

Symbolical Reasoning about Numerical Data: A Hybrid Approach

CHRISTOPH S. HERRMANN

TH Darmstadt, FB Informatik, FG Intellektik, Alexanderstr. 10, 64283 Darmstadt, Germany

herrmann@informatik.th-darmstadt.de

Abstract. By combining methods from artificial intelligence *and* signal analysis, we have developed a hybrid system for medical diagnosis. The core of the system is a fuzzy expert system with a dual source knowledge base. Two sets of rules are acquired, automatically from given examples and indirectly formulated by the physician. A fuzzy neural network serves to learn from sample data and allows to extract fuzzy rules for the knowledge base. A complex signal transformation preprocesses the digital data a priori to the symbolic representation. Results demonstrate the high accuracy of the system in the field of diagnosing electroencephalograms where it outperforms the visual diagnosis by a human expert for some phenomena.

Keywords: expert systems, fuzzy logic, hybrid systems, medical diagnosis, neural networks

1. Introduction

The emerging need to evaluate a vast variety of electronic patient data, often in the form of multi-dimensional signals, raises the demand for automated diagnosis methods. We have modelled the cognitive process of diagnosing an electroencephalogram (EEG) (see Fig. 1) in a hybrid system incorporating methods from Artificial Intelligence (AI). To optimally relate the automatic diagnosis to the human model, we combined an AI-system with a signal analysis approach for preprocessing.

At our institute of intellectics, the science that incorporates research about AI and cognition [1], we have modelled the cognitive process of diagnosing EEGs in a hybrid system. Here, we will describe a novel hybrid system consisting of a fuzzy expert system for rule-based reasoning with a fuzzy neural network for the automatic acquisition of knowledge in addition to the indirectly acquired knowledge from an expert (Section 3). The automatic acquisition of rules by the network is implemented in parallel with the classical formulation of expert rules. All components of the system are based on a fuzzy representation (Subsection 3.2), serving as an interface notation between the components and making a fuzzification of input data necessary. In Subsection 3.3, a very effective

mapping technique will be introduced, transforming fuzzy variables into a neural representation.

To visualize the processing of real medical data in such a system, we will demonstrate the diagnosis of electroencephalograms (Section 2). Other types of data will also be appropriate for diagnosis in this system. We will describe the system without relating it too much to the specific application. Also, related topics are proposed that offer similar processing to visualize other possible applications of the system.

With our results (Section 4) we show how the automatic system outperforms the accuracy of a human expert in one topic of EEG diagnosis. In addition, we found out that the automatic acquisition of rules by learning from examples generates not only quantitatively more rules but also rules with qualitatively higher precision than the indirect method of formulating expert rules. A pruning algorithm is proposed to reduce the rule base in case of overfitting due to too many rules of negligible rule strengths. Also, we will give a brief description of how to extend the capabilities of the system by exchanging existing components or integrating new ones.

Finally, a discussion will summarize the pros and cons of our method in Section 5 and some Internet pointers are provided in Section 6 to access parts of the hybrid system.




Figure 1. Ten seconds of EEG, containing three bulbus artifacts (eye movements) in seconds 0, 7.5 and 8.

2. Electroencephalography

Electroencephalography is predestinated to automatic analysis, since the data is registered electronically and need not be transformed into electronic format any more. On the contrary, for human visual inspection EEGs would need to be visualized on paper or on a screen. In a routine EEG recording approximately 16 to 20 traces of EEG are registered at certain locations on the skull. The trace of an EEG consists of so-called graphoelements, the smallest meaningful graphical elements, and is sometimes superimposed by artifacts. The main activity in healthy persons is the basic rhythm being a sinusoidal waveform of spindle-shaped modulated amplitude in the range of eight to twelve Hertz (cycles per second). Pathological graphoelements are waves that differ from this continuous activity and mostly occur in form of slower frequency components (slow waves) or spontaneously high amplitudes (spikes). Artifacts are also deviations from the basic rhythm but do not result from brain electric potentials but from other physiological signals (e.g., eye movements or muscle contraction).

The diagnosis of electroencephalograms evaluates the basic rhythm and interprets all occurring graphoelements and artifacts. The basic rhythm is described in form of its amplitude and frequency and the deviation over time of the two. Physicians use certain attributes to address different frequency ranges of EEG phenomena as shown in Table 1. Similarly the terminology to describe the amplitude of an EEG component is somewhat

Table 1. Overview of the medical frequency bands.

Band name	Frequency	Possible interpretation
Delta(δ)	<4 Hz	Artifacts and pathologically slow activity
Theta(θ)	ca. 4–8 Hz	Usually pathologically slow activity
Alpha(α)	ca. 8–12 Hz	Basic rhythm
Beta(β)	>12 Hz	Cognitive processes

Table 2. Medical definition of verbal attributes to describe EEG amplitudes.

Verbal description	Amplitude
Low	<20 μ V
Medium	20 to 60 μ V
High	>60 μ V

standardized even though attributes from everyday language are used (see Table 2).

Graphoelements have to be differentiated from artifacts of similar wavelike shape. This is sometimes possible only with knowledge of the location of the graphoelement *and* knowledge about the activity in other locations. Certain graphoelements are reliable indicators for diseases but have different meanings depending on the age and mental condition of the patient. For the above reasons, a number of requirements have to be met by our system that may also be important for different domains:

- The processing of EEGs needs to be robust against noisy raw data.
- Reliable features need to be extracted from preprocessed data.
- Numerical data has to be represented symbolically in order to do symbolical reasoning.
- A suitable symbolical representation for the description of signals *and* non-numerical EEG parameters has to be supplied.
- The vagueness of medical rules must be regarded.
- An inferencing mechanism for vague logical reasoning is needed.

3. Why Use a Hybrid?

Medical data that is measured and stored electronically is very well suited for automatic processing since it need not be converted to an electronic representation any more. This is not only true for the demonstrated EEG example but also for other electrophysiologic time

series (ECG, EMG, etc.) and medical image data (CT, MRI, PET, etc.).

The evaluation of an EEG is based on time series of numerical values. To deal with numerical problems of this sort, various signal analysis approaches are applied in medicine to process the electronically acquired data. Instructive insights were taken from transforming the time domain EEG into its frequency domain representation by the Fourier transform [2]. But, all signal analysis methods transform the data into a new representation or calculate new parameters. They are not able to automatically reason about the newly generated data. Whenever computerized *reasoning* is to be carried out on the resulting data, a transformation into a *symbolic* representation has to be performed prior to the actual diagnosis process. In our hybrid system, we show a new method of preprocessing and symbolizing medical data for subsequent reasoning with techniques from artificial intelligence.

Most popular among AI methods in medicine are knowledge based systems [3], modeling the diagnostic behaviour of experts. An explanation component makes a generated diagnosis transparent to the user. A variety of such expert systems is being used in everyday practice of physicians since Shortliffe introduced MYCIN [4], an expert system designed to diagnose infections of the human blood. One of the greatest difficulties in designing a convenient expert system is the acquisition of the knowledge base, since rule-based systems are usually unable to learn. Our hybrid system introduces a new approach where a dual source knowledge base is generated by indirect and automatic knowledge acquisition.

