Journal of Diabetes Science and Technology
Volume 7, Issue 1, January 2013
© Diabetes Technology Society

ORIGINAL ARTICLE

Hypoglycemia Early Alarm Systems Based on Recursive Autoregressive Partial Least Squares Models

Elif Seyma Bayrak, B.S.,[1] Kamuran Turksoy, B.S.,[2] Ali Cinar, Ph.D.,[1,2] Lauretta Quinn, Ph.D., R.N.,[3] Elizabeth Littlejohn, M.D.,[4] and Derrick Rollins, Ph.D.[5]

Abstract

Background:

Hypoglycemia caused by intensive insulin therapy is a major challenge for artificial pancreas systems. Early detection and prevention of potential hypoglycemia are essential for the acceptance of fully automated artificial pancreas systems. Many of the proposed alarm systems are based on interpretation of recent values or trends in glucose values. In the present study, subject-specific linear models are introduced to capture glucose variations and predict future blood glucose concentrations. These models can be used in early alarm systems of potential hypoglycemia.

Method:

A recursive autoregressive partial least squares (RARPLS) algorithm is used to model the continuous glucose monitoring sensor data and predict future glucose concentrations for use in hypoglycemia alarm systems. The partial least squares models constructed are updated recursively at each sampling step with a moving window. An early hypoglycemia alarm algorithm using these models is proposed and evaluated.

Results:

Glucose prediction models based on real-time filtered data has a root mean squared error of 7.79 and a sum of squares of glucose prediction error of 7.35% for six-step-ahead (30 min) glucose predictions. The early alarm systems based on RARPLS shows good performance. A sensitivity of 86% and a false alarm rate of 0.42 false positive/day are obtained for the early alarm system based on six-step-ahead predicted glucose values with an average early detection time of 25.25 min.

Conclusions:

The RARPLS models developed provide satisfactory glucose prediction with relatively smaller error than other proposed algorithms and are good candidates to forecast and warn about potential hypoglycemia unless preventive action is taken far in advance.

J Diabetes Sci Technol 2013;7(1):206–214

Author Affiliations: [1]Department of Chemical and Biological Engineering, Illinois Institute of Technology, Chicago, Illinois; [2]Department of Biomedical Engineering, Illinois Institute of Technology, Chicago, Illinois; [3]College of Nursing, University of Illinois at Chicago, Chicago, Illinois; [4]Department of Pediatrics, Section of Endocrinology, Biological Sciences Division, University of Chicago, Chicago, Illinois; and [5]Department of Chemical and Biological Engineering, Iowa State University, Ames, Iowa

Abbreviations: (CGMS) continuous glucose monitoring sensor, (FPE) final prediction error, (PH) prediction horizon, (PLS) partial least squares, (RARPLS) recursive autoregressive partial least squares, (RMSE) root mean square error, (SSGPE) sum of squares of glucose prediction error, (T1DM) type 1 diabetes mellitus

Keywords: hypoglycemia alarms, partial least squares regression, recursive algorithm, type 1 diabetes mellitus

Corresponding Author: Ali Cinar, Ph.D., Department of Chemical and Biological Engineering, Illinois Institute of Technology, 10 W. 33rd St., Chicago, IL 60616; email address *cinar@iit.edu*

Introduction

Prevention of hypoglycemia is a major challenge for people with diabetes who use insulin treatment to manage their blood glucose. People with type 1 diabetes mellitus (T1DM) may have several episodes of hypoglycemia per week. Untreated hypoglycemia may lead to unconsciousness, seizures, and death. Availability of continuous glucose monitoring sensors (CGMSs) has been a big motivation for studies designed to monitor, predict, and control blood glucose in T1DM.

Early alarms are essential and should provide enough time to take the necessary action to prevent hypoglycemia. Pump suspension has been reported based on early detection.[1–3] This article reports a novel approach for the prediction of future glucose values and early alarms to warn about impending hypoglycemia. Further evaluation of the algorithm for pump suspension or rescue carbohydrate is not studied in this paper. There are various hypoglycemia alarm algorithms in the literature based on different prediction methods and horizons.[2–6] Many alarm systems are based on interpretation of recent trends in glucose values[6,7] by tracking the slope of successive glucose values or extrapolating the current value. Reliable glucose prediction models are required to implement early alarms to reduce risk of hypoglycemia.

