|\theta)$ is not defined.

We are interested in knowledge bases (CADIAG-2-like knowledge bases) that consist of finite collections of *IF-THEN* rules. We will identify such rules with triples of the form $\langle \theta, \phi, \eta \rangle$, where $\theta \in SL$ is the antecedent (or, more in keeping with probabilistic terminology, the conditioning event), $\phi \in SL$ is the consequent (or the uncertain event) and $\eta \in [0,1]$ is the degree to which θ confirms ϕ, which we will interpret as the conditional probability of ϕ given θ.

Let $\mathcal{K} = \{\langle \theta, \phi, \eta \rangle \mid \theta, \phi \in SL,\ \ \eta \in [0,1]\}$ be the set of all such triples in the language L.

For the next two definitions let $\Phi \subset \mathcal{K}$.

Definition 2 We say that a probability function ω on L *(probabilistically) satisfies* Φ if, for all $\langle \theta, \phi, \eta \rangle \in \Phi$, we have that $\omega(\theta) > 0$ and $\omega(\phi|\theta) = \eta$.

Definition 3 We say that Φ is *satisfiable* (or consistent) if there exists a probability function ω on L that satisfies Φ.[2]

Next we define a partial ordering relation relevant in the context of the inference mechanism of CADIAG-2.

Definition 4 Let $\preceq$ be the partial ordering relation on $[0,1]$ defined as follows: for $a, b \in [0,1]$, $a \preceq b$ if and only if $0 < a \leq b$ or $0 \leq a < 1$ and $b = 0$.

We define the strict partial ordering $\prec$ from $\preceq$ in the conventional way.

As we will see later, the definition of the ordering $\preceq$ responds to the use of both 0 and 1 as maximal values in

[1] Here (and throughout) $\models$ represents classical entailment.
[2] We will be using the terms 'satisfiable' and 'consistent' indistinguishably.

CADIAG-2. The value 0 denotes certainty in the non-occurrence of an event or falsity of a statement and the value 1 denotes certainty in its occurrence or its truth.

For the next definition let

$$\mathbb{D} = [0,1] \times [0,1] - \{(0,1),(1,0)\}.$$

Definition 5 The function $\max^* : \mathbb{D} \longrightarrow \mathbb{R}$ is defined as follows, for $(a,b) \in \mathbb{D}$:

$$\max^*(a,b) = \begin{cases} a & \text{if } b \prec a \\ b & \text{otherwise} \end{cases}$$

The definition of $\max^*$ is extended to more than two arguments in its trivial way.

3 THE INFERENCE IN CADIAG-2

In this section we briefly describe the inference mechanism in CADIAG-2 when restricted to the set of binary rules, which from now we will denote by Φ_{CB}. A more detailed description can be found in (Ciabattoni and Vetterlein, 2009).

We have three different types of rules in Φ_{CB} according to the typology defined in (Ciabattoni and Vetterlein, 2009):

– Type *confirming to the degree d* (**c**). A rule of this type is formalised in our settings as $\langle \phi, p, \eta \rangle$ for $p \in L$, $\phi \in L_{Lit}$ and $\eta \in (0,1]$. The one that follows is an example of a rule of this type in Φ_{CB}:
  IF suspicion of pancreatic tumor by computerized tomography
  THEN pancreatic cancer
  with degree of confirmation 0.85
– Type *mutually exclusive* (**me**). A rule of type **me** is formalised as $\langle q, \neg p, 1 \rangle$, for $p, q \in L$. We give an example of a rule of this type:
  IF positive rheumatoid factor
  THEN not seronegative rheumatoid arthritis
  with degree of confirmation 1
– Type *always occurring* (**ao**). A rule of type **ao** is formalised as $\langle \neg q, \neg p, 1 \rangle$, for $p, q \in L$. What follows is an example of a rule of this type in Φ_{CB}:
  IF not chorea minor
  THEN not polymyositis
  with degree of confirmation 1

The types **me** and **ao** are *classical* in the sense that the degree of confirmation for the rules of these types is 1 and that the antecedent of such rules (or evidence in our settings) needs to be confirmed (it needs to be given value 1 by the system, see below) in order for these rules to be triggered by the inference engine.

We will denote by Φ^c, Φ^{me} and Φ^{ao} the subsets of Φ_{CB} containing the rules of types **c**, **me** and **ao** respectively.

