4-13-2015

Towards a Model-Based Meal Detector for Type I Diabetics

Sanjian Chen
University of Pennsylvania, sanjian@seas.upenn.edu

James Weimer
University of Pennsylvania, weimerj@seas.upenn.edu

Michael R. Rickels
University of Pennsylvania, rickels@mail.med.upenn.edu

Amy Peleckis
University of Pennsylvania, amy.peleckis@uphs.upenn.edu

Insup Lee
University of Pennsylvania, lee@cis.upenn.edu

6th Workshop on Medical Cyber-Physical Systems (MedicalCPS 2015) http://workshop.medcps.org/ in conjunction with CPS Week 2015 http://www.cpsweek.org/2015/ Seattle, WA, April 13, 2015

Towards a Model-Based Meal Detector for Type I Diabetics

Abstract
Blood glucose management systems are an important class of Medical Cyber-Physical Systems that provide vital everyday decision support service to diabetics. An artificial pancreas, which integrates a continuous glucose monitor, a wearable insulin pump, and control algorithms running on embedded computing devices, can significantly improve the quality of life for millions of Type 1 diabetics. A primary problem in the development of an artificial pancreas is the accurate detection and estimation of meal carbohydrates, which cause significant glucose system disturbances. Meal carbohydrate detection is challenging since post-meal glucose responses greatly depend on patient-specific physiology and meal composition.

In this paper, we develop a novel meal-time detector that leverages a linearized physiological model to realize a (nearly) constant false alarm rate (CFAR) performance despite unknown model parameters and uncertain meal inputs. *Insilico* evaluations using 10, 000 virtual subjects on an FDA-accepted maximal physiological model illustrate that the proposed CFAR meal detector significantly outperforms a current state-of-the-art meal detector that utilizes a voting scheme based on rate-of-change (RoC) measures. The proposed detector achieves 99.6% correct detection rate while averaging one false alarm every 24 days (a 1.4% false alarm rate), which represents an 84% reduction in false alarms and a 95% reduction in missed alarms when compared to the RoC approach.

Keywords
Time series analysis, Medical Information Systems

Disciplines
Computer Engineering | Computer Sciences

Comments
6th Workshop on Medical Cyber-Physical Systems (MedicalCPS 2015) http://workshop.medcps.org/ in conjunction with CPS Week 2015 http://www.cpsweek.org/2015/ Seattle, WA, April 13, 2015

Towards a Model-Based Meal Detector for Type I Diabetics*

Sanjian Chen James Weimer
Dept. of Computer & Information Science
University of Pennsylvania
{sanjian, weimerj}@seas.upenn.edu
Michael R. Rickels Amy Peleckis
Division of Endocrinology, Diabetes & Metabolism
University of Pennsylvania Perelman School of Medicine
rickels@mail.med.upenn.edu,amy.peleckis@uphs.upenn.edu
Insup Lee
Dept. of Computer & Information Science
University of Pennsylvania
lee@cis.upenn.edu

ABSTRACT

Blood glucose management systems are an important class of Medical Cyber-Physical Systems that provide vital everyday decision support service to diabetics. An artificial pancreas, which integrates a continuous glucose monitor, a wearable insulin pump, and control algorithms running on embedded computing devices, can significantly improve the quality of life for millions of Type 1 diabetics. A primary problem in the development of an artificial pancreas is the accurate detection and estimation of meal carbohydrates, which cause significant glucose system disturbances. Meal carbohydrate detection is challenging since post-meal glucose responses greatly depend on patient-specific physiology and meal composition.

In this paper, we develop a novel meal-time detector that leverages a linearized physiological model to realize a (nearly) constant false alarm rate (CFAR) performance despite unknown model parameters and uncertain meal inputs. *In-silico* evaluations using $10,000$ virtual subjects on an FDA-accepted maximal physiological model illustrate that the proposed CFAR meal detector significantly outperforms a current state-of-the-art meal detector that utilizes a voting scheme based on rate-of-change (RoC) measures. The proposed detector achieves 99.6% correct detection rate while averaging one false alarm every 24 days (a 1.4% false alarm rate), which represents an 84% reduction in false alarms and a 95% reduction in missed alarms when compared to the RoC approach. [1]

Categories and Subject Descriptors

G.3 [**Probability and Statistics**]: Time series analysis; J.3 [**Life and Medical Sciences**]: Medical Information Systems

1. INTRODUCTION

Approximately 29 million people in the U.S. have diabetes, among which about 5% have Type 1 diabetes [14], an auto-immune disease that destroys a person's pancreas' ability to release insulin. Type 1 diabetics depend on everyday insulin infusion or injection to maintain their glucose level within the acceptable range where too much insulin can cause life-threatening hypoglycemia (extremely low glucose level) and too little insulin can cause nerve-damaging hyperglycemia (high glucose level) [2]. Unfortunately, meal carbohydrates are a major disturbance to one's blood glucose level, and therefore every Type 1 diabetic faces a lifelong control challenge: he/she has to carefully titrate insulin doses for every meal so that post-meal hyperglycemia is effectively controlled without risking hypoglycemia.

In recent years, Continuous Glucose Monitoring (CGM) technology has become more popular, which drives a whole class of Medical Cyber-Physical System (MCPS), most notably the *artificial pancreas* (AP), that aims to facilitate glucose management for Type 1 diabetics. At the AP system's core are a CGM sensor, a wearable insulin pump for insulin infusion and boluses, and algorithms that control insulin titration [8]. Reliably predicting meals is difficult in real-life situations, thus all AP systems depend on certain kinds of meal declaration/detection mechanisms. Meal detection is a safety critical problem, where an incorrectly identified meal may trigger the system to either deliver too much insulin unnecessarily or deliver too little insulin, both of which have harmful (if not deadly) consequences.

