

UNCLASSIFIED

Defense Technical Information Center
Compilation Part Notice

ADP012694

TITLE: Diagnostic Reasoning with Multilevel Set-Covering Models

DISTRIBUTION: Approved for public release, distribution unlimited

This paper is part of the following report:

TITLE: Thirteenth International Workshop on Principles of Diagnosis
[DX-2002]

To order the complete compilation report, use: ADA405380

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, etc. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:

ADP012686 thru ADP012711

UNCLASSIFIED

Diagnostic Reasoning with Multilevel Set-Covering Models

Joachim Baumeister¹ and Dietmar Seipel¹

Abstract. We consider multilevel set-covering models for diagnostic reasoning: though a lot of work has been done in this field, *knowledge acquisition* efforts have been investigated only insufficiently. We will show how set-covering models can be build incrementally and how they can be refined by knowledge enhancements or representational extensions. All these extensions have a primary characteristic: they can be applied without changing the basic semantics of the model.

Keywords: set-covering diagnosis; model-based diagnosis; qualitative modeling; knowledge acquisition; abductive reasoning

1 Introduction

In this paper we will present a new interpretation of set-covering models [1] which is a suitable representation for the manual development of knowledge-based systems. Because of its simple semantics set-covering models are rapidly understood by the experts, but still maintain a well-known model-based interpretation. In [2] we showed how knowledge-based diagnostic systems can be developed incrementally with set-covering models, thus supporting rapid prototyping of such systems. In this paper we will extend this approach to multilevel set-covering models, and we will describe how simple set-covering models can be enhanced by representational extensions. Practical experience has shown that these additions facilitated the development of a real world example from a medical ICU domain.

A *set-covering model* consists of a set of diagnoses, a set of findings (observations) and covering relations between the elements of these two sets. There exists a covering relation between a diagnosis and a finding, iff the diagnosis implies the observation of the finding. We can define covering relations between diagnoses as well, iff a diagnosis implies the observation of another diagnosis. The basic idea of set-covering diagnosis is the detection of a reasonable set of diagnoses which can explain the given observations. To do this, we propose an abductive reasoning step: Firstly, hypotheses are generated in order to explain the given observations. Secondly competing hypotheses are ranked using a *quality measure*.

Reasoning with set-covering models has got a long tradition in diagnostic reasoning: Early work was done by Patil [3] with his system ABEL, which implemented a comprehensive set-covering representation including causal, associational and grouping relations. Reggia et al. [1] contributed a formal approach to set-covering models and addressed the problem of hypothesis generation with a pre-

cise description of *generator sets*. Later [4] they introduced the integration of Bayesian probabilities in set-covering models. With the system MOLE [5] Eshelman focussed on the problem of acquiring set-covering knowledge. He proposed an interactive process that allows for refining previously acquired knowledge after a reasoning step to differentiate between conflicting hypotheses. Console et al. [6] showed with the system CHECK how to combine heuristic and causal knowledge. There heuristic knowledge was used to find reasonable hypotheses for a given observation. In a second step the causal knowledge was used to generate abductive explanations for the hypotheses. Long [7] extended covering models with probabilities and a rich syntax of temporal and non-temporal causation events. Since knowledge acquisition is a cost sensitive task, reuse of existing knowledge is another emerging aspect. Puppe [8] showed how set-covering knowledge can be combined with other classes of knowledge like heuristic rules, case-based knowledge or decision trees.

Most of these approaches only investigated syntax and semantics of the reasoning process, but did not consider the knowledge engineering process. Eshelman's MOLE system [5] differs from our knowledge acquisition approach, since there knowledge refinement is performed by adding new covering relations to the model. In our paper we will present (multilevel) set-covering models and show how to enrich these simple models with knowledge enhancements like *similarities* and *weights* or *representational extensions* for more complex covering relations. A primary characteristic of the presented extensions is the incrementality: each extension can be applied independently from other enhancements and will not change the basic semantics of the model, but refine special aspects of it.

The rest of the paper is organized as follows: In Section 2 we will introduce the basic concepts of set-covering models and show how to enrich set-covering models with additional knowledge like similarities and weights. Beyond that we will introduce representational extensions of set-covering models in Section 3 that will enable us to formulate exclusions, necessary relations and complex covering relations (conjunctions, disjunctions, cardinalities). In Section 4 we will shortly summarize the problem of hypothesis generation and we will introduce constraints that shrink the exponentiell size of possible hypotheses. We will conclude this paper in Section 5 with an overview of the work we have done so far and promising directions we are planning to work on in the future.

2 Set-Covering Models

A set-covering model consists of a set of diagnoses, a set of findings (observations) and covering relations between the elements of these two sets. There exists a covering relation between a diagnosis and

¹ University of Würzburg, Department of Computer Science, Am Hubland, 97074 Würzburg, Germany, email: {baumeister, seipel}@informatik.uni-wuerzburg.de

a finding, iff the diagnosis predicts the observation of the finding. Furthermore we can define covering relations between two diagnoses to state that a diagnosis implies another diagnosis. In this way we can build a *covering-tree* for a diagnosis, where we postulate that the leafs of the covering-tree have to be observable findings. So each covering path will start with a diagnosis and lead to an observable finding.

