

Emerging Applications for Intelligent Diabetes Management

**Cindy Marling, Matthew Wiley
and Razvan Bunescu**

School of Electrical Engineering and
Computer Science, Ohio University
Athens, Ohio 45701, USA

Jay Shubrook and Frank Schwartz

Appalachian Rural Health Institute
Diabetes and Endocrine Center
Ohio University College of Osteopathic Medicine
Athens, Ohio 45701, USA

Abstract

Diabetes management is a difficult task for patients, who must monitor and control their blood glucose levels in order to avoid serious diabetic complications. It is a difficult task for physicians, who must manually interpret large volumes of blood glucose data to tailor therapy to the needs of each patient. This paper describes three emerging applications that employ AI to ease this task and shares difficulties encountered in transitioning AI technology from university researchers to patients and physicians.

Keywords: case-based reasoning, machine learning, diabetes management

Diabetes Management: A Difficult Challenge

Type 1 diabetes (T1D) is an autoimmune disease in which the pancreas fails to produce insulin, an essential hormone needed to convert food into energy. It is a chronic disease, which cannot be cured, but which must be treated and managed over time. T1D patients at the Appalachian Rural Health Institute Diabetes and Endocrine Center are treated with insulin pump therapy. A mechanical pump infuses the patient with insulin, attempting to mimic and replace normal pancreatic function. The management goal is for the person with diabetes to achieve and maintain blood glucose (BG) levels close to those of a person without diabetes. It has been experimentally determined that good BG control can help delay or prevent serious long-term diabetic complications, including blindness, amputations, kidney failure, strokes, and heart attacks (Diabetes Control and Complications Trial Research Group 1993). Avoiding complications improves quality of life for patients, while reducing the financial burden of health care cost expenditures.

Diabetes management is a challenging task for patients, who must monitor their BG levels and daily activities, and for physicians, who recommend therapeutic adjustments based on the monitoring data. Task complexity stems from: (a) a wide variability among individual patients in terms of sensitivity to insulin, response to lifestyle factors, propen-

sity for complications, adherence to physician recommendations, and response to treatment; and (b) voluminous BG data, which is automatically collected through sensors, but which must be manually analyzed and interpreted. Since 2004, we have been conducting clinical research studies with T1D patients in order to develop and evaluate software tools for intelligent diabetes management.

AI Solutions to Diabetes Management Problems: Work in Progress

In the AI in Medicine tradition, real-world medical problems have provided fertile ground for AI research, driving research directions in search of practical solutions. We began with case-based reasoning (CBR) and then synergistically incorporated machine learning (ML) approaches.

Case-Based Decision Support

The 4 Diabetes Support SystemTM (4DSS) aims to: (a) automatically detect problems in BG control; (b) propose solutions to detected problems; and (c) remember which solutions are effective or ineffective for individual patients. It can assist busy clinicians managing multiple T1D patients, and it might eventually be embedded in insulin pumps or smart phones to provide low-risk advice to patients in real time. CBR was selected as the initial approach because: (a) diabetes management guidelines are general in nature, requiring personalization; (b) a wide range of both physical and lifestyle factors influence BG levels; and (c) CBR has been successfully applied to managing other chronic medical conditions (Holt et al. 2005; Bichindaritz 2008).

The first step in developing 4DSS was to build a case base as a central knowledge repository. Although abundant BG data was initially available, usable cases were not. This is because the life-events coinciding with BG levels, used by physicians to determine appropriate therapy, were not routinely recorded. To acquire contextualized cases for the system, a clinical research study was conducted, involving 20 T1D patients. Each patient participated for six weeks, manually entering daily BG, insulin, and life-event data into an experimental database via a Web-based interface. Physicians reviewed the data, detecting BG control problems and recommending therapeutic adjustments. Patients implemented the recommended adjustments (or not), and physicians re-

viewed subsequent data to evaluate the clinical outcomes, in an iterative cycle. Problems, solutions and outcomes were structured into cases and stored in the case base. Figure 1 shows a sample case from 4DSS. Fifty such cases were acquired during the clinical research study.

