System Overview: The 4 Diabetes Support System

Cindy Marling[1], Razvan Bunescu[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, bunescu@ohio.edu

[2] The Diabetes Institute
Heritage College of Osteopathic Medicine
Ohio University, Athens, Ohio 45701, USA
shubrook@ohio.edu, schwartf@ohio.edu

Abstract. This paper presents an overview of the 4 Diabetes Support System (4DSS). The 4DSS is a hybrid case-based reasoning system that aims to help patients with type 1 diabetes on insulin pump therapy achieve and maintain good blood glucose control. These patients and their professional health care providers have access to a plethora of blood glucose data. However, there is a paucity of automated data analysis to interpret this data and make it actionable. This paper describes how the 4DSS seeks to fill this gap by providing intelligent decision support for diabetes management.

1 Introduction

There are an estimated 346 million people worldwide who have diabetes [13]. Approximately 20 million of them have type 1 diabetes (T1D), the most severe form, in which the pancreas fails to produce insulin. Because insulin is an essential hormone needed to convert food into energy, T1D patients depend upon external supplies of insulin. T1D patients at the Diabetes Institute at Ohio University are treated with insulin pump therapy. An insulin pump continuously infuses the patient with basal insulin. The patient may instruct the pump to deliver additional insulin boluses to account for meals or blood glucose excursions.

While diabetes can not yet be cured, it is actively managed through blood glucose (BG) control. Good BG control is known to help delay or prevent long-term diabetic complications, including blindness, amputations, kidney failure, strokes, and heart attacks [3]. Effective BG control entails vigilant self-monitoring of BG levels. T1D patients prick their fingers from 4 to 6 times a day and use glucometers to measure their blood. They may also wear continuous glucose monitoring (CGM) devices, which produce BG readings every 5 minutes.

BG monitoring data is relayed to physicians, who must manually interpret it to find BG control problems and recommend appropriate therapeutic adjustments. While voluminous BG data contributes to data overload for physicians,

data concerning life events that impact BG levels is not routinely maintained. Physicians may feel, paradoxically, that they have too much data and yet not enough data at the same time.

The 4 Diabetes Support System (4DSS) is a hybrid case-based reasoning (CBR) system that aims to assist T1D patients and their physicians by detecting problems in BG control and suggesting personalized therapeutic adjustments to correct them. It represents work in the CBR in the Health Sciences tradition [1]. While the 4DSS is still a research platform, commercialization efforts are currently underway [6]. The purpose of this paper is to provide a succinct system overview. Additional 4DSS references are available [4, 5, 7, 8, 10–12].

2 Overview

Figure 1 shows a graphical overview of the 4DSS. The system is data driven. Blood glucose data comes from glucometers and CGM devices. Insulin data comes from the patient's pump. The patient uploads BG and insulin data to Medtronic's proprietary CareLink system [9], where it is extracted and transferred to the 4DSS database. The patient manually enters data about life events that impact BG levels, including food, exercise, sleep, work, stress and illness. Originally entered via computer-based browsers, life-event data is now entered via smart phones.

The situation assessment module scans patient data. Traditionally in CBR systems, situation assessment begins with a static description of the current problem. It is different in this domain, because BG control problems continue over time, and because patients are not necessarily aware of problems when they occur. The 4DSS situation assessment module has three major components: problem detection, glycemic variability classification, and BG prediction. These components were built using rule-based reasoning, machine learning algorithms, and time series prediction techniques, giving the 4DSS its hybrid character.

The problem detection component contains 18 rule-based routines that incorporate physician strategies for finding problems in patient data. At a high level, these routines look for problems involving: (1) hyperglycemia, or high BG, which contributes to long-term diabetic complications; (2) hypoglycemia, or low BG, which may result in severe immediate reactions, including weakness, dizziness, seizure or coma; (3) fluctuations between hyper- and hypoglycemia; and (4) lapses in diabetes self-care.