The inference engine in CADIAG-2 gets started with a set of symptoms and diseases occurring in Φ_{CB} present in the patient. Let $\Gamma \subset L$ be the set of such medical entities.

CADIAG-2 starts with an assignment v_0 on Γ that gives a value in the interval $[0,1]$ to each entity in Γ. The values assigned by v_0 are, according to most of the literature on CADIAG-2, intended to stand for membership degrees in the context of fuzzy set theory and represent the degrees to which such entities are present in the patient –see (Adlassnig et al., 1986),(Adlassnig et al., 1985) or (Adlassnig, 1986) for more on the intended interpretation of v_0–. However, other interpretations could also be suitable, at least to some extent. In fact, when defining the system *CadPL* in the next section, the interpretation to which we will commit will be probabilistic.

The assignment v_0 is defined for the negation of the entities in Γ and logical equivalents according to the following rule:

If $v_0(\theta) = \eta$ then $v_0(\neg\theta) = 1 - \eta$, for $\theta \in SL$ and $\eta \in [0,1]$.[3]

After the initial assignment the inference rules in Φ_{CB} come into play. All the rules triggered by the sentences in Γ and its negations are used during the inference process. At each step in the inference process a rule is applied (that is done, in principle, in no particular order).

At the first step in the inference a rule of the form $\langle \theta, \phi, \eta \rangle$ in Φ_{CB} is triggered, with $\eta \in [0,1]$ and $\theta, \phi \in L_{Lit}$. In order for that to happen θ or its negation needs to be in Γ and the value $v_0(\theta)$ has to be strictly positive if $\langle \theta, \phi, \eta \rangle$ is a rule in Φ^c and equal to 1 if the rule is in Φ^{me} or Φ^{ao}. The application of the rule $\langle \theta, \phi, \eta \rangle$ generates a new assignment, v_1, on $\{\phi\}$. The value assigned to ϕ by it is calculated as the minimum between η and $v_0(\theta)$ and the value assigned to $\neg\phi$ and logical equivalents (if necessary for the inference) is calculated from $v_1(\phi)$ as mentioned above for v_0.

At the n^{th} step in the inference process a new rule of the form $\langle \theta, \phi, \eta \rangle$ in Φ_{CB} will be triggered, for $\eta \in [0,1]$ and $\theta, \phi \in L_{Lit}$. At the n^{th} step the application of this new rule will generate a new assignment on $\{\phi\}$ that will give ϕ the minimum between η and the value of θ considered for triggering the rule at this step in the inference. The value of θ needs to be strictly positive if the rule is of the type **c** and equal to 1 if it is of one

[3] Notice that such a rule is compatible with a probabilistic interpretation of v_0.

of the types **me** or **ao**. If the strictly positive values generated for θ before the n^{th} step are multiple then the inference mechanism in CADIAG-2 will call the rule $\langle \theta, \phi, \eta \rangle$ again in further steps, if it has not done so previously, until all the values for θ have been used and all the possible values for ϕ generated. The assignment v_n is defined for $\neg\phi$ as mentioned above for v_0.

The inference process goes on until all the rules triggered by all the sentences in Γ and its negations have been used and all the possible assignments for the sentences involved in the inference have been generated. CADIAG-2 yields as an outcome of the inference the set of medical entities in L occurring in the rules triggered by the initial evidence in Γ along with the *maximal value* (with respect to the ordering $\preceq$ defined in Section 2) assigned to them during the inference. If a sentence is assigned both value 0 and 1 along the inference process the system generates an error message.

It is worth mentioning that the original inference process in CADIAG-2 works in a slightly different way. The update in the value of the distinct sentences involved in the inference is done as soon as two different values for the same sentence are produced by the system. The value chosen in the update for atomic sentences in L is the *maximal* one (with respect to the ordering $\preceq$). Notice though that this feature has a highly undesirable result (unless further restrictions on the rules or on the order in which the rules are applied are imposed), which is that the outcome of a run of the inference mechanism can depend on the order in which the rules are applied. Such a drawback is easily avoided by assuming that the update is only done at the end of the process. There are other several undesirable features about the inference in CADIAG-2, most of them related to the maximal value 0 and negated propositions. Maybe the most evident concerning the maximal value 0 is that a medical entity that at some step along the inference process is assigned value 0 (that is to say, it is considered false with certainty or impossible) triggers any rules in which it occurs as the conditioning event if any other value other than 0 is assigned to it along the inference process. For a deeper analysis of such aspects of the inference process in CADIAG-2 see (Ciabattoni and Vetterlein, 2009).