Problem: Nocturnal hypoglycemia. BG levels are dangerously low all night. The patient reports feeling “totally out of it” when she wakes up. She does not eat anything to correct the hypoglycemia until noon. She had not eaten a bedtime snack the night before.

Solution: The patient should always have a mixed-nutrient snack before bed. She should lower her overnight basal rate. The combination of more food and less insulin will prevent overnight lows.

Outcome: The patient reports eating mixed nuts and crackers before bed. She sets the basal rate in her pump as advised. BG data for subsequent weeks shows that the problem is resolved.

Figure 1: A sample case

The sample case records an actual problem of nocturnal hypoglycemia. Hypoglycemia, or low BG, leads to weakness, confusion, dizziness, sweating, shaking, and, if not treated in time, seizures, coma, or death. Hyperglycemia, or high BG, contributes to long-term diabetic complications. Extremely high BG levels can cause diabetic ketoacidosis, a serious condition leading to severe illness or death. It is important to note that patients do not know when problems are impending and are frequently unaware of problems even once they occur. Problems that occur when patients are asleep, as in the sample case, are especially dangerous.

Typically in CBR systems, reasoning begins with a known problem that can be readily described and elaborated. Solving a given problem entails finding and adapting the most similar, or most useful, case in the case base. In this domain, problems are not usually given, or known *a priori*, but must be detected in continuous patient data. Our approach was to model automated problem detection routines on physician problem detection strategies. Rule-based routines were implemented to detect 12 common BG control problems identified by physicians. A 4DSS prototype was built to (a) detect BG control problems in patient data; (b) display detected problems to the physician, who would select those of interest; (c) retrieve, for each selected problem, the most applicable case in the case base; and (d) display the retrieved case to the physician as decision support in determining appropriate therapy to correct the problem.

Evaluation and feedback were obtained through a patient exit survey and two structured sessions in which diabetes practitioners evaluated system capabilities. 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 BG, insulin and life-event data helped them to identify BG trends more readily and adjust therapy more effectively. Conclusions were: (1) the 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 BG control problems; and (3) data entry time demands on the patient must be reduced. Results of this study were reported in (Schwartz, Shubrook, and Marling 2008; Marling, Shubrook, and Schwartz 2008; 2009).

A second clinical research study, involving 26 T1D patients, was conducted to (a) reduce patient time demands and (b) re-evaluate the 4DSS prototype. BG and insulin data stored in the patient’s pump were uploaded to the experimental database rather than entered by the patient. Patients were asked for their typical daily schedules, and these were used to approximate actual daily life-events. Patients were not required to supply continuous glucose monitoring (CGM) data, but it was uploaded for patients who normally used it as part of routine care. Data that could not be automatically transferred or approximated was omitted from consideration. During evaluation, approximately half as many problems were detected per patient per week as were detected in the first clinical study. This finding was statistically significant ($p = .017$), although there were no statistically significant differences between the two patient populations and no reason to suspect that patients were actually experiencing fewer problems.

An adverse event that occurred during this study highlights the potential for 4DSS to impact health and wellbeing. A participating patient experienced a problem in which his pump failed and stopped delivering insulin. He was aware that his BG was high, and he instructed the pump to deliver more insulin. However, he did not know that the pump was not functioning, and his BG 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 the patient was hospitalized. Had the system been running in real time, the patient might have been alerted to make a simple adjustment before experiencing DKA.

Conclusions from this study were: (1) lack of life-event and CGM data impairs the ability to detect clinical problems; and (2) extending system capabilities to predict and prevent problems presents new research challenges and new opportunities to improve health outcomes. Results of this study were published in (Schwartz et al. 2010).

A third clinical research study is currently underway with the goals of enlarging the case base, developing additional problem detection routines, and automatically adapting past solutions to meet specific needs of current patients. Twelve T1D patients have already completed a 3-month protocol in which they: (a) upload insulin pump and CGM data weekly; and (b) supply otherwise unavailable life-event data via a Web browser on a daily basis. To date: (1) 30 new cases

have been added to the case base; (2) six new problem detection routines have been developed; and (3) a case adaptation module has been implemented.

