MUNIN - A Causal Probabilistic Network for Interpretation of Electromyographic Findings*

Steen Andreassen[1,2], Marianne Woldbye[1], Bj0rn Falck[1,3] and Stig K. Andersen[1,2].

[1] Nordjysk Udviklingscenter, Badehusvej 23, DK-9000 Aalborg
[2] Institute for Electronic Systems, Aalborg University, Strandvejen 19, DK-9000 Alborg
[3] Turku University Hospital, SF-20520 Turku

ABSTRACT

Experience gained through building a causal network for interpretation of electromyographic findings has shown that probabilistic inference is a realistic possibility in networks of non-trivial size. The use of nodes with many internal states has made it possible to make a conceptually simple and compact representation of knowledge. "Deep knowledge" in the form of pathophysiological models are used to reduce the problem of estimating thousands of conditional probabilities to a manageble size. The network has built-in mechanisms that will detect when the network is confronted with a situation outside the limits of its own knowledge and it handles conflicting evidence in a simple and consistent way.

INTRODUCTION

In some medical expert systems, causal networks have proven themselves as a helpful tool in the organisation of the knowledge of a domain (for references see e.g. Szolovits 1982). Concepts in the domain are represented as nodes in the network and their interaction is represented as causal links between the nodes.

We have chosen electromyography (EMG), the diagnosis of muscle and nerve diseases through analysis of bioelectrical signals from muscle and nerve tissue, as our application domain. Out of 13 distinct types of knowledge, required to carry out an EMG examination, 7 can conveniently be represented in a causal network (Andersen et al. 1986).

Only few expert systems have used a probabilistic approach, and to our knowledge none of the systems based on causal networks have exploited recent progress in algorithms for propagation of probabilistic evidence (Kim and Pearl, 1983; Pearl, 1986).

This paper reports on experience gained by building a small, but non-trivial prototype expert system for EMG, using a causal probabilistic network, MUNIN * - MUscle and Nerve Inference Network.

After outlining the functions and limitations of the prototype several issues will be discussed:

* This work is supported in part by the EEC ESPRIT programme, project P599

** Coincidentally, according to Norse mythology Munin is one of two ravens, whispering intelligence into the ear of the god, Odin.

- the knowledge representations
- the experience gained with knowledge acquisition and verification
- the interaction of empirical (surface) knowledge and causal understanding as exposed in textbooks (deep knowledge)
- handling of conflicting evidence
- performance of the system, when confronted with evidence lying outside the knowledge of the network

FUNCTIONS AND LIMITATIONS OF THE PROTOTYPE

In earlier publications (Andersen et al. 1986, Andreassen et al. 1986) it was proposed that a causal network contains the information necessary for a unified approach to three of the main tasks of a medical expert system: diagnosing, planning of data acquisition, and explanation of the systems reasoning. In this paper we shall only deal with diagnosis, while implementation and some aspects of planning and explanation are considered by Jensen et al. (1987). At the beginning of a diagnostic session the disease node is initialised with a priori probabilities corresponding to the observed frequencies of the diseases in patients referred for EMG examinations (figure 1). The number of diseases is restricted to three, each with two to four states, corresponding to gradations and/or different varieties of the diseases. In addition the patient may be in one of the states normal (no disease with neuro-muscular symptoms) or other (a neuro-muscular disorder other than the three mentioned diseases), giving a total of eleven different "disease" states. An algorithm for propagation of evidence in causal networks was developed by Kim and Pearl (1983). The algorithm was adapted to this network and supplemented by a method for coherent initialisation of probabilities (Jensen et al. 1987).

Using these methods the a priori distribution of diseases can be used to generate expectations for the pathophysiological changes caused by the diseases. The pathophysiological changes in a given muscle are described by eight pathophysiological nodes, MU.LOSS through DENERVATION, each of them with from two to nine states. Since almost half of the patients are "normal", the expectations for the pathophysiological nodes are largely normal. The pathophysiological nodes in turn generate expectations for the 15 findings nodes through their causal links, either directly or through an intermediate node (MUP.CONCLUSION). MUP.CONCLUSION does not have a natural patho-physiological interpretation. It only serves the purpose of integrating the information from the three findings MUP.AMPLITUDE, MUP.DURATION and MUP.POLY-PHASIC. These three findings are all obtained from the same EMG test: analysis of Motor Unit Potentials.

