

Symposium Paper

Hypoglycemia Prediction and Detection Using Optimal Estimation

CESAR C. PALERM, Ph.D.,^{1,*} JOHN P. WILLIS, Ph.D.,² JAMES DESEMONE, M.D.,³
and B. WAYNE BEQUETTE, Ph.D.¹

ABSTRACT

Patients with diabetes play with a double-edged sword when it comes to deciding glucose and A_{1c} target levels. On the one side, tight control has been shown to be crucial in avoiding long-term complications; on the other, tighter control leads to an increased risk of iatrogenic hypoglycemia, which is compounded when hypoglycemia unawareness sets in. Development of continuous glucose monitoring systems has led to the possibility of being able not only to detect hypoglycemic episodes, but to make predictions based on trends that would allow the patient to take preemptive action to entirely avoid the condition. Using an optimal estimation theory approach to hypoglycemia prediction, we demonstrate the effect of measurement sampling frequency, threshold level, and prediction horizon on the sensitivity and specificity of the predictions. We discuss how optimal estimators can be tuned to trade-off the false alarm rate with the rate of missed predicted hypoglycemic episodes. We also suggest the use of different alarm levels as a function of current and future estimates of glucose and the hypoglycemic threshold and prediction horizon.

INTRODUCTION

THE EFFICACY OF INTENSIVE TREATMENT in preventing complications of diabetes has been established by the Diabetes Control and Complications Trial¹ and the United Kingdom Prospective Diabetes Study.² In both trials the treatment regimens that reduced average glycosylated hemoglobin A_{1c} to approximately 7% (normal range is 4–6%) were associated with fewer long-term microvascular complications.

Recent evidence even suggests that these target levels might not be low enough.³

Alas, tight glycemic control carries with it a threefold increase in the danger for hypoglycemic episodes. Serious hypoglycemia can result in altered consciousness, convulsions (resulting in injury to the patient or others), coma, and even death.⁴ Iatrogenic hypoglycemia is the main barrier that prevents people with diabetes from achieving normal A_{1c} values. Iatrogenic hypoglycemia results from the inter-

¹ Department of Chemical and Biological Engineering; ² Department of Chemistry and Chemical Biology, Rensselaer Polytechnic Institute, Troy, New York.

³ Albany Medical Center, Albany, New York.

*Now at: Department of Chemical Engineering, University of California Santa Barbara, Santa Barbara, California.

play of absolute or relative insulin excess and compromised glucose counterregulation in type 1 diabetes.^{5,6}

Frequent hypoglycemic events are also linked to the hypoglycemia-associated autonomic failure syndrome, also known as hypoglycemia unawareness. The reduced awareness and response lead to worsening episodes of severe hypoglycemia, in which the patient is unable to take remedial action on her own.⁵⁻⁷

The increased risk, together with the advent of continuous glucose monitoring systems, have spurred interest in the detection and prediction of hypoglycemia.⁸⁻¹¹ In hypoglycemia prediction, the goal is to give enough warning to the user and thus avoid the condition entirely. While simple extrapolation of two recent measurements can be used to predict future glucose values, significant measurement noise can result in poor predictive performance.

Bequette¹² develops an optimal estimation approach that is applied to subcutaneous or interstitial glucose sensor current signals, and presents a structure to estimate the unmeasured blood glucose as well. Knobbe and Buckingham¹³ present another approach using an extended Kalman filter formulation that, besides interstitial fluid (ISF) and blood glucose, estimates the time lag between the two and the sensor sensitivity factor.

These glucose estimates can then be used to predict hypoglycemia. In this paper we use subcutaneous glucose measurements and estimates exclusively; this is done to keep the further challenges of blood glucose estimation out of the analysis. Furthermore, Choleau et al.⁸ suggest that ISF glucose measurements might be good enough for successfully preventing hypoglycemia.

