Causal probabilistic network modeling—An illustration of its role in the management of chronic diseases

by R. Hovorka
S. Andreassen
J. J. Benn
K. G. Olesen
E. R. Carson

This paper describes the role of the novel technique of causal probabilistic network (CPN) modeling as an approach to tackling control system problems typified by that of the administration of treatment to the patient suffering from a chronic disease such as diabetes. Three roles of a CPN are discussed. First, since diabetes arises as a consequence of impaired control of carbohydrate metabolism, the ability of a CPN to represent the uncertainty of a physiologically-based model is described. Second, its ability to make robust estimates of the parameters of the metabolic model is presented, and finally, in conjunction with decision theory approaches, its ability to compare alternative therapies and advise on insulin therapy for patients with insulin-dependent diabetes mellitus is illustrated.

The management of chronic noncommunicable diseases such as diabetes (diabetes mellitus), raised blood pressure (hypertension), and elevated levels of cholesterol poses some difficult challenges for the clinician. In most cases, from an engineering or systems perspective, such diseases can be viewed as arising from a partial or complete failure of one or more of the multitude of feedback control loops of the human organism. The management of such diseases requires regular disease status monitoring, and therapeutic interventions (control actions) are required to minimize the progression of the disease and to minimize the risk of long-term complications.

Information technology offers the capability of supporting clinical decision-making and management in the context of such chronic diseases in a variety of ways. Over and above database technology, which can support more consistent and reliable patient records, knowledge-based systems can be used to structure the clinician-patient consultation, thus potentially enabling the clinician to function at a higher level of expertise. Equally, algorithmic, knowledge-based, and model-based approaches can be used to assist in the context of treatment planning and adjustment, advising on dosage adjustment, and, by means of simulation, predicting outcome in response to change of therapy.

During the last decade, several new approaches for knowledge representation have been introduced. Causal probabilistic networks (CPNs), also known as belief networks,[1] belong to a family of methods that allow representation and processing of *uncertain* knowledge. CPNs have a solid theoretical basis in Bayesian probability theory. The updating of probabilities in CPNs is, however, a computationally hard problem (technically it is an NP-complete problem[2]), and, until recently, the computational issues limited the use of this approach. However, the recent development of computationally more efficient algorithms for networks with certain topologies[1,3–5] makes CPNs highly appropriate for the implementation of clinically relevant diagnostic expert systems,[6] as well as for other applications.[7,8] When combined with approaches based on decision theory, a general framework for planning and advisory systems emerges, which has already found one application in ventilator therapy planning.[9] The CPN approach enables population-specific and patient-specific information to be mixed within a formally coherent framework. This approach also provides a facility for dealing with uncertainties over time. Our paper describes the merits of this powerful modeling approach in tackling control system problems typified by the treatment of the patient suffering from chronic disease.

Causal probabilistic networks allow representation and processing of uncertain knowledge.

Diabetes

The CPN modeling technique will be illustrated by considering the planning and adjustment of insulin therapy for a diabetic patient. Diabetes is a major chronic disease in the industrialized world. It affects over 2 percent of the population of Europe and approximately one hundred million people worldwide.[10] Diabetes is a lifelong condition and can give rise to a variety of life-threatening complications. For example, it is the most common cause of blindness in people under the age of 65 in the United Kingdom and accounts for over 40 percent of lower-limb amputations carried out within the National Health Service.[11]

The incidence and severity of such complications can be reduced with good clinical management directed toward effective control of a patient's blood glucose level.[12] Control requires a careful balance between diet, physical activity, and insulin therapy. A high level of clinical expertise is required in order to achieve this result. While such expertise will be found in specialist hospital diabetic units and in some primary care practices with an interest in diabetes, it is not always to be found in other sectors of the health service. One way of making this clinical expertise more widely available, including the home setting of the diabetic patient, is through the appropriate use of information technology.[13]

Insulin-treated diabetes (known clinically as type 1 diabetes) results from partial or complete failure of the pancreas to produce the hormone insulin in response to elevated blood glucose levels. In normal health the glucose-insulin interaction is a classical example of control exerted by an effective negative feedback loop. From a control engineering perspective such a diabetic patient can be regarded as a multiinput-multioutput physiological system that contains several controllable and measurable variables as well as a number of other factors that are not directly observable and that, as such, are beyond the clinician's control. In this complex system the diet and administered insulin dose can be considered as control variables that need to be adjusted in order to achieve the therapeutic objective, namely to maintain a balance between energy supply and expenditure at blood glucose levels set by the clinician.

A number of computer-based approaches have been attempted in order to assist in insulin dosage adjustment and, more generally, in the treatment or long-term management of diabetic patients. The methods and techniques used include computer algorithms for advising on insulin dosage adjustment,[14–16] knowledge-based systems to advise on patient management in out-patient clinics,[17,18] and mathematical models as a means of simulating and predicting blood glucose level in response to change in therapy.[19–23] It is a novel extension to this last category, as applied in the context of managing a chronic disease, that is pro-

vided by CPN-based modeling as described in the remainder of this paper.

Causal probabilistic network—General aspects

A causal probabilistic network consists of nodes and directed links. Nodes represent random variables: state variables, conditions, measurements, etc. For the sake of simplicity, we will restrict ourselves to discrete random variables. Links represent causal relations between nodes. Ancestors of a node are called its parents. Figure 1 gives an example of a CPN containing a representation of two diseases. The topology of the network determines the dependencies and independencies between the variables in the network. For example, the links in the network indicate that Disease A affects the presence of Symptom C; Disease B affects the presence of Symptom C and Symptom D. The absence of a link between Disease A and Disease B indicates that Disease A and Disease B are independent in the absence of information about Symptom C. The absence of a link between Disease A and Symptom D indicates that Disease A and Symptom D are conditionally independent given Disease B. In other words, when the status of Disease B is known, the probability of Symptom D can be calculated irrespective of the status of Disease A.

As far as the network structure is considered, one restriction holds. The network must be a directed acyclic graph. Examples of legal structures are shown in Figure 2. In the top panel of Figure 2, the directed tree has a single root node and at most one directed path between any two nodes. The middle panel shows a directed multitree structure with several root nodes and at most one directed path between any two nodes. In the bottom panel a general directed acyclic graph is depicted with more root nodes and possibly more than one path between two nodes.

The relation between a node and the complete set of its parents has to be specified in the form of conditional probabilities. In our example, the relation between Disease B and Symptom D is fully specified by giving the conditional probabilities

P(D = PRESENT | B = PRESENT)
P(D = PRESENT | B = ABSENT)
P(D = ABSENT | B = PRESENT)
P(D = ABSENT | B = ABSENT)

Figure 1 An example of a causal probabilistic network

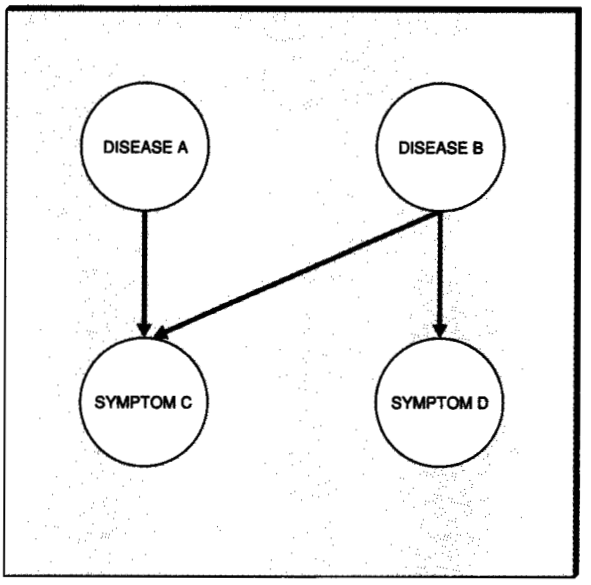

Nodes with an empty set of parents are called root nodes. Having assigned *a priori* probability distributions to root nodes, a probability model is fully defined. As in our example, assigning $P(A) = p_A$ and $P(B) = p_B$, where p_A and p_B are the probabilities of incidence of Disease A and Disease B among the population respectively, the probability of any possible combination of A, B, C, and D, e.g., P(A = PRESENT, B = ABSENT, C = PRESENT, D = ABSENT), is specified unambiguously. It is then possible to calculate the elements of a joint probability table, which in our little example is P(A, B, C, D).