We represent sentences together with the assignments generated for them at each step in the inference by pairs in $SL \times [0, 1]$ along with a subscript indicating the step in the process at which such pairs have been generated. As mentioned above, a step in the inference process is given by the application of a rule in Φ_{CB} and the new assignments that it generates for the sentences involved in the rule.

Let $p \in L$ and $\eta \in [0, 1]$ be the highest assignment to p in a run of the inference mechanism in CADIAG-2. We will use the subscript max* on the pair (p, η) –that is to say, $(p, \eta)_{\max^*}$– to denote that η is the *maximal* value assigned during the inference process for p (with respect to the ordering $\preceq$).

4 THE FORMAL SYSTEM *CADPL*

As seen in the previous section, the inference mechanism in CADIAG-2 gets started by assigning to the initial symptoms and diseases present in the patient (Γ) values in the interval $[0, 1]$ (the initial assignment v_0 to the entities in Γ and their corresponding negations). As mentioned earlier, such values stand in principle for fuzzy membership within the context of fuzzy set theory and are motivated by the *vague* nature of some medical entities that occur in CADIAG-2. In this paper though the interpretation that we will give to the assignment v_0 will be probabilistic.

Let us consider the medical entity '*reduced glucose in serum*'. Let us assume that the value assigned by the evaluation system in CADIAG-2 (i.e., v_0) to the statement '*Patient A has reduced glucose in serum*' out of the evidence given by the corresponding measurement of the amount of glucose in Patient A is η, for some $\eta \in [0, 1]$. As an example, we could interpret such value as the *degree of belief* that a medical doctor has in the truth of the statement given the evidence. As such η could be interpreted as a probability. The probabilistic interpretation is certainly favoured by the *discretization* applied to medical concepts in CADIAG-2 (for example, the concept '*glucose in serum*' generates five distinct medical entities in CADIAG-2: '*highly reduced glucose in serum*', '*reduced glucose in serum*', '*normal glucose in serum*', '*elevated glucose in serum*' and '*highly elevated glucose in serum*'). Notice that such an interpretation places us within the subjective probabilistic frame and thus, for the sake of coherence, the knowledge base Φ_{CB} should also be interpreted subjectively. Other interpretations are also possible though. For example, one could regard such values as the ratio given by the number of doctors that agree on the truth of the statement out of all the doctors involved in the assessment. In order to accommodate such values into a coherent probabilistic frame along with the statements in Φ_{CB} one could justify them as being *subjective* probabilities assessed by a *group* of experts –see (Genest and Zidek, 1986) or (Osherson and Vardi, 2006) for an analysis and justification of such concept–.

Let $p \in L$ represent a medical entity present in the patient and assume that $\eta \in [0, 1]$ is the initial value assigned to it by v_0 at the start of a run of the inference mechanism in CADIAG-2. We can formalise this by means of a probabilistic conditional statement represented by a triple of the form $\langle \kappa, p, \eta \rangle$, where $\kappa \in SL$ is the evidence that supports the presence of p in the patient.

Next we will define the formal system $CadPL$. Recall that the ultimate goal when doing so is to define a system which represents the inference process in CADIAG-2 when interpreted from a probabilistic point of view. Although the inference in CADIAG-2 can be closely related to probability theory (given the nature of the rules of inference in Φ_{CB}) it is not based on probabilistic methods and so the degree of freedom when choosing the rules of the system $CadPL$ is high. We have chosen the rules by interpreting *in the most natural way* the steps along the inference process within a probabilistic frame. The main idea behind such interpretation consists of the identification of the inference process with the *propagation of evidence* facilitated by the rules in Φ_{CB}. For example, from $\langle \kappa, \phi, \eta \rangle$, where $\kappa \in SL$ is evidence supporting the presence of ϕ in the patient, and $\langle \phi, \theta, \zeta \rangle$ in Φ_{CB} we would infer $\langle \kappa, \theta, \min(\eta, \zeta) \rangle$, where $\min(\eta, \zeta)$ is the value (probability) assigned to θ given the evidence κ. We would have a *propagation* process of this nature for each single piece of evidence. The evidence would then be brought together in CADIAG-2 by what we call the *right conjunction* rule: given two outcomes of the inference process, say $\langle \kappa_1, p, \eta \rangle$ and $\langle \kappa_2, p, \zeta \rangle$, for $p \in L$ and $\kappa_1, \kappa_2 \in SL$, CADIAG-2 combines the evidence given by κ_1 and κ_2 by computing $\langle \kappa_1 \wedge \kappa_2, p, \max^*(\eta, \zeta) \rangle$. The inference rules of $CadPL$ that we next present formalise this interpretation.