Optimal estimation using the Kalman filter to predict glucose levels and its rate of change was proposed,[4] and 30 min prediction horizons (PHs) using the sensitivity and specificity were reported as 90% and 79%, respectively.[8] Glucose prediction using artificial neural networks has also been studied.[9] The performance of the models has been reported in terms of root mean square error (RMSE) and model delay for different PHs. The RMSE is around 10, 18, and 27 mg/dl for 15, 30, and 45 min of PH, respectively. No results for detecting hypoglycemia and the alarm performance were reported.

Glucose insulin dynamics show intersubject/intrasubject variability. Metabolic and glycemic changes due to meal consumption or physical activity may lead to further variation in glucose–insulin dynamics. These variations can be captured with subject-specific recursive models. The recursive autoregressive partial least squares (RARPLS) algorithm uses CGMS readings to predict future glucose concentrations. Partial least squares (PLS) is a widely used multivariate regression method for modeling and monitoring in chemometrics, supervision of chemical process operations and social sciences. It is extended by using autoregressive terms to capture dynamic variations in data. The recursive updates are used to eliminate the influence of old data from the model.

Methods

Partial Least Squares

Partial least squares is a multivariate regression method, especially convenient for large number of highly correlated data sets. The PLS models summarize the original data matrix (input variables X) to extract the most predictive information for the response variable (Y) and maximize the covariance between X and Y. It derives its usefulness from its ability to analyze data with many, noisy, collinear, and even incomplete variables in both X and Y.[10] The "outer relations" for X and Y blocks are, respectively,

$$X = TP^T + E = \sum_{i=1}^{a} t_i p_i^T + E \quad (1)$$

$$Y = UQ^T + F = \sum_{i=1}^{a} u_i q_i^T + F \quad (2)$$

where E and F represent the residual matrices, t_i and u_i are latent vectors, and p_i and q_i are the loadings vectors for X and Y blocks, respectively. The inner relation can be built between u and t for every component. The model for inner relation is

$$u_h = b_h t_h \quad (3)$$

where $b_h = u_h' t_h / t_h' t_h$ is the regression coefficient.[11]

Latent vectors can be calculated using nonlinear iterative PLS,[12] Kernel,[13] or simple PLS[14] algorithm; no significant difference among these algorithms has been reported.

The appropriate number of latent variables that gives the best balance between the data fit and prediction is determined by various techniques such as cross validation.

Recursive Partial Least Squares

In order to adapt the changes in the system, it is necessary to avoid the influence of old data as new data become available. A recursive PLS algorithm with a moving window[15] is used for this purpose with a few modifications.

When new data are available at sampling time k, the PLS regression is performed using the updated data matrices:

$$X_{k+1}^T = [x_{k-w+1} \cdots\cdots x_k] \tag{4}$$

$$Y_{k+1}^T = [y_{k-w+1} \cdots\cdots y_k] \tag{5}$$

where w is the window size; the oldest data at sampling time k-w is excluded from the data matrices.

Partial least squares modeling works best when the data are fairly symmetrically distributed and have a fairly constant "error variance."[10] In a time-varying process, the mean levels of the variables may be changing with time.[16] Therefore, for the recursive application presented in this article, current mean and variance is calculated for the current window, and data are mean centered and scaled with variance at every sampling step.

Recursive and Autoregressive Partial Least Squares

The use of previous values of glucose concentration for building an autoregressive model captures dynamic variations in data better.[17] Recursive PLS regression is combined with autoregression to improve prediction results.

Y matrix consists of 1 to n steps ahead prediction (y_{k+n}, y_{k+n-1}, y_{k+n-2}, … , y_{k+1}), and X matrix includes the previous glucose measurement from CGMSs for the a previous sampling times:

$$X = [y_k y_{k-1} y_{k-2} \cdots y_{k-a}] \tag{6}$$

where a is the order of autoregression for glucose. Different model parameters are computed for each PH.

