The 4 Diabetes Support System:
A Case Study in CBR
Research and Development

Cindy Marling[1], Matthew Wiley[1], Tessa Cooper[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, mw350004@ohio.edu,
tc273702@ohio.edu, bunescu@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 the 4 Diabetes Support System[TM] (4DSS) project as a case study in case-based reasoning (CBR) research and development. This project aims to help patients with type 1 diabetes on insulin pump therapy achieve and maintain good blood glucose control. Over the course of seven years and three clinical research studies, a series of defining cases altered the research and development path. Each of these cases suggested a new, unanticipated research direction or clinical application. New AI technologies, including naive Bayes classification and support vector regression, were incorporated. New medical research into glycemic variability and blood glucose prediction was undertaken. The CBR research paradigm has provided a strong framework for medical research as well as for artificial intelligence (AI) research. This new work has the potential to positively impact the health and wellbeing of patients with diabetes. This paper shares the 4DSS project experience.

Keywords: medical decision support, diabetes management, CBR research paradigm

1 Introduction

The 4 Diabetes Support System[TM] (4DSS) project began in 2004 with the goal of producing a case-based decision support system for diabetes management. Seven years and three clinical research studies later, the research and development path has diverged considerably from that originally planned. While the medical goals have not changed, we now envision different AI integrations and new clinical applications, with potentially far greater impact on health and wellbeing. We hypothesize that the credit for discovering these new directions, or

perhaps the blame for straying from the original path, lies squarely with the case-based reasoning (CBR) research paradigm. At each fork in the road was a case suggesting a new direction.

This paper presents the 4DSS project as a case study in CBR research and development. The next section describes the diabetes management domain, within which this work was conducted. Section 3 recaps the first clinical research study, during which the original 4DSS prototype was built. Section 4 presents the case of the bouncing blood glucose, which motivated the integration of naive Bayes classification into the 4DSS situation assessment module. The case of the just-too-late problem detection, presented in Section 5, led to the integration of support vector regression for real-time situation assessment. Section 6 describes the case of the one-fingered typist, which led to new work in user interface design. The remainder of the paper briefly touches upon future plans and related work, followed by the summary and conclusion.

2   The Diabetes Management Domain

The World Health Organization (WHO) estimates that over 220 million people have diabetes worldwide [30]. From 5 to 10% of these people have type 1 diabetes (T1D), the most severe kind. In T1D, the pancreas fails to produce insulin, an essential hormone needed to convert food into energy. Therefore, T1D patients must depend on exogenous supplies of insulin to live. T1D can not, at present, be prevented or cured; however, it can be treated and effectively managed. At the Appalachian Rural Health Institute Diabetes and Endocrine Center, T1D patients are treated with insulin pump therapy. In insulin pump therapy, a patient is continuously infused with a basal amount of insulin via a pump at all times. To account for food intake or other daily activities, the patient may instruct the pump to deliver additional boluses of insulin.

The cornerstone of diabetes management is blood glucose control. It was experimentally determined, in a landmark 1993 study, that good blood glucose control can help delay or prevent serious long-term complications of diabetes, including blindness, amputations, kidney failure, strokes, and heart attacks [11]. Achieving and maintaining good blood glucose control is a difficult task for patients, who must continuously monitor their blood glucose levels and daily activities. It is a difficult task for physicians, who must review copious quantities of blood glucose and life event data, looking for problems and making therapeutic adjustments to correct them.

Providing intelligent decision support to facilitate good blood glucose control is an endeavor in the CBR in the Health Sciences tradition [7]. The general nature of clinical diabetes management guidelines, coupled with a wide variability among individual patients, means that therapy must be customized to the needs of each individual. The strong influence of social context, including qualitative lifestyle preferences and daily life events, on quantitative outcomes means that combinations of multiple diverse features must be taken into account. The nature of chronic disease management precludes the derivation of a single solution;

rather, each patient must be followed and managed over time. Other medical domains with such characteristics have provided fertile ground for CBR research and development [8]. That said, the lessons learned from the 4DSS project do not only apply in health sciences domains. We can draw analogies, for example, to continuously monitoring the "health" of an oil rig to detect problems and making "therapeutic" adjustments to correct detected problems.

3   The First Clinical Research Study

A preliminary clinical study was conducted between 2006 and 2007 to assess the feasibility of creating an intelligent decision support system for patients with T1D on insulin pump therapy. The results of this study have been previously published [15, 16, 24]; a synopsis is presented here. The goal was to design, build and evaluate a case-based decision support system prototype. While plentiful blood glucose data was initially available, usable cases were not. This is because the life events impacting blood glucose levels are not routinely maintained. To acquire cases for 4DSS, 20 patients were enrolled in a study, and 12 of them completed a 6-week protocol. Each patient entered daily blood glucose, insulin, and life-event data into a database via a Web browser. Each patient used a continuous glucose monitoring (CGM) system for three 3-day periods to capture supplemental blood glucose readings at 5-minute intervals.

The research team met weekly to review patient data. Physicians identified problems in blood glucose control and recommended therapy adjustments for each patient. As patients made therapy adjustments, physicians monitored to evaluate the effectiveness of the changes. These problems, solutions and outcomes were structured into 50 cases for the 4DSS system prototype. One of these cases is shown, in simplified form, in Figure 1. This figure shows a case of nocturnal hypoglycemia. Hypoglycemia, or low blood glucose, may cause weakness, dizziness, confusion, sweating, and if not promptly treated, seizures, coma, or death. Hypoglycemia that occurs while the patient is asleep, as in the sample case, is especially dangerous. Hyperglycemia, or high blood glucose, contributes to the long-term complications of diabetes. Extremely high blood gluocose levels may trigger diabetic ketoacidosis, a serious condition causing severe acute illness or death.

A significant research challenge was detecting patient problems so that useful past solutions could be retrieved and reused. Patients do not always know when problems are impending and are frequently unaware of them even after they occur. Twelve commonly occurring problems were identified by physicians, and rule-based situation assessment routines were developed to automatically recognize these problems in patient data. A 4DSS prototype was then built to: (1) detect problems in blood glucose control; (2) display detected problems to the physician, who would select the problems of interest; (3) retrieve, for each selected problem, the most similar case in the case base; and (4) display the retrieved cases as decision support for determining appropriate therapeutic adjustments.

Fig. 1. A sample case from 4DSS

Evaluation and feedback were obtained through a patient exit survey and two structured sessions in which diabetes practitioners evaluated the problem detection and case retrieval capabilities of the system. Patients indicated that they would willingly accept automated decision support, but noted that the time required for data entry was a deterrent. Physicians noted that the integration of blood glucose, insulin and life-event data helped them to identify glucose trends more readily and adjust therapy more effectively. Conclusions from this study were: (1) the 4DSS system prototype provides proof of concept that intelligent decision support can assist in diabetes management; (2) additional problem/solution/outcome cases are needed to provide solutions for more blood glucose control problems; and (3) data entry time demands on the patient must be reduced.

4 The Case of the Bouncing Blood Glucose

A second clinical research study was conducted between 2008 and 2009 to enhance and evaluate the 4DSS prototype. The primary goal of this study was to streamline the user interface to reduce time demands on patients and then re-evaluate system performance. A secondary goal was to expand system functionality by implementing additional problem detection routines. Twenty-six patients with T1D on insulin pump therapy enrolled, and 23 patients completed the 5-week protocol.

Not long into this study, we encountered the case of the bouncing blood glucose. To date, we could automatically detect 12 types of blood glucose control problems, all involving hypoglycemia or hyperglycemia. Now we had a case where a patient bounced back and forth between hypo and hyperglycemia, in a roller coaster pattern, as shown in Figure 2. We thought that it would be straightforward to implement a new rule-based situation assessment routine to detect this problem. However, several attempts met with failure.


Fig. 2. Bouncing blood glucose

We had encountered *glycemic variability*, a current and controversial topic of diabetes research [9, 13, 19]. Excessive glycemic variability has been linked to hypoglycemia unawareness, an acutely dangerous condition, and to oxidative stress, which contributes to long-term diabetic complications. Its automated detection would be a major contribution to clinical diabetes management. Researchers have proposed numerous metrics for characterizing glycemic variability, but none are widely used in clinical practice [21, 22]. This was a research challenge, and we were off on a detour.

While experts disagree on how to measure glycemic variability, physicians readily recognize it when they see it in blood glucose plots. We therefore considered the quantifiable aspects of glycemic variability as they relate to physicians' perception. We began with the best accepted glycemic variability metric, the Mean Amplitude of Glycemic Excursion (MAGE) [26]. MAGE captures the distance between the local maxima and minima of a blood glucose plot. We then devised two new metrics, distance traveled and excursion frequency, to capture aspects of glycemic variability not accounted for by MAGE. Distance traveled captures overall daily fluctuation, and excursion frequency counts the number of significant glucose excursions in a day.

Two physicians (JS and FS) classified blood glucose plots as excessively variable or not, based on their gestalt perceptions. The same plots were scored for

MAGE, distance traveled and excursion frequency. The physician classifications and metric scores for 218 blood glucose plots were used to train multiple machine learning algorithms, via the Weka machine learning toolkit [29]. During testing, a naive Bayes classifier was able to match physician ratings 85% of the time. The preliminary results are in press [17].

In ongoing work, we are exploring other machine learning algorithms, additional glycemic variability features, and data smoothing techniques. While work continues to improve performance, machine learning classification is now an integral part of 4DSS. Within the original 4DSS framework, glycemic variability classification extends the situation assessment module. When excessive glycemic variability is detected as a problem, case retrieval can be invoked to find an applicable solution. Due to the clinical importance of glycemic variability, the new classifier also has potential as a standalone clinical assessment tool.

5 The Case of the Just-Too-Late Problem Detection

Another patient in the second clinical research study presented the next pivotal case. This patient's pump failed and stopped delivering insulin. He was aware that his blood glucose was high, and he instructed the pump to deliver more insulin. However, he did not know that the pump was not functioning, and his blood glucose continued to climb. He went into diabetic ketoacidosis (DKA) and was admitted to the hospital, where he experienced a (non-fatal) heart attack. When his data was scanned retroactively, the system automatically detected the pump problem eight hours before he was admitted to the hospital. Had the system been running in real time, the patient might have been alerted to change his infusion set, a simple adjustment that could have prevented the DKA. This case highlighted the need to predict and prevent problems instead of just detecting and correcting them. Off we went in another new research direction.

Real-time blood glucose data is not currently available from CGM systems. This is not a technical limitation; rather, the U.S. Food and Drug Administration (FDA) has not yet approved its use. There is, however, an active diabetes research program, the Artificial Pancreas project [12], which is currently seeking FDA approval for real-time problem prediction and intervention. The immediate goal (among many long-range plans) is to predict when blood glucose will drop to hypoglycemic levels and alert the patient in time to take preventive action. Blood glucose prediction is typically approached from a control theory perspective or a physiological pharmacokinetic modeling perspective. We believe that our CBR perspective gives us an advantage, in that the contextual features impacting blood glucose levels have so far been largely ignored.

Given the natural temporal ordering of blood glucose measurements, we approach blood glucose prediction as a time series forecasting problem. Here, the task is to estimate the future value of a target function based on current and past data samples. We chose Support Vector Regression (SVR) models [28] for this task, as they can easily incorporate contextual features, without any as-

30 Minute Predictions							
Method	E_{RMS}	R^2	A	B	C	D	E
SVR	**18.0**	**0.92**	**93.0**	7.0	0.0	0.0	0.0
Baseline	25.1	0.84	87.8	11.8	0.0	0.4	0.0

60 Minute Predictions							
Method	E_{RMS}	R^2	A	B	C	D	E
SVR	**30.9**	**0.76**	**81.0**	18.1	0.4	0.5	0.0
Baseline	43.2	0.52	74.5	21.5	2.2	1.8	0.0

Table 1. SVR and baseline results

sumptions of feature independence. Furthermore, SVR models have been used successfully in numerous time series prediction problems [23].

In a preliminary experimental evaluation, an SVR model was trained to predict the blood glucose level of a T1D patient. We had three months of data available, including blood glucose measurements recorded at 5-minute intervals, insulin dosages, and contextual life events. To create training and testing examples, a date was selected approximately one month into the 3-month data set. Data for seven days before that date was used as training data, while data on that date and the two subsequent days was used as test data. Testing and training examples were represented as feature vectors including: blood glucose level; the exponentially smoothed rate of change in blood glucose level; insulin dosage; carbohydrate intake; and exercise time and intensity. Two separate SVR models were trained and tested to predict blood glucose levels 30 and 60 minutes into the future. The SVR models were trained with a linear kernel; the LibSVM [10] implementation of support vector machines for regression was used.

In Table 1, we compare the performance of the two SVR models with a baseline that uses the present blood glucose level as the prediction for any future blood glucose value. This baseline, while simple, was used because it outperformed more complex moving average and rate of change baselines. We report the root mean square error E_{RMS}, the coefficient of determination R^2, and the percentage of predictions falling in the 5 areas from A to E in the Clarke Error Grid Analysis (CEGA) [14]. CEGA is a domain specific standard for assessing the accuracy of blood glucose measurements or predictions. It was originally designed to assess the quality of blood glucose sensors. Area A is the desired region: it corresponds to measured/predicted values that are within 20% of actual/observed values. Areas B through E were created in recognition that it is clinically more important for measurements or predictions to be accurate in some regions than in others. For example, if a patient's blood glucose level is predicted to be 400, but turns out to be 360, that would be an acceptable error. Both levels are very high, and the same medical treatment would apply. However, if the patient's blood glucose level is predicted to be 100, but turns out to be 60, that would be a serious error, because a hypoglycemic episode requiring immediate treatment would be masked.
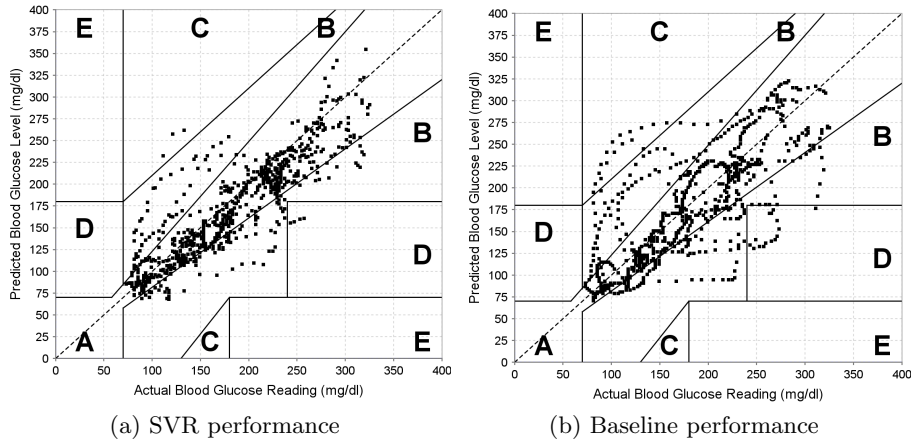
(a) SVR performance   (b) Baseline performance

Fig. 3. CEGA plots for 60 minute prediction of blood glucose levels

The SVR models are promising, as they outperform the baseline on all performance measures. The two CEGA plots in Figure 3 show the performance of SVR and the baseline, respectively, for the 60 minute prediction. Overall, the plots show the learned SVR model making predictions that are closer to the ideal diagonal line.

Our research in blood glucose prediction is currently focused on making the learned models more robust in the presence of measurement noise and data anomalies. We have created a modified framework for regularized spline smoothing in which blood glucose measurements obtained through traditional self-glucose monitoring (SGM) override the less accurate CGM measurements. This approach is based on the way physicians read blood glucose plots containing both CGM and SGM data. By replacing the raw blood glucose values used during training with smoothed values, we expect to mitigate the effects of noise on the final prediction performance.

When integrated into the 4DSS situation assessment module, blood glucose prediction will be an important step toward real-time case-based decision support. A preemptive solution could be retrieved from the case base as soon as a problem is predicted, 30 to 60 minutes before it might otherwise occur. Again, due to the clinical value of blood glucose prediction, this functionality has potential uses beyond the 4DSS. For example, prediction could be incorporated directly in insulin pumps, as a safety feature, to trigger alarms and warnings. Prediction would also be useful in a standalone educational tool with what-if analysis, to help patients understand the impact of their actions on blood glucose control.

6 The Case of the One-Fingered Typist

The second 4DSS clinical research study was completed in 2009. Results were not all positive [25]. Streamlining data entry entailed the following changes: (1) patients no longer interactively entered data; (2) blood glucose and insulin data stored in the pump was automatically uploaded to the database; and (3) daily life events were approximated by typical daily schedules. Time requirements were reduced for patients, but only half as many problems were detected as when the old data entry system was used. This finding was statistically significant (p=0.017), although there were no statistically significant differences between the patient populations and no reason to suspect that patients were actually experiencing fewer problems.

A third clinical research study was conducted between 2009 and 2011. The primary aims of this study were to: (1) enlarge the case base; (2) develop additional problem detection routines; and (3) add an adaptation module to tailor past solutions to the specific needs of current patients. To balance the need to collect all relevant data with the need to minimize patient time demands, it was necessary to build yet another data entry interface before enrolling patients. Seventeen patients enrolled, and twelve patients completed a 3-month protocol. For this study, patients: (1) collected CGM data for the full three months; (2) uploaded insulin pump and CGM data weekly; and (3) supplied daily life-event data that would otherwise be unavailable via a Web browser.

During this study, a patient enrolled who typed with one finger. This patient conscientiously tried to enter all required data, but the process was painful for him, and the data captured was error prone and incomplete. We had long anticipated that, when the system moved from the research laboratory to clinical practice, the user interface would be embedded in medical devices, electronic health records, and/or cell phones. We had not intended to do any research or development in user interface design. Yet, here was a case with a problem in need of a solution.

We are currently in the process of designing and implementing a smart phone data capture system for 4DSS. This system will be used and evaluated during the next clinical research study. The system is being built for a mobile browser, to accommodate any brand of smart phone. The over-arching goal is to create an interface that requires only one finger to enter data. Smart phone touch screens lend themselves to this goal. The evolving interface contains mostly buttons, check boxes, and drop down lists, all of which can be selected by the touch of a finger. The home screen for the new smart phone interface is shown in Figure 4.

We envision other potential advantages to this mode of data capture. Patients will be able to enter data as events occur, rather than waiting until the end of the day to supply a whole day's data. For example, a patient can touch a "Going to Work" button as he or she is actually going to work. Because cell phones can automatically obtain the time, the patient will no longer need to manually enter it, improving data accuracy. Furthermore, using the smart phone platform, physicians will be able to text therapy recommendations to patients, rather than relying on email. This could potentially reduce the time it takes for patients to
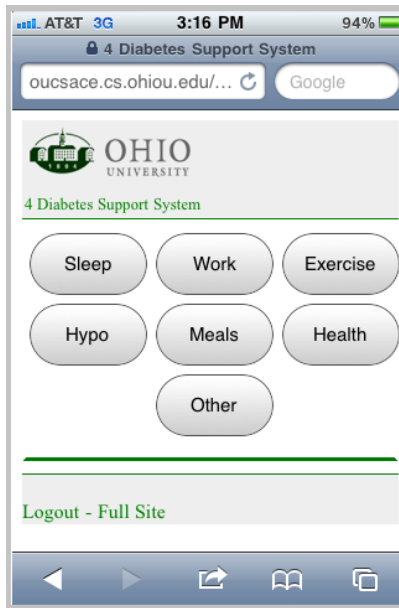
Fig. 4. New 4DSS smart phone data entry interface

implement physician recommendations. It is interesting to note that the first reported use of CBR for diabetes management was in support of a telemedicine application [5]. Now, telemedicine technology may support case-based diabetes management.

7 Case-Based Decision Support Redux

Patients have concluded their participation in the third 4DSS clinical research study, but associated system development activities are ongoing. To date: (1) 30 new cases have been added to the case base; (2) six new problem detection routines have been developed; (3) an adaptation module was built; and (4) a new backend interface was built for physicians to review cases and to view system recommendations. A fourth 4DSS clinical research study is now being planned; the protocol is under development. This study will support the original project goals for case-based diabetes management, the more recent goals of detecting excessive blood glucose variability and predicting blood glucose levels, and the as-yet-to-be-determined research goals that will undoubtedly arise along the way. Parallel efforts are underway to facilitate transfer of 4DSS techology from the research laboratory to clinical use. The new smart phone interface may be viewed as part of this effort. New techology transfer collaborators are working on market assessment and product placement. This will be another chapter for the evolving 4DSS project.

8   Related Work

Twenty years ago, Stephen Slade wrote "Case-Based Reasoning: A Research Paradigm" [27]. He described CBR as a paradigm for reasoning from experience that addresses two core AI research agendas: understanding human thought, and building intelligent systems. In short, he proposed CBR as an ideal paradigm for conducting AI research. Today, it is clear that CBR is a useful paradigm for conducting diabetes research, as well. The intuition for CBR as a medical research paradigm is deeply rooted in medical history. Over 100 years ago, when Dr. Alois Alzheimer first encountered a puzzling new brain disease, he wrote that, when confronted with unknown disease patterns, "a further histological examination must be effected to determine the characteristics of each single case" [18]. While synergies between CBR research and medical research are well documented [8], the view of a CBR framework for medical research is atypical. A notable exception was the early MNAOMIA project [6], which had medical research in the field of psychiatric eating disorders as a full-fledged project goal.

A current research and development effort with many parallels to the 4DSS project is the Mälardalen stress project [3]. This project, which began in 2002 as a collaboration between CBR researchers and leading clinicians in stress diagnosis, has evolved over time to meet domain specific challenges. The initial research direction was case-based stress diagnosis, with physiological features used as defining case parameters. Discrete wavelet transformation was used to automatically extract features from the raw cardio and pulmonary signals used by clinicians in manual diagnosis [20]. The project has since grown and expanded in many directions. When it became difficult to acquire enough cases, fuzzy rules were used to generate artificial cases for the case base, which improved diagnostic performance [1]. When imprecise sensor measurements, noise, and human error impeded the performance of the similarity metric, fuzzy similarity matching was introduced to improve robustness [4]. When it became desirable to consider the social context surrounding a patient's stress, textual features were introduced. Natural language processing was therefore integrated, using WordNet and domain specific ontological knowledge [2]. The current multi-modal system incorporates CBR, rule-based reasoning, fuzzy logic and textual information retrieval to aid in the diagnosis and treatment of stress [3].

9   Summary and Conclusion

The 4 Diabetes Support System^{TM} project stands as a case study in CBR research and development. It illustrates the power of the CBR research paradigm as a framework for medical research. 4DSS was originally envisioned as a CBR system to help patients with type 1 diabetes on insulin pump therapy achieve and maintain good blood glucose control. Over the course of seven years and three clinical research studies, a series of defining cases altered the original research and development path. After the case of the bouncing blood glucose, the medical research goals expanded to include glycemic variability measurement and automated detection of excessive glycemic variability. Meanwhile, the AI approach

expanded to incorporate machine learning classification. After the case of the just-too-late problem detection, 4DSS goals were extended to real-time problem prediction and prevention. Support vector regression models for blood glucose prediction were explored and integrated. After the case of the one-fingered typist, development efforts extended to smart phone interfaces for automated data capture. A concentrated focus on the salient characteristics of each individual case has led to the discovery of new research and development challenges. Meeting these challenges could positively impact the health and wellbeing of patients with diabetes.

10   Acknowledgments

References

1. Ahmed, M.U., Begum, S., Funk, P., Xiong, N.: Fuzzy rule-based classification to build initial case library for case-based stress diagnosis. In: Hamza, M.H. (ed.) 9th IASTED International Conference on Artificial Intelligence and Applications (AIA). pp. 225–230 (2009)
2. Ahmed, M.U., Begum, S., Funk, P., Xiong, N., von Schéele, B.: Case-based reasoning for diagnosis of stress using enhanced cosine and fuzzy similarity. Transactions on Case-Based Reasoning on Multimedia Data 1(1), 3–19 (2008)
3. Ahmed, M.U., Begum, S., Funk, P., Xiong, N., von Schéele, B.: A multi-module case based biofeedback system for stress treatment. Artificial Intelligence in Medicine 51(2), 107–115 (2011)
4. Begum, S., Ahmed, M.U., Funk, P., Xiong, N., von Schéele, B.: A case-based decision support system for individual stress diagnosis using fuzzy similarity matching. Computational Intelligence 25(3), 180–195 (2009)
5. Bellazi, R., Montani, S., Portinale, L., Riva, A.: Integrating rule-based and case-based decision making in diabetic patient management. In: Althoff, K.D., Bergmann, R., Branting, L.K. (eds.) Case-Based Reasoning Research and Development: Third International Conference on Case-Based Reasoning, ICCBR-99. pp. 386–400. Springer, Berlin (1999)
6. Bichindaritz, I.: Case-based reasoning adaptive to several cognitive tasks. In: Veloso, M., Aamodt, A. (eds.) Case-Based Reasoning Research and Development: First International Conference, Proceedings ICCBR-95. pp. 391–400. Springer, Berlin (1995)
7. 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)

8. Bichindaritz, I., Marling, C.: Chapter 7: Case-based reasoning in the health sciences: Foundations and research directions. In: Bichindaritz, I., Jain, L.C., Vaidya, S., Jain, A. (eds.) Computational Intelligence in Healthcare 4: Advanced Methodologies. Springer, Berlin (2010)

9. Ceriello, A., Ihnat, M.A.: Glycaemic variability: A new therapeutic challenge in diabetes and the critical care setting. Diabetic Medicine 27(8), 862–867 (2010)

10. Chang, C.C., Lin, C.J.: LIBSVM: A library for support vector machines (2001), software available at http://www.csie.ntu.edu.tw/~cjlin/libsvm, accessed April, 2011

11. 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)

12. Juvenile Diabetes Research Foundation: Artificial pancreas project (2011), http://www.artificialpancreasproject.com/, accessed April, 2011

13. Kilpatrick, E.S., Rigby, A.S., Atkins, S.L.: For debate. Glucose variability and diabetes complication risk: We need to know the answer. Diabetic Medicine 27(8), 868–871 (2010)

14. Kovatchev, B.P., Gonder-Frederick, L.A., Cox, D.J., Clarke, W.L.: Evaluating the accuracy of continuous glucose-monitoring sensors: Continuous glucose-error grid analysis illustrated by TheraSense Freestyle Navigator data. Diabetes Care 27(8), 1922–1928 (2004)

15. 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)

16. 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)

17. 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 (2011), in press

18. Maurer, K., Volk, S., Gerbaldo, H.: Chapter 1: Auguste D. The history of Alois Alzheimer's first case. In: Whitehouse, P.J., Maurer, K., Ballenger, J.F. (eds.) Concepts of Alzheimer Disease: Biological, Clinical and Cultural Perspectives. Johns Hopkins University Press, Baltimore, MD (2000)

19. Monnier, L., Colette, C.: Glycemic variability. Diabetes Care 31(Supplement 2), S150 (2008)

20. Nilsson, M., Funk, P., Olsson, E.M.G., von Schéele, B., Xiong, N.: Clinical decision support for diagnosing stress-related disorders by applying psychophysiological medical knowledge to an instance-based learning system. Journal of Artificial Intelligence in Medicine 36(2), 156–176 (2005)

21. Rodbard, D.: Interpretation of continuous glucose monitoring data: glycemic variability and quality of glycemic control. Diabetes Technology & Therapeutics 11(s1), S–55–S–67 (2009)

22. Rodbard, D.: New and improved methods to characterize glycemic variability using continuous glucose monitoring. Diabetes Technology & Therapeutics 11(9), 551–565 (2009)

23. Sapankevych, N.I., Sankar, R.: Time series prediction using support vector machines: A survey. IEEE Computational Intelligence Magazine 4(2), 24–38 (2009)

24. 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)
25. 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)
26. Service, F., Molnar, G., Rosevear, J., Ackerman, E., Gatewood, L., Taylor, W.: Mean amplitude of glycemic excursions, a measure of diabetic instability. Diabetes 19(9), 644–655 (1970)
27. Slade, S.: Case-based reasoning: A research paradigm. AI Magazine 12(1), 42–55 (1991)
28. Smola, A.J., Scholkopf, B.: A tutorial on support vector regression. Tech. Rep. TR-98-030, NeuroCOLT2 Technical Report Series (1998)
29. Witten, I.H., Frank, E., Trigg, L., Hall, M., Holmes, G., Cunningham, S.J.: Weka: practical machine learning tools and techniques with Java implementations. In: Proceedings ICONIP/ANZIIS/ANNES99: Future Directions for Intelligent Systems and Information Sciences. pp. 192–196 (1999)
30. World Health Organization: Diabetes (2011), http://www.who.int/mediacentre/factsheets/fs312/en/index.html, accessed April, 2011