2.1 The Basic Model

The basic idea of set-covering diagnosis is the detection of a reasonable set of diagnoses which can explain the given observation of findings. In an *abductive reasoning* step hypotheses are firstly generated in order to explain the given observations (*hypothesis generation*). In a second step, we define a quality measure for ranking competing hypotheses (*hypothesis testing*). Set-covering models describe relations like:

A diagnosis D predicts that the parameters A_1, \dots, A_n are observed with corresponding values v_1, \dots, v_n .
 A diagnosis D predicts the diagnoses D_1, \dots, D_m .

We call each of these relations *covering relations* and we denote them by

$$r_i = D \rightarrow A_i : v_i, \quad 1 \leq i \leq n,$$

$$r'_i = D \rightarrow D_i, \quad 1 \leq i \leq m.$$

Covering models can be visually described like in Figure 1. In this

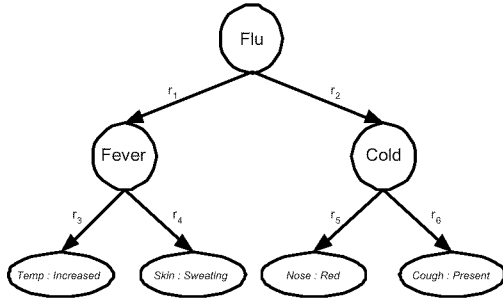


Figure 1. Basic set-covering model for diagnoses *Flu*, *Fever* and *Cold*.

example the model states that diagnosis *Flu* implies the observation of the diagnoses *Fever* and *Cold*. Diagnosis *Fever* itself forces the observation of the attributes *Temperature* and *Skin* with their corresponding values *Increased* and *Sweating*.

The basic algorithm for set-covering diagnosis is very simple: Given a set of observed findings, it uses a simple hypothesize-and-test strategy, which generates hypotheses (coined from diagnoses) in the first step and tests them against the given observations in a second step. The test is defined by calculating a quality measure, which expresses the covering degree of the hypothesis regarding the observed findings. The generation and evaluation of the hypotheses is an iterative process, which stops when a satisfying hypothesis has been found or all hypotheses have been considered. Usually the algorithm will look at single diagnoses, compute the corresponding quality measure, and then it will generate hypotheses with multiple diagnoses, if needed.

In the worst case this procedure will generate 2^n candidates for n diagnoses. So heuristics are needed to keep the method computationally tractable (c.f. Section 4).

The basic sets for this task are the following: We define Ω_D to be the set of all diagnoses and Ω_A the set of all observable parameters (attributes). To each parameter $A \in \Omega_A$ a range $dom(A)$ of values is assigned, and $\Omega_V = \bigcup_{A \in \Omega_A} dom(A)$ is the set of all possible values for the parameters. If a parameter A is assigned to a value v , then we call $A : v$ a *finding*.

$$\Omega_F = \{ A : v \mid A \in \Omega_A, v \in dom(A) \}$$

is the set of all findings. Furthermore we call an element $S \in \Omega_S = \Omega_D \cup \Omega_F$ a *state*.

A *covering relation* r between a diagnosis D and a state S ($S \neq D$) is denoted by $r = D \rightarrow S$. We say that “ D predicts S ” or that “ D covers S ”. Then $c_r = D$ is called the *cause* and $e_r = S$ is called the *effect*. We define Ω_R to be the set of all covering relations contained in the model. Then $D^+ \in \Omega_R$ is the set of all covering relations with diagnosis D as the cause, i.e. $D^+ = \{ r \in \Omega_R \mid c_r = D \}$. E.g., for the model in Figure 1 we obtain $c_{r_1} = Flu$ and $e_{r_1} = Fever$, $Cold^+ = \{ r_5, r_6 \}$.

Since S can be a diagnosis itself, we are able to build *multilevel* set-covering models. A state S *transitively covers* another state S' , if either S covers S' or S covers another state S'' that transitively covers S' .

We call $\mathcal{F}_O \subset \Omega_F$ the set of *observed findings* and a set $\mathcal{H} \subseteq \Omega_D$ of diagnoses a *hypothesis*. A finding that is not transitively covered by the hypothesis \mathcal{H} is called *isolated*, and the set of all observed findings that are isolated will be denoted by $\mathcal{F}_{\mathcal{H},O}^{isolated} \subseteq \mathcal{F}_O$. E.g. for a hypothesis $\mathcal{H} = \{ D_1 \}$ and $\mathcal{F}_O = \{ A_1 : v_1, A_2 : v_2, A_4 : v_4 \}$ we obtain $\mathcal{F}_{\mathcal{H},O}^{isolated} = \{ A_2 : v_2 \}$.