Evaluation Criteria

Prediction error is expressed in terms of RMSE:

$$RMSE = \sqrt{\frac{\Sigma(y - \hat{y})^2}{n}} \tag{7}$$

where $\hat{y}$ is the predicted glucose concentration (mg/dl) by the model and n is the data length. The sum of squares of glucose prediction error (SSGPE) is

$$SSGPE\% = \sqrt{\frac{\Sigma(y - \hat{y})^2}{\Sigma y^2}} \times 100 \tag{8}$$

Early Hypoglycemia Alarms

N-step-ahead predicted glucose concentration values from the RARPLS algorithm is used for early hypoglycemia alarm. In clinical trials, it is observed that accuracy of CGMS data becomes more important in lower glucose readings that lead to hypoglycemia. To be on the safe side and avoid immediate hypoglycemia, our alarm algorithm suggests validating CGMS data with finger stick or YSI measurements when the predicted glucose value crosses the safety threshold (90 mg/dl).

The alarm algorithm first checks the current data, and if the glucose concentration is under the hypoglycemic threshold, an *immediate hypoglycemia alarm* is triggered. If a glucose value is higher than the threshold of 90, the algorithm checks for predictions of future glucose values to determine the need to trigger an early hypoglycemia alarm. When the *n*-step-predicted value crosses the hypoglycemia threshold (70 mg/dl), an *early hypoglycemia alarm* is raised. The flow chart of the alarm algorithm is illustrated in **Figure 1**.

