Towards Case-Based Reasoning for Diabetes Management

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, Ohio 45701, USA
marling@ohio.edu
[2] Appalachian Rural Health Institute, Diabetes and Endocrine Center
College of Osteopathic Medicine
Ohio University, Athens, Ohio 45701, USA
shubrook@ohio.edu, schwartf@ohio.edu

Abstract. This paper presents preliminary work in using case-based reasoning (CBR) for diabetes management. The long range goal of the project is to provide intelligent decision support to people with Type 1 diabetes on insulin pump therapy. Case-based reasoning (CBR) was selected for this application 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. Progress to date includes: (a) construction of software tools for data collection and visualization; (b) compilation of a case library; and (c) implementation of code for automated problem detection to identify new problem cases in raw patient data. Work continues on a CBR therapy advisor that will propose individualized solutions for detected problems in blood glucose control.

1 Introduction

People with Type 1 diabetes are unable to produce their own insulin, and so 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 [1]. 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 pre-determined amount of insulin three or four times a day. With 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. Therefore, to provide individualized decision support that can help each patient maintain good blood glucose control, we propose a case-based approach. While we are not the first to employ case-based reasoning (CBR) in the diabetes domain [2], we are the first to use it as the primary reasoning modality for supporting patients with Type 1 diabetes on insulin pump therapy. Our long-term goal is to build a CBR therapy advisor 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 five minutes, 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 [3–6], 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 [7], 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 [8].

2 Using Cases to Represent Diabetes Management Knowledge

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 [9]. According to Kolodner [9], 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 (usually every three days); 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 Acquisistion

To acquire cases, we designed a preliminary study involving 20 people 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 three to four 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 six weeks. At the beginning of the study, background data was collected from the patient and entered into an Oracle 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 Continuous Glucose Monitoring System (CGMS) three separate times, for three days at a time. CGMS provides retrospective blood glucose readings every five minutes during the three day sensing period, greatly expanding upon the data available from 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 six 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 in order 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 a sample case. The patient reported that he had a hypoglycemic episode 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. A short time later, both CGMS and finger stick data showed 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 grams of carbohydrates needed to restore his blood glucose level to normal. This is an important problem to correct in order to avoid a "roller coaster" pattern of highs and lows. Such a pattern was evident in the CGMS and finger stick data for this patient in the early days of the study.

As can be seen in Figure 1, the physician recommended a change in the treatment of hypoglycemia. The patient was advised to suspend his pump for

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 two to three 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.

Fig. 1. A Sample Case

15 minutes to stop the insulin infusion, to take another finger stick reading 15 minutes later, 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 again due to lack of insulin. He then had to use additional 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.

2.4 Case Reuse

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, so 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 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.

While work continues in this area, we have already implemented routines to detect the following conditions: hyperglycemia upon wakening, hypoglycemia upon wakening, over corrections for hyperglycemia, over corrections for hypoglycemia, over-boluses for meals, pre-waking CGMS lows, lows after exercise, pre-meal hyperglycemia, pre-meal hypoglycemia, post-meal hyperglycemia, post-meal hypoglycemia, and pump problems. Newly detected problems can then be compared to problems stored in the care base to find relevant therapy adjustments.

3 Current Status and Future Work

To date, 50 cases have been compiled, covering a broad range of problems experienced by patients with Type 1 diabetes on insulin pump therapy. Twelve types of problems requiring therapy adjustment can be automatically detected in raw patient data. We are currently at work on our advisory system prototype, developing and refining the modules for case retrieval and adaptation.

Once our prototype is complete, our next step will be to conduct another, larger study, 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. Should the early promise be fulfilled, 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, diabetics who become pregnant, and teenaged diabetics.

4 Related Research

4.1 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 [2, 10, 11]. 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 [2].

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. Case-based reasoning was introduced to specialize the behavior of rules, by tuning rule parameters. The ensuing integration was of the slave-master variety, as defined in [12], 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 [11]. 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 as 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.

4.2 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 [13]. 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. Our work may be viewed in the context of extending this functionality to provide an intelligent "Therapy Wizard."

A major research thrust in building diabetes support systems is telemedicine [2, 14–16]. 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 [14]. DIABTel enables physicians to manage diabetes patients remotely, providing graphical data visualization tools, but no intelligent decision support as of yet [15]. 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 [16]. 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 [16] is on knowledge management: giving patients and physicians the information they need to manually make their own decisions about diabetes management.

4.3 CBR for Managing Other Long-Term Medical Conditions

CBR research and development in the health sciences has been fruitful and extensive. For the past four 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 [8, 17, 18]. 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 [3], CARE-PARTNER [4], RHENE [5] and the Auguste Project [6]. 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.

5 Summary and Conclusions

We have presented preliminary work in building a case-based therapy advisor for people with Type 1 diabetes on insulin pump therapy. In this domain, we must deal with: (a) considerable variability among individual patients; (b) general guidelines that must be tailored to the needs of each patient; (c) the combined effects of physical and lifestyle factors on blood glucose control; and (d) the continuous nature of data over time. We have conducted a preliminary study involving 20 patients with Type 1 diabetes on insulin pump therapy. Through this study, we have built a case base of 50 problems in blood glucose control with their associated solutions and outcomes. We can automatically detect 12 types of problems in raw patient data. Work continues on a prototypical therapy advisor, which will provide therapy recommendations for detected problems.