Beyond the already mentioned reductions on the number of diseases, the prototype is also restricted in other ways: multiple simultaneous diseases are not considered and the network does not handle measurement of nerve signals, which are as important as measurement of muscle signals. Furthermore, it only considers findings from one muscle. The interesting interaction between the causal network and the ncuroanatomical knowledge, which has a mainly topological nature and is not readily represented in a causal network, is therefore not considered in this paper. The authors think that solutions to each of these problems can be found, without violating the rigorous probabilistic approach used for the prototype.

The diagnostic task consists of adjusting the probabilities in all nodes as the findings are entered into the findings nodes. In figure 2 findings corresponding to a typical case of "moderate chronic axonal neuropathy" have been entered. A finding entered into a findings node is indicated by a broken horizontal 100% bar. The network correctly indicates a large probability for moderate to severe axonal neuropathy, it generates distributions for the pathophysiological nodes that are consistent with "moderate chronic axonal neuropathy" and offers predictions of the outcomes of the remaining findings, should the physician chose to perform the appropriate EMG-tests.

REPRESENTATION OF UNCERTAIN KNOWLEDGE

The medical knowledge is embedded in the causal network in three ways. Through the choice of the number and character of the nodes, through the assignment of causal links between the nodes, and since we have chosen a probabilistic approach, through the conditional probabilities associated with the causal links.

As an example, consider the causal links from the disease node to the pathophysiological node, loss of motor units (MU.LOSS), which reflects the percentage of nerve fibres that still survive and reach the muscle. A motor unit is the muscle fibres, typically several hundred, that are innervated by one and the same nerve fibre. In figure 3 the conditional probabilities for MU.LOSS, given some of the states in the disease node are shown.

The table states that 92% of all "normal" patients will have "no" loss of motor units, 5% "moderate" and 1% "severe" loss of motor units, while virtually nobody (0.1%) will have "total" loss of motor units. 2% will present with a different picture that can not be described in terms of "MU.LOSS"(see section on IGNORANCE AND CONFLICTS).