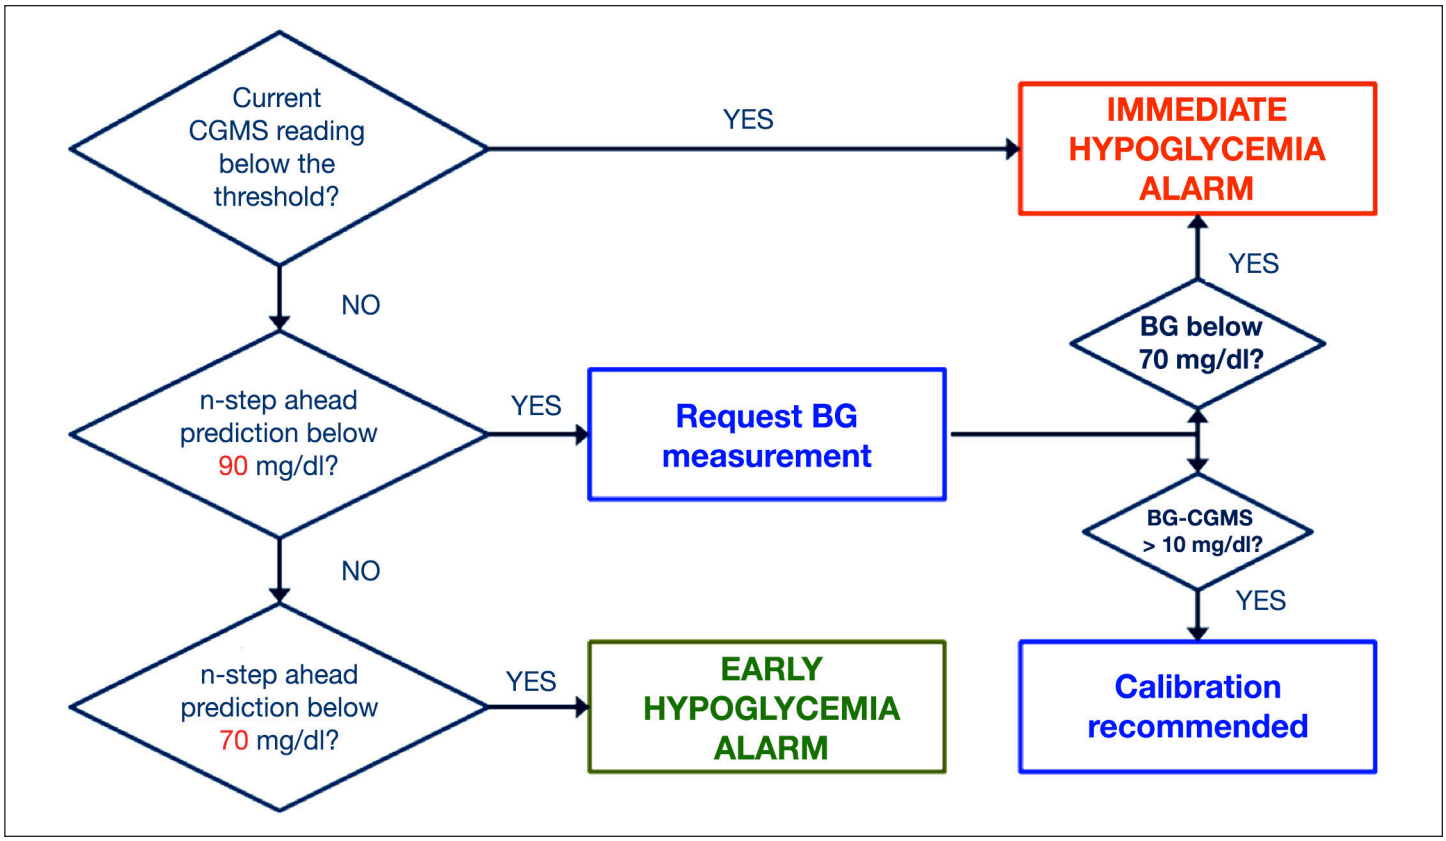

Figure 1. Early and immediate hypoglycemia detection algorithm flow chart. BG, blood glucose.

Results

Subject Data

The prediction and alarm algorithm is tested retrospectively with data from 17 subjects, age varying between 18 and 25 years (data collected at University of Illinois Chicago, College of Nursing, and Iowa State University). A Medtronic Guardian Real-Time Continuous Glucose Monitoring System is used for subcutaneous glucose data acquisition (5 min data interval). The data analyzed in this study are a reflection from the daily life of T1DM subjects; no insulin-induced hypoglycemia is included in the data. However, each data set included at least one or two episodes of hypoglycemia. One-week-long data were available to test the algorithm. For most subjects, interruptions occurred in data. The alarm prediction algorithm was used for sections of data that did not include any interruptions. Consequently, data lengths of half day to three days were used in testing the prediction algorithm. Another data set from the Diabetes Research in Children Network[18] is also used to assess the alarm algorithm. *In silico* study is performed by using the University of Virginia/Padova metabolic simulator with 20 subjects (10 adults and 10 adolescents).

Preprocessing of the Data

Even though continuous glucose sensors have embedded analog filters, the output signal is still noisy and should be filtered in order to enhance their signal-to-noise ratio.[19] Real-time filtering of CGMS values reduces the inaccuracy and noise in data but causes delay in prediction.[20] An adaptive Kalman filtering algorithm is proposed by Facchinetti

and coauthors[21] for real-time denoising. In this study, two filtering approaches were considered. First, a noncausal Savitzky–Golay smoothing filter is used to eliminate the noise.[22] The second approach is a real-time adaptation of the Savitzky–Golay filter. In the real-time algorithm, a window consisting of (f - 1) previous data and the current data is used by fitting a first-order polynomial. This data window is updated with each new CGMS datum that becomes available. A window size of 9 (f = 9) and a first-order polynomial (n = 1) are selected to obtain good filtering by targeting small RMSE and small filtering delay in the prediction. **Figure 2** shows the plot of raw CGMS data and filtered data after implementing the real-time algorithm, illustrating that filtering did not cause significant delay in glucose concentration information.

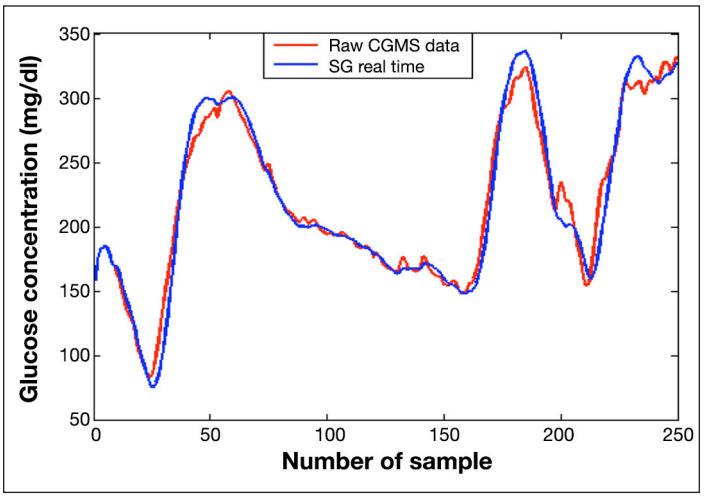

Figure 2. Raw CGMS data (blue line) is filtered to eliminate noise in the signal. SG, sensor glucose.

