

Bilattice CADIAG-II: Theory and Experimental Results

Paolo Baldi, Agata Ciabattoni, and Klaus-Peter Adlassnig

Abstract CADIAG-II is a functioning experimental fuzzy expert system for computer-assisted differential diagnosis in internal medicine. To overcome the current limitations of the system, we propose an extension based on bilattices. The proposed changes were implemented and reviewed in a retrospective evaluation of 3,131 patients with extended information about patient’s medical history, physical examination, laboratory test results, clinical investigations and – last but not least – clinically confirmed discharge diagnoses.

1 Introduction

1.1 Background

Computer-based support of medical diagnosis and treatment has a long tradition. Early approaches were based on statistical methods such as Fisher’s discriminant analysis to classify symptom patterns into diseased or non-diseased categories [39]. Others used Bayes’ theorem to assign probabilities to the possible presence of dis-

Paolo Baldi
Department of Philosophy, University of Milan, Via Festa del Perdono 7, 20122 Milan, Italy, e-mail: paolo.baldi@unimi.it

Agata Ciabattoni
Institute of Logic and Computation, Vienna University of Technology, Favoritenstrasse 9, 1040 Vienna, Austria, e-mail: agata@logic.at

Klaus-Peter Adlassnig
Section for Artificial Intelligence and Decision Support, Center for Medical Statistics, Informatics, and Intelligent Systems, Medical University of Vienna, Spitalgasse 23, 1090 Vienna, Austria, e-mail: klaus-peter.adlassnig@meduniwien.ac.at
and
Medexter Healthcare, Borschkegasse 7/5, 1090 Vienna, Austria, e-mail: kpa@medexter.com

eases [20, 38]. The first medical expert system was MYCIN [34], whose purpose was to give advice for the diagnosis and the treatment of patients with infectious diseases. Equipped with a well-grounded heuristic rule-based approach to determine diagnostic and therapeutic proposals, MYCIN was extensively tested; its performance was comparable to that of humans [41]. A variety of logical and probabilistic reasoning approaches in medical diagnosis have been compared in [36]. Artificial intelligence methods and systems for medical consultation were discussed, among others, in [11, 25]; an extended threaded bibliography was provided in [30, 40]. CADIAG-I and CADIAG-II were also early approaches to provide differential diagnostic support. Based on logical approaches described in the seminal paper by Ledley and Lusted [26], CADIAG-I was extended in several subsequent versions [4]. CADIAG-II employs fuzzy sets and fuzzy logic, as described in this report, and gave rise to a variety of refined modeling approaches [13–17, 22, 31, 32, 37]. It was extensively tested in various fields of clinical application [3, 6, 27–29]. Recent approaches to clinical decision support for the selection of diagnosis and therapy mainly consist of machine learning and “big” data approaches. Successful applications include image pattern recognition in fields such as radiology [10] and pathology [33]. IBM’s Watson for oncology is one of many recent machine learning system approaches; its aim is to provide recommendations for the treatment of breast cancer. However, its success appears to be limited [35].

1.2 CADIAG Systems

Computer-Assisted DIAGnosis (CADIAG) systems are data-driven, rule-based expert systems for computer-assisted consultation in internal medicine [1, 4, 7, 12]. Their development dates back to the early 1980s at the University of Vienna Medical School (now Medical University of Vienna). The systems provide diagnostic hypotheses as well as confirmed and excluded diagnoses in response to the input of a list of symptoms, signs, laboratory test results, and clinical findings pertaining to a patient. When possible, they also explain the indicated conclusions and propose further useful examinations.

The first system of the family – CADIAG-I – dealt with three-valued logical variables (present, unknown, absent) and IF-THEN relationships between given three-valued input on the one hand and diagnoses on the other. Kleene’s logic provides all the necessary formal definitions, see also [4]. However, the real-world patient’s input (symptoms, interpreted signs, laboratory test results, and clinical findings) is usually inherently (linguistically) vague and necessarily includes borderline cases. Moreover, a large part of the given medical knowledge about definitional, causal, statistical, and heuristic relationships between a patient’s input and described diseases is intrinsically uncertain. Measurements are sometimes imprecise, linguistic categories are characterized by fuzzy borderlines, the co-occurrence of symptoms and diseases is stochastically uncertain, and both medical data and medical knowledge are often incomplete. Therefore, computer systems for medical decision-making

usually cannot generate clinically accurate results when based on formal systems whose objects can only be either absolutely true, absolutely false, or unknown (as in Kleene's logic). The successor system CADIAG-II can process both definite and uncertain information. CADIAG-II is based on fuzzy set theory [43] to deal with linguistic medical terms, and on fuzzy logic to define and process weighted IF-THEN rules [14].

Despite this improvement, CADIAG-II has been criticized for its inability to deal with negative evidence [16], and with rules diminishing the certainty of a particular diagnosis apart from complete exclusion.

1.3 Objective

The aim of the present report is to introduce an extension of CADIAG-II, which includes negative knowledge, and experimentally evaluate the presented proposal. An earlier attempt confined to theory was published in [17]. Here we introduce Bilattice CADIAG-II, an extension of CADIAG-II based on bilattices, and validate the theoretical results by means of a retrospective evaluation in a newly-programmed CADIAG-II implementation. The data of 3,131 patients, including extended medical histories, physical examinations, and laboratory and clinical test results were analyzed with the original CADIAG-II and with Bilattice CADIAG-II. The results were compared with the corresponding clinically confirmed diagnoses at the time of discharge. All underlying real patient data including the corresponding clinical discharge diagnoses originate from a hospital near Vienna. In addition to preserving all the inference results of CADIAG-II, Bilattice CADIAG-II was able to infer the absence of 679 diseases which could not be inferred by CADIAG-II previously.

We believe that creating a new knowledge base explicitly designed for Bilattice CADIAG-II, which would make extensive use of negative rules and counter-evidence for a medical conclusion other than total exclusion, could still further improve the (already very good) performance of Bilattice CADIAG-II.