If we assume that P(A, B, C, D) is available, a variety of queries about probabilities in the network can be answered. For instance, should the probability of the presence of Symptom C be requested assuming no evidence is available, it can be calculated from the joint probability table as

$$P(C = PRESENT) = \sum_{A,B,D} P(A, B, C = PRESENT, D)$$

When the evidence that assigns a value to a node has been obtained, it is natural to require the probability of the remaining nodes to be updated. With use of CPN terminology, the evidence is propagated throughout the network. Having observed the presence of Symptom C in our exam-

Figure 2 Graphical structures allowed for CPN representation

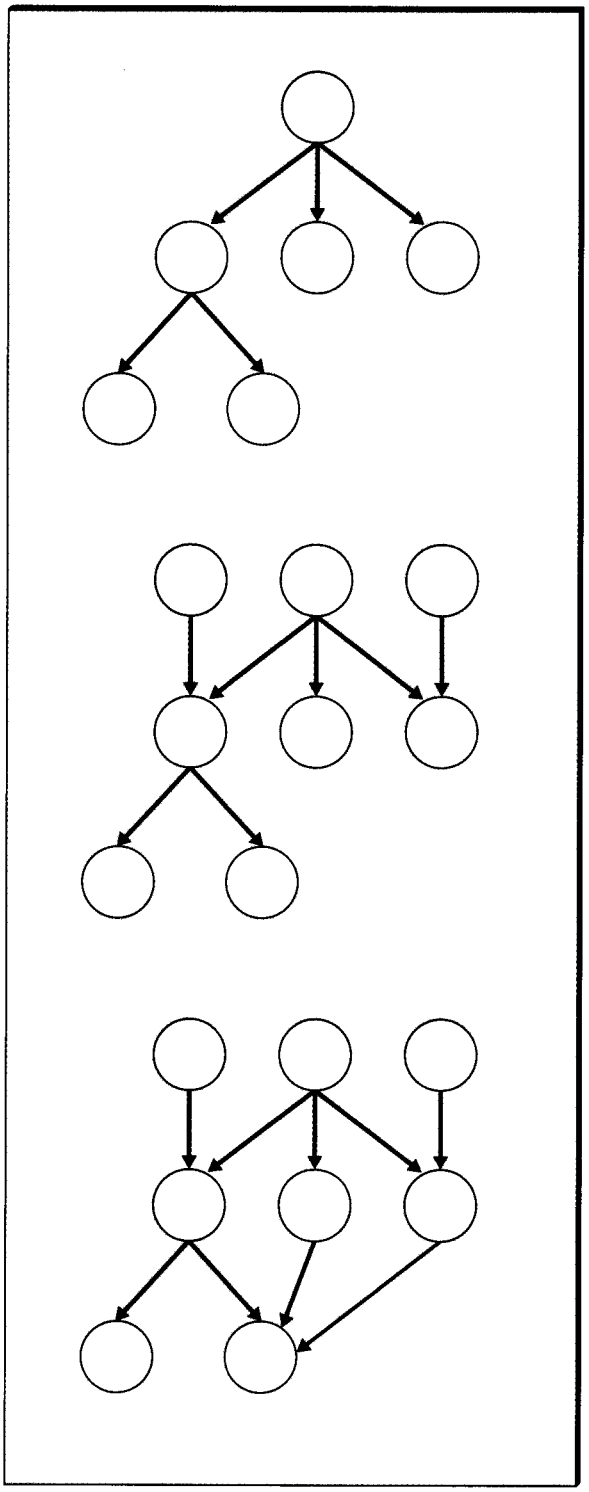

ple, we see that the probability of a patient having Disease A is obtained as the conditional probability P(A = PRESENT | C = PRESENT). This result can be computed from the joint probability table as

$$P(A = PRESENT \mid C = PRESENT) = \frac{P(A = PRESENT, C = PRESENT)}{P(C = PRESENT)} = \frac{\sum_{B,D} P(A = PRESENT, B, C = PRESENT, D)}{\sum_{A,B,D} P(A, B, C = PRESENT, D)}$$

The remaining question concerns the initial calculation of the joint probability table P(A, B, C, D). This table can be calculated from the known conditional probabilities and the known *a priori* probabilities, using the independencies apparent from the network

$$P(A, B, C, D) = P(C|A, B, D)\, P(A, B, D) = P(C|A, B)\, P(A, B, D)$$

since C and D are independent given B. Further

$$P(A, B, C, D) = P(C|A, B)\, P(D|A, B)\, P(A, B) = P(C|A, B)\, P(D|B)\, P(A, B)$$

since D and A are independent given B. Finally,

$$P(A, B, C, D) = P(C|A, B)\, P(D|B)\, P(A)\, P(B)$$

since A and B are independent.

The joint probability table of our example contains sixteen ($2^4 = 16$) elements. For practical problems, the size of the table would exceed the limits imposed by current technology. A network with 100 nodes, each node having five states, would require a table with more than 10^{69} ($5^{100} \sim 8 \times 10^{69}$) elements. The calculations illustrated above increase exponentially with the size of a network. Clearly, alternative methods are required to allow network initialization (calculation of the *a priori* probability of each node) and evidence propagation (calculation of the *a posteriori* probability of each node).

The power and utility of the CPN is that knowledge about the geometric structure of a network can be used to facilitate more efficient evidence propagation. In case of tree and multitree graph structures, initialization and evidence propagation can

be performed through local operations by each node on information provided by its parents and children. These operations can be carried out by using an algorithm that computes all probabilities in a number of steps linearly related to the number of nodes in a network.[1] For the case of networks with loops, i.e., networks where two or more paths exist between two nodes, no such simple algorithm is available to perform the calculation. However, recent theoretical developments[3] followed by a practical implementation[4,5] have allowed calculations to be performed efficiently in networks of considerable size. MUNIN,[6] an EMG (electromyograph) diagnostic assistant, contains more than one thousand nodes, and propagation takes just a few seconds.

CPNs have been primarily used in *diagnostic* systems. The novelty of the approach described in this paper lies in the introduction of the *time domain* in knowledge representation. A discrete-time physiologically-based model of carbohydrate metabolism is built using a CPN, and the ability of this approach to represent uncertain knowledge, to perform robust parameter estimation, and to compare competing therapies is described.

Insulin dosage adjustment in insulin-dependent diabetic patients

A CPN-based model of carbohydrate metabolism. Considerable knowledge about carbohydrate metabolism has been obtained from clinical experiments. Carbohydrate metabolism has been studied under various conditions, and both qualitative and quantitative information are now available. The basic structure of carbohydrate metabolism is clear. Glucose enters the body and is distributed, excreted, or utilized. Control mechanisms exist that modulate both the appearance and disappearance of glucose in the organism. The quantification of these processes always includes some degree of uncertainty reflecting the variations that exist between individuals and the temporal variations that occur in each individual.

The scheme of the model is given in Figure 3. Glucose enters the plasma circulation from two possible sources. Either the carbohydrate contained in the "MEAL" enters the "STOMACH," is absorbed from the gastrointestinal tract as glucose ("GUT ABSORPTION"), and appears in plasma ("BG"), or glucose is released by the liver ("GLUCOSE PRODUCTION") at a rate that depends on glucose and insulin levels. The model includes the representation of three processes that are involved in glucose removal from plasma. These processes are insulin-dependent glucose utilization ("INSULIN-DEPENDENT UPTAKE"), insulin-independent utilization ("INSULIN-INDEPENDENT UPTAKE"), and "RENAL EXCRETION." A significant part of the overall glucose uptake is known to be insulin-mediated and is thus named insulin-dependent glucose utilization. Insulin has a stimulatory effect on glucose uptake by cells, especially in the muscles. Glucose uptake by some tissues, e.g., the central nervous system, is independent of insulin levels and is controlled predominantly by plasma glucose concentration itself. When plasma glucose reaches an elevated threshold concentration, a portion of that glucose is filtered by the kidney and excreted in urine.

The novelty of this approach lies in the introduction of the time domain in knowledge representation.

In normal subjects, insulin is produced by the pancreas. In insulin-dependent diabetic patients, insulin secretion by the pancreas is severely impaired and has to be replaced externally by insulin injections. Insulin is usually injected into subcutaneous tissue. It is then slowly absorbed ("INSULIN ABSORPTION") and enters the plasma circulation. The rate of absorption varies for different insulin types. The time course of the effect of insulin on glucose production and uptake is delayed from that of the time course of insulin concentration in plasma. Thus "ACTIVE INSULIN" was introduced to represent the correct time course of the insulin stimulatory potential.

High variations exist between diabetic patients in terms of quantity of insulin administered. The insulin stimulatory potential can be significantly decreased, and patients often demonstrate insulin

Figure 3 Model of carbohydrate metabolism

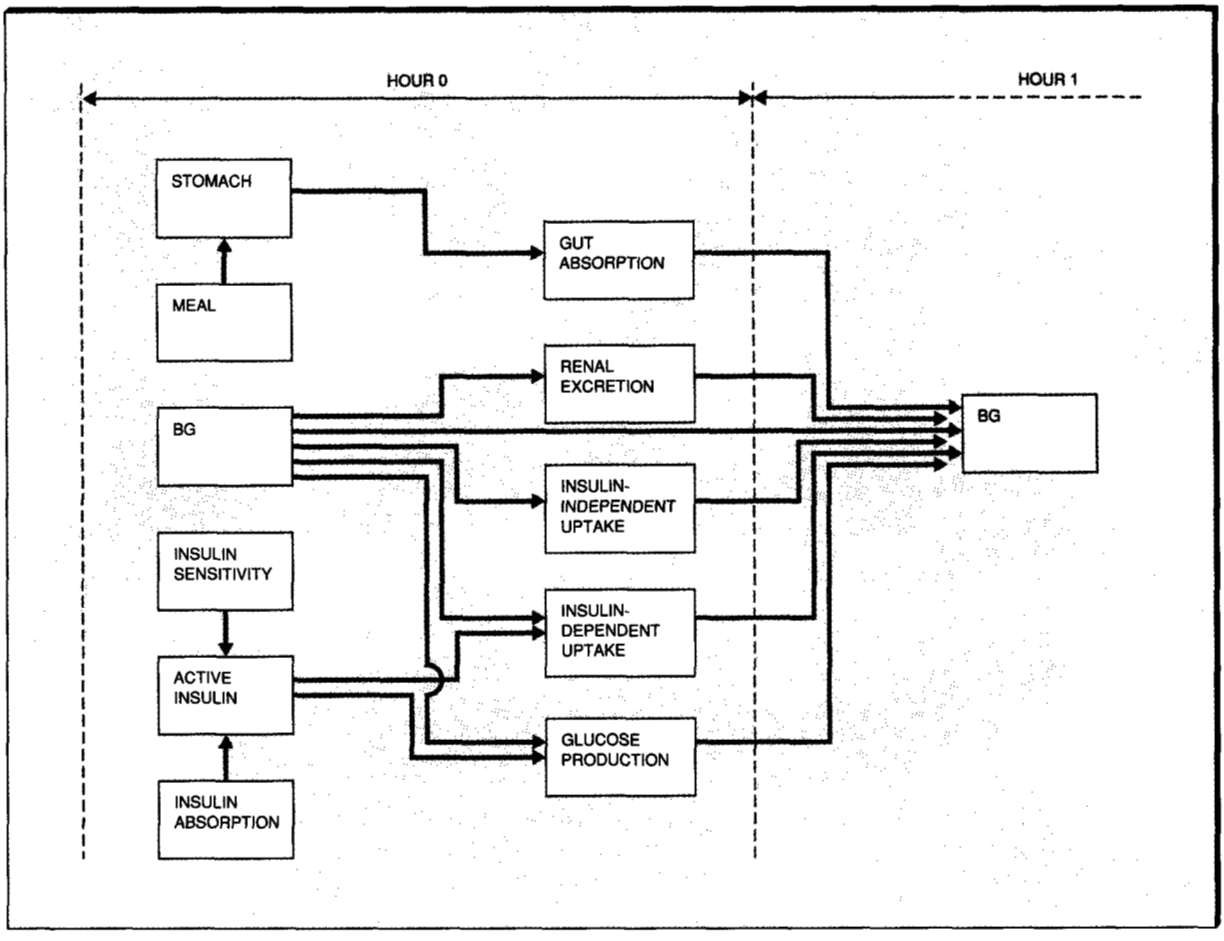

Adapted from Andreassen et al.[28]

resistance. Our model represents insulin resistance using an "INSULIN SENSITIVITY" scaling factor that scales down "ACTIVE INSULIN" according to the level of insulin resistance. Variations in insulin resistance are known to be responsible for a great deal of the overall variation exhibited by diabetic subjects. In our model, "INSULIN SENSITIVITY" is treated as a model parameter and is estimated for each individual.

The model is a discrete time model with a one-hour step. Figure 3 shows one time slice that is repeated 24 times in the model. The model thus covers a one-day period. The arrows represent the causal links between the processes involved in plasma glucose control. The quantification of the relations is given in the form of conditional probabilities, the derivation of which is illustrated in Figure 4. The top panel of the figure shows the relation between the quantity of glucose in the stomach and the rate of glucose absorption as compiled from data reported in medical literature.[24,25] The data have the typical format: mean value and a measure of variance (in this case standard deviation) of the dependent variable are given for several discrete values of the independent variable. This knowledge can be translated into conditional probabilities. The dependent variable is assumed to be a discrete stochastic variable, and by sampling a normal distribution specified by the mean and variance it is possible to calculate its probability. The bottom panel of

Figure 4 Derivation of conditional probabilities to specify linkage between a node and its parent

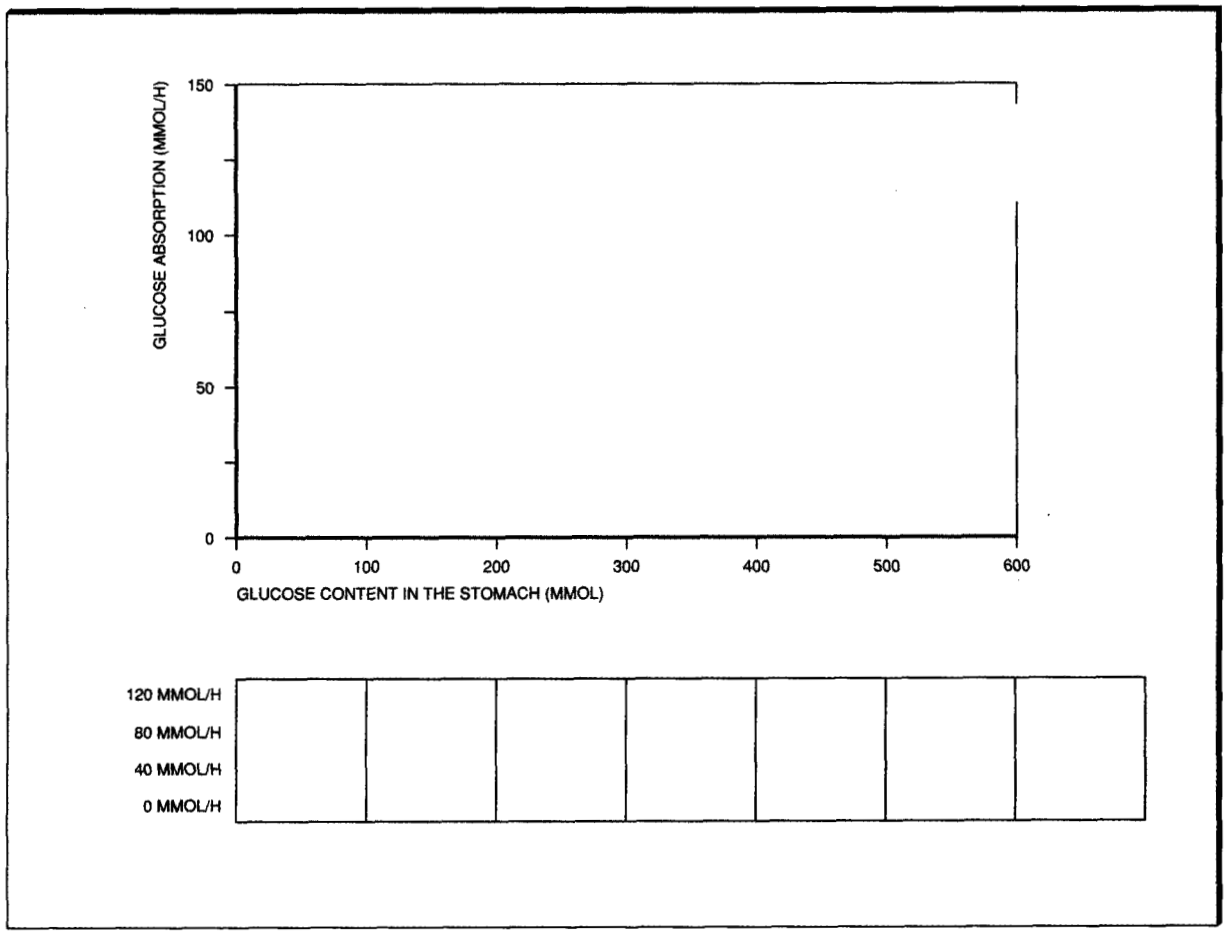

Adapted from Andreassen et al.[28]

Figure 4 shows the conditional probabilities derived in this manner from the data shown in the top panel. For example, given the empty stomach, the probability of zero absorption is 100 percent (full bar in 0 millimoles per hour [mmol/h] row under 0 mmol label) and zero percent for absorption rates of 40, 80, and 120 mmol/h.

The discretization of what in reality are continuous variables is carried out in a manner that enables the whole clinical spectrum of values to be covered. For example, blood glucose concentration is divided into a finite number of ranges to which the clinician would apply qualitative labels such as "low," "normal," "high," "very high," etc. Each of these ranges is then represented by a discrete value. For instance, "normal" blood glucose concentration could be represented by the value of 6 mmol/l.

This then means that there are two dimensions of uncertainties attached to any discrete value. The first is the quantization error arising from the representation of, say, the normal blood glucose concentration range by a discrete value. This form of error is not explicitly considered. The second is the uncertainty as to whether the qualitative label (e.g., "normal" blood glucose concentration) is appropriate in a given situation or whether an adjacent range (which qualitatively would be labeled "high" or "low") is more probable. It is this second dimension of uncertainty that is repre-

Figure 5 Parameter estimation using CPN

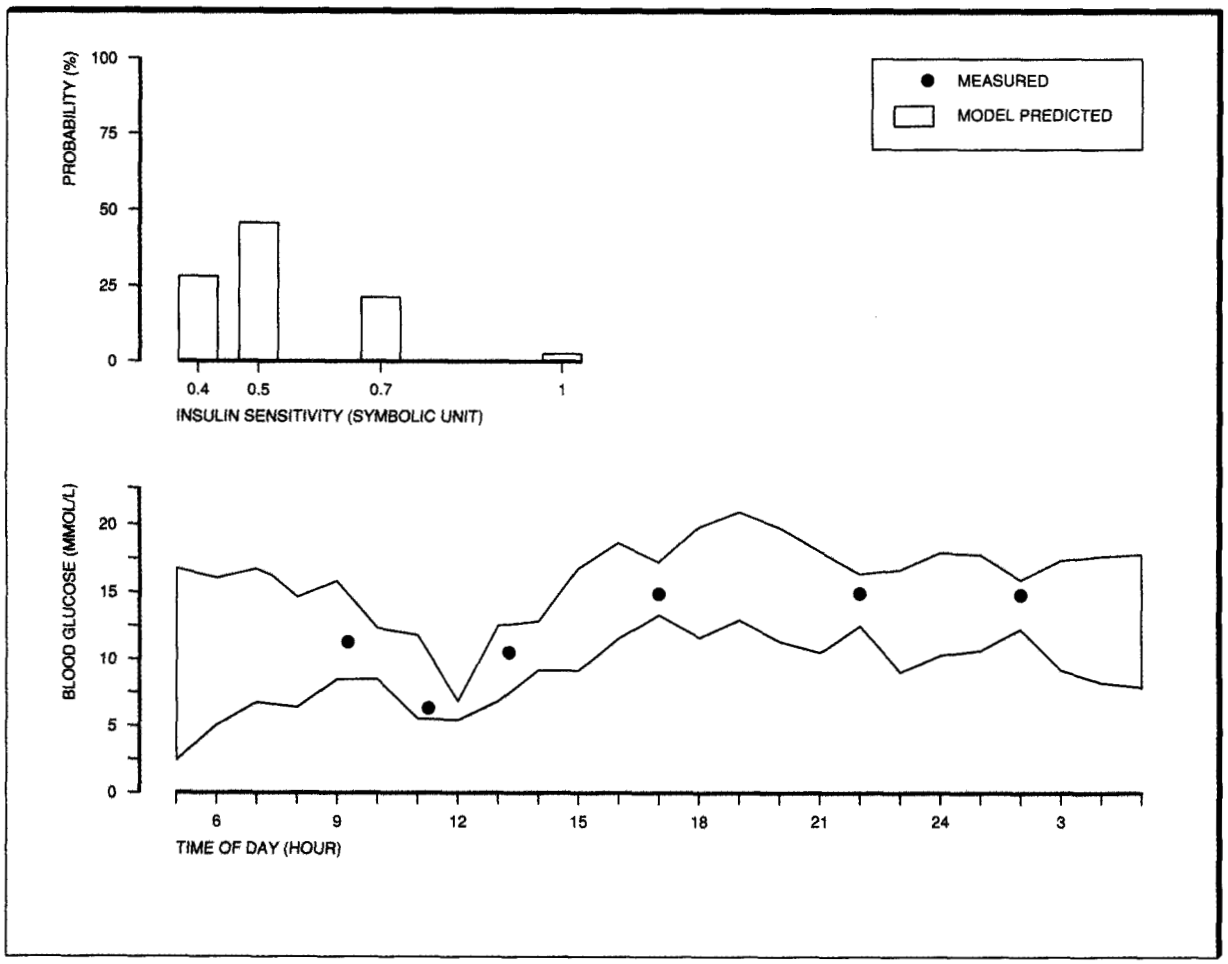

sented by the conditional probabilities such as shown in the lower panel of Figure 4.

Adaptation of the model for an individual patient. Insulin therapy and carbohydrate content of the meal are model inputs, plasma glucose is a model output, and insulin sensitivity is a model parameter. After data are observed from an individual patient, the model parameter has to be updated to reflect the observed input-output relation. The results of the parameter estimation process are illustrated in Figure 5. The upper panel shows the *a posteriori* probability distribution of insulin sensitivity calculated by the system from the data obtained from an individual patient. It should be noted that the estimation does not result in a single value output as would be the case if using standard least-square parameter estimation techniques. Rather a probability distribution is generated that indicates the extent to which individual parameter values can explain the observed data. Strictly speaking, as the estimation is based on Bayesian theory, the starting, *a priori* probability distribution of insulin sensitivity is also taken into account when the updated, *a posteriori* probability distribution is calculated. The patient exhibits decreased insulin sensitivity. Normal sensitivity has a numeric value of 1; the patient's insulin sensitivity is spread around a value of 0.5. The lower panel of the figure shows the model fit to measured blood glucose values. The mean ±SD region of the model prediction is given. This type

Figure 6 Comparison of blood glucose measurements and the mean ± SD region of model prediction

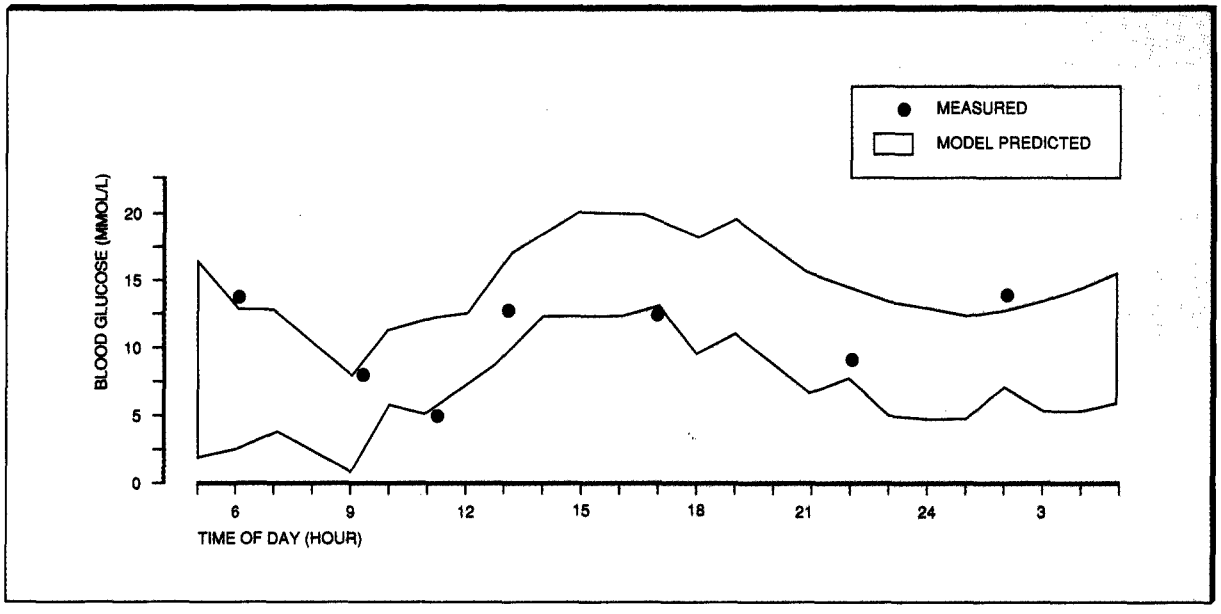

of information is quite unique and cannot be directly obtained from deterministic modeling approaches.

Model-based prediction of blood glucose. After the insulin sensitivity from observed data is estimated, the model can be used to predict blood glucose for various diet and insulin regimens. This prediction can be carried out either for regimens actually administered or for purely hypothetical regimens. In the former case, the precision of the model prediction can be evaluated. Model-based prediction of blood glucose is compared with values actually measured. In the latter case, a therapy recommendation can be generated as discussed in the next subsection. An example of the assessment of the precision of model prediction is shown in Figure 6. The figure shows mean ±SD region for predicted blood glucose and also the actual blood glucose measurements, allowing visual assessment of the precision of the prediction. The model prediction was calculated using the data about diet and insulin regimens actually administered during the period indicated, and using the insulin sensitivity parameter estimated for this patient on a previous occasion. Using more formal techniques, e.g., by calculating mean square prediction error, a single measure of the precision of the prediction can be computed and used during the evaluation phase of the system development.

Therapy planning. The ability of the model to predict blood glucose as an outcome of insulin therapy can be used to generate advice on therapy. The prediction is not, however, sufficient for advice generation. A performance measure has to be adopted to indicate the benefit (or loss) associated with a therapy being administered. Alternative therapies can then be evaluated by comparing the performance measures arising from them. The therapy with the extremal value of performance measure can then be recommended.

From a medical perspective, low and high blood glucose values are not desirable. A low blood glucose value (hypoglycemia) is perceived by the diabetic patient as being unpleasant and may result in an acute loss of consciousness with a risk of chronic brain damage. High blood glucose is regarded as a major cause of late diabetic complications such as blindness or impaired kidney function. Clinically, the penalization of low and high glucose values (reflecting their undesirability) in relative terms can be expressed by the M-value.[26] These penalties resulted from a sub-

Figure 7 Penalties associated with different levels of blood glucose

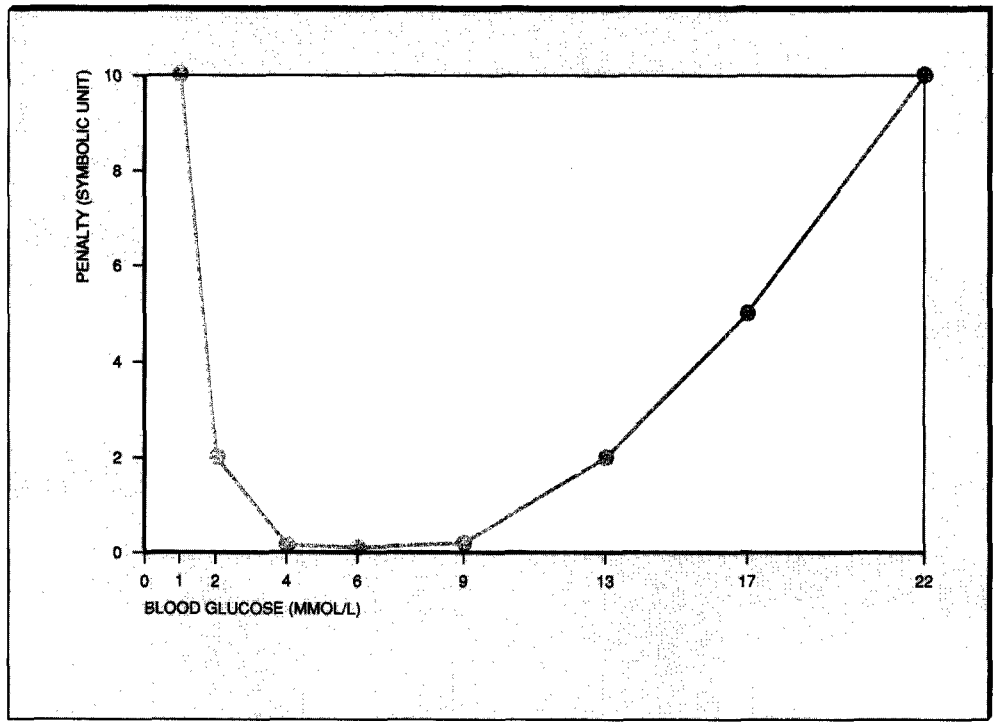

Adapted from Andreassen et al.[28]

jective assessment being carried out by a panel of diabetes experts. We have adopted a similar scheme of penalties shown in Figure 7. It should be noted that the penalties are not symmetric, and this scheme is thus essentially different from a standard square-law approach to penalties.[27]

The penalties were elicited from the clinical experts participating in this study. They represent the experts' subjective assessment of the risk or inconvenience or both for the patient. For the purpose of the elicitation of the penalties, the penalty curve shown in Figure 7 was expressed in hours of life lost for each hour that glucose was maintained at any particular level. For the high blood glucose levels the penalties mainly represented the risk of premature death and qualitative corrections made for loss of quality of life due to complications such as blindness, kidney failure, etc. The currently available epidemiological data only allow these assessments to be made in a qualitative and subjective way. The penalties associated with low blood glucose levels contain a substantial element of discomfort and social embarrassment as perceived by the patient as well as contributions from the risk of accidents, brain damage, or other organ damage during hypoglycemic episodes. The penalty curve thus represents an attempt to produce a single utility measure that takes several of the patient's dimensions of utility into account. We only consider the penalty curve to be a qualitative and subjective representation of some aspects of the patient's total utility, but it is interesting that it turned out to be quite similar to the M-value. The only major difference is that our penalty curve penalizes high blood glucose a bit more severely than the M-value does.

Blood glucose is predicted by the model as a probability distribution. To calculate the performance measure for an insulin regimen, ideas from utility theory are employed. A weighted mean of penalties is computed using the probability distribution generated by the system. The calculation process is illustrated in Table 1. In essence, the product of the probability of each outcome (i.e.,

Figure 8 The comparison of an administered therapy (default therapy) and the therapy with minimum performance measure (advised therapy)

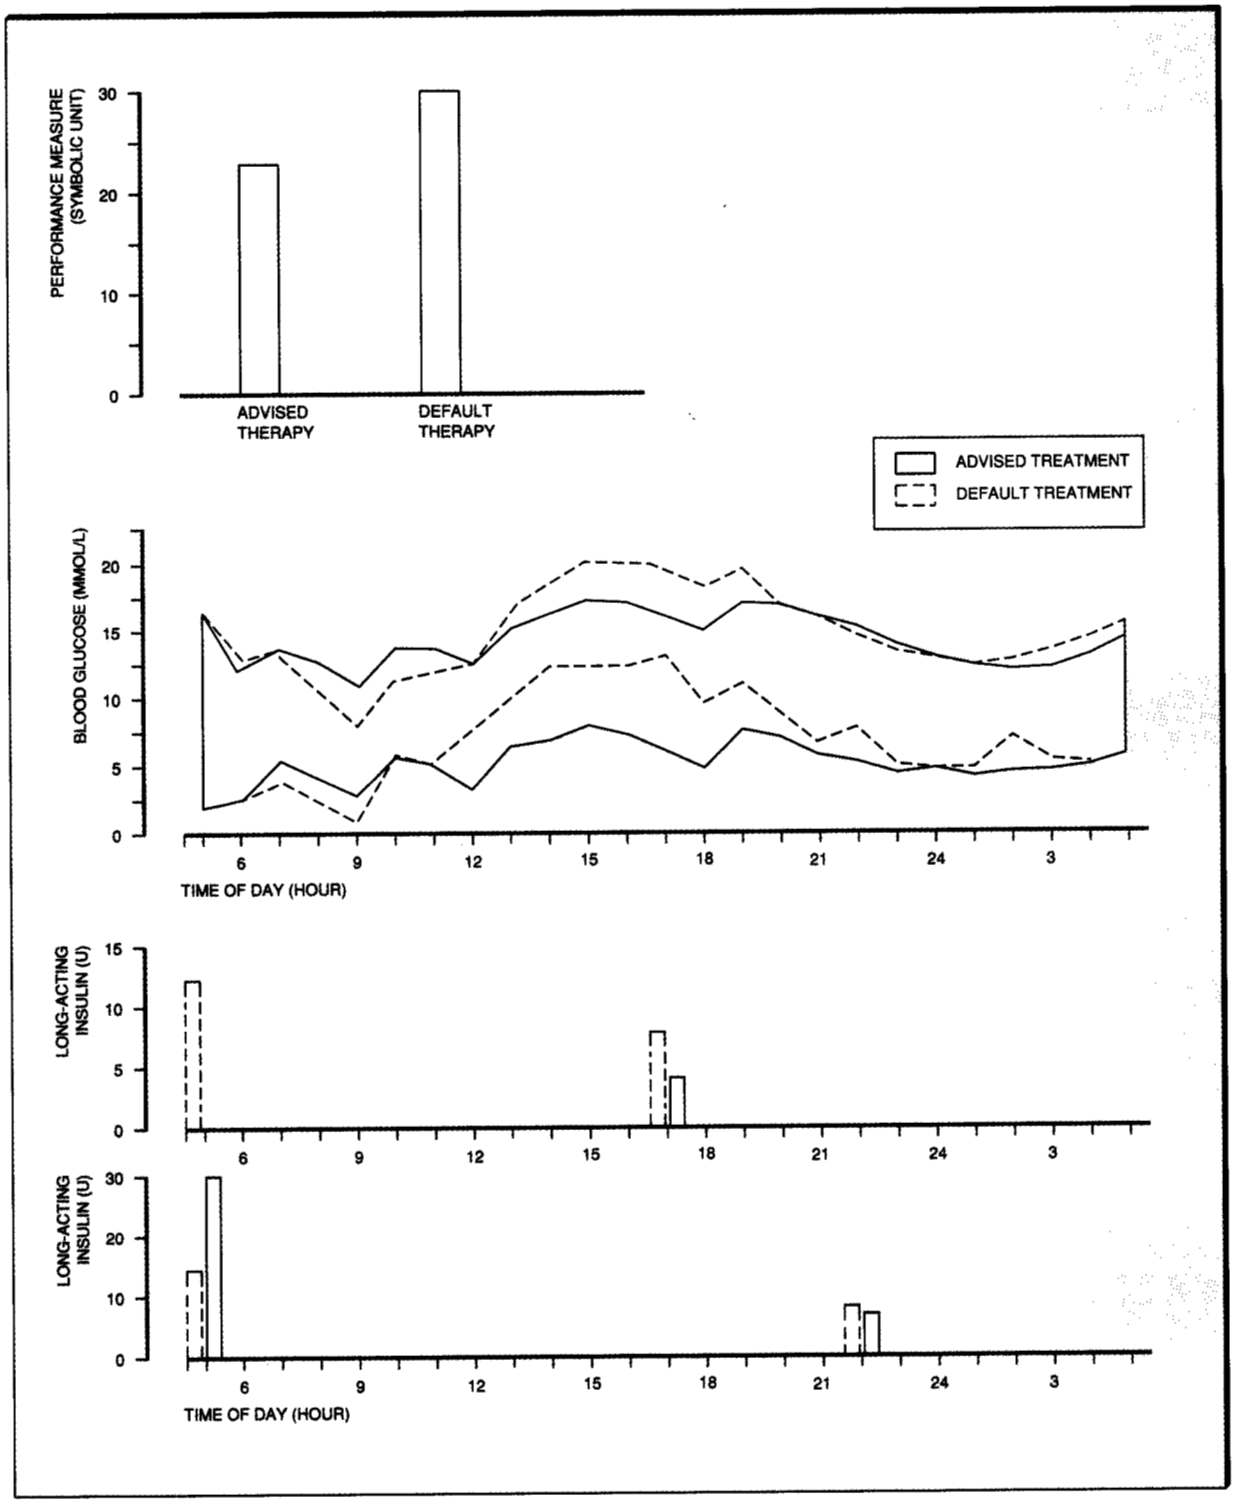

Table 1 Calculation of the penalty associated with a blood glucose value when blood glucose value is specified as a probability distribution (the penalty function shown in Figure 7 is employed in the calculation).

Blood Glucose (mmol/l)	Probability (percent)	Penalty Function	Product
1	0	10.0	0
2	0	2.0	0
4	0	0.1	0
6	1	0	0
9	6	0.1	0.006
13	20	2.0	0.400
17	36	5.0	1.800
22	37	10.0	3.700
			Σ 5.906

a given blood glucose) and the penalty attached to the outcome is summed over all possible blood glucose outcomes. This process gives the total penalty attached to the blood glucose probability distribution at a specific time. To obtain the performance measure of the therapy, the penalties of 24 blood glucose probability distributions are summed (i.e., values predicted at 24 consecutive one-hour intervals).

With the performance measure defined, alternative insulin therapies can be compared. The therapy with the minimum performance measure is assumed to be optimal and is recommended. To find a therapy with minimum performance measure, the n-dimensional space is searched for what is hopefully a global minimum, where n is the number of insulin injections per day. Several methods can be employed to perform the search. We have adopted a gradient method which usually converged in less than 10 iterations. Figure 8 shows the comparison between default therapy (the therapy that was actually being administered) and the therapy with minimum performance measure (the one that would be recommended). The top panel shows the improvement in the performance measure as predicted by the system. The next panel gives the mean ±SD regions of blood glucose predicted for a default insulin therapy and for therapy with minimum performance measure. The last two panels indicate the alterations in doses of short-acting insulin and long-acting insulin suggested by the system.

Implementation. These techniques have been employed to build SWAN,[28] a system that advises on insulin therapy for patients who require this form of treatment. The system runs on SUN-based workstations and uses HUGIN[4] to handle the probabilities in the network. Prior to calculations, HUGIN compiles the network, rearranging the network structure and creating partial joint-probability tables. The conditional probability tables required for the specification of one time-slice of the network as shown in Figure 3 hold about 13 000 numbers. After compilation the tables required to hold one time-slice hold about 500 000 numbers corresponding to 2 megabytes (Mb). The total compiled SWAN network with 24 time-slices takes up about 50 Mb. The SWAN program itself occupies less than 1 Mb. A single propagation in the network requires approximately 10 seconds of CPU time; the search for optimum therapy requires 3 to 6 minutes of CPU time. The current activities in this project are focused on further refinement of the network by including a more detailed representation of the glucose-insulin dynamics, on speeding up the calculations, and on clinical evaluation.

Concluding remarks

The novel technique of causal probabilistic network modeling has been described and employed to generate advice on insulin therapy for diabetic patients. CPN models have the intrinsic capability of representing uncertainty in the model specification and are substantially novel in their ability to estimate model parameters as probability distributions rather than as single values. Predictions generated by the system are also produced in the form of probability distributions. Combining CPN models and decision theory approaches allows the assessment of competing therapies, and hence, an optimum therapy based on the chosen criterion can be identified. Clinically, this approach to insulin dose adjustment, based on probabilities, has many more similarities to the approach of the physician caring for patients with diabetes than deterministic models. Both the physician and the CPN model seek insulin doses that maintain as near normal a blood glucose as is possible with an acceptably small risk of a dangerously low value of blood glucose.

Acknowledgment

This work was carried out as part of the EC-AIM EURODIABETA Project No. A1019 and was also supported by a grant from The Wellcome Trust to

RH. The clinical data were collected by clinical partners of the EURODIABETA project.