Machine Learning Classification of BG Plots

During 4DSS development, we encountered a type of BG problem that we could not readily detect by encoding physician problem detection strategies in rules. This was excessive glycemic variability, a bouncing back and forth between hypo and hyperglycemia, as illustrated by Figure 2. Glycemic variability is an active area of current diabetes research (Ceriello and Ihnat 2010; Kilpatrick, Rigby, and Atkins 2010). Excessive glycemic variability has been linked to hypoglycemia unawareness, an acutely dangerous condition, and to oxidative stress, which contributes to long-term diabetic complications (Monnier et al. 2006). Its successful detection would enable routine screening for all T1D patients, a valuable clinical application in and of itself.

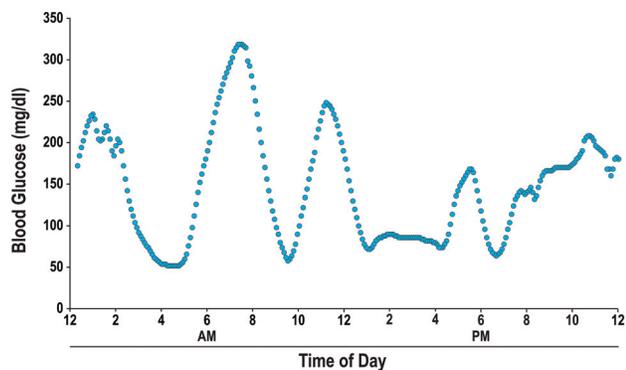


Figure 2: An actual patient’s daily blood glucose plot, exhibiting excessive glycemic variability

Although glycemic variability is difficult to measure or to formalize, physicians know it when they see it in BG plots, like the one shown in Figure 2. Therefore, we considered the quantifiable aspects of glycemic variability as they relate to physicians’ perception. We began with the best-accepted existing metric, the Mean Amplitude of Glycemic Excursion (MAGE), which captures the distance between the local maxima and minima (peaks and nadirs) of a BG plot (Service et al. 1970). We then designed two new metrics to capture aspects of variability not accounted for by MAGE. These are distance traveled, which captures overall daily fluctuation, and excursion frequency, which counts the number of significant glucose excursions in a day.

Three hundred BG plots were reviewed by two physicians (JS and FS), who characterized each plot as excessively variable or not, based on their gestalt perceptions of the plots. They were in agreement on 218 of the plots, which were then scored for MAGE, distance traveled, and excursion frequency. The scores and physician ratings were used to train ML algorithms to classify BG plots, using the Weka machine learning toolkit (Witten et al. 1999). Physicians then rated another 100 BG plots as excessively variable or not, for use in evaluating the ML classifiers. The best performing

ML algorithm, a naive Bayes classifier, matched concordant physician ratings 85% of the time. This preliminary work, which serves as proof of concept, was reported in abstract form (Marling et al. 2010); a full report is in press (Marling et al. 2011). We believe that a clinically viable screen for excessive glycemic variability can be built by: (a) obtaining glycemic variability ratings from many more physicians; (b) encoding additional measurable aspects of glycemic variability; (c) smoothing the BG data to reduce the effects of noise; and (d) training and evaluating additional ML classification algorithms.

Support Vector Regression for BG Prediction

Detecting BG problems, as in 4DSS and the screen for excessive glycemic variability, allows corrective action to be taken. The ability to predict impending BG problems before they occur would enable preemptive intervention. This would not only improve overall BG control, but could greatly impact patient safety. For example, the sleeping patient in the sample case (Figure 1) could be awakened and advised to eat before becoming hypoglycemic. Then she would not lie in a dangerous state all night long. Consequently, a significant part of our current research effort is directed towards designing ML models that can be trained on available clinical patient data to predict BG levels.