After introducing the backgrounds of the CADIAG-II system, we discuss the basics of bilattices and present Bilattice CADIAG-II. We then provide an overview of its implementation, describe the results of the retrospective evaluation, and discuss the performance of the presented extension.

2 Background – CADIAG-II

2.1 Overall Consultation Process

Inferring a diagnosis from a given set of patient's medical data in all CADIAG implementations is achieved in four steps, which are shown in Fig. 1.

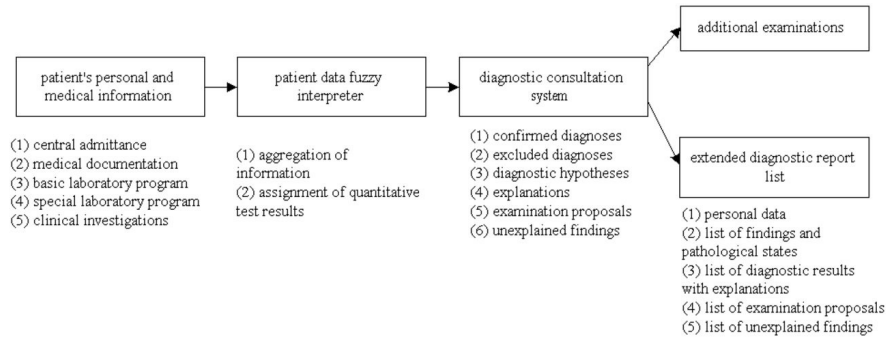


Fig. 1 The consultation process in CADIAG (from [7], p. 208)

Step 1: The physician (or some allied medical personnel) enters personal and medical data about the patient. These usually consist of detailed observational data, such as the medical history, signs from physical examinations, quantitative laboratory test results, and the outcome of clinical investigations (e.g., X-ray and ultrasonography). CADIAG makes a clear distinction between patient-recounted, physician-reported, and laboratory-measured data and their abstraction as clinical, usually linguistic terms applied in diagnostic discourse.

Step 2: A transformation step named data-to-symbol conversion abstracts or aggregates patient information into clinical terms [8]. Aggregation combines one or more documentation items from the electronic health record into an abstract symptom, sign, laboratory or clinical result using logical operators. Here, two-valued Boolean logic is applied. Abstraction is used to transform quantitative test results into abstract medical concepts, and give them a particular evidence value $\in [0, 1]$. An example of an abstracted symptom is ‘elevated serum glucose level’, which is set according to the quantitative result of the glucose test and the definition of *elevated*. The formal modeling of semantic medical concepts such as ‘elevated’ that considers their inherent unsharpness of boundaries in linguistic concepts, visible in their gradual transition to adjacent medical concepts, is based on fuzzy set theory. Fuzzy sets are defined by membership functions, which assign to every symptom S_i a degree of membership μ_{S_i} . These degrees express the level of compatibility of the measured concrete value with the semantic concepts under consideration. They range from zero to unity, wherein zero stands for ‘not compatible’ and unity for ‘fully compatible’ (see Fig. 2).

Step 3: Starting with the set of medical entities and their corresponding evidence values generated by data-to-symbol conversion, CADIAG infers sets of confirmed diagnoses, diagnostic hypotheses, excluded diagnoses, and unexplained findings. The basic concept CADIAG-II’s inference mechanism relies upon is the compositional rule of fuzzy inference [42], which allows inference under uncertainty. The rules contained in the knowledge base are iteratively applied to the set of medical entities pertaining to the patient until a fixpoint is reached.




Fig. 2 Symbolic representation of medical entities using fuzzy sets (from [7], p. 211)

Step 4: In addition to the diagnostic results, CADIAG proposes a list of useful examinations that will possibly confirm or exclude some of the generated diagnostic hypotheses. The generated diagnostic results are explained in detail by a separate explanatory system.

2.2 Knowledge Representation

Definitional, causal, statistical, or heuristic relationships between single and compound fuzzy logical antecedents (left-hand side) and consequences (right-hand side) are represented as IF-THEN rules. Rules with a single medical entity as antecedent, such as a symptom or an abstracted laboratory test result, express associations between two medical entities. Compound antecedents are represented as combinations of medical entities connected by *and*, *or*, and *not*, as well as the operators *at least* and *at most*. They permit the definition of pathophysiological states as well as the incorporation of specific complex, but medically well-known criteria for diagnosing diseases. The associations between the IF- and the THEN-part of the rules are characterized by two kinds of relationships: the frequency of occurrence (*FOO*) of the antecedent with the consequence, and the strength of confirmation (*SOC*) of the antecedent for the consequence.

2.2.1 Rating of Medical Entities and Data-to-Symbol Conversion

Reported and measured medical data are always assigned their *natural* data type, i.e., integers or real numbers for laboratory findings, and one of the Boolean values TRUE or FALSE for binary data. In data-to-symbol conversion, CADIAG-II assigns

a real number $[0,1]$ or a ‘strength of evidence’ to every symptom by applying the two mechanisms described in *Step 2* (abstraction and/or aggregation), wherein a value of 1 means that the corresponding symptom is fully present, while values in $]0, 1[$ mean that the symptom is present in the patient to a certain degree. Symptoms that can definitely be excluded are assigned a value of 0. A value of ϵ is assigned to non-examined medical entities. Since data-to-symbol aggregation rules only operate on Boolean data items, the operators *and*, *or*, *not*, *at least* and *at most* are interpreted and applied in their natural manner.