Currently, most Type 1 diabetics who use CGM sensors and wearable insulin pumps manually input the time and estimated carb count of each meal into the device, which then calculates a suggested insulin dose. Unfortunately, self-reported meal information is inherently unreliable [12]. Thus, more dependable meal detection methods are necessary to ensure patient safety. This paper proposes a novel meal detection method that leverages a linear model of glu-

*This research was supported in part by NSF CNS-1035715, NSF IIS-1231680, Public Health Services research grant R01-DK091331 (to M.R.R.), and in part by the DGIST Research and Development Program of the Ministry of Science, ICT and Future Planning of Korea (CPS Global Center). The authors would like to acknowledge support from the Human Metabolism Resource of the Institute for Diabetes, Obesity Metabolism at Penn.

cose and insulin responses that is inspired by a first-principle minimal physiological model [3].

The proposed detector exhibits consistent detection performance in *in-silico* evaluations using an FDA-accepted Type 1 diabetes physiological model across a wide range of virtual subject parameter settings. Specifically, it achieves a 99.6% detection rate (one miss every 90 days) with a low false alarm rate of once every 24 days. This is a significant improvement compared to a state-of-the-art voting algorithm that detects meals based on glucose rate-of-change (RoC) [12]: in the same *in-silico* experiment, the voting algorithm detects 93.2% of meals (one miss every 5 days) while raising a false alarm every 4 days.

The remainder of this paper is organized as follows: Section 2 reviews related work; Section 3 gives the problem formulation; Section 4 presents two major classes of glucose-insulin physiological models; Section 5 introduces the Constant False Alarm Rate (CFAR) meal detector; Section 6 compares the performance of our detector with a state-of-the-art meal detector in an *in-silico* evaluation using an FDA-accepted physiological model.

2. RELATED WORK

Input detection via hypothesis testing is a mature research area where approaches can be roughly classified as model-driven [6, 36] or data-driven [4, 10, 23, 26, 27, 28]. When applied to MCPS problems, model-driven approaches (e.g., likelihood ratio testing) suffer due to modeling nonlinearities and unknowns. In contrast, data-driven MCPS approaches (e.g., machine learning classification) are considered promising [1, 20], but suffer from over-fitting when accurately annotated real-world training data is scarce.

Specific to the problem of meal detection and estimation, various algorithms have been proposed [12, 5, 19]. Most existing algorithms detect meals based on the RoC of glucose values, more precisely, first and/or second derivatives derived from CGM readings [12, 19]. The RoC metrics are estimated in real-time and are either directly compared to pre-tuned thresholds to make a decision, or used to estimate meal impulses. One of the most prominent RoC meal detectors is described in [12]. It utilizes a voting algorithm to detect meal disturbances using combinations of glucose derivatives and Kalman filtering. A pilot evaluation on real patients using the RoC detector suggests an average delay of 30 minutes between the onset of a meal and detection [12].

Although RoC-based meal detectors are easy to understand and implement, they have some fundamental limitations considering everyday diabetes management. The post-meal glucose response of a Type 1 diabetic depends on many factors and can greatly vary across different patients. Most insulin-dependent Type 1 diabetics are instructed to take insulin boluses around meal times, where both the timing and the amount of meal boluses can change postprandial plasma glucose responses [9]. Moreover, the exact breakdown of protein, fat, and carbohydrates in a meal also impacts glucose responses [37]. Additionally, the complex interaction between glucose and insulin metabolism systems can trigger RoC-based meal detection which was actually a drop in "insulin on board" [13]. Therefore, threshold based RoC meal detectors are likely to generate numerous false alarms and can miss true meal events when the postprandial glucose rise is not as prominent.

Other approaches to meal detection use physiological models to generate a population of possible post-meal glucose shapes [5]. Based on a residual generated using the difference between a no-meal realization of the population models and the actual CGM measurements, near-perfect meal detection can be achieved using population-based approaches but requires significantly longer to make an accurate decision when compared to RoC-based approaches [12, 5, 19]. Moreover, this performance is dependent on an accurate estimate of patient specific parameters (e.g. insulin time action profile), which in real-world scenarios requires periodic controlled studies on each patient where the meal events are accurately logged.

An alternative approach to event detection that tends to minimize the detection time while maximizing performance in problems with scarce data and uncertain models is the CFAR detector [30, 32]. CFAR detectors leverage parameterized model structure to generate sufficient statistics that are maximally invariant to unknown parameters such that a threshold test achieves a (nearly) constant false alarm rate [30]. Although not previously applied to the diabetic meal detection problem, these detectors have been shown to work well in other CPS applications with structured dynamics and unknown parameters, specifically in detecting faults in networked systems [34, 35], building heating, ventilating and air conditioning (HVAC) systems [31], smart grids [33], and most recently in a MCPS application to detect critical pulmonary shunts in infants [16].

3. PROBLEM FORMULATION

All AP systems need accurate estimates of the meal carbohydrate disturbances. To estimate the meal carbohydrate disturbances requires an accurate (and timely) estimate of when meals occur. Current approaches to meal detection lack detection accuracy (RoC-based approaches) or lack timely detection guarantees (population-based approaches). Thus, the following summarizes our problem statement:

3.1 Problem Statement

This paper addresses the meal detection problem, where given recent glucose level measurements and insulin inputs, design a run-time monitor that accurately and quickly detects the ingestion of carbohydrates.

4. PHYSIOLOGICAL MODELS

First-principle models of glucose physiology broadly fall into two categories: maximal models and minimal models [7]. Maximal models use fine-grain compartmental submodels to describe the dynamics of glucose and insulin. These models are mostly used for simulation purposes, since controller design for non-linear models with unknown parameters is difficult. On the other hand, the minimal models use only a few coarse-grain compartments to model the physiology, and they have a simple structure that is convenient for linearization and control design [15]. This section introduces an FDA-accepted maximal model, which will later be used in *in-silico* evaluations, and a minimal model, which inspires the process modeling in designing our CFAR detector.

