Control and Cybernetics

vol. **39** (2010) No. 2

Medical diagnosis support by the application of associational cognitive maps*

by

Wojciech Froelich and Alicja Wakulicz-Deja

Institute of Computer Science, Silesian University, Sosnowiec, Poland

Abstract: The objective of the presented research is to construct a model of a patient's health that is based on the idea of cognitive map, a graphical knowledge-representation tool. The application of the proposed model for medical diagnosis is the practical goal of the research. Initially, we provide a brief review of the related works on medical decision support systems and cognitive maps. Afterwards, we sketch the general idea of the conceptual approach to the representation of medical knowledge and provide a new formulation of the medical diagnosis problem. Then, we define our model based on associational cognitive maps and show how it can be applied to diagnosis support. Due to the relative ease of understanding of cognitive map, the model can be easily interpreted and used, thereby making medical knowledge widely available through computer consultation systems. The application example presented is based on a relatively simple, real medical case.

Keywords: medical diagnosis, decision support, cognitive maps.

1. Introduction to medical decision support systems

Many medical decision support systems have been developed in recent years. Therefore, we decided to present only a general, descriptive introduction to this area. To relate our research to some of the existing systems, two general approaches to their design are presented. The first approach is the attribute-value approach represented, e.g., by traditional expert systems such as MYCIN (Buchanan and Shortliffe, 1984). These types of systems are usually equipped with a rule knowledge base and a corresponding inference mechanism. On the basis of rules, it is possible to classify new instances of medical observations by matching the set of observed symptoms (feature vector) to the conditional part of a rule and then to use logical deduction to achieve the diagnosis or construct a plan for the therapy. The diagnosis is usually understood as a crisp or approximate classification of symptoms. Note also that, in rule-based systems,

*Submitted: January 2009; Accepted: September 2009.

the association between diverse types of medical symptoms with respect to their consequences is achieved within the conditional part of a rule. This association is, in fact, interpreted as a possible co-occurrence of values (or subsets of values) of the conditional attributes. The time interval, during which this co-occurrence was experienced, is usually not considered. The inference mechanism relies on the relationship between the conditional and decision parts of the rules. However, the semantics of this dependence is, in many cases, intuitive and formally not clear e.g., it does not express any formal, well-defined, cause-and-effect relationship. The representation of causality (Pearl, 2000) is, in fact, left to human intuition. Note also that mapping between rules and medical natural language has some difficulties due to the incompatibility between the structures of information expressed in rule knowledge base and the semantics of medical language. The attempts to overcome some of the limitations of the traditional, rule-based systems can lead to the use of probabilistic models (e.g., Bayesian networks), non-classical logic, or hybrid solutions with supplemental annotation of rules in natural language. The second approach, which is somewhat the opposite of the attribute-based solutions approach, is the conceptual modeling approach (Chen, Fuller and Friedman, 2005). The main premise for the representation of conceptual knowledge is taken from the biology of the human mind and cognitive science. It was inspired by the biological and psychological observations that human knowledge is structured in the form of concepts and the associations among them. Another biologically-inspired hypothesis contends that the human learning process is, in fact, a process of relating a new observation to the existing structure of associations among concepts. The biological inspirations led initially to the general and intuitive understanding of such terms as 'concept' or 'associations among concepts,' which later were assigned different meanings in the computer science literature. Conceptual modeling is a research area that has attracted increasing interest in recent years. There are knowledge-representation methods that can be considered, in general, to exemplify the conceptual modeling approach, e.g., ontologies, semantic networks, formal concept analysis, or cognitive maps. We can mention the Unified Medical Language System (UMLS) (McCray, 2003), which contains a controlled medical vocabulary and a semantic network of biomedical concepts. Due to the lexical nature of ontologies, their automatic acquisition attempts concentrate mainly on mining medical texts (medical corpora). A visual map of medical terms extracted from different textual sources was presented in Lin et al. (2007). The introduction of dynamics to conceptual modeling in medicine starts with the analysis of temporal sequences of symptoms and interventions. Temporal reasoning in medicine has been investigated for almost three decades. Among many existing solutions, we can mention the creation of temporal clinical databases and the possibility of performing predictions on this basis. A review of temporal representation and reasoning in medicine can be found in Adlassnig et al. (2006). The relationship between temporal knowledge and medical natural language can be found in Zhou and Hripcsak (2007). In addition to temporal information, a medical expert system has to deal with uncertain knowledge (Straszecka, 2006). The uncertainty in medical science relates, e.g., to imprecise classification of symptoms or vague knowledge about temporal dependencies. The uncertainty is also involved in linguistic descriptions (Zhou and Hripcsak, 2007) provided by a patient or physician. The above problems require the application of appropriate methods for representing uncertain or approximate knowledge. Among many examples, we can mention a hybrid system with fuzzy reasoning that supports medical diagnosis proposed by Kuncheva (Kuncheva et al., 1993). There are also examples of expert systems (Steimann, 1997, 2001) developed on the basis of fuzzy sets theory that combine most of the above-mentioned features, i.e., conceptual knowledge representation and temporal and uncertain reasoning. Medical diagnosis, monitoring of the patient's condition, and planning treatment have been integrated in a sophisticated decision support system proposed by Zhou and Hripscak (Zhou and Hripcsak, 2007). One contemporary direction in research related to medical decision support systems is the application of cognitive maps (CMs), which may be a less-known conceptual approach to the representation of medical knowledge. The main advantage of CMs is their ability to incorporate and combine concepts, which have been heterogeneously and approximately described, with their mutual causal relations. As opposed to ontologies, which operate mostly on a symbolic level (i.e., using symbolic identifiers of concepts and relations), cognitive maps can be directly related to data. The structure of a CM is usually limited to modelling causal relations between concepts, but they can be easier understood and then applied while performing classification and reasoning on the basis of new observations. An introduction to the theory of cognitive maps and the relationship of existing CM-based systems to our model are presented in the following section.

