

A. Fred, J. Filipe, M. Partinen, T.Paiva, "PSG-Expert: An Expert System for the Diagnosis of Sleep Disorders". In European Neurological Network, T. Paiva and T. Penzel (Eds), IOS Press, no 78, series Studies in Health Technology and Informatics, pp 127-147, 2000.

PSG-EXPERT – AN EXPERT SYSTEM FOR THE DIAGNOSIS OF SLEEP DISORDERS

Ana FRED

*Instituto de Telecomunicações / Instituto Superior Técnico
Av. Rovisco Pais, 1049-001 Lisboa, Portugal
Email: afred@lx.it.pt*

Joaquim FILIPE

*Escola Superior de Tecnologia / Instituto Politécnico de Setúbal
R. Vale de Chaves, Estefanilha, 2910 Setúbal, Portugal*

Markku PARTINEN

*Sleep Disorders Clinic and Research Center
Helsinki, Finland*

Teresa PAIVA

*Centro de Estudos Egas Moniz / Hospital de Santa Maria
Lisboa, Portugal*

Abstract. This paper describes PSG-EXPERT, an expert system in the domain of sleep disorders exploring polysomnographic data. The developed software tool is addressed from two points of view: (1)- as an integrated environment for the development of diagnosis-oriented expert systems; (2)- as an auxiliary diagnosis tool in the particular domain of sleep disorders.

Developed over a Windows platform, this software tool extends one of the most popular shells – CLIPS (C Language Integrated Production System) with the following features: backward chaining engine; graph-based explanation facilities; knowledge editor including a fuzzy fact editor and a rules editor, with facts-rules integrity checking; belief revision mechanism; built-in case generator and validation module. It therefore provides graphical support for knowledge acquisition, edition, explanation and validation.

From an application domain point of view, PSG-Expert is an auxiliary diagnosis system for sleep disorders based on polysomnographic data, that aims at assisting the medical expert in his diagnosis task by providing automatic analysis of polysomnographic data, summarising the results of this analysis in terms of a report of major findings and possible diagnosis consistent with the polysomnographic data. Sleep disorders classification follows the International Classification of Sleep Disorders. Major features of the system include: browsing on patients data records; structured navigation on Sleep Disorders descriptions according to ASDA definitions; internet links to related pages; diagnosis consistent with polysomnographic data; graphical user-interface including graph-based explanatory facilities; uncertainty modelling and belief revision; production of reports; connection to remote databases.

1. INTRODUCTION

The problem of medical diagnosis can be stated as follows: given a set of *symptoms* (clinical data) and *signals*, or *test results* (tests performed on the patient), appraise pathological situations, identifying which *diseases* justify the particular findings.

A variety of computer assisted techniques have been proposed to help solving the diagnostic problem [1,7,12,15]. Commonly known as *expert systems*, approaches, undertaking distinct formalisms for modeling domain knowledge, include: rule-based or production systems, frame-based systems, semantic networks, neural networks, and Bayesian belief networks. Explanatory mechanisms are more easily implemented with rule-based systems and graphical models making these preferred solutions to the diagnostic problem. Since the pioneering work by Feigenbaum (with the mass spectrogram interpreter DENDRAL, 1971), and the MYCIN project (a computer system that diagnosis bacterial infections of the blood and prescribes suitable drug therapy) by Shortliffe in 1976, the theory of expert systems has progressively evolved, accompanied by the appearance of expert systems development tools.

This paper describes PSG-EXPERT, an expert system in the domain of sleep disorders. This auxiliary diagnosis system aims at assisting the medical expert in his diagnosis task by providing automatic analysis of polysomnographic data, summarising the results of this analysis in terms a report of major findings and possible diagnosis consistent with the polysomnographic data. Sleep disorders classification follows the International Classification of Sleep Disorders - ICSD [5].

The development of knowledge based systems to be used as ancillary diagnosis systems, usually undertakes a software engineering methodology based on the spiral model. Using this paradigm a series of increasingly refined prototypes are produced until an end-user version is finalised and delivered. Each development cycle comprises three main phases, namely: knowledge acquisition and refinement; software implementation; and system validation. Although the first two phases are of major importance, validation plays a decisive role in systems quality and acceptance, having also a great impact in the development time. Yet, many development tools or methodologies [2] fail to give much attention to this phase or completely overlook it.

The first part of the paper (section 3) presents essentially the technical aspects of the PSG-EXPERT, as an integrated environment, which provides an adequate graphical support for every phase in the expert system development cycle, from knowledge acquisition and edition to reasoning explanation and knowledge-base validation, with emphasis on the last phase. The PSG-EXPERT environment constitutes a specialization of a general Artificial Intelligence technology for our application domain.