Figure 2. Basic set-covering model for diagnosis D

Now we will describe the computation of the precision of a state for a given observation. The precision $\pi(S)$ of a state S provides a real value between 0 and 1 to describe the degree of accuracy the covered states of S are observed.

Bottom-Up Computation of Precisions. Given the set \mathcal{F}_O of observed findings, the precision π of each state is computed bottom-up starting with the findings:

$$\pi(A : v) = \begin{cases} 1, & \text{if } A : v \in \mathcal{F}_O \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

The precision $\pi(D)$ of a diagnosis D can be computed as soon as the precisions of all its successors S are known. For this we define

$$\begin{aligned} D_{\geq c}^+ &= \{ r \in D^+ \mid \pi(e_r) \geq c(e_r) \}, \\ D_{>0}^+ &= \{ r \in D^+ \mid \pi(e_r) > 0 \}, \end{aligned}$$

as the sets of all *relevant* covering relations, i.e. relations that predict states with a precision greater than a user defined threshold function.

$$\pi(D) = \begin{cases} \frac{\sum_{r \in D_{\geq c}^+} \pi(e_r)}{|D_{>0}^+|}, & \text{if } D_{>0}^+ \neq \emptyset \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

The denominator counts all successor states of D with a positive precision, which gives us the maximally achievable score. The nominator sums up the precision of all successor states with a precision, that is greater than or equal to the completeness value, which gives us the actually achieved score.

The *completeness value* $c(D)$ of a diagnosis is specified by the modeler and is motivated by the fact, that a covering model for a diagnosis will contain more states than the diagnosis will cause in an average case. Nevertheless in most cases the observation of a percentage of the modeled states will legitimate the validation of this diagnosis. To emphasize this percentage the modeler has to specify a completeness value $c(D)$. Unless this factor is reached by the observation set in the current case, the diagnosis may neither be considered as a validly observed state, nor will it be considered as a valid hypothesis candidate.

Since we also want to consider multiple faults, i.e. hypotheses containing more than one diagnosis, we define

$$\mathcal{H}^+ = \bigcup_{D \in \mathcal{H}} D^+ \quad \mathcal{H}_{>0}^+ = \bigcup_{D \in \mathcal{H}} D_{>0}^+ \quad \mathcal{H}_{\geq c}^+ = \bigcup_{D \in \mathcal{H}} D_{\geq c}^+$$

The covering relations $r \in \mathcal{H}_{\geq c}^+$ are called *relevant* for \mathcal{H} . Observe, that relevancy depends on \mathcal{F}_O , since the precisions have been computed based on \mathcal{F}_O .

Quality Measures. The quality measures are used to rank the possible hypotheses with respect to the given observation. As we already introduced the precision of a single diagnosis we now will define the *quality* of a hypothesis, which can contain multiple diagnoses. The quality of a hypothesis provides a real value between 0 and 1 to describe the degree of accuracy with which the hypothesis \mathcal{H} can explain the given observation \mathcal{F}_O .

Definition 2.1 (Quality Measure) The *quality* $\varrho(\mathcal{H})$ of a hypothesis \mathcal{H} is given by

$$\varrho(\mathcal{H}) = \frac{\sum_{r \in \mathcal{H}_{\geq c}^+} \pi(e_r)}{|\mathcal{H}_{>0}^+| + |\mathcal{F}_{\mathcal{H},O}^{\text{isolated}}|}. \quad (3)$$

Notice that, in contrast to the precision, the quality measure does not evaluate a single diagnosis with respect to the transitively observed predictions, but assesses a hypothesis (containing possibly multiple diagnoses) on the basis of the transitively predicted and observed findings and the unexplained (isolated) findings.

We see that $\varrho(\mathcal{H}) \in [0, 1]$ for any hypothesis $\mathcal{H} \in \Omega_{\mathcal{H}}$: The lower bound 0 is obtained, if $\mathcal{H}_{\geq c}^+ = \emptyset$. The upper bound 1 is obtained,

if all predictions are fully observed, i.e. $\mathcal{H}_{\geq c}^+ = \mathcal{H}_{>0}^+$, and the set $\mathcal{F}_{\mathcal{H},O}^{\text{isolated}} = \emptyset$.

Example. For the covering relation given in Figure 2, the set

$$\mathcal{F}_O = \{ A_2 : v_2, A_3 : v_3, A_4 : v_4, A_5 : v_5, A_6 : v_6 \}$$

of findings, and the hypothesis $\mathcal{H} = \{D_1\}$, we obtain $\pi(D_2) = 1$, $\pi(D_3) = 1$ (with $c(D_2) = c(D_3) = 0.7$). Since we obtain $\mathcal{H}^+ = \{r_1, r_2, r_3\}$ for hypothesis \mathcal{H} we can calculate

$$\begin{aligned} \mathcal{H}_{\geq c}^+ &= \{r_1, r_2\}, \\ \mathcal{F}_{\mathcal{H},O}^{\text{isolated}} &= \{A_2 : v_2, A_3 : v_3\}. \end{aligned}$$

Up to now we presented the basic representation for set-covering models containing diagnoses and findings connected with covering relations. Of course this simple representation might not always meet the requirements of real world applications. Therefore we will shortly present knowledge extensions of set-covering models. In [2] we showed how to apply these extensions in an incremental way.