2. Cognitive map as a knowledge representation - related works

Cognitive map is a biologically-inspired (Tolman, 1948; Okeef'e and Nadel, 1978), knowledge-representation method introduced initially in the social sciences by Axelrod (Axelrod, 1976). The initial theoretical model of CMs was targeted at the representation of cause-and-effect relationships among concepts described in natural language. The intuitive and easy-to-interpret representation of a CM is a directed graph, consisting of concept nodes and the cause-and-effect relations (directed arcs) among the concept nodes. The concepts, which represent sets of observations within the domain of interest, can be described in diverse ways. In a case, in which a particular concept has been observed, it is assumed as activated and can afterwards influence the activation of other concepts in the graph. This means that concepts can be activated by means of observation or by stimulation from their neighbors in the graph of the CM. The arcs indicate the directions of causal dependencies between source and target concepts. In the generic model of a CM, the causal dependencies can be positive (the respective arc is labeled with a "+" sign) or negative (the respective arc is labeled with a "–" sign). For positive causal dependencies, the activation of one concept causes the activation of another concept, and for negative causal dependencies, the activation of a concept deactivates the other concept. The dynamics of CMs is thus expressed by arcs that can be used to represent the behavior of the CMs over time. Note that the initial intuitive model of the CM used in social sciences and economics did not define precisely many important features of the CM, e.g., what does it precisely mean that a concept is active, or what does a causal relationship among concepts mean formally. The behaviors of generic CMs are not obvious in cases in which opposite types of influences (positive and negative) exert impact upon one concept. It is not clear whether or under what circumstances the compatible influences should be summed.

Among many extensions to the generic CM model, the most influential are the fuzzy cognitive maps (FCMs) introduced by Kosko (Kosko, 1986). The concepts in this model are mapped to the real-valued activation level from the continuous closed interval $[0, 1] \in R$, where 0 means no activation and 1 means full activation. In FCMs, the arcs are labeled with the weights $w_{ij} \in [-1, 1]$ that represent the strength of the causality, where -1 means fully-negative causal influence and $+1$ means fully-positive causal influence. The interpretation of the weights depends on the assumed definition of the causal relationship or the assumed method of learning (adaptive or evolutionary). While performing adaptive type of learning (the DHL-Differential Hebbian Learning), the so-called activation values, a_i, of all concepts are observed primarily at two consecutive time steps. For every concept, one can compute: $\Delta a_i = a_i(t) - a_i(t-1)$. Afterwards, the modification of weights is based on the rule, i.e., $w_{ij}(t+1) = w_{ij}(t) + c(t)[\Delta a_i \Delta a_j - w_{ij}(t)]$, where $c(t) = 0.1[1 - t/1.1q]$, is the learning coefficient and $q \in N$ is the parameter. While applying the evolutionary type of learning, the weights are modified by the genetic operators during the evolutionary process. The aim of the optimization is to achieve values of weights that will enable the simulation of the changes of activation of concepts observed in the environment. In this case, the computation of the fitness function involves the evaluation of the entire FCM by measuring the sum of the differences between observed and actual activity values for all concepts. The above learning methods impose some important limitations on the possible understanding of causality, especially in the light of the sophisticated analysis presented in Pearl (2000).

Having an activation state of concepts (with only a subset of them active at the same time), it is possible to compute iteratively the activation values of concepts in the future, using, e.g., the equation: $a_j(t+1) = \gamma(a_j + \sum_{i=1, i \neq j}^{n} w_{ij} a_i)$, where γ is the threshold function that serves to confine unbounded values to a strict range. The computation of the next states of nodes is interpreted as forward reasoning. Backward reasoning is also possible and leads to searching of possible causes on the basis of known effects. Other features, such as the additive combination of influences of arcs on nodes or the incorporation of time delays in the maps, are still being discussed among researchers. The problem of modeling time delays between concepts has been addressed in Park and Kim (1995). The proposal of HO-FCMs (Stach et al., 2006) overcomes the problem of modeling high-order dynamics by adding memory to the concept nodes. The drawback of the generic CM is its inability to deal with co-occurrence of concepts, such as those expressed by logical 'and' conditions. A survey of FCM extensions is presented in Aguilar (2005). Despite some problems with the initial model of the FCM, it has become very influential and has been used in many successful applications. In the last decade, the main stream of applications of cognitive maps are decision support systems (DSSs). FCMs have already been applied successfully in medical diagnosis support, including language impairment disorders. The solution presented in Stylios and Georgopoulos (2008) assumes that medical diagnosis is a single concept in CM. The reasoning is performed forward, towards the stated goal of achieving the activation of one of the expected diagnosis nodes. In spite of operating with a relatively large number of concepts compared to other known FCM applications, the system presented in Georgopoulos, Malandraki and Stylios (2003) considers only a one-step, forward-reasoning process. The decomposition of FCM to the group of local maps has been proposed in Stylios and Georgopoulos (2008). The application of fuzzy cognitive maps to medical diagnosis was also proposed by Innocent and John (2004). The integration of FCM and case-base reasoning (CBR) in medical decision support systems seems to be very promising approach (Georgopoulos and Stylios, 2005).

3. Conceptual approach to medical knowledge representation

The first problem during medical investigations is the necessity of rapid classification of symptoms. Symptoms are manifestations of a disease and are comprised of the observations reported by patients or made by doctors while examining patients. The identification of the underlying causes of symptoms is crucial and improves the chance of proper diagnosis of the disease and prescribing the correct treatment. In many cases, however, the directly-observed symptoms provide insufficient information for the proper identification of the underlying cause of the symptoms. Therefore, some medical examinations and tests must be conducted to provide additional data and information for use in providing the appropriate diagnosis and treatment. Sometimes, the initial treatment must be based on symptoms alone in order to treat severe symptoms, such as excessive body temperature. Unfortunately, when this is done without treating the underlying cause, the symptoms are very likely to recur. This shows the dynamics of the process of diagnosis and treatment that has to be embedded in the representation of medical knowledge.