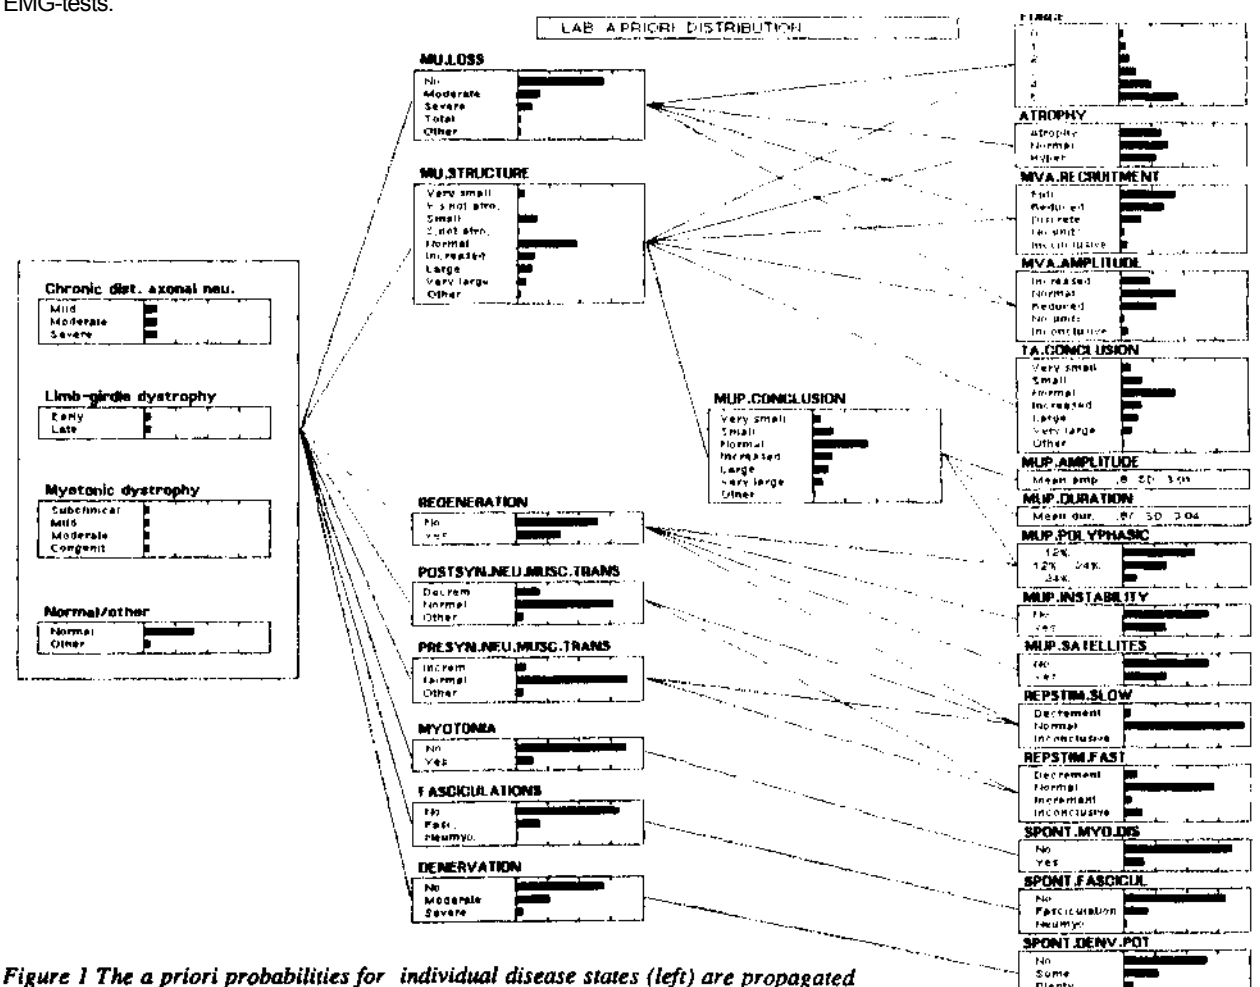

Figure 1 The a priori probabilities for individual disease states (left) are propagated through the layer of pathophysiological nodes (middle) to the findings nodes (right). The length of the horizontal bars indicates the probability of the states in the node.

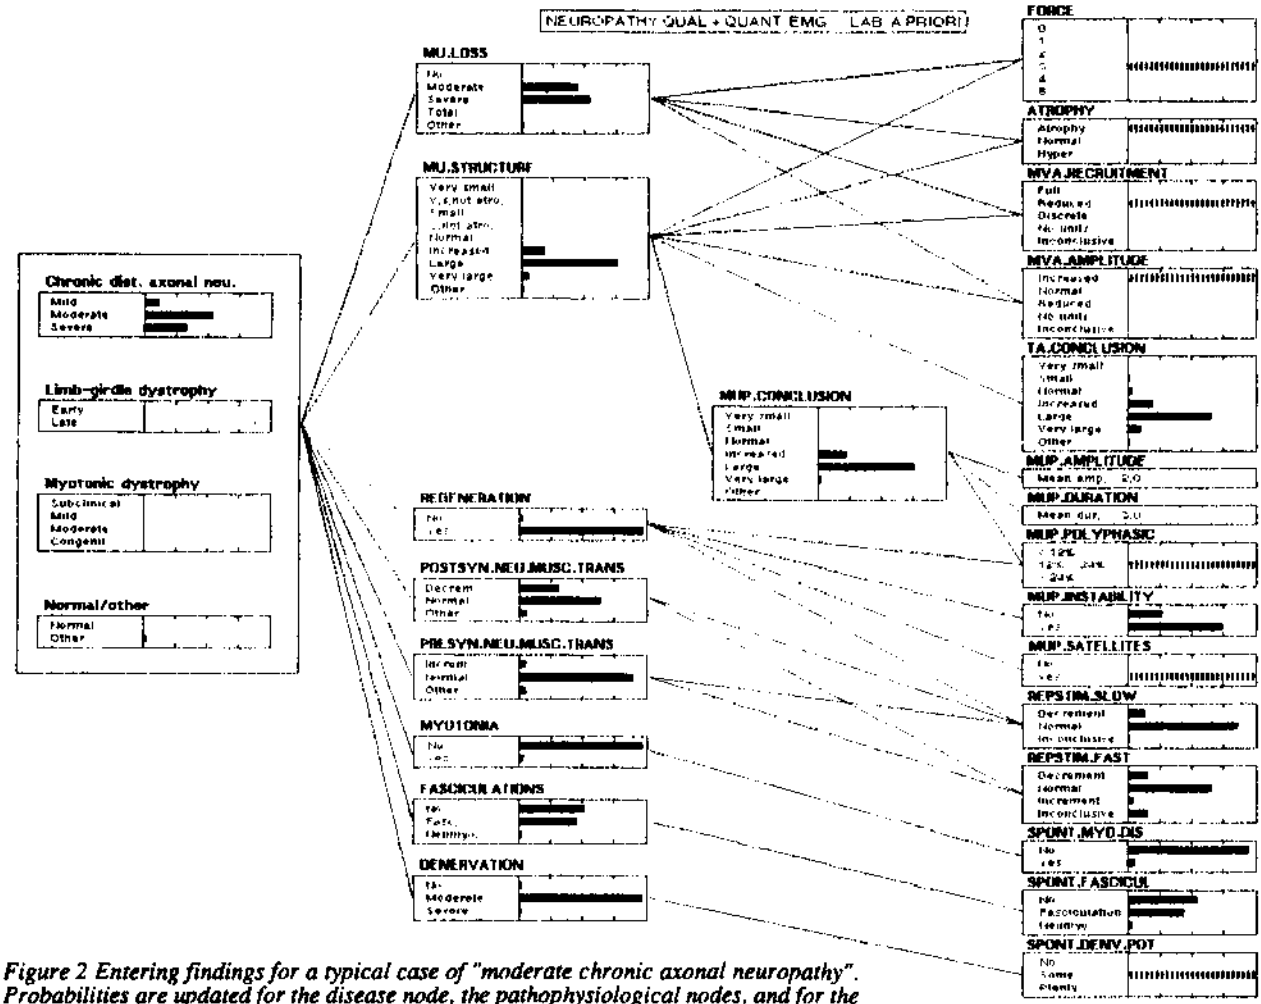

Figure 2 Entering findings for a typical case of "moderate chronic axonal neuropathy". Probabilities are updated for the disease node, the pathophysiological nodes, and for the remaining findings.

	MU.LOSS					Σ
DISORDER	no (%)	mod (%)	sev (%)	total (%)	other (%)	(%)
Normal	91.9	5.0	1.0	0.1	2.0	100
Chronic axonal neuropathy:						
mild	60.0	33.0	4.0	1.0	2.0	100
moderate	16.0	64.0	16.0	2.0	2.0	100
severe	2.0	25.0	61.0	10.0	2.0	100
Other	33.0	30.0	30.0	5.0	2.0	100

Figure 3 The conditional probabilities P(MU.LOSS I DISORDER) for the states of MU.LOSS, given 5 different disorders.

As can be seen from the second line, patients with "mild chronic axonal neuropathy" are much more prone to loss of motor units, with only 60% having "no" loss of motor units. Moderately and severely affected patients are even worse off. The last line "other" represents the expected loss of motor units, when something else but unknown is wrong with the patient. The relative lack of information is reflected in a relatively even distribution of the probabilities over all the possible states. Examples of how the conditional probabilities are acquired and verified are given in the next section of knowledge acquisition and verification.

For the findings with continuous outcomes the probabilities are replaced by probability densities. In figure 4 the conditional probability densities which are assumed to have a normal distribution are plotted for each of the states in MUP.CONCLUSION.

Some of the nodes have more than one parent, i.e. there are two or more "parent nodes" with causal links to a "child node". Examples are "FORCE" and "ATROPHY" that both have MU.LOSS and MU.STRUCTURE as parents. This requires filling out a conditional probability matrix that specifies P(FORCE I MU.LOSS, MU.STRUCTURE) for all 6x5x9 combinations of FORCE, MU.LOSS, and MU.STRUCTURE. From the point of view of knowledge acquisition it is not a trivial task to determine the 270 conditional probabilities in this three dimensional matrix.

