A Web-based System for the Intelligent Management of Diabetic Patients[*]

Alberto Riva, Riccardo Bellazzi[*], Mario Stefanelli[*]
IRCCS Policlinico San Matteo, Pavia, Italy
[*]Department of Computer Science, University of Pavia, Italy

A recent statement by the American Diabetes Association [1] has defined a basic set of standards of medical care for diabetic patients. These standards have a two-fold aim: a) allowing health care professionals to set treatment goals and assess the quality of treatment they provide, and b) allowing patients to evaluate their current clinical state and the quality of care they receive. The reduction of the blood glucose level to near-normal values is indicated as the primary therapeutic goal, and the use of intensive treatment programs is suggested for this purpose [2]. An intensive treatment program for insulin-dependent diabetes mellitus (IDDM) usually includes frequent self-monitoring of blood glucose, regular meal and exercise planning, physiologically-based insulin regimens, continuing patient education, and the periodic assessment of treatment goals.

The Medical Informatics Laboratory of the Department of Computer Science of the University of Pavia, in collaboration with the San Matteo University Hospital of Pavia, is currently working on the design and implementation of a computer-based system to support the activities of both patients and physicians in the field of IDDM management, according to the above cited guidelines. The system we are designing will be able to to perform in an automated way a whole range of routine medical tasks (such as data collection and interpretation, alarm generation, etc.) and to assist the physician in making timely and informed therapeutic decisions when necessary. It will also increase the rate of information exchange between the patient and the health care center, thereby reducing the need for frequent visits and the risk of unhandled dangerous events. The system is being developed in the World-Wide Web (WWW) environment, in order to take advantage of the pervasiveness and low-cost features of current Internet-based technologies [3].

The following sections will describe in more detail the medical context of IDDM management, the architecture of the system and the most important implementation choices. Finally, we will present some snapshots of the current prototype

---

[*]Reprinted from: *M. D. Computing*, 14(5):360-364, nov 1997

of the system.

1 The medical context

Insulin-dependent diabetes mellitus is a widespread chronic illness that accounts for a very large part of the health care budget in developed countries. The common therapy of IDDM uses exogenous insulin administrations to control the blood glucose level, trying to prevent oscillations that can lead to dangerous events (e.g. hyperglycemic crises) and to invalidating long-term complications. The recent Diabetes Control and Complications Trial has conclusively demonstrated a strong correlation between the quality of blood glucose control that can be obtained through insulin therapy and the chance to develop long-term complications [2]. On the other hand, diabetic patients are normally not hospitalized, and are therefore required to play an active, although controlled, role in the therapeutic and decisional process. More in detail, the patient is required to follow a predefined therapeutic protocol that specifies the total insulin requirement, the number of daily injections, the insulin type and dose of each injection. The patient is also allowed to modify the insulin doses in response to events such as an out-of-bounds level of blood glucose or a change in life-style, using simple adjustment rules provided by the physician along with the protocol. Finally, the patient must monitor and record the values of several physiological parameters, such as blood glucose level and glycosuria. The physician regularly evaluates the state of the patient by periodic visits during which the data gathered since the previous visit are reviewed and the therapeutic protocol is modified if needed.

2 System design

As the first step in the design of the system, we have performed a thorough analysis of the therapeutic and decisional processes involved in the management of IDDM. The results of the analysis pointed out the most relevant characteristics of the disease from the point of view of information flows and task allocation.

Information management in the therapy of IDDM is a critical task. The patient's self-monitoring activity generates a very large amount of data: on the average, three measurements are collected on each day, just for blood glucose. On the other hand, the quality of the data can be quite poor, due to the presence of missing or unreliable data and, in any case, the data are not sufficient to reconstruct the actual dynamics of the physiological variables being measured. Moreover, the complexity of glucose metabolism makes it extremely difficult to determine a reliable mathematical model on which to base simulations and predictions. *Ad hoc* techniques are therefore needed to handle the data and to cope with its low information content.

From the organizational point of view, the management of IDDM is a *distributed* activity: since the patients are usually not hospitalized, the planning and the implementation of the therapy take place in different places and at different times. It is a *collaborative* activity, since the decision tasks typical of IDDM management are performed in a cooperative fashion by the physician and the patient himself. It is an activity involving a large number of specialized tasks (data collection and analysis, protocol revision, etc), that are carried out by several different agents, who possess different skills and amounts of knowledge.