2.2.2 Interpretation of *FOO* and *SOC* and Type of Rules

Relationships between medical entities are represented as rules being attributed with the frequency of occurrence *FOO* and the strength of confirmation *SOC*. The interpretation for *FOO* and *SOC* as proposed in [7] is the following: given a set of patients P

$$FOO = \frac{\sum_a \min\{\alpha(a), \beta(a)\}}{\sum_a \beta(a)} \quad SOC = \frac{\sum_a \min\{\alpha(a), \beta(a)\}}{\sum_a \alpha(a)}$$

where $\alpha(a)$ and $\beta(a)$ are the degrees to which the entities α and β apply to a patient a , and the sum \sum_a ranges over all patients in P . The patient database associated with CADIAG-II did not contain enough patients for calculating all numbers *FOO* and *SOC* by the above formulas. But this is true for any patient database, even for those of large hospitals – one does not have enough data to calculate *all* associations between *all* symptoms and *all* diseases! For this reason, most of these values were estimated by clinical experience of physicians and taken from published data in medical text books and scientific medical journals. Both, *FOO* and *SOC* are real numbers in $[0,1]$. Similar to evidence values, the values 0 and 1 are also specifically interpreted in CADIAG-II. $SOC = 1$ ensures that the right-hand side of the rule holds, if the IF-part is true. $FOO = 1$ means that the left-hand side has to occur with the right-hand side, otherwise, the right-hand side is excluded. $FOO = 0$ and $SOC = 0$ says that the left-hand side never occurs with the right-hand side in this rule (and vice versa). If the IF-part is true, the right-hand side must be excluded.

According to these definitions, rules in CADIAG-II may express the following IF-THEN relationships between two expressions α and β :

1. α implies β to the degree $d \in (0, 1]$
2. α excludes β
3. the exclusion of α implies the exclusion of β

Thus, a distinction is made between three groups of rules. This classification is based on the following interpretation of *FOO* and *SOC*:

c_d , representing ‘confirming to the degree d ’
 (c_d) when $0 < SOC = d \leq 1$ and $0 < FOO < 1$

me, representing ‘mutually exclusive’
(me) $SOC = 0$ and $FOO = 0$

ao, representing ‘always occurring’
(ao) when $0 < SOC < 1$ and $FOO = 1$

A prototype of a CADIAG-II rule would be:

$D77$: $SYC7$ with $FOO = 1.0$, $SOC = 1.0$
 $SYC7$: $(D1 \wedge S602) \wedge \neg((S1001 \vee S758) \vee S761)$,

where $D77$ is ‘seropositive rheumatoid arthritis, stage I’, $D1$ is ‘rheumatoid arthritis’, $S602$ is ‘Waalser Rose test, positive’, $S1001$ is ‘X-ray, joints, symptoms of arthritis, erosions’, $S758$ is ‘X-ray, joints, partial dislocation’, and $S761$ is ‘X-ray, joints, ankylosis of the peripheral joints’.

This rule, which is of the type (c_d) , with $d = 1$, is interpreted as follows

```

IF   rheumatoid arthritis AND Waaler Rose test, positive AND
      NOT (
          X-ray, joints, symptoms of arthritis, erosions OR
          X-ray, joints, partial dislocation OR
          X-ray, joints, ankylosis of the peripheral joints
      )
THEN seropositive rheumatoid arthritis, stage I

```

The left-hand side of the rule confirms the right-hand side or may confirm it to a certain degree, while the left-hand side obligatorily occurs with the right-hand side of the rule. Thus, if the IF-part is evaluated to 0, the right-hand side will be excluded.

2.2.3 Inference and Operator Usage

The central concept of CADIAG-II’s inference is the compositional rule of fuzzy inference [42]. Using the strength of evidence of medical entities after data-to-symbol conversion and all rules from the knowledge base as starting point, the inference mechanism calculates the degree of evidence μ_{PD} for a patient P and a particular disease D_j using the following equations:

For hypotheses generation and confirmation:

$$\mu_{PD}^1(P, D_j) = \max_{S_i} \min\{\mu_{PS}(P, S_i), \mu_{SD}^{SOC}(S_i, D_j)\} \text{ for rules of type } (c_d)$$

For exclusion by present, but excluding symptoms:

$$\mu_{PD}^2(P, D_j) = \max_{S_i} \min\{\mu_{PS}(P, S_i), 1 - \mu_{SD}^{SOC}(S_i, D_j)\} \text{ for rules of type } (me)$$

For exclusion by absent, but obligatory symptoms:

$$\mu_{PD}^3(P, D_j) = \max_{S_i} \min\{1 - \mu_{PS}(P, S_i), \mu_{SD}^{FOO}(S_i, D_j)\} \text{ for rules of type (ao)}$$

Here, μ_{PS} denotes the strength of evidence for patient P and a particular symptom S_i , μ_{SD} denotes the *FOO* and *SOC* relationships, resp., between symptom S_i and disease D_j . For every symptom-disease relationship, i.e., for every rule in the knowledge base, the minimum of μ_{PS} and μ_{SD} is interpreted as the strength of evidence implied by a particular symptom. The overall strength of evidence for a particular disease is calculated as the maximum of all evidences from rules indicating this particular disease. There is one exception to this procedure: if at least one rule infers an evidence of 0 (or exclusion), then the evidence of the corresponding disease is always set to 0 and is thus excluded.

An additional evidence value ω is used during inference to represent contradictions. If a medical entity has been proven by the inference process, i.e., set to 1, and another rule infers exclusion or evidence of 0, then the evidence of the involved medical concept is set to ω and the inference process is stopped due to this contradiction for the involved entities. All other inferences continue to be processed.

Inference steps applying all possible rules to the available evidence are repeated until the change of every evidence value within one inference step is less than a given threshold (e.g., 0.01), i.e., until a fixpoint is reached. Symptom-symptom, symptom combination-disease, and disease-disease relationships also exist, thus $\mu_{SS}(S_i, S_j)$, $\mu_{SCD}(SC_i, D_j)$, and $\mu_{DD}(D_i, D_j)$ are part of the extensive CADIAG-II knowledge base (for details, see [2]).

For the evaluation of the truth values of complex antecedents in inference rules, *and* is calculated as *min*, or as *max*, not as *complement* ($1-x$), *at least i of n* uses the i -th smallest of n evidence values, and *at most i of n* uses the i -th largest of n evidence values.

3 Bilattice CADIAG-II

CADIAG-II can only express the total exclusion of a medical entity and is unable to provide for so-called negative evidence, i.e., indicating the absence of a particular medical entity not only with certainty but also to a certain degree. Moreover, the syntax of CADIAG-II rules impedes the definition of rules giving graded evidence against a medical entity, and the compositional rule of inference will always prefer the higher rating of an entity over a lower rating (except in case of exclusion). These properties are sometimes listed as weaknesses of the CADIAG-II system [16].