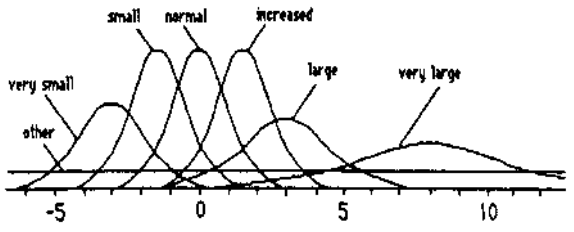

Figure 4 Conditional probability distribution for

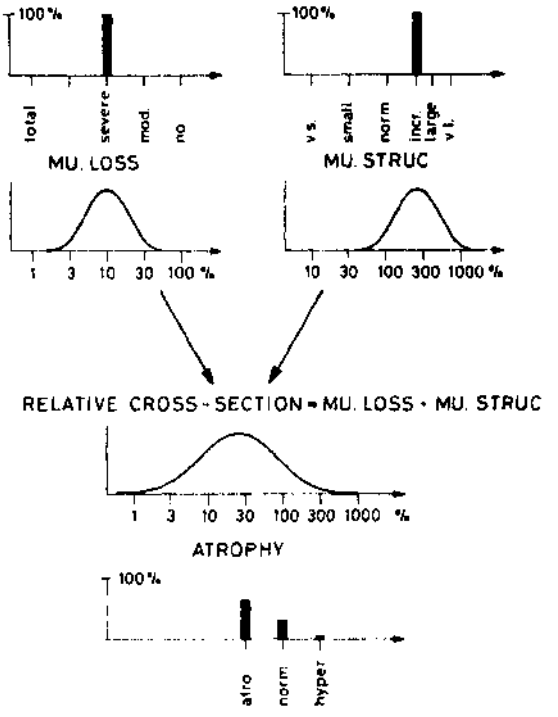

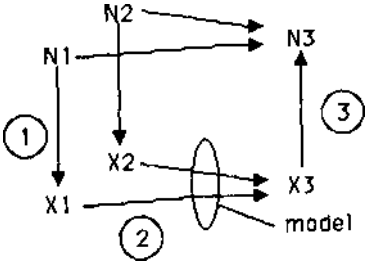

Once models have been developed for all nodes, we have tried to verify the network in two different ways. One way is to generate expectations corresponding to a single disorder. For example, the expectations corresponding to "moderate chronic axonal neuropathy" may be presented to a medical expert (figure 7) who is asked to identify discrepancies between his expectations and the expectations generated by the network. The other way of verifying the network is by entering findings typical of different diseases, and asking the medical expert to identify differences between the probabilities computed by the network and his own diagnosis. In figure 2 findings corresponding to a case of chronic axonal neuropathy were entered. In that case the probabilities computed by the network were satisfactory.

Discrepancies between the network and the medical experts lead to revision of the model parameters and occationally to revision of inaccurate or incomplete models. Occationally, it may even be necessary to modify the structure of the network, adding or deleting states or nodes. The current version of the network represents the third major revision. In our hands the revision process has so far been "benign", with improvement in performance for each revision and without unexpected side-effects of revisions. When the network is expanded to handle a range of clinically realistic cases, a more formal testing involving a number of clinical experts is planned.

IGNORANCE AND CONFLICTS

Even if a dedicated effort is made by the builders of the MUNIN system to collect and describe a large number of diseases to the system, there will still be a residual number of diseases that arc either unknown or incorrecdy described to the system. How should the system behave when confronted with a disease unknown to it? If there is a disease known to the system that fits the findings of the case poorly, although with a better fit than all other diseases, then Bayesian systems tend to give a strong statement in favour of the disease with the least poor fit. To avoid this behaviour we have introduced the state "other", both in the disease node and in some of the pathophysiological nodes. In figure 8 a set of findings have been entered. The network is confused by these findings (so are the authors), and the network indicates its reservations by giving "other" a high probability.