Section 4 addresses the tool from an application domain perspective, as an auxiliary diagnosis environment for sleep disorders. Adopting a users point of view, the versatility and distinctive features of the system are illustrated through a set of interaction examples.

2. SYSTEM ARCHITECTURE AND FUNCTIONALITY

Many tools (shells) have been proposed to assist the development team of expert systems in the design and implementation phase, in order to achieve fast prototyping. These shells are often text-based, with the exception of very expensive tools [8] [13][14]. Some of the more flexible and expensive tools were developed for dedicated environments such as Lisp-machines, although nowadays these machines have almost disappeared, due to their specificity and high cost [1]. Therefore, the AI community has been watching with some degree of expectation, the development of an academic tool, called CLIPS [4] that was created at the NASA's Johnson Space Centre. CLIPS has the advantage of being C-based,

instead of LISP-based and it is a very low-cost expert system development tool. However, there are some difficulties, concerning mainly two problems: the lack of a backward-chaining module (CLIPS has a forward-chaining inference engine based on the well-known OPS5 architecture) and the lack of an explanation facility (either text-based or graphic-based).

Besides their particular difficulties and/or limitations as discussed above, existing tools neglect the important phase of system validation. Some actually provide incipient rule integrity checking mechanisms, but none seems to address performance evaluation issues.

On the contrary, PSG-EXPERT includes a dedicated software module that has been developed in order to facilitate validation in a distributed environment, where several experts ought to be consulted at different physical locations. This software module is inspired in modern organisational workflow and groupware methodologies [3].

Several levels of validation can be defined: overall system evaluation concerning functionality, robustness, user-friendliness; knowledge integrity and correctness; diagnosis performance. All these levels are approached by the aforementioned software module, hereafter designated by CASE-BUILDER.

PSG-EXPERT is based on and extends the CLIPS expert system development tool with a set of features which fill in a gap that hampered the development of CLIPS-based expert systems since its creation, in 1984. Besides having a backward chaining mode and an explanation facility, which facilitate the development of diagnosis-oriented applications, PSG-EXPERT completely encapsulates CLIPS, under a Windows-based graphical user interface (GUI). The new tool includes graph-based graphical explanation facilities; a knowledge editor including a fuzzy fact editor and facts-rules integrity checking; a belief revision mechanism; a built-in case generator and a validation module. The graphical interface was developed in Delphi, a Windows rapid application development (RAD) tool. The interface connects to CLIPS using pipes for data transfer and semaphores for synchronisation.

Figure 1 summarises the system architecture and its main features. It comprises three main conceptual modules, namely: knowledge acquisition, diagnosis and validation. Each conceptual module corresponds to a different software product that can be used

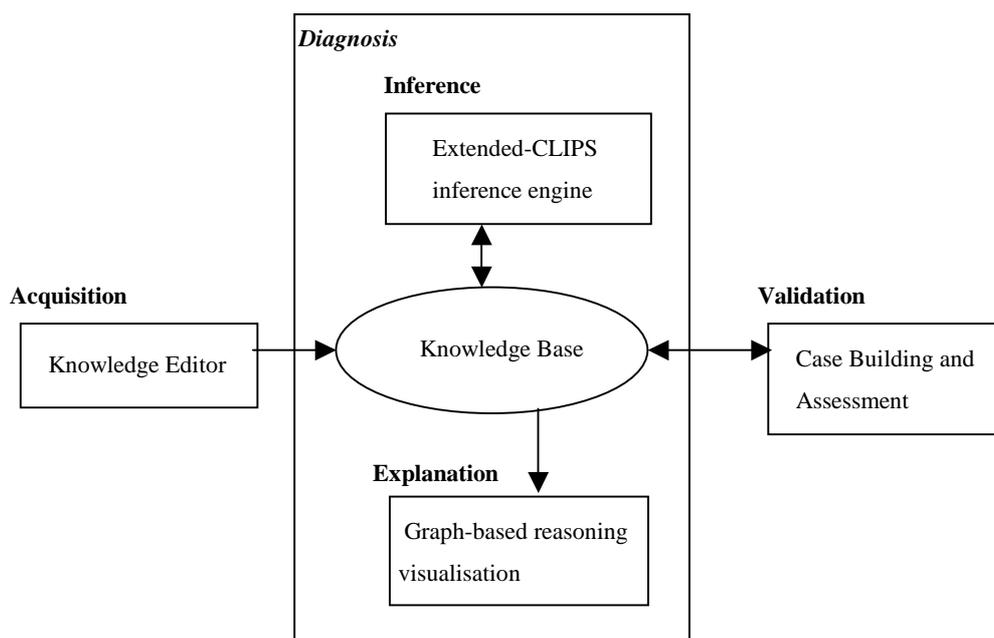


Figure 1. The PSG-EXPERT architecture.

independently of the other two.