This preliminary work has already yielded the following benefits. First, the data collection and visualization tools have provided a new resource for physicians that enables recognition of abnormal glucose patterns and the life events most likely to have caused them. This facilitates physician recommendations for therapy changes. Current commercial products do not provide the detailed lifestyle data we collect, which can explain excursions in blood glucose control. Second, the automated problem detection routines provide a time-saving tool for physicians burdened by data overload. Third, the cases themselves have proven to be useful educational tools for medical students and for patients. Patients are better able to manage their own diabetes when they have a greater understanding of the factors affecting their blood glucose levels. We believe that CBR is especially well suited to this domain and has enabled the progress to date. We look forward to future system development, evaluation and extension.

6 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. We would also like to thank Tony Maimone and Wes Miller for their software development and knowledge engineering contributions.

References

1. The Diabetes Control and Complications Trial Research Group: 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** (1993) 977–986
2. Bellazzi, R., Larizza, C., Montani, S., Riva, A., Stefanelli, M., d'Annunzio, G., Lorini, R., Gómez, E.J., Hernando, E., Brugués, E., Cermeno, J., Corcoy, R., de Leiva, A., Cobelli, C., Nucci, G., Prato, S.D., Maran, A., Kilkki, E., Tuominen, J.: A telemedicine support for diabetes management: the T-IDDM project. Computer Methods and Programs in Biomedicine **69** (2002) 147–161

3. Bichindaritz, I.: MNAOMIA: Improving case-based reasoning for an application in psychiatry. In: Artificial Intelligence in Medicine: Applications of Current Technologies, Stanford, CA, Working Notes of the AAAI-96 Spring Symposium (1996)
4. Bichindaritz, I., Kansu, E., Sullivan, K.M.: Case-based reasoning in CAREPARTNER: Gathering evidence for evidence-based medical practice. In Smyth, B., Cunningham, P., eds.: Advances in Case-Based Reasoning: 4th European Workshop, Proceedings EWCBR-98, Berlin, Springer-Verlag (1998) 334–345
5. Montani, S., Portinale, L., Leonardi, G., Bellazzi, R.: Applying case-based retrieval to hemodialysis treatment. In McGinty, L., ed.: Workshop Proceedings of the Fifth International Conference on Case-Based Reasoning, Trondheim, Norway (2003)
6. Marling, C., Whitehouse, P.: Case-based reasoning in the care of Alzheimer's disease patients. In Aha, D.W., Watson, I., eds.: Proceedings of the Fourth International Conference on Case-Based Reasoning, ICCBR-01, Berlin, Springer (2001) 702–715
7. American Diabetes Association: American Diabetes Association Complete Guide to Diabetes. 3 edn. Bantam, New York, NY (2003)
8. Bichindaritz, I., Marling, C.: Case-based reasoning in the health sciences: what's next? Artificial Intelligence in Medicine **36** (2006) 127–135
9. Kolodner, J.: Case-Based Reasoning. Morgan Kaufman, San Mateo, CA (1993)
10. Montani, S., Bellazzi, R.: Supporting decisions in medical applications: The knowledge management perspective. International Journal of Medical Informatics **68** (2002) 79–90
11. Montani, S., Magni, P., Bellazzi, R., Larizza, C., Roudsari, A.V., Carson, E.R.: Integrating model-based decision support in a multi-modal reasoning system for managing type 1 diabetic patients. Artificial Intelligence in Medicine **29** (2003) 131–151
12. Aha, D., Daniels, J.J., eds.: Case Based Reasoning Integrations: Papers from the 1998 Workshop, Madison, WI, AAAI Press (1998) Technical Report WS–98–15.
13. Medtronic MiniMed: Medtronic MiniMed Solutions Pump & Meters Software. Medtronic MiniMed, Northridge, CA (2005)
14. Popow, C., Horn, W., Rami, B., Schober, E.: VIE-DIAB: A support program for telemedical glycaemic control. In Dojat, M., Keravnou, E., Barahoma, P., eds.: Proceedings of the 9th Conference on Artificial Intelligence in Europe, AIME 2003, Berlin, Springer-Verlag (2003) 350–354
15. Gómez, E.J., Hernando, M.E., García, A., del Pozo, F., Cermeno, J., Corcoy, R., Brugués, E., de Leiva, A.: Telemedicine as a tool for intensive management of diabetes: the DIABTel experience. Computer Methods and Programs in Biomedicine **69** (2002) 163–177
16. Hernando, M.E., Gómez, E.J., Gili, A., Gómez, M., García, G., del Pozo, F.: New trends in diabetes management: Mobile telemedicine closed-loop system. In Duplaga, M., Zielinski, K., Ingram, D., eds.: Transformation of Healthcare with Information Technologies, Amsterdam, IOS Press (2004)
17. Schmidt, R., Montani, S., Bellazzi, R., Portinale, L., Gierl, L.: Case-based reasoning for medical knowledge-based systems. International Journal of Medical Informatics **64** (2001) 355–367
18. Nilsson, M., Sollenborn, M.: 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, Menlo Park, California, AAAI Press (2004) 178–183