The architecture of the system has therefore been modeled on the characteristics that we have just outlined. In particular, the system must be based on telemedicine techniques in order to allow a frequent and reliable communication between the patient and the care center. It must be extremely usable and it must be integrated with the current practice of IDDM management, so that it can be easily accepted both by health care professionals in a medical setting and by unskilled patients at home.

The result is a distributed architecture centered around two main components: a *Patient unit* and a *Medical unit*. The patient unit is located on a personal computer at the patient's home, while the medical unit resides on a workstation in the health care center. The patient unit is able to periodically establish a communication with the medical unit during which the patient's self monitoring data are uploaded to the medical unit, and a revised therapeutic protocol is received, if needed. The same link can also be used by the patient and the physician to exchange messages and to signal alarms.

The following section will describe in more detail the services that are provided by the two unit in the current prototype of the system.

2.1 The patient unit

Database The most important task of the patient unit is the short term storage of the data collected by the patient. Blood glucose measurements are usually obtained through a reflectometer, and can subsequently be downloaded to the patient unit. When this is not possible, the data must be manually input by the patient in the electronic equivalent of the patient's diary, which is today the most widely used data recording system. Other data to be collected include glycosuria measurements, the insulin injection doses and timings, and qualitative information on the meals. All these data are stored in a temporary database, until they are uploaded to the medical unit.

Data analysis Although the complete interpretation of the collected data is the responsibility of the medical unit, the patient unit must be able to perform a preliminary data analysis in order to detect situations that require an immediate intervention. Due to the limited computational capabilities of the patient unit, only simple analysis techniques based on local trend detection can be employed at this stage [4].

Reasoning The patient unit should assist the patient in the day-to-day implementation of the insulin therapy, by suggesting adjustments to the doses prescribed by the protocol in response to the current value of the blood glucose. This is done by using a simple fuzzy controller whose parameters are set by the medical unit.

2.2 The medical unit

Database The medical unit must collect and store the data coming from all the patients undergoing insulin therapy. This is needed in order to be able to perform long-term analysis of the clinical state of a patient, as well as population-wide statistical studies. Moreover, the database is used to store patients' anagraphic data, predefined therapeutic protocols, information on insulin types, messages, etc.

Data analysis The medical unit employs a variety of data analysis techniques to extract the largest possible amount of knowledge from the available data. Tools that can be used range from simple statistical indicators (e.g. average BGL over one day) to temporal abstractions (detection of trends and other significant temporal patterns [5]), to high-level probabilistic descriptors (e.g. modal day [6]). The main purpose of the analysis is to detect and point out possible problems associated with the patient's clinical state.

Reasoning The medical unit includes a reasoning module that assists the physician in the periodic protocol revision task. Using a combination of heuristic and rule-based techniques, the system generates a set of suggestions on possible modifications to the protocol that the patient is following. The suggestions aim at solving the problems detected during the data analysis task, and can regard the insulin doses, the diet plan, or the amount of physical exercise to perform. For example, given the presence of recurrent morning hyperglycemia, a possible suggestion could be to increase the dose of intermediate insulin at dinner time. Applying all the suggestions to the current protocol leads to a number of alternative protocols that are presented to the physician. The physician can then accept one of the suggested protocols, or revise the current one. At the end of the process, the new therapeutic protocol is stored and will be transmitted to the patient unit during the next connection [7].

2.3 Communication

Communication issues are a fundamental aspect of a distributed architecture. The system components must be able to establish reliable, bidirectional communication links, and to exchange information using a common protocol. In a more abstract sense, communication requires the presence of a shared *ontology* that defines the concepts all the system components refer to [8]. For example, the meaning of "high BGL level" is defined in the ontology as a qualitative abstraction that maps a specified range of numerical values to a symbolic one. The ontology is also used to generate the structure of the database tables, so that the information circulating in the system can be stored permanently without modifications.

3 Implementation choices

The design of the system has been guided by a set of constraints that we regard as essential for the effectiveness of the project. First, the system architecture must support the easy integration of additional components: as new services are developed, possibly in response to feedback from the users, they must be made available without disrupting the functionalities of the rest of the system. We have therefore selected a distributed architecture based on the concept of *intelligent agent* [9] as the most suitable one for our project. The various services are represented as software agents able to offer a well-defined set of functionalities to other agents, and to exploit the other functionalities available in the distributed environment to accomplish their task. For example, the data analysis service relies on a *Temporal abstractions server*, that in turn resorts to the *Ontology server* to obtain the definition of the abstractions to perform. It is important to note that the whole process is completely transparent to the user, who only sees the overall result of the distributed reasoning process.