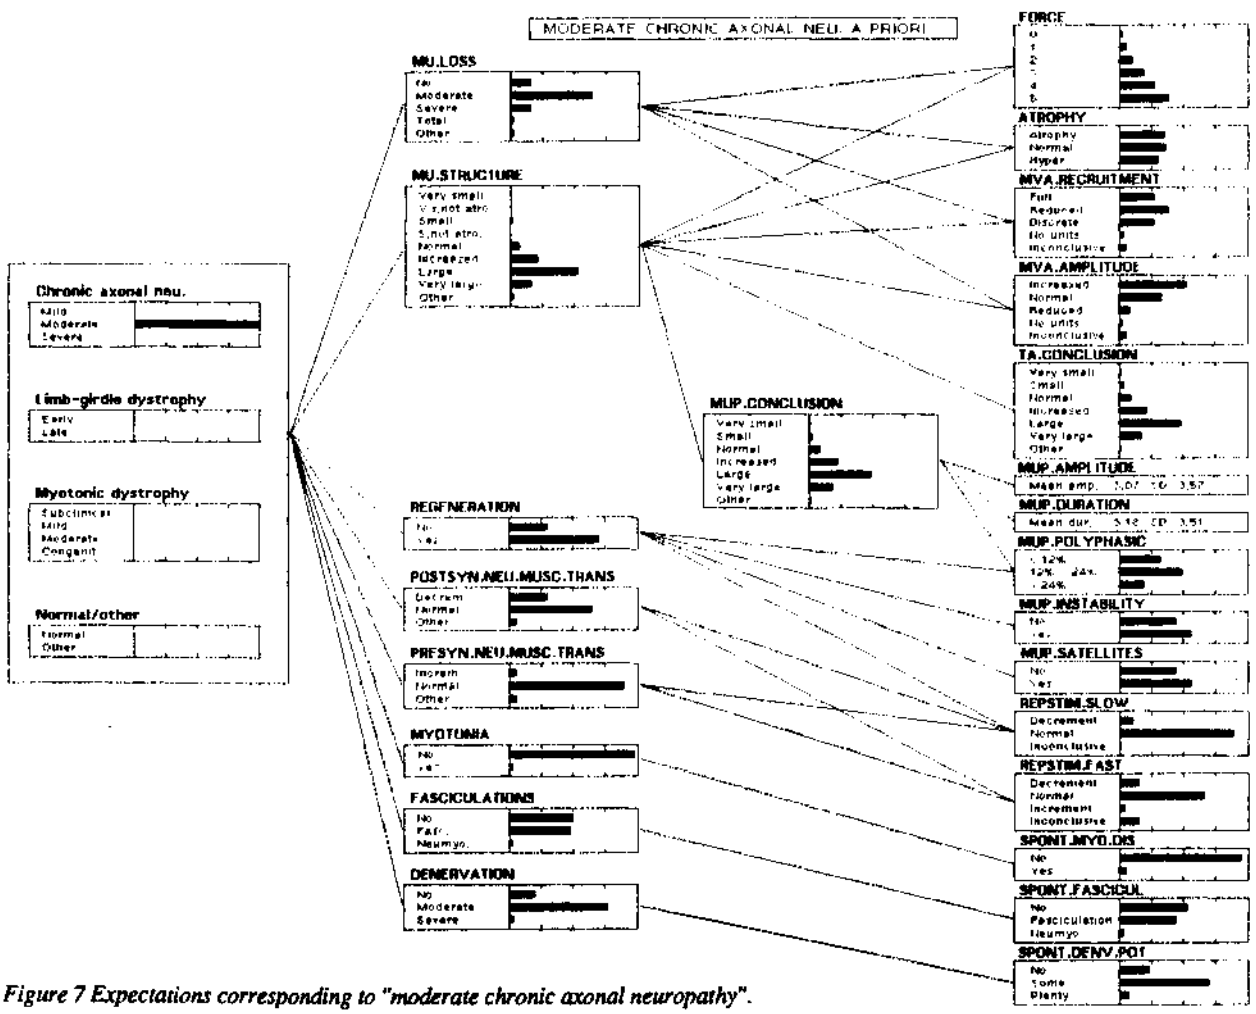

Figure 7 Expectations corresponding to "moderate chronic axonal neuropathy".

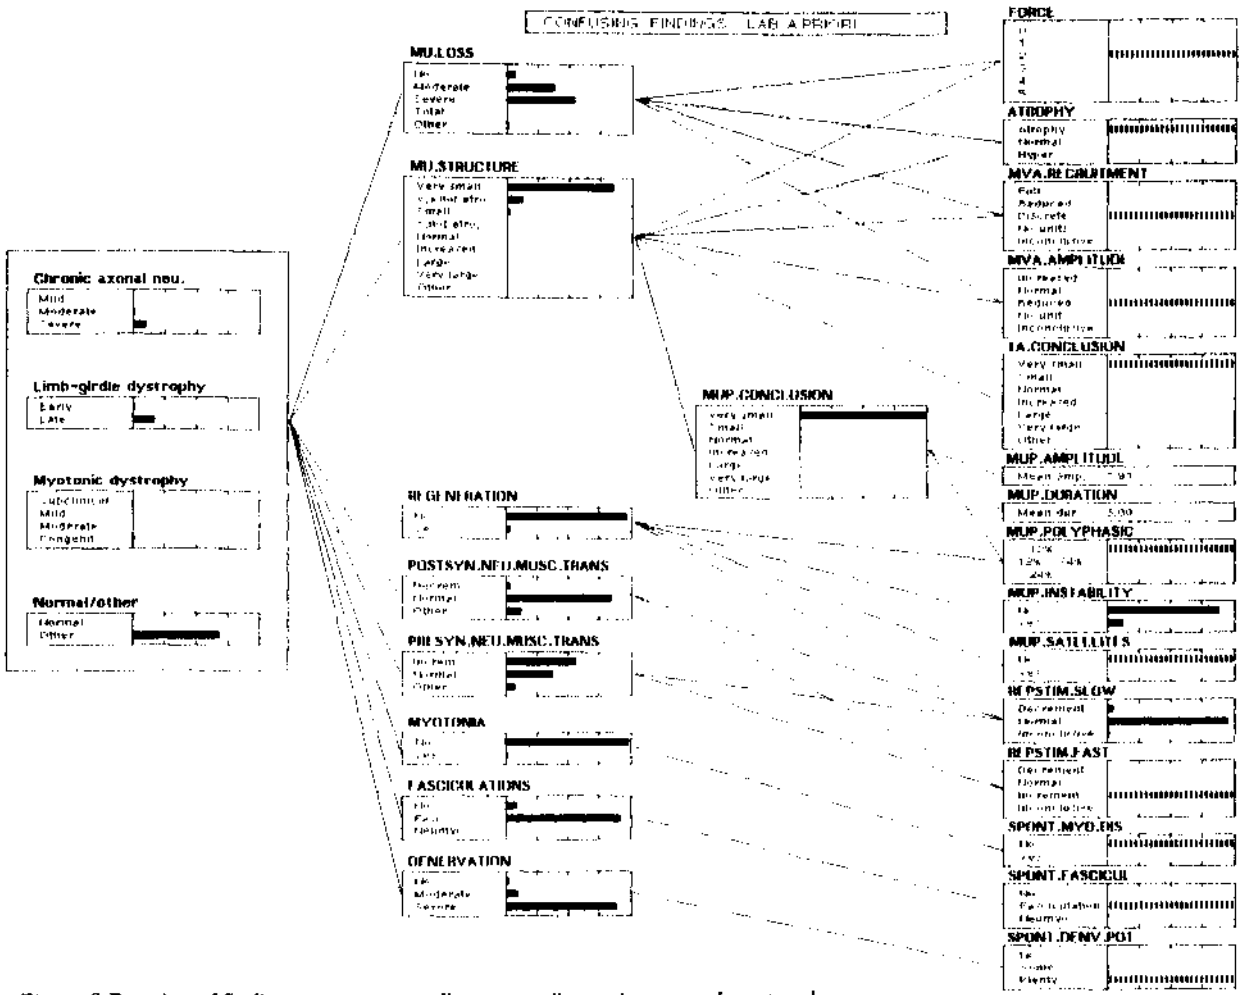

Figure 8 Entering of findings not corresponding to any disease known to the network.

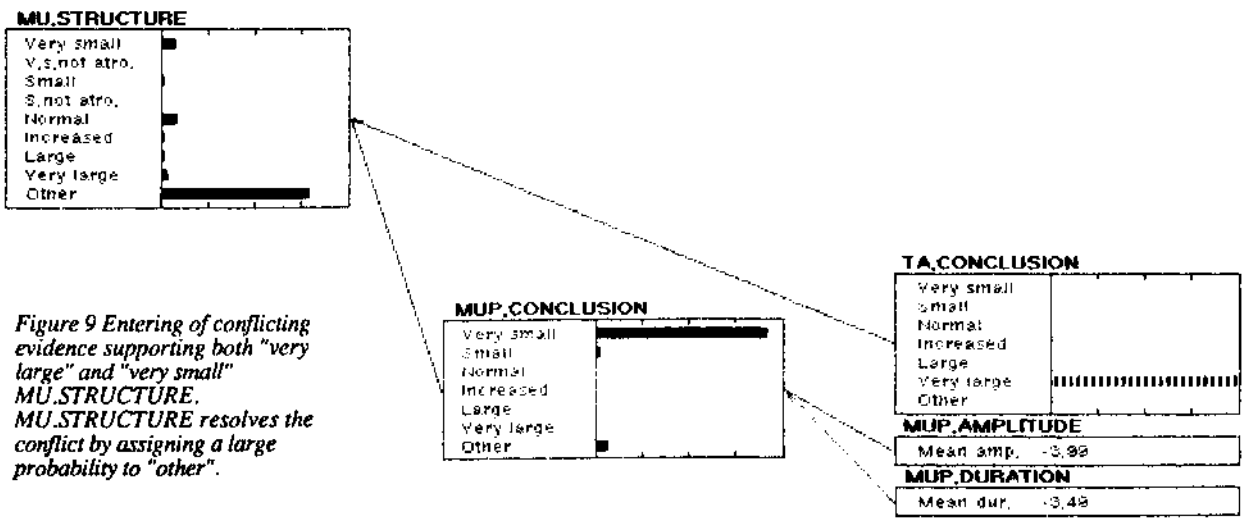

Figure 9 Entering of conflicting evidence supporting both "very large" and "very small" MU.STRUCTURE. MU.STRUCTURE resolves the conflict by assigning a large probability to "other".