A theory $\mathcal{T}$ in $CadPL$ is a finite subset of $\mathcal{K}$. A theory in our settings is not necessarily assumed to be consistent. In fact, Φ_{CB} is not consistent as a probabilistic knowledge base, it contains a large number of conflicts (i.e., *minimal* inconsistent subsets) –see our companion paper (Klinov et al., 2010) for more on the conflicts in Φ_{CB}–. However, the inference engine in CADIAG-2 does not *explode* in the presence of such conflicts.

For what follows let $\mathcal{T} = \Omega \cup \Phi_1 \cup \Phi_2$ be a theory, with

$$\Omega = \{\langle \kappa_1, \phi_1, \eta_1 \rangle, ..., \langle \kappa_m, \phi_m, \eta_m \rangle\},$$

for some $m \in \mathbb{N}$, and Φ_2 a set of triples of the form $\langle \theta, \phi, 1 \rangle$, for $\theta, \phi \in SL$.

Let $\mathcal{C}_\Omega = \{\kappa_1, ..., \kappa_m\}$ and $\Gamma = \{\phi_1, ..., \phi_m\}$. We will assume $\mathcal{C}_\Omega$ to be a consistent set (i.e., there exists a classical valuation that satisfies $\mathcal{C}_\Omega$).

Within the context of CADIAG-2, Γ is intended to represent the set of medical entities present in the patient and Ω the initial assignment (v_0) in the inference process given $\mathcal{C}_\Omega$, the evidence in support of the presence in the patient of the corresponding medical entities in Γ. The set Φ_1 represents the rules of type **c** and Φ_2 the rules of types **me** and **ao**.

The formal system $CadPL$ is defined by the following inference rules:

- Reflexivity rule
  For $\phi \in SL$, $\kappa \in \mathcal{C}_\Omega$ and $\eta \in [0, 1]$,
  $$\frac{\langle \kappa, \phi, \eta \rangle \ \in \Omega}{\mathcal{T} \vdash \langle \kappa, \phi, \eta \rangle}$$
- Negation rule
  For $\phi, \psi \in SL$ and $\eta \in [0, 1]$,
  $$\frac{\mathcal{T} \vdash \langle \psi, \phi, \eta \rangle}{\mathcal{T} \vdash \langle \psi, \neg\phi, 1 - \eta \rangle}$$
- Equivalence rule
  For $\psi, \phi, \theta \in SL$ and $\eta \in [0, 1]$,
  $$\frac{\psi \equiv \phi \qquad \mathcal{T} \vdash \langle \theta, \phi, \eta \rangle}{\mathcal{T} \vdash \langle \theta, \psi, \eta \rangle}$$
- Minimum rules
  For $\theta, \phi \in SL$, $\kappa \in \mathcal{C}_\Omega$, $\eta \in (0, 1]$ and $\zeta \in [0, 1]$,
  First rule:
  $$\frac{\mathcal{T} \vdash \langle \kappa, \theta, \eta \rangle \qquad \langle \theta, \phi, \zeta \rangle \ \in \Phi_1}{\mathcal{T} \vdash \langle \kappa, \phi, \min(\eta, \zeta) \rangle}$$
  Second rule:
  $$\frac{\mathcal{T} \vdash \langle \kappa, \theta, 1 \rangle \qquad \langle \theta, \phi, 1 \rangle \ \in \Phi_2}{\mathcal{T} \vdash \langle \kappa, \phi, 1 \rangle}$$
- Right conjunction rule
  For $\phi \in SL$, $C_1, C_2 \subseteq \mathcal{C}_\Omega$ and $\eta, \zeta \in [0, 1]$,
  $$\frac{\mathcal{T} \vdash \langle \bigwedge C_1, \phi, \eta \rangle \qquad \mathcal{T} \vdash \langle \bigwedge C_2, \phi, \zeta \rangle}{\mathcal{T} \vdash \langle \bigwedge \{C_1 \cup C_2\}, \phi, \max^*(\eta, \zeta) \rangle}$$
- Exhaustivity rule
  For $\phi \in SL$, $\kappa \in \mathcal{C}_\Omega$, $C \subseteq \mathcal{C}_\Omega$ and $\eta \in [0, 1]$,
  $$\frac{\mathcal{T} \vdash \langle \bigwedge C, \phi, \eta \rangle \qquad \forall \zeta \in [0, 1]\ \mathcal{T} \nvdash \langle \kappa, \phi, \zeta \rangle}{\mathcal{T} \vdash \langle \kappa \wedge \bigwedge C, \phi, \eta \rangle}$$