Neural networks have also made their way into diagnosis. They are able to learn relationships between data sets by simply having sample data fed into their input and output layers. In the field of pattern recognition in medical data, neural network based approaches have led to quite remarkable results, for example in processing EEG traces [5]. But, in the domain of medical reasoning, it is strictly required that a diagnosis is understandable to the physician (e.g., via an explanation component). This requirement is not met by neural networks, since their processing can not retrospectively be examined. Hence, we are not using them for the diagnosing task, but for automatic rule generation.

Fuzzy logic [6] also makes its appearance in medicine, dealing with the vagueness of verbal expressions. Terms like *many*, *few* or *probably* are hard to model with conventional logic. The linguistic variables

offered by fuzzy representations allow pseudo-verbal descriptions close to natural human expressions.

Since in humans the complex processing of multi-modal information is based on applying multiple processes, researchers now try to apply hybrid architectures to achieve the same [7]. Single approaches have brought up remarkable results in AI but they are not yet able to deal with more complex tasks¹. Thus, there are three main reasons to apply hybrid techniques: *technique enhancement*, *multiplicity of application tasks* and *multi-functionality* [7]. *Technique enhancement* integrates multiple methods in order to overcome the limitations of the single approaches. A technique that has a weakness in a certain aspect is combined with a technique that has its strength in this aspect. The *multiplicity of application tasks* describes the fact that a problem consists of many different subproblems that cannot all be solved by the same method. Thus, different methods are to be applied—each addressing one of the subproblems. If a single architecture is capable of performing multiple information processing tasks it will be considered *multi-functional*. Based on these factors a classification scheme for hybrid systems was proposed in [7]:

Function-replacing hybrids: This class of hybrids replaces the principle function of a single approach by another processing technique according to the principle of *technique enhancement*. E.g., the weight change calculation of a neural net by a genetic algorithm would belong to this class.

Intercommunicating hybrids: Due to the *multiplicity of application tasks* it is sometimes required to apply multiple autonomous processes that communicate with each other. Such *intercommunicating hybrids* collectively solve a complex task and are co-ordinated by a control mechanism. An expert system calling a neural net to search patterns could be a possible application.

Polymorphic hybrids: If a single method can achieve the functionality of different processing techniques according to the principle of *multi-functionality*, it is considered a *polymorphic hybrid*. A neural net applied for symbolic reasoning would belong to this class of hybrids.

Our hybrid system combines a signal analysis approach for EEG preprocessing, fuzzy logic to represent vagueness, a neural network for automatic knowledge acquisition and an expert system for symbolic

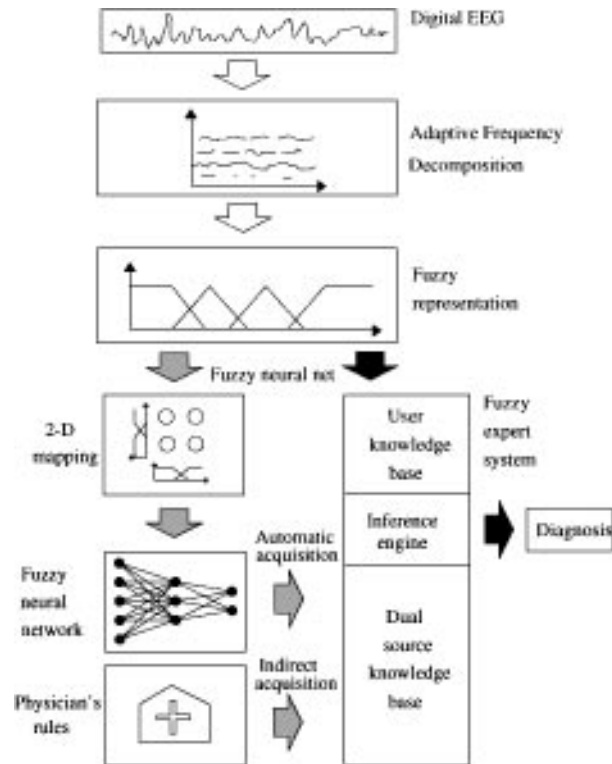


Figure 2. Schematic drawing of the hybrid system.

reasoning according to the *multiplicity of application tasks*. Since all four components are interconnected and need to communicate to solve the complex task of diagnosing an EEG the classification as an *intercommunicating* hybrid seems appropriate. Also, since the fuzzy representation is intended to overcome the weakness of crisp reasoning in the expert system, the aspect of *technique enhancement* is addressed by the system. The resulting assembly of four single approaches to form one hybrid architecture is shown in Fig. 2.

The top of the diagram indicates the preprocessing by a signal analysis (Adaptive Frequency Decomposition) that is application specific and serves as a basis for the extraction of fuzzy features (light arrows). Within the lower part of the diagram, grey arrows indicate a learning or knowledge acquisition phase and solid arrows indicate execution phase. During the learning phase, two sets of rules are combined to form the dual source knowledge base of the fuzzy expert system. On the one hand, rules are gathered from physicians and on the other hand, rules are extracted from the fuzzy neural network. For the latter purpose, the net is at first trained to perform the physician's diagnosis task on sample data [9]. For this purpose the fuzzy values have to be

mapped to the input neurons via the two-dimensional mapping scheme. After completed training, rules describing the diagnosis process are extracted from the net. When the acquisition of knowledge is finished, the execution phase can be run, actually diagnosing medical data. Now the fuzzy features of a sample (e.g., an EEG trace) are compared with the formerly acquired knowledge base via the inference engine. The fuzzy expert system then comes up with a diagnosis.

The following five subsections describe the implementation of the components of the hybrid system in more detail and show the advantages and avoidance of disadvantages.

3.1. EEG Preprocessing

Numerous signal analysis approaches are frequently applied to medical data (e.g., [10]), helping physicians to interpret electronic patient data. E.g., the following advantages are of highest importance for the topic under discussion. Features may be extracted from the signal that express certain phenomena, leading to a data reduction if only those features are stored, rather than

the whole signal. Also, raw data can be transformed into a more meaningful representation that contains the attributes of the human verbal description of the data. Nevertheless, signal analysis by itself can not serve as a reasoning diagnosis but has to be considered as a pre-processing mechanism and needs to be combined with other methods.

Discretized values of a signal, whether they represent a time series, a spatial mapping or any other medical data, are relatively meaningless to a human interpreter. For visual inspection these values are brought into an expressive graphical representation. When a physician describes phenomena within this representation, he uses features that are not explicitly contained in the raw data but generated by his interpretation. In order to provide a representation of such data suitable for reasoning with physician-formulated rules, the same features have to be extracted from the raw data.

For the case of time series in general and EEGs in particular, a description by means of frequency and amplitude is very common. To allow for such a description,

the signal has to be represented in its frequency domain rather than in its time domain. This is achieved by any frequency transformation like the Fast Fourier Transformation (FFT) [2].

The so-called alpha rhythm is a phenomenon that usually occurs in any healthy human, expressed by a sinusoidal waveform of approximately 10 Hz as can be seen in Fig. 3(a) as a time signal. Figure 3(b) shows the same data in frequency representation where the main peak results from the continuous 9.5 Hz activity.

When additional waveforms are contained in the time signal of an EEG, like the bulbus artifact (eye movement) in Fig. 4(a), also the frequency spectrum will contain additional peaks. In Fig. 4(b), the 2 Hz peak represents the slow but high-amplitude wave (bulbus artifact), the smaller camelback-shaped maximum in the range of 7 to 10 Hz results from the alpha-wave at the end of the signal and the 16 Hz component represents the tiny ripples within the signal.