To overcome these limitations, we propose an extension of CADIAG-II which was mainly inspired by peculiar algebraic structures known as bilattices. We provide an introduction to the subject before explaining the proposal and its implementation.

3.1 Algebraic Preliminaries on Bilattices

Let us start by briefly recalling the definition of lattices, which are among the most important algebraic structures in logic [9].

Definition 1 Let S be a non-empty set and \leq an order relation on S . (S, \leq) is known as a *lattice* if, given any $x, y \in S$, there exist in S both the infimum (the greatest element smaller than x, y according to the order \leq) and the supremum (the smallest element greater than x, y according to the order \leq) of $\{x, y\}$ with respect to \leq .

The operations \wedge and \vee defined by $x \wedge y = \inf\{x, y\}$ and $x \vee y = \sup\{x, y\}$ are known as *lattice operations*. A unary operation \neg is named a *negation* on a lattice if, for all $x, y \in S$

- $\neg\neg x = x$,
- $x \leq y \Rightarrow \neg y \leq \neg x$

Definition 2 A lattice (S, \leq) is considered *bounded* when the maximum and minimum element exist in (S, \leq) , i.e., elements in S which are greater (or smaller) than any other elements of S . A bounded lattice (S, \leq) is *complete* if, for every non-empty $X \subseteq S$, $\inf X$ and $\sup X$ belong to S .

Example 1 A very natural example of a complete lattice with negation is the structure $([0, 1], \leq, \neg, 0, 1)$ where $[0, 1]$ is the interval of real numbers between 0 and 1; \leq is the usual ordering of the real numbers, and $\neg x = 1 - x$. We will refer to this structure as a *standard real lattice*.

Bilattices were introduced by Ginsberg [21] as a general framework for various applications in artificial intelligence. The underlying concept is to deal with two order relations. The first represents a ‘degree of truth’ and is, in fact, just a generalization of the usual ordering of truth values in classical and in multi-valued logic. The second ordering is meant to represent the quantity of information obtained for a proposition. Degrees of knowledge permit representation of the difference between ‘not knowing if a proposition is true or false’ (the proposition is evaluated with the minimum of the knowledge order) and ‘knowing that a proposition is false’ (the proposition is evaluated with the minimum of the truth order). More formally, we have the following:

Definition 3 Let $B_t = (B, \leq_t, False, True)$ and $B_k = (B, \leq_k, \perp, \top)$ be complete lattices, where B is a non-empty set, $False, True$ are the minimum and maximum for \leq_t , and \perp and \top are the minimum and maximum for \leq_k .

We refer to the structure $\mathcal{B} = (B, \leq_t, \leq_k, False, True, \perp, \top)$ as *bilattice*.

A *negation* over \mathcal{B} is a unary operation \neg such that:

- $\neg\neg x = x$
- $x \leq_t y \Leftrightarrow \neg y \leq_t \neg x$
- $x \leq_k y \Leftrightarrow \neg x \leq_k \neg y$

As B_t and B_k are lattices, for each one of them we will have two corresponding lattice operations, denoted with \wedge_t, \vee_t and \wedge_k, \vee_k respectively.

Note that the intended meaning of the orderings is only revealed by the notion of negation: the truth order \leq_t is indeed reverted by negation, while the knowledge order \leq_k is preserved.

A prominent example of bilattice, which will be used in the sequel, is the so called *product bilattice*; see [19]. The elements of this structure are pairs, which are intended to represent *reasons for* and *reasons against* the truth of a given proposition.

Example 2 Let $\mathcal{L} = (L, \leq, 0, 1)$ be a complete lattice. We refer to the following structure as the *product bilattice* over L

$\mathcal{B}(\mathcal{L}) = (L \times L, \leq_t, \leq_k, (0, 1), (1, 0), (0, 0), (1, 1))$ where:

- $(x, y) \leq_t (x', y') \Leftrightarrow x \leq x' \text{ and } y' \leq y$
- $(x, y) \leq_k (x', y') \Leftrightarrow x \leq x' \text{ and } y \leq y'$
- $(0, 1)$ and $(1, 0)$ are minimum and maximum, respectively, for \leq_t
- $(0, 0)$ and $(1, 1)$ are minimum and maximum, respectively, for \leq_k

We may introduce a negation over $\mathcal{B}(L)$ by letting:

$$\neg(x, y) = (y, x).$$

Informally, given two elements of a product bilattice (i.e., two pairs of values) a and b , the example above says that “ a is *less true* than b ” when for a there are fewer *reasons for* and more *reasons against* than for b , while “ a is *less known* than b ” when for a there are both, fewer *reasons for* and fewer *reasons against* than for b .

From the relation between bilattices and lattice orderings, it is easy to establish how bilattice operations in a product bilattice relate to the original lattice ones. Indeed, we have:

- $(x, y) \wedge_t (x', y') = (x \wedge x', y \vee y')$
- $(x, y) \vee_t (x', y') = (x \vee x', y \wedge y')$
- $(x, y) \wedge_k (x', y') = (x \wedge x', y \wedge y')$
- $(x, y) \vee_k (x', y') = (x \vee x', y \vee y')$

3.2 An Extension of CADIAG-II Based on Bilattices – bCADIAG-II

We have now introduced all prerequisites to describe the proposed extension of CADIAG-II based on bilattices (hence the name Bilattice CADIAG-II or bCADIAG-II, for short). By applying the concept of product bilattices, we simply associate each basic entity with not just a single degree in $[0, 1]$ but a pair, representing *reasons for* and *reasons against* the truth of the entity. The interpretation of these values is as follows: a value of 0 means that we do not have evidence (or counter-evidence) of this medical entity, while a value of 1 is interpreted as full evidence (or full exclusion). Intermediate values denote insufficient evidence to either fully confirm or fully exclude the entity in question.

Since data-to-symbol conversion in CADIAG-II uses fuzzy sets and rules for more or less evident (borderline) symptoms (and none for more or less excluded symptoms), we do not have any direct initial evaluation of counter-evidence that we could directly use in an extended version of the system. Therefore, whenever the data-to-symbol conversion issues the value c to a particular entity, we associate with that entity the pair of evidence and counter-evidence $(c, 1 - c)$.