Another important requirement was the ability to obtain a working and reliable prototype in a relatively short time. The number and variety of services needed made it impossible to develop all of them from scratch; therefore, we have tried to rely on widely available and standard solutions where possible. For example, long-term data storage can be provided by any relational data base, the only necessary feature being the ability to receive SQL queries over the network. More important, this requirement has led to the choice to base the user interface of the whole system on the WWW paradigm. We have developed a specialized Web server (written in the Common Lisp language) that is able to access all the services available in the distributed environment and to generate HTML pages displaying the results [10] (HTML, HyperText Markup Language, is the language used to write pages in the WWW environment). The user connects to the server using a standard Web browser, and all the interaction takes places under the form of HTML pages and forms. We have also used the Java and JavaScript languages to enhance the capabilities of HTML when used as a tool to build graphical user interfaces. Our experience in this field has shown that the combination of modern network-based technologies and AI-based tools constitutes the basis for a very powerful and flexible development environment, that offers great advantages in terms of usability and availability [11]. The system can be accessed from any location on the network and from any platform, without the need for specialized interfaces or expensive hardware resources. Moreover, a single system can handle a large number of users, and can become the basis of a regional care network based on the Internet, supporting the exchange of clinical information and medical knowledge across different health care centers.

4 Examples

To conclude, we present a few snapshots of the medical unit interface, as it would appear to a user connecting to it with a normal web browser (in this case, Netscape Navigator). Figure 1 shows the page that gives access to the data contained in the patients' daily diary. After selecting the desired patient, the user can choose the variable to display among the ones stored in the database; the visualization can also be restricted to the data collected between specified start and end dates, and to a particular time of the day (e.g. Breakfast). Figure 2 shows the graph of the blood glucose measurements; the data were retrieved from the database according to the conditions specified in the previous page, and were then plotted by a Java applet. Finally, Figure 3 shows the form used to assign a predefined insulin administration regimen to the patient. The interface displays the insulin dose percentages for each insulin type and each time period during the day. The user can also modify the patient's total insulin requirement and the actual insulin types to be used.

5 Conclusions

We have described the design and implementation of a distributed, computer-based system for diabetes management. The system is the result of the integration of a large number of components, ranging from databases to data analysis tools to automated reasoning modules. The *intelligent agents* approach enabled us to develop the system in a flexible and reliable way. Interoperability between the various agents is ensured by a common communication protocol and by a shared medical ontology.

The systems is designed to assist physicians and patients in the routine tasks typical of IDDM management, integrating itself in a seamless way in the normal decision and information management processes. In order to satisfy the usability and acceptability requirements that are critical for the success of medical informatics applications, the system runs in the World-Wide Web environment: all the services can be accessed using a standard web browser and share the same, HTML-based graphical user interface.

The system is being developed in close collaboration with the medical staff of the University Hospital of Pavia, in order to ensure from the beginning its adequacy from the medical point of view, and will soon be tested on a set of voluntary patients.

6 Acknowledgments

The authors wish to thank Stefano Fiocchi, Cristiana Larizza, Stefania Montani and Cristiano Canobbio for their helpful collaboration. This work is supported by EU grant HC-1047, "T-IDDM – Telematic management of Insulin-Dependent Diabetes Mellitus".

References

[1] American Diabetes Association. Standards of medical care for patients with diabetes mellitus. *Diabetes Care*, 20(Supplement 1):S5–S13, January 1997.

[2] The Diabetes Control and Complication Trial Research Group. The effect of intensive treatment of diabetes on the development and progression of long-term complications in insulin-dependent diabetes mellitus. *The New England Journal of Medicine*, 14–329:977–986, 1993.

[3] J. J. Cimino, S. A. Socratous, and P. D. Clayton. Internet as a clinical information system: application development using the World–Wide Web. *Journal of the American Medical Informatics Association*, 2(5):273–284, September 1995.

[4] R. Bellazzi, C. Larizza, A. Riva, A. Mira, S. Fiocchi, and M. Stefanelli. Distributed intelligent data analysis in diabetic patient management. In J. Cimino, editor, *Proceedings of the 1996 AMIA Annual Fall Symposium*, pages 194–198, Philadelphia, 1996. Hanley and Belfus.

[5] C. Larizza, R. Bellazzi, and A. Riva. Temporal abstractions for diabetic patients management. In E. Keravnou, C. Garbay, R. Baud, and J. Wyatt, editors, *Proceedings of AIME 97*, number 1211 in Lecture Notes in Computer Science, pages 319–330. Springer, 1997.