To cope with some peculiarities of EEG signals that might not occur in other data, we developed the




Figure 3. One second of alpha rhythm in its time domain (a) and frequency domain (b) representation.




Figure 4. A bulbus artifact in its time domain (a) and frequency domain (b) representation.




Figure 5. Membership functions $\mu_\delta, \mu_\vartheta, \mu_\alpha$ and μ_β of the fuzzy variable *frequency*.

Adaptive Frequency Decomposition [11] to extract meaningful features from the signals. We will not go into the technical details of this specialized method but would like to mention that it has been successfully introduced to neurologists [12] and is being used for clinical applications [13, 14]. The following methods will be based on this frequency domain representation of EEG data.

Prior to this preprocessing, analog EEG signals were band-filtered with a bandwidth from 0.1 Hz to 64 Hz to eliminate direct currents and aliasing effects. Also, a notch filter for the German line frequency of 50 Hz was applied. Then, the signal was digitized into 12-bit data format at a sampling rate of 512 Hz.

3.2. Fuzzy Representation

Since we want a description of vagueness and a fuzzy interfacing mechanism, data has to be transformed into a fuzzy representation before being diagnosed by our system. This process depends on the type of data and shall be demonstrated via our EEG example.

Figure 5 shows the fuzzification of the variable *frequency* into membership values with respect to electroencephalography. Disjunctive frequency bands *delta* (δ), *theta* (ϑ), *alpha* (α) and *beta* (β) have been defined for this domain [15]. In combination with the fuzzification of the *amplitude* into the terms *zero*, *low*, *mid* and *high* this allows to formulate rules like the following that will be equally interpreted by all domain experts.

‘A *bulbus artifact* is a wave with a *delta* frequency and *high* amplitude.’

Fuzzification for Reasoning: Inference, the mechanism to draw conclusions from the knowledge base, depends strongly upon the representation of the knowledge itself [16]. Hence, effectiveness and faultiness of a system are also affected by the representation paradigm used. Here, we introduce a new representation paradigm consisting of a set of rules

about how to represent rules. Since these meta-rules are not relevant for pure control applications and are thus often neglected they shall be mentioned in the sequel.

For the set \mathcal{M} denoting all possible membership terms $\{\mu_1, \dots, \mu_k\}$ of a fuzzy variable in some universe of discourse \mathcal{X} we get:

Normalization: A widely accepted design criterion is the normalization criterion [17]:

$$\forall \mu_i \in \mathcal{M} \quad \exists x \in \mathcal{X} : \mu_i(x) = 1$$

This criterion is used for modelling existing but not precisely measurable values and demands that at least one point of each membership function is assigned the degree of membership 1 which is also considered as the degree of acceptance. (Fuzzy variables obeying this criterion can also be interpreted as possibility distributions.)

Cover Full Range: Since any point in a universe of discourse has a certain meaning, it has to have at least one attribute describing it. In human reasoning there will never be a gap of description within the range of one variable. Thus,

$$\forall x \in \mathcal{X} \quad \exists \mu_i \in \mathcal{M} : \mu_i(x) > 0$$

represents the requirement to fully specify the entire range of a universe of discourse to avoid values that can never fulfill any premise of a rule. Membership designs like in Fig. 6 that include gaps between attributes, will be excluded by this criterion. Also, for the purpose of machine learning the entire range has to be specified, since it is not known a priori at which points in the universe of discourse the learning algorithm will lay its premise or conclusion attributes.




Figure 6. Forbidden gap between membership functions.

A simple criterion that implies the above criteria is the sum-of-1-criterion (\sum 1-criterion):

$$\forall x \in \mathcal{X} : \sum_{\mu_i \in \mathcal{M}} \mu_i(x) = 1 \quad (1)$$

The fuzzy variable *frequency* (depicted in Fig. 5) obeys this criterion by letting the sum of all membership functions at any point in the universe of discourse being 1.

In contrast to natural language variables, where membership functions mostly are *low*, *medium* and *high*, for scientific variables there often exists a common terminology within a certain domain. For example, a frequency of a couple of Hertz would be considered *high* for a seismic domain, whereas some Mega Hertz might still be called *low* in Satellite Communication. Most often there exist terms like Ultra High Frequency (UHF) or Very High Frequency (VHF) to describe a range of a variable. This type of scientific terms bear the nature of being well separated, e.g., a frequency can never be UHF and VHF to the same extent. This disjunctive property of attributes assigned to variable ranges is often implied when experts formulate rules and should be taken into regard by the representation. The \sum 1-criterion ensures this type of disjunction. Whenever a membership value decreases from 1, another value has to increase from 0 (e.g., when one thinks of a temperature not to be warm any more with a degree of 1, one assumes it to be cold with a degree above 0). Thus, the sum of membership values to be assigned to a point in the universe of discourse is always 1, since one does not lose the confidence to describe the point completely—but rather chooses two different descriptions. This is expressed by the criterion of full acceptance (see below).

Three additional features can be derived from the \sum 1-criterion that make good sense in reasoning purposes:

Full acceptance: The degree of acceptance does not have value 1 only at one point in the range of a variable, as demanded by the normalization-criterion, but at every point:

$$\begin{aligned} \forall \mu_i \in \mathcal{M} : \mu_i(x) < 1 \\ \Rightarrow \exists \mu_j \in \mathcal{M}, j \neq i : \mu_j(x) > 0 \end{aligned}$$

Negations: It is now possible to express negations not only by $\mu_{\neg i}(x) = 1 - \mu_i(x)$ but by means of the remaining membership functions:

$$\forall \mu_i \in \mathcal{M} : \mu_{\neg i}(x) = \sum_{\mu_j \in \mathcal{M}, j \neq i} \mu_j(x)$$

No three at a time: The abovementioned disjunctive property also implies that in case of uncertainty over which membership to assign, a maximum of two adjacent membership functions may be activated:

$$\forall x \in \mathcal{X} : |\{\mu_i \in \mathcal{M} \mid \mu_i(x) > 0\}| \leq 2$$

A reproach often made to fuzzy logic is that membership functions are retuned after first tests have failed and that they relate in no way to any plausible point of view. To avoid this, we have designed the terms in correlation with a number of related publications [18]. In areas where there is agreement about the attribute of a variable we have assigned the value 1. Where we found disagreement, we have put the transitions of our membership functions.

Summarizing the fuzzification aspects, we would like to point out that whenever inference is applied to Fuzzy Logic, certain criteria should be regarded when designing the membership functions. This will assure adequate modelling not only of expert knowledge but also of the expert domain and its implicit inference rules (e.g., disjunctive membership terms). Design of membership functions should be carried out only once and prior to their usage:

- Relate the ranges of membership functions to scientific results in literature.
- Use common terminology from related work to name your membership functions.
- Obey to the \sum 1-criterion (Eq. (1)).

3.3. Mapping Fuzzy Variables to Neural Networks

In order to learn a pattern recognition or diagnosis process in a neural network, the fuzzy features have to be mapped into the input layer of the net. Here, we demonstrate a new method of mapping two-dimensional fuzzy variables into neural networks, minimizing the number of neurons in the network and keeping their number constant.

The eight terms of the two four-term fuzzy variables *frequency* and *amplitude* must be mapped into the input layer neurons of our neural net. In a straight-forward manner eight neurons would be implemented in the input layer each of which would be assigned the membership value of one of the eight fuzzy terms. This would allow the representation of exactly one spectral phenomenon. But in every time-slice of EEG there exist various different phenomena. For example, the main

Table 3. Calculation of the neural activity from fuzzy membership values.

Neuron	Membership functions	Values	Neural activation
$N_{\text{delta} \times \text{zero}}(f, a) =$	$\mu_{\text{delta}}(f) \cdot \mu_{\text{zero}}(a)$	$= 0.6 \cdot 0.6 =$	0.36
$N_{\text{delta} \times \text{low}}(f, a) =$	$\mu_{\text{delta}}(f) \cdot \mu_{\text{low}}(a)$	$= 0.6 \cdot 0.4 =$	0.24
$N_{\text{theta} \times \text{zero}}(f, a) =$	$\mu_{\text{theta}}(f) \cdot \mu_{\text{zero}}(a)$	$= 0.4 \cdot 0.6 =$	0.24
$N_{\text{theta} \times \text{low}}(f, a) =$	$\mu_{\text{theta}}(f) \cdot \mu_{\text{low}}(a)$	$= 0.4 \cdot 0.4 =$	0.16
		Sum of activation:	1.00

rhythm should be present in any healthy subject, described as a wave with an alpha frequency and a mid amplitude. Also there might exist artifacts at the same time—mostly bulbus artifacts resulting from eye movements described as waves with delta frequency and high amplitude. If more than one feature were encoded in eight neurons one could no longer decide which amplitude belongs to which frequency (called the binding problem, [19] pp. 88–90). Therefore, to represent multiple phenomena eight neurons would be required for each phenomenon. Since the number of features contained in every time-slice varies through the EEG (not every sample is deranged by an artifact), a representation is needed which is capable of coding multiple phenomena in a constant number of neurons.

Figure 7 shows such a coding scheme with 16 neurons, each representing a combination of two fuzzy terms. The 16 encoded combinations can be regarded as the cross product of the two fuzzy variables (similar to the tensor product [20]). Since now the two variables are put into relation, one can activate multiple neurons and still differentiate the features that are being

represented. So the two-step process of transforming points from a frequency spectrum to neural activity, represents a quantization that preserves the topography (e.g., a spectral component will have the same localisation in the spectrum and in the neural map).

The example shows the fuzzification of a spectral component. A frequency $f = 3.8$ Hz is fuzzified to the membership functions $\mu_{\text{delta}}(f) = 0.6$ and $\mu_{\text{theta}}(f) = 0.4$ while an amplitude $a = 5 \mu V$ fuzzifies to $\mu_{\text{zero}}(a) = 0.6$ and $\mu_{\text{low}}(a) = 0.4$. This is indicated in Fig. 7 without values and in Table 3 the calculation of the neural activity is shown. The resulting cross product of membership values is mapped into the 16 neurons $N_{\text{frequency} \times \text{amplitude}}$.

Depending on the values of the membership functions involved, one, two or four neurons can be activated by a single feature of this two-dimensional domain. For the sake of an equal interpretation of the rules aquired after learning, an equal representation is needed for every feature resulting in a need for an equal sum of activation. In analogy to the $\sum 1$ -criterion for membership functions (Eq. (1)), we demand that the sum of activation in all neurons resulting from one spectral phenomenon with frequency f and amplitude a equals 1. We call this the *sum-of-1-criterion for neural activity*:

$$\forall f \in \mathcal{F}, a \in \mathcal{A} : \sum_{\substack{i \in \{\text{delta}, \dots, \text{beta}\}, \\ j \in \{\text{zero}, \dots, \text{high}\}}} N_{i \times j}(f, a) = 1$$

where \mathcal{F} and \mathcal{A} denote the two universes of discourse for the two fuzzy variables *frequency* and *amplitude*. In general, this leads to:

$$\forall x \in \mathcal{X}, y \in \mathcal{Y} : \sum_{\substack{i \in \mathcal{M}_i, \\ j \in \mathcal{M}_j}} N_{i \times j}(x, y) = 1$$

where \mathcal{M}_i and \mathcal{M}_j denote the sets of possible membership terms of two fuzzy variables with universes of discourse \mathcal{X} and \mathcal{Y} .

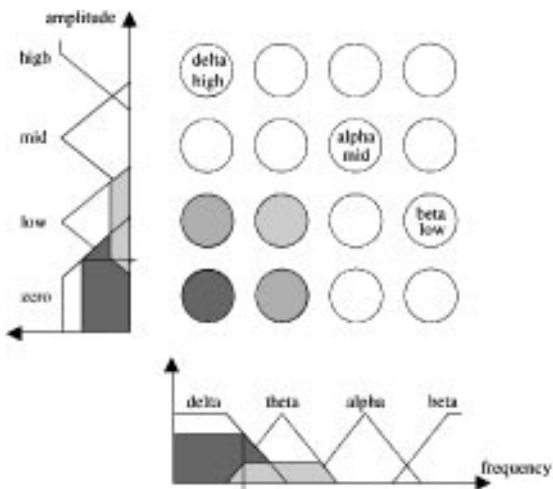


Figure 7. Two-dimensional mapping scheme.

In order to fulfill this criterion, we chose the algebraic product to implement the *and* combination of two membership terms i and j in each neuron $N_{i \times j}$. Only by this implementation, we can guarantee to satisfy the \sum 1-criterion for neural activity.

Proposition 1. *Let there be two fuzzy variables $V1$ and $V2$ in two universes of discourse \mathcal{X}_1 and \mathcal{X}_2 with sets of membership terms \mathcal{M}_1 and \mathcal{M}_2 , satisfying their \sum 1-criteria for membership functions (1).*

$$\forall x_1 \in \mathcal{X}_1. \sum_{i \in \mathcal{M}_1} \mu_i(x_1) = 1 \quad (2)$$

$$\forall x_2 \in \mathcal{X}_2. \sum_{i \in \mathcal{M}_2} \mu_i(x_2) = 1 \quad (3)$$

Then the elements of the cross product of $V1$ and $V2$ will sum up to 1.

Proof:

$$\begin{aligned} \forall x_1 \in \mathcal{X}_1, x_2 \in \mathcal{X}_2 : & \sum_{\substack{i \in \mathcal{M}_1, \\ j \in \mathcal{M}_2}} \mu_i(x_1) \cdot \mu_j(x_2) \\ = & \sum_{i \in \mathcal{M}_1} (\mu_i(x_1) \cdot \sum_{j \in \mathcal{M}_2} \mu_j(x_2)) \\ = & \sum_{i \in \mathcal{M}_1} \mu_i(x_1) \cdot 1 && \text{due to (3)} \\ = & \sum_{i \in \mathcal{M}_1} \mu_i(x_1) \\ = & 1 && \text{due to (2)} \quad \square \end{aligned}$$

Clearly, we can use the same argumentation to prove the correctness of the \sum 1-criterion for neural activity for any arbitrary number n of dimensions instead of this $n = 2$ example. Hence, the mapping scheme and all of its advantages and possibilities, overcoming the binding problem with a limited and constant number of input neurons, applies for any multi-dimensional variable that can be represented in the above formalism.

For two fuzzy variables, one having n membership terms and the other having m membership terms, this results in a fixed number of $n \cdot m$ required neurons as demanded above. In the abovementioned straightforward mapping $(n + m) \cdot f$ neurons would be required for a number of f features occurring. Thus, there is a certain number of features above which this method also results in a data compression (break-even point). If the number of features exceeds the break-even number, f_{BE} , the number of neurons required is less if our

mapping technique is applied.

$$f_{BE} = \frac{n \cdot m}{n + m}$$

In the case of *amplitude* and *frequency*, each having four membership terms, f has to be above 2 to result in a reduction of neurons needed. Due to the fact that at least two spectral components occur in every EEG sample (α -basic rhythm and cognitive β -activity), the mapping technique also shrinks the network size.

3.4. Fuzzy Neural Network

The fuzzy features, presented to the neural network by the two-dimensional mapping method, are then trained to be detected by the net. Among the great variety of neural network architectures, ranging from “vanilla” backpropagation to more enhanced radial basis function networks (RBF), we chose a fuzzy-neural architecture. We compared different types of fuzzy networks, that allow us to extract rules after learning, to conventional approaches in [21] and decided on a network called FuNe (for details on implementation see [22]).² The special multilayer perceptron architecture is trained with a gradient descent algorithm. There exist three types of neurons in the hidden layer grouped together topographically. One group of neurons can perform only the *or* function of multiple inputs while another group only performs the *and* function. As there may also be unary rules, composed of simple one-term-premises, there is a third group of neurons having single inputs and single outputs. The output neuron simply acts as an *or* function of all middle neurons. In the initial state, the fully interconnected network represents all possible logical functions of *or*-premises, *and*-premises, and the 16 *unary* premises. The total number of hidden neurons, n , in (4) can be reduced by limiting the number of premise terms, p .

$$n = 2 \cdot \sum_{i=2}^p \binom{16}{i} + 16 \quad (4)$$

After the learning process, all connections below a certain threshold are eliminated. A similar pruning method has been proposed by Le Cun in 1990 [24] in order to increase learning speed—but, it is also useful for limiting the number of resulting rules. This is of special interest to us, because we will extract exactly these rules after the training and want to avoid




Figure 8. Neural network after completed training.

rules with negligible rule strength for the sake of better interpretation (see Section 4.3).

A network trained to detect bulbus artifacts is shown in its final state in Fig. 8. As an example, the pattern at the input neurons represents a bulbus artifact (ba) in a simplified manner³. The delta frequency component is high while all other frequency components are low. The artifact is detected in the output neuron, shown by its activation of 1. Some of the inputs no longer contribute to the detection task at all, since their low-weight connections have been pruned.

Rules R1 to R3 are extracted from the three shaded neurons of the network and build up part of the knowledge base for the fuzzy expert system⁴.

- R1: IF delta mid OR delta high THEN ba
- R2: IF delta zero THEN -ba
- R3: IF delta high THEN ba

Not all of these rules contribute to the result in the same way, which is expressed by different rule strengths. The strength of a rule is computed from the weight of the connection between the neuron representing the rule and the output neuron. Rule (2) has a negative strength indicating a negation of the conclusion (dashed line in Fig. 8).

3.5. Fuzzy Expert System

The heart of our hybrid system is a fuzzy expert system, called FuzzyCLIPS [25], derived from CLIPS [26]. This expert system consists mainly of three components (see Fig. 2, 'Fuzzy expert system'), the *dual source knowledge base* containing the combined rule sets from indirect and automatic knowledge acquisition, the *user knowledge base* where the actual phenomena of a patient are entered, and the *inference mechanism* concluding a diagnosis from the former two.

The calculation of a rule premise in our fuzzy expert system is shown in Fig. 9 in a graphical representation for the rule

IF alpha mid THEN basic rhythm

The premise of the rule is considered as the *conjunction* of the two membership functions *alpha* and *mid* in the linguistic variables *frequency* and *amplitude* as indicated by the grey pyramid.




Figure 9. Two dimensions of a spectral component as two linguistic variables *frequency* and *amplitude* showing the conjunction $\mu_{\alpha}(f)$ AND $\mu_{\text{mid}}(a)$. The resulting membership value of the grey pyramid has to be projected to the conclusion *basic rhythm* that is defined in a different universe of discourse.

Correctly, this would be formalized as fuzzy modus ponens for approximative reasoning [27]:

$$\frac{p \Rightarrow c, p'}{c'}$$

reading as *from an implication $p \Rightarrow c$ and a premise p' , one can infer the conclusion c'* . Other than in classical logic, here, premise and conclusion are not simply *true* (0) or *false* (1) but have membership values in the range [0, 1]. Thus, the premise will be written as $p = \mu_p(x)$ and the conclusion as $c = \mu_c(y)$ where x and y are elements of two universes of discourse \mathcal{X} and \mathcal{Y} . Additionally, there is a membership function μ_R of the rule or implication,

$$\forall (x, y) \in \mathcal{X} \times \mathcal{Y} : \mu_R(x, y) = \text{MIN}(\mu_p(x), \mu_c(y))$$

Given the membership value $\mu_{p'}(x)$ of a fact p' , now we can infer the grade of membership $\mu_{c'}(y)$ of the fuzzy conclusion c' being the maximum of all possible minima in \mathcal{X} .

$$\begin{aligned} \forall (x, y) \in \mathcal{X} \times \mathcal{Y} : \mu_{c'}(y) \\ = \text{MAX}_{x \in \mathcal{X}}(\text{MIN}(\mu_{p'}(x), \mu_R(x, y))) \end{aligned}$$

When acquiring rules for the knowledge base of an expert system, two major paradigms apply: indirect and automatic knowledge acquisition. For the case of indirect knowledge acquisition the system is taught which rules lead to the desired decision capability. Automatic knowledge acquisition is applied when this knowledge is incomplete but examples may serve as a teacher for machine learning techniques. Our knowledge base consists of two differently acquired sets of rules. The first one is acquired indirectly from an expert by having him explain the rules that he uses to solve the diagnosis task. The second one is acquired automatically by machine learning in a neural network. The resulting dual source knowledge base integrates two autonomous rule sets. They might contain rules that are equal in premises and conclusion but different in the degree of certainty, which results in contradiction in the worst case. As proposed by Holland in 1986 [28], we do allow these contradictory rules, postulating it as a natural feature of decision processes. The following example might shed some light upon this matter:

R3: IF delta high THEN BA
 R4: IF delta high THEN FIRDA

The first rule, R3, is the formalized representation of the verbal rule for a bulbus artifact (BA) from Section 3.2 while the second rule, R4, describes so-called *frontal intermittend rhythmic delta activity*. If the two rules R3 and R4 have the same importance (equal rule strengths) one certainly needs a further rule to make a decision, like

R5: IF occurrence rhythmic THEN FIRDA

which contributes to one of the former conclusions. The same is supposed for our dual source knowledge base. There may very well be phenomena that exist in sample data that lead to rules while the same phenomenon is described differently by the expert rules. In this case the resulting uncertainty has to be solved by a higher rule strength for the more probable rule or a further rule, manifesting one of the possible conclusions.

Besides an explanation component it is also important for this medical application that a *don't know* conclusion exists, telling the user if abnormal phenomena were detected but could not be interpreted.

4. Results

EEGs were recorded from a total of 24 patients during routine neurologic medical check-up (13 males, 11 females) plus one extra dummy EEG to train the physician in handling the system. Actual diagnoses by the physicians were ranking from healthy (volunteer test person) over severe epilepsy to brain death. Each patient recording lasted approximately 20 minutes. According to the international 10–20 system⁵ 16 traces of bipolar electrode recordings were considered for our experiments. EEG raw data was split into two separate data sets for learning and evaluation. The system was trained with manually selected data samples from the learning set and evaluated with all samples from the evaluation set. The fuzzy neural network was trained with preprocessed one-second EEG samples from different traces. Since every second of EEG contains multiple relevant phenomena, several thousand data samples were used for training and evaluation.

For the task of evaluation, we worked with two independant neurologists. One of them based his diagnosis solely on the visual analysis of the EEG raw data. The other one used the list of detected phenomena and was only shown selected EEG samples by the system. The presented samples were used by the explanation component as justification for the inferred diagnosis

results. Since the neurologist using the automatic diagnosis system saw only about 0.5% of the actual EEG, the application of the system resulted in a significant decrease of needed diagnosis time—neglecting the time for the system to run the rule inferencing process⁶. The experiments were conducted across all patients.

When talking about results in medical diagnosis, we would first like to mention one major point on which to put the main focus. Other than in common detection tasks it is not sufficient to simply watch the overall performance of the system in terms of average error, like it would be appropriate for character recognition, for instance. The diagnosis performance must be split up into the missing of phenomena, called false negatives (FN), and the accidental finding of phenomena which are not actually present, called false positives (FP). For certain diagnosis types, it is very important to have absolutely no false positives whereas a few false negatives are still acceptable, like in diagnosing cerebral death. Vice versa, in other cases it is important to have no false negatives, like in diagnosing epileptic seizures for emergency purposes. (Both examples were chosen from the class of EEG diagnosis tasks.)

4.1. Medical Phenomena

Our system is not intended to come up with a final diagnosis including medication and further treatment but it is supposed to support the physician with objective measures about the EEG. As a first decision criterion, the doctor is supplied with qualitative and quantitative information about the basic rhythm of the patient. This information is given in form of localisation, spatial extent and temporal persistence as well as

the frequency and amplitude of the basic rhythm. Additionally, certain graphoelements are detected such as bulbus artifacts, electrode artifacts, transpiration artifacts, slow activity in the delta or theta band and epileptic seizures (see 5.1 for spikes and sharp waves). Their temporal and spatial occurrence is then evaluated and higher-order phenomena are reported to the medical expert (Berger effect, focal pathologic slow activity, generalized pathologic slow activity and hyperventilation responses).

Except in one case, for all of the patients under investigation, our system was able to detect enough of the patient's EEG-phenomena to lead the two independent physicians to diagnoses that were identical in medical relevance. They both agreed on the conformity of their diagnoses. But, a lack of EEG activity, as in the examined case of brain death, also resulted in a lack of applicable rules, thus leading to a failure of the system. This was the reason for the abovementioned *don't know* conclusion that we implemented retrospectively.

In the sequel, the potential performance accuracy of the system is demonstrated with respect to detecting bulbus artifacts. Figure 10 shows two diagrams illustrating false positives (dashed line) versus false negatives (solid line) as a function of a certain threshold ϵ (FP/FN-diagram) for the detection of bulbus artifacts. The decision of the system, whether it is a bulbus artifact or not, is represented by an analog value in the range from 0 to 1. For a binary decision, a threshold (ϵ) is needed, above which a BA is assumed. If ϵ is chosen too high, not all BAs are detected, resulting in false negatives. For too low a threshold, phenomena which aren't BAs are detected—thus representing false positives. In Fig. 10(a) there is a wide gap between false positives and false negatives (0.3–0.6). A threshold




Figure 10. False positive BAs (dashed) and false negative BAs (solid) of a patient record containing no other artifacts than bulbus artifacts (a) and of a patient record with additional electrode artifacts (b).

centered in this gap will assure a high probability for neither diagnosing false positives nor false negatives.

The example of Fig. 10(b) nicely demonstrates the degree of accuracy of the system. The diagram shows only a narrow gap for the threshold, proposing possible misses in future data. The system can correctly diagnose all existing bulbus artifacts in the patient record with ease, shown in the narrow band of false negatives (rightmost peak in Fig. 10(b)). Nevertheless, the wide scope of apparently false positives gives rise to the assumption that the diagnosis performance is of minor quality, since with a threshold of $\epsilon = 0.5$, 10% diagnosed BAs seem to be false positives. Retrospectively examined, it showed that those false positives were indeed BAs, but of minimal amplitude and thus overlooked in the visual diagnosis by the human expert. Only the 1.8% false positives that remained up to a threshold of $\epsilon = 0.94$ were not BAs but electrode artifacts, being graphoelements that are somewhat related to BAs in terms of frequency and amplitude and can only be kept apart by an inter-electrode examination that we implemented in an additional rule. Thus, we showed with these results that the system's precision is of high accuracy.

The comparison of the two rule sets, serving as basis for the dual source knowledge base, brought up two interesting results.

- The automatically generated rule set, extracted from the network, contained more rules than the one indirectly acquired from the expert. These extra rules seem redundant at first glance, because the expert can easily distinguish between different phenomena on the basis of more obvious criteria. But, for the electronic system designed to diagnose a vast variety of phenomena it may very well be of great importance to 'know' every extra description neglected by the expert.
- The extracted rules were more precise in their degree of fuzzy membership, since they were actually calculated from examples.

For our example of the bulbus artifact, the physician formulated the rule:

IF freq = delta AND ampl = high THEN ba

The automatically generated rule base contained a whole group of rules with different rule strengths. The premises, conclusions and rule strengths are shown below (compare to the network connections in Fig. 8). Only the rules with rule strengths (RS) above 1.0

are listed, while rules with lower rule strengths were pruned (see Section 4.3).

Premise:	Conclusion:	RS:
delta mid OR delta high	ba	1.75
delta zero	ba	1.51
delta high	ba	2.73

The detailed rule base is the reason why the system was able to detect low-amplitude artifacts (see Fig. 10(b)). This would not have been possible with the mere expert rule.

4.2. Overcoming the Disadvantages of Single Approaches

The fact that a rule base does not fully cover all possible phenomena in sample data is taken into regard in our hybrid system by splitting the rule base into two parts—one formulated by a human expert, and another one generated by machine learning. A fuzzy neural network with the possibility of extracting knowledge in the form of rules achieves the automatic acquisition. Also, the fuzzy expert system we use can handle the mentioned problem of representing vagueness by applying linguistic fuzzy variables from Fuzzy Logic.

Our fuzzy neural net is easily interpretable due to a limited but logically defined architecture which solves the *black-box phenomenon*. Of course, this is bought with a decrease of generalization ability. Limiting the neuron's functionality to logical operations may result in not being able to solve a detection task that would well be solvable with a standard backpropagation algorithm⁷. But, if the FuNe network error does converge to zero, the knowledge that is represented within the network architecture can be extracted in form of rules. Since in our network every neuron represents the premise of a rule, it is possible to weaken, strengthen or freeze a premise or a whole rule, allowing interaction with the learning process.

By integrating the signal analysis into our hybrid reasoning architecture, we are now able to interpret the numerical values automatically that previously had to be judged by a human expert.

4.3. Knowledge Base Pruning

The complete set of rules, derived from neural network learning contains a large number of different rules and is too complex to be evaluated by an expert. Thus, the correctness can not be judged by the physician.

Table 4. Pruning results.

Pruned	Simple	AND	OR	False—	Low amplitude BAs
<30%	Increasing			0	0
30%	17	56	16	0	1
40%	4	8	1	0	4
50%	3	5	1	0	5
60%	1	4	1	0	6
>60%	Decreasing			>0	6

Also, too large a number of rules, which have been adapted exactly to the sample data, results in a loss of generalization capability which is called the overfitting problem. A pruning method for neural networks has been proposed by [24] to improve learning speed by cutting down the number of weights in the net which also results in better generalization. The number of weights differing from zero more than a value θ is proportional to the number of rules extracted from our network FuNe2. It has been shown that a *neural network can solve difficult problems efficiently if it is given more neurons initially than it needs subsequently* [30]. We applied this result by training a fully interconnected network and extracting *all* resulting rules before pruning off those rules with rule strength smaller than a value θ . This threshold θ is increased, until a fixed percentage of rules is pruned (see first column of Table 4). When pruning more than 60%, the system starts to miss some bulbus artifacts (false negatives).

Table 4 shows how the pruning results in a decrease of the three rule types (Simple, AND and OR) leading to a better interpretation. It also leads to an increase of detected low-amplitude bulbus artifacts which add to the number of high-amplitude bulbus artifacts that remained stable⁸. Since these low-amplitude artifacts were initially overlooked by the visual examiner, this demonstrates how the increased generalization ability of the system outperforms the human expert.

5. Discussion

By combining three major methodologies of Artificial Intelligence and a signal analysis approach into a single hybrid system we managed to combine most of their advantages, avoiding some of their disadvantages at the same time.

The hybrid system, described in this article, introduces the following new paradigms of modeling the cognitive task of diagnosis:

- Instead of acquiring the whole knowledge base automatically from examples (inductive learning [31]), or refining a rough knowledge base [32], we are using a dual source knowledge base. This knowledge base consists of two sets of rules, coexisting with equal importance, one generated indirectly from rules formulated by an expert and another one generated automatically by machine learning in a neural network. This offers the full range of benefits from neural network learning. Additionally, the automatically acquired rules can be supervised and adapted in the expert system, eliminating the *black-box phenomenon* (see Subsection 4.2).
- The coexistence of the two resulting rule bases, with the possibility of competing knowledge, is not only permitted but desired. The competition of rules and the mechanism of taking further rules into account is part of the cognitive decision task.

Besides the integration of the existing fuzzy neural network and fuzzy expert system, some intelligent interfacing techniques are introduced, that might equally be used for different applications:

- A two-dimensional mapping technique, used to assign the membership values of linguistic variable terms to input layer neurons of a network. By this means, it is possible to input existing fuzzy representations into neural networks in order to autonomously acquire knowledge from sample data (see Subsection 3.3).
- An application-specific fuzzification of spectral EEG data that will work for most other multi-dimensional data, especially other frequency domain data, like voice spectra in speech recognition, as well as all kinds of medical images (see Subsection 3.2).

Also, a paradigm for the design of membership functions of fuzzy variables is introduced that is most useful for reasoning purposes. In Section 4.2, we pointed out how we were able to overcome some of the disadvantages of the combined approaches by combining them to a hybrid system. Our results (see Section 4) point out some promising features:

- We proved the possible precision of our system, being more accurate than the human expert, when applied to the task of diagnosing bulbus artifacts in EEGs.
- Comparing the two rule sets of our dual source knowledge base it showed that the automatically

Table 5. The WWW addresses of some related web pages.

Topic	URL
The author	http://aida.intellektik.informatik.th-darmstadt.de/~chris
FuNe	http://www.microelectronic.e-technik.th-darmstadt.de/rych/neuro.html
CLIPS	http://www.jsc.nasa.gov/~clips/CLIPS.html
FzCLIPS	http://ai.iit.nrc.ca/fuzzy/fuzzy.html

acquired set was more extensive and of higher precision than the one indirectly acquired.

- By pruning our automatically acquired rule base, we managed to increase the generalization of the system and to avoid the overfitting problem (Section 4.3).
- Adapting our system to specific diagnosis tasks is supported by easy interfacing. The exchanging and adding of components is very easy, due to the modular architecture (see Fig. 2).

5.1. Limitations and Extensions

The system proved to work well to diagnose those phenomena that can be represented in the frequency domain by the Adaptive Frequency Decomposition preprocessing. But, since frequency transformation is only suitable for continuous waveforms, sharp transients such as spikes etc., can not be represented [33]. For this reason, we investigated the possibility of implementing a wavelet analysis, that is widely used to detect such transients [10].

In [34] we introduced a wavelet preprocessing alternative to the Adaptive Frequency Decomposition that is supposed to work in parallel with that component (compare Fig. 2). This also nicely demonstrates the possibility of extending and varying the hybrid system, since the architecture uses a standardized fuzzy interface to interconnect the components. By a wavelet analysis, a signal is generated from a spiky EEG that indicates the occurrence of sharp transients, e.g., a high value of the signal represents sharp transients while low values represent continuous slow frequency EEG segments. The performance of the wavelet preprocessing has been evaluated in [35] who detected 85.2% sharp transients with a standard neural network and 88.8% with the FuNeI architecture in the signal generated by the wavelet analysis.

In future we plan to investigate the interference effects of a dual source knowledge base. They occur when the two bases contain contradictory or similar rules with different rule strengths. For every case of perfect contradiction of two rules, possibly resulting in

no conclusion, an additional rule could be generated by a supervising mechanism to indicate the occurrence of a conflict in the rule base to the user.

6. Internet Availability

Currently we are setting up an internet distribution of relevant parts of the system. For information about this, access the authors WWW-page. For sources of system components by others than the author, corresponding WWW-pages are listed in Table 5.

Acknowledgments

The author expresses his thanks to Kung-Kiu Lau who polished the English of this article and Antje Strohmaier who gave valuable comments on an earlier version.

Notes

1. "The grander claims for rule-based AI (Artificial Intelligence) in the 80's seemed to fail in part because specialist practitioners either failed to understand or failed to flag the limitations of their single approach. . . . there is no one definitive approach which will explain cognition or solve complex problems . . . instead, there are a number of tools and models that can be applied under different circumstances" (from [8], p. vii).
2. For a survey and critique of rule extraction from neural networks in general see [23]).
3. For sake of simplicity, neuron activations are either 0 or 1. In reality they are assigned values in the range [0, 1] depending on the membership functions.
4. The '-' represents a negation in rule R2.
5. See [29] for details.
6. Since EEGs are not diagnosed by the person who records them, but by a different person at a later time, this neglectation seems allowed.
7. This did not occur for our relatively simple EEG signal analysis task. But, when we applied our fuzzy neural network to some of the problems of the UCI machine learning database, it showed that we were not able to reduce the FuNe network error to nearly zero which was possible for a backpropagation net.
8. High-amplitude bulbus artifacts were easily detected both by the physician and by the system due to their significant shape (compare Fig. 1).

References

1. W. Bibel, "Intellectics," in *Encyclopedia of Artificial Intelligence*, pp. 705–706, 1992.
2. B.H. Jansen, J.R. Bourne, and J.W. Ward, "Spectral decomposition of EEG intervals using walsh and fourier transforms," *IEEE Transactions on Biomedical Engineering*, vol. 28, pp. 836–838, 1981.
3. B.G. Buchanan and E.H. Shortliffe, *Rule-Based Expert Systems*, Addison Wesley, 1985.
4. E.H. Shortliffe, *Computer-based medical consultation: MYCIN*, American Elsevier, 1976.
5. A.N. Mamelak, J.J. Quattrochi, and J.A. Hobson, "Automated staging of sleep in cats using neural networks," *Electroencephalography and Clinical Neurophysiology*, vol. 79, pp. 52–61, 1991.
6. L.A. Zadeh, "Fuzzy sets," *Information and Control*, vol. 8, pp. 338–353, 1965.
7. S. Goonatilake and S. Khebbal, "Intelligent hybrid systems: Issues, classifications and future directions," in *Intelligent Hybrid Systems*, pp. 1–20, 1995.
8. S. Goonatilake and S. Khebbal (eds.), *Intelligent Hybrid Systems*, John Wiley & Sons, 1995.
9. C.S. Herrmann, "A fuzzy neural network for detecting graphoelements in EEGs," in *Supercomputers in Brain Research: From Tomography to Neural Networks*, edited by H.J. Herrmann, D.E. Wolf, and E. Pöppel, World Scientific Publishing Company, pp. 193–198, 1995.
10. S.J. Schiff, A. Aldroubi, M. Unser, and S. Sato, "Fast wavelet transformation of EEG," *Electroencephalography and Clinical Neurophysiology*, vol. 91, pp. 442–455, 1994.
11. C.S. Herrmann, "Mechanism and System for Analysing Electroencephalographic Recordings," *PCT Patent Application PCT/EP 96/02309 in Europe, USA and Japan*, 1996.
12. C.S. Herrmann, H.P. Hundemer, and H.C. Hopf, "Adaptive Frequency Decomposition for EEG-Analysis," in *40th Annual German Conference on Neurophysiology*, 1995.
13. H.P. Hundemer, C.S. Herrmann, and H.C. Hopf, "Clinical Use of a New EEG-Analysis System," in *40th Annual German Conference on Neurophysiology*, 1995.
14. L. Schuster, C.S. Herrmann, H.P. Hundemer, and H.C. Hopf, "Automatic focus detection in digital EEG by a knowledge-based system" (*European Congress on Clinical Neurophysiology*), *Electroencephalography and Clinical Neurophysiology*, vol. 99, no. 4, p. 317, 1996.
15. A.S. Gevins and A. Rémond (eds.), *Methods of Analysis of Brain Electrical and Magnetic Signals (Handbook of Electroencephalography and Clinical Neurophysiology)*, Elsevier Science Publishers, vol. 1, 1987.
16. W. Bibel and J.-M. Nicolas, "The role of logic for data and knowledge bases—A brief survey," in *Foundations of Knowledge Base Management*, edited by Joachim W. Schmidt and C. Thanos, Springer: Berlin, pp. 3–22, 1988.
17. R. Kruse, J. Gebhardt, and F. Klawonn, *Foundations of Fuzzy Systems*, Wiley, 1994.
18. E. Niedermeyer and F. Lopes da Silva, *Electroencephalography, Basic Principles, Clinical Applications and Related Fields*, William & Wilkins, 1993.
19. D.E. Rumelhart and J.L. McClelland, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, MIT Press, 1986.
20. P. Smolensky, "Tensor product variable binding and the representation of symbolic structures in connectionist systems," *Artificial Intelligence*, vol. 46, pp. 159–216, 1990.
21. C.S. Herrmann, S.K. Halgamuge, and M. Glesner, "Comparison of fuzzy rule based classification with neural network approaches for medical diagnosis," in *European Congress on Fuzzy and Intelligent Technologies (EUFIT)*, edited by H.-J. Zimmermann, Wissenschaftsverlag Mainz, pp. 1664–1667, 1995.
22. S.K. Halgamuge, W. Pöschmüller, S. Ting, M. Höhn, and M. Glesner, "Identification of underwater sonar images using fuzzy-neural architecture FuNeI," in *International Conference on Artificial Neural Networks (ICANN)*, Springer, 1993, pp. 922–925.
23. R. Andrews, J. Diederich, and A.B. Tickle, "A survey and critique of techniques for extracting rules from trained artificial neural networks," *Knowledge-Based Systems*, vol. 8, no. 6, 1995.
24. Y. Le Cun, J.S. Denker, and S.A. Solla, "Optimal brain damage," in *Proceedings of the 2nd Conference on Advances in Neural Information Processing Systems (NIPS-90)*, edited by R.P. Lippman, J.E. Moody, and D.S. Touretzky, 1990, pp. 598–605.
25. National Research Council Canada. *FuzzyCLIPS User's Guide Version 6.02A*, Knowledge Systems Laboratory, 1994.
26. NASA, CLIPS Reference Guide, *Artificial Intelligence Section*, 1993.
27. S. Russell and P. Norvig, *Artificial Intelligence: A modern Approach*, Prentice Hall, 1995.
28. J.H. Holland, "Escaping brittleness: The possibilities of general-purpose learning algorithms to parallel rule-based systems," in *Machine Learning*, edited by R.S. Michalski, J.G. Carbonell, and T.M. Mitchell, Morgan Kaufmann, vol. 2, chap. 20, pp. 593–623, 1986.
29. H.H. Jasper, "The ten twenty electrode system of the international federation," *Electroencephalography and Clinical Neurophysiology*, vol. 10, pp. 371–375, 1958.
30. G.A. Brown, C. Hulme, P.D. Hyland, and I.J. Mitchell, "Cell suicide in the developing nervous system: A functional neural network model," *Cognitive Brain Research*, vol. 2, pp. 71–75, 1994.
31. R.S. Michalski, "A theory and methodology of inductive learning," in *Machine Learning*, edited by R.S. Michalski, J.G. Carbonell, and T.M. Mitchell, Tioga, chap. 4, pp. 83–134, 1983.
32. D. Ourston and R.J. Mooney, "Theory refinement combining analytical and empirical methods," *Artificial Intelligence Journal*, vol. 66, pp. 273–309, 1994.
33. S.D. Stearns, *Digital Signal Analysis*, Hayden Book Company, 1975.
34. K. Jahnke, "Preprocessing of EEGs by means of wavelet-analysis for subsequent processing in a neural network," Master's Thesis, TH Darmstadt, FG Mikroelektronische Systeme/FG Intellektik, 1995.
35. C. Nickel, "Investigating brain electric potentials with intelligent systems," Master's Thesis, TH Darmstadt, FG Mikroelektronische Systeme/FG Intellektik, 1995.
36. S.C. Shapiro, *Encyclopedia of Artificial Intelligence*, John Wiley, 1992.

Christoph S. Herrmann is an associate researcher in the field of Artificial Intelligence. His research activities are focused on Soft-Computing techniques and medical applications. Dr. Herrmann received a Master's degree in electrical engineering in 1992 and a PhD in Computer Science in 1996, both from Darmstadt University of Technology.