2.2 Extension by Similarities and Weights

Similarities between findings and weights for states provide significant knowledge extensions for set-covering models. In the following we will show how to include these enhancements into the quality measures given above.

Similarities. Consider a parameter A with the domain

$$\text{dom}(A) = \{ no, si, mi, hi \},$$

with the meanings normal (*no*), slightly increased (*si*), medium increased (*mi*), and heavily increased (*hi*), where $A : hi$ is predicted. We clearly see that the observation $A : mi$ deserves a better precision than the observation $A : no$. Nevertheless the simple quality measure considers both observations as unexplained findings and makes no difference between the similarities of the parameter values. For this reason we want to define *similarities* as an extension to set-covering models.

We define the similarity function

$$\text{sim} : \Omega_{\mathcal{V}} \times \Omega_{\mathcal{V}} \rightarrow [0, 1]$$

to capture the similarity between two values assigned to the same parameter. The value 0 means no similarity and the value 1 indicates two equal values. In cluster analysis problems this function is also called *distance function* (cf. [9]).

With similarities we need to adapt Equation (1) for computing the precision of findings.

$$\pi(A : v) = \text{sim}(\text{Val}_{\mathcal{H}}(A), \text{Val}_{\mathcal{F}_O}(A)),$$

where Val returns the value of a specified attribute contained in a specified set of states.

$$\text{Val} : 2^{\Omega_S} \times \Omega_A \rightarrow \Omega_{\mathcal{V}}.$$

If no special similarity is included in the model, then we get the simple quality measure by defining the *default similarity* $\text{sim}(v, v') = \delta_{v,v'}$, where $\delta_{v,v'} = 1$, if $v = v'$, and $\delta_{v,v'} = 0$, otherwise.

Weights. The introduction of weights for covered states is another common generalization of the basic covering model. Here we apply

a weight function $w : \Omega_S \rightarrow \mathbb{N}_+$, to emphasize that some states (findings and diagnoses) have a more significant pathological importance than other states.

When applying weights to the model we need to adapt Equation (2) which calculates the precision for a given diagnosis:

$$\pi(D) = \begin{cases} \frac{\sum_{r \in D_{\geq c}^+} w(e_r) \cdot \pi(e_r)}{\sum_{r \in D_{>0}^+} w(e_r)}, & \text{if } D_{>0}^+ \neq \emptyset, \\ 0, & \text{otherwise.} \end{cases}$$

Like for the precision of a diagnosis, we need to adapt Equation (3) to calculate the quality of a given hypothesis:

$$\varrho(\mathcal{H}) = \frac{\sum_{r \in \mathcal{H}_{\geq c}^+} w(e_r) \cdot \pi(e_r)}{\sum_{r \in \mathcal{H}_{>0}^+} w(e_r) + \sum_{F \in \mathcal{F}_{\mathcal{H}, \mathcal{O}}^{\text{isolated}}} w(F)}$$

If all states have the same weight, i.e., $w(S) = 1$ for all $S \in \Omega_S$, then the model reduces to the simple covering model.

In addition to similarities and weights we already have introduced uncertain covering relations and causal effect functions as possible extensions (cf. [2]).

3 Complex Covering Relations

In the previous section we introduced the basic set-covering model and extensions that allow for the refinement of set-covering knowledge build with basic covering relations. In this section we propose some further extensions of the representation, AND-, OR- and [MIN, MAX]-relations.

To keep the interpretation of covering models simple, we only allow these extensions for covering relations between diagnoses and (directly observable) findings.

3.1 Conjunction of Covering Relations

It is desirable to be able to represent conjunctions between covering relations. An AND-covering relation

$$D \rightarrow_{\text{AND}} \{F_1, \dots, F_n\}$$

denotes the characteristic, that all covering relations $D \rightarrow F_i$ have to be fulfilled simultaneously.