Notice that the *exhaustivity* rule does not have any bearing on the *decidability* of whether $\langle \kappa, \phi, \zeta \rangle$ is provable from $\mathcal{T}$ or not for $\zeta \in [0, 1]$, $\phi \in SL$ and $\kappa \in \mathcal{C}_\Omega$. The *exhaustivity* rule can only be applied after its provability or non-provability has been decided.

Given a theory $\mathcal{T}$ of $CadPL$ and a triple $\Theta \in \mathcal{K}$, a proof of Θ from $\mathcal{T}$ is defined as a finite sequence of *sequents* of the form

$$\mathcal{T} \vdash \Theta_1, ..., \mathcal{T} \vdash \Theta_n$$

with $\Theta_n = \Theta$ and where, for $i \in \{1, ..., n\}$, each Θ_i in $\mathcal{T} \vdash \Theta_i$ follows from $\mathcal{T}$ by the application of one of the rules above, from Θ_j in a previous sequent (with $j < i$) or from Θ_j, Θ_k in previous sequents (with $j, k < i$) by one of the rules above.

For what follows let Θ be the triple $\langle \theta, \phi, \eta \rangle$, for some $\eta \in [0,1]$ and $\theta, \phi \in SL$.

Definition 6 We say that there exists a *maximal* proof of Θ from $\mathcal{T}$ if there exists a proof of Θ from $\mathcal{T}$ and there is no proof from $\mathcal{T}$ of $\langle \theta, \phi, \zeta \rangle$ with $\eta \prec \zeta$.

Definition 7 We say that Θ follows *maximally* from $\mathcal{T}$ (denoted by $\mathcal{T} \vdash_{CadPL} \Theta$) if there exists a maximal proof of Θ from $\mathcal{T}$.

For the next proposition let $\mathcal{T} = \Omega \cup \Phi_1 \cup \Phi_2$, with

$$\Omega = \{\langle \kappa_1, q_1, \eta_1 \rangle, ..., \langle \kappa_m, q_m, \eta_m \rangle\},$$

$\Phi_1 = \Phi^c$ and $\Phi_2 = \Phi^{me} \cup \Phi^{ao}$, $\mathcal{C}_\Omega = \{\kappa_1, ..., \kappa_m\} \subset SL$ and $\Gamma = \{q_1, ..., q_m\} \subset L$ a set of medical entities occurring in Φ_{CB}.

Proposition 1 *Let $p \in L$ and $\eta \in [0,1]$. We have that*

$$\mathcal{T} \vdash_{CadPL} \langle \bigwedge \mathcal{C}_\Omega, p, \eta \rangle$$

if and only if $(p, \eta)_{\max^*}$ is the outcome of a run of the inference engine of CADIAG-2 on $\mathcal{T}$.

Proof [4] In order to prove the left implication let us consider a run of the inference engine of CADIAG-2 on $\mathcal{T}$. The inference starts from pairs of the form $(q, \eta)_0$ and $(\neg q, 1 - \eta)_0$ for some $\eta \in [0,1]$ for all $q \in \Gamma$. In $CadPL$ a pair of the form $(q, \eta)_0$, for $q \in \Gamma$, corresponds to a sequent of the form $\mathcal{T} \vdash \langle \kappa, q, \eta \rangle$, for $\kappa \in \mathcal{C}_\Omega$. The pair $(\neg q, 1 - \eta)_0$ corresponds to the sequent $\mathcal{T} \vdash \langle \kappa, \neg q, 1 - \eta \rangle$. The former corresponds to an application of the *reflexivity* rule. The latter follows from the first one by an application of the *negation* rule.