4.1 A Maximal Model

A maximal model that describes the glucose-insulin responses with meals has been proposed [21, 22], based on which the UVa/Padova Type 1 Diabetes Mellitus Metabolic

Simulator (T1DMS) has been developed [18]. This model is an FDA-accepted substitute for animal testing in pre-clinical trials when evaluating certain control algorithms [11]. It consists of a set of continuous-time differential equations with 13 state variables and 32 physiological parameters. The model includes three sub-systems: the insulin subsystem, the meal glucose absorption, and the glucose kinetics. Due to space constraints, this section only sketches the model equations with brief explanations of the state variables. Extensive details of the model, including the modeling rationale and meanings of the variables & parameters, can be found in a series of publications [21, 22, 18, 11].

The insulin sub-system describes the transportation of insulin from the subcutaneous injection site to other compartments of the body such as the liver, plasma, and tissues. This subsystem has seven state variables which evolve according to the follow equations [21, 22]:

$$\begin{aligned}\dot{I}_p(t) = &-(m_2+m_4)I_p(t) + m_1 I_l(t) + k_{a1}S_1(t) \\ &+ k_{a2}S_2(t)\end{aligned} \tag{1a}$$

$$\dot{X}(t) = P_{2U}/V_i I_p(t) - P_{2U}X(t) - P_{2U} * I_b \tag{1b}$$

$$\dot{I}_1(t) = k_i/V_i I_p(t) - k_i I_1(t) \tag{1c}$$

$$\dot{I}_d(t) = k_i I_1(t) - k_i I_d(t) \tag{1d}$$

$$\dot{I}_l(t) = m2 * I_p(t) - (m_1+m_3)I_l(t) \tag{1e}$$

$$\dot{S}_1(t) = -(k_{a1}+k_d)S_1(t) + u(t) \tag{1f}$$

$$\dot{S}_2(t) = k_d S_1(t) - k_{a2}S_2(t). \tag{1g}$$

In the above equations, I_p represents the mass of plasma insulin. $X(t)$ is a remote insulin signal that also appears in the glucose kinetics. I_1 and I_d represent a delayed insulin signal that governs the endogenous glucose production. I_l represents the liver insulin. S_1 and S_2 represent a two-compartment subcutaneous insulin process. $u(t)$ is the subcutaneous insulin injection/infusion input (wearable insulin pumps for Type 1 diabetics inject insulin into the subcutaneous tissue). Table 2 lists the physiological entities that the parameters represent[2]. Parameters listed as "rate parameters" in Table 2 are those simply govern the flux rate in/out of a compartment. Details about the parameters can be found in the maximal modeling literature [21, 22, 18].

The meal absorption sub-system models how meal carbohydrates pass the stomach, intestine and finally becomes glucose appearing in the plasma [22]. The stomach is represented by two compartments: one for the solid phase and the other for the liquid phase. The dynamics are modeled by the following equations [22]:

$$Q_{sto}(t) = Q_{sto1}(t) + Q_{sto2}(t) \tag{2a}$$

$$\dot{Q}_{sto1}(t) = -k_{gri}Q_{sto1}(t) + m(t) \tag{2b}$$

$$\dot{Q}_{sto2}(t) = -k_{empt}(Q_{sto})Q_{sto2}(t) + k_{gri}Q_{sto1}(t) \tag{2c}$$

$$\dot{Q}_{int}(t) = -k_{abs}Q_{int}(t) + k_{empt}(Q_{sto})Q_{sto2}(t) \tag{2d}$$

$$\begin{aligned}k_{empt}(Q_{sto}) = &\, k_{min} + (k_{max}-k_{min})/2 \\ &\times \begin{pmatrix} \tanh(\alpha * (Q_{sto} - b * D)) \\ -\tanh(\beta * (Q_{sto} - d * D)) + 2 \end{pmatrix}.\end{aligned} \tag{2e}$$

$$\alpha = 5/(2 * D * (1-b)), \quad \beta = 5/(2 * D * d) \tag{2f}$$

Q_{sto} is the amount of glucose in the stomach. Q_{sto1} and Q_{sto2} represent glucose in the solid phase and liquid phase, respectively. Q_{int} is the amount of glucose in the intestine. $k_{empt}(Q_{sto})$ is a non-linear function that represents the rate of carbohydrates emptying from the stomach. D is the total amount of ingested glucose in the last meal. $m(t)$ is the input of meal carbohydrates. The physiological meanings of parameters can be found in Table 2.

The insulin and absorbed meal glucose interact through the glucose kinetics, which is modeled by three state variables. The equations are given as follows [22]:

$$R_a(t) = f * k_{abs} * Q_{int}/BW \tag{3a}$$

$$\begin{aligned}\dot{G}_p(t) = &-k_1 * G_p(t) + k_2 * G_t(t) \\ &+ \max(0, k_{p1} - k_{p2} * G_p - k_{p3} * I_d(t)) \\ &- F_{snc} - \max(0, k_{e1} * (G_p(t) - k_{e2})) + R_a(t)\end{aligned} \tag{3b}$$

$$\begin{aligned}\dot{G}_t(t) = &-\frac{(V_{m0} + V_{mx} * X(t)) * G_t(t)}{K_{m0} + G_t(t)} \\ &+ k_1 * G_p(t) - k_2 * G_t(t)\end{aligned} \tag{3c}$$

$$\dot{G}_m(t) = -k_{sc} * G_m(t) + k_{sc}/V_g * G_p(t). \tag{3d}$$

G_p represents the plasma glucose concentration. G_t represents glucose in the rapidly equilibrating tissue. G_m represents the subcutaneous glucose. Parameters are explained in Table 2. Note that the insulin action on glucose is modeled by $X(t)$ and $I_d(t)$ appearing in the $\dot{G}_p(t)$ and $\dot{G}_t(t)$ equations, and the meal glucose rate of appearance $R_a(t)$ is calculated from the meal sub-system state Q_{int}.

The maximal model of the T1DMS consists of all the 13 differential equations presented above. The insulin subsystem is a linear model. The meal sub-system contains a non-linear parameter $k_{empt}(Q_{sto})$. The glucose kinetics sub-system has several non-linear terms, such as the max operators and state product $X(t)G_t(t)$. Most of the model parameters, as listed in Table 2, are not easily identifiable on patients. Because of the non-linearity and unknown parameters, it is very difficult to directly use the maximal model in control design. Instead, the model is used primarily for simulation purposes. The FDA accepted 300 virtual subjects, each of which is a realization of the entire parameter vector, that are sampled from a joint distribution. The parameter distribution was drawn from clinical experimental data obtained from individuals who underwent a triple tracer meal protocol [22].

4.2 Minimal Models

Minimal models represent another class of first-principle glucose-insulin models. The basic idea is to lump together compartments to describe the dominating dynamics of the glucose-related physiology in a minimal number of compartments. One of the most commonly accepted minimal models is described in [3], and referred to as the Bergman model. The Bergman model uses a single lumped compartment to model insulin and another lumped compartment to model glucose in plasma. Insulin governs the changes of glucose level either directly or through another remote compartment. Under this compartmentalization scheme, at most three state variables are needed describe the glucose-insulin physiology: plasma glucose, plasma insulin, and insulin in the remote compartment. Seven minimal models are proposed in [3], from the simple insulin-independent models (models No. 1 to 3), to the more elaborated forms (models No. 4 to 7). Not all the models use all of the three states.

[2]Due to space, Table 2 is located on the final page.

The Bergman model No. 4, whose equations are given as follows, has a linear form and explicitly describes insulin-dependent glucose uptake,

$$\dot{G}(t) = p_1 G(t) + p_2 * I(t) + p_3 \tag{4}$$

where $G(t)$ and $I(t)$ represent plasma glucose and insulin, respectively, and p_1, p_2, and p_3 are unspecified model parameters.

While the Bergman model describes the plasma glucose-insulin dynamics, a second order, two-compartment meal model is presented in [15]:

$$\begin{aligned} \dot{g}(t) &= -\frac{1}{t_G} g(t) + \frac{A_G}{t_G} D_G(t) \\ \dot{m}(t) &= -\frac{1}{t_G} m(t) + \frac{1}{t_G} g(t), \end{aligned} \tag{5}$$

where $g(t)$ represents glucose in the first compartment and $m(t)$ represents the plasma glucose appearance, which is an input to the Bergman model. $D_G(t)$ is the meal carbohydrate input. A_G is the carbohydrate bioavailability and t_G is the time of maximum glucose rate of appearance.

Lastly, the insulin pathway from subcutaneous tissue to plasma can be modeled by the following second order process [24, 17].

$$\begin{aligned} \dot{x}(t) &= -k_a x(t) + u(t-\tau) \\ \dot{I}(t) &= -k_e I(t) + \frac{k_a}{V_d} x(t), \end{aligned} \tag{6}$$

where $x(t)$ and $I(t)$ are insulin in the subcutaneous compartment and plasma, respectively, k_a and k_e are rate parameters, V_d is the insulin volume, and $u(t-\tau)$ represents the insulin input with a time delay τ.[3]

Combining equations 4, 5, and 6 results in a fifth-order linear model that describes the glucose-insulin kinetics given meal carbohydrate inputs and subcutaneous insulin inputs. This linear model will be used in the CFAR detector design.

5. CFAR MEAL DETECTOR DESIGN

In this section, we introduce the CFAR meal detector with a bin-counting heuristic which counts the number of decisions generated using dual parameter-invariant statistics evaluated at each time step. The remainder of this section details the dual parameter-invariant statistics and the bin-counting heuristic, respectively.

5.1 Parameter Invariant Design

The design of CFAR detectors utilizing parameter invariant statistics originates from statistical signal processing [29]. Our previous work [32] proposes a general framework of parameter invariant design in medical systems. In this paper, we apply the general methodology in the context of the meal detection problem.

Our primary goal lies in designing a detector (or monitor) that has consistent detection performance on a large population of Type 1 diabetics given a wide range of meal and insulin inputs. Because many of the patient-specific physiological parameters (e.g., those in the maximal models) cannot be identified, the detector is designed to be invariant to model parameters. In addition, the amount of meal carbohydrates is manually reported by users in current Type 1 diabetes management systems, which can be unreliable for a variety of reasons, e.g., the patient may forget to report or miscount the meal portion. Therefore, the detector is also designed to be invariant to the exact magnitude of meal inputs.[4] Insulin bolus times and doses are used by our detector since those are pump logged information.

We start with the linear minimal model in equations 4, 5, and 6. Following standard control theory techniques [25], the state-space linear model can be transformed into a z-domain representation and then written in a discrete time matrix form $\boldsymbol{y} = \boldsymbol{F\theta} + \boldsymbol{Gv} + \boldsymbol{\sigma n}$, in which $\boldsymbol{y}$ represents the outputs, $\boldsymbol{F}$ represents the process model, $\boldsymbol{\theta}$ represents model parameters, $\boldsymbol{G}$ represents input response, $\boldsymbol{v}$ represents inputs, and $\boldsymbol{\sigma n}$ are zero-mean Gaussian noises.