Figure 3. Covering relation $D \rightarrow_{\text{AND}} \{F_2, F_3\}$

Then the weights of the AND-connected findings F_i will only contribute to the precision of D if *all* of these findings are observed. If not all findings are observed, then D cannot explain the findings and we have to check if another diagnosis from the hypothesis can explain these observations. All remaining findings – so far unexplained – will be added to the set of isolated findings $\mathcal{F}_{\mathcal{H}, \mathcal{O}}^{\text{isolated}}$. This will decrease the quality measure for the current hypothesis, since $\mathcal{H}_{>c}^+$ will not contain relations covering the unexplained observations. Given an AND-covering relation of the form

$$r = D \rightarrow_{\text{AND}} \{F_1, \dots, F_n\}$$

we define for each $F_i \in \{F_1, \dots, F_n\}$:

$$\pi_r(F_i) = \begin{cases} \pi(F_i), & \text{if for all } F_j \in e_r : \pi(F_j) > 0 \\ 0, & \text{otherwise} \end{cases}$$

We try to explain all findings F_i with $\pi_r(F_i) = 0$ but $\pi(F_i) > 0$ by other diagnoses $D' \in \mathcal{H} \setminus \{D\}$. All remaining findings F_i , which cannot be explained by other diagnoses are added to $\mathcal{F}_{\mathcal{H}, \mathcal{O}}^{\text{isolated}}$.

Example. Assume that we have the covering model of Figure 3, where $c(D) = 0.5$, and we observe the set $\mathcal{F}_{\mathcal{O}} = \{F_1, F_2\}$. Then $\pi(F_3) = 0$, since F_3 is not in $\mathcal{F}_{\mathcal{O}}$. Therefore not all precisions of the AND-covered findings are greater than 0, and we define $\pi_r(F_2) = \pi_r(F_3) = 0$. We obtain $\mathcal{F}_{\mathcal{H}, \mathcal{O}}^{\text{isolated}} = \{F_2\}$ for hypothesis $\mathcal{H} = \{D\}$. Notice, that F_3 is not in $\mathcal{F}_{\mathcal{H}, \mathcal{O}}^{\text{isolated}}$, since it is not observed.

3.2 Disjunction of Covering Relations

We also can express alternative covering relations with disjunction. Here we can distinguish between inclusive (OR) and exclusive (XOR) disjunctions.

In Figure 4 we can see two different disjunctive covering relations for diagnosis D : in the left one the findings F_2, F_3 are connected with the OR-covering relation $D \rightarrow_{\text{OR}} \{F_2, F_3\}$, whereas at the right side the findings are connected with an XOR-covering relation $D \rightarrow_{\text{XOR}} \{F_2, F_3\}$. These OR/XOR-relations state, that only one




Figure 4. OR-/XOR-covering relations.

of the connected finding has to be observed to fulfill the relation. Of course we need to consider the different semantics in covering models. When computing the quality measures we have to take the following three cases into account:

1. If *none* of the predicted findings is observed, then nothing has to be done. The covering relations connected with the OR/XOR-condition cannot contribute to the quality measure of the parent state.

2. If *one* of the predicted findings is observed, then we simply cut all other states connected by OR/XOR-relations from the model. When computing the quality measure we only take the observed finding into account.
3. If *more than one* of the predicted findings are observed (e.g. $\{F_2, F_3\} \subseteq \mathcal{F}_O$), then we have to differentiate between OR and XOR relations. For both we take the finding with the maximal contribution; e.g. regarding the weighted precision

$$\pi_w(F) = \pi(F) \cdot w(F).$$

For OR-relations we simply ignore the remaining observations for assessing the quality. They will neither contribute to the quality of the hypothesis nor will they need to be explained by other diagnoses.

For XOR-relations the observations left over still have to be explained. Like for the AND-relations we try to explain them with the other diagnoses contained in the current hypothesis. All remaining findings, that cannot be explained by other diagnoses, are added to the set of isolated findings $\mathcal{F}_{\mathcal{H},\mathcal{O}}^{\text{isolated}}$.

We see that we carefully have to use OR/XOR-relations, because of their different interpretation of the observation. For example, multiple observations of one XOR-covering relation are taken negatively into account (i.e., they are assumed to be unexplained findings of the current hypothesis), whereas in ordinary OR-relations they will not contribute in any way.

As shown for AND-covering relations we also have to locally define the precision for OR/XOR-covered findings in context of the given diagnosis: Consider an OR-relation (analogous for XOR):

$$r = D \rightarrow_{\text{OR}} \{F_1, \dots, F_n\}.$$

We select a finding $F_{\text{max}} \in \{F_1, \dots, F_n\}$, such that $\pi_w(F_{\text{max}}) = \max(\pi_w(F_i), 1 \leq i \leq n)$. Then we say that

$$\pi_r(F_i) = \begin{cases} \pi(F_i), & \text{if } F_i = F_{\text{max}} \\ 0, & \text{otherwise.} \end{cases}$$

If there is more than one F_i with maximum weighted precision $\pi_w(F_i)$, then all but one (randomly selected) finding will set to the precision $\pi_r(F_i) = 0$.

When we compute the precision $\pi(D)$ of a diagnosis D , then the precisions of the findings F_i that are covered by an OR/XOR-covering relation contribute with the measure $\pi_r(F_i)$ and not with the usual precision measure $\pi(F_i)$.

For XOR-relations we have to explain the remaining findings by other diagnoses contained in the hypothesis or add them to $\mathcal{F}_{\mathcal{H},\mathcal{O}}^{\text{isolated}}$.