In this paper we demonstrate the effect of sampling frequency, threshold selection (glucose level for hypoglycemia), and prediction horizon (how far ahead the glucose is being predicted) on the sensitivity and specificity of the prediction of hypoglycemia. For characteristic glucose changes we first show the inherent limitations in performance at predicting future glucose values, even in the limit of frequent measurements and no noise. We then discuss how optimal estimators can be tuned

to trade-off the false alarm rate with the rate of missed predicted hypoglycemic episodes. We also suggest the use of different alarm levels as a function of current and future estimates of glucose and the hypoglycemic threshold.

OPTIMAL ESTIMATION

Noisy signals and unmeasured states are a fact control systems engineers have to face in many applications. Measured variables are used to estimate the unmeasured states, using a model that relates the two. There is also a trade-off between the probability that a change in the measured variable is a change of the real state or is due simply to the measurement noise.

A common approach known as the Kalman filter¹⁴ uses recursive maximum likelihood for state estimation. The Kalman filter can be used to estimate blood glucose from subcutaneous measurements (for an overview, see Bequette¹²). For our purpose though, we will only use ISF glucose. What follows is a brief overview of the technique, starting with an intuitive derivation of the model used in this particular application.

We have a noisy measurement of the ISF glucose, which we denote by y_k , meaning the measurement at time step k . We know that this is a linear combination of the actual ISF glucose state (which we denote by g_k) and the measurement noise (v_k):

$$y_k = g_k + v_k \quad (1)$$

Besides using the Kalman filter to estimate the state g_k , we also want to estimate its rate-of-change, and the rate-of-change of the rate-of-change, so that we have trend information with which to perform our predictions.

We can express the ISF glucose state in terms of its rate-of-change (denoted by d_k):

$$g_{k+1} = g_k + d_k \quad (2)$$

this is, the new ISF glucose at time step $k + 1$ is the ISF glucose at the previous time step k plus the amount it changed over that time period.

Similarly, we can express the rate-of-change at time step $k + 1$ as the rate-of-change at the previous time step k plus the amount it changed (i.e., the rate-of-change of the rate-of-change), which we denote by f_k . Thus,

$$d_{k+1} = d_k + f_k \quad (3)$$

In our formulation we define f_k as a stochastic signal, meaning that its value changes by a random amount; this is what we call the process noise. Mathematically, we can write this as

$$f_{k+1} = f_k + w_k \quad (4)$$

where w_k is the process noise.

The measurement and process noise are considered stochastic processes; thus each has an associated covariance. In the context of the Kalman filter, Q is the covariance of the process noise, and R is the covariance of the measurement noise. These in general are not known, and thus become tuning parameters of the Kalman filter; the relative weighing shifts the trade-off between trusting the model prediction or the actual measurement. More details are presented in the Appendix, and the method used to select the Q/R ratio is presented in the section Limiting Factors in Prediction Reliability.

The Kalman filter is a two-step process; it first calculates the estimate based on the model, given the previous data, and then updates this estimate based on the actual measurement. The predictor and corrector equations can be combined; in the case of a 1-min sampling period with the Q/R ratio obtained for the noisy data ($Q/R = 2.1268 \times 10^{-6}$), we have the following set of equations to calculate the estimates:

$$\hat{g}_k = 0.7734 (\hat{g}_{k-1} + \hat{d}_{k-1}) + 0.2266y_k \quad (5)$$

$$\hat{d}_k = -0.025 \hat{g}_{k-1} + 0.975\hat{d}_{k-1} + \hat{f}_{k-1} + 0.025y_k \quad (6)$$

$$\hat{f}_k = -0.0013 (\hat{g}_{k-1} + \hat{d}_{k-1}) + \hat{f}_{k-1} + 0.0013y_k \quad (7)$$

where the coefficients are determined in part by the Kalman filter gain calculated for the given Q and R .

Using this formulation we then have estimates of the ISF glucose $\hat{x}_k$, its rate-of-change $\hat{d}_k$, and the rate-of-change of the rate-of-change $\hat{f}_k$, which we can then use to make predictions of what the glucose level will be into the future if the current trend persists.

SETTING THE FRAMEWORK

Before we can use the Kalman filter estimates to predict, or even detect, hypoglycemia, we need to know how hypoglycemia is defined in terms of glucose concentrations. This, unfortunately, varies depending on who you ask. Gregory and Aynsley-Green¹⁵ outline the different possible ways to define hypoglycemia, which include statistical, functional, and symptomatic. Of these, the functional definition of hypoglycemia based on the evidence of a counterregulatory response in healthy individuals seems to be gaining favor. This puts the threshold for blood glucose concentration in a range of 65–70 mg/dL (3.6–3.9 mM).^{6,7}

From the point of view of preventing iatrogenic hypoglycemia this makes sense, as through either an artificial pancreas or patient intervention in response to an alarm, we want to mimic as close as possible what the response for a normal individual would be. For our purpose then, we define hypoglycemia as a blood glucose concentration below 70 mg/dL (3.9 mM). This is also the threshold level used in other studies.^{8,11}

We have already mentioned that this study uses only subcutaneous signals, thus removing the added complexity of estimating blood glucose from the analysis. By the same token, there are some other assumptions underlying our work. We are assuming that we have a reasonably reliable sensor, whose behavior is linear throughout the range of interest; what we do include is noise, even cases with extreme levels that would not be expected in a working system. We are aware that there will be other issues involved in the deployment of such systems that need to be addressed, such as liability and legal concerns, among others. Our approach in this presentation is to overlook such issues and to focus on what is possible, even if other considerations end up in-

roducing other constraints in an actual implementation.

LIMITING FACTORS IN PREDICTION RELIABILITY

A simulation approach to examine how different factors limit the reliability of our predictions is required, as it allows easy replication of the study and it is easy to exclude confounding elements that would be present using *in vivo* data. In particular, we are interested in the effect of sampling frequency, prediction horizon, and threshold level.

A sinusoid is used as a simple description of the characteristic fluctuations of blood glucose

through a period of 3.5 h, going from a mean of 140 mg/dL (7.8 mM) to a peak of 240 mg/dL (13.3 mM) and then down to 40 mg/dL (2.2 mM). The maximum rate of change of this signal is 3 mg/dL/min (0.17 mM/min), which is consistent with observed clinical data.¹⁶ It is the same signal used in the study presented by Tamada et al.¹⁷ This is, in a way, more challenging than what the predictor would encounter in an actual patient under typical circumstances.¹⁶

Figure 1 shows the sinusoid signal used, together with three of the actual noisy “measurement” signals used in the simulations. A set of $n = 100$ such noisy signals was generated, with a mean of 12 mg/dL (0.7 mM) for the noise standard deviation, and a standard deviation of 5 mg/dL (0.3 mM). This way the




FIG. 1. Sample of the simulated glucose profile data. *Top panel:* Baseline signal, which has a maximum rate of change of 3 mg/dL/min. *Bottom panel:* Three of the noisy signals from the full data set ($n = 100$), showing the variation in noise levels among the signals.

data set contains signals with varying levels of noise, including some extreme cases.

From basic system theory, we know that in order to capture enough of a signal to be able to reconstruct it from the discrete measurements, the sampling frequency must be at least twice that of the highest frequency contained in the signal (this is the Nyquist–Shannon sampling theorem^{18,19}). To illustrate this, Figure 2 shows two sinusoids of different frequencies, both sampled at the same intervals. It is clear that we cannot distinguish between these two signals with this sampling. The other issue related to sampling is resolution; we need to be able to pick up significant changes in glucose concentrations in a reasonable time frame. If we sample at long intervals (e.g., every 20 min) we could miss significant changes in the trend;

fewer measurements would also limit the ability of the Kalman filter to determine if the change is due to noise or an actual change in glucose concentration.

The other factor is the prediction horizon—how far into the future are we estimating what the glucose concentration will be if the current trend persists. It is intuitive that the shorter the prediction horizon, the better our prediction will be, as it is less likely that the trend will change over a shorter period of time. Figure 3 illustrates this; using the noiseless sinusoid and a 1-min sampling period, it shows a comparison of predictions made over different horizons. It is clear that shorter is better.

