DIABETES TECHNOLOGY & THERAPEUTICS
Volume 5, Number 3, 2003

Research Use of the AIDA www.2aida.org Diabetes Software Simulation Program: A Review—Part 1. Decision Support Testing and Neural Network Training

ELDON D. LEHMANN, M.B. B.S., B.Sc.[1,2]

ABSTRACT

The purpose of this two-part review is to overview research use of the AIDA diabetes software simulator. AIDA is a diabetes computer program that permits the interactive simulation of plasma insulin and blood glucose profiles for teaching, demonstration, and self-learning purposes. It has been made freely available, without charge, on the Internet as a noncommercial contribution to continuing diabetes education. Since its launch in 1996 over 300,000 visits have been logged at the main AIDA Website—www.2aida.org—and over 60,000 copies of the AIDA program have been downloaded free-of-charge. This review describes research projects and ventures, undertaken for the most part by other research workers in the diabetes computing field, that have made use of the freeware AIDA software. Relevant research work was identified in three main ways: (i) by personal (e-mail/written) communications from researchers, (ii) via the ISI Web of Science citation database to identify published articles that referred to AIDA-related papers, and (iii) via searches on the Internet. In a number of cases research students who had sought advice about AIDA, and diabetes computing in general, provided copies of their research dissertations/theses upon the completion of their projects. The two reviews highlight some of the many and varied research projects that have made use of the AIDA diabetes simulation software to date. A wide variety of diabetes computing topics have been addressed. In Part 1 of the review, these range from testing decision support prototypes to training artificial neural networks. In Part 2 of the review, issues surrounding dietary assessments, developing new diabetes models, and performance monitoring of closed-loop insulin delivery devices are considered. Overall, research projects making use of AIDA have been identified in Australia, Italy, South Korea, the United Kingdom, and the United States. These reviews confirm an unexpected but useful benefit of distributing medical software, like AIDA, for free via the Internet—demonstrating how it is possible to have a synergistic benefit with other researchers—facilitating their own research projects in related medical fields. The reviews highlight a variety of these projects that have benefited from the free availability of the AIDA diabetes software simulator. In a number of cases these other research projects simply would not have been possible without unrestricted access to the AIDA

[1]Academic Department of Radiology, St. Bartholomew's Hospital; and Department of Imaging (MRU), Imperial College of Science Technology & Medicine, NHLI Royal Brompton Hospital, London, United Kingdom.
[2]Dr. Lehmann is a co-developer of the AIDA diabetes simulator.
The AIDA software referred to in this review is an independent, noncommercial development which is being made available free-of-charge via the Internet—at a dot org (.org) not-for-profit Website—as a noncommercial contribution to continuing diabetes education.

software and/or technical descriptions of its workings. In addition, some specific common themes begin to emerge from the research ventures that have been reviewed. These include the use of simulated blood glucose data from the AIDA program for preliminary computer-lab-based testing of other decision support prototypes. Issues surrounding such use of simulated data for separate prototype testing are discussed further in Part 2 of the review.

INTRODUCTION

AIDA IS A FREEWARE COMPUTER PROGRAM that permits the interactive simulation of plasma insulin and blood glucose profiles for demonstration, teaching, and self-learning purposes. It has been made freely available, without charge, via the Internet as a noncommercial contribution to continuing diabetes education. In the 7+ years since its original World Wide Web launch in March/April 1996 over 300,000 visits have been logged to the main AIDA Web pages at www.2aida.org and over 60,000 copies of the program have been downloaded, *gratis*.

The AIDA software has been previously described in detail elsewhere in this journal.[1] Briefly, it incorporates a compartmental model that describes glucose–insulin interaction in patients completely lacking endogenous insulin secretion. It contains a single extracellular glucose compartment into which glucose enters via both intestinal absorption and hepatic glucose production. The AIDA model also contains separate compartments for plasma and "active" insulin[2,3]—the latter being responsible for glycaemic control, while insulin is removed from the former by hepatic degradation. The actual mathematics underlying the model have been fully documented elsewhere.[2] Details of the AIDA model are also accessible from within the AIDA software package, and can be viewed and printed separately via the Internet (from www.2aida.org/technical). Examples of the sort of simulations that AIDA can offer can be found elsewhere in this journal,[1,4–6] and the literature,[7–10] as well as at www.2aida.org/demo on the Internet.

In addition to its possible educational, teaching, self-learning, and/or demonstration uses AIDA may also be able to support research work into the application of computers in diabetes care. For the current two-part review we had three main aims: (i) to identify research projects that have made use of the AIDA diabetes simulation software; (ii) to overview the use of AIDA in these research ventures; and (iii) to identify common themes that might be of use to other researchers planning to undertake similar diabetes computing research projects. The reviews in particular highlight areas where other research groups have wanted to make use of simulated blood glucose data in their own research work. Necessarily some of this research relies on novel, technical computational approaches.

METHODOLOGY

Relevant research work for the reviews were identified in three main ways: (i) by personal (e-mail/written) communications from researchers who had made use of AIDA, (ii) via the ISI Web of Science citation database to identify published articles that referred to AIDA-related papers, and (iii) via searches on the Internet. In a number of cases research students who had sought advice about AIDA, and diabetes computing in general, kindly provided copies of their research dissertations/theses upon the completion of their projects.

Some of the information described in the reviews has previously been made freely available on the Internet at www.2aida.org/aida/research2.htm, and selected researchers were directed to this Web page to check the accuracy of the basic material provided.