3.3 Cardinalities in Covering Relations

Another enrichment of the set-covering representation is the connection of covering relations by cardinality constraints. We express such cardinalities by [MIN, MAX]-covering relations. Consider the example in Figure 5. The covering relation between diagnosis D and the findings F_1, F_2, F_3, F_4 and F_5 means, that between 2 and 4 of the predicted findings have to be observed. We denote such relations by

$$r = D \rightarrow_{[2,4]} \{F_1, F_2, F_3, F_4, F_5\}.$$

When we interpret [MIN, MAX]-relations $r = D \rightarrow_{[\text{MIN}, \text{MAX}]} \mathcal{F}$, then we have to consider three possible cases for the number $k = |\mathcal{F} \cap \mathcal{F}_O|$ of relevant findings:




Figure 5. A [MIN, MAX]-covering relation.

1. If $k \in [\text{MIN}, \text{MAX}]$, then all findings in $\mathcal{F} \cap \mathcal{F}_O$ will contribute.
2. If $k > \text{MAX}$, then let $\mathcal{F}_{\text{max}} \subseteq \mathcal{F} \cap \mathcal{F}_O$ be the MAX findings with the maximum weighted precisions among the findings in \mathcal{F} (i.e. $|\mathcal{F}_{\text{max}}| = \text{MAX}$). We explain the findings in \mathcal{F}_{max} by D . Then we try to explain the findings in $(\mathcal{F} \cap \mathcal{F}_O) \setminus \mathcal{F}_{\text{max}}$ by other diagnoses also contained in the hypothesis. These findings $(\mathcal{F} \cap \mathcal{F}_O) \setminus \mathcal{F}_{\text{max}}$, which we cannot explain by other diagnoses $D' \in \mathcal{H} \setminus \{D\}$, are added to $\mathcal{F}_{\mathcal{H},\mathcal{O}}^{\text{isolated}}$.
3. If $k < \text{MIN}$, then we try to explain all findings in $\mathcal{F} \cap \mathcal{F}_O$ by other diagnoses $D' \in \mathcal{H} \setminus \{D\}$. Findings, which cannot be explained by other diagnoses, are added to $\mathcal{F}_{\mathcal{H},\mathcal{O}}^{\text{isolated}}$.

We integrate [MIN, MAX]-relations into set-covering models by locally defining the precision for findings connected by a [MIN, MAX]-relation $r = D \rightarrow_{[\text{MIN}, \text{MAX}]} \mathcal{F}$. Then we say that for each $F \in \mathcal{F}$:

$$\pi_r(F) = \begin{cases} 0, & \text{if } k < \text{MIN} \\ \text{or} & \text{if } k > \text{MAX} \wedge F \notin \mathcal{F}_{\text{max}} \\ \pi(F'), & \text{if } k \in [\text{MIN}, \text{MAX}] \\ \text{or} & \text{if } k > \text{MAX} \wedge F \in \mathcal{F}_{\text{max}} \end{cases}$$

where \mathcal{F}_{max} is again the set of the MAX findings with the best weighted precisions among the findings in \mathcal{F} .

When calculating the quality measure for a diagnosis or hypothesis we apply the precision $\pi_r(F)$ for all findings F connected by the relation r . Findings F with $\pi_r(F) = 0$ but $\pi(F) > 0$ need to be explained by other diagnoses contained in the hypothesis or will be added to $\mathcal{F}_{\mathcal{H},\mathcal{O}}^{\text{isolated}}$.

It is worth mentioning that ordinary covering relations for a diagnosis are following a similar concept, since we only will consider predicted findings that are also observed but not all predicted findings of the diagnosis. But as opposed to [MIN, MAX]-relations all observed predictions will contribute to the quality. In [MIN, MAX]-relations only MAX observed findings will contribute; more than MAX findings have to be explained by other diagnoses. In general, an ordinary covering model for a diagnosis D with n covered findings is comparable to a $[c(D) \cdot n, n]$ -relation connecting the n findings.

3.4 Bounded Covering Relations

The introduction of similarities for finding values is a useful knowledge extension. Nevertheless in some situations the expert wants to express that a relation is only fulfilled if a covered parameter is observed with exactly the predicted value, rather than a similar value. Therefore we supplement necessary covering relations, disjunctive,

conjunctive and constrained covering relations with the optional label *bounded*. We obtain the required behaviour by locally defining the *default similarity* measure for bounded relations:

$$\text{sim}(Val_{\mathcal{H}}(A), Val_{\mathcal{F}_O}(A)) = \delta_{Val_{\mathcal{H}}(A), Val_{\mathcal{F}_O}(A)}.$$

I.e., only if a parameter A is observed with the predicted value, then 1 is assigned to its precision.