The previous illustrations show the case when we have no noise, and we can analytically solve for the different parameters from the




FIG. 2. Two sinusoids, sampled with the sample period, have the same values; one has a period that does not satisfy the Nyquist sampling theorem (dashed line). At this sampling frequency for the solid sinusoid, it is possible to reconstruct the signal, but there is a lot that happens in between samples, so although the sampling period satisfies the Nyquist condition, it is not good for prediction and detection of blood glucose.




FIG. 3. Comparing prediction horizons: (a) illustration of predicting into the future using only the current value and the rate of change at $t = 120$ min, for various prediction horizons and a noiseless case; (b) the predicted value over the whole signal for predictions using the rate-of-change of the rate-of-change. Note that shorter prediction horizons would be almost on top of the actual values in this perfect estimate and noiseless case; this uses a 1-min sampling period.

equation of the sinusoid. So the question now is, what happens when the signal is not known to the estimator/predictor in advance, and also contains noise that is obscuring the actual glucose concentration values?

In order to keep the simulations consistent, the estimator was tuned for each of the selected sampling times (from 1 min to 30 min) over the full data set, minimizing the norm of the difference between the estimated signal and the true underlying baseline sinusoid (Fig. 4 shows a sample of an estimated signal from the noisy measurements). The estimator is then run on the noisy data, and used to predict future glucose values. The prediction is updated every minute, using an up-sampling technique that keeps the rate-of-change of the rate-of-change constant in between the actual sample times. This was done for prediction horizons of 1 min up to 30 min into the future.

Using the corresponding predicted values, a decision is made with the set glucose threshold. If the predicted value is below this threshold, then a suitable alarm can be set off. Data were generated for hypoglycemic thresholds from 40 mg/dL (2.2 mM) up to 100 mg/dL (5.6 mM) in 5 mg/dL (0.3 mM) increments. These can then be compared with the actual underlying signal to determine the sensitivity (defined as the percentage of true hypoglycemic

measurements correctly predicted out of the total number of true hypoglycemic instances) and specificity (defined as the percentage of correctly identified non-hypoglycemic measurements; this is the complement of the percentage of false-positives). For the noiseless case the sensitivity and specificity at a sampling period of 1 min are 98% across all prediction horizons and thresholds, and 99.9% for prediction horizons of 5 min or less.

Sensitivity and specificity will change depending on the sampling period used, as well as the prediction horizon and threshold selected. In Figure 5a and b, respectively, the sensitivity and specificity are shown as a function of prediction horizon for the case of a sampling period of 5 min and a hypoglycemia threshold of 70 mg/dL. It shows the results of running the estimator on the noiseless case, which serves as an indication of the upper bound that could be achieved in practice. For the full set of noisy data, the mean with the corresponding error bars (1 standard deviation) is also shown. We lose both sensitivity and specificity the further into the future we try to predict (longer prediction horizon). These plots correspond to slices of the plots in Figure 6c and d.

Figure 6 shows the mean sensitivity and specificity as a function of prediction horizon and hypoglycemic threshold for sampling pe-




FIG. 4. Sample of the simulated glucose profile data with the corresponding Kalman filter estimate.

riods of 1, 5, and 20 min. The transparent meshes correspond to the results when the estimator/predictor is used on the noiseless case; this corresponds to an upper limit on the performance that can be expected for the given conditions. The colored meshes correspond to the mean of the sensitivity and specificity (as the case might be) across the full data set, with the color indicating the standard deviation.

There are several things to note in these plots. For the case of a sampling period of 1 min (Fig. 6a and b), the performance is good; even with noise, the sensitivity and specificity obtained are exceptional, even with long prediction horizons. As we increase the sampling period, there is a steady deterioration. Take the case of the 5-min sampling period (Fig. 6c and d); the specificity that can be achieved, even for the noiseless case, decreases significantly, with some degradation in the sensitivity as well.

Note also that now, as the prediction horizon increases there is also a loss of performance, even though the other factors remain constant. Going to a longer sampling period of 20 min shows an even more pronounced degradation (Fig. 6e and f); longer prediction horizons suffer even more than with the 5-min sampling period.