Let us assume now that we are at the n^{th} step of the inference process and that a rule of the form $\langle \theta, \phi, \zeta \rangle$ is triggered, for some $\zeta \in [0,1]$ and $\theta, \phi \in SL$. Let us suppose that we have $(\theta, \mu)_{n-t}$, the pair that triggers the rule at the n^{th} step of the process, for $\mu \in (0,1]$ and $t \leq n - 1$. In $CadPL$ that would correspond to a sequent of the form $\mathcal{T} \vdash \langle \kappa, \theta, \mu \rangle$ derived from a previous step in the inference, for $\kappa \in \mathcal{C}_\Omega$. The inference mechanism in CADIAG-2 produces the pairs $(\phi, \min(\zeta, \mu))_n$ and $(\neg\phi, 1 - \min(\zeta, \mu))_n$ which, in the formal system $CadPL$, corresponds to the sequents $\mathcal{T} \vdash \langle \kappa, \phi, \min(\zeta, \mu) \rangle$ and $\mathcal{T} \vdash \langle \kappa, \neg\phi, 1 - \min(\zeta, \mu) \rangle$ respectively, which follow by an application of one of the *minimum* rules and, for the latter, an application of the *negation* rule on the former.

At the end of the process CADIAG-2 generates the pair $(p, \eta)_{\max^*}$ for each sentence $p \in L$ involved in the inference, where η is the maximal value (with respect to the ordering $\preceq$) among those assigned to p along the inference. This maximisation process is achieved in $CadPL$ by means of repeated applications of the *right conjunction* rule. Instances of the *exhaustivity* rule (if necessary) complete the inferential counterpart of CADIAG-2 in $CadPL$.

In order to prove the right implication let us suppose that we have a maximal proof of the form

$$\mathcal{T} \vdash \Theta_1, ..., \mathcal{T} \vdash \Theta_m,$$

where Θ_m is the triple $\langle \bigwedge \mathcal{C}_\Omega, p, \eta \rangle$, for some $\eta \in [0,1]$ and $p \in L$.

The first sequent of the proof needs to respond to an instance of the *reflexivity* rule, $\mathcal{T} \vdash \langle \kappa, q, \eta \rangle$, for some $q \in \Gamma$, $\kappa \in \mathcal{C}_\Omega$ and $\eta \in [0,1]$. The corresponding counterpart of this sequent in CADIAG-2 is the pair $(q, \eta)_0$.

Let us move now to the n^{th} sequent, with $n \leq m$. The n^{th} sequent can be an instance of the *reflexivity* rule, $\mathcal{T} \vdash \langle \kappa, q, \eta \rangle$, for some $q \in \Gamma$, $\eta \in [0,1]$ and $\kappa \in \mathcal{C}_\Omega$. The counterpart for this sequent in CADIAG-2 is the pair $(q, \eta)_0$.

The n^{th} sequent can follow from a previous one in the proof by an instance of the *negation* rule. Let us suppose that the n^{th} sequent is $\mathcal{T} \vdash \langle \theta, \neg\phi, 1 - \eta \rangle$ for some $\eta \in [0,1]$ and $\theta, \phi \in SL$ and that there is a sequent $\mathcal{T} \vdash \Theta_i$, for some $i < n$, of the form $\mathcal{T} \vdash \langle \theta, \phi, \eta \rangle$. The latter corresponds to a pair of the form $(\phi, \eta)_t$ in CADIAG-2 and the former to the pair $(\neg\phi, 1 - \eta)_t$, where t is the step in the inference process at which such pairs have been generated.

The n^{th} sequent can follow from a previous one by an instance of one of the *minimum* rules. Let us assume that the n^{th} sequent is $\mathcal{T} \vdash \langle \kappa, \phi, \min(\eta, \zeta) \rangle$, for some $\phi \in SL$, $\kappa \in \mathcal{C}_\Omega$, $\eta \in [0,1]$ and $\zeta \in (0,1]$, that $\mathcal{T} \vdash \langle \kappa, \psi, \zeta \rangle$ is a previous sequent in the proof and that $\langle \psi, \phi, \eta \rangle$ is a rule in Φ_{CB}. The latter corresponds in CADIAG-2 to the pair $(\psi, \zeta)_t$ and the former to the pair $(\phi, \min(\eta, \zeta))_{t+k}$, where $t, t+k$ indicate the steps at which the pairs have been generated by the inference process.

The n^{th} sequent can follow from previous sequents by an application of the *right conjunction* rule. The counterpart in CADIAG-2 of such an outcome consists of the maximisation process at the end of the inference. Instances of the *exhaustivity* rule are irrelevant to the inference in CADIAG-2.

