Computational Intelligence, Volume 25, Number 3, 2009

TOWARD CASE-BASED REASONING FOR DIABETES MANAGEMENT: A PRELIMINARY CLINICAL STUDY AND DECISION SUPPORT SYSTEM PROTOTYPE

CINDY MARLING,[1] JAY SHUBROOK,[2] AND FRANK SCHWARTZ[2]

[1]*School of Electrical Engineering and Computer Science, Russ College of Engineering and Technology, Ohio University, Athens*

[2]*Appalachian Rural Health Institute, Diabetes and Endocrine Center, College of Osteopathic Medicine, Ohio University, Athens*

This paper presents a case-based decision support system prototype to assist patients with Type 1 diabetes on insulin pump therapy. These patients must vigilantly maintain blood glucose levels within prescribed target ranges to prevent serious disease complications, including blindness, neuropathy, and heart failure. Case-based reasoning (CBR) was selected for this domain because (a) existing guidelines for managing diabetes are general and must be tailored to individual patient needs; (b) physical and lifestyle factors combine to influence blood glucose levels; and (c) CBR has been successfully applied to the management of other long-term medical conditions. An institutional review board (IRB) approved preliminary clinical study, involving 20 patients, was conducted to assess the feasibility of providing case-based decision support for these patients. Fifty cases were compiled in a case library, situation assessment routines were encoded to detect common problems in blood glucose control, and retrieval metrics were developed to find the most relevant past cases for solving current problems. Preliminary results encourage continued research and work toward development of a practical tool for patients.

Key words: case-based reasoning, diabetes management, intelligent decision support.

1. INTRODUCTION

People with Type 1 diabetes are unable to produce their own insulin, and thus they must depend on exogenous insulin to survive. Maintaining good blood glucose control is a difficult task for patients with Type 1 diabetes, who must keep their blood glucose levels as close to normal as possible to maintain their health and avoid serious complications (The Diabetes Control and Complications Trial Research Group 1993). Too little insulin results in elevated blood glucose levels, called hyperglycemia, which can lead to numerous complications over time, including blindness, neuropathy, and heart failure. Too much insulin results in depressed blood glucose levels, or hypoglycemia, during which patients may experience weakness, confusion, dizziness, sweating, shaking, and if not treated promptly, loss of consciousness or seizure.

A person with Type 1 diabetes tries to maintain blood glucose levels between assigned high and low targets and monitors actual levels through finger stick testing from four to six times a day. A patient on insulin pump therapy is continually infused with basal insulin at varying rates throughout the day and can instruct the pump to deliver additional doses, called boluses, before meals and to correct for hyperglycemia. The amount of bolus insulin depends on the amount of carbohydrate to be consumed, the current blood glucose level, the anticipated level of physical activity and the patient's own historical response to insulin. Insulin pump therapy differs from conventional insulin therapy, in which patients use syringes to inject themselves with a predetermined amount of insulin three or four times a day. With

This is an update and expansion of the paper "Toward Case-Based Reasoning for Diabetes Management," which was presented at the Workshop on Case-Based Reasoning in the Health Sciences at the International Conference on Case-Based Reasoning (ICCBR-07).

Address correspondence to Cindy Marling, Russ College of Engineering and Technology, Ohio University, Athens, OH 45701; e-mail: marling@ohio.edu

conventional insulin therapy, there is less opportunity to fine-tune control or to account for variations in daily routine.

In our experience, physical and lifestyle factors frequently combine in complex ways to impact blood glucose levels in people with Type 1 diabetes. Consider, for example, a patient of the third author who is on his high school wrestling team. Exercise, including participation in competitive sports, is known to be beneficial in controlling blood glucose levels. However, this patient's blood glucose could be too high, too low, or within range when he wrestled, depending on (a) whether he was at practice or at a competitive meet and (b) whether he won or lost his match. Therefore, to help each patient maintain good blood glucose control, despite individual variations and complex physical/social interactions, we are exploring case-based decision support. While we are not the first to employ case-based reasoning (CBR) in the diabetes domain (Bellazzi et al. 2002), we are the first to use it as the primary reasoning modality for supporting patients with Type 1 diabetes on insulin pump therapy. Our goal is to build a CBR therapy adviser that will help patients improve and maintain their control while reducing the data analysis workload on physicians.

New continuous glucose monitors (CGM) that record glucose values every 5 min, coupled with insulin pump technology, make it possible for patients to electronically record and send their data to physicians on a daily or weekly basis. Because data is automatically collected, but not automatically analyzed, this puts a tremendous burden on the physician to review blood glucose records and then make appropriate adjustments in insulin dosages. The third author, for example, has over 700 patients on insulin pump therapy and may make over 20 therapeutic adjustments per day. To reduce this data overload, we began investigating ways to automatically analyze these records and provide appropriate therapeutic recommendations. We envision that such recommendations would initially be provided to physicians for review, but might eventually, once proven safe and effective, be provided directly to patients.

CBR, which has been successfully applied to managing other long-term medical conditions (Bichindaritz 1996; Bichindaritz, Kansu, and Sullivan 1998; Montani et al. 2003b; Marling and Whitehouse 2001), seems especially appropriate for this domain. Type 1 diabetes is marked by a wide variability among patients in terms of sensitivity to insulin, response to environmental and lifestyle factors, propensity for complications, compliance with physician's orders, and response to treatment. Guidelines for managing the disease are well established (American Diabetes Association 2003), but they are general in nature and must be tailored to the needs of each patient. The opportunity for cases to complement guidelines in such situations was presented in (Bichindaritz and Marling 2006). We have built a case-based decision support system prototype to assist patients with Type 1 diabetes on insulin pump therapy. Preliminary results are encouraging, and work continues toward a practical tool for patients.

2. REPRESENTING DIABETES MANAGEMENT KNOWLEDGE IN CASES

2.1. Case Design

A case represents knowledge by storing (a) the description of a specific problem that has occurred; (b) the solution that was applied to that particular problem; and (c) the outcome of applying the solution to that problem (Kolodner 1993). According to (Kolodner 1993), in representing a problem, it is necessary to include all information that is typically used to describe such a problem as well as all information that is explicitly taken into account by a human problem solver in solving the problem. Typical information used to describe diabetes management problems includes blood glucose target levels, actual blood glucose levels, insulin sensitivity, carbohydrate ratios, type of insulin used, basal rates of insulin infusion, bolus doses of insulin with food consumption and/or for correction, type of bolus wave, meal times, and amount of carbohydrate consumed at each meal. Information explicitly taken into account in solving diabetes management problems includes: time of change of insulin infusion set; location of insulin infusion set; mechanical problems with the pump; actions taken to self-correct for hypoglycemia; specific foods consumed at each meal; alcohol consumption; time, type, duration, and intensity of exercise; work cycles; sleep cycles; menstrual cycles; stress and illness.

Solutions to problems usually, but not always, involve changes in insulin dosage. Such changes may be to the amount of basal insulin taken at different times of the day, depending on the amount of physical activity during particular time periods, the amount of bolus insulin used for each meal or correction, or the waveform of a bolus to suit particular foods consumed. Solutions may also involve changes in nutrition, exercise, treatment for hypoglycemia, alcohol consumption, the timing of insulin infusion set changes, the site of insulin infusion set placement, or other lifestyle factors.

The outcome of a proposed solution may be to (a) fix a problem; (b) improve, but not entirely resolve, a problem; or (c) fail to resolve a problem. When a proposed solution fails to fix a problem, we must consider the role of patient compliance. For example, one patient, advised to increase her bolus dosage, refused to do so, afraid of potential hypoglycemia. Increasing the bolus dosage is still an appropriate recommendation, but for this patient, must be followed up with additional education and reassurance. This additional education and reassurance may be viewed as a modification of, or repair to, the original unsuccessful solution. In general, when a solution is unsuccessful, it may be repaired or replaced by an alternate solution.

2.2. Case Acquisition

To acquire cases, we conducted a preliminary clinical study involving 20 patients with Type 1 diabetes on insulin pump therapy. This study was approved by Ohio University's institutional review board (IRB) and required that each participant sign a formal informed consent form. Note that it would not have been possible to retroactively construct cases from existing clinical records. For one thing, much of the data required for decision making is not routinely maintained. For another, clinical visits are scheduled every 3 to 4 months, while blood glucose levels fluctuate continuously. To observe the effects of recommended therapy adjustments requires frequent data updates. In the ideal situation, data would be captured in real time.

Each patient participated in the study for 6 weeks. At the beginning of the study, background data was collected from the patient and entered into an Oracle[1] database. This included personal data, occupational information, pump information, insulin sensitivity, carbohydrate ratios, HbA1c (a measure of long-term blood glucose control), complications of the disease, other chronic diseases, medications, family history of diabetes, and typical daily schedules for work, exercise, meals, and sleep. During the study, each patient wore the Medtronic MiniMed[2] Continuous Glucose Monitoring System (CGMS) three separate times, for 3 days at a time. CGMS provides retrospective blood glucose readings every 5 min during the 3-day sensing period, greatly expanding upon the data available from

[1]Oracle is a leading relational database management system (RDBMS) marketed by the Oracle Corporation (Oracle Corporation 2008).

[2]Medtronic MiniMed is a leading manufacturer of insulin pump and glucose monitoring equipment (Medtronic MiniMed 2008).

routine daily finger sticks, which typically average six per day in insulin pump patients. The CGMS data was directly downloaded into the database.

Every day during the 6-week period, the patient manually entered his or her actual daily data into the database. This included six to ten daily blood glucose readings from finger sticks, bolus dosages and waveforms, basal rates, work schedules, sleep schedules, exercise, meals, infusion set changes, hypoglycemic episodes, menstrual cycles, stress, and illness. In addition, patients were allowed and encouraged to enter information about any miscellaneous events that they felt could be impacting their blood glucose levels. The data entry system developed for the study was Web-based, allowing anytime, anywhere access and an intuitive feel for users of ordinary Web browsers.

Once a week, knowledge engineers met with physicians to review the patient data. The data was displayed in both text and graphical form to facilitate physician review. Physicians identified new problems and recommended therapy adjustments. They also monitored and evaluated the effectiveness of therapy adjustments made earlier in the study. Following the meetings, physicians contacted patients with any questions or recommendations for changes in therapy. Knowledge engineers structured the findings into cases.

2.3. A Sample Case

Figure 1 shows the graphical representation of daily data for a patient. The horizontal axis indicates time, and the primary vertical axis indicates blood glucose level. A secondary vertical axis marks the units of basal insulin infused per hour. Basal rates are indicated by the dark blue line at the bottom of the display. The central blue curve represents CGMS data. The red markers interspersed with the curve indicate blood glucose measurements obtained from finger sticks. At the top of the graph are markers for other types of data collected, as indicated by the key at the bottom of the screen. Clicking on any marker will display additional information. For example, clicking on a Hypo Event marker displays the time of a hypoglycemic episode, the blood glucose level, and the symptoms experienced by the patient. Clicking on a Hypo Action marker displays what the patient did to correct the hypoglycemia. The Hypo Event shown in Figure 1 occurred at 7:50 PM on February 16, 2006. The patient's blood glucose reading was 55, and the patient's symptoms were confusion, dizziness, weakness, and feeling sleepy. The patient treated his hypoglycemia by consuming orange juice, yogurt, and whole wheat sesame snacks. This data can quickly be viewed within the context of the patient's entire day, facilitating an understanding of the patient's status and the identification of problems.

Figure 2 shows a sample case.[3] The context for this case can be seen in Figure 1. Shortly after the patient's hypoglycemic episode, described above, both CGMS and finger stick data show the patient to be hyperglycemic. In describing his self-treatment for hypoglycemia, the patient provided evidence for the likely cause of his ensuing hyperglycemia. The patient ate and drank more than the recommended 15 to 30 g of carbohydrates needed to restore his blood glucose level to normal. This is an important problem to correct to avoid a "roller coaster" pattern of highs and lows. Such a pattern was evident for this patient in the early days of the study.

As can be seen in Figure 2, the physician recommended a change in the treatment of hypoglycemia. The patient was advised to suspend his pump for 15 min, to take another

[3]It should be noted that the format of the sample case shown in Figure 2 is the human readable format used in knowledge engineering sessions with physicians. A machine readable case appears as an object of a hierarchical Java class containing over 140 data fields. Also note that the generalizations and differentiations shown are for use in case retrieval and/or eventual use in case adaptation. Only actual cases are stored in the case base, not generalized prototypes.

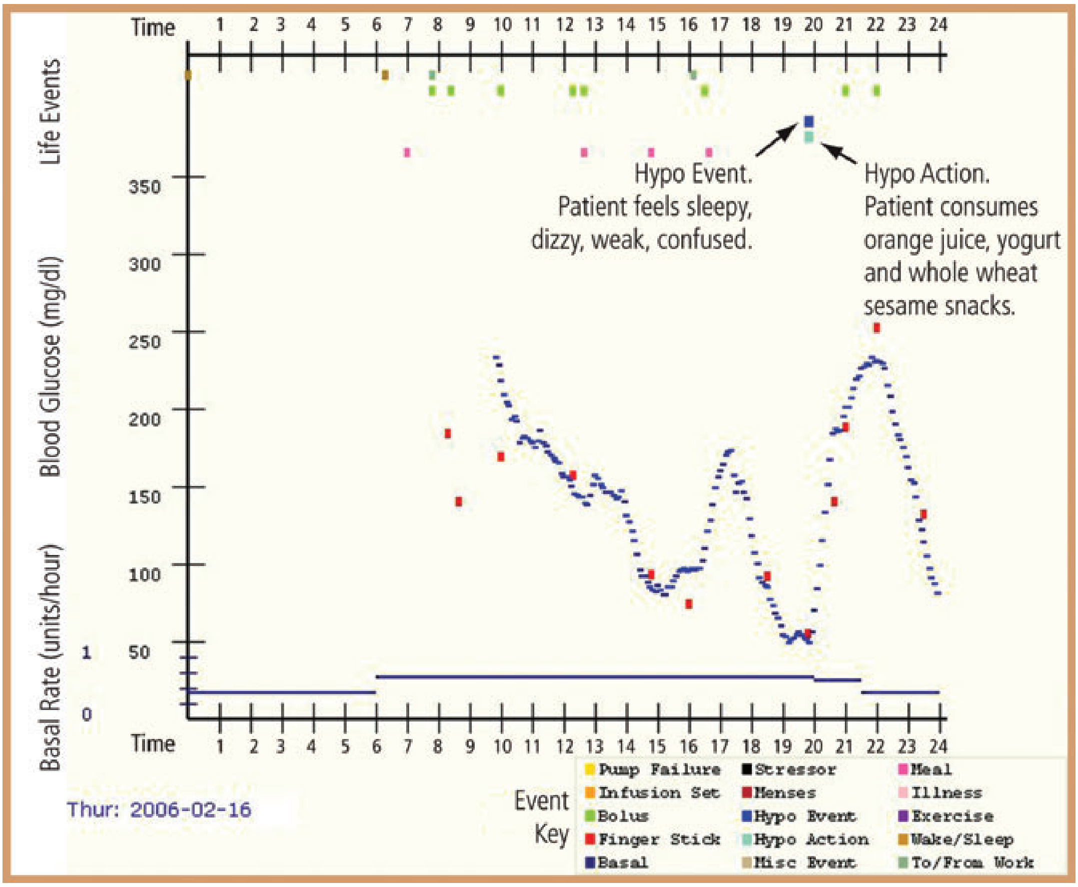

FIGURE 1. Graphical representation of daily data for a patient.

finger stick reading, and to reconnect his pump if his blood glucose level is then within his target range. He was also advised to consume orange juice only, not yogurt and whole wheat sesame snacks in addition. The patient initially complied with this advice, with an unintended ill effect. He suspended his pump and drank his orange juice, but he forgot to reconnect his pump until he became hyperglycemic. He then had to use bolus insulin to correct the hyperglycemia, which was precisely the problem he was trying to avoid. The solution for this case was repaired accordingly, to advise the patient to set an alarm signaling the time to recheck his blood glucose and reconnect his pump. As for the outcome, the patient was no longer willing to risk disconnecting his pump, but he did bring his carbohydrate intake into line. As the study progressed, the data showed he experienced less hyperglycemia following treatment for hypoglycemia.

3. USING CASES FOR INTELLIGENT DECISION SUPPORT

Fifty cases were compiled for the case-based decision support system prototype during the preliminary study, as described earlier. Figure 3 provides a high level overview of how a physician interacts with the system. The details of system operation are as follows.

Problem: Patient Overcorrected for Hypoglycemia
Patient: Patient 1
Physician: Physician 1
Date: February 16, 2006
How Detected: Patient reports Hypo Event at 7:50 PM. Symptoms were confusion, dizziness, weakness, and feeling sleepy. Finger stick was 55. Patient reports intake of orange juice, yogurt, and whole wheat sesame snacks. Patient appears high on CGMS data 2 to 3 hours after Hypo Event.
Generalization: Use finger stick data instead of CGMS.
Generalization: Patient reports intake of more than 30 carbs for treating hypoglycemia.
Differentiation: When patient does not report detecting and correcting a Hypo Event, but pattern of low followed by high occurs, you may have Somogyi effect instead.

Solution: For hypoglycemia, patient should suspend pump for 15 minutes, recheck blood glucose, and reconnect pump if blood glucose is within the target range. Patient should drink the orange juice, but not have all the other food.
Generalization: Suspend pump as above, and set food intake at 30 carbs.
Repair: Remind patient to set an alarm for 15 minutes to signal time to reconnect the pump.
Reason for Repair: Patient forgot to reconnect pump on February 22, 2006, resulting in high blood glucose. Patient suspended at 5:27 PM and reconnected at 6:51 PM with a blood glucose level of 176.

Outcome: Patient complied with most of this. He was not comfortable suspending his pump again, but he did adjust his carbs. This helped. The advice is sound for future patients.

FIGURE 2. A sample case.

3.1. Situation Assessment

The first step in solving a new problem is to assess the current problem situation. In most CBR applications, a problem can be readily described, and the challenge is to find and adapt the most similar and/or useful solution(s). In our domain, the patient and physician are not always aware of the problem specifics, or even that there is a particular problem, thus the first challenge is to identify problems that warrant changes in diabetes management. This involves monitoring the raw patient data to detect patterns that have resulted in therapy adjustments in the past. Because the data continues over time, and there are no predefined

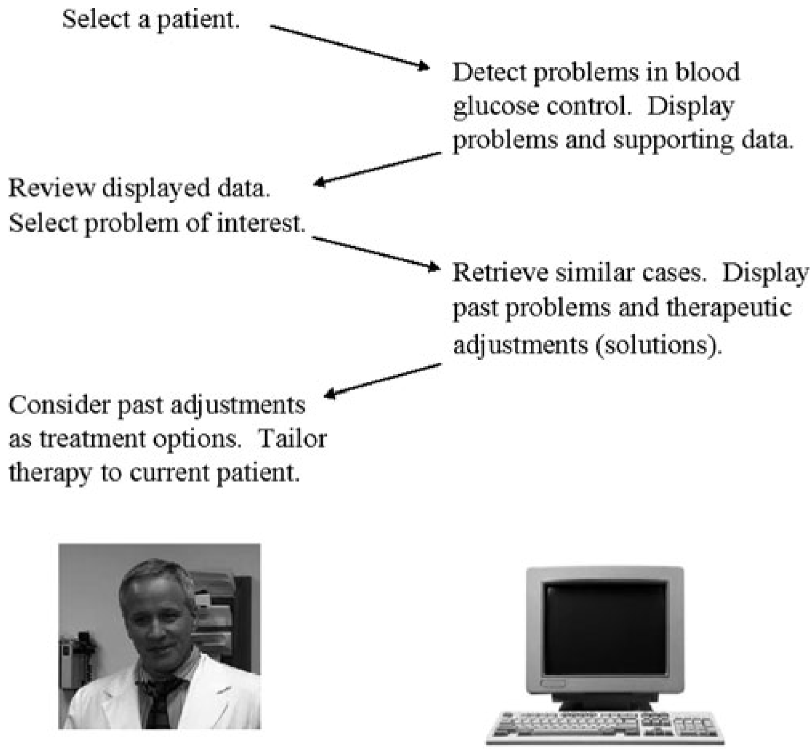

FIGURE 3. High level overview of physician interaction with system prototype.

start or stop points for individual problems, this task bears more resemblance to looking for cases in continuous video than to ordinary CBR situation assessment.

For the system prototype, the user, a physician, first selects the desired patient. The system then searches the glucose and life-event database for problems encountered by this patient and displays the problems detected. The system detects the following common problems in blood glucose control, which were identified during the preliminary study:

1. Hyperglycemia upon wakening;
2. Hypoglycemia upon wakening;
3. Overcorrection for hyperglycemia;
4. Overcorrection for hypoglycemia;
5. Over-boluses for meals;
6. Pre-waking CGMS lows;
7. Lows after exercise;
8. Pre-meal hyperglycemia;
9. Pre-meal hypoglycemia;
10. Post-meal hyperglycemia;
11. Post-meal hypoglycemia;
12. Possible pump or infusion set malfunction.

3.2. Case Retrieval

Next, the physician selects a problem of interest from among those detected for the patient. The selected problem is then compared to the problems stored in the case base. A standard two-step nearest neighbor algorithm is employed, as described in Kolodner (1993). First, a subset of cases most relevant to the current problem is selected by considering the problem type. Because different types of problems have different types of associated treatments, this allows us to focus on only those cases that are possibly relevant in the current situation. For example, if the patient is experiencing hypoglycemia, there is no need to further consider cases involving hyperglycemia.

In the second step, the nearest neighbor metric computes a similarity score for each potentially relevant case, as identified earlier. Cases are compared on 18 different features that were (a) identified by physicians as relevant to treating problems in blood glucose control and (b) empirically determined to help differentiate cases. These features were selected from among the over 140 components comprising each case.

A numeric threshold is used to determine if the highest scoring cases are similar enough to contain useful therapeutic suggestions for the current problem. The system cannot offer decision support when the top scoring case is below the threshold, as this indicates a competence problem to be addressed by growing the case base in the future. The most similar cases above the threshold are displayed with their solutions as decision support to the physician. As similar problems often have similar solutions, therapeutic adjustments made for similar problems in the past may aid the physician in determining appropriate therapy for the current patient (Kolodner 1991). The physician can determine whether or not, and in what form, to pass the retrieved advice on to the patient.

Figure 4 shows the high-level problem description for a sample input case. Below the sample case are the high-level problem descriptions and actual physician recommended solutions for (a) a nonmatching case and (b) the sample case's nearest neighbor. The nonmatching case is ruled out in the first step of the retrieval process, because the input case deals with hypoglycemia while the nonmatching case deals with hyperglycemia. Because the input case and the best matching case both deal with hypoglycemia, these cases move on to the second

Sample Input Case:
Problem: The patient has nocturnal lows.

Nonmatching Case:
Problem: The patient is consistently a little high after breakfast.
Solution: The patient should change the breakfast carb ratio from 1:13 to 1:12.

Best Matching Case:
Problem: The patient had an overnight low.
Solution: The patient should have at least a small bedtime snack, perhaps a glass of milk.

FIGURE 4. A sample input case with a nonmatching case and its closest match. Note that for illustrative purposes, only high-level problem and solution descriptions are shown.

step of the retrieval process. In this step, the input case and the best matching case differ on only two significant features: (a) the input problem occurs on a regular basis while the problem in the best matching case does not; and (b) the input problem was detected by the situation assessment routine for pre-waking CGMS lows, while the problem in the best matching case was not. The two cases share many features in common, including time of day, hypoglycemia correction techniques, and relationship to meals, exercise, boluses, and stress. The actual physician recommended solution for the input case was, "The patient should lower the early morning basal rate from 0.8 to 0.7 units between midnight and 7:00 AM. The patient should always have a bedtime snack." This solution overlaps with the solution given for the best matching case, which was deemed by physicians to be a beneficial treatment for the problem in the input case.

4. EVALUATION AND FEEDBACK

The feasibility of using case-based decision support for patients with Type 1 diabetes on insulin pump therapy was evaluated through a patient exit survey and two structured sessions in which diabetes practitioners evaluated the problem detection and case retrieval capabilities of the system prototype.

4.1. Patient Exit Survey

Twelve of the 20 patients who participated in the preliminary study completed the full 6-week protocol. Patients who did not complete the study cited personal problems and/or lack of time as the primary reasons for discontinuation. An exit survey was given to the 12 patients who completed the initial study concerning time requirements, ease of use, and potential applicability. Patients were not asked about the system prototype per se, as it was built after patient participation in the study. Figure 5 shows patient responses to six multiple-choice survey questions.[4]

Patient responses indicate that a user interface should be designed and implemented that would require less time to use and be more readily available to patients than their laptop or desktop computers. This is consistent with our long-range goal of integrating the decision support software into the patients' own medical devices. The most positive feedback gleaned from the survey is that patients accept the idea of an automated decision support system and expect such a system to be beneficial in managing their diabetes. They are open to the idea of accepting advice from the system, giving the same responses when asked how likely they would be to adopt therapy adjustments recommended by their doctors or by automated systems.

4.2. Situation Assessment

Situation assessment routines were developed during the course of the preliminary study to detect common problems in blood glucose control. Initial feedback was obtained by running the routines on data from patients participating in the latter part of the study and asking physicians to verify outputs. Once all patient data was collected, the routines were run on the entire database of blood glucose and life-event data. The system detected a total of 352

[4]Write-in responses are not shown in Figure 5. For Item 1, two patients wrote in "30 to 60 min," and one patient wrote in "45 min." For Item 3, two patients wrote in "A or D," and one patient wrote in "C or D."

1. Approximately how many minutes each day did it take you to use the Internet-based computer data entry system?							
Choice A	N	Choice B	N	Choice C	N	Choice D	N
60 minutes	2	30 minutes	2	15 minutes	3	less than 15 minutes	2
2. The Internet computer data entry system was:							
Choice A	N	Choice B	N	Choice C	N	Choice D	N
very easy to use	3	fairly easy to use	8	fairly difficult to use	1	very difficult to use	0
3. In what format would you prefer to use a computerized therapy adjustment wizard?							
Choice A	N	Choice B	N	Choice C	N	Choice D	N
a Web browser on a desktop or laptop computer	0	a small handheld computer (Palm) or BlackBerry	3	included on my meter	0	included on my pump	6
4. Do you feel that immediate feedback with advice concerning your blood glucose levels by an automated system would be beneficial in managing your diabetes?							
Choice A	N	Choice B	N	Choice C	N	Choice D	N
very beneficial	7	fairly beneficial	5	not very beneficial	0	confusing & intrusive	0
5. How likely are you to adopt a therapy adjustment recommended by your doctor?							
Choice A	N	Choice B	N	Choice C	N		
very likely	10	somewhat likely	2	not very likely	0		
6. If a computerized therapy adjustment wizard were to recommend a therapy adjustment, how likely would you be to adopt it?							
Choice A	N	Choice B	N	Choice C	N		
very likely	10	somewhat likely	2	not very likely	0		

FIGURE 5. Patient survey responses.

problems over 6 weeks for the 12 patients who completed the study. Knowledge engineers reviewed all 352 problem detections to be sure that all results appeared reasonable from a programming perspective before soliciting feedback from health care professionals.

Ten of the 352 problem detections were randomly selected for evaluation by a panel of three physicians and one advance practice nurse specializing in diabetes. The graphic visualization tool was used to give the evaluators an overview of each problem. Evaluators were then asked to complete a structured review form concerning the correctness and usefulness of each problem detection. The resulting feedback, shown in Figure 6, was encouraging.

1. This is a correct identification of a problem.

Agree:	77.5%
Mixed Feelings:	15.0%
Disagree:	7.5%

2. It would be useful to call this problem to the attention of the patient.

Agree:	87.5%
Mixed Feelings:	10.0%
Disagree:	2.5%

3. It would be useful to call this problem to the attention of the physician.

Agree:	90.0%
Mixed Feelings:	7.5%
Disagree:	2.5%

FIGURE 6. Feedback on automated problem detection.

1. The problem in the original case and the problem in the matching case are:

Similar:	80.0%
Dissimilar:	20.0%

2. Applying the matching case's solution to the original problem would be:

Beneficial:	70.0%
Neither Beneficial nor Detrimental:	23.0%
Detrimental:	7.0%

FIGURE 7. Feedback on case retrieval.

4.3. Case Retrieval

Leave-one-out testing was performed to evaluate the similarity metric and case retrieval capabilities. Leave-one-out testing is a standard technique in which one case at a time is removed from, or “left out” of, the case base. The problem description of the removed case is then used as input to the system, and its nearest neighbor is found among the remaining cases in the case base. During leave-one-out testing, thresholding was turned off, so that the closest match was always returned, even when there was no usefully similar case in the case base. Tests for all 50 cases were initially reviewed by knowledge engineers, and weights were adjusted to refine the performance of the similarity metric. Then, a panel of three physicians specializing in diabetes reviewed a random selection of 10 cases with their nearest neighbors.

For each test case, physicians were given the problem descriptions and recommended solutions of the original case and its nearest neighbor. Then they were asked to evaluate case similarity and the benefit of transfering the solution from the nearest neighbor to the original case. Results are shown in Figure 7. As can be seen from this figure, there is clearly room for improvement. However, it should be noted that some cases collected during the study were not substantially similar to any of the other 49 cases in the case base. It is a central tenet of CBR that the quality of reasoning depends on the experiences stored in the case base (Kolodner 1993), so improvement is anticipated as more cases are added to fill in these competence gaps. Because it took 16 months to collect 50 cases, and additional cases are not readily available, enlarging the case base to improve system performance awaits future work.

5. FUTURE WORK

A second clinical study has been designed and has been approved by the IRB. Twenty-eight patients will participate for 3 months each. In addition to the data collected during the preliminary study, these patients will also provide CGMS data every day for 3 months. The goals of this study are to significantly grow the case base, to incorporate patient similarity as well as problem similarity into the case retrieval process, and to explore the automation, or at least semi automation, of case acquisition and outcome assessment. The next step will be to conduct a multi site evaluation of the tool that results from the second clinical study. The purpose of this study will be to empirically evaluate the effect of case-based therapy recommendations on blood glucose levels and to assess their effect on quality of life for patients. Longer term, we anticipate many opportunities to extend our work to patients with different types of diabetes, patients on different therapy regimens, and patients with special needs, including elite athletes with diabetes, pregnant women with diabetes, and teenagers with diabetes. Additional research and development, including extensive safety testing, may eventually allow patients to directly access the system for continuous monitoring and daily decision making support.

6. RELATED RESEARCH

6.1. CBR in Diabetes Management

The first research project to investigate CBR for diabetes management was the Telematic Management of Insulin-Dependent Diabetes Mellitus (T-IDDM) project (Bellazzi et al. 2002; Montani and Bellazzi 2002; Montani et al. 2003a). The goals of this project were to (a) support physicians in providing appropriate treatment for maintaining blood glucose control; (b) provide remote patients with tele monitoring and tele consultation services; (c) provide cost-effective monitoring of large numbers of patients; (d) support patient education; and (e) allow insulin therapy customization (Bellazzi et al. 2002).

Initially, T-IDDM's decision support module was rule based. The rule-based reasoning (RBR) system analyzed the blood glucose data sent by the patient, identified problems in control, and recommended revisions to insulin therapy. CBR was introduced to specialize the behavior of rules, by tuning rule parameters. The ensuing integration was of the slave-master variety, as defined in (Aha and Daniels 1998), in that RBR remained the primary reasoning modality, while CBR was used for support. In a later effort, a probabilistic model of the effects of insulin on blood glucose over time was added to T-IDDM (Montani et al. 2003a). This model became the primary reasoning modality, with RBR and CBR used when the model could not provide satisfactory results.

Our work, while sharing common goals with T-IDDM, has taken a different approach. Because T-IDDM patients were on conventional insulin therapy, rather than insulin pump therapy, therapy adjustments were limited to the insulin protocol, the amount of insulin taken for regularly scheduled daily injections. The data input to the probabilistic model for a patient was the insulin protocol plus three to four blood glucose measurements per day. The model used by T-IDDM was a steady-state model that did not account for daily variations in diet or lifestyle, but treated them as stochastic occurrences, or noise. This approach makes sense when the data obtained and the therapy adjustment options are limited, as in conventional insulin therapy. We expect to see greater benefits of CBR for patients on insulin pump therapy, who can fine-tune a wider range of insulin and lifestyle parameters to help them live as healthfully and normally as possible. Thus, CBR is our central reasoning modality, and future integrations will be of the master–slave variety, with CBR playing the leading role.

6.2. Other Computer Systems for Diabetes Support

The state of the art in software that is commercially available to people with Type 1 diabetes on insulin pump therapy is exemplified by that provided by the pump and glucometer manufacturer Medtronic MiniMed (Medtronic MiniMed 2005). Patients can download the blood glucose and insulin dosage data stored in their devices to a central site, where they and their physicians can review it. Data may be viewed in log form or in various graphical representations. No attempt is made to automatically interpret the data or to provide therapeutic advice. The software does include a "Bolus Wizard," which uses a numeric formula to recommend individual bolus dosages based on carbohydrate intake, carbohydrate ratio, blood glucose level, blood glucose target range, insulin sensitivity and active insulin.

A major research thrust in building diabetes support systems is telemedicine (Bellazzi et al. 2002; Popow et al. 2003; Gómez et al. 2002; Hernando et al. 2004). Telemedicine aims to enable remote-access health care, reducing face-to-face office visits between patients and physicians, while maintaining, or improving, the quality of care. For example, VIE-DIAB uses mobile phones to transfer data and therapy recommendations between patients and physicians (Popow et al. 2003). DIABTel enables physicians to manage diabetes patients remotely, providing graphical data visualization tools, but no intelligent decision support as of yet (Gómez et al. 2002). A personal Smart Assistant (SA), planned as part of the Intelligent Control Assistant for Diabetes (INCA) project, is the telemedicine system most closely related to our work (Hernando et al. 2004). They collect and consider data specifically applicable to patients with Type 1 diabetes on insulin pump therapy, although not at the level of detail we do. The focus of (Hernando et al. 2004) is on knowledge management: giving patients and physicians the information they need to manually make their own decisions about diabetes management.

6.3. CBR for Managing Other Long-Term Medical Conditions

CBR research and development in the health sciences has been fruitful and extensive. For the past 5 years, workshops have been held to showcase this work at the International and European Conferences on Case-Based Reasoning. Extensive overview articles are available in Bichindaritz and Marling (2006)s, Schmidt et al. (2001), and Nilsson and Sollenborn (2004). The work most closely related to our own involves CBR for managing other long-term, or chronic, medical conditions. Important considerations for such domains include data that varies over time, individual variability among patients, and the need to tailor general guidelines to the requirements of each patient. Related projects include MNAOMIA (Bichindaritz 1996), CARE-PARTNER (Bichindaritz et al. 1998), RHENE (Montani et al. 2003b), and the Auguste Project (Marling and Whitehouse 2001). MNAOMIA is a multi modal reasoning system that operates in the domain of psychiatric eating disorders. CARE-PARTNER provides decision support for the follow-up care of stem cell transplant patients. RHENE provides case-based retrieval for monitoring hemodialysis sessions for end-stage renal disease patients. The Auguste Project provides decision support for planning the ongoing care of Alzheimer's disease patients.

7. SUMMARY AND CONCLUSION

A case-based decision support system prototype has been built to assist patients with Type 1 diabetes on insulin pump therapy. Fifty cases were compiled during a preliminary clinical study to assess the feasibility of constructing such a system. Patients who completed the study provided feedback via an exit survey. The system prototype's situation assessment and case retrieval modules were evaluated by diabetes practitioners. Preliminary results are encouraging and suggest that ongoing research and development could lead to a practical tool for patients.

From a physician's perspective, patients need good blood glucose control to avoid diabetic complications. Systems are commercially available to collect large quantities of blood glucose data, but physicians are left to analyze this data manually. Surveys have shown that some physicians feel overwhelmed by data overload, which can contribute toward a "clinical inertia" in which physicians do not attempt to adjust therapy for patients with diabetes during each office visit (Grant, Buse, and Meigs 2005; Shah et al. 2005). Case-based decision support is an enabling technology that could ease the physician's workload while helping patients achieve and maintain better blood glucose control.

ACKNOWLEDGMENTS

The authors gratefully acknowledge support from Medtronic MiniMed, Ohio University's Russ College Biomedical Engineering Fund, and the Ohio University Osteopathic College of Medicine Research and Scholarly Affairs Committee. They would also like to thank Tessa Cooper, Eric Flowers, Thomas Jones, Tony Maimone, Wes Miller, Stan Vernier, and Don Walker for their software development and knowledge engineering contributions.

REFERENCES

AHA, D., and J. J. DANIELS, *Edited by*. 1998. Case Based Reasoning Integrations: Papers from the 1998 Workshop. Technical Report WS-98-15, AAAI Press, Madison, WI.

AMERICAN DIABETES ASSOCIATION. 2003. American Diabetes Association Complete Guide to Diabetes (3rd ed.). Bantam, New York, NY.

BELLAZZI, R., C. LARIZZA, S. MONTANI, A. RIVA, M. STEFANELLI, G. D'ANNUNZIO, R. LORINI, E. J. GÓMEZ, E. HERNANDO, E. BRUGUÉS, J. CERMENO, R. CORCOY, A. DE LEIVA, C. COBELLI, G. NUCCI, S. D. PRATO, A. MARAN, E. KILKKI, and J. TUOMINEN. 2002. A telemedicine support for diabetes management: The T-IDDM project. Computer Methods and Programs in Biomedicine, **69**:147–161.

BICHINDARITZ, I. 1996. MNAOMIA: Improving case-based reasoning for an application in psychiatry. *In* Artificial Intelligence in Medicine: Applications of Current Technologies. Working Notes of the AAAI-96 Spring Symposium, Stanford, CA.

BICHINDARITZ, I., E. KANSU, and K. M. SULLIVAN. 1998. Case-based reasoning in CAREPARTNER: Gathering evidence for evidence-based medical practice. *In* Proceedings of the 4th European Workshop in Advances in Case-Based Reasoning (EWCBR-98). *Edited by* B. Smyth and P. Cunningham. Springer-Verlag, Berlin, pp. 334–345.

BICHINDARITZ, I., and C. MARLING. 2006. Case-based reasoning in the health sciences: What's next? Artificial Intelligence in Medicine, **36**:127–135.

GÓMEZ, E. J., M. E. HERNANDO, A. GARCÍA, F. DEL POZO, J. CERMENO, R. CORCOY, E. BRUGUÉS, and A. DE LEIVA. 2002. Telemedicine as a tool for intensive management of diabetes: The DIABTel experience. Computer Methods and Programs in Biomedicine, **69**:163–177.

GRANT, R. W., J. B. BUSE, and J. B. MEIGS. 2005. Quality of diabetes care in U.S. academic medical centers: Low rates of medical regimen change. Diabetes Care, **28**:337–442.

HERNANDO, M. E., E. J. GÓMEZ, A. GILI, M. GÓMEZ, G. GARCÍA, and F. DEL POZO. 2004. New trends in diabetes management: Mobile telemedicine closed-loop system. *In* Transformation of Healthcare with Information Technologies. *Edited by* M. Duplaga, K. Zielinski, and D. Ingram. IOS Press, Amsterdam, pp. 70–79.

KOLODNER, J. 1991. Improving human decision making through case-based decision aiding. AI Magazine, **12**:52–68.

KOLODNER, J. 1993. Case-Based Reasoning. Morgan Kaufman, San Mateo, CA.

MARLING, C., and P. WHITEHOUSE. 2001. Case-based reasoning in the care of Alzheimer's disease patients. *In* Proceedings of the Fourth International Conference on Case-Based Reasoning, ICCBR-01. *Edited by* D. W. Aha and I. Watson. Springer, Berlin, pp. 702–715.

MEDTRONIC MINIMED. 2005. Medtronic MiniMed Solutions Pump and Meters Software. Medtronic MiniMed, Northridge, CA.

MEDTRONIC MINIMED. 2008. Medtronic: About Us, http://www.minimed.com/about/index.html (accessed May 2008).

MONTANI, S., and R. BELLAZZI. 2002. Supporting decisions in medical applications: The knowledge management perspective. International Journal of Medical Informatics, **68**:79–90.

MONTANI, S., P. MAGNI, R. BELLAZZI, C. LARIZZA, A. V. ROUDSARI, and E. R. CARSON. 2003a. Integrating model-based decision support in a multi-modal reasoning system for managing type 1 diabetic patients. Artificial Intelligence in Medicine, **29**:131–151.

MONTANI, S., L. PORTINALE, G. LEONARDI, and R. BELLAZZI. 2003b. Applying case-based retrieval to hemodialysis treatment. *In* Workshop Proceedings of the Fifth International Conference on Case-Based Reasoning. *Edited by* L. McGinty. Trondheim, Norway.

NILSSON, M., and M. SOLLENBORN. 2004. Advancements and trends in medical case-based reasoning: An overview of systems and system development. *In* Proceedings of the Seventeenth International Florida Artificial Intelligence Research Society Conference—Special Track on Case-Based Reasoning. AAAI Press, Menlo Park, CA, pp. 178–183.

ORACLE CORPORATION. 2008. Oracle, http://www.oracle.com (accessed May, 2008).

POPOW, C., W. HORN, B. RAMI, and E. SCHOBER. 2003. VIE-DIAB: A support program for telemedical glycaemic control. *In* Proceedings of the 9th Conference on Artificial Intelligence in Europe, AIME 2003. *Edited by* M. Dojat, E. Keravnou, and P. Barahoma. Springer-Verlag, Berlin, pp. 350–354.

SCHMIDT, R., S. MONTANI, R. BELLAZZI, L. PORTINALE, and L. GIERL. 2001. Case-based reasoning for medical knowledge-based systems. International Journal of Medical Informatics, **64**:355–367.

SHAH, B. R., J. E. HUX, A. LAUPACIS, B. ZINMAN, and C. VAN WALRAVEN. 2005. Clinical inertia in response to inadequate glycemic control: Do specialists differ from primary care physicians? Diabetes Care, **28**:600–606.

THE DIABETES CONTROL AND COMPLICATIONS TRIAL RESEARCH GROUP. 1993. The effect of intensive treatment of diabetes on the development and progression of long-term complications in insulin-dependent diabetes mellitus. New England Journal of Medicine, **329**:977–986.