It is interesting to note that changing the threshold does not change things much; only at the lower end is there noticeable change, but this is more due to the fact that the sine wave reaches its lowest value at 40 mg/dL and changes direction at that point. The above analysis is done with the sensitivity and specificity calculated with respect to the selected threshold, rather than the fixed threshold used to define hypoglycemia. This is useful in showing how the predictor performs in relation to the changing signal and different noise levels, but




FIG. 5. Sensitivity (a) and specificity (b) for different prediction horizons using a sampling period of 5 min and a hypoglycemia threshold of 70 mg/dL. The red dash-dot line is the noiseless case using the estimator, and the blue dashed line is the mean for the full noisy data set, with the error bars corresponding to the standard deviation. As the predictions are made further into the future, the less reliable they become; thus shorter prediction horizons give us better performance.

says nothing as to what the trade-off would be in selecting a different threshold to trigger an alarm with respect to “actual” hypoglycemia. We discuss this in the next section.

RESPONDING TO HYPOGLYCEMIA PREDICTIONS

The current methods used for hypoglycemia prediction use a single alarm level; that is, there will be hypoglycemia or there will not be, and this is done at a fixed prediction horizon. Based on the scatter plot (Fig. 7) and the corresponding statistics (Fig. 8) we propose a multitiered alarm system. Using the 20-min horizon, we can issue a mild warning, more for information purposes than a call for corrective action. A shorter prediction horizon of 10 min can be used for a higher level of warning, asking the user to keep close watch. If the 5-min prediction detects hypoglycemia, then we can issue an alarm requesting corrective action.

Related to the blood glucose concentration threshold used to define hypoglycemia is the issue of patient tolerance to false alarms. Currently, the consensus seems to be that going below the specified threshold is something that has to be avoided at all costs. While this makes sense with current treatment where action is

taken when symptoms appear, this might not necessarily be the same when predicting hypoglycemia.

In our proposed method, there are two parameters that can be modified that will affect sensitivity and specificity besides the sampling period. One is the more generally assumed, of the detection threshold; the other is the prediction horizon. In combination, together with the minimal risk (if any) of having the patient go below 70 mg/dL by 5–10 mg/dL, we suggest that the detection threshold should be kept the same as the threshold defining hypoglycemia (as long as this is not changed to coincide with the definition of appearance of symptoms). Therefore, use only the prediction horizon as the element that will be adjusted by the patient (and/or their physician) to balance their comfort versus the sensitivity and specificity obtained.

CONCLUSIONS

We have shown a Kalman filter implementation to estimate the first and second rates of change of an ISF glucose concentration measurement. Using these estimates, we can carry out predictions of future values of ISF glucose. Applying this method to simulated data, we




FIG. 6. Specificity and sensitivity as a function of prediction horizon and threshold, for sampling periods of 1 (a and b), 5 (c and d), and 20 min (e and f). The clear meshes are for the noiseless case using the estimator (i.e., signal is unknown to the predictor).

have shown how the sensitivity and specificity of the predictions change as a function of sampling period and prediction horizon.

Proper selection of sampling periods and multitiered alarm levels can provide a hypoglycemic alarm with high specificity and sensitivity, saving the user the aggravation of false alarms and the danger of missing true events. Optimal estimation is the key component in being able to accomplish performance that approaches that of an ideal noiseless situation.

APPENDIX

Underlying the Kalman filter is the discrete-time model

$$x_{k+1} = \Phi x_k + \Gamma u_k + \Gamma^w w_k \quad (8a)$$

$$y_k = C x_k + v_k \quad (8b)$$

where w_k and v_k are, respectively, the input and measurement noise vectors. They have covariance matrices Q and R (respectively), which in practice are not known, so they become tuning parameters. If the measurements have a high degree of uncertainty and we want to trust the model more than the measurement, then R is weighted more relative to Q . Conversely, weighting Q high relative to R implies that the measurements are to be trusted more than the model predictions.

The predictor-corrector equations of the estimator are

$$\hat{x}_{k|k-1} = \Phi \hat{x}_{k-1|k-1} + \Gamma u_k \quad (9a)$$

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + L_k (y_k - C \hat{x}_{k|k-1}) \quad (9b)$$

where $\hat{x}$ is the estimate of the states. The subscript notation $k|k-1$ indicates the estimate at time step k using measurements up to and including time step $k-1$. Equation 16 is used to propagate the state estimate (using the model) from the most recent time step $k-1$ to the current time step k . Equation 17 is then used to update the estimate, where L_k is the Kalman gain.

This Kalman filter gain is found by solving the equations

$$P_k = \Phi P_{k-1} \Phi^T + \Gamma^w Q \Gamma^{wT} - \Phi P_{k-1} C^T (C P_{k-1} C^T + R)^{-1} C P_{k-1} \Phi^T \quad (10a)$$

$$L_k = P_k C^T (C P_k C^T + R)^{-1} \quad (10b)$$

using an initial state covariance P_0 . A common assumption is that the Kalman gain is constant; it can be found by solving the non-linear steady-state Riccati equation

$$P_\infty = \Phi [P_\infty - P_\infty C^T (C P_\infty C^T + R)^{-1} C P_\infty] \Phi^T + \Gamma^w Q \Gamma^{wT} \quad (11)$$

In this paper Eq. 11 is solved using routines in MATLAB (The MathWorks, Inc., Natick, MA).

In our formulation, the rate-of-change of the rate-of-change is assumed to vary in a random fashion. As this changes, it affects the rate-of-change of ISF glucose and thus ISF glucose as well. We can combine Eqs. 2, 3, and 4 together with Eq. 1 by using vector notation. Then,

$$\begin{bmatrix} g_{k+1} \\ d_{k+1} \\ f_{k+1} \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} g_k \\ d_k \\ f_k \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} w_k \quad (12a)$$

$$y_k = [1 \quad 0 \quad 0] \begin{bmatrix} g_k \\ d_k \\ f_k \end{bmatrix} + v_k \quad (12b)$$

or more compactly, by defining the state-vector

$$x_k = \begin{bmatrix} g_k \\ d_k \\ f_k \end{bmatrix} \quad (13)$$

and by comparing to the standard notation for the process model, we can identify

$$\Phi = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \quad (14a)$$

$$\Gamma^w = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \quad (14b)$$

$$C = [1 \quad 0 \quad 0] \quad (14c)$$




FIG. 7. Scatter plot of predicted versus actual glucose for prediction horizons of 5, 10, and 20 min; in the ideal case all would fall on the diagonal dashed line. The sampling period is 1 min. s.c., subcutaneous.




FIG. 8. The statistics over the full data set for the case of a 1-min sampling period. Sensitivity (a) and specificity (b) improve with shorter prediction horizons. Box plots show the median and spread, while the circle is the mean and the solid vertical lines the standard deviation.

so we have

$$x_{k+1} = \Phi x_k + \Gamma^w w_k \quad (15a)$$

$$y_k = Cx_k + v_k \quad (15b)$$

Whenever we have a new measurement, any change can be due to the process noise (i.e., a change in w_k), or to the measurement noise v_k . This is where the Kalman filter comes in to do a trade-off between these two possibilities in an optimal fashion.

The Kalman filter uses a two-step process. It first calculates the estimate of the state-vector using the model (estimates are denoted by the $\hat{\cdot}$ symbol), using the known information up to that time, so at time step k we have

$$\hat{x}_{k|k-1} = \Phi \hat{x}_{k-1|k-1} + \Gamma u_k \quad (16)$$

where $\hat{x}_{k|k-1}$ is the estimate at time k using information up to time step $k-1$, and $\hat{x}_{k-1|k-1}$ is the estimate at time $k-1$ using all information up to time step $k-1$.

The second step is to correct this estimate, by incorporating the newest measurement at time k :

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + L(y_k - C\hat{x}_{k|k-1}) \quad (17)$$

where L is the steady-state Kalman gain, and $(y_k - C\hat{x}_{k|k-1})$ is the difference between the measured output and the expected output using the estimate from Eq. 16. Solving for the constant Kalman gain, we can substitute it together with Eq. 16, into Eq. 17; this gives us Eq. 5, 6, and 7 as presented above.

REFERENCES

1. The Diabetes Control and Complications Trials Research Group: The effect of intensive treatment of diabetes on the development and progression of long-term complications in insulin-dependent diabetes mellitus. *N Engl J Med* 1993;329:977-986.
2. UK Prospective Diabetes Study Group: Intensive blood-glucose control with sulphonylureas or insulin compared with conventional treatment and risk of complications in patients with type 2 diabetes (UKPDS 33). *Lancet* 1998;352:837-853.
3. Khaw K, Wareham N, Luben R, Bingham S, Oakes S, Welch A, Day N: Glycated haemoglobin, diabetes, and mortality in men in Norfolk cohort of European Prospective Investigation of Cancer and Nutrition (EPIC-Norfolk). *BMJ* 2001;322:15-18.
4. American Diabetes Association: Implications of the Diabetes Control and Complications Trial. *Diabetes Care* 2002;25(Suppl 1):S25-S27.
5. Cryer PE: Hypoglycemia: the limiting factor in the glycaemic management of type I and type II diabetes. *Diabetologia* 2002;45:937-948.
6. Dagogo-Jack S: Hypoglycemia in type 1 diabetes mellitus—pathophysiology and prevention. *Treat Endocrinol* 2004;3:91-103.
7. Bolli GB: Treatment and prevention of hypoglycemia and its unawareness in type 1 diabetes mellitus. *Rev Endocrinol Metab Disord* 2003;4:335-341.
8. Choleau C, Dokladal P, Klein J, Ward WK, Wilson GS, Reach G: Prevention of hypoglycemia using risk assessment with a continuous glucose monitoring system. *Diabetes* 2002;51:3263-3274.
9. Heise T, Koschinsky T, Heinemann L, Lodwig V: Hypoglycemia warning signal and glucose sensors: requirements and concepts. *Diabetes Technol Ther* 2003;5:563-571.
10. Ward WK: The role of new technology in the early detection of hypoglycemia. *Diabetes Technol Ther* 2004;6:115-117.
11. Bode B, Gross K, Rikalo N, Schwartz S, Wahl T, Page C, Gross T, Mastrototaro J: Alarms based on real-time sensor glucose values alert patients to hypo- and hyperglycemia: the Guardian continuous monitoring system. *Diabetes Technol Ther* 2004;6:105-113.
12. Bequette BW: Optimal estimation applications to continuous glucose monitoring. In: *Proceedings of the American Control Conference*. Boston: IEEE, 2004: 958-962.
13. Knobbe EJ, Buckingham B: The extended Kalman filter for continuous glucose monitoring. *Diabetes Technol Ther* 2005;7:15-27.
14. Stengel RF: *Optimal Control and Estimation*. New York: Dover Publications, 1994.
15. Gregory JW, Aynsley-Green A: The definition of hypoglycemia. *Baillieres Clin Endocrinol Metab* 1993; 7:587-590.
16. Heinemann L, and Anderson JH Jr: Measurement of insulin absorption and insulin action. *Diabetes Technol Ther* 2004;6:698-718.
17. Tamada J, Eastman R, Desai S, Wei C, Lesho M: Hypoglycemia detection and prediction. Presented at The Third Diabetes Technology Meeting, San Francisco, November 2003.
18. Nyquist H: Certain topics in telegraph transmission theory. *Trans AIEE* 1928;47:617-644.
19. Shannon CE: Communication in the presence of noise. *Proc Inst Radio Engineers* 1949;37:10-21.

Address reprint requests to:

B. Wayne Bequette, Ph.D.

Department of Chemical

and Biological Engineering

Rennselaer Polytechnic Institute

110 8th Street

Troy, NY 12180-3590

E-mail: bequette@rpi.edu