[6] T. Deutsch, E.D. Lehmann, E.R. Carson, A.V. Roudsari, and K.D. Hopkins. Time series analysis and control of blood glucose levels in diabetic patients. *Computer Methods and Programs in Biomedicine*, 41:167–182, 1994.

[7] A. Riva and R. Bellazzi. High level control strategies for diabetes therapy. In P. Barahona, M. Stefanelli, and J. Wyatt, editors, *Proceedings of the Fifth Conference on Artificial Intelligence in Medicine Europe*, number 934 in Lecture Notes in Computer Science, pages 185–196. Springer, June 1995.

[8] M. A. Musen. Dimensions of knowledge sharing and reuse. *Computers and Biomedical Research*, 25:435–467, 1992.

[9] S. Russell and P. Norvig. *Artificial Intelligence: a modern approach.* Prentice Hall, 1995.

[10] A. Riva and M. Ramoni. Lispweb: a specialized HTTP server for distributed AI applications. *Computer Networks and ISDN Systems*, 28(711):953–961, 1996.

[11] A. Riva and R. Bellazzi. Distributed AI technologies for patient management. In *Proceedings of the AAAI Spring Symposium on Artificial Intelligence in Medicine*, pages 144–148, Stanford, CA, March 1996. AAAI.

Figure captions

1. The HTML page through which the user can select the data to retrieve from the database and how to display them (the name of the patient has been changed to ensure privacy).

2. A plot of the blood glucose values at breakfast, obtained through the form in Figure 1. The plot is generated by a Java applet, using the data retrieved from the database.

3. The HTML page that allows the user to view and modify an insulin plan. Percentages of the total insulin dose can be specified for each time-slice and for each insulin type.

Telematic Management of Insulin–Dependent Diabetes Mellitus

Diary: Dragoni Alessandro

424 BGL measures available, from 01–09–1993 to 28–02–1994.

Show:

BGL time series	NPH time series	Glycosuria time series	Ketonuria time series
●	○	○	○

Time Slice:

● All	○ Breakfast	○ Lunch
○ Mid Afternoon	○ Dinner	○ Bedtime

Time Period:

Start Date	End Date

Output as:

○ Text	● Graphics	○ Text & Graphics

Press Submit to display the results or Reset to clear the above fields.

Figure 1:

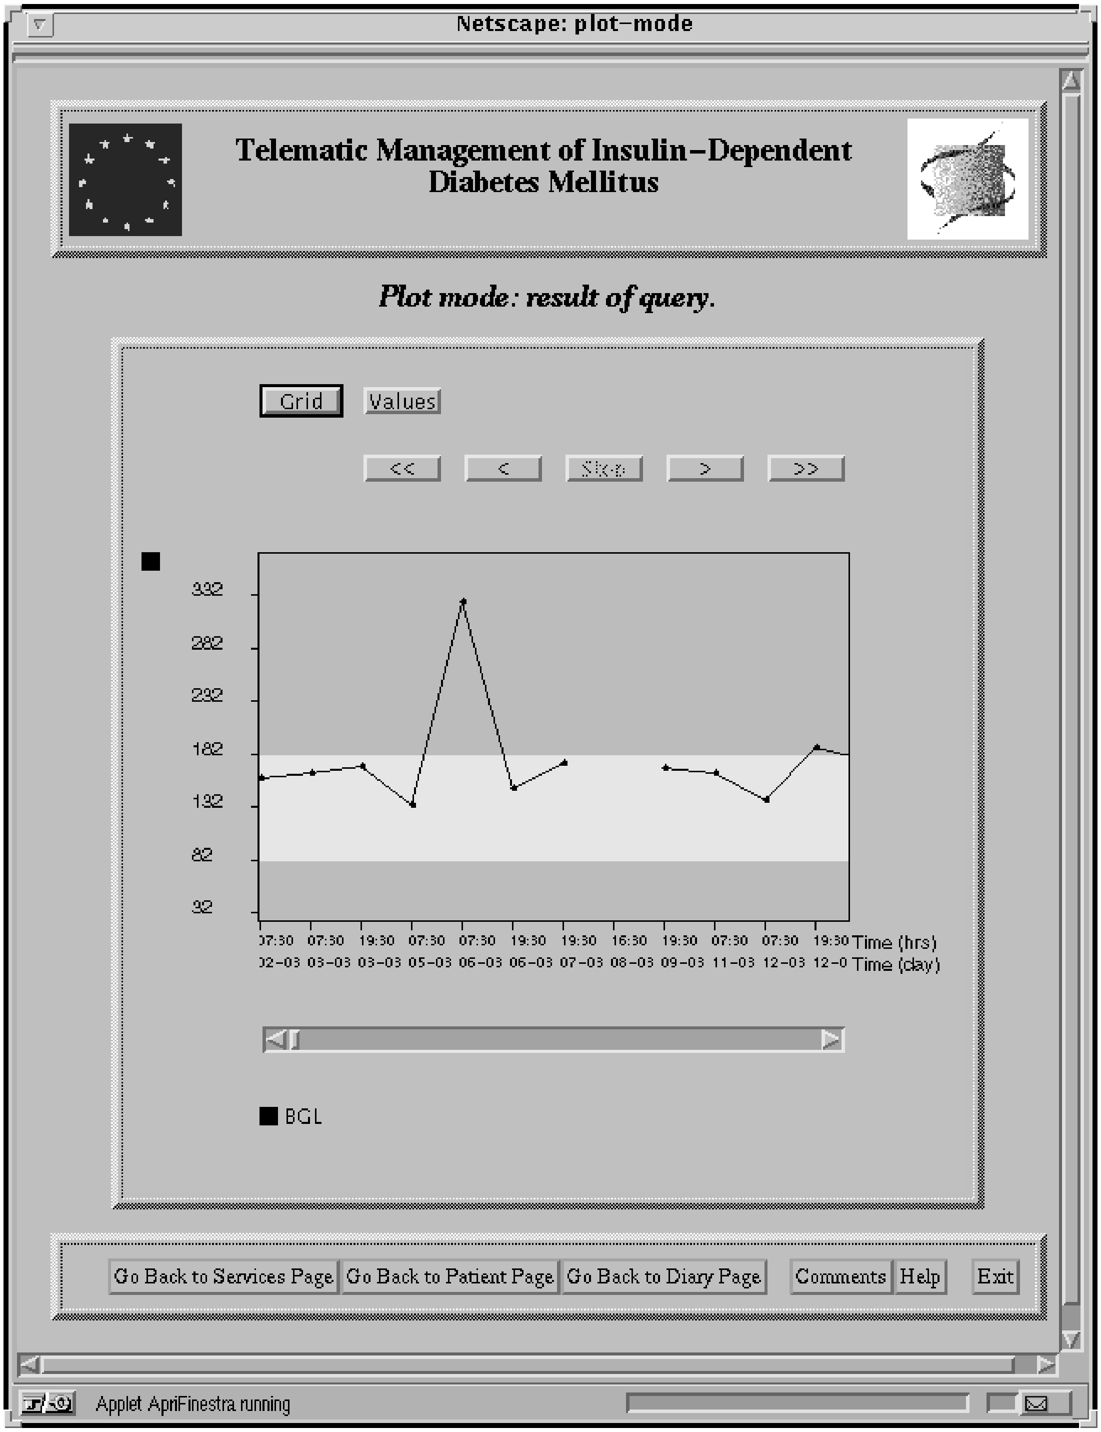

Telematic Management of Insulin-Dependent Diabetes Mellitus

Plot mode: result of query.

Applet ApriFinestra running

Figure 2:

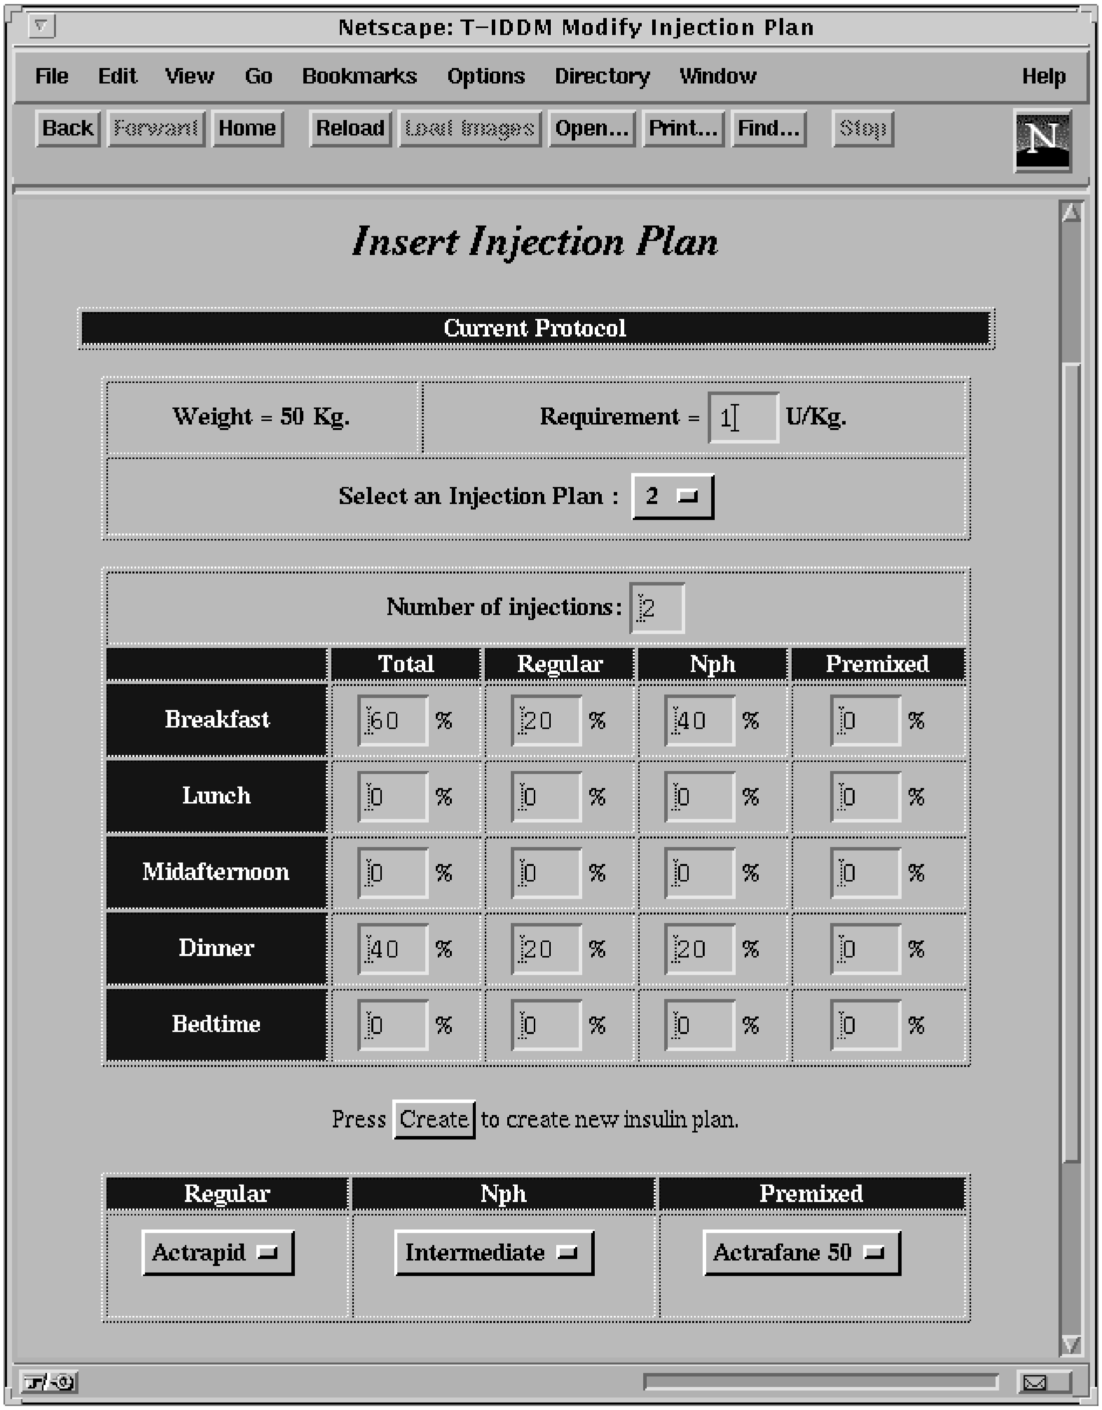

Insert Injection Plan

Current Protocol

Weight = 50 Kg.

Requirement = 1 U/Kg.

Select an Injection Plan : 2

Number of injections: 2

	Total	Regular	Nph	Premixed
Breakfast	60 %	20 %	40 %	0 %
Lunch	0 %	0 %	0 %	0 %
Midafternoon	0 %	0 %	0 %	0 %
Dinner	40 %	20 %	20 %	0 %
Bedtime	0 %	0 %	0 %	0 %

Press Create to create new insulin plan.

Regular	Nph	Premixed
Actrapid	Intermediate	Actrafane 50

Figure 3: