Proceedings of the 19th World Congress
The International Federation of Automatic Control
Cape Town, South Africa. August 24-29, 2014

Sensitivity Analysis of a Predictive Pump Suspension System to Treat People with Type 1 Diabetes

Fraser Cameron*, Bruce A. Buckingham, Darrell M. Wilson**, B. Wayne Bequette***

*Rensselaer Polytechnic Institute, Troy, NY 12180 USA (e-mail: fmccamer@gmail.com)
**Stanford University, Stanford, CA 94305 USA

Abstract: The primary goal of a low glucose suspend system is to reduce the risk of overnight hypoglycemia (low blood glucose) in individuals with type 1 diabetes by reducing/suspending insulin infusion. We have developed a Kalman filter-based algorithm, combined with a number of safety rules, to implement a predictive low glucose suspend system that shuts off an insulin pump based on a prediction of hypoglycemia 30-70 minutes in the future. This system has been studied in over 2,000 nights in an outpatient-home environment. In this paper, based on an analysis of this data, we isolate the effects of the individual rules in part by simulating their removal from the existing data. Specifically, we decompose the basal insulin into small boluses and, using a model of insulin pharmacodynamic action (the time effect of insulin on blood glucose), alter the real data corresponding to the addition or removal of basal insulin via simulation. Our results show that limiting the total suspension to 180 minutes per night prevents excessive suspension in cases where the average calibration is an excessive 58 mg/dl, above the mean of 18 mg/dl. Further, we also show that a simple threshold algorithm that suspends below 100 mg/dl if the glucose level is flat or falling, is comparable in performance. Lastly, we show that the Kalman filter at the heart of this algorithm reduces the time spent below 70 mg/dl by 50% at the expense of a mean rise of 12 mg/dl in morning glucose levels.

Keywords: **Biomedical systems, Biomedical control, Pump Suspension, Diabetes, Artificial Pancreas**

1. INTRODUCTION

Type 1 diabetes (T1D) is an autoimmune disease that directly destroys the body's ability to produce insulin and indirectly the body's ability to regulate blood glucose concentrations. Individuals with T1D must use either multiple daily injections of insulin (one bolus of long-acting insulin each day, and boluses of rapid-acting insulin at meal/snack time and when blood glucose needs to be reduced), or continuous infusion of rapid-acting insulin using an insulin pump. While insulin therapy can lead to lower blood glucose levels and reduce the risk of complications due to high blood glucose, there is an underlying risk of hypoglycemia (low blood glucose) if insulin is over-administrated. Indeed, one of the greatest fears of a parent of a child with T1D is extended overnight hypoglycemia, which could lead to a coma or, in rare cases, death (known as "dead in bed" syndrome). The development of continuous glucose monitoring (CGM) technology, allowing a near continuous measurement of glucose levels, enabled the use of alarms to warn individuals of low (or high) glucose levels. Unfortunately, these alarms have been found to be insufficient, since individuals and their caregivers often sleep through the alarms (Buckingham et al., 2005).

Low glucose suspend (LGS) or pump shut-off (PSO) systems have been developed specifically to shut-off pumps to reduce the risk of hypoglycemia, based on real-time CGM signals; they are also a natural first step towards the development of a fully closed-loop artificial pancreas (Kowalski, 2009; Harvey et al., 2010; Cobelli et al., 2011; Bequette, 2012). Initial LGS systems were threshold-based (the pump is turned off when the threshold is violated), while much current effort has been on predictive low glucose suspend (PLGS) systems that turn-off a pump when a hypoglycemic event is predicted to occur (usually 30-70 minutes in the future; Bequette, 2014).

The PLGS algorithm that we have developed involves the use of a Kalman filter predict future glucose values, combined with a set of rules to reduce the risk of prolonged periods of pump shut-off. This algorithm was first tested in in-clinic studies (Cameron et al., 2012), followed by extensive out-patient (in-home) studies (Maahs et al., 2014). The objective of this paper is to analyze the results from over 2,000 nights of out-patient studies to understand the effect of various algorithm parameters and rules on the blood glucose control. We first review the results of the outpatient study, then describe our hybrid experiment/simulation approach, and finally discuss the results.

2. OUTPATIENT STUDY DATA

The data used in this paper comes from an outpatient clinical trial of the described pump suspension algorithm. The idea is that by predicting impending hypoglycemia and suspending insulin delivery that the body's natural release of glucose into

the blood stream would mitigate or prevent the hypoglycemia. For each of the 45 patients there was a run-in phase where the algorithm was enabled each night and then a 42-night phase where the algorithm was randomly and blindly enabled or disabled. This design isolated the effectiveness of the algorithm from any behavioural adjustments the patients might make for a final system. This resulted in 925 control nights and 1125 intervention nights.

The system consisted of a continuous glucose monitor (CGM) and an insulin pump communicating with a bedside laptop computer that contained the pump suspension algorithm. Each of the nights has at least 4 hours of CGM data post activation, a morning blood glucose measurement, records of any snacks from bedtime until morning, records of exercise in the previous day, basic demographics, and a full insulin history.