TESTING DECISION SUPPORT PROTOTYPES #1

Bellazzi and colleagues from the Dipartimento di Informatica e Sistemistica, University of Pavia, Pavia, Italy have described using simulated blood glucose data from AIDA to test out a number of computer decision support prototypes under development in their labora-

tory. In Dr. Bellazzi's own words (personal communication, 1999):

> "I started working on the use of information technology tools in diabetes mellitus in 1994, and, due to my background in engineering, one of my first steps was to understand the medical problem and its pathophysiology. At the same time I also studied papers that could add a mathematical basis to the medical concepts I was learning. During these activities I read papers about the AIDA model and downloaded the AIDA program that, at that time, was running under the DOS operating system.
>
> Since that period I have exploited the diabetes simulator implemented in AIDA in a large number of research projects, and the model was used as the basis for several Masters theses[11–13] at the School of Engineering, where I am teaching. Moreover, AIDA was used by one of my coworkers at the Department of Pediatrics of the Policlinico S. Matteo Hospital in Pavia.
>
> To mention some of the more interesting activities we did, we exploited the AIDA model for:
>
> - Testing of different strategies (algorithms) for insulin optimization in routine care. In this area we compared different approaches, from rule-based systems, to Fuzzy-rule based systems and finally to model-based strategies.
> - Testing of a new model, based on Fuzzy systems and qualitative modeling, for blood glucose forecasting in patients with type I diabetes mellitus. In this area we used the patient simulator to obtain data and test the predictions in different situations and insulin protocols.
> - Education of Masters students of the school of engineering about diabetes. To this end we used AIDA to quickly introduce diabetes to people that wanted to Master in diabetes modeling.
>
> Moreover, together with the University of Padova, it was decided to use AIDA for comparison in the development of a new simulator of type I diabetes mellitus within the European Union IV Framework funded project, called T-IDDM (Telematic Management of Diabetes Mellitus patients).
>
> As a researcher, I am therefore in debt to the AIDA model for its help in my research activities, that fostered new approaches and ideas. As a teacher, I believe that it is a valuable instrument for teaching students that do not come from the school of medicine and who quickly need to learn about diabetes regulation mechanisms."

Dr. Bellazzi and colleagues clearly have found the AIDA diabetes simulator of some use in their own research work. However, it may be of interest to review their studies in some greater detail.

Qualitative models and fuzzy systems: an integrated approach for learning from data

In their report entitled "Qualitative models and fuzzy systems: an integrated approach for learning from data" Bellazzi et al.[14] have described a prototype method for the identification of the dynamics of nonlinear systems in diabetes care by trying to learn from data. The key idea that underlies their approach consists of the integration of qualitative modelling techniques with fuzzy logic systems. The resulting hybrid method exploits the *a priori* structural knowledge of the system to try and initialise a fuzzy inference procedure that determines, from the available experimental data, a functional approximation of the system dynamics that can be used as a reasonable predictor of the patient's future state. The major advantage that is believed to result from such an integrated framework lies in a significant improvement in both the efficiency and robustness of identification methods based on fuzzy models that learn an input–output relation from the data provided. As a benchmark for the methodology, the authors have considered the problems of identifying the response to insulin therapy for insulin-dependent (type 1) diabetic patients.

For this work, Bellazzi et al.[14] used simulated data from AIDA, in which the noise level could be suitably manipulated for test purposes, as well as the number of missing data varied/in-

creased. The overall procedure occurred in two steps:

(1) Training phase. Over a simulated period of 24 h, with different sampling times ranging from 15 min to 8 h, the authors have simulated patient responses to an injection of regular insulin followed by a meal. The performance of a function approximator, $y(x)$, obtained using the prototype has been compared, through root mean square error calculations, with that of a separate function approximator identified just from the data.

(2) Validation phase. Evaluation of the predictive accuracy of the prototype's function approximator [$y(x)$] has taken place when dealing with a new dataset. In particular the new data were obtained by simulating the patient response to a typical daily insulin protocol, composed of two injections of NPH insulin and two injections of regular insulin, followed by a meal.

Bellazzi et al.[14] have tested both phases in different experimental settings: (i) training and validation with noise-free data and a minimum sampling time (of 15 min); (ii) training and validation with noisy data and a minimum sampling time; and (iii) training and validation with a maximum sampling time. The preliminary results obtained using AIDA simulation data have motivated the authors to do further work with their approach, moving towards an evaluation of the method with real patient data.

Learning from data through the integration of qualitative models and fuzzy systems

In a further report, entitled "Learning from data through the integration of qualitative models and fuzzy systems," Bellazzi et al.[15] attempted to build on the earlier work from this research group—presenting a methodology for the identification of the dynamics of nonlinear pathophysiological systems once again by trying to learn from data. The key idea that underlies this approach consists, once more, of the integration of qualitative modelling methods with fuzzy logic systems. The major perceived advantage that derives from such an integrated framework lies in its capability both to represent the structural knowledge of the system at study and to exploit the available experimental data. As a result a functional approximation of the system dynamics can be determined and used as a predictor of the patient's future state. As a testing ground for their method the authors have considered the problem of identifying the response to insulin therapy in diabetes—using simulated data from AIDA to test out their approach.

QSIM (Fig. 1) is the qualitative simulator that was developed by exploiting the physiological knowledge available in the literature, and in particular by referring to the studies presented in the main AIDA model paper[2] and the report of Berger and Rodbard.[3] Based on this information Bellazzi et al.[15] initialized the membership functions of the fuzzy systems.

As such, this work[15] describes a novel approach to the identification of nonlinear dynamic systems, which integrates fuzzy systems and qualitative models. The simulations obtained from a set of qualitative differential equations have been used to automatically encode the available knowledge in a fuzzy rule-based system. The system is then tuned to a set of experimental data (from AIDA). The results obtained so far demonstrate that the presented framework generates fuzzy systems that may be used for a quick and reliable identification of nonlinear systems in diabetes care.[15]

How to improve fuzzy-neural system modeling by means of qualitative simulation

In further work by the same group the authors[16] have reported that the main problem in

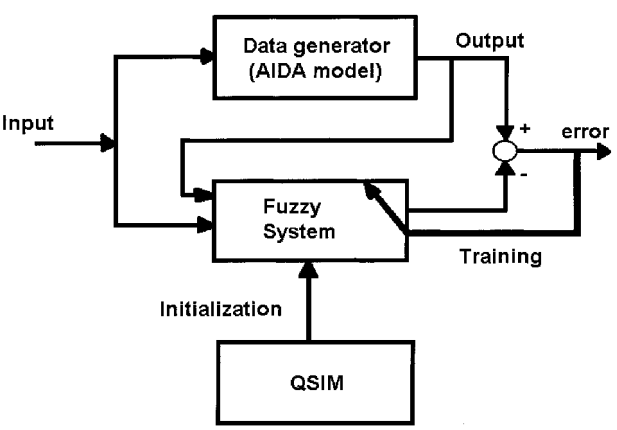

FIG. 1. Summary of the scheme used for fuzzy system identification. Derived from Bellazzi et al.[15]

efficiently building robust fuzzy-neural models of nonlinear systems lies in the difficulty of defining a "meangingful" fuzzy rule-base. The authors' solution is based on a hybrid method that integrates fuzzy systems with qualitative models. Bellazzi and colleagues introduce qualitative models to exploit the available, although incomplete, *a priori* physical knowledge of the system with the goal to infer, through qualitative simulations, all of its possible behaviours. The authors show that a rule base, which captures all of the distinctions in the system states, is automatically generated by encoding the knowledge of the system dynamics described by the outcomes of its qualitative simulation. Such a rule-base properly initializes a fuzzy identifier, which is then tuned to a set of experimental data.[16]

Once again, simulated data coming from AIDA was used for training and testing, with data sampled from the AIDA program every 15 min.

Adaptive controllers for intelligent monitoring

In their report entitled "Adaptive controllers for intelligent monitoring," Bellazzi et al.[17] have also described an approach designed around the usual scheme of diabetes outpatient management. This is based on (i) a period evaluation of patients' metabolic control performed by the physician, and (ii) patient-tailored tables for self-adjustment of insulin dosages. Following this scheme the authors have defined a system built on a two-level architecture. The High Level Module exploits both medical knowledge and clinical information in order to assess an insulin protocol, defined in terms of insulin timing, type, and total amount. The High Level Module exchanges information with the Low Level Module in order to define the control actions to be taken at the low level, as well as to periodically evaluate protocol adequacy on the basis of patient data. The goal of the Low Level Module, whose characteristics can be adaptively modified by the High Level Module, is to suggest the next insulin dosage depending on the actual blood glucose measurement and a certain predefined insulin delivery protocol. The Low Level Control Module is based on an adaptive controller, consisting of a fuzzy set controller and an ARX (Autoregressive eXogenous input) Model. The scheme presented by the authors[17] may be conveniently viewed in a telemedicine context, in which the low-level controller is implemented on a portable device communicating with the high-level controller, implemented on a remote computer.

The low-level controller for this work[17] was tested using the AIDA diabetes simulation package. Bellazzi et al.[17] simulated a patient weighing 40 kg. They considered the following protocol: three injections per day of regular insulin, in correspondence to each meal, and two injections of NPH (intermediate-acting) insulin, at lunchtime and at bedtime. The total amount of daily insulin was 1 U/kg. The proportion of the total amount for each dosage followed this scheme: before breakfast (7 a.m.) 25% of regular insulin, before lunch (1 p.m.) 20% of regular insulin and 10% of NPH insulin, before dinner (6 p.m.) 20% of regular insulin, and at nighttime (10 p.m.) 25% of NPH (intermediate-acting) insulin.

The authors compared the performance of four different control strategies for the implementation of the regulator. The first strategy (a) just used a protocol without low-level control (i.e., an open-loop strategy). The second strategy (b) was an implementation via fuzzy rules of a decision rule set routinely used by patients in their self-monitoring activity. In other words it was a fuzzy controller having as an input variable the difference between the actual and the desired blood glucose level (i.e., a rule-based strategy). The third strategy (c) was based on a fuzzy controller exploiting the ARX model predictions. The fourth strategy (d) used the fuzzy controller operating with "perfect" predictions (i.e., the blood glucose level that would be obtained if the open-loop strategy was followed).

The authors simulated the control system over 192 h (8 days) with control actions and measurements at each meal (i.e., three times per day, at 7 a.m., 1 p.m., and 6 p.m.). Bedtime was taken to be 10 p.m. Finally, it was assumed that the output measurements would be affected by random Gaussian noise.

It was observed that the rule-based control (strategy b) involved insulin adjustments of positive and negative signs, whereas the fuzzy

controller with the ARX model (strategy c) and the fuzzy controller with "perfect" predictions (strategy d) had only positive arguments. The authors concluded that these observations were related to the bad control performed with strategy (b)—which produced oscillations in control actions as well as in the patient's blood glucose level. While initial prototype testing using data from AIDA simulations has been useful and encouraging, further testing of these approaches with real patient data is clearly required.

A distributed system for diabetic patient management

In their report entitled "A distributed system for diabetic patient management," Bellazzi et al.[18] describe a telemedicine-based prototype for diabetes patient management. They present its architecture, the technical solutions adopted, and the methodologies on which it is based. The system is designed to provide decision support in a distributed environment, and is composed of two modules: (1) a Patient Unit and (2) a Medical Unit, connected by telecommunications services. The authors outline how the two modules can interact to perform effective monitoring and cooperative control of glucose metabolism. In particular, Bellazzi and colleagues detail the data analysis tasks performed by the two units and how the results are used to assist patients and physicians in revising and adjusting the therapeutic protocol. The reported prototype implementation uses the HyperText Transfer Protocol (HTTP) as the communications protocol and HyperText Markup Language (HTML) pages as the graphical user interface.

The authors describe how the output of the reasoning module of the system is an ordered list of alternative diabetes regimens that should be able to solve the metabolic problems detected by the system. These alternative regimen protocols are presented to the physician who can then try them out using the AIDA simulator[2] and choose the most suitable one.

Protocol-based reasoning in diabetic patient management

In their report entitled "Protocol-based reasoning in diabetic patient management," Montani et al.[19] propose a system for teleconsultation in the management of patients with insulin-dependent diabetes mellitus, accessible through the use of the Internet. The prototype is able to collect blood glucose monitoring data, analyse them through a set of tools, and suggest therapy adjustments in order to tackle the identified metabolic problems and fit these to the patient's needs. The program tries to generate advice and employs it to modify the current therapeutic protocol, presenting the physician with a set of feasible solutions, from which he/she can choose the most appropriate one. For this work, in order for the prototype to be able to calculate the effectiveness of a given food intake, food "activity" was calculated using the AIDA model approach.[2]

In all three Masters theses,[11–13] and subsequent publications from this group,[14–19] AIDA was used as a simulator of blood glucose dynamics to test out the decision support prototypes.

Summary of experience

With all these reports it is interesting to consider if any particular conclusions can be drawn from these studies. This is especially the case as a number of different (but related) computational approaches have been trialled. However, from the actual papers it is not immediately clear what the conclusion of all this effort has been. For instance, might one computer-based approach be better or more promising than another? Or is the "jury still out" on all of this? This question was put to Dr. Bellazzi, who replied (personal communication, 2002) as follows:

> "Sometimes we have applied techniques in the diabetes field just to test methods (for instance our work on qualitative modeling and fuzzy systems), while other times it was the application that motivated the use of the particular methods. In our experience the current state of the art seems to be the following:
>
> (1) For decision support we have implemented a rule-based system, that appears reasonably stable and seems to work properly. It has still to be established who should be using it (patients or physicians). We also make use of a case-based reasoning system.

For this the case base needs to be chosen in order to span a good variety of cases. Only in this way does case based reasoning seem to add value. More recently we have implemented a probabilistic model for modal day propagation, that seems interesting, although at present we cannot say for sure whether this will be really useful.

(2) What is currently in every day use in our Multi-Access Services for the Management of Diabetes Mellitus (M2DM) Project is a warning system that automatically analyzes data by exploiting the modal day concept and simple rules. The system then triggers automatic messages and uses SMS (mobile phone technology) to notify patients and/or physicians, highlighting problems and priorities. This system appears to work well and seems a good way to cope with the information overload problem always present in telemedicine systems."

As well as the group from Pavia, other researchers have also made use of AIDA for testing out their decision support prototypes.

TESTING DECISION SUPPORT PROTOTYPES #2

McCausland and colleagues from the University of Melbourne, Melbourne, Victoria, Australia have been developing an expert system to advise on insulin dosage adjustment in diabetes. The approach uses a combination of general rules, and rules that can be extracted from patient data—the idea being to produce a knowledge-based system that is able to "learn" from the data and as a result automatically fine-tune the rules. These researchers are using simulated data from AIDA for initial testing of their decision-support prototype (L. McCausland, personal communication, 1999).

TESTING DECISION SUPPORT PROTOTYPES #3

Staite from University College Northampton in England has been developing a rule-based expert system to try and assist in patient insulin therapy self-management.[20] The program uses IF . . . THEN . . . type production rules and concentrates on insulin dosage adjustment for insulin-dependent (type 1) diabetic patients. The prototype is based around the user inputting blood glucose values taken at pre-determined times of the day (e.g., before each meal). Staite has made use of simulated blood glucose data from AIDA for initial testing of this decision support approach.

ARTIFICIAL NEURAL NETWORKS—BACKGROUND

An artificial neural network (ANN) is an information processing model inspired by the way the densely interconnected parallel structure of the brain is thought to process information. In information technology terms an ANN is a computer system made up of a number of simple, highly interconnected processing elements that process information by their dynamic state responses to external inputs. ANNs are loosely modelled on the neuronal structure of the mammalian cerebral cortex, but obviously on a much smaller and less complex scale.

ANNs are believed to be of use in situations that involve the identification of highly nonlinear and/or empirical systems, situations that conventional computer systems are less reliable at solving.

A simple example of an ANN making use of a three-layer architecture is shown in Figure 2. These layers are constructed from a number of interconnected nodes that contain an activation function. Patterns are presented to the network via the input layer, which communicates to one or more hidden layers where the actual pro-

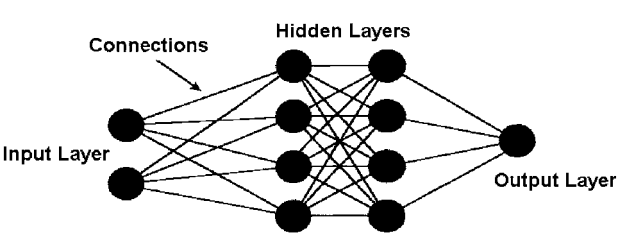

FIG. 2. Multilayered ANN. Derived from Pender.[21]

cessing is done via a system of weighted connections. The hidden layers then link to an output layer (Fig. 2).[21]

Most ANNs contain some form of learning rule that modifies the weights of the connections according to the input patterns that are presented. Although there are many types of training algorithms used by neural networks, much of the work described in this review has made use of the "back-propagation" training algorithm.

Using this technique input vectors and the corresponding output vectors are used to train a network until it can approximate a function, associating input vectors with specific output vectors. It is believed that networks with biases, a sigmoid layer, and a linear output layer should be capable of approximating most functions with a finite number of discontinuities.[21]

Apparently, well-trained back-propagation networks tend to give reasonable answers when presented with inputs that they have never seen before. Typically a new input will lead to an output similar to the correct output for input vectors that were used in training, and that are similar to the new input being presented. This generalisation property makes the approach of interest to researchers in the diabetes field. As such, theoretically, it should be possible to train a neural network on a representative set of input/target pairs and get reasonable results without necessarily training the network on all possible input/output pairs.

ARTIFICIAL NEURAL NETWORKS—UNIVERSITY OF STRATHCLYDE PROTOTYPES

Dr. Bill Sandham and his group from the Institute of Communications and Signal Processing (Department of Electronic and Electrical Engineering) at the University of Strathclyde (Glasgow, Scotland, U.K.) have developed a number of ANN prototypes, and have done some interesting preliminary evaluation work. For the following research the MATLAB Neural Networks Toolbox software was used to create various ANN prototypes, which were then trained using data from AIDA and/or clinical data. Using simulator-generated data was described as removing "all the ethical and practical problems associated with collecting data from real patients,"[21] and seems to have facilitated the rapid development/prototyping of a number of novel neural networks.

Prototype #1

The first ANN prototype, developed by Pender[21] under Dr. Sandham's supervision, made use of a hidden layer of 10 tan-sigmoid neurons to receive inputs directly and then broadcast their outputs to a layer of linear neurons that computed the network output (Fig. 3). All weights and biases within this ANN were initialised with random values, and the network was trained using training parameters and an