If quantitative data are available as records of measurements (e.g., the measurement of the number of white cells in the patient's blood), there is usually a need to map them to meaningful categories (concepts), e.g., mapping the number of actual white cells in the blood to acceptable or dangerous levels of white cells. The classification of quantitative data to the qualitative medical concept depends frequently on diverse factors, e.g., factors related to the patient's age or the history of the disease. These factors can be seen as the parameters of the classification process. In some cases, a medical concept can be understood as a generalization of a subset of observations (measurements), but other concepts used in medical science may be based only on the doctor's direct observations and can be very complex. Therefore, the heterogeneous description of medical observations and concepts within a model of therapy is appreciated. The identification of medical concepts on the basis of available observations (or medical examinations) can be quite complex. The concepts can involve not only similarity-based clusters of observations but also temporal patterns that can be observed within data (e.g., a rapid temperature increase in a short period of time). In medical practice, the classification problem is left to the doctors and depends on their skills, specializations, and medical experience. On the other hand, in many cases, the construction of medical concepts and classification of medical data can be solved by well-known methods, e.g., by supervised learning of classifiers on the basis of examples (extracted from disease histories). Intuitive understanding of the relationships between concepts, the patient's condition, the electronic health record, and natural medical language has been sketched in Fig. 1. Medical concepts are usually mutually associated, and the knowledge of these associations should be represented and involved within the medical model. The associations among medical concepts include relations that are usually temporal, approximate and involve the dynamics of causes and symptoms with their mutual positive or negative influences. It is important to distinguish two types of concepts. One type consists of the concepts that are purely observational, i.e., they are related to symptoms or measurements. The other type consists of the concepts that represent medical interventions (e.g., prescribing drugs or surgery). Interventional concepts cannot be influenced by any other concepts within the model. Note that the concepts can be expressed by a doctor in modern language or in Latin to avoid the ambiguity of medical terms. We would like to refer to the example that has been given in Lin et al. (2007). If we look for terms related to the term "allergy," we will find a large set of possible words, e.g., "hypersensitivity," "allergens," "asthma," "cytokines," "eosinophils," and "occupational diseases." The set of associated terms can include synonyms, related diseases, and diverse types of therapies. The semantics of conceptual knowledge can be defined by the application of fuzzy relations (Tamir and Kandel, 1995). It is possible to consider the relations between concepts that can be expressed in form of symbolic terms, e.g., the cause-and-effect relationship can be expressed by: 'A causes B,' where A and B denote the names assigned to concepts, and the term 'causes' refers to the name of the relationship between them.

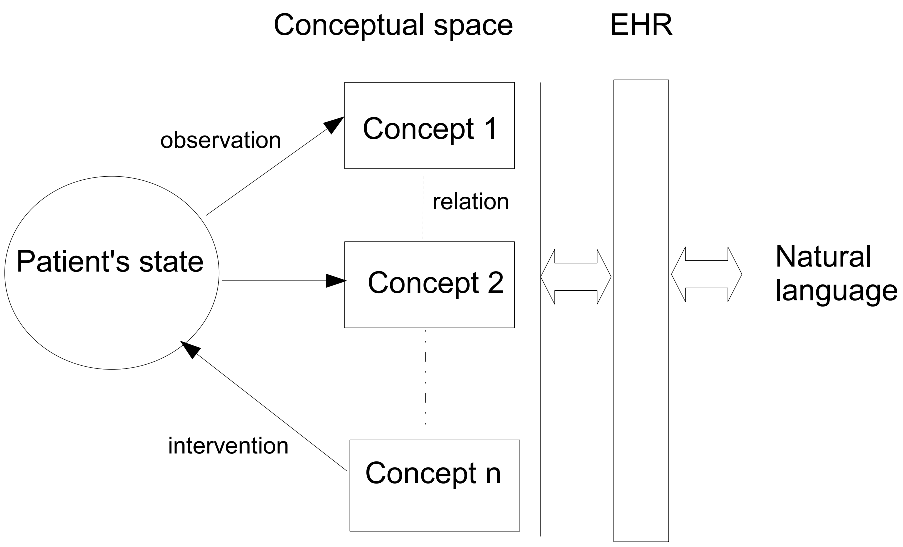

Figure 1. A general scheme of the conceptual approach to medical knowledge representation

In medical praxis, it would be appreciated if concepts and their mutual relationships were mapped to standardized terminology. This can be done by the application of electronic health records (EHRs). The initiative openEHR (Beale and Heard, 2007) provides schemes for storing medical data that can be captured within concepts and linked to the standard vocabulary. It would be the role of decision support system to interpret stored observations and provide the doctor with helpful suggestions that are compatible with the EHRs. Even a brief overview of the openEHR initiative is far beyond the scope of this paper. The complexity of mapping from the set of available terms to medical language depends on the assumed language (formal, controlled, or natural). In our opinion, the semantic interpretation of medical language should be placed within the conceptual space, where the construction of concepts and relations among them can be directly interpreted within medical data.

4. A problem of medical diagnosis

The correct diagnosis of a disease should lead to prescribing an appropriate treatment. The intent of treatment is to improve the health of the patient and eliminate the symptoms of the disease, but, in fact, treatment may lead to unexpected effects that require appropriate and rapid interventions. Continuous monitoring of the patient's health and immediate adaptation of treatment can be crucial for the success of the entire therapy. Medical therapy can be seen as a complex, dynamic process that should be anticipated and planned beforehand by doctors. As we have sketched, the problem of medical diagnosis requires the construction of a complex representation scheme to capture static and dynamic phenomena that can be encountered during the therapy. We would like to emphasize that the key feature of such knowledge representation is its capability of providing precise representations of diverse types of data and the associations within them. This involves:

1. qualitative and quantitative descriptions of data (symptoms or interventions) in the generalized form of concepts,
2. representation of relations between concepts,
3. the possibility of evidential (identification of causes on the basis of symptoms) and causal reasoning (simulation of effects of interventions),
4. identification of patterns within the conceptual representation scheme.