Recursive Autoregressive Partial Least Squares

Dynamic (recursive) versus time-invariant models are analyzed for the same data set. For the time-invariant model, day-long CGMS data are used to construct the model, and the rest of the data of the same subject is used for model validation, whereas the dynamic model is updated at each sampling time with a moving window length of 1 day.

Akaike's final prediction error (FPE) criteria are used for the selection of autoregressive model order (a).[23] The maximum model order is arbitrarily set to 9. **Table 1** displays the FPE values for each order. Minimum FPE is obtained while a = 4.

The effect of autoregressive order on glucose prediction with different PH is also analyzed using minimum RMSE (mg/dl) criteria. For PHs (PH = 1, 2 , …, 10), minimum RMSE is obtained when the order of autoregression is 4, which is consistent with the FPE criteria.

Table 1.
Autoregressive Order Selection Using Minimum Final Prediction Error Criteria

Autoregressive order (a)	1	2	3	4	5	6	7	8	9
FPE	13.38	8.16	8.18	8.06	8.13	8.18	8.19	8.16	8.21

Table 2 shows the RMSE and SSGPE values for both recursive and time-invariant autoregressive PLS algorithm using the noncausal smoothing filter and real-time filtering algorithm. The dynamic model is capable of capturing variations in a subject's glucose concentration and minimizing the error better. Even though two consecutive data sets from the same subject are used for both constructing and validating the model, the static model yields higher errors for each PH.

Recursive autoregressive partial least squares provided the minimum RMSE and %SSGPE values for future glucose prediction. This model can be considered as a reliable glucose prediction model for early hypoglycemia detection, and it is used for the alarm algorithm in the rest of this article.

Hypoglycemia Detection and Alarms

A hypoglycemic event is defined as continuous sequence of data below the selected threshold instead of individual data points. If there are more than two steps between the groups of data that are below 70 mg/dl, they are considered two different hypoglycemia events. An alarm is also defined as a continuous event and considered true positive if it is issued up to 60 min before a hypoglycemic event, and the alarms raised during the event are not counted as early hypoglycemia alarm since our focus is in early detection (**Table 3**). False positive region is defined as the region where the alarm is triggered out of true positive region (long before or after a hypoglycemic event). An alarm is

Table 2.
Root Mean Square Error and Percentage Sum of Squares of Glucose Prediction Error Values for Time-Invariant Models and Dynamic Models with Moving Window Using Continuous Glucose Monitoring Sensor Data

PH	Time-invariant autoregressive PLS algorithm		RARPLS algorithm	
	%SSGPE	RMSE (mg/dl)	%SSGPE	RMSE (mg/dl)
Savitzky–Golay smoothing filter				
2	0.76	1.06	0.67	0.46
4	2.33	3.2	2.07	1.95
6	4.55	6.17	4.04	4.45
8	7.27	9.81	6.43	7.67
10	10.32	13.92	9.16	11.29
Real-time first-order filter				
2	3.83	3.47	1.66	1.78
4	8.24	7.48	4.06	4.32
6	12.14	11.03	7.35	7.79
8	16.06	14.63	11.22	11.84
10	20.74	18.95	15.19	15.92

Table 3.
Event Definitions for Early Hypoglycemia Alarm Algorithm

True positive	If alarm is raised and no true hypoglycemia is recorded in the following 60 min
False positive	If alarm is raised and no true hypoglycemia is recorded in the following 50 min
False negative	Early hypoglycemia alarm is not issued 60 min in advance of a true hypoglycemic event
Time to detection	Time between the first alarm and the true hypoglycemic event

considered false negative if it is not raised up in true positive region (60 min advance of true hypoglycemic event). Time to detection is another important parameter in hypoglycemia alarms and is defined as the time between the first alarm raised *within the true positive region* and the beginning of the true hypoglycemic event.