4 Constraints for Hypothesis Generation

As mentioned in the introduction of Section 2, the problem of hypothesis generation is exponential, since for n diagnoses we need to consider about 2^n hypotheses in the worst case for an observation. In the following we want to sketch some heuristics to restrict the hypothesis space.

In a first step, we will filter all diagnoses $D \in \Omega_{\mathcal{D}}$, that are *relevant*, i.e. having the minimum precision. For this, we define the set of relevant diagnoses

$$\Omega_{\mathcal{D}}^{rel} = \{D \in \Omega_{\mathcal{D}} \mid \pi(D) \geq c(D)\}.$$

Then, only diagnoses $D \in \Omega_{\mathcal{D}}^{rel}$ will be taken into account, when generating hypotheses. Before describing concepts to shrink the set of hypotheses, we will define *generators* as a compact representation for sets of hypotheses, which had been introduced by Reggia et al. [1].

Definition 4.1 (Generator) A generator $\mathcal{G}_I = \{G_1, \dots, G_n\}$ consists of non-empty pairwise-disjoint subsets $G_i \subseteq \Omega_{\mathcal{D}}^{rel}$. The hypotheses $\mathcal{H}_{\mathcal{G}_I}$ generated by \mathcal{G}_I is defined as

$$\mathcal{H}_{\mathcal{G}_I} = \{\mathcal{H} \subseteq \Omega_{\mathcal{D}} \mid |\mathcal{H} \cap G_i| \leq 1, \text{ for all } 1 \leq i \leq n\}.$$

For $\mathcal{G}_I = \emptyset$, it holds that $\mathcal{H}_{\mathcal{G}_I} = \{\emptyset\}$. We can see, that $\mathcal{H}_{\mathcal{G}_I}$ is analogous to a cartesian set product.

For example, for the set-covering model defined in Figure 1 and $\mathcal{F}_O = \{temp : inc, skin : sweat, nose : red\}$, we obtain $\mathcal{G} = \{\mathcal{G}_1, \mathcal{G}_2\}$ with $\mathcal{G}_1 = \{\{cold\}, \{fever\}\}$ and $\mathcal{G}_2 = \{\{flu\}\}$. So we can compute $\mathcal{H}_{\mathcal{G}} = \{\emptyset, \{cold\}, \{fever\}, \{cold, fever\}, \{flu\}\}$ to be the set of interesting hypotheses.

A method for computing and updating generator sets is extensively described in [4]. Generators are used to efficiently generate hypotheses in an incremental manner: In a first step, sets of generators describing higher level diagnoses (concepts) are created. For hypotheses containing higher level diagnoses and having a high quality measure, we build sets of generators containing underlying specialized diagnoses and test them with their corresponding quality measure. In the following, we introduce two basic knowledge extension, that additionally shrink the space of generated hypotheses.

4.1 Exclusion Constraints

We can define *exclusion constraints* to filter diagnoses from the process of hypotheses generation. In general, two kinds of constraints are possible:

$$\neg(D \wedge F_1 \wedge \dots \wedge F_n)$$

If findings F_1, \dots, F_n are observed, then remove generated hypotheses, containing diagnosis D .

$$\neg(D_1 \wedge \dots \wedge D_m)$$

Remove generated hypotheses, containing all the diagnoses D_1, \dots, D_m at the same time.

Thus, we create hypotheses using generator sets and check each generated hypothesis against the available exclusion constraints. If one exclusion constraint evaluates true, the hypothesis is discarded.

It is worth noticing, that the modification of generator sets with respect to exclusion constraints yields a combinatorial size of generators and therefore is not reasonable. An evaluation of the generated hypotheses according to existing exclusion constraints has been proven to be more efficient.

4.2 Necessary Covering Relations

A stronger type of covering relations are *necessary covering relations*. A necessary covering relation between a diagnosis D and a finding F_1 means, that D *necessarily covers* F_1 and that F_1 always has to be observed if D is hypothesized. We depict a necessary covering relation with $D \xrightarrow{neg} F_1$ as shown in Figure 6.




Figure 6. Necessary Covering relation for a diagnosis D .

For applying necessary covering relations we introduce an adapted definition of the precision π_{neg} for each diagnosis $D \in \Omega_{\mathcal{D}}$:

$$\pi_{neg}(D) = \begin{cases} 0, & \text{if } \exists r \in \Omega_{\mathcal{R}} : r = D \xrightarrow{neg} F \text{ with} \\ & F \in \Omega_{\mathcal{F}} \wedge \pi(F) < \tau \\ \pi(D), & \text{otherwise} \end{cases}$$

where $\tau \in [0, 1]$ is a specified threshold, which defines when a finding is sufficiently observed (e.g. $\tau = 0.8$).

Therefore a diagnosis D does not propagate any contribution to its parent states until all necessarily covered findings are (sufficiently) observed. Consequently, D will not appear in any generator and thus will not be included in any hypothesis.

5 Conclusions and Future Work

After describing the basic structures of set-covering relations we have shown how to enrich the model with additional knowledge like similarities or weights. We also considered the computation of quality measures of these parts. Furthermore, we have shown representational extensions to the set-covering model to facilitate necessary, disjunctive, conjunctive or constrained covering relations. An important characteristic of all these extensions is the incrementality: some enhancements can be added to refine special aspects of the model but will not change its basic semantics; others are used to guide the process of candidate generation.

In the future we are planning to work on the following fields: Incremental development requires restructuring the model from time to time. We are currently working on restructuring methods for set-covering models that do not alter the basic semantics but improve the design of the diagnosis knowledge. In software engineering *refactoring* [10, 11] has been emerged as the corresponding method. In general we have to look at *validation techniques* for set-covering models besides simple case testing. Because of the special structure of the model we also have to consider static verification techniques for the set-covering representation. For a survey in this field we refer to [12, 13, 14, 15].

In this paper we presented a hand-driven development of set-covering models. But it seems to be possible to *learn* coarse models *automatically* from a small number of available cases. Later on these models should be refined by the developer with additional knowledge. With such a semi-automatic development step, the initial costs of knowledge acquisition can be reduced conveniently. Some work in this field has been done by Thompson et al. [16] and Wang et al. [17]. This step is not considered if we have a sufficiently large set of data, since then traditional machine learning methods (e.g. learning neural networks, learning Bayes networks) seem to be more appropriate.

ACKNOWLEDGEMENTS

The authors would like to thank Frank Puppe for his helpful suggestions and comments.

REFERENCES

- [1] James A. Reggia, Dana S. Nau, and Pearl Y. Wang. Diagnostic Expert Systems Based on a Set Covering Model. *Journal of Man-Machine Studies*, 19(5):437–460, 1983.
- [2] Joachim Baumeister, Dietmar Seipel, and Frank Puppe. Incremental Development of Diagnostic Set-Covering Models with Therapy Effects. In *Proceedings of the KI-2001 Workshop on Uncertainty in Artificial Intelligence*, Vienna, Austria, 2001.
- [3] Ramesh S. Patil, Peter Szolovits, and William B. Schwartz. Modeling Knowledge of the Patient in Acid-Base and Electrolyte Disorders. In: Szolovits, P. (Ed.). *Artificial Intelligence in Medicine*, Westview Press, Boulder, Colorado, 1982.
- [4] Yun Peng and James A. Reggia. *Abductive Inference Models for Diagnostic Problem-Solving*. Springer, Berlin, 1990.
- [5] Larry Eshelman. *Mole: A Knowledge-Acquisition Tool for Cover-and-Differentiate Systems*, pages 37–79. In: Sandra Marcus (ed.): *Automating Knowledge Acquisition for Expert Systems*. Kluwer Academic Publishers, 1988.
- [6] Luca Console, Luigi Portinale, Daniele Theseider Dupr, and Pietro Torasso. Combining Heuristic and Causal Reasoning in Diagnostic Problem Solving. In Jean-Marc David, Jean-Paul Krivine, and Reid Simmons, editors, *Second Generation Expert Systems*, pages 46–68. Springer, 1993.
- [7] William J. Long. Temporal Reasoning for Diagnosis in a Causal Probabilistic Knowledge Base. *Artificial Intelligence in Medicine*, 8(3):193–215, 1996.
- [8] Frank Puppe. Knowledge Reuse among Diagnostic Problem-Solving Methods in the Shell-Kit D3. *Int. J. Human-Computer Studies*, 49:627–649, 1998.
- [9] Jiawei Han and Micheline Kamber. *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers, San Mateo, California, 2000.
- [10] William F. Opdyke. *Refactoring Object-Oriented Frameworks*. PhD thesis, University of Illinois, Urbana-Champaign, IL, USA, 1992.
- [11] Martin Fowler. *Refactoring. Improving the Design of Existing Code*. Addison-Wesley, 1999.
- [12] Anca Vermesan and Frans Coenen. *Validation and Verification of Knowledge Based Systems. Theory, Tools and Practice*. Kluwer Academic Publishers, 1999.
- [13] Alun Preece. Building the Right System Right. In *Proceedings of KAW'98 Eleventh Workshop on Knowledge Acquisition, Modeling and Management*, 1998.
- [14] Frans Coenen and Trevor Bench-Capon. *Maintenance of Knowledge-Based Systems*. Academic Press, 1993.
- [15] Marc Ayel and Jean-Pierre Laurent. *Validation, Verification and Test of Knowledge-Based Systems*. Wiley, 1991.
- [16] Cynthia A. Thompson and Raymond J. Mooney. Inductive Learning for Abductive Diagnosis. In *Proceedings of the AAAI-94, Vol. 1*, pages 664–669, 1994.
- [17] Xue Z. Wang, M.L. Lu, and C. McGreavy. Learning Dynamic Fault Models based on a Fuzzy Set Covering Method. *Computers in Chemical Engineering*, 21:621–630, 1997.