From now on, each basic entity of the system will be represented as a $\alpha (s, t)$, where α is an atomic formula and (s, t) an element of the product bilattice, where the value s stands for *reasons for* α , and the value t for *reasons against* α . Recall that compound rules of CADIAG-II deal with complex logical formulas. Therefore, after an initial evaluation of the entities, i.e., an association of a pair of values to them, compound formulas will be obtained as follows:

For any α, β basic entities in CADIAG-II, and $(s, t), (u, v) \in B(L)$

$$(\wedge) \frac{\alpha (s, t) \quad \beta (u, v)}{(\alpha \wedge_k \beta)(s \wedge u, t \vee v)}$$

$$(\neg) \frac{\alpha (s, t)}{\neg \alpha \neg(s, t) = (t, s)}$$

Let us now focus on the IF-THEN rules of the system. We will focus on their role in transmitting knowledge (in the sense of bilattices) rather than merely truth. We will represent each such rule in the format

$$\text{IF } \alpha \text{ THEN } \beta (x, y)$$

where α and β are compound formulas, denoting symptoms and disease, respectively, and (x, y) is a pair of values in the product bilattice.

Given a compound formula α evaluated with the pair (u, v) in the bilattice product, and an IF-THEN rule as the one above, we will then compute the value of the conclusion β simply as follows:

$$(k) \frac{\text{IF } \alpha \text{ THEN } \beta (s, t) \quad \alpha (u, v)}{\beta (s, t) \wedge_k (u, v)}$$

The form of the rule might suggest a sort of ‘knowledge modus ponens’, where the value associated with the antecedent and the one associated with the rule are combined in order to obtain the value of the consequent, by means of the operation \wedge_k . However, it should be noted that the pair of values attached to IF α THEN β should not be regarded as ‘reasons for’ and ‘reasons against’ an implication $\alpha \rightarrow \beta$, but rather as a measure of how much, *from the values of evidence and counter-evidence of the antecedent, we can infer evidence and counter-evidence of the consequent, respectively*. The crucial issue is then to associate such pairs of values with each of the IF-THEN rules, since such values are not immediately provided from the dataset. Depending on the type of rules, we proceed as follows:

- We represent rules of type c_d and ao , as:

$$\text{IF } \alpha \text{ THEN } \beta \text{ (SOC, FOO).}$$

Note that rules of type ao , i.e., those rules where the full exclusion of the premise implies the full exclusion of the conclusion, are just a particular case of rules of type c_d with $FOO = 1$.

- We generalize rules of type me as

$$\text{IF } \alpha \text{ THEN } \neg\beta \text{ (1 - SOC, 1 - FOO).}$$

Note that the mutually exclusive rule of CADIAG-II is a particular case of the me rule above, with $SOC = 0$ and $FOO = 0$. For each of the above rules c_d , ao , and me , we may use (k) to combine the pair associated with the premises and the one associated with the IF-THEN rule in order to obtain the relevant pair associated with the conclusion.

The use of FOO in the c_d rules is a new feature in our proposal. It expands the inferential power of our system with respect to CADIAG-II: values of FOO different from 1 were indeed present in rules of type c_d of CADIAG-II, but were not previously used at all. The use of FOO is justified by the following consideration: (1) FOO is a generalization of the conditional probability of the premise of the rule, given the conclusion; (2) such a value is directly proportional to the conditional probability of the negation of the conclusion, given the negation of the premises; (3) the latter is a measure of how much the exclusion of the premises allows inference of the exclusion of the conclusion. Note that, as a limit case, the rules of type (ao) , i.e., those rules where the full exclusion of the premise implies the full exclusion of the conclusion, are actually those with $FOO = 1$.

Finally, the use of the pair $(1 - SOC, 1 - FOO)$ rather than (SOC, FOO) for the rules of type me , is justified by the fact that we are using $\neg\beta$ rather than β as the conclusion of the rule.

So far, we have presented a different way of dealing with rules and inferences for CADIAG-II. Since the bilattice operation \wedge_k is used for combining the premise in the rule (k) , the focus of the inference process will no longer be on truth, but on knowledge order, aiming to maximize the latter. In this spirit, it appears reasonable to require that, for any entity β , all pairs of values produced for β by the system via applications of (k) are then combined through \vee_k .

Remark 1 All logical operations, including negation are monotone with respect to knowledge order, so that a fixpoint can be found for each entity.

Remark 2 The generalization of the rules of type me will only have an effect if negative rules are incorporated into knowledge bases.

Let us recall that, in CADIAG-II, the value 0 (for falsity) was treated in a different way than other values, because it was given preference over higher non-zero results, while the highest value was always chosen for all remaining truth values. This shows that a knowledge order was already implicitly involved there. The value 0, which

stands for ‘totally false’, was indeed taken to provide more knowledge than other intermediary values, namely the full exclusion of a given entity. This is treated in a more elegant and coherent way in our approach.

4 Implementation and Experimental Results

CADIAG-II was originally developed and run on an IBM host computer system [18] (at the Department of Medical Computer Sciences, University of Vienna Medical School) which is no longer in operation. In order to evaluate the described improvements, both CADIAG-II and bCADIAG-II were implemented within the PC-based medical expert system shell MedFrame [24], and comparatively tested against a set of patients with clinically confirmed discharge diagnoses.

4.1 MedFrame, CADIAG-II, and bCADIAG-II

MedFrame [24] is an expert system shell designed especially for implementing medical expert systems using fuzzy concepts. It provides the medical knowledge engineer with

- various knowledge representation formalisms to store medical knowledge and reflect adequate inference mechanisms,
- concepts for modeling and handling uncertainty in medical terminology and medical relations, with special emphasis on fuzzy methodologies,
- mechanisms for storing patient data and history in a standardized manner,
- concepts for representing medical knowledge in a standardized way, and
- utilities for implementing inference mechanisms easily and rapidly.

The rheumatological knowledge base of the original CADIAG-II [5] which currently contains 1,126 symptoms and 170 documented diagnoses was imported into MedFrame’s knowledge database, resulting in 658 fuzzy sets and 2,143 rules for data-to-symbol conversion, as well as 21,470 IF-THEN rules for inference (982 symptom-symptom, 368 disease-disease, 61 symptom-combination-disease, and 20,041 symptom-disease relationships). MedFrame’s utilities for developing inference mechanisms were used to re-implement CADIAG-II based upon the original IBM host implementation. In combination, the transferred knowledge base as well as the newly implemented inference mechanisms entirely comply with all the approaches described in Sect. 2.

In addition, the modified inference process of bCADIAG-II described in Sect. 3.2 was implemented. The operators were rendered capable of dealing with unknown entities analogous to CADIAG-II. Therefore, bCADIAG-II incorporates the following improvements:

- Inclusion of negative evidence. In addition to the strength of evidence, for every medical concept the strength of counter-evidence is also maintained in form of a product bilattice.
- Advanced handling of *FOO*. In addition to (me) rules, *FOO* and \wedge_k are also used in the evaluation of (c_d) rules.
- Application of \vee_k for calculating the overall evidence of a medical concept.

4.2 Evaluation and Results

For evaluating the performance of bCADIAG-II compared to that of CADIAG-II, the data of 3,131 anonymized hospitalized rheumatic patients were imported into MedFrame's patient database, including extended clinical data from the patients' histories, physical examinations, laboratory tests, and clinical investigations. Furthermore, all the 8,445 clinically confirmed discharge diagnoses of these patients were also transferred to MedFrame and used as a diagnostic gold standard. The number of discharge diagnoses in the set of available patient data ranged from 0–9 (mean 2.59, median 2).

The 3,131 patients were then analyzed by both implementations, CADIAG-II and bCADIAG-II, applying the same knowledge base. This was done in batch mode. Since the patient data cases only contained clinically confirmed existing diseases and no information about definitely absent diseases, the evaluation was focused on opposing the discharge diagnoses to confirmed diagnoses and diagnostic hypotheses. In this context, confirmed is equivalent to a strength of evidence of 1.0, and hypothetical is equivalent to a strength of evidence between 0.4 and 0.99. Moreover, for bCADIAG-II the strength of counter-evidence of a concept was required to be less than 0.4 in order to be considered hypothetical.

The interpretation for each patient was assigned to one of six classes, as shown in Table 1. The results of the evaluation are listed in Table 2.

4.3 Discussion of Results

The original CADIAG-II/RHEUMA implementation was evaluated several times focusing on confirmation of the correctness and soundness of the generated diagnostic results in retrospective and prospective studies [23,28]. In contrast, the evaluation at hand did not check the results for correctness, but concentrated solely on the impact of the undertaken change of the underlying inference process on the outcome.

Table 2 clearly shows that bCADIAG-II performs just as well as CADIAG-II. The inferred results are identical, except for differences in the number of generated excluded diagnoses.

Since neither of the modifications has an impact on the calculation of positive evidences, it is no wonder that the inference results are identical with respect to

Table 1 Classification of interpretations

Class	Categorization
Full Hit	All discharge diagnoses were contained in the diagnostic results as confirmed or hypothetical
75–99%	Between 75% and 99% of all discharge diagnoses were contained in the diagnostic results as confirmed or hypothetical
50–74%	Between 50% and 74% of all discharge diagnoses were contained in the diagnostic results as confirmed or hypothetical
25–49%	Between 25% and 49% of all discharge diagnoses were contained in the diagnostic results as confirmed or hypothetical
1–24%	Between 1% and 24% of all discharge diagnoses were contained in the diagnostic results as confirmed or hypothetical
No Hit	None of the discharge diagnoses were contained in the diagnostic results as confirmed or hypothetical

Table 2 Interpretation results

Class	CADIAG-II	bCADIAG-II
Full Hit	937/29.89%	937/29.89%
75–99%	209/6.68%	209/6.68%
50–74%	1,217/38.87%	1,217/38.87%
25–49%	367/11.72%	367/11.72%
1–24%	15/0.48%	15/0.48%
No Hit	387/12.36%	387/12.36%
Number of confirmed diagnoses	569	569
Number of diagnostic hypotheses	20,777	20,779
Number of excluded diagnoses	50,789	50,765
Mean/median/maximum of confirmed diagnoses	0.17/0/3	0.17/0/3
Mean/median/maximum of diagnostic hypotheses	6.37/6/22	6.37/6/22
Mean/median/maximum of excluded diagnoses	15.57/16/36	15.57/16/36

confirmed and hypothetical diagnoses. Apart from the 50,765 excluded diagnoses, bCADIAG-II additionally infers 703 hypothetical absent diagnoses, i.e., diagnoses with a strength of evidence less than 0.4 and a strength of counter-evidence more than 0.4. Twenty-four of these cases are the reason for the difference in the number of excluded diagnoses (50,789 in CADIAG-II and 50,765 in bCADIAG-II). In bCADIAG-II, the ‘knowledge modus ponens’ assigns strength of evidence to these concepts, which is an improvement of the inference process. Apart from these 24 cases, an additional 679 hypothetical absent diagnoses for 258 patients are provided by bCADIAG-II.

These numbers demonstrate the high potential of using negative evidence, especially in the process of differential diagnosis. A computer-based differential diagnostic system, accordingly equipped, would provide the physician with information about diseases which are most likely not present, and thus direct the physician's attention to other diseases. Yet, since the CADIAG-II/RHEUMA knowledge base does not utilize these concepts (except for full exclusion) and the improvements in the results are only due to advancements in the inference process, there is a clear need to re-design the respective knowledge base. It should include the use of the concept of negative evidence.

Apart from comparative results, the evaluation confirms the results of previous studies. bCADIAG-II was able to infer at least one of the available discharge diagnoses for 88% of the reference patients, and more than 75% of all discharge diagnoses for more than 36% of them. While the given evaluation employed a threshold of 0.4 for diagnostic hypotheses, it was set to 0.2 in [28]. An evaluation of 3,131 reference patients with bCADIAG-II and a threshold of 0.2 resulted in the detection of at least one of the available discharge diagnoses for 95.5% of the reference patients, and more than 75% of all discharge diagnoses for more than 68% of patients.

5 Conclusions

After some further steps in the formalization of the CADIAG-II inference process [12], bCADIAG-II should be another measure towards putting CADIAG-II onto an extended solid formal basis. By applying the concept of product bilattices, the prerequisites for including negative evidence (i.e., rules diminishing the certainty of a particular diagnosis) into the CADIAG-II system was established. The experimental results proved the identical behavior of CADIAG-II and bCADIAG-II and confirmed the quality of inference results in comparison with former evaluations. In addition, bCADIAG-II increased the quantity of generated results in the form of indications to absent diseases other than those excluded with the existing knowledge base. Nevertheless, the evaluation clearly showed that significant improvements can only be achieved by re-creating the knowledge base and make extensive use of negative rules and counter-evidence other than total exclusion.

Acknowledgements This research was partly supported by the Vienna Science and Technology Fund (WWTF), Grant no. MA07-016. We are indebted to Andrea Rappelsberger for her extended assistance in formatting and finalizing this report.

References

1. Adlassnig, K.-P.: A fuzzy logical model of computer-assisted medical diagnosis. *Methods of Information in Medicine* **19**(3), 141–148 (1980)

2. Adlassnig, K.-P.: Fuzzy set theory in medical diagnosis. *IEEE Transactions on Systems, Man, and Cybernetics* **SMC-16**(2), 260–265 (1986)
3. Adlassnig, K.-P., Akhavan-Heidari, M.: CADIAG-2/GALL: An experimental expert system for the diagnosis of gallbladder and biliary tract diseases. *Artificial Intelligence in Medicine* **1**(2), 71–77 (1989)
4. Adlassnig, K.-P., Grabner, G.: The Viennese computer-assisted diagnostic system. Its principles and values. *Automedica* **3**, 141–150 (1980)
5. Adlassnig, K.-P., Kolarz, G.: CADIAG-2: Computer-assisted medical diagnosis using fuzzy subsets. In: Gupta, M.M., Sanchez, E. (eds) *Approximate Reasoning in Decision Analysis*, pp. 219–247, North-Holland Publishing Company, Amsterdam (1982)
6. Adlassnig, K.-P., Scheithauer, W.: Performance evaluation of medical expert systems using ROC curves. *Computers and Biomedical Research* **22**(4), 297–313 (1989)
7. Adlassnig, K.-P., Kolarz, G., Scheithauer, W., Grabner, G.: Approach to a hospital-based application of a medical expert system. *Medical Informatics* **11**(3), 205–223 (1986)
8. Boegl, K., Leitich, H., Kolousek, G., Rothenfluh, T., Adlassnig, K.-P.: Clinical data interpretation in MedFrame/CADIAG-4 using fuzzy sets. *Biomedical Engineering – Applications, Basis & Communications* **8**(6), 488–495 (1996)
9. Burris, S., Sankappanavar, H.P.: *A Course in Universal Algebra*. Graduate Texts in Mathematics. Springer, New York (1981)
10. Choy, G., Khalilzadeh, O., Michalski, M., Do, S., Samir, A.E., Panykh, O.S., Geis, J.R., Pandharipande, P.V., Brink, J.A., Dreyer, K.J.: Current applications and future impact of machine learning in radiology. *Radiology* **288**(2), 318–328 (2018)
11. Chunyan, A., Shunshan, Y., Hui, D., Quan, Z., Liang, Y.: Application and development of artificial intelligence and intelligent disease diagnosis. *Current Pharmaceutical Design* **26**(26), 3069–3075 (2020)
12. Ciabattoni, A., Vetterlein, T.: On the (fuzzy) logical content of CADIAG-2. *Fuzzy Sets and Systems* **161**(14), 1941–1958 (2010)
13. Ciabattoni, A., Vetterlein, T., Adlassnig, K.-P.: A formal framework for Cadiag-2. In: Adlassnig, K.-P., Blobel, B., Mantas, J., Masic, I. (eds) *Medical Informatics in a United and Healthy Europe – Proceedings of MIE 2009 – The XXIIInd International Congress of the European Federation for Medical Informatics, Studies in Health Technology and Informatics 150*, pp. 648–652, IOS Press, Amsterdam (2009)
14. Ciabattoni, A., Picado-Muñoz, D., Vetterlein, T., El-Zekey, M.: Formal approaches to rule-based systems in medicine: The case of CADIAG-2. *International Journal of Approximate Reasoning* **54**(1), 132–148 (2013)
15. Daniel, M.: Remarks on a cyclic inference in the fuzzy expert system CADIAG-IV. In: Phuong, N.H., Ohsato, A. (eds) *VJFuzzy '98: Vietnam-Japan Bilateral Symposium on Fuzzy Systems and Applications*, Halong Bay, Vietnam, 30th September – 2nd October, 1998, Proceedings, pp. 619–627, Hanoi (1998)
16. Daniel, M.: Theoretical comparison of inference in CADIAG and MYCIN-like systems. *Tatra Mountains Mathematical Publications* **16**(2), 255–272 (1999)
17. Daniel, M., Hájek, P., Nguyen, P.H.: CADIAG-2 and MYCIN-like systems. *Artificial Intelligence in Medicine* **9**(3), 241–259 (1997)
18. Fischler, F.: *Die Wissensbasis und der Inferenzprozeß des medizinischen Expertensystems CADIAG-II/E*. Diploma thesis, University of Vienna, Vienna (1994)
19. Fitting, M.: Kleene's logic, generalized. *Journal of Logic and Computation* **1**(6), 797–810 (1990)
20. Fryback, D.G.: Bayes' theorem and conditional nonindependence of data in medical diagnosis. *Computers and Biomedical Research* **11**(5), 423–434 (1978)
21. Ginsberg, M.: Multivalued logics: A uniform approach to inference in artificial intelligence. *Computer Intelligence* **4**(3), 265–316 (1988)
22. Hajek, P., Phuong, N.H.: Möbius transform for CADIAG-2. *Journal of Computer Science and Cybernetics* **13**(3), 103–122 (1997)
23. Kolarz, G., Adlassnig, K.-P.: Problems in establishing the medical expert systems CADIAG-1 and CADIAG-2 in rheumatology. *Journal of Medical Systems* **10**(4), 395–405 (1986)