We only considered early hypoglycemia alarms to evaluate our algorithm's performance; alarms held during the event are not counted as true positive, as the continuous sensors are already equipped with immediate alarms for the current data point. Sensitivity, false positive ratio, and time to detection are reported to determine alarm performance. Sensitivity is used as the measure of correctly identified positives, and false positive ratio is defined to quantify false alarm rate per day:

$$Sensitivity = \frac{TruePositive}{TruePositive + FalseNegative} \tag{9}$$

$$False\ Alarm\ Ratio = \frac{Number\ of\ FalsePositive}{Time(Data\ Length)} \tag{10}$$

Different PHs are analyzed to find better settings for early alarms. The hypoglycemia threshold was held constant at 70 mg/dl, and 4-, 6-, 8-, and 10-step-ahead prediction used the RARPLS algorithm. Although using 10-step-ahead prediction provides slightly better time to detection, it resulted in significantly lower sensitivity. As the PH decreases, better sensitivity, lower false alarm ratio, and relatively shorter time to detection are obtained. Sixty-nine hypoglycemia

events existed in the data analyzed. Performance results for different PHs with prefiltering and real-time filtering are illustrated in **Table 4**. Time to detection is calculated as an average time of all true positive alarms triggered.

Table 4.
Alarm Performance Evaluation for Different Prediction Horizons

PH (steps ahead)	Sensitivity	False alarm rate (false positives/day)	Time to detection (min)
Savitzky–Golay smoothing filter			
10	0.59	0.56	34.71
8	0.68	0.48	32.73
6	0.90	0.36	28.25
4	0.91	0.31	20.76
Real-time first-order filter			
10	0.56	0.61	32.34
8	0.67	0.52	31.08
6	0.86	0.42	25.25
4	0.89	0.35	18.83

There is a slight difference in sensitivity obtained with four- and six-step-ahead PHs as expected. The 25.25 min advanced hypoglycemia detection time can provide a sufficient period for preventive action. Therefore we selected the six-step-ahead PH optimal for our alarm algorithm. **Figure 3** captures two days of a subject who is exposed to two hypoglycemic episodes, and an alarm was triggered 30 and 20 min in advance of the events, respectively.