Cited references

1. J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, San Mateo, CA (1988).
2. G. F. Cooper, "The Computational Complexity of Probabilistic Inference Using Bayesian Belief Networks," *Artificial Intelligence* **42,** 393–405 (1990).
3. S. L. Lauritzen and D. J. Spiegelhalter, "Local Computations with Probabilities on Graphical Structures and Their Application to Expert Systems (with Discussion)," *Journal of the Royal Statistical Society B* **50,** 157–224 (1988).
4. S. K. Andersen, K. G. Olessen, F. V. Jensen, and F. Jensen, "HUGIN—A Shell for Building Bayesian Belief Universes for Expert Systems," *Proceedings of IJCAI 89* (1989), pp. 1080–1085.
5. F. V. Jensen, S. L. Lauritzen, and K. G. Olessen, "Bayesian Updating in Causal Probabilistic Networks by Local Computations," *Computational Statistics Quarterly* **4,** 296–282 (1990).
6. S. Andreassen, F. V. Jensen, S. K. Andersen, B. Falck, U. Kjaerulff, M. Woldbye, A. R. Sorensen, A. Rosenfalck, and F. Jensen, "MUNIN—An Expert EMG Assistant," J. E. Desmedt, Editor, *Computer-Aided Electromyography and Expert Systems*, Elsevier, Amsterdam (1989), pp. 255–277.
7. B. Abramson, "ARCO1: An Application of Belief Networks to the Oil Market," *Proceedings of the 7th Conference on Uncertainty in Artificial Intelligence*, Los Angeles (1991), pp. 1–8.
8. E. J. Horvitz, J. S. Breese, and M. Henrion, "Decision Theory in Expert Systems and Artificial Intelligence," *Journal of Approximate Reasoning* **2,** 247–302 (1988).
9. G. Rutledge, G. Thomsen, et al., "Combining Qualitative and Quantitative Computation in a Ventilator Therapy Planner," *Proceedings of the 13th Annual Symposium on Computer Applications in Medical Care*, IEEE, New York (1988), pp. 315–319.
10. Eurodiabeta, "Information Technology for Diabetes Care in Europe: The EURODIABETA Initiative," *Diabetic Medicine* **7,** 639–650 (1990).
11. J. L. Day and M. Spathis, "District Diabetes Centres in the United Kingdom," *Diabetic Medicine* **5,** 372–380 (1988).
12. R. Engerman, J. M. B. Bloodworth, and S. Nelson, "Relationship of Microvascular Disease in Diabetes to Metabolic Control," *Diabetes* **26,** 760–769 (1977).
13. "Impact of Computer Simulation on Diabetes Research and Practice," editorial in *Diabetes Care* **12,** 748–749 (1989).
14. A. M. Albisser, A. Schiffrin, M. Schulz, J. Tiran, and B. S. Leibel, "Insulin Dosage Adjustment Using Manual Methods and Computer Algorithms: A Comparative Study," *Medical and Biological Engineering and Computing* **24,** 577–584 (1986).
15. T. Shimauchi, N. Kugai, N. Nagata, and O. Takatani, "Microcomputer-Aided Insulin Dose Determination in Intensified Conventional Insulin Therapy," *IEEE Transactions on Biomedical Engineering* **35,** 167–171 (1988).
16. J. S. Skyler, D. L. Skyler, D. E. Seigler, and M. J. O'Sullivan, "Algorithms for Adjustment of Insulin Dosage by Patients Who Monitor Blood Glucose," *Diabetes Care* **4,** 311–318 (1981).
17. E. R. Carson, S. Carey, F. E. Harvey, P. H. Sonksen, S. Till, and C. D. Williams, "Information Technology and Computer-Based Decision Support in Diabetic Management," *Computer Methods and Programs in Biomedicine* **32,** 179–188 (1990).
18. M. Fieschi, M. Joubert, D. Fieschi, G. Botti, C. Michel, and H. Proudhon, "The Sphinx Project," I. De Lotto and M. Stefanelli, Editors, *Artificial Intelligence in Medicine*, North-Holland, Amsterdam (1985), pp. 83–94.
19. M. P. Berger and D. Rodbard, "Computer Simulation of Plasma Insulin and Glucose Dynamics after Subcutaneous Insulin Injection," *Diabetes Care* **12,** 725–736 (1989).
20. D. G. Cramp and E. R. Carson, "The Dynamics of Short-Term Blood Glucose Regulation," C. Cobelli and R. N. Bergman, Editors, *Carbohydrate Metabolism: Quantitative Physiology and Mathematical Modelling*, Wiley, Chichester (1981), pp. 349–367.
21. E. R. Carson, C. Cobelli, and L. Finkelstein, *The Mathematical Modeling of Metabolic and Endocrine Systems*, John Wiley & Sons, New York (1983).
22. R. Hovorka, S. Svacina, E. R. Carson, C. D. Williams, and P. H. Sonksen, "A Consultation System for Insulin Therapy," *Computer Methods and Programs in Biomedicine* **32,** 303–310 (1990).
23. E. D. Lehmann, T. Deutsch, E. R. Carson, and P. H. Sonksen, "Combining Rule-Based Reasoning and Mathematical Modelling in Diabetes Care: System Overview," *Artificial Intelligence in Medicine* (1992); (in press).
24. J. J. Benn, S. J. Bozzard, D. Kelley, A. Mitrakou, T. Aoki, J. Sorensen, J. E. Gerich, and P. H. Sonksen, "Persistent Abnormalities of the Metabolism of an Oral Glucose Load in Insulin-Treated Type 1 Diabetics," *Metabolism* **38,** 1047–1055 (1989).
25. M. A. Boroujerdi, M. S. Leaning, and K. Nicolosi, "Knowledge-Rich Representations of Carbohydrate Metabolism for Diabetic Expert Systems," *Medical Informatics 88: Computers in Clinical Medicine*, British Medical Informatics Society, London (1988), pp. 89–96.
26. J. Schlichtkrull, O. Munck, and M. Jersild, "The M-Value, an Index of Blood Sugar Control in Diabetics," *Acta Medica Scandinavia* **177,** 95–102 (1965).
27. J. O. Berger, *Statistical Decision Theory and Bayesian Analysis*, Springer-Verlag, New York (1985).
28. S. Andreassen, R. Hovorka, J. J. Benn, K. G. Olessen, and E. R. Carson, "A Model-Based Approach to Insulin Adjustment," M. Stefanelli, A. Hasman, M. Fieschi, and J. Talmon, Editors, *Lecture Notes in Medical Informatics* **44,** *AIME 91*, Springer-Verlag, Berlin (1991), pp. 239–248.

Accepted for publication June 22, 1992.

Roman Hovorka *Centre for Measurement & Information in Medicine, Department of Systems Science, City University, Northampton Square, London EC1V 0HB, United Kingdom (electronic mail: r.hovorka@city.ac.uk).* Dr. Hovorka received M.Sc. and Ph.D. degrees in mathematical informatics in 1984 and 1989, respectively, at Charles University, Prague. He is currently a postdoctoral research Fellow supported by the pharmaceutical company Glaxo. His professional interests include simulation and modeling of carbohydrate metabolism and metabolic systems in general, and use of those methods in therapy planning and drug development. Since

participating in the Eurodiabeta project funded by the European Community, he is also interested in applications of causal probabilistic networks. He publishes regularly in scientific journals and conference proceedings.

Steen Andreassen *Department of Medical Informatics and Image Analysis, Aalborg University, Fredrik Bajers Vej 7D, DK-9220 Aalborg East, Denmark (electronic mail: sa@miba.auc.dk).* Dr. Andreassen obtained his M.Sc. in electronic engineering from the Technical University of Denmark and his Ph.D. from the Technical University of Denmark and the University of Copenhagen. Currently he is Reader in Medical Informatics in his department at Aalborg University. He has published extensively in physiological measurement and control, biomedical engineering, and medical informatics. A major current research interest is the use of causal probabilistic network modeling techniques in clinical decision support systems.

Jonathan J. Benn *Department of Medicine, U.M.D.S., St. Thomas' Hospital, London SE1 7EH, United Kingdom.* Dr. Benn received his degree in natural sciences at Trinity College, Cambridge, and subsequently went to St. Thomas' Hospital, qualifying in medicine in 1978. He continued his medical career at St. George's Hospital and then at St. Thomas' Hospital, where he has been a lecturer in the Department of Medicine since 1987. His research interests include the study of intermediary metabolism with particular reference to diabetes. He is shortly to take a post as consultant physician with an interest in diabetes and endocrinology in Burton-on-Trent.

Kristian G. Olesen *Department of Medical Informatics and Image Analysis, Aalborg University, Fredrik Bajers Vej 7D, DK-9220 Aalborg East, Denmark (electronic mail: kgo@miba.auc.dk).* After receiving a B.Sc. in mathematics in 1982 and an M.Sc. in computer science in 1985 from Aalborg University, Dr. Olesen's primary research interest has been centered around methodological issues on model-based expert systems. He has held positions at Aalborg University and as project manager on the ESPRIT Project P599, An Intelligent Assistant for Electromyography, at Nordjysk Udviklingscenter. He received his Ph.D. in 1992 and is currently assistant professor at Aalborg University.

Ewart R. Carson *Centre for Measurement & Information in Medicine, Department of Systems Science, City University, Northampton Square, London EC1V 0HB, United Kingdom (electronic mail: E.R.Carson@city.ac.uk).* Dr. Carson obtained a B.Sc. in electrical engineering from the University of St. Andrews, an M.Sc. in instrument and control engineering and a Ph.D. and D.Sc. from City University. He is currently Director of the Centre for Measurement and Information in Medicine and Professor of Systems Science at City University. He is also Visiting Professor in Medical Informatics at U.M.D.S., St. Thomas' Hospital, London. He is active in IFAC and IMEKO affairs and has published widely, particularly in the fields of modeling and applications of control system approaches in physiology and medicine and in the development of medical decision support systems.

Reprint Order No. G321-5490.