The knowledge acquisition facility is implemented through a knowledge editor which enables the easy introduction and modification of knowledge, including facts, rules and criteria, as described in section 2.2. The diagnosis module includes the knowledge base, the inference engine and a graphic-based explanatory facility. Easy case building and assessment is possible through the validation module. The underlying knowledge representation scheme is described next.

2.1 Knowledge Representation

PSG-EXPERT is a rule-based expert system development tool. Besides the typical if-then rules, PSG-EXPERT has a special construct – the *Criterion* – that represents an expression of the form “if at least *n* conditions are satisfied then...”. A criterion can actually be rewritten as a set of if-then rules, but the abbreviation has been considered very useful by the domain experts because it matches the way their knowledge is commonly expressed. The structure of a *criterion* is shown in figure 2, together with the set of if-then rules it replaces.

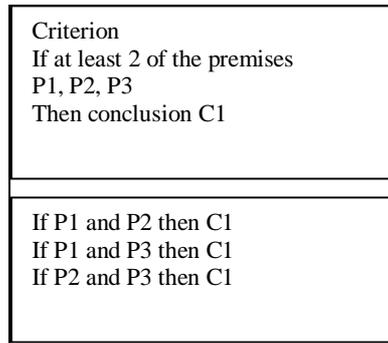


Figure 2. A Criterion in PSG-EXPERT is a new knowledge structure that replaces a set of if-then rules.

Similarly to CLIPS, PSG-EXPERT has an uncertainty management mechanism based upon *certainty factors* manipulation [1]. With respect to criteria, special combination formulas had to be used in order to process certainty factors, so that the resulting value is the same as if the set of equivalent if-then rules were used.

Any rule entered by the knowledge engineer can be used in forward and backward chaining, unlike some expert system development tools, such as GoldWorks [13], which require a specification of the reasoning direction in which the rules are going to be used. An extension to CLIPS was developed in order to support backward chaining.

In diagnosis it is often necessary to combine rules in favour of a certain conclusion with others which are against it. Moreover, there may be rules that completely preclude a diagnosis in spite of a lot of evidence in favour. This diagnosis-specific framework is provided in PSG-EXPERT by means of a built-in rule that combines the facts in favour and against a given conclusion, therefore establishing a model structure, and a task structure. This gives a certain guidance for knowledge acquisition and representation, unlike generic expert system development tools which leave all the structuring decisions to the expert. This model-based development methodology is characteristic of 2nd generation expert systems [2].

2.1.1 Facts

Facts in PSG-EXPERT actually represent the grounded beliefs contained in the knowledge base, being either crisp or fuzzy, and follow from the CLIPS fact representation scheme. A fact can be asserted in the knowledge base as a simple fact or as a structured fact. Structured facts are similar to frames in other representation languages. The constituents of structured facts are slots, with a name and a value. A default value may be associated with each slot. Similarly to the frame paradigm there is an inheritance mechanism, which propagates default values defined in the frame structures to their instances.

In order to accommodate the "conditions in favour" and "conditions precluding" philosophy, the system embodies a special type of structured facts: final diagnoses (disorders). From the user point of view these facts have the form

disorder <name> is <value>

with <value> being one of the following:

- likely – meaning that some condition in favour of a given diagnosis was satisfied;
- precluded – there are some conditions met that preclude the given diagnosis;
- present – the result of combining *likely* and *precluded* intermediate results by an internal rule.

2.2 The Knowledge Acquisition Facility

The system provides a user-friendly interface for the introduction, revision and maintenance of the knowledge-base which enables the application of adequate knowledge elicitation principles [6]. A structured environment for the edition of facts (crisp and fuzzy), rules and criteria was designed that performs basic rules-facts integrity checking. Figure 3 illustrates this module by showing a typical screen for the edition of fuzzy variables.

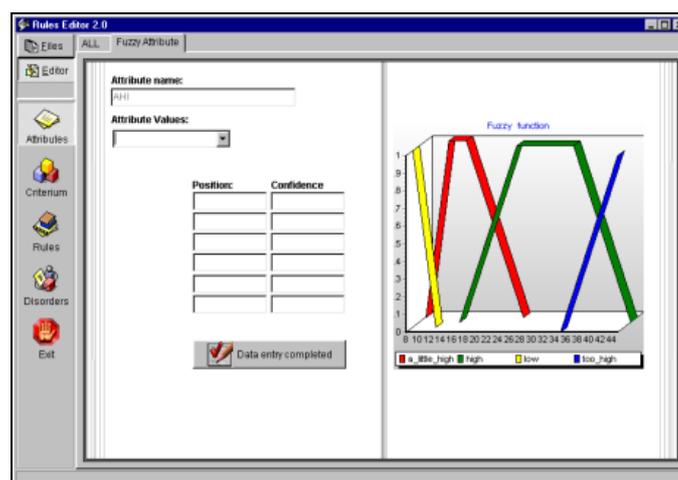


Figure 3. The Knowledge Editor main screen.