A typical intervention night is shown in Fig. 1. The x's indicate the glucose concentration as measured by the continuous glucose monitor and provided to the algorithm. When these trend downwards at the start of the dataset the algorithm triggers suspensions, as indicated by the orange triangles and the zeroing out of the basal rate. Later in the evening two periods of sensor noise also trigger pump suspensions. Eliminating those is the subject of further study.

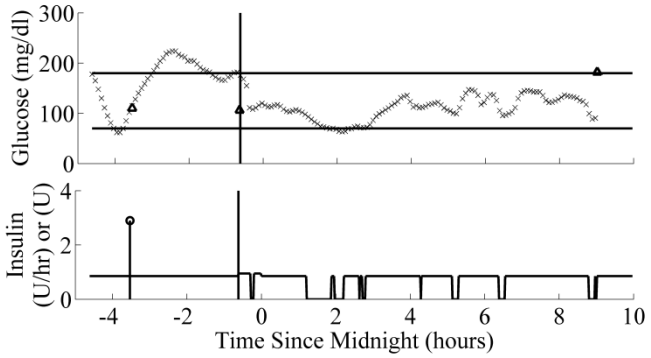

Fig 1. Sample Intervention Night. The top plot shows the CGM values (x), reference/calibration values (triangle), and bounds of desirable glucose values (horizontal line). The bottom plot shows the basal insulin (line) and boluses (bubble). The vertical line is the time when the system was activated.

3. SIMULATOR

Running a separate clinical trial to test the effect of a change to or removal of a rule would be prohibitive in terms of cost, and noisy due to inter- and intra-patient variability. Instead we simulate from the existing data. Specifically, we assume that the patients' insulin sensitivity can be calculated according to the 1800-rule:

$$IS = \frac{1800}{TDD}$$

which is a common heuristic used in clinical practice. TDD is the patients' total daily insulin dose in units/day. Then we use an average of published insulin time action profiles (Frohnhauer et al., 2001; Heinemann and Steiner, 1997; Swan et al., 2009) shown in Fig. 2, to approximate the effect of any insulin subtracted or added. The multiplication of $IS$with the curve in Fig. 2 represents the convolution model of insulin action, $\vec{I}$. Given a glucose and insulin profile represented by glucose values $\vec{g}$ and $\vec{u}$ at regular time intervals the simulator morphs the original profile into a simulated ones by looping over the following steps. 1) $u^* = f(u_{1\ldots k}, g_{1\ldots k})$ where u_i is the i$^{\text{th}}$ element of the vector $\vec{u}$ and $f(\vec{i}, \vec{y})$ is a controller that determines what new input to command given the past history of inputs $\vec{i}$ and outputs $\vec{y}$. 2) $\Delta u = u^* - u_k$ 3) $u_k = u^*$ 4) $g_{k+1\ldots \inf} = g_{k+1\ldots \inf} + \vec{I}\Delta u$ which simulates the effect of the changed insulin administration and 5) $k = k + 1$. After each repetition of the above steps the $\vec{g}$ and $\vec{u}$ vectors represent the simulated value of glucose $\vec{g}$ corresponding to the provided insulin $\vec{u}$. This simple explanation ignores issues corresponding to vector lengths and missing glucose readings that can be easily fixed in practice.

This simulator only makes assumptions about insulin action. It does not make any assumptions about meals, exercise, sleep, or anything else. Consequently, the inaccuracy of the simulator stems only from estimating the effects of large changes to the administered insulin. For the vast majority of simulated cases the glucose levels are changing only within $\pm$ 10 mg/dl, a range for which the assumption of locally linear insulin action likely holds.

An example simulation for the base algorithm and one where the prediction horizon is extended from 30 to 70 min is shown in Fig. 3. Here, the trial night begins at the vertical black line. The blue is the simulated closed-loop insulin delivery and resultant glucose values. Increasing the prediction horizon leads to earlier suspension and so higher glucose values and less hypoglycemia.

Because we can get negative values in simulation when removing particularly important rules, we use a modified risk measure that is fitted to the Kovatchev risk profile, but that allows for negative values (Cameron, 2010; Cameron et al., 2011).

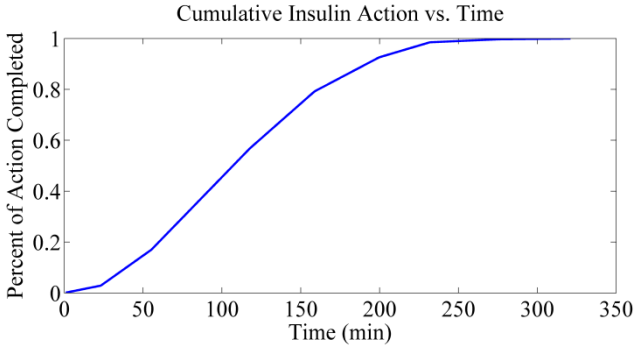

Fig. 2. Cumulative Insulin Action vs. Time