[4] For the sake of brevity we will deal with sentences as if they were equivalence classes. If anything applies to a sentence of the form $\neg\phi$, with $\phi \in SL$, we also assume that it applies to any logical equivalent of $\neg\phi$ without mentioning it.

This completes the proof.

It is worth commenting on some features of the inference rules of $CadPL$ in connection with probability theory. *Soundness* with respect to probabilistic semantics of the *reflexivity*, *negation*, and *equivalence* rules is clear. Soundness of the second *minimum* rule is also clear. The first *minimum* rule is certainly not sound with respect to such semantics. However, the *right conjunction* rule is not sound and the *exhaustivity* rule assumes some probabilistic *independence* among sentences that may not actually be independent. Overall, $CadPL$ does not score well within probability theory. This is no surprise. The computation of conditional probabilistic statements in a compositional way, as done by CADIAG-2 primarily by means of the min and max* operators, is clearly bound to be probabilistically unsound. One may wonder though what could be done in order to improve the inference on probabilistic grounds from a knowledge base like Φ_{CB}. The answer seems to be 'not much'. Certainly a Φ_{CB}-like knowledge base (i.e., a knowledge base given by some binary probabilistic conditional statements) is not the most convenient for inferential purposes in probability theory for medical applications like CADIAG-2. As is well known, there are other knowledge-base structures better suited for that purpose, Bayesian networks being the most celebrated among them, see (Castillo et al., 1997) or (Pearl, 1988).

In terms of consistency, it is worth noting that $CadPL$ satisfies what we can call *weak consistency* –called *weak soundness* in (Hajek, 1988)–, defined as follows: if there is a maximal proof in $CadPL$ of a statement of the form $\langle \bigwedge \Delta, \phi, 1\rangle$ (or $\langle \bigwedge \Delta\phi, 0\rangle$) from a certain theory $\mathcal{T}$, with $\phi \in SL$ and $\Delta \subseteq SL$ then, if there is a maximal proof in $CadPL$ of a statement of the form $\langle \bigwedge \Delta^*, \phi, \eta\rangle$, with $\Delta \subset \Delta^*$, then $\eta = 1$ (or $\eta = 0$ respectively). That is to say, if $CadPL$ concludes certainty about the occurrence of some event or about the truth or falsity of some sentence then adding new evidence does not alter this certainty. Weak consistency is provided in $CadPL$ and so in the inference mechanism of CADIAG-2 by the operator max* defined over the ordering $\preceq$.

It is also worth noting that one could guarantee consistency (i.e., satisfiability) by relaxing the interpretation of Φ_{CB} and consider η in each triple $\langle \theta, \phi, \eta\rangle \in \Phi_{CB}$ a lower-bound probability threshold rather than a point-valued probabilitiy and by restricting the system to a *positive fragment* of L_{Lit} (i.e., only one of p, $\neg p$ can occur in Φ_{CB}). This way consistency is trivially guaranteed for Φ_{CB} together with any outcomes produced by the system during the inference process.

In terms of soundness there does not seem to be much that one can do in order to improve the inference mechanism for knowledge bases like Φ_{CB}, or at least not much that one can do that does not come at the price of generating probabilistic statements with very low probabilistic bounds (when working with lower-bound probability thresholds), which would make CADIAG-2 potentially useless for practical purposes. There is some room for improvement for some steps in the inference that come by the addition of independence assumptions among some of the medical entities in Φ_{CB}. Under such independence assumptions a *product rule* instead of the *minimum rule* (i.e., composition through the product operator instead of the minimum one) could yield soundness in the context of lower-bound probability thresholds for some inference steps.

5 CONCLUSION

CADIAG-2 is a reasonably well-performing medical expert system (Adlassnig et al., 1986), but how it is so is far from clear. The inference engine of CADIAG-2 was built with methods of approximate reasoning in fuzzy set theory but, as such, it was not based on any logical formalism or theory embedded with a clear semantics. This fact motivated the main aim of this paper, which was no other than the *understanding* of CADIAG-2 in a *logical* way.

The natural interpretation of the inference rules of CADIAG-2 (i.e., probabilistic) placed us upon the attempt of interpreting the inference itself probabilistically. We formalised this interpretation by means of the system $CadPL$, the logical (probabilistic) counterpart of the inference engine of CADIAG-2. The unsoundness of some of the rules of $CadPL$ (and thus of some inference steps in CADIAG-2) and the inconsistency of the calculus (and thus of the inference process in CADIAG-2) was made clear. Apart from these drawbacks, otherwise expected, some other aspects of $CadPL$ were also stressed and analysed. At the end of the paper some possibilities for an improvement of CADIAG-2 in terms of soundness and consistency were also mentioned.