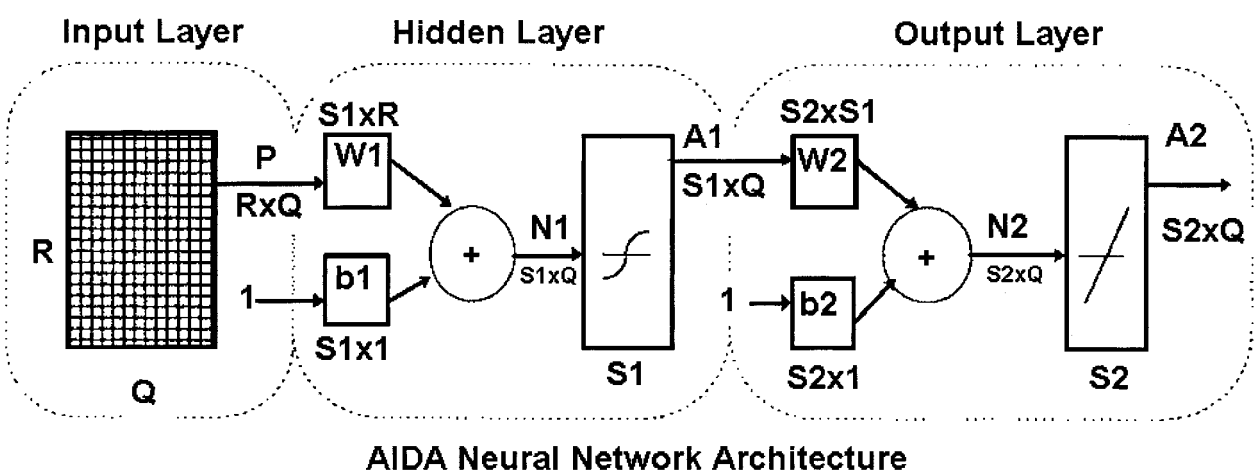

FIG. 3. AIDA ANN architecture. Derived from Pender.[21]

TABLE 1. AVERAGE AND RELATIVE ERRORS OF AN ANN COMPARED WITH SIMULATED "PATIENT" DATA FROM THE AIDA DIABETES SIMULATOR

Trained with	*Tested with*	*Average error |AIDA–ANN| [mmol/L (mg/dL)]*	*Relative error |AIDA–ANN|/ AIDA (%)*
Patients 1 + 2	Patient 1	0.4 (7.2)	5.0
Patient 1	Patient 1	0.25 (4.5)	3.0
Patient 1	Patient 2	1.18 (21.2)	12.4

Derived from Pender.[21]

error goal. The error was simply the difference between the neuron response and the desired or target vector. This network prototype was trained using an approximation of Newton's method called the Levenberg–Marquardt technique, which is meant to be more powerful and sophisticated than the more commonly applied alternative "gradient descent" method.

This ANN was trained to predict blood glucose levels 2 h ahead using the data from case scenario 0001 ("Joy Wilson") in the AIDA database. A prediction time of 2 h was adopted because it was felt that a longer time could become more inaccurate because of meals. Input data used for this approach included (i) blood glucose and (ii) plasma insulin levels, as well as the AIDA glucose fluxes, (iii) peripheral glucose uptake, (iv) net hepatic glucose balance, (v) carbohydrate absorption from the gut, and (vi) renal glucose excretion. The data were read off the AIDA graphs every 15 min. The data taken at 15, 30, and 45 min past each hour were used as the training set. The remaining data were kept aside to be used as test data (previously unseen by the ANN). The main findings are shown in Table 1.

This ANN prototype was trained with a representative set of one patient's diet and insulin schedule and compared with the simulated original values. It predicted blood glucose levels to within 3% of the actual blood glucose level. However, when the neural network was trained on the data from one simulated subject (patient 1), less accurate predictions were made for another separate (different) subject (patient 2).[21] These results are shown graphically in Figure 4.

It will be interesting to see how truly generalisable these ANNs are, in practice, when trained with a large amount of data and tested against a new, different patient—separate from the training set.

Prototype #2

Study #1. This research has been taken forward in the same laboratory by Sandham et al.[22] who have been investigating making clinical blood glucose predictions using more sophisticated ANNs.[23,24]

As highlighted above, ANNs are said to be particularly useful in situations that involve the

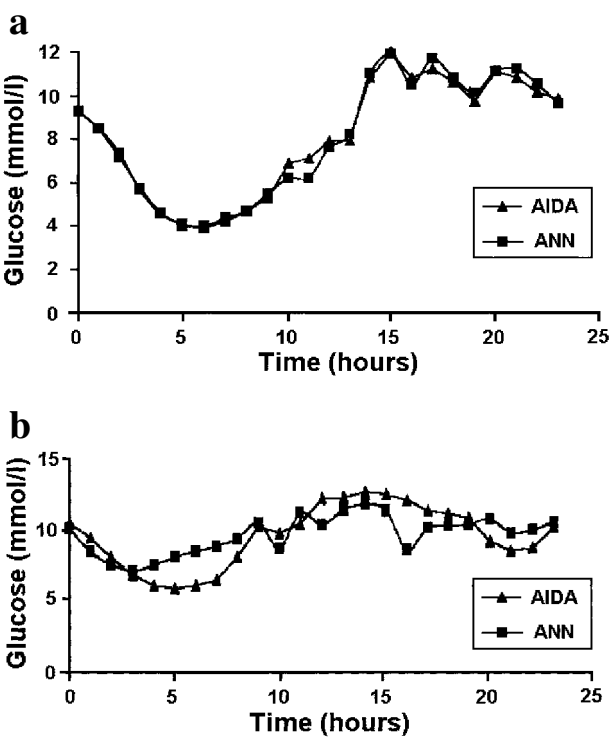

FIG. 4. **a:** Results from an ANN trained and tested with data from the same AIDA simulated "patient." Derived from Pender.[21] **b:** Results from an ANN tested with data from an AIDA simulated "patient" different from that used for ANN training. Derived from Pender.[21]

identification of highly nonlinear and/or empirical systems. They need to be trained on sets of patterns that display typical features of the system, but, once trained, are meant to be generalisable using a variety of other data; the knowledge/experience acquired through training is embedded in the weight matrices of the ANN. Furthermore, ANNs are meant to be able to assimilate information on a continuous basis, through a dynamic learning process.

For the second prototype a recurrent ANN was adopted, since this has been reported to demonstrate superior performance for prediction problems with short-term predictive accuracies ranging from 70% to 90% (in other fields).[25]

The recurrent ANN, as introduced by Elman[26] (Fig. 5), has delays in the feedback loops at the outputs of the recurrent layer, which enable previous time-step values to be used in the current time step.

Basically, the Elman recurrent network is a two-layer network with feedback from the first layer output to the first layer input. This recurrent connection allows the Elman network to both detect and generate time-varying patterns.

Training for this recurrent ANN was performed using back-propagation incorporating a momentum term and an adaptive learning rate. Two separate activation functions were employed: neurons in the recurrent layer used a tan-sigmoidal function, whereas neurons in the output layer used a linear function. By inspection it was found that 95 recurrent layer neurons gave the best results.

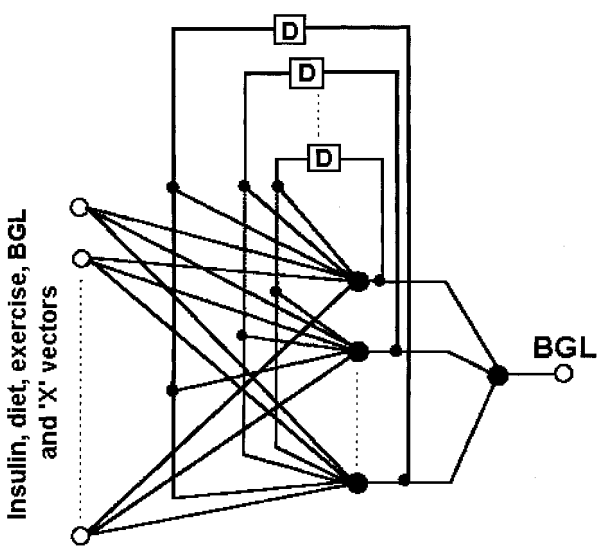

FIG. 5. Architecture of the Elman recurrent ANN[26] used for blood glucose prediction. BGL = blood glucose level. Derived from Sandham et al.[22]

For testing out this approach, initially six patients from the Diabetic Outpatient Department of Glasgow Royal Infirmary were selected. However, because of the number of blood glucose measurements required, only the results from two patients could eventually be used.

Comparisons of the ANN with actual measured data for these two patients showed that most of the ANN predictions were very close to the measured values provided by the patients' blood glucose meters [differences of 1.5 mmol/L (27 mg/dL) or less]. However, given that data were only available from two patients, a further study has been undertaken by Dr. Sandham using the AIDA diabetes simulator to provide a much larger amount of simulated patient data to more fully train and test the ANN.

Study #2. The aims and objectives of this further study were threefold: (i) to harvest diabetic patient data using the AIDA diabetes simulation package, (ii) to train the recurrent ANN with the generated data, and (iii) to compare AIDA's simulated blood glucose levels with predicted blood glucose levels from the ANN (W.A. Sandham, personal communication, 2002).

In total 50 data sets were harvested from the PC AIDA software (all from one case scenario), with six samples per data set. Each sample comprised an AIDA simulated blood glucose level, a carbohydrate intake (meal) in grams, two doses of (short- and intermediate-acting) insulin, and a target predicted blood glucose level a set time after the sample.

Carbohydrate intake and insulin dose were changed for each data set. All other parameters were kept constant. The first 40 data sets were used to train the network, and the final 10 data sets were used to test it. The average error of this approach was reported as being 0.7 mmol/L (13 mg/dL) (a relative error of 11%). It is important to stress, however, that these data are based on only one AIDA case scenario, albeit with multiple simulations.

Nevertheless, with these test results the authors have been encouraged to expand the AIDA-based ANN to include all the variable parameters incorporated within the AIDA diabetes software simulation package. The next step would then be to train this new network and test it to see if a similar standard of prediction could be obtained with a more comprehensive ANN. As intimated above, it will also be interesting to see how well such an ANN trained with, say, data from 30 separate simulated AIDA patients manages when presented with a completely new case.

Looking ahead, the interest in being able to predict blood glucose levels accurately is based on the fact that if such reliable and accurate blood glucose predictions were possible, then therapy planning opportunities would clearly arise. With an ANN approach, the rationale is that a therapy optimiser could be encoded alongside the blood glucose predictor, within the ANN (Fig. 6).

The ANN predictor enables the predicted blood glucose level at time $k + 1$ (BGL_{k+1}) to be derived from the actual blood glucose at time k (BGL_k), and the anticipated diet, exercise, and insulin regimen at time k (DEI_k). In addition, previous values of these parameters, via the delay units (D), are used to optimise performance. All these parameters are then employed, together with the target BGL (BGL_T) at the time $k + 1$, to produce an optimum therapy plan in terms of diet, exercise, and insulin regimen (DEI_{k+1}). After suitable ANN training, weighting factor matrices can be obtained for each patient, for either glycaemic prediction or attempted therapy optimisation.

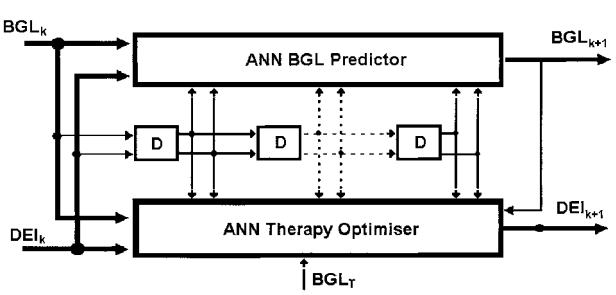

FIG. 6. Summary of the schematic of an ANN blood glucose predictor and ANN therapy optimser. BGL = blood glucose level; DEI = diet, exercise, and insulin regimen; D = delay units. See text for explanation. Derived from Sandham et al.[22]

One word of caution, however, should be sounded. This is to reinforce the fact that the AIDA model is a steady-state model of glucose–insulin interaction.[2] As a result the model does not simulate the transient conditions that might result following a change in the insulin or dietary regimen. Rather the program simulates the longer-term effects of those changes, some 48–72 h later. Given this, while the AIDA program is fine to be used for initial lab/bench-testing of various prototypes, it will be self-evident that there is still a very clear need eventually for proper clinical testing with real patient data.

In this respect, when used as described in this review, AIDA should theoretically be able to facilitate the rapid prototyping, development, and initial testing of new/novel computational approaches; although real patient data for formal testing will still be required thereafter. However, there is no reason that simulated data cannot be used to initially form the basis of (or "underpin") the experience/knowledge encoded within an ANN.

ARTIFICIAL NEURAL NETWORKS #2

Haque[27] from Brunel University in London, England has also used AIDA to provide blood glucose data to train an ANN of human carbohydrate metabolism in type 1 diabetes mellitus. Using a standard back-propagation network algorithm he developed an ANN to try and predict blood glucose levels for insulin-dependent diabetic patients following changes in either insulin therapy or carbohydrate intake. The network was trained using data collected from AIDA. The choice regarding the number of hidden nodes and system parameters like the initial weighting factor distribution, learning rate, etc., were determined by investigating the network through the training process.

Input parameters for the ANN included the blood glucose level, carbohydrate intake, and insulin regimen at time, t. Patient-specific parameters were not included. The output value was the predicted blood glucose level at some future time, $t + 1$. Training the ANN required 150 data sets. To expedite this process AIDA online—accessible directly at www.2aida.org/

online—was used as this allows easy access to the required data in an electronic form. Furthermore random values for the input carbohydrate range (0–80 g) and insulin range (0–40 units) were used—generated with the RAND() function in Microsoft Excel™—to permit as wide a variety of different cases to be simulated.[27]

The ANN was batch-trained with the computer left running for approximately 40 h to try all the hidden nodes and other system parameters until the predictive error was reduced to about 10%. Each node was trained for up to 3,000 cycles before trying the next node. The best results were obtained using 23 hidden nodes.[27]

Following training, the ANN was tested, first using 50 data sets that the neural network had seen before, and then using a further 50 data sets that the neural network had never seen before. The output result of each set was compared with the actual data obtained from the AIDA model. The ANN gave correct results (to the nearest integer) for 44 out of the 50 known data sets (12% error). For the 50 unknown data sets the ANN gave correct results for 38 cases (24% error).

Haque has done some interesting work. As he highlights in his report[27] there are various ways that the predictive capabilities of the AIDA-trained ANN might be improved. He suggests using some knowledge or algorithm such as Jacob's Enhanced Back Propagation method to automatically adjust parameters such as the initial weighting factor and the momentum factor—–rather than just setting these by inspection. Additional training could also be done—maybe up to 10,000 cycles for each hidden node—rather than just 3,000. Additional hidden layers might also be added. Furthermore it is important to note that not all components of the AIDA model were represented in the ANN. In particular, patient-specific parameters for the kidneys (such as the renal threshold of glucose and renal function), insulin sensitivity (liver and peripheral), and the patient's weight were not included. It is to be expected that incorporating these clinical pa-

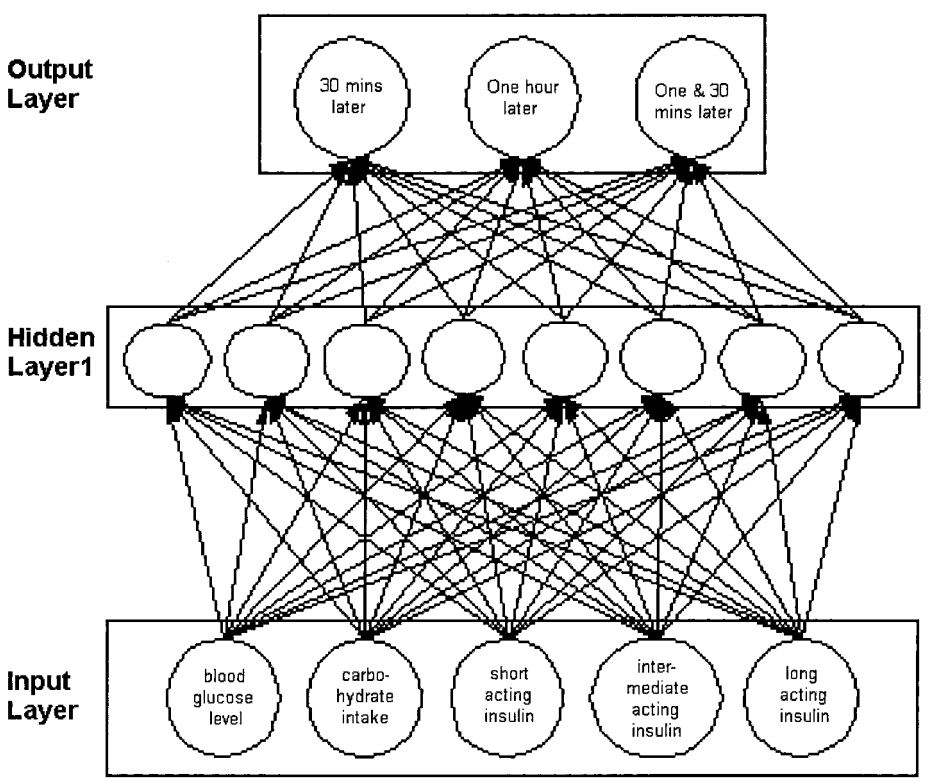

FIG. 7. Structure of a neural network model. Derived from Chang et al.[28]

rameters in future work could lead to an improvement in the predictive accuracy of the ANN.

ARTIFICIAL NEURAL NETWORKS #3

Chang et al.[28] have developed an Internet-based home monitoring prototype for diabetes care. Their approach proposes a telemedicine expert system to communicate the results of home monitoring of diabetes to a central hospital database, and to the physician in charge of the patient. The expert system planned for decision support uses a neural network.

The strategy of the expert system uses a mapping method with back-propagation training of the neural network. With the back-propagation method, the weighting factors are controlled to reduce the performance function—that is the difference between the actual and desired network outputs. The basic back-propagation learning method adopted by Chang et al.[28] aims to update the network weighting factors and biases in the direction in which the performance function decreases most rapidly.

In this network, there are five inputs: (i) present blood glucose level, (ii) carbohydrate intake, and the amounts of (iii) short-, (iv) intermediate-, and (v) long-acting insulin. The outputs from this network are the blood glucose levels that are predicted a few hours later. Figure 7 overviews the structure of the neural network.

The function of this approach is to establish patients' response patterns and try and estimate their blood glucose levels in the near future. The system was evaluated by Chang et al.[28] using test case scenario data obtained from the AIDA diabetes simulator.

WHAT TO READ IN PART 2?

In Part 2 of this review,[29] issues surrounding dietary assessments, developing new diabetes models, and performance monitoring of closed-loop insulin delivery devices are considered. Attempts to try and apply a "learned space of parameter interactions" using simulated blood glucose data are overviewed. A business plan utilising telehealth-care technology within the home, and the development of a Web-based educational diabetes simulator are also considered. On-going research using AIDA is mentioned, and issues surrounding the generation of simulated blood glucose data for training/validation of other diabetes prototypes are discussed. Finally, evaluation/validation research with AIDA and issues surrounding the placing of technical information on the Web are considered.[29]

ACKNOWLEDGMENTS

The author would like to thank the researchers whose work is overviewed in this review article and on the Web at www.2aida.org/aida/research2.htm—many of whom have provided details of their on-going research. Particular thanks are due to Dr. Riccardo Bellazzi and colleagues (Pavia, Italy), Dr. Bill Sandham and colleagues (Glasgow, Scotland, U.K.), Mr. Abdul Haque (London, England, U.K.), and Dr. Seok-Cheol Chang and colleagues (South Korea).

REFERENCES

1. Lehmann ED: Experience with the Internet release of AIDA v4—an interactive educational diabetes simulator. Diabetes Technol Ther 1999;1:41–54.
2. Lehmann ED, Deutsch T: A physiological model of glucose-insulin interaction in type I diabetes mellitus. J Biomed Eng 1992;14:235–242.
3. Berger M, Rodbard D: Computer simulation of plasma insulin and glucose dynamics after subcutaneous insulin injection. Diabetes Care 1989;12:725–736.
4. Tatti P, Lehmann ED: Utility of the AIDA diabetes simulator as an interactive educational teaching tool for general practitioners (primary care physicians). Diabetes Technol Ther 2001;3:133–140.
5. Tatti P, Lehmann ED: Use of the AIDA diabetes simulation software—www.2aida.org—as an interactive educational tool for teaching student nurses. Diabetes Technol Ther 2001;3:655–664.
6. Tatti P, Lehmann ED: Using the AIDA—www.2aida.org—diabetes simulator. Part 1: Recommended guidelines for health-carers planning to teach with the software. Diabetes Technol Ther 2002;4:401–414.
7. Lehmann ED: Preliminary experience with the Internet release of AIDA—an interactive educational diabetes simulator. Comput Methods Programs Biomed 1998;56:109–132.

8. Lehmann ED: Interactive educational simulators in diabetes care. Med Inf (Lond) 1997;22:47–76.
9. Lehmann ED: The freeware AIDA interactive educational diabetes simulator—http://www.2aida.org—(1) A download survey for AIDA v4.0. Med Sci Monit 2001;7:504–515.
10. Lehmann ED: The freeware AIDA interactive educational diabetes simulator—http://www.2aida.org—(2) Simulating glycosylated haemoglobin (HbA_{1c}) levels in AIDA v4.3. Med Sci Monit 2001;7:516–525.
11. Guglielmann R: Non linear systems identification by using qualitative models and fuzzy systems [in Italian] [M.Sc. Dissertation]. Pavia, Italy: University of Pavia, 1996.
12. Bossi G: Neuro-fuzzy system identification based on qualitative models: an application to physiological systems [in Italian] [M.Sc. Dissertation]. Pavia, Italy: University of Pavia, 1996.
13. Telecco P: Metabolic control in diabetic patients: decision support systems [in Italian] [M.Sc. Dissertation]. Pavia, Italy: University of Pavia, 1994.
14. Bellazzi R, Ironi L, Guglielmann R, Stefanelli M: Qualitative models and fuzzy systems: an integrated approach for learning from data. Artif Intell Med 1998;14:5–28.
15. Bellazzi R, Ironi L, Guglielmann, Stefanelli M: Learning from data through the integration of qualitative models and fuzzy systems. Lecture Notes Artif Intell 1997;1211:501–512.
16. Bellazzi R, Guglielmann R, Ironi L: How to improve fuzzy-neural system modeling by means of qualitative simulation. IEEE t-Neural Networks 2000;11: 249–253.
17. Bellazzi R, Siviero C, Stefanelli M, De Nicolao G: Adaptive controllers for intelligent monitoring. Artif Intell Med 1995;7:515–540.
18. Bellazzi R, Riva A, Larizza C, Fiocchi S, Stefanelli M: A distributed system for diabetic patient management. Comput Methods Programs Biomed 1998;56: 93–107.
19. Montani S, Bellazzi R, Larizza C, Riva A, d'Annunzio G, Fiocchi S, Lorini R, Stefanelli M: Protocol-based reasoning in diabetic patient management. Int J Med Inform 1999;53:61–77.
20. Staite R: The use of IT-mediated patient self management in insulin therapy management [B.Sc. Dissertation]. Northampton, UK: University College Northampton, 1999.
21. Pender JE: Modelling of blood glucose levels using artificial neural networks [Dissertation]. Glasgow: University of Strathclyde, 1997.
22. Sandham WA, Nikoletou D, Hamilton DJ, Patterson K, Japp A, MacGregor C: Blood glucose prediction for diabetes therapy using a recurrent artificial neural network. In: Proceedings, EUSIPCO-98, IX European Signal Processing Conference, Rhodes Island, Greece, Vol. 11. 1998:673–676.
23. Sandham WA, Hamilton DJ, Japp A, Patterson K: Neural network and neuro-fuzzy systems for improving diabetes therapy. In: Proceedings, 20th International Conference of the IEEE Engineering in Medicine and Biology Society, 29 October–1 November 1998, Hong Kong, Vol. 20, Part 3/6. Hong Kong: IEEE Engineering in Medicine and Biology Society, 1998:1438–1441.
24. Sandham WA, Hamilton DJ, Japp A, Patterson K, Nikoletou D, MacGregor C: Use of artificial neural networks for improved diabetes therapy. In: Proceedings, ISSC-98, Irish Signals and Systems Conference. Dublin: Dublin Institute of Technology, 1998.
25. Patterson WD: Artificial Neural Networks—Theory and Applications. Singapore: Prentice Hall, 1996.
26. Elman JL: Distributed representations, simple recurrent networks and grammatical structure. Machine Learn 1991;7:195–225.
27. Haque A: Modelling human metabolism using neural network [Dissertation]. London: Brunel University, 1999.
28. Chang SC, Ryoo SM, Yu SY, Ahn BH: Development of an internet-based home monitoring system for diabetes mellitus. In: Proceedings, 4th IEEE International Symposium on Consumer Electronics (ISCE'99). Melaka, Malaysia: IEEE, 1999.
29. Lehmann ED: Research use of the AIDA www.2aida.org diabetes software simulation program: a review—Part 2. Generating simulated blood glucose data for prototype validation. Diabetes Technol Ther 2003;(in press).

Address reprint requests to:
Dr. Eldon D. Lehmann
c/o www.2aida.org Diabetes Simulator Development Team
P.O. Box 46104
London, EC2Y 8WN, UK

E-mail: www@2aida.org