24. Kopecky, D., Adlassnig, K.-P.: A framework for clinical decision support in internal medicine. In: Schreier, G., Hayn, D., Ammenwerth, E. (Eds.) *Tagungsband der eHealth2011 – Health Informatics meets eHealth – von der Wissenschaft zur Anwendung und zurueck, Grenzen ueberwinden – Continuity of Care*, 26.–27. Mai 2011, Wien, pp. 253–258, Oesterreichische Computer Gesellschaft, Wien (2011)
25. Kulikowski, C.A.: Artificial intelligence methods and systems for medical consultation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PAMI-2**(5), 464–476 (1980)
26. Ledley, R.S., Lusted, L.B.: Reasoning foundations of medical diagnosis. Symbolic logic, probability, and value theory aid our understanding of how physicians reason. *Science* **130**(3366), 9–21 (1959)
27. Leitich, H., Adlassnig, K.-P., Kolarz, G.: Development and evaluation of fuzzy criteria for the diagnosis of rheumatoid arthritis. *Methods of Information in Medicine* **35**(4-5), 334–342 (1996)
28. Leitich, H., Kiener, H.P., Kolarz, G., Schuh, C., Graninger, W., Adlassnig, K.-P.: A prospective evaluation of the medical consultation system CADIAG-II/RHEUMA in a rheumatological outpatient clinic. *Methods of Information in Medicine* **40**(3), 213–220 (2001)
29. Leitich, H., Adlassnig, K.-P., Kolarz, G.: Evaluation of two different models of semiautomatic knowledge acquisition for the medical consultant system CADIAG-II/RHEUMA. *Artificial Intelligence in Medicine* **25**(3), 215–225 (2002)
30. Miller, R.A.: Medical diagnostic decision support systems—Past, present, and future: A threaded bibliography and brief commentary. *Journal of the American Medical Informatics Association* **1**(1), 8–27 (1994)
31. Picado Muiño, D., Ciabattoni, A., Vetterlein, T. (2013) Towards an interpretation of the medical expert system CADIAG 2. In: Seising, R., Tabacchi, M.E. (eds) *Fuzziness and Medicine: Philosophical Reflections and Application Systems in Health Care*, *Studies in Fuzziness and Soft Computing* 302, pp. 323–338, Springer, Berlin (2013)
32. Rusnok, P., Vetterlein, T., Adlassnig, K.-P. (2009) Cadiag-2 and Fuzzy Probability Logics. In: Adlassnig, K.-P., Blobel, B., Mantas, J., Masic, I. (eds) *Medical Informatics in a United and Healthy Europe – Proceedings of MIE 2009 – The XXIIInd International Congress of the European Federation for Medical Informatics*, *Studies in Health Technology and Informatics* 150, p. 773, IOS Press, Amsterdam (2009)
33. Serag, A., Ion-Margineanu, A., Qureshi, H., McMillan, R., Saint Martin, M.-J., Diamond, J., O’Reilly, P., Hamilton, P.: Translational AI and deep learning in diagnostic pathology. *Frontiers in Medicine* **6**:185 (2019)
34. Shortliffe, E.H.: *Computer-Based Medical Consultations: MYCIN*. *Artificial Intelligence Series 2*, Elsevier, New York (1976)
35. Somashekhar, S.P., Sepúlveda, M.-J., Puglielli, S., Norden, A.D., Shortliffe, E.H., Rohit Kumar, C., Rauthan, A., Arun Kumar, N., Patil, P., Rhee, K., Ramya, Y. Watson for oncology and breast cancer treatment recommendations: Agreement with an expert multidisciplinary tumor board. *Annals of Oncology* **29**(2), 418–423 (2018)
36. Szolovits, P., Pauker, S.G.: Categorical and probabilistic reasoning in medical diagnosis. *Artificial Intelligence* **11**(1-2), 115–144 (1978)
37. Vetterlein, T. Adlassnig, K.-P.: The medical expert system CADIAG-2, and the limits of reformulation by means of formal logics. In: Schreier, G., Hayn, D., Ammenwerth, E. (Eds.) *eHealth2009 – Health Informatics meets eHealth – von der Wissenschaft zur Anwendung und zurück, Tagungsband eHealth2009 & eHealth Benchmarking 2009*, pp. 123–128, Österreichische Computer Gesellschaft, Wien (2009)
38. Warner, H.R., Toronto, A.F., Veasey, L.G., Stephenson, R.: A mathematical approach to medical diagnosis. Application to congenital heart disease. *Journal of the American Medical Association* **177**(3), 177–183 (1961)
39. Wernecke, K.D.: On the application of discriminant analysis in medical diagnostics. In: Bock, H.H., Lenski, W., Richter, M.M. (eds) *Information Systems and Data Analysis. Studies in Classification, Data Analysis, and Knowledge Organization*, pp. 267–279, Springer, Berlin (1994)

40. Yanase, J., Triantaphyllou, E.: A systematic survey of computer-aided diagnosis in medicine: Past and present developments. *Expert Systems with Applications* **138**:112821, (2019)
41. Yu, V.L., Fagan, L.M., Wraith, S.M., Clancey, W.J., Scott, A.C., Hannigan, J., Blum, R.L., Buchanan, B.G., Cohen, S.N.: Antimicrobial selection by a computer. A blinded evaluation by infectious diseases experts. *Journal of the American Medical Association* **242**(12), 1279–1282 (1979)
42. Zadeh, L.A.: Outline of a new approach to the analysis of complex systems and decision processes. *IEEE Transactions on Systems, Man, and Cybernetics* **SMC-3**(1), 28–44 (1973)
43. Zadeh, L.A.: The concept of a linguistic variable and its application to approximate reasoning—I. *Information Sciences* **8**(3), 199–249 (1975)