The CFAR detector runs in a sliding window fashion. It has a few critical time windows which are illustrated in Figure 1. At each time step, the detector is given a vector of T past measurements $\boldsymbol{y} \in \boldsymbol{R}^T$ and insulin bolus inputs within the time window wt. To be invariant to the unknown model parameters, the CFAR detector first eliminates the effects of parameters $\boldsymbol{\theta}$ by projecting $\boldsymbol{y}$ onto the null space of $\boldsymbol{F}$ so that the term $\boldsymbol{F\theta}$ becomes zero.

The core of the detector is a bi-directional meal hypothesis test. The detector first hypothesizes that a meal happened wd steps back from the current time (meals are treated as impulses in the detector model). wd is a detector parameter that stays constant at run-time once it is chosen. The null hypothesis H_0 states that a meal indeed happened in a time window around next to the hypothesized meal time (the H_0 window in Figure 1). The event hypothesis H_1 states that a meal actually happened in an even earlier time window (the H_1 window in Figure 1). The sizes of H_0 and H_1 windows are design parameters to be chosen. The input response matrices G_0 and G_1 represent the hypothesized meal time windows of H_0 and H_1, respectively. When testing H_0 in the direction of H_1, the detector eliminates the effect of H_0 input by projecting G_1 onto the null space of G_0. It then calculates a statistic $t_0(y)$ as a ratio, where the numerator represents the remaining energy in the glucose measurements explained by signal under H_1 and the denominator represents the energy that is not explained by H_1. Then the detector tests H_1 in the direction of H_0 and calculates $t_1(y)$ in exactly the same way but in the opposite direction.

The statistic $t_0(y)$ assumes H_0 is true and tests H_1. The statistic $t_1(y)$ assumes H_1 is true and tests H_0. To guarantee a minimum level of performance, the detector rejects H_0 when $t_0(y) > \eta_0$, where η_0 is related the probability of false alarms. And similarly, the detector rejects H_1 when $t_1(y) > \eta_1$, where η_1 is related to the probability of missed detection.

The detector makes a meal detection decision based on the bi-directional hypothesis tests. When H_1 is rejected and H_0 is not, it claims a meal happened in the H_0 zone. When H_0 is rejected and H_1 is not, it claims a meal happened in the H_1 zone. When both are not rejected, it means there is not enough power in the signal to make a decision. When both are rejected, it means there are residual energy in the measurements that are explained by both hypothesises' signals.

Figure 1 demonstrates how the CFAR detector works on simulated scenarios generated by the FDA-accepted maximal model. The CGM measurements are sampled at one minute time steps. The true meal happens around time 2160

[3]The results in this paper were obtained assuming $\tau = 0$. Studying the effect of non-zero τ values is planned as future work.

[4]In our future work, we will explore how to leverage an estimated range of the meal input for prediction and control purposes.

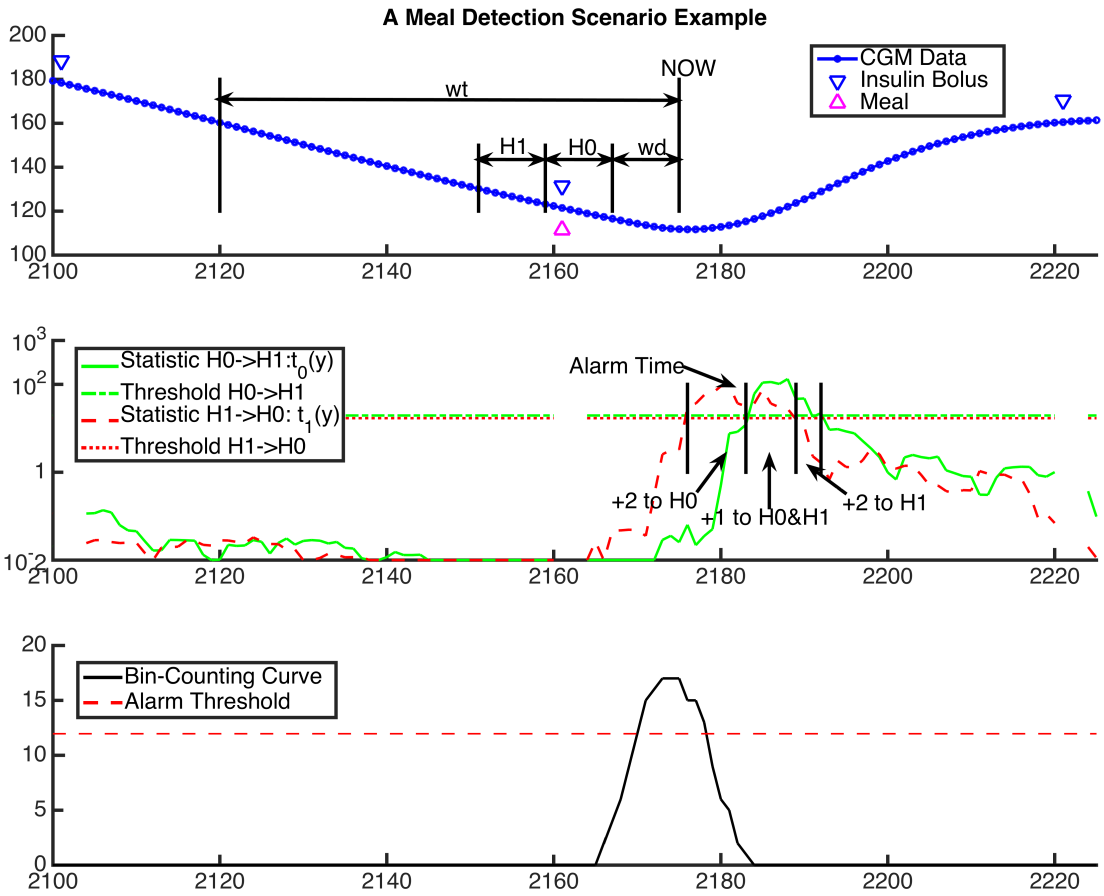

Figure 1: A meal detection example of the CFAR detector.