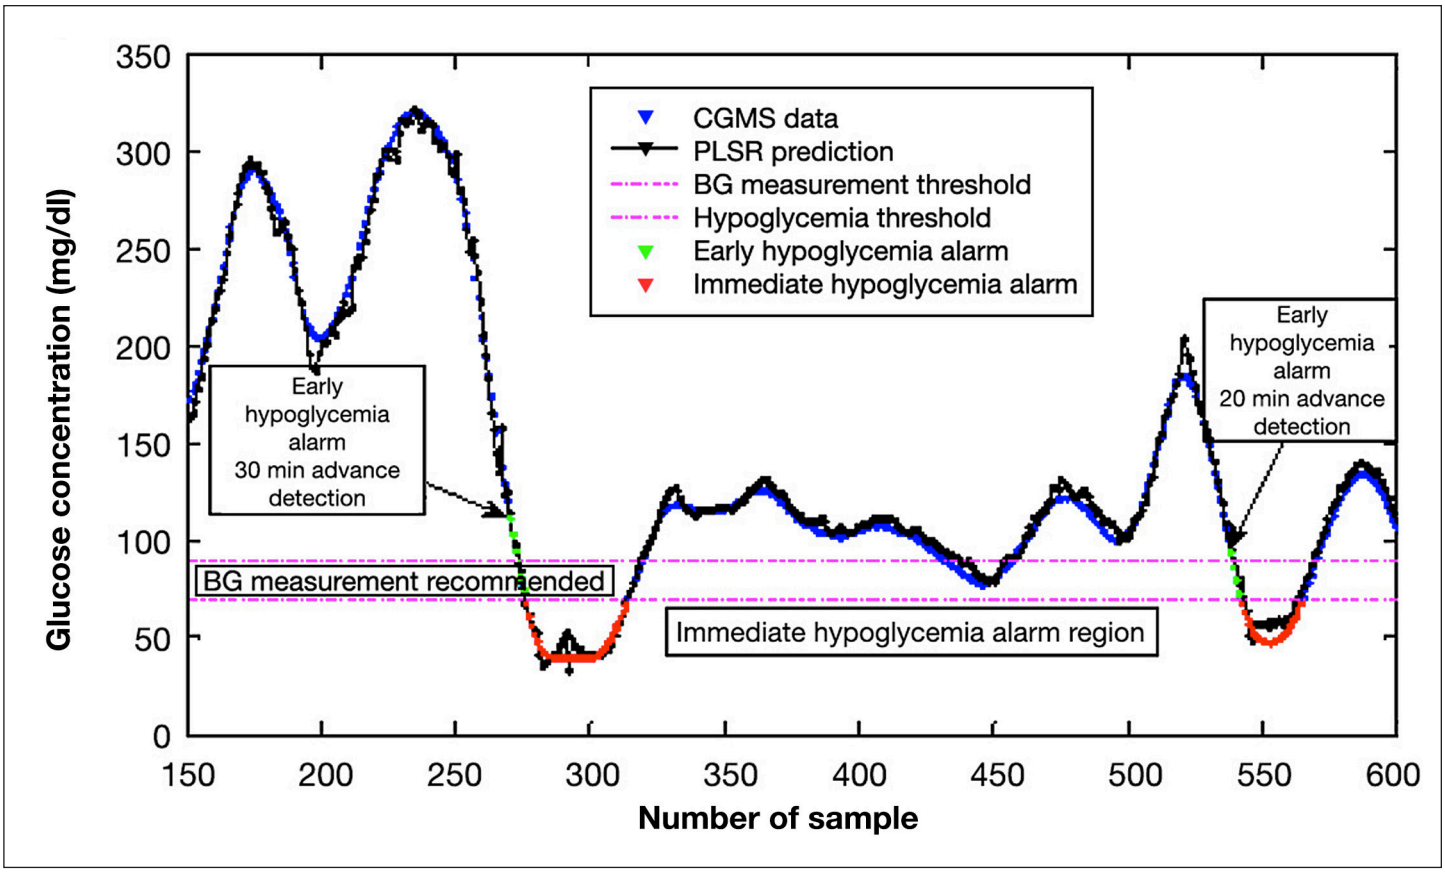

Figure 3. Early hypoglycemia alarms are issued, on average, 25 min in advance of a hypoglycemia event. PLSR, partial least squares regression; BG, blood glucose.

Discussion

We proposed a RARPLS algorithm to model and predict future glucose concentration of a subject with T1DM. We found modeling the variations in glucose concentration as a better approach than tracking only the recent changes in the glucose concentration. Glucose concentration is highly variable, and therefore, using time-invariant models would not provide satisfactory results, whereas our recursive algorithm that updated the model in every sampling step was capable of capturing the high variations.

The RARPLS models developed were then used in hypoglycemia warning alarms. One of the most important parameters for hypoglycemia alarms is the PH. We could predict up to 10-step-ahead future glucose concentration, and prediction results are satisfactory; maximum 11.29 and 15.92 RMSE are obtained for 10-step-ahead predictions with noncausal and online filters, respectively. However, because of this relatively higher error, the sensitivity was poor for longer PH. Six-step-ahead (30 min) PH increased the sensitivity significantly and provided enough time for patients or care providers to take preventive action. A real-time algorithm with the first-order filter provided satisfactory results. A sensitivity of 86% and a 0.42 false positive rate are obtained based on six-step-ahead prediction. The filter can be improved further by balancing accuracy and prediction delay.

Recursive autoregressive partial least squares is a strong candidate for developing reliable linear glucose prediction models and can be used in early hypoglycemia prediction and warning systems.

Funding:

This work was supported by National Institutes of Health Grant DK 085611-01.