The network may also get confused at the pathophysiological level, if conflicting findings are entered. In figure 9 MU. STRUCTURE receives input from TA.CONCLUSION supporting "very large" and from "MUP.CONCLUSION" supporting "very small". MU.STRUCTURE states its confusion by assigning a high probability to "other". There are two interpretations of the "other" state. Maybe the findings entered into the network faithfully represents the states of the muscle, although this would be very difficult to imagine from a medical point of view. In that case a "hole" in the knowledge of the network has been uncovered and one or more of the models in the network will have to be modified. It may even be necessary to add new states to excisting nodes, to add new nodes or to add new causal links.

The alternative interpretation is that erroneous findings have been entered. Generally, the appearance of "other" in a node signals that conflicting and possibly erroneous input has been given to the network.

Whichever interpretation is correct the situation is handled in a reasonable way. Once the status of MU.STRUCTURE is "other", it does not lend support to any disease. Basically, the node is largely ignored until the conflict is resolved.

CONCLUSION

A network for the interpretation of EMG finding has been constructed. We expect a network of this type to be an important building block in an expert system for EMG. Although the network is small and in its current form has only limited functions, it has allowed us to reach a number of conclusions:

1) With present algorithms for propagation of evidence in causal probabilistic networks, probabilistic inference is a feasible approach. Since the computation time of the algorithms is increasing approximately linearly with the number of states in the network, we expect that probabilistic inference can also be used in networks considerably larger than the current network.

2) The shift from nodes with only two states (yes,no) to nodes with multiple states has given a conceptual simplicity that makes knowledge acquisition and verification easier. It also makes the knowledge representation very compact.

3) The use of "deep knowledge" in the form of models has reduced the almost intractable problem of estimating thousands of probabilities to the much more tractable problem of adjusting a much smaller number of model parameters. The models have the added virtue that they can be explained through pathophysiological reasoning similar to the reasoning done by an expert.

4) Lack of knowledge in the system and conflicting evidence is handled in a simple and consistent way by adding the state "other" to some of the nodes. This way the network can signal, when it reaches the limits of its knowledge.

ACKNOWLEDGEMENTS

The work on the MUNIN expert system is part of ESPRIT project P599, An Expert Assistant for Electromyography, and is partially supported by the European Commision. All members of the MUNIN group, Annelise Rosenfalck at NUC (Aalborg), and Finn Verner Jensen, Uffe Kjamilff and Anders Rieck Sorensen at Judex Datasystems (Aalborg) are thanked for their valuable contributions. The other partners in the ESPRIT project, Computer Resources International (Copenhagen), National Hospital of Nervous Diseases (London) and Logica (Cambridge) are thanked for valuable discussions.

REFERENCES

Andersen, S.K., Andreassen, S., Woldbye, M., "Knowledge Representation for Diagnosis and Test Planning in the Domain of Electromyography." Proc. 7th European Conference on Artificial Intelligence. Brighton, UK, 1986, pp 357-368.

Andreassen, S., Andersen, S.K., Woldbye, M., "An Expert System for Electromyography." Proc. of Novel Approaches to the Study of Motor Systems. Banff, Alberta, Canada, 1986, pp 2-3.

Kim, J.H. and Pearl, J. "A Computational Model for Causal and Diagnostic Reasoning in Inference Systems." Proc 8th Joint International Conference on Artificial Intelligence. Karlsruhe, West Germany, 1983, pp 190-193.

Jensen, F.V., Andersen, S.K., Kjaerulff, U., Andreassen, S. "A Causal Network Prototype in the Domain of Electromyography - Which Functionality can be Achieved by Using a Probabilistic Approach?" Research Report R-87-8. Institute of Electronic Systems, Aalborg University, Denmark. 1987.

Pearl, J. "Fusion, Propagation, and Structuring in Belief Networks." Artificial Intelligence. 29, 1986, pp 241-288.

Szolovits, P. Artificial Intelligence in Medicine, edited by Peter Szolovits, Westview Press, Boulder Colorado , AAAS Selected Symposium 51. 1982.