Acknowledgements This research has been supported by the Vienna Science and Technology Fund ($WWTF$), Grant $MA07-016$. I would like to thank A. Ciabattoni, T. Vetterlein and P. Rusnok for useful comments on previous drafts of this paper and for technical support with CADIAG-2.

REFERENCES

[Adlassnig, 1986]Adlassnig, K. (1986). Fuzzy set theory in medical diagnosis. *IEEE Transactions on Systems, Man and Cybernetics*, 16(2):260–265.

[Adlassnig et al., 1985]Adlassnig, K., Kolarz, G., Effenberger, W., and Grabner, H. (1985). Cadiag: Approaches to computer-assisted medical diagnosis. *Computers in Biology and Medicine*, 15:315–335.

[Adlassnig et al., 1986]Adlassnig, K., Kolarz, G., Scheithauer, W., and Grabner, H. (1986). Approach to a hospital-based application of a medical expert system. *Informatics for Health and Social Care*, 11(3):205–223.

[Buchanan and Shortliffe, 1984]Buchanan, B. and Shortliffe, E. (1984). *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Addison-Wesley, Reading, MA.

[Castillo et al., 1997]Castillo, E., Gutiérrez, J., and Hadi, A. (1997). *Expert Systems and Probabilistic Network Models*. Springer-Verlag, New York.

[Ciabattoni and Vetterlein, 2009]Ciabattoni, A. and Vetterlein, T. (2009). On the fuzzy (logical) content of cadiag2. *Fuzzy Sets and Systems (to appear shortly)*.

[Genest and Zidek, 1986]Genest, C. and Zidek, J. (1986). Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 1(1):114–135.

[Hajek, 1988]Hajek, P. (1988). Towards a probabilistic analysis of mycin-like expert systems. In *COMPSTAT 88 proceedings*. Physica-Verlag Heidelberg.

[Hajek, 1989]Hajek, P. (1989). Towards a probabilistic analysis of mycin-like expert systems 2. In *Artificial Intelligence and Information-Control Systems of Robots*. North-Holland.

[Hajek, 1998]Hajek, P. (1998). *Metamathematics of Fuzzy Logic*. Kluwer, Dordrecht.

[Hajek and Valdés, 1994]Hajek, P. and Valdés, J. (1994). An analysis of mycin-like expert systems. *Mathware and Soft Computing*, 1:45–68.

[Halpern, 2003]Halpern, J. (2003). *Reasoning About Uncertainty*. MIT Press.

[Heckerman, 1986]Heckerman, D. (1986). Probabilistic interpretations for mycin's certainty factors. In *Uncertainty in Artificial Intelligence*, pages 167–196. North-Holland.

[Klinov et al., 2010]Klinov, P., Parsia, B., and Picado-Muino, D. (2010). The consistency of the cadiag-2 knowledge base: A probabilistic approach. In *LPAR*.

[Klir and Folger, 1988]Klir, G. and Folger, T. (1988). *Fuzzy Sets, Uncertainty and Information*. Prentice-Hall International.

[Leitich et al., 2002]Leitich, H., Adlassnig, K., and Kolarz, G. (2002). Evaluation of two different models of semiautomatic knowledge acquisition for the medical consultant system cadiag2/rheuma. *Artificial Intelligence in Medicine*, 25:215–225.

[Osherson and Vardi, 2006]Osherson, D. and Vardi, M. (2006). Aggregating disparate estimates of chance. *Games and Economic Behavior*, 56:148–173.

[Pearl, 1988]Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufman, San Mateo, California.

[Shortliffe, 1976]Shortliffe, E. (1976). *Computer-Based Medical Consultations: MYCIN*. Elsevier, New York.

[Valdés, 1992]Valdés, J. (1992). An integrated environment for uncertainty processing using graphical models. In *IPMU 92 proceedings*.

[Zadeh, 1965]Zadeh, L. (1965). Fuzzy sets. *Information and Control*, 8:338–353.

[Zadeh, 1975]Zadeh, L. (1975). Fuzzy logic and approximate reasoning. *Synthese*, 30:407–428.

[Zimmermann, 1991]Zimmermann, H. (1991). *Fuzzy Set Theory and its Applications*. Kluwer Academic Publisher, Boston - Dordrecht - London.