(the pink upper triangle in the figure). As the H_0 window approaches the true meal event (the detector never knows when a meal actually happened and runs hypothesis tests at every step), the statistic $t_1(y)$ (testing H_1 in the direction of H_0; the red dashed line in the figure) starts rising and becomes separated from $t_0(y)$. This indicates that H_1 is rejected and the detector claims H_0. Then as the detector moves further ahead, the true meal enters the H_1 window, and $t_0(y)$ starts rising and $t_1(y)$ starts falling, indicating that H_0 is rejected and the detector claims H_1.

5.2 Bin-Counting Meal Alarms

In the previous section, we develop a detector that generates a decision at each time step. At run-time, the detector runs in a sliding window fashion, with the relative positions of the windows fixed once the detector parameters are chosen (see Figure 1). As the detector approaches a true meal event (the ground truth meal times are unknown to the detector), it will first pass through the H_0 window and then the H_1 window. Therefore, one meal event will accumulate a few H_0 claims and then some H_1 claims as the detector windows slide through. We leverage this phenomenon and develop a bin-counting heuristic to generate more robust meal detection alarms.

Table 1: Credit addition rules for the bin-counting process.

	$t_0(\boldsymbol{y}) > \eta_0$	$t_0(\boldsymbol{y}) \leq \eta_0$
$t_1(\boldsymbol{y}) > \eta_1$	Both Rejecting (+1 to each)	Accept H_0 (+2 to H_0)
$t_1(\boldsymbol{y}) \leq \eta_1$	Accept H_1 (+2 to H_1)	Insufficient Power (assign no credits)

The bin-counting process creates one bin per each sample time and registers a number of "meal hits" under it. Table 1 presents the credit adding rules. At every step, if the detector claims H_0, then two credits are added to every bin in the H_0 window; similarly, if the detector claims H_1, two credits are added to every bin in the H_1 window. If both H_0 and H_1 are rejected, then one credit is added to every bin in both H_0 and H_1 windows. If both H_0 and H_1 are accepted, it means neither H_0's nor H_1's statistic is greater than the threshold, and thus no credit is assigned. The credit adding rules are also highlighted in Figure 1.

A peak in the bin-count signal means that the detector makes a number of decisions at different time steps that all point to the same meal time, indicating a strong positive

hit of a meal. On the other hand, if a bin only receives one or two counts (a typical width of H_0 and H_1 windows is 5 sample steps), it means the detector did not generate consistent decisions as the bin passed through the H_0 and H_1 windows, indicating a possibility of false positive.

Two design parameters, a threshold t and a minimum width w, are used to define positive peaks in the bin-count curve: a peak is characterized by at least w consecutive bin counts that are above t. At each time step, the detector examines the bin-count signal (up until the current step) and raise a meal alarm if a new peak emerges. In Figure 1, the bin-count curve crosses the threshold of $t = 12$ around time 2170 (the widths of wd, H_0 and H_1 windows are all set to 5 sample periods) and stays above it for at least $w = 3$ steps, therefore, a meal alarm is declared by the CFAR detector around the time 2182. The bin-count peaks emerge with a time lag from the detector's perspective because of how the CFAR detector works: it adds bin counts to past time windows, H_0 and H_1, according to current hypothesis test results. The example also demonstrates that the alarming sensitivity of our detector can be tuned by varying t and w: the detector may raise the meal alarm sooner with smaller t and w, at the cost of more false alarms. We also note that there is a few steps delay between the actual meal time and when the bin-count curves starts rising. This delay phenomenon is consistently observed in the simulation studies and we believe this is related to the fact that there is a delay from the onset of eating to CGM starts changing, which makes sense physiologically: in the maximal model, meal carbohydrates have to pass several digestion compartments before affecting the plasma glucose. In our future work we will incorporate the meal delay into the detector design in order to accurately identify the onset of a meal.

6. CASE STUDY

We evaluate the CFAR meal detector on an in-silico diabetes database that is generated by the FDA-accepted maximal model described in Section 4.1. We compare the performance of our detector with a state-of-the-art RoC-based voting meal detection algorithm published in [12].

6.1 In-Silico Diabetic Database

The academic version of the FDA-accepted Type 1 Diabetic Simulator [18] (T1DMS) comes with ten virtual subjects that are drawn from the same parameter distribution with the FDA-accepted virtual subject population. Each virtual subject is realization of the 32 parameters of the physiological model (see Table 2 for the complete list of parameters). To thoroughly test the performance of our detector across a wide range of possible patient parameters, we randomly sampled $10,000$ virtual subjects from the parameter space spanned by the ten T1DMS virtual subjects. The $10,000$ virtual subjects are sampled from an independent uniform distribution in the hypercubes spanned by the T1DMS subjects. Because real patients' parameters are actually from a joint distribution based on clinical study data [22], we over-sample the parameter space by assuming no knowledge about the joint distribution, which is good for testing the robustness of the detectors.

We run simulations of the FDA-accepted maximal model (with the virtual subjects' parameters) by feeding meals and insulin inputs that mimic the real-life scenario of a Type 1 patient. Three meals with random amounts of carbohydrates are fed to a virtual subject every simulation day. Whenever a meal is given, a meal bolus that is calculated based on a randomized meal insulin ratio is also fed into the simulator. In addition, we set a checkpoint once every hour except overnight time (22PM - 6AM) to mimic the scenario that a patient may check his/her own blood sugar and take correction insulin boluses when the glucose reading is too high. The dose of each correction bolus is calculated based on a randomly drawn insulin sensitivity value. The simulation settings, including the amount of carbohydrates in each meal, meal insulin ratio, and insulin sensitivity are independently and randomly sampled from ranges that are consistent with de-identified clinical data on over 60 Type 1 diabetics' CGM readings and pump logs, which are collected in a collaboration study with the Penn Diabetes Center.