References:

1. Agrawal P, Welsh JB, Kannard B, Askari S, Yang Q, Kaufman FR. Usage and effectiveness of the low glucose suspend feature of the medtronic paradigm veo insulin pump. J Diabetes Sci Technol. 2011;5(5):1137–41.
2. Buckingham B, Chase HP, Dassau E, Cobry E, Clinton P, Gage V, Caswell K, Wilkinson J, Cameron F, Lee H, Bequette BW, Doyle FJ 3rd. Prevention of nocturnal hypoglycemia using predictive alarm algorithms and insulin pump suspension. Diabetes Care. 2010;33(5):1013–7.
3. Buckingham B, Cobry E, Clinton P, Gage V, Caswell K, Kunselman E, Cameron F, Chase HP. Preventing hypoglycemia using predictive alarm algorithms and insulin pump suspension. Diabetes Technol Ther. 2009;11(2):93–7.
4. Palerm CC, Willis JP, Desemone J, Bequette BW. Hypoglycemia prediction and detection using optimal estimation. Diabetes Technol Ther. 2005;7(1):3–14.
5. Eren-Oruklu M, Cinar A, Quinn L. Hypoglycemia prediction with subject-specific recursive time-series models. J Diabetes Sci Technol. 2010;4(1):25–33.
6. Dassau E, Cameron F, Lee H, Bequette BW, Zisser H, Jovanovic L, Chase HP, Wilson DM, Buckingham BA, Doyle FJ 3rd. Real-time hypoglycemia prediction suite using continuous glucose monitoring: a safety net for the artificial pancreas. Diabetes Care. 2010;33(6):1249–54.
7. Cameron F, Niemeyer G, Gundy-Burlet K, Buckingham B. Statistical hypoglycemia prediction. J Diabetes Sci Technol. 2008;2(4):612–21.
8. Palerm CC, Bequette BW. Hypoglycemia detection and prediction using continuous glucose monitoring-a study on hypoglycemic clamp data. J Diabetes Sci Technol. 2007;1(5):624–9.
9. Pérez-Gandía C, Facchinetti A, Sparacino G, Cobelli C, Gómez EJ, Rigla M, de Leiva A, Hernando ME. Artificial neural network algorithm for online glucose prediction from continuous glucose monitoring. Diabetes Technol Ther. 2010;12(1):81–8.
10. Eriksson L, Johansson E, Kettaneh-Wold N, Wold S. Multi- and megavariate data analysis: principles and applications. Umeå: Umetrics; 2001.
11. Geladi P, Kowalski BR. Partial least squares regression: a tutorial. Anal Chim Acta. 1986;185:1–17.
12. Wold H. Soft modeling: the basic design and some extensions. Systems under Indirect Observation. 1982;2: 589–91.
13. Lindgren F, Geladi P, Wold S. The kernel algorithm for PLS. J Chemometrics. 1993;7(1):45–59.
14. De Jong S. SIMPLS: An alternative approach to partial least squares regression. Chemom Intell Lab Syst. 1993;18(3): 251–63.
15. Qin SJ. Recursive PLS algorithms for adaptive data modeling. Comp Chem Eng. 1998;22(4-5): 503–14.
16. Dayal BS, MacGregor JF. Recursive exponentially weighted PLS and its applications to adaptive control and prediction. J Process Control. 1997;7(3):169–79.

17. Eren-Oruklu M, Cinar A, Quinn L, Smith D. Estimation of future glucose concentrations with subject-specific recursive linear models. Diabetes Technol Ther. 2009;11(4):243–53.

18. The Diabetes Research in Children Network (DirecNet) Study Group. The effect of basal insulin during exercise on the development of hypoglycemia in children with type 1 diabetes. 2005. Available from: http://direcnet.jaeb.org/Studies.aspx?RecID=161.

19. Bequette BW. Continuous glucose monitoring: real-time algorithms for calibration, filtering, and alarms. J Diabetes Sci Technol. 2010;4(2):404–18.

20. Sparacino G, Facchinetti A, Cobelli C. "Smart" continuous glucose monitoring sensors: on-line signal processing issues. Sensors (Basel). 2010;10(7):6751–72.

21. Facchinetti A, Sparacino G, Cobelli C. An online self-tunable method to denoise CGM sensor data. IEEE Trans Biomed Eng. 2010;57(3):634–41.

22. Savitzky A, Golay MJ. Smoothing and differentiation of data by simplified least squares procedures. Anal Chem. 1964;36(8):1627–39.

23. Akaike H. Fitting autoregressive models for prediction. Annal Inst Stat Math. 1969;21(1):243–7.