The glycemic variability classication component assesses problems involving BG fluctuation. It detects the problem of excessive glycemic variability, which is believed to presage long-term complications caused by oxidative stress. When expert rules proved inadequate for detecting this problem, machine learning algorithms, including multi-layer perceptrons and support vector machines, were introduced. These algorithms classify 24-hour BG plots by variability level to match physician gestalt perception of such plots. This 4DSS component has stand-alone clinical applicability beyond its role in 4DSS situation assessment.

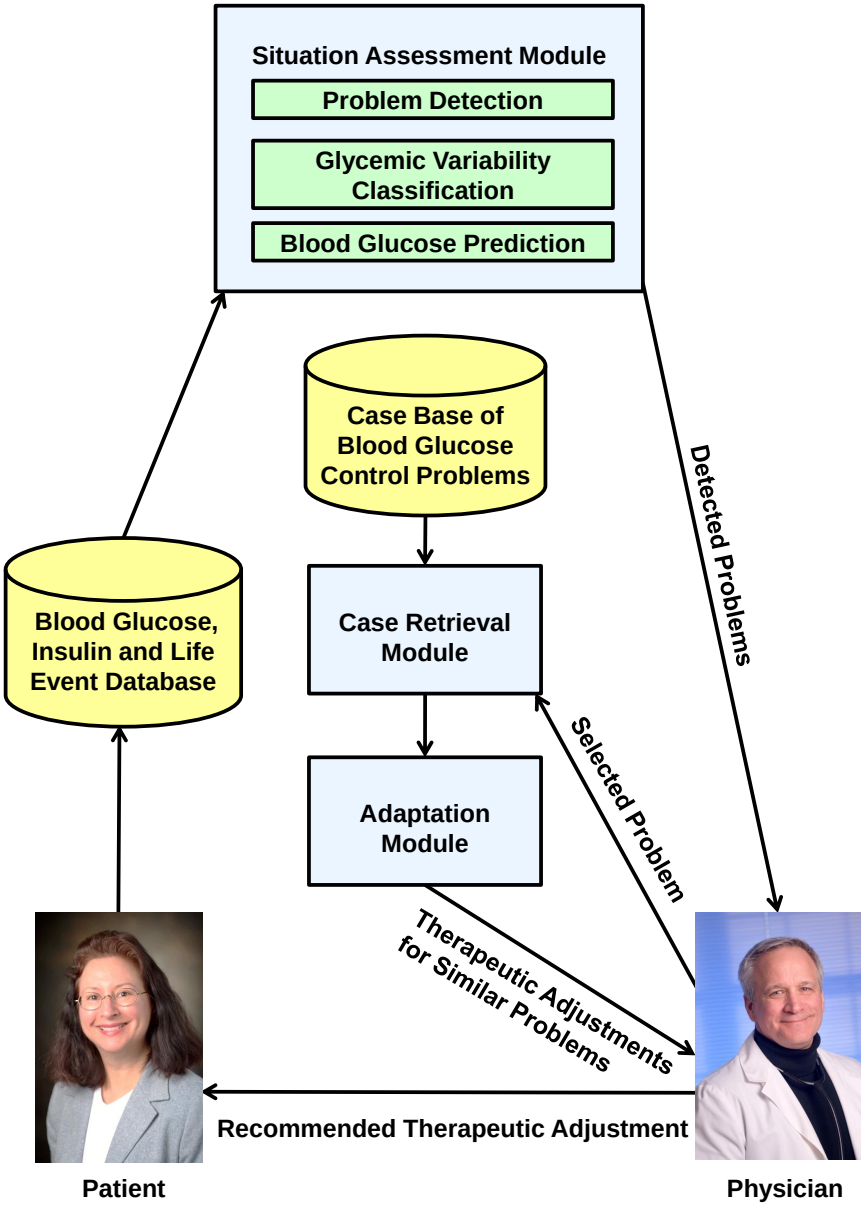

Fig. 1. Overview of the 4 Diabetes Support System

The BG prediction component, currently under construction, incorporates time series prediction techniques to anticipate problems before they occur. While BG data is not currently available in real time and must be scanned retrospectively, we are preparing for its near-term availability. Predicting problems even 30 minutes in advance would give patients time to take preventative actions. This 4DSS component also has a role beyond the 4DSS, enhancing patient safety through eventual incorporation in medical devices.