The complete *in-silico* database consists of simulated glucose data sampled at a one minute period and the corresponding insulin & meal inputs. The database currently includes two sets of experimental data: 3-day simulation data of $10,000$ virtual subjects (data set No. 1) and 15-day simulation data of $1,000$ virtual subjects (data set No. 2). The first set is to test a detector's performance across a large number of virtual subjects and the second set can be used to test a detector's performance for a prolonged per-subject simulation period.

6.2 In-Silico Experimental Results

We test the CFAR meal detector and the RoC-based three out of four voting meal detector [12] on the *in-silico* diabetic databases described before. The two detectors are ran on the exactly same set of virtual subjects' simulated data and the CGM measurements are sampled at a one-minute period. A receiver operating characteristic (ROC) curve is generated for each detector by changing the detection sensitivity. For the CFAR meal detector, the bin-counting threshold t is varied from 16 to 11 to change its sensitivity. The rest of the CFAR detector parameters are chosen as follows: the training data window width wt is 300 samples (or 5 hours), the lengths of the wd, H_0 and H_1 windows are all 5 samples, and the bin-counting w is set to be 3. For the three out of four RoC voting meal detector, the RoC threshold is varied.

Figure 2 shows the ROC curves of two detectors when evaluated on the first experiment data set of $10,000$ virtual subjects (the ROC curves of two detectors running on the second data set of $1,000$ virtual subjects are highly consistent with the results on the $10,000$ data set; the two are almost identical). The CFAR detector's best operation point (circled out in Figure 2) is at the 99.6% detection rate and 1.4% false alarm rate. That is on average about one missed detection every 90 days and one false alarm every 24 days. To achieve the same detection rate, the RoC voting detector has a 39.6% false alarm rate, which is two false alarms every day, 48 times higher than our detector. To achieve the same false alarm rate, the RoC voting detector has a 78.5% detection rate, which is about one missed detection every 1.5 days, 60 times higher than our detector. At one of its best operation points (circled out in Figure 2), the RoC voting detector has a 93.2% detection rate and a 8.4% false alarm rate; that is on average about one miss every 5 days and one false alarm every 4 days, i.e., 18 times more missed detection and 6 times more false alarms than our detector. The CFAR detector significantly outperforms the RoC voting detector. In practice, a false alarm made by the meal detector may trigger an AP system to unnecessarily deliver insulin boluses, putting the patient at a greater risk of hypo-

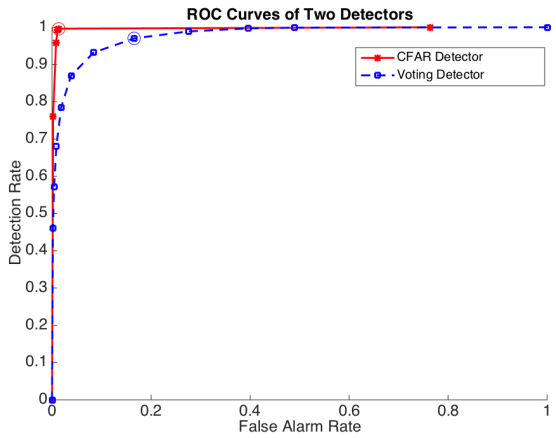

Figure 2: False Alarm Rate vs. Detection Rate of the two detectors on 10,000 virtual subjects.

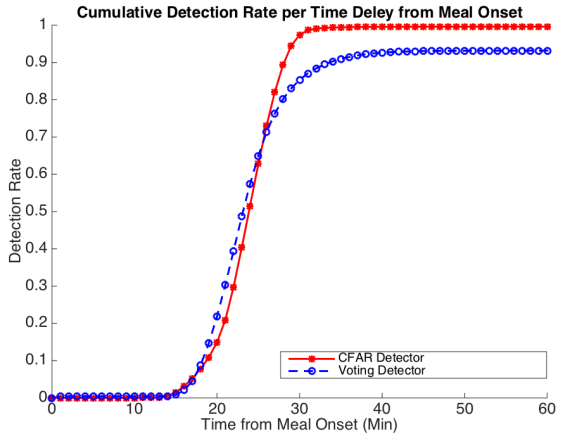

Figure 3: Cumulative detection rates over time from the onset of meals on 10,000 virtual subjects.

glycemia; on the other hand, a missed detection may result in insufficient insulin to compensate post-meal glucose rises and leave the patient suffer from hyperglycemia.

Figure 3 shows the cumulative detection percentage over the time delay since the onset of a meal, a measure of how soon each detector raises an alarm after a true meal. The curves in Figure 3 are obtained when both detectors are running at their operating points marked in Figure 2. Initially, the CFAR detector's detection time is about less than one minute (that is less than one sample time) late than the rate-of-change voting detector, an insignificant difference given that the glucose physiology changes slowly on a minute scale (most commercial CGM sensors transmit one reading every 5 minutes). After the crossing point at around 70%, the CFAR detector quickly reaches its 99.6% detection rate before the 30 minute delay, whereas the RoC voting detector reaches its maximum 93.2% detection rate around the 40 minute delay mark.