Since BG measurements have a natural temporal ordering, we approach the task of predicting BG levels as a time series forecasting problem. In time series prediction, the task is to estimate the future value of a target function based on current and past data samples. Numerous prediction problems in a wide array of domains ranging from finance (e.g. stock market (Kim 2003)), currency exchange rates (Giles, Lawrence, and Tsoi 2001), to medicine (e.g. sleep apnea (Aguirre, Barros, and Souza 1999)), environment (e.g. air quality (Perez and Reyes 2001), rainfall rate (Toth and Montanari 2000)), or power systems (e.g. electric utility load (Chen, Chang, and Lin 2004)) have been approached in the past as time series forecasting problems.

We have conducted a preliminary experimental evaluation in which a Support Vector Regression (SVR) model (Smola and Scholkopf 1998) was trained to predict the BG levels of a T1D patient. An arbitrary pivot date was selected about one month into the experimental data. Then 7 days before the pivot date were used to create training data, while test data was created from the 3 days following (and including) the pivot date. Since BG measurements are recorded by CGM systems every 5 minutes, one day may contribute up to 288 training or testing examples. Two separate SVR models were trained and tested to predict the BG levels for 30 and 60 minutes into the future. Training and testing examples were represented as feature vectors using the following set of features:

1. The BG level of patient x at present time t_0 .
2. A simple moving average over 4 past points from, and including, t_0 .
3. An exponentially smoothed rate of change in BG level over 4 past points from, and including, t_0 .

4. Bolus dosage totals starting 30 minutes before prediction time, computed for durations of 30 minutes and 10 minutes respectively. The bolus dosage refers to insulin that is injected before meals and/or to correct for hyperglycemia.
5. Basal rate averages starting right before prediction time, over 5 or 15 minute time intervals. This is the rate at which insulin is slowly and continuously infused into the patient by the pump. The basal rate changes throughout the day to accommodate changing insulin needs.
6. Meal carbohydrate amounts starting 30 minutes before prediction time, for durations of 30 minutes and 15 minutes respectively.
7. Exercise intensity averages starting right before prediction time, over 5, 30 or 60 minute time intervals. Exercise tends to amplify the effect of insulin. This effect influences BG levels during and after exercise; the length of the effect depends on the length and the intensity of the exercise.

The influence that each type of event exerts on the BG level is known to vary with time. This specific time dependent variability was taken into account through the offset and the length of the various time intervals that were used to define the features above. For example, the effects of exercise are strongest while the patient is exercising, but they may persist for several hours, especially if exercise is intense. This is why exercise features are computed in shorter 5-minute intervals close to the time of exercise, with intervals lengthening to 30 and 60 minutes as exercise recedes into the past. The SVR models were trained with a linear kernel, using a capacity parameter $C = 100$, and a default tube width $\epsilon = 1.0$. We used the LIBSVM implementation of SVMs for regression (Chang and Lin 2001). In Table 1, we compare the performance of the SVR models trained to predict BG level for 30 and 60 minutes into the future with the simple baseline $BGL(x, t_0)$ that uses the present BG level to predict any future BG level value. We use this simple baseline for comparison only because it was found to outperform more complex moving average and rate of change baselines.

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
$BGL(x, t_0)$	25.1	0.84	87.8	11.8	0.0	0.4	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
$BGL(x, t_0)$	43.2	0.52	74.5	21.5	2.2	1.8	0

Table 1: SVR and baseline $BGL(x, t_0)$ results

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) (Kovatchev et al. 2004). CEGA is a standard for evaluating the accuracy of BG measurement that is normally used to assess the quality of blood glucose sensors.

As shown in Figures 3 and 4, the Clarke Error Grid breaks a scatter plot into five regions:

- A. Points within 20% of the actual BG value
- B. Points that are more than 20% off but that would not lead to inappropriate treatment
- C. Points leading to unnecessary, but not harmful, treatment
- D. Points that obscure hypoglycemia or hyperglycemia, leading to a lack of necessary treatment
- E. Points misclassifying hypoglycemia as hyperglycemia, or vice versa, leading to harmful treatment

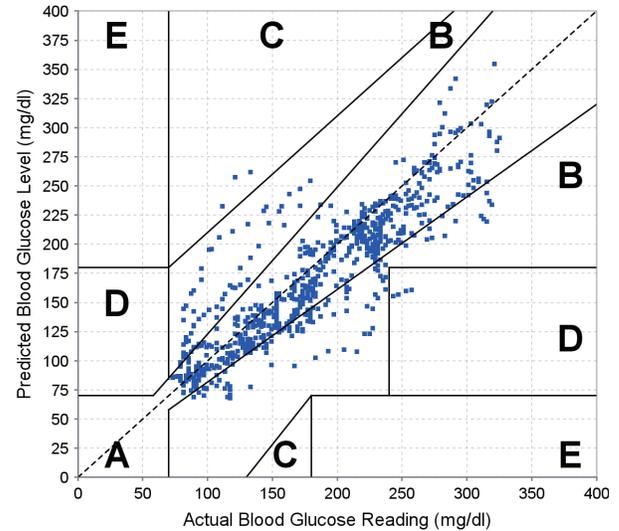


Figure 3: Performance of SVR for 60 minute prediction

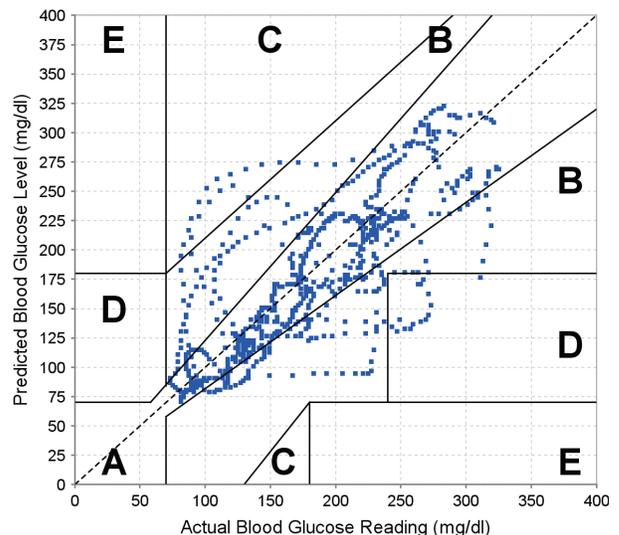


Figure 4: Performance of baseline for 60 minute prediction

The SVR models are promising, as they outperform the baselines on all performance measures. The two CEGA plots

in Figures 3 and 4 show the performance of SVR and the $BGL(x, t_0)$ baseline respectively for the 60 minute prediction time. The plots clearly show that, overall, the learned SVR model makes predictions that are closer to the ideal diagonal line.

To account for individual patient differences, a predictive model is trained for each patient. We also plan to explore transfer learning approaches that effectively exploit data coming from multiple patients in order to improve the model predictions, which will be especially useful for patients with limited historical data. Trained prediction models will be stored in a new case base of models, so that we may further consider the possibilities of adapting past models to bootstrap predictions for new patients. Since the patient data is often inaccurate or incomplete, we are investigating learning methods that are robust in the presence of missing or uncertain data and that can also identify data anomalies automatically.

AI Technology Transfer: Another Difficult Challenge

In this section, we share observations on the difficulties of transitioning AI technology from the university setting to the real world, based on the experiences of two of the authors (CM and FS). We do this from the perspective of faculty members intent upon keeping our “day jobs” of conducting academic research, educating students, and treating patients (FS). We recognize that: (a) issues differ for AI researchers working for companies or starting up their own companies; and (b) issues are more easily raised than resolved. Some lessons learned are shared.

University technology transfer: The goal of the university technology transfer office (TTO) is to facilitate commercialization of intellectual property and to ensure that the university benefits financially from the ideas of its faculty. TTO goals for economic development have not always harmonized with the faculty-held tenet of broad dissemination of knowledge or the physician’s desire to improve the health and wellbeing of patients. A high turnover of key personnel in the TTO at our university has also hindered progress. *Lesson learned:* TTO personnel will conscientiously do their jobs, and that means extra work for academics. An executive overview presentation (just 10 slides) can help to provide focus and save time in TTO meetings.

Technology leaks: Ideas that we have published and/or discussed with industry representatives have subsequently been incorporated into marketed products without attribution. This has been good for patient care and good for the commercial products, but bad for researcher financial support and researcher morale. *Lesson learned:* Transitioning technology does not necessarily equate to monetizing it; there may be intangible benefits.

Patents: The TTO filed a patent application on our behalf. Should software be patented? Many computer scientists would argue otherwise. Arguments in favor were: (a) companies would not invest in unpatented technologies; and (b) freeware is not suitable for safety-critical medical applications.

Safety: When AI technology is to be used directly by patients in the United States, it must first be approved by the U.S. Food and Drug Administration (FDA). While critical for ensuring patient safety, the FDA approval process entails extensive investments of time and money, making it infeasible for academics. As a consequence, we have directed our focus to tools that help physicians manage patients, rather than tools used directly by patients. This enables us to ensure patient safety by keeping a professional in the loop, which removes FDA concerns but limits the avenues of research and application. *Lesson learned:* Know the regulatory agencies in your application domain. A funding proposal was rejected when a reviewer found it “astounding” that we planned to run software in a medical device.

Conflict of interest: Patents, licenses, and new products resulting in economic benefit to a faculty member can be viewed as conflicts of interest by the university. This can be so even when rights are co-owned by the university. In particular, funding graduate students with money raised for commercialization and/or supervising graduate students whose work could contribute to product development may lead to legal, as well as ethical, conflicts. In such instances, faculty may need to leave the university to complete product development without conflict. There is no conflict, however, if a company funds university research and then uses the results in its products.

Lack of conflict-free money: A CEO of a medical device company, reviewing our research, shared that patients expect better and more intelligent software with each new hardware release. However, patients also expect the software to be included for free. Free is good for patients, and good for medical device companies, but bad for funding AI research.

Forging Ahead

Despite technology transfer difficulties, we are forging ahead with plans to make intelligent diabetes management a reality for patients and physicians. We have a waiting list of patients who have volunteered to participate in clinical research studies. They are counting on us to translate the research into practical tools they can use. We envision the following potential avenues of commercialization and use:

- Software could be marketed directly to physicians for office use
- Software could be included in electronic health record (EHR) systems
- Software could be embedded in insulin pumps and/or smart phones for patient use
- Software could be incorporated in continuous glucose monitoring (CGM) systems, so that all BG plots would come with associated analyses
- BG Control Centers could be established, where BG data could be uploaded, analyzed, and monitored by advanced practice nurses, who would forward appropriate findings to physicians and patients

In summary, diabetes management is more than a challenging domain for AI research. It is an opportunity for AI ap-

plications to positively impact the health and wellbeing of people with diabetes.

Acknowledgments

We gratefully acknowledge support from Medtronic, the Ohio University Russ College Biomedical Engineering Fund, the Ohio University College of Osteopathic Medicine Research and Scholarly Affairs Committee, and the Ohio University Diabetes Research Initiative. We thank Tessa Cooper, Eric Flowers, Thomas Jones, Anthony Maimone, Wesley Miller, Stanley Vernier and Donald Walker for software development and knowledge engineering. We also thank our research nurses, Cammie Starner, Lynn Petrik, Barbara Pellegrin, and Regina Reed.

References

- Aguirre, L. A.; Barros, V. C.; and Souza, A. V. 1999. Nonlinear multivariable modeling and analysis of sleep apnea time series. *Computers in Biology and Medicine* 29(3):207–228.
- Bichindaritz, I. 2008. Case-based reasoning in the health sciences: Why it matters for the health sciences and for CBR. In Althof, K.-D.; Bergmann, R.; Minor, M.; and Hanft, A., eds., *Advances in Case-Based Reasoning: 9th European Conference, ECCBR*, 1–17. Berlin: Springer.
- Ceriello, A., and Ihnat, M. A. 2010. Glycaemic variability: A new therapeutic challenge in diabetes and the critical care setting. *Diabetic Medicine* 27(8):862–867.
- Chang, C.-C., and Lin, C.-J. 2001. *LIBSVM: A library for support vector machines*. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>, accessed April, 2011.
- Chen, B.-J.; Chang, M.-W.; and Lin, C.-J. 2004. Load forecasting using support vector machines: A study on EUNITE competition 2001. *IEEE Transactions on Power Systems* 19(4):1821–1830.
- 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(14):977–986.
- Giles, L. C.; Lawrence, S.; and Tsoi, A. C. 2001. Noisy time series prediction using a recurrent neural network and grammatical inference. *Machine Learning* 44(1/2):161–183.
- Holt, A.; Bichindaritz, I.; Schmidt, R.; and Perner, P. 2005. Medical applications in case-based reasoning. *The Knowledge Engineering Review* 20(3):289–292.
- Kilpatrick, E. S.; Rigby, A. S.; and Atkins, S. L. 2010. For debate. Glucose variability and diabetes complication risk: We need to know the answer. *Diabetic Medicine* 27(8):868–871.
- Kim, K.-J. 2003. Financial time series forecasting using support vector machines. *Neurocomputing* 55:307–319.
- Kovatchev, B. P.; Gonder-Frederick, L. A.; Cox, D. J.; and Clarke, W. L. 2004. 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.
- Marling, C. R.; Shubrook, J. H.; Vernier, S. J.; Wiley, M. T.; Cooper, T. L.; and Schwartz, F. L. 2010. New metrics for glucose variability using continuous glucose monitoring data. Tenth Annual Diabetes Technology Meeting, Diabetes Technology Society, Bethesda, MD, abstract number A97.
- Marling, C. R.; Shubrook, J. H.; Vernier, S. J.; Wiley, M. T.; and Schwartz, F. L. 2011. Characterizing blood glucose variability using new metrics with continuous glucose monitoring data. *Journal of Diabetes Science and Technology*. In press.
- Marling, C.; Shubrook, J.; and Schwartz, F. 2008. Case-based decision support for patients with type 1 diabetes on insulin pump therapy. In Althof, K.-D.; Bergmann, R.; Minor, M.; and Hanft, A., eds., *Advances in Case-Based Reasoning: 9th European Conference, ECCBR*, 325–339. Berlin: Springer.
- Marling, C.; Shubrook, J.; and Schwartz, F. 2009. Toward case-based reasoning for diabetes management: A preliminary clinical study and decision support system prototype. *Computational Intelligence* 25(3):165–179.
- Monnier, L.; Mas, E.; Ginet, C.; Michel, F.; Villon, L.; Cristol, J.; and Colette, C. 2006. Activation of oxidative stress by acute glucose fluctuations compared with sustained chronic hyperglycemia in patients with type 2 diabetes. *Journal of the American Medical Association* 295(14):1681–1687.
- Perez, P., and Reyes, J. 2001. Prediction of particulate air pollution using neural techniques. *Neural Computing and Applications* 10(2):165–171.
- Schwartz, F. L.; Vernier, S. J.; Shubrook, J. H.; and Marling, C. R. 2010. 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.
- Schwartz, F. L.; Shubrook, J. H.; and Marling, C. R. 2008. 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.
- Service, F.; Molnar, G.; Rosevear, J.; Ackerman, E.; Gatewood, L.; and Taylor, W. 1970. Mean amplitude of glycemic excursions, a measure of diabetic instability. *Diabetes* 19(9):644–655.
- Smola, A. J., and Scholkopf, B. 1998. A tutorial on support vector regression. Technical Report TR-98-030, NeuroCOLT2 Technical Report Series.
- Toth, E.; Brath, A., and Montanari, A. 2000. Comparison of short-term rainfall prediction models for real-time flood forecasting. *Journal of Hydrology* 239:132–147.
- Witten, I. H.; Frank, E.; Trigg, L.; Hall, M.; Holmes, G.; and Cunningham, S. J. 1999. Weka: practical machine learning tools and techniques with java implementations. In *Proceedings ICONIP/ANZIIS/ANNES99: Future Directions for Intelligent Systems and Information Sciences*, 192–196.