The situation assessment module reports the problems it finds to a physician, who must then select problems of interest. A selected problem triggers the case retrieval module of the 4DSS. The case retrieval module obtains the most similar cases from the 4DSS case base.

The case base includes 80 cases, each of which contains a specific BG control problem experienced by a T1D patient, a physician's recommended therapeutic adjustment for that problem, and the clinical outcome for the patient after making the therapeutic adjustment. These cases were compiled during clinical research studies in which: (1) T1D patients contributed BG, insulin and life event data; (2) physicians manually identified BG control problems and recommended solutions to patients; (3) patients made the recommended therapeutic adjustments (or not); and (4) physicians examined subsequent patient data to determine the efficacy of the solutions.

To retrieve the most similar cases from the case base, the case retrieval module employs a traditional two-step process. First, a subset of potentially similar cases is identified, and then the nearest neighbors are selected from that subset. In the first step, cases are partitioned by problem type. In the second step, a standard nearest neighbor metric is used. Domain specific similarity functions are combined with empirically determined weights to obtain an overall score for each case. Cases scoring above a similarity threshhold are forwarded to the adaptation module.

The adaptation module personalizes a solution from a retrieved case to fit the situation of a current patient. It begins with the most similar case, but if the solution in that case is not adaptable, it considers the next most similar case, and so on. A solution is a therapeutic adjustment composed of one or more actions that a patient can take. During adaptation, individual actions may be deleted or modified. For example, one possible action is to have a bedtime snack. If the current patient is already having an adequate bedtime snack, this action could be removed from the recommendation. In other situations, the advice could be modified so that the patient eats more or less food before bed, eats a different type of food before bed, or has a snack at a different time of day.

The adapted therapeutic adjustment is relayed to the physician as decision support. The physician decides whether or not to relay the recommendation to the patient. It has long been a goal to provide low-risk advice directly to patients, in real-time, as well as to their physicians. However, this must remain a future goal until the safety and efficacy of the system is proven through clinical trials and approval is obtained from governmental regulatory agencies.

3 Evaluation

Each component of the 4DSS has been evaluated. These evaluations have provided proof of concept, illuminated system strengths and weaknesses, and guided system development. Note that a definitive clinical trial, assessing system impact on patient outcomes, remains to be conducted.

The problem detection component was evaluated after the first and second 4DSS clinical studies. In the first evaluation, a panel of diabetes practitioners rated a sampling of problem detections, and in the second, each patient's physician rated all of the problem detections for the patient. In the first test, 77.5% of problem detections were rated as correct [4], while in the second, 97.9% were rated as correct [11].

The glycemic variability classification component was also evaluated twice. Here, ten-fold cross validation was used to determine the accuracy, sensitivity and specificity of each potential classifier, where correctness is defined as matching physician classifications. In an early test, a naive Bayes classifier matched physicians 85% of the time [8]. The current best classifier, a multi-layer perceptron, has accuracy, sensitivity and specificity of 93.8%, 86.6%, and 96.6%, respectively [12].

The BG prediction component, which is still under construction, is evaluated by examining the difference between predicted BG values and actual BG values. We compute the root mean square error (RMSE), as well as domain specific metrics, for different BG models. Here, evaluation is intertwined with model construction, as the RMSE and other metrics drive model refinement.

The case retrieval module was evaluated by a panel of diabetes practitioners after the first and second clinical studies. Leave-one-out testing was used to provide a sampling of case retrievals for evaluation. In the first test, evaluators rated the retrieved cases as similar to test cases 80% of the time and rated the retrieved solutions as beneficial for test patients 70% of the time [4]. In the second test, they rated retrieved cases as similar 79% of the time and retrieved solutions as beneficial 82% of the time [8].