As opposed to other approaches, we take the position that a medical diagnosis is more frequently a problem of identification of a process than the classification of a static observation (e.g., feature vector) to the predefined class. To formalize this idea, we assume that we have available a set of raw medical data $O = S + I$, which consists of symptoms S and interventions I. Every observation or intervention will be classified to the concept $c \in C$. For every concept, we will assign its classification function:

$$\psi_i : O \times C \rightarrow [0,1] \in R, \tag{1}$$

where $[0,1]$ denotes a closed, continuous interval within R, and the subscript i refers to a particular concept. The function ψ_i can operate in approximate manner, reflecting by its value the uncertainty that can emerge during classification (e.g., due to incomplete observational data). Note that assuming $c \in C$ as a fuzzy set and interpreting ψ_i as a fuzzy membership function do not automatically imply the application of fuzzy logic while performing reasoning within the concept space.

Let us now define a mapping $a : C \rightarrow [0,1] \in R$ of every concept to the so-called activation value that represents its current state. Intuitively, $a(c_i)$ specifies whether a corresponding observation or intervention is experienced. On the basis of $\psi_i(o,c)$ we will compute the activation value $a(c_i) = \psi_i(o,c)$. Let $T = \{t_0, t_1, ..., t_n\}$ be an ordered set of time labels pointing to time moments. For the purpose of this paper, we assume discrete time flow, i.e., $t_i - t_{i-1} = \triangle t_i$ is a constant equal to 1. If the time has to be considered, the respective activation values of concepts will be complemented by the letter t. The flow of medical observations and interventions can be represented by a temporal sequence of concept activation values, i.e., by a subset $F \subset T \times C \times A$. The constraint is that the activation values are uniquely assigned to concepts given a particular time step. Note that any sequence from $f \in F$ can reflect in its trivial form a single static observation of one concept, e.g., $f = \{< t_s, c_1, a_v >\}$, where t_s is the time label, c_1 is the concept label, and a_v is the corresponding activation value.

Table 1. Activation of concepts in time

Time	**Activations of concepts**
t_1	...
..	...
t_s	a_1=0.3, a_2=0.6
t_o	a_4=0.4, a_8= 0.82, a_9=0.63
t_e	a_5= 0.92
...	...
t_n	...

We can also write this in the form: $a(t_s, c_1) = a_v$. Assuming that every concept label is mapped to the unique index, we have $a_1(t_s) = a_v$. The exemplary sequence of concept activations has been sketched in Table 1. In some cases, we would like to consider only concepts with binary activation values. In such case, the uncertainty involved in classification of symptoms is neglected. For every active concept c_i for which $a_i > a_{tre}$ holds, we will write simply $a_i(t)$ instead of specifying the exact activation value. Thus, for $a_{tre} = 0.5$, the sequence from Table 1 is reduced and can be denoted as: $\{a_2(t_s), a_8(t_0), a_9(t_0), a_5(t_e)\}$. Afterwards, we look for possible relations between concepts that govern their activations over time. Later, we will identify the concepts that co-occur (are active) at the same time as associated concepts. The activation of some concepts can cause (in its intuitive meaning) the activation of other concepts. The association and causality are two overlapping relations that will be assumed within concept space C. As will be shown, both of these relations can play an important role in analyzing the state of the patient's health.

Now, let us present our understanding of medical diagnosis. Let D be a set of known sequences of concept activations such that every $f_d \in D$ represents the dynamics of one particular disease, from the first negative symptoms at time t_s, through the set of medical interventions that lead to the last symptoms at time t_e and final recovery at $t_e + 1$. The symptoms and interventions are represented as activations of concepts. Let us assume that, at time period $[t_s, t_0]$, we observed a sequence $f_0 \in F$, where t_0 denotes a current time moment.

We define a medical diagnosis as the mapping $\delta : F \rightarrow D$ such that $\delta(f_0) = f_d$. The problem of defining similarities or patterns for sequences within D is left here as a challenge for future research. Note that the diagnostic sequence $f_d \in D$ should extend f_0 in time with the plan of therapy. Thus, the understanding of the patient's recovery is in accordance with our intuition and means the elimination of symptoms and the return of the patient's health to a stable state. For the purposes of this paper, we assume that the diagnostic sequences from D have to be provided by a physician.

5. Associational cognitive maps

For the targeted medical diagnosis support system, we propose an extended model of the cognitive map. Apart from representing causal dependencies between concepts, the proposed ACM (associational cognitive map) enables to represent a phenomenon that is well known in medical science, i.e., the co-occurrence of symptoms that are not linked by causal dependencies. We define the associational cognitive map as the triple:

Definition 1

$$ACM =< C, \alpha, \kappa >,$$

where C is the finite set of concepts, α is the association relation, and κ is the causal relation.

The understanding of concepts as approximated subsets of observations was presented in Section 4. Now, we explain the interpretation of relations α and κ. We assume that $\alpha \subseteq C \times C$ is a binary association relation that expresses the co-occurrence of concepts within C. The relation $c_i \propto c_j$ holds when:

$$\forall t \in T. (a_i (t) > a_{tre}) \Leftrightarrow (a_j (t) > a_{tre}), \tag{2}$$

where $a_{tre} \in [0,\ 1]$ is a pre-defined constant, the activation threshold. Note that, in comparison to the idea of association rules based on frequent itemsets (Agrawal and Imielinski, 1993), our association relation differs according to some essential features: 1) the association refers to concepts within the cognitive map; 2) the α relation is symmetric; and 3) the information about time (co-occurrence) is explicitly involved in its definition. As mentioned earlier, the generic model of the cognitive map is intended to play the role of a causal graph with edges representing the causal relations between concepts. However, the formal interpretation of causal relations in CM depends to a great extent on the approach used during setting or learning their weights. Due to existing controversies in understanding of causality (Pearl, 2000), in our ACM we prefer a generalized definition of the causal relationship that will be used within a cognitive map. Let $\kappa \subseteq C \times C$ be a binary relationship 'is-a-cause': $c_1 \kappa c_2$ denotes that c_1 is a cause of c_2. Since the relationship κ should reflect human causal knowledge (in a medical diagnosis problem), it should possess some of features experienced in common reasoning. Without doubt, it should also impose some constraints on the construction of the cognitive map. We assume that a concept cannot be the cause of itself. If a particular concept c_1 is a cause of c_2, then the reverse situation is not possible at the same time. For the two concepts c_1, c_2, we assume anti-reflexivity: $\forall c_1 \in C \backsim (c_1 \kappa c_1)$ and anti-symmetry: $\forall c_1, c_2 \in C\ c_1 \kappa\ c_2 \Rightarrow \backsim (c_2 \kappa c_1)$. Indirect cycles through other concepts are possible. Let us now consider the semantics of relation κ. We would like to propose its interpretation by applying the characteristic function μ, namely:

$$c_1 \kappa c_2 \Leftrightarrow \mu(c_1, c_2) \in Z, \tag{3}$$

where $\mu : C \times C \to Z$ is the characteristic function. The set Z is not restricted to numerical values and can involve linguistic utterances, e.g., “weak” or “strong". For example, if we assume $Z = [-1, 1] - \{0\}$ the characteristic function can be defined (such as in DHL algorithm) as $\mu(c_i, c_j) = \gamma(\sum_T \Delta a_i \Delta a_j)$, where γ is the normalization function and $\Delta a_i = a_i(t) - a_i(t-1)$ for any $t \in T$. However, the computation of $\mu(c_i, c_j)$ is not restricted to taking into account only the changes of activation values of two concepts in consecutive moments of time and can be improved by the BDA algorithm (Huerga, 2002), for example. Depending on the available data (type of environment, e.g., observational, temporal, or experimental), the attempts for proper construction of μ and thus interpretation of κ can be performed in diverse ways.

The more information that we are able to acquire from raw data, the more relevant to our requirements is the μ we could try to construct. Based on our assumptions, we now expect that the associational and causal relations will influence the activation of concepts within the set of sequences F. For a concept that is at some time active, we expect its association with or causal dependency upon other concepts. The knowledge about these dependencies can be stored in the proposed ACM model. On this basis, we will search for possible effects and causes for the groups of active concepts by forward and backward reasoning within the ACM. The intention is to use the current observations and knowledge stored in ACM to reconstruct possible concept activations in the future or in the past. The goal is to identify within F any matching diagnostic sequence stored in the set D.

6. Reasoning within ACM

From the point of view of medical diagnosis, evidential reasoning is essential. Before planning the therapy, it is necessary to find any hidden causes of the observed symptoms. In many cases, the association of current symptoms is the effect of a common cause. Therefore, there is a need to acquire as many symptoms as possible in order to identify the common cause correctly. This situation can recur while considering the diagnosis as the identification of a dynamic process. Note that due to the nature of biological processes, many diseases are manifested only partially by symptoms that are easily observable. In some cases, the hidden part of the disease can last for years. Since medical investigations and observations are usually incomplete, searches for causes that are based only on a set of observed symptoms can, in many cases, be insufficient. The general idea of the proposed solution is for both types of relations among symptoms, i.e., associational (representing co-occurrence) and causal, to be used during medical reasoning.

At time t_0, let us assume that the activation values of all concepts are computed on the basis of available observations (current symptoms) and the classification functions, i.e., $\forall c_i \in C (a_i(t_0) = \psi_i(o, c_i))$. Let $C_{obs} = \{c_i \in C : a_i(t_0) > a_{tre}\}$ (where $a_{tre} \in [0, 1]$ is the activation threshold) be the set of concepts (related to current symptoms) observed at t_0. Initially, we would like to use knowledge about associations stored in the ACM to extend the set C_{obs}. The goal of this α-extension is to take into account symptoms that usually co-occur with those represented in C_{obs} but that, in fact, were not observed. This extension will be done on the basis of the α relationship. For all $c_i \in C_{obs}$, we construct $C_\alpha = \{c_j \in C : c_j \alpha c_i\}$. Thus, we achieve at time t_0 the set $C_0 = C_{obs} + C_\alpha$ of concepts active within the ACM.

In fact, the cognitive map acts as a causal model, so the backward reasoning related to the direction of the causal relation κ leads to abduction. We look for the set $C_{b,\kappa} \subset C$ of possible causes of activations of concepts from C_0. For every $c_j \in C_0$, we search backward for the set of causes $C_{b,\kappa} \subset C$ such that, for every $c_i \in C_{b,\kappa}$, the relationship, $\forall c_j \in C_0 \; \exists \; c_i \in C_{b,\kappa} c_i \kappa c_j$, holds. The computation of activation values of concepts from $C_{b,\kappa}$ must be done thoughtfully. It can be considered that the intuitive, additive mechanism of causality used within cognitive maps can be also applied during backward reasoning. We assume that the strength of the causality measured by the function μ remains the same while performing reasoning in reverse direction. Under those assumptions, we can apply the following formula to compute the activation values: $\forall c_i \in C_{b,\kappa} (a_i = \gamma(\sum_j (a_j \mu (c_i, c_j) \;))$, where γ serves to reduce the unbounded values to a strict [0, 1] range. The next step is to extend the set $C_{b,\kappa}$ on the basis of the α relationship, in the same way as this was done for C_0. Finally, we obtain $C_b = C_{b,\kappa} + C_\alpha$, which constitutes the set of concepts anticipated as active at a previous time step (at time $t_0 - 1$). The next step of backward reasoning can be made the same way, assuming C_b plays the role of C_0.

The forward (in accordance with the direction of the relationship κ) reasoning within the ACM is straightforward. In fact, it is a simulation of possible effects in the future, i.e., the prediction of the possible progress of the disease. For every $c_i \in C_0$, we search for the set of effects $C_{f,\kappa} \subset C$ such that $\forall c_i \in C_0 \; \exists \; c_j \in C_{f,\kappa} c_i \kappa c_j$. Thus, the set $C_{f,\kappa}$ specifies all possible effects of C_0. The activation level of concepts from $C_{f,\kappa}$ can be computed from $\forall c_j \in C_{f,\kappa} (a_j = \gamma(a_j + \sum_i (a_i \mu \; (c_i, c_j) \;))$. Now, it is possible to extend the set $C_{f,\kappa}$ on the basis of the relationship α. We will obtain the set of active concepts $C_f = C_{f,\kappa} + C_\alpha$.

Above, we have shown the procedures for one-step backward and forward reasoning processes based on the proposed ACM model. Note that, as opposed to rule-based systems or probabilistic models based on acyclic graphs, there are no explicit stopping conditions (a kind of goal concepts) for the reasoning process within cognitive maps. The reasoning stops due to the limited scope of the ACM (limited knowledge) or when the maximum number of steps of interest to the user have been performed. Also, the backward reasoning stops when all possible causes of symptoms are found. Performing backward reasoning, starting from the set of currently observed symptoms, leads to the reconstruction of a possible temporal sequence of concept activations $f_r = \{C_b(t_s), \ldots, C_b(t_0 - 1), C_0\}$. After such reconstruction, there is a need to make a diagnosis of f_r within the set of diagnostic sequences D. The found sequence $f_{r,d} = \delta(f_r)$ will include the concepts corresponding to future interventions and consequences of their applications. In our opinion, due to the nature of medical diagnoses, the solely automatic classification of f_r is not recommended and must involve additional medical knowledge provided by a physician. The diagnostic sequence from D and the dynamics of concept activations within ACM, which correspond to possible medical interventions, should be reviewed and assessed by a medical doctor.

7. Case study

We have tested the ACM using a real, but simple, medical case. In the case we have chosen, a patient came to the eye specialist and reported that after few days of playing tennis he had the symptom of corneal ulcer. The patient had reported also allergic reactions to various allergens in the past. He did not have a fever. The doctor suspected three probable causes: an injury after the tennis match, an allergy based on previous knowledge of the patient, and a bacterial infection of the eye. The first proposed treatment was related to the suspected injury, i.e., a simple surgical removal of the ulcer. Because the patient denied having sustained an injury to his eye, he was given *alergocon*, an anti-allergy drug in the form of eye drops and *dicortinef*, a mixture of antibiotics also in the form of eye drops. Unfortunately, all the diagnoses and treatments were wrong. After three weeks, there was no improvement in the patient's symptoms. Where was the mistake? The physician did not examine the patient's state. Subsequently, the patient went to another doctor, who conducted a bacterial examination of the patient's throat with negative results. The examination was repeated a few times without identifying any bacteria in the throat, in spite of the fact that visual examination clearly suggested a bacterial infection in the patient's throat. Even blood analyses detected no bacteria that could be causing the sore throat. In spite of the disappointing investigations, this time the diagnosis was right. There was a body-wide bacterial infection that had caused both the sore throat that were reported few days later and the corneal ulcer. The lack of fever also suggested a bacterial infection as opposed to a viral infection, which usually causes high fever. The first proposed medicine was *biotraxon*, an antibiotic from the third generation of the *cephalosporin* group of antibiotics, but, unfortunately, this was another mistake. The patient had a severe allergic reaction to the antibiotic, including high fever, muscle pain, and rigor. Note that the patient had reported his past allergies, but this information was neglected by the physician. Surprisingly, the physician did not want to accept the fact he had made a mistake and that the antibiotic had caused the patient's allergic reaction. The third doctor made a simple decision and changed the prescribed medicine to *xorimax*, an antibiotic in the second generation group of *cephalosporin*, which was known to be an excellent antibiotic that could easily be tolerated by most patients. This turned out to be a very good choice, taking into account lack

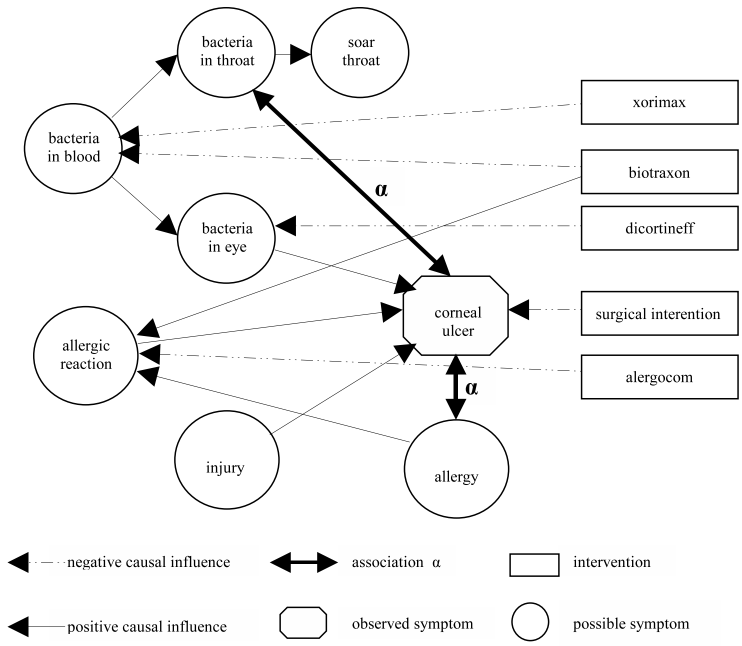

Figure 2. Associational cognitive map as a knowledge base used for medical diagnosis

of identification of particular bacteria. The medical case described above was simulated in our system. The set of concepts C was divided into symptoms S and interventions I. The set S consists of M (quantitative measurements) and Q (qualitative observations). We assigned symbolic labels to all considered concepts, in particular: $M = \{$ m_1 - 'bacteria in throat', m_2 - 'bacteria in eye', and m_3 - 'bacteria in blood'$\}$; $Q = \{q_1$ – 'sore throat', q_2 – 'allergic reaction', q_3 – 'injury', q_4 – 'corneal ulcer', and q_5 – 'allergy'$\}$; and I $= \{i_1$ – 'xorimax', i_2 – 'biotraxon', i_3 – 'dicortineff' (acetil cefuroxim), i_4 – 'surgical intervention', and i_5 – 'alergocon'$\}$. The drugs (of the given trade names) are available in pharmacies in Poland.

The ACM used for the considered medical case is presented in Fig. 2. In order to show the advantage of our associational extension in comparison to traditional (used only as a causal graph) cognitive maps, we performed the simulation of the previously described medical history. The set of initially reported symptoms were $C_{obs} = \{q_4,\ q_5\}$. The doctor expected three activation sequences of concepts. The first one was $f_1 = \{q_3(t_0 - 1), q_4(t_0)\}$. The diagnosis $\delta(f_1)$ led to the diagnostic sequence $f_{1,d} = \{q_3(t_0 - 1), q_4(t_0), i_4(t_0 + 1)\}$. The activation of the concept i_4 - 'surgical intervention' led to direct deactivation of the concept q_4 - 'corneal ulcer' and recovery. The reported association between

q_4 and q_5 was not considered by the doctor. The next analyzed alternative led to the diagnostic sequence $f_{2,d} = \{q_2(t_0 - 1), q_4(t_0), q_5(t_0), i_5(t_0 + 1)\}$. Due to consideration of the known association $q_4 \alpha q_5$, the activation value of q_2 was high. The application of i_5 – 'alergocon' led to the recovery of the patient. The third considered diagnostic sequence $f_{3,d} = \{m_2(t_0 - 1), q_4(t_0), q_5(t_0), i_3(t_0 + 1)\}$ suspected the possible isolated infection in eye with the prescription of *dicortineff* eye drops. Note that the found causes of 'corneal ulcer' are not associated in our ACM, therefore every diagnostic sequence of $f_{1,d}, f_{2,d}, f_{3,d}$ was considered separately, as has been done by physicians. Assuming we do not have additional knowledge on the features of the disease, it could be really difficult to find any reason to search backward within the ACM for all possible causes of allergy or bacteria in the eye. There is also no reason for a physician to perform numerous additional investigations, because, as in our case, some of these investigations do not provide any helpful information. Another problem is the growing specialization of medical doctors (e.g., eye specialists do not usually investigate throat problems). On the basis of analysis of other medical cases not described in this paper, we would like to stress the fact that the lack of the associational knowledge is in many cases the main cause of diagnostic failure and leads to serious health consequences. Note that the ACM sketched in Fig. 2 is substantially limited and without doubt does not involve many possible causes and associations. The above-described simple case is only used as an illustration of a possibly complex problem. Now let us consider that the physician takes into account the additional information on possible association between q_4 - 'corneal ulcer' and m_1 - 'bacteria in throat.' The α-extension of C_{obs} is performed and the set $C_\alpha(t_0) = \{m_1\}$. After activating concepts from $C_0 = C_{obs} + C_\alpha = \{q_4, q_5, m_1\}$, backward reasoning within ACM was performed until the corresponding sequence of concepts was found. The diagnoses were performed and led to the right diagnostic sequence $f_{4,d} = \{m_3(t_0 - 2), m_2(t_0 - 1), q_4(t_0), q_5(t_0), m_1(t_0), q_1(t_0 + 1), i_1(t_0 + 1)\}$. Note that the alternative sequence $f_d = \{m_3(t_0 - 2), m_2(t_0 - 1), q_4(t_0), q_5(t_0), m_1(t_0), q_1(t_0 + 1), i_2(t_0 + 1)\}$ with prescription of *biotraxon* was eliminated due to the high risk of allergic reaction and leading to an unstable state of the ACM. Unfortunately, as you may recall, this alternative sequence was applied by one of the doctors and caused a really dangerous health situation for the patient.

In the preliminary version of the implemented application, we assumed that the observed symptoms were given by a physician. The physician decides in a crisp way whether a particular symptom occurs or not. The physician uses natural language, e.g. he writes "This patient has got corneal ulcer". The above assumption simplified the implementation in different ways. First, the classification functions psi_i were constructed on the basis of simple matching of string patterns between input text given by a doctor and corresponding labels of concepts stored in knowledge base. Thus, the results of the classification of symptoms are binary i.e. $\forall c_i \in C_0 \; a_i = 0 \vee a_i = 1$, the threshold $a_{tre} = 1$. The consequence of the above simplification is such that the set of concepts C can be represented by the simple table in relational database with two fields, the name of symptoms and the binary classification result, respectively. Note that the construction of approximate medical concept on the basis of raw data is a sophisticated problem involving also medical domain knowledge. In the preliminary application we did not addressed this problem (the concepts were crisp and the measurement data were not used). We prepared only the application for the future use of the approximated classification method by involving the thresholding process.

In order to test the proposed reasoning process we filled the table of concepts with symptoms of several diseases including those presented in this paper. The pairs of concepts belonging to the relations (α, κ) are also stored in the table within relational database. Every pair of concepts is complemented by the corresponding weight that reflects the strength of the relation, i.e. the set $Z = < -1, 1 >$ that is used for the computation of μ is a continuous interval. For the experiments we used the weights given by physicians. We also decided to store the state of the reasoning process (and the data for thresholding) as tables in relational database. On the one hand, such solution decreased substantially the speed of reasoning process, on the other hand it enabled testing in a convenient way. The result of the reasoning process is the sequence f. The final diagnosis $\delta : F \to D$ was made in a crisp way. The reconstructed sequence f was simply matched to the set of sequences stored in D. The content of D was prepared using available medical knowledge. We are aware of the fact that the above sketched assumptions i.e. the mixture of crisp concepts and approximate reasoning, led without doubt to the partial solution of the addressed problem. The other issue is whether all classes of medical diagnostic tasks can be solved using the proposed solution. Without doubt, incorporation of more uncertainty in the classification an reasoning can decrease the applicability of the proposed method. The answer to this question seems to be impossible without acquisition of more medical data. Unfortunately, at this stage of research, the scalability of the proposed method cannot be practically verified. The main problem with the analysis of larger medical problems is the availability of reliable medical knowledge. It should involve concepts, relations and the diagnostic sequences, and therefore is hardly available. Development of automatic acquisition method of the required knowledge from textual and numerical data sources might be considered.

8. Conclusions

The presented solution has been motivated by some of the requirements imposed by the targeted application: the temporal and causal association of disease symptoms that seem to be crucial for the right medical diagnosis. We have proposed the associational extension of cognitive maps and a corresponding model called ACM. In the ACM, we introduced the possibility of associating co–occurring concepts that can be identified as a kind of higher-level concept within a cognitive map. We showed how the reasoning within the ACM is performed. We have also sketched the exemplary illustrative problem of medical diagnosis and its simulation using a preliminary version of the targeted application. Due to the complexity of medical knowledge, we have decided to impose several simplifications to the presented analysis. Given the imposed simplifications and problems we have mentioned, it is apparent that some of them are challenges for future research.

Acknowledgments
A part of the presented research on cognitive maps has been supported by the grant N N516 368334 from the Polish Ministry of Science and Higher Education.

References

Adlassnig, K.-P., Combi, C., Das, A.K., Keravnou, E.T. and Pozzi, G. (2006) Temporal representation and reasoning in medicine: Research directions and challenges. *Artificial Intelligence in Medicine* **38** (2), 101–113.

Agrawal, R. and Imielinski, T. (1993) Mining association rules between sets of items in large databases. In: *Proc. of 1993 ACM-SIGMOD Int. Conference on Management of Data.* ACM Press, 207–216.

Aguilar, J. (2005) Survey about Fuzzy Cognitive Maps Papers. *International Journal of Computational Cognition* **3** (2).

Axelrod, R. (1976) *Structure of Decision – The Cognitive Maps of Political Elites.* Princeton University Press.

Beale, T. and Heard, S., eds. (2007) Open EHR, architecture overview. *http://www.openehr.org/svn/specification/TAGS/Release-1.0.1/ publishing/archite cture/overview.pdf.*

Buchanan, B.G. and Shortliffe, E.H. (1984) *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project.* AAAI Press.

Chen, H., Fuller, S.S. and Friedman, C.P. eds. (2005) *Medical Informatics: Knowledge Management and Data Mining in Biomedicine.* Integrated Series in Information Systems, Springer.

Georgopoulos, V.C., Malandraki, G.A. and Stylios, C.D. (2003) A fuzzy cognitive map approach to differential diagnosis of specific language impairment. *Artificial Intelligence in Medicine*, **29**(3), 261–278.

Georgopoulos, V.C. and Stylios, C.D. (2005) Augmented Fuzzy Cognitive Maps Supplemented with Case Base Reasoning for Advanced Medical Decision Support. In: M. Nikravesh, L.A. Zadeh and J. Kacprzyk, eds., *Soft Computing for Information Processing.* Springer, 389–398.

Huerga, A.V. (2002) A balanced differential learning algorithm in fuzzy cognitive maps. In: *Proceedings of the 16th International Workshop on Qualitative Reasoning.*

Innocent, P.R. and John, R.I. (2004) Computer aided fuzzy medical diagnosis. *Inf. Sci.* **162** (2), 81–104.

Kosko, B. (1986) Fuzzy Cognitive Maps. *International Journal of Man-Machine Studies* **24**(1), 65–75.

Kuncheva, L., Zlatev, R.Z., Neshkova, S.N. and Gamper, J. (1993) A combination Scheme of Artificial Intelligence and Fuzzy Pattern Recognition in Medical Diagnosis. In: *FLAI '93: Proceedings of the 8th Austrian Artificial Intelligence Conference on Fuzzy Logic in Artificial Intelligence*. Springer-Verlag, London, UK, 157–164.

Lin, F., Ying, H., MacArthur, R.D., Cohn, J.A., Barth-Jones, D.C. and Crane, L.R. (2007) Decision making in fuzzy discrete event systems. *Inf. Sci.*, **177** (18), 3749–3763.

McCray, A.T. (2003) An Upper Level Ontology for the Biomedical Domain. *Comp. Funct. Genom.* **1** (4), 80–84.

Okeef'e, J. and Nadel, L. (1978) *The Hippocampus as Cognitive Map*. Clarendon Press, Oxford.

Park, K.S. and Kim, S.H. (1995) Fuzzy cognitive maps considering time relationships. *International Journal Human-Computer Studies* **42**, 157–168.

Pearl, J. (2000) *Causality, Model Reasoning and Inference*. Cambridge University Press.

Stach, W., Kurgan, L.A., Pedrycz, W. and Reformat, M. (2006) Higher-order Fuzzy Cognitive Maps. In: *NAFIPS 2006 – 2006 Annual Meeting of the North American Fuzzy Information Processing Society*, 166–171.

Steimann, F. (1997) Fuzzy set theory in medicine. *Artificial Intelligence in Medicine*, **11**(1), 1–7.

Steimann, F. (2001) Survey on the use and usefulness of fuzzy sets in medical AI. *Artificial Intelligence in Medicine* **21** (1), 131–137.

Straszecka, E. (2006) Combining uncertainty and imprecision in models of medical diagnosis. *Inf. Sci.*, **176**(20), 3026–3059.

Stylios, C.D. and Georgopoulos, V.C. (2008) Fuzzy Cognitive Maps Structure for Medical Decision Support Systems. *Studies in Fuzziness and Soft Computing* **218**, 151–174.

Tamir, D.E. and Kandel, A. (1995) Fuzzy Semantic Analysis and Formal Specification of Conceptual Knowledge. *Inf. Sci.* **82** (3-4), 181–196.

Tolman, E.C. (1948) Cognitive Maps in Rats and Men. *The Psychological Review* **55** (4), 189–208.

Zhou, L. and Hripcsak, G. (2007) Temporal reasoning with medical data - A review with emphasis on medical natural language processing. *Journal of Biomedical Informatics* **40** (2), 183–202.