2.3 *The Diagnosis Task*

Diagnosis is a knowledge-intensive task. It requires both a knowledge-base including all relevant facts and rules, as well as an inference machine capable of logical reasoning and an adequate human-machine interface in order to justify the derived conclusions [1,12].

2.3.1 *The Inference Engine*

As indicated previously, PSG-EXPERT extends the forward engine provided by CLIPS and supports backward chaining. This allows a goal-directed type of reasoning which is often necessary for diagnosis purposes.

The backward chaining engine has been implemented as a modification of the basic forward chaining engine, adding the capability to manipulate a special structured fact called GOAL and a set of special facts called RULES that are actually the translation of forward chaining rules into a special structured fact syntax, to be used by the backward chaining engine. The translation is internal and transparent to the knowledge engineer.

2.3.2 *Dealing with Missing Data*

There are two ways of handling missing data under the PSG-EXPERT environment, that must be selected before triggering the diagnosis mechanism: (i)- assume normality; (ii)- ask the user. When the first option is selected, default values are set for each un-instantiated fact, as defined through the knowledge editor. While reasoning according to the “ask user” option, only those facts whose knowledge is essential to prove/disprove a given conclusion are prompted to the user through an adequate interface that enables questioning the system *why* that piece of information is being asked.

2.3.3 *The Explanation Facility*

PSG-EXPERT provides a graphical interface for explaining the reasoning process followed during either backward chaining or forward chaining. If, during the backward chaining process, the user is asked to provide the value of some variable he/she can ask the traditional “*why*” question, which the system will answer showing the rule which it is trying to prove. After both backward chaining or forward chaining the user can ask “*how*” the conclusions were deduced [6]. In that case a reasoning graph is presented.

Given the diagnosis orientation of this tool the reasoning graph produced by the explanation mechanism has some special nodes that represent diagnoses (or disorders – in the medical domain). The criteria used in the reasoning are not expanded into a set of if-then rules in the explanation graph because criteria reflect the way experts actually reason. Therefore, to differentiate criteria from if-then rules, different symbols are used in the graph. Facts are also shown explicitly in the reasoning chain using different symbols. Lines connecting all these elements show the conceptual links that have been established during the reasoning performed by the expert system.

The user can then, not only visualise the conceptual graph representing the reasoning process followed by the expert system during the diagnosis, but also inspect the details of the different elements of the graph (facts, rules, criteria and final diagnoses).

2.4 *The Validation Facility*

Validation is probably one of the most difficult parts of an expert system development life-cycle [2]. PSG-EXPERT includes a tool, called CASE-BUILDER, to facilitate the system validation by domain experts. This is mainly a set of organised screens through which the experts can provide all the relevant data – filling in the values of available variables – in order to describe a case or situation. Additionally, the same experts are required to indicate a diagnosis for each case they have described.

Case data or sample problems are critical to the evaluation of a knowledge-based system.

Test cases provide the major vehicle for testing and demonstrating the capabilities of a knowledge-based system. Unlike tracing, which is directed toward determining how the system arrived at its conclusions, test cases are directed toward determining whether the system can reach the proper conclusions.

Testing consists of feeding data into the system and comparing the system conclusions with what was expected.

These data can be of three types: (i) very specific cases which are required to test that specific situations are handled correctly. These data can be “manufactured”; (ii) other cases are required to permit the evaluation of the application’s handling of typical or common situations. Such data could be collected from real problems presented to human experts for solution; (iii) still other cases might be generated somewhat randomly to represent the breadth of situations the application may have to handle in the future.

The apparently best source of cases is historical data. This ensures that cases are realistic, representative of real problems and they are not contrived. However, using historical data often has three significant disadvantages: (i) the data are rarely complete; (ii) the historical solution might not have been adequate – e.g. not analysed by an expert; (iii) historical data are frequently on paper rather than in electronic form.

Generating new cases from scratch, directly into electronic format, has the advantage that data can be created fairly inexpensively and can be tailored to test particular components of an application. This approach may however have two serious drawbacks: (i) solutions must be created for each test case generated. The expert must indicate the solution – the diagnosis can not be “generated” as the input data for the case was. Furthermore, (ii) since these cases can be “unrealistic”, the expert may not be eager to invest the time and personal effort to develop the solutions.

The conclusion is that, unfortunately, generating case data is much easier than generating the solutions that go with those data. Therefore the case-data problem should be viewed as the case-solution problem; the ability to generate solutions and not data is what tends to limit the case data that can be developed for an application.

CASE-BUILDER provides the advantages of case generation described above, minimising the two main problems that have been pointed out. Firstly, this software tool provides not only the means for data input but also an easy and fast way to identify both the solution and the respective justification. It also allows experts to provide advice “on the side” which can later on be used by the development team to refine the system knowledge base. Secondly, since the data is developed by an expert, it seldom is “unrealistic”.

CASE-BUILDER provides a mechanism that allows multi-expert validation of a knowledge system. The first step in this process consists of an expert entering the data and providing his/her own diagnosis. Then, using CASE-BUILDER mechanisms for case communication, the case that was simulated by the first expert circulates amongst a group of experts who can provide their own evaluation of the case but without each of them being able to see the others’ diagnoses. All the communication is made using a special e-mail utility embedded in CASE-BUILDER. Sending cases only requires a button click. After all the experts have analysed a case it then is sent to the development team, who can use the

case solutions to assess the accuracy of the system and to improve it if possible, by adopting some of the experts' recommendations.

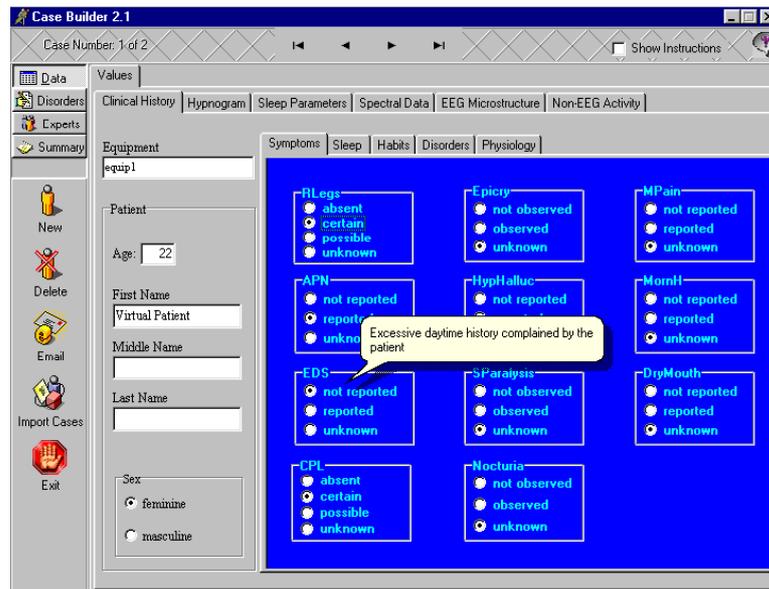


Figure 4. Data entry screen in Case-BUILDER.

Although this tool is specific for supporting the validation of expert systems made for diagnosis purposes, it is not specific of the medical domain. All examples given here are though extracted from an application on the medical domain.

The list of variables presented in CASE-BUILDER, and their respective possible values, is easily configurable through a database, avoiding the need for modification and recompilation of the program every time a new variable is added or removed.

Figure 4 shows the data entry screen. Each case must be assigned one or more diagnosis consistent with the data (see figure 5). Introduction of the diagnosis is made by clicking on the *Disorders* icon on the left column and selecting one of the available diagnosis, on the right side of the screen.

There are also dedicated spaces where the expert can enter free text to explain the diagnosis (*Justifications*) and to include additional comments considered of interest (*Observations*) useful for the knowledge engineers further develop the knowledge base.

Furthermore, cases thus entered are saved in a local database, being immediately accessible in the diagnosis module. The expert can hence simulate particular situations to be evaluated by the system and then analyse its reasoning in these case-studies, a process particularly useful and efficient for knowledge evaluation and revision.

This provides a means for evaluating diagnosis consistency both in between experts and by the system.

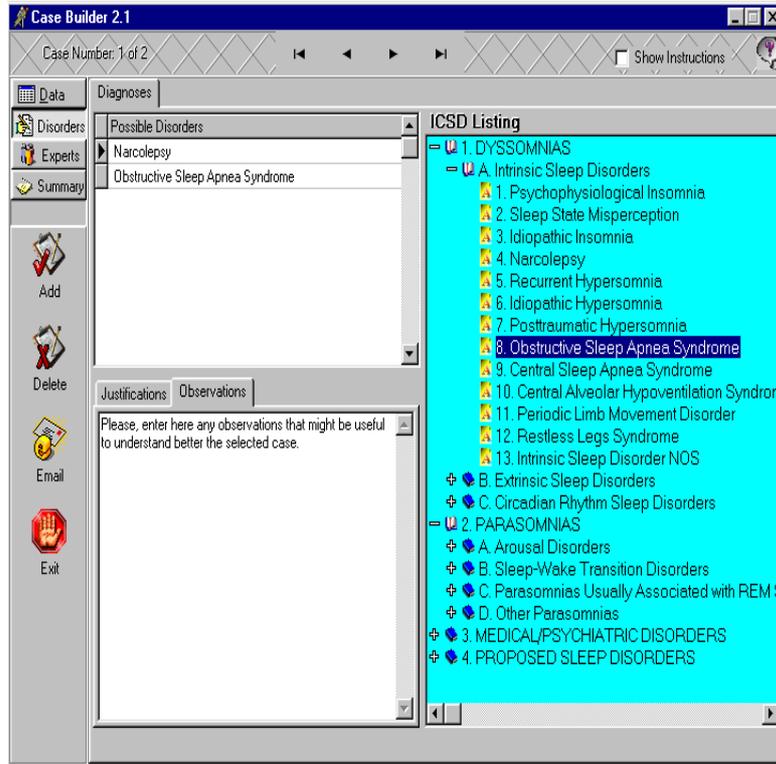


Figure 5. Diagnosis entry screen in Case-BUILDER.

3. DIAGNOSING WITH PSG-EXPERT: THE USER'S PERSPECTIVE

The previously described specialized environment was developed to be used as the technical basis for developing an expert system for the diagnosis of sleep disorders using mainly polysomnographic and clinical data [9,10,11].

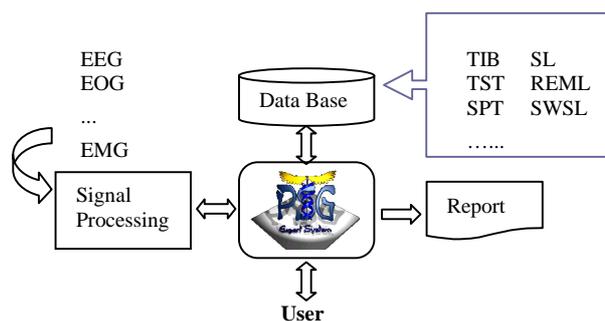


Figure 6: Inputs and outputs of the PSG Expert.

Figure 6 gives a schematic description of the expert system environment. Biological signals are automatically scored by a core of signal processing algorithms. The resulting polysomnographic data and additional clinical data are stored in a neurological database. Given a particular recording, the system identifies the set of possible sleep disorders consistent with the data and produces a report similar to what is produced nowadays by human experts.

The major features of the system are as follows:

- Browsing of patients information and viewing of polysomnographic data records, from either local or remote databases (see fig. 7).
- Structured navigation on Sleep Disorders descriptions according to ASDA definitions.

- Internet link to related html pages.
- Diagnosis of sleep disorders consistent with polysomnographic data, supporting:
 - Fuzzy rule-based reasoning operating both in data-driven (forward chaining) and goal-driven (backward chaining) modes and including a belief revision mechanism
 - HOW and WHY explanation of the reasoning process
 - Graph-based graphical explanation facilities
 - Handling of missing data and uncertainty
- Production of a report summarising the major findings on data records.
- Built-in CaseBuilder facilities for the introduction of new patient records into the database or to data exchange and system validation through transparent email communications.

The final version of this expert system was submitted to a validation phase, using the CASE-BUILDER tool described in the previous section.