The adaptation module was evaluated by showing physicians sample problems, with both original and adapted solutions, and eliciting feedback on a questionnaire. Physicians rated the original solutions as being fine without adjustment 47% of the time, needing minor adjustment 40% of the time, and needing major adjustment 13% of the time. They judged the adapted solutions to be better than the original solutions 83% of the time [2].

4 Acknowledgments

This work is supported, in part, with funding from the National Science Foundation under the Smart Health and Wellbeing program (award IIS-1117489). We gratefully acknowledge support from Medtronic, the Ohio University Russ College Biomedical Engineering Fund, the Ohio University Heritage College of Osteopathic Medicine Research and Scholarly Affairs Committee, and the Ohio

University Diabetes Research Initiative. We would also like to thank the participating patients, research nurses, and graduate research assistants for their contributions to this work.

References

1. Bichindaritz, I.: Case-based reasoning in the health sciences: Why it matters for the health sciences and for CBR. In: Althoff, K.D., Bergmann, R., Minor, M., Hanft, A. (eds.) Advances in Case-Based Reasoning: 9th European Conference, ECCBR. pp. 1–17. Springer, Berlin (2008)
2. Cooper, T.: Case Adaptation for an Intelligent Decision Support System for Diabetes Management. Master's thesis, Ohio University, Athens, Ohio (2010)
3. 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(14), 977–986 (1993)
4. Marling, C., Shubrook, J., Schwartz, F.: Case-based decision support for patients with type 1 diabetes on insulin pump therapy. In: Althoff, K.D., Bergmann, R., Minor, M., Hanft, A. (eds.) Advances in Case-Based Reasoning: 9th European Conference, ECCBR. pp. 325–339. Springer, Berlin (2008)
5. Marling, C., Shubrook, J., Schwartz, F.: Toward case-based reasoning for diabetes management: A preliminary clinical study and decision support system prototype. Computational Intelligence 25(3), 165–179 (2009)
6. Marling, C., Wiley, M., Bunescu, R., Shubrook, J., Schwartz, F.: Emerging applications for intelligent diabetes management. AI Magazine 33(2), 67–78 (2012)
7. Marling, C., Wiley, M., Cooper, T., Bunescu, R., Shubrook, J., Schwartz, F.: The 4 Diabetes Support System: A case study in CBR research and development. In: Ram, A., Wiratunga, N. (eds.) Case-Based Reasoning Research and Development: 19th International Conference on Case-Based Reasoning, ICCBR 2011 Proceedings. pp. 137–150. Springer, Berlin (2011)
8. Marling, C.R., Shubrook, J.H., Vernier, S.J., Wiley, M.T., Schwartz, F.L.: Characterizing blood glucose variability using new metrics with continuous glucose monitoring data. Journal of Diabetes Science and Technology 5(4), 871–878 (2011)
9. Medtronic: Carelink personal software (2012), http://www.medtronicdiabetes.net/products/carelinkpersonalsoftware, http://www.medtronicdiabetes.net/products/carelinkpersonalsoftware, accessed May, 2012
10. Schwartz, F.L., Shubrook, J.H., Marling, C.R.: Use of case-based reasoning to enhance intensive management of patients on insulin pump therapy. Journal of Diabetes Science and Technology 2(4), 603–611 (2008)
11. Schwartz, F.L., Vernier, S.J., Shubrook, J.H., Marling, C.R.: Evaluating the automated blood glucose pattern detection and case-retrieval modules of the 4 Diabetes Support System. Journal of Diabetes Science and Technology 4(6), 1563–1569 (2010)
12. Wiley, M., Bunescu, R., Marling, C., Shubrook, J., Schwartz, F.: Automatic detection of excessive glycemic variability for diabetes management. In: Proceedings of the Tenth International Conference on Machine Learning and Applications (ICMLA-11). pp. 148–154. IEEE Press, New York, NY (2011)
13. World Health Organization: Diabetes (2012), http://www.who.int/mediacentre/factsheets/fs312/en/index.html, accessed May, 2012