3.1 *Polysomnographic Signals and Polysomnographic Data*

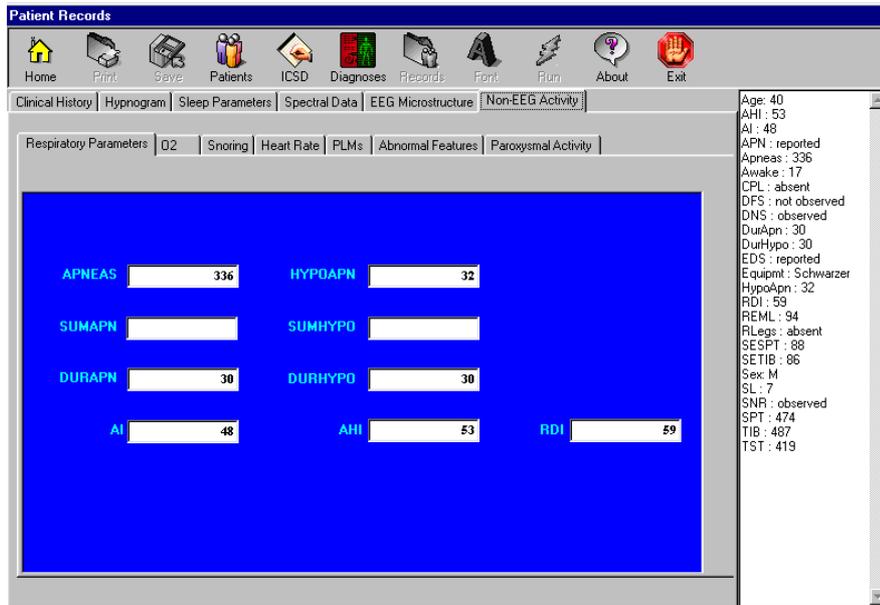
Biosignals most frequently used to describe human sleep are: the electrooculogram (EOG) - used to monitor eye movements; the electroencephalogram (EEG), two or more channels, from which brain electrical activity is assessed; the electromyogram (EMG) - chin and legs - used to detect limb movements and muscle tonus; respiration related signals, such as the respiratory flux and effort, O₂ saturation, or snoring; and the electrocardiogram (ECG). Analysis of these signals enables the discrimination between the three distinct physiological states: wakefulness, rapid-eye-movement sleep (REM) and slow-wave or non-REM sleep.

The global features of sleep dynamics, i.e. sleep macrostructure, have been summarized in a graphical representation, known as the *hypnogram*, of the evolution of sleep along the night as a sequence of sleep stages identified according to standard criteria established by Rechtschaffen and Kales. From this data global parameters characterizing sleep quality, sleep structure and cyclicity have been defined. It is recognized, however, that reliable diagnosis cannot be achieved based only on this type of global and limited information. Within each stage the EEG exhibits a rich variety of patterns of transient and background activity, that compose the sleep EEG microstructure. Parameters derived from EEG spectral analysis and time domain analysis have therefore been included as important information to consider for diagnosis purposes. Non EEG activity derived parameters complete the set of variables obtained from the processing of the above polysomnographic signals used for diagnosis of sleep disturbances.

The types of data incorporated by the PSG-EXPERT are variables are organized into the following categories: clinical history; hypnogram data; sleep parameters; spectral data; EEG time related activity; non EEG activity.

Figure 4 showed a typical screen exhibiting the structured presentation of the above types of information, presented in the phase of data entry through the CASE-BUILDER sub-module. Figure 7 depicts another example concerning the visualization of a patient's record.

Figure 7: Screen for visualization of patient data concerning respiratory parameters.



3.2 Application Case-Study

The case-study concerns the analysis of a patient case which data are partly in a database. Two diagnostic strategies are possible: testing a specific disorder (hypothesis) – *backward chaining reasoning* - or finding all possible diagnosis consistent with the given – *forward chaining reasoning*. When forward chaining is performed no additional data is used; however, in the case of backward chaining reasoning, the system may ask for some additional data, directly correlated with the diagnostic hypothesis under analysis, which may be entered by the doctor, as needed.

Figure 7 shows the explanation graph produced after a diagnosis. In this case the system performed a forward chaining reasoning and it identified a disorder (obstructive sleep apnea) with the certainty factor that is indicated on the upper-left area of the figure.

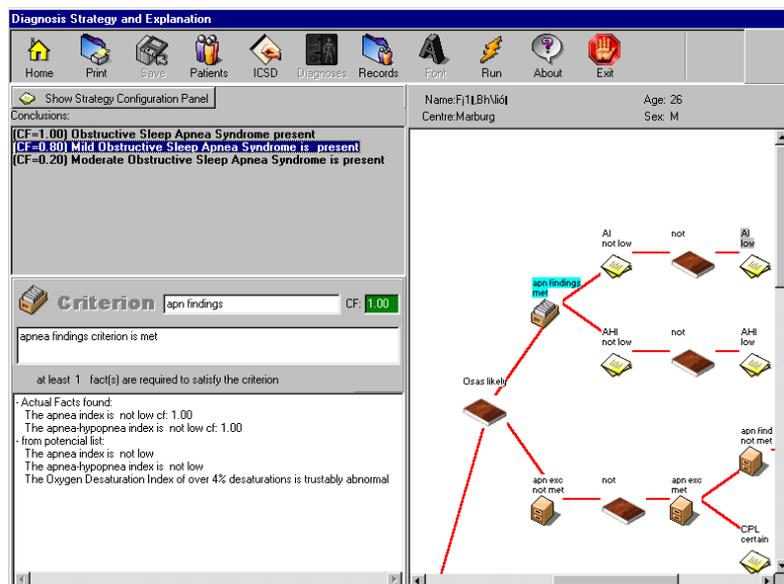


Figure 5. Observing the diagnosis result and the reasoning process.

The complete reasoning chain is shown in a scrollable window, on the right part of the screen. When the user clicks on a graph element, detailed information about it can be seen at the bottom left part of the screen.

5. CONCLUSIONS

This paper presented technical aspects of PSG-EXPERT, an integrated environment for the development of diagnosis-oriented knowledge based systems applied to sleep disorders. This system provides graphical support for every phase of the development cycle, from knowledge acquisition and edition to reasoning explanation and knowledge-base validation.

Using the Delphi language and encapsulating CLIPS, PSG-EXPERT is a Windows based software system that extends the later shell with the following features: backward chaining engine; graph-based explanation facilities; knowledge editor including fuzzy facts and rules editing with rules-facts integrity checking; belief revision mechanism; and a built-in case generator and validation module.

In addition to the “if-then” structures, common to all rule-based systems, knowledge representation supports a new type of concept –the *criterion* – that enables a more compact and natural way of expressing rules of the form: “if at least n conditions are satisfied then...”.

Furthermore, the system is designed with built-in structuring decision rules that permits the combination of evidence in favour and precluding a given diagnosis, thus establishing a model structure and a task structure.

Unlike usual commercial or public domain shells, a key feature of this system is the validation module, an interactive subsystem that enables fast creation and interchange of commented cases that serve the purpose of diagnosis consistency and performance evaluations. This enhancement over traditional systems promotes the use of groupware techniques to expert system development and validation.

Concerning its application as an auxiliary diagnosis system in the domain of sleep disorders, versatility of the system was partially illustrated through a set of interaction examples. Distinctive features of the system include: the exploration of objective polysomnographic data in the reasoning process, complementing and/or confirming clinical data; representation of expert knowledge in terms of natural language like diagnosis rules; a graphical interface exposing the reasoning process; distributed environment for data storage and retrieval; internet links to related pages; uncertainty modelling and belief revision; production of reports; a case building and validation module promoting groupware techniques; a global user-friendly interface supporting three main types of system usage – design and knowledge maintenance, diagnosing and validation.

REFERENCES

- [1] Buchanan, B. and E. Shortliffe (Eds), *Rule-Based Expert Systems*, Addison-Wesley, 1984.
- [2] Burnet, E. and N. Scharenborg, The KADS method: Industrial Applications and Prospects for Expansion, in *the 12th International Conference Avignon'92*, 1992.
- [3] Brenner, W. Agents as Tools of the Information Society, in *Intelligent Software Agents*, Brenner, Zarnekow and Wittig (Eds), Springer-Verlag, 1998.
- [4] CLIPS Version 6.04, CLIPS Reference Manual, Software Technology Branch, Lyndon B. Johnson Space Centre, 1997.
- [5] Diagnostic Classification Steering Committee, Thorpy M. J. (Chairman), *International Classification of Sleep Disorders: Diagnostic and Coding Manual (Revised)*, Rochester, Minnesota: American Sleep Disorders Association, 1997.
- [6] Diaper D.(Ed.), *Knowledge Elicitation: Principles, Techniques and Applications*, Ellis Horwood, Ltd, 1989.
- [7] Duda, R., and .P. E. Hart (Eds), *Pattern Classification and Scene Analysis*, John Wiley & Sons, 1973.
- [8] EXSYS Professional, User's Guide, EXSYS Inc., 1994.
- [9] Filipe, J., Fred A., Paiva, T., Partinen M. A Second-Generation Expert System for the Diagnosis of Sleep Disorders using Polysomnographic Data, in *Proc. Int'l Conf. on Information Systems Analysis and Synthesis*, pp. 296-300, 1996.
- [10] Filipe, J., Fred A., Fernandes M. GDOS – A Graphical Diagnostic-Oriented Expert System Development Tool, in *Proc. Int'l Conf. on Enterprise Information Systems*, pp. 211-218, 1999.
- [11] Fred, A. and Filipe, J. GMED – A Case-Study of a Specialized Graphical Medical Diagnosis Shell Application, in *Proceedings of IASTED International Conference on Artificial Intelligence and Soft Computing*, 1999.
- [12] Giarratano, J. and G. Riley (Eds), *Expert Systems: Principles and Programming*, PWS Publishing Company, Boston, 1994.
- [13] GoldWorks, User's Guide, Gold-Hill Inc., 1993.
- [14] KEE – The Knowledge Engineering Environment, Intellicorp, 1987.
- [15] Russel, S. and P. Norvig (Eds), *Artificial Intelligence: A Modern Approach*, Prentice-Hall, Inc., 1995.