Figure 4 compares the number of per-subject missed meals of the two detectors. The results demonstrate the consistency of detection performance on different subjects, i.e., whether the detector performs particularly bad for certain

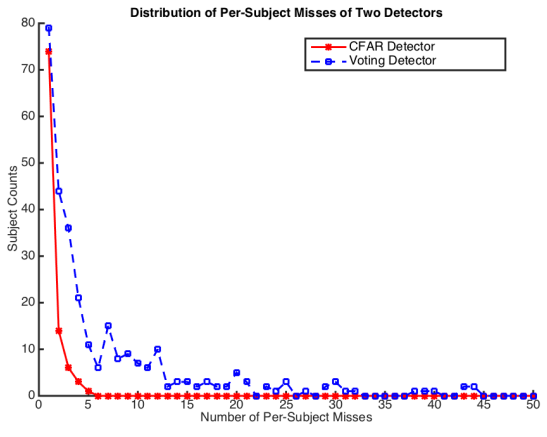

Figure 4: Distribution of per-subject misses on 1,000 virtual subjects.

subjects. In this test we used the data set No. 2 (1,000 subjects, 15 simulation days on each) since we want to examine the per-subject performance of each detector. Over 15 simulation days, each subject has 44 meals to be tested (3 meals a day; the very first meal is excluded as the simulator is still warming up). The CFAR detector has 137 total misses on 98 subjects, out of which 74 subjects (or 76%) have only one miss. The maximum number of per-subject misses is 5. The RoC voting detector has 2063 total misses on 298 subjects, out of which 79 subjects (or 26%) have one miss, and 21 subjects have over half of the meals missed, including 2 subjects on whom all meals are missed. The results clearly demonstrate the power of parameter invariant design: the detector exhibits a much more consistent performance over a large virtual population, comparing to the RoC voting detector.

The *in-silico* evaluation shows that the CFAR detector achieves a high detection rate with a low false alarm rate across a large virtual patient population and a wide range of meal/insulin simulation settings. It significantly outperforms the existing RoC voting detector in the same evaluation, and it also has a much more consistent performance in terms of per-subject missed detections.

7. CONCLUSION

In this paper, we develop a CFAR meal detector based on an augmented version of the Bergman minimal model. *In-silico* evaluation using the FDA-accepted maximal physiological model shows that the detector significantly outperforms the existing RoC based voting meal detector. The new detector achieves a near-perfect 99.6% detection rate across a wide range of physiological parameters and simulation settings, with a low average false alarm rate of once every 24 days.

This preliminary work is intended to motivate the study of providing guarantees regarding meal (disturbance) detection and ultimately the prediction of near future glucose values to inform critical actions that aim at preventing unsafe conditions such as hypoglycemia. Our future work includes extensive evaluation of the meal detector on clinical data and developing techniques to formally bound the performance of the detector.

8. REFERENCES

[1] K.-P. Adlassnig, C. Combi, A. K. Das, et al. Temporal representation and reasoning in medicine: research directions and challenges. *Artificial Intelligence in Medicine*, 38(2):101–113, 2006.

[2] A. D. Association et al. Hypoglycemia (low blood glucose). *URL http://www. diabetes. org/living-with-diabetes/treatment-and-care/blood-glucose-control/hypoglycemia-low-blood. html*, 2011.

[3] R. N. Bergman, Y. Z. Ider, C. R. Bowden, and C. Cobelli. Quantitative estimation of insulin sensitivity. *American Journal of Physiology-Endocrinology And Metabolism*, 236(6):E667, 1979.

[4] A. Burgos, A. Goñi, A. Illarramendi, and J. Bermúdez. Real-time detection of apneas on a pda. *IEEE Transactions on Information Technology in Biomedicine*, 14(4):995–1002, 2010.

[5] F. Cameron, G. Niemeyer, and B. A. Buckingham. Probabilistic evolving meal detection and estimation of meal total glucose appearance. *Journal of diabetes science and technology*, 3(5):1022–1030, 2009.

[6] J. Chen and R. J. Patton. *Robust model-based fault diagnosis for dynamic systems*. Springer Publishing Company, Incorporated, 2012.

[7] C. Cobelli, C. D. Man, G. Sparacino, L. Magni, G. D. Nicolao, , and B. P. Kovatchev. Diabetes: Models, signals, and control. *Biomedical Engineering, IEEE Reviews in*, 2, 2009.

[8] C. Cobelli, E. Renard, and B. Kovatchev. Artificial pancreas: past, present, future. *Diabetes*, 60(11):2672–2682, 2011.

[9] E. Cobry, K. McFann, L. Messer, V. Gage, B. VanderWel, L. Horton, and H. P. Chase. Timing of meal insulin boluses to achieve optimal postprandial glycemic control in patients with type 1 diabetes. *Diabetes technology & therapeutics*, 12(3):173–177, 2010.

[10] V. A. Convertino, S. L. Moulton, G. Z. Grudic, et al. Use of advanced machine-learning techniques for noninvasive monitoring of hemorrhage. *Journal of Trauma-Injury, Infection, and Critical Care*, 71(1):S25–S32, 2011.

[11] C. Dalla Man, F. Micheletto, D. Lv, M. Breton, B. Kovatchev, and C. Cobelli. The uva/padova type 1 diabetes simulator new features. *Journal of diabetes science and technology*, 8(1):26–34, 2014.

[12] E. Dassau, B. W. Bequette, B. A. Buckingham, and F. J. Doyle. Detection of a meal using continuous glucose monitoring implications for an artificial β-cell. *Diabetes care*, 31(2):295–300, 2008.

[13] C. Ellingsen, E. Dassau, H. Zisser, B. Grosman, M. W. Percival, L. Jovanovič, and F. J. Doyle. Safety constraints in an artificial pancreatic β cell: an implementation of model predictive control with insulin on board. *Journal of diabetes science and technology*, 3(3):536–544, 2009.

[14] C. for Disease Control, Prevention, et al. National diabetes statistics report: estimates of diabetes and its burden in the united states, 2014. *Atlanta, GA: US Department of Health and Human Services*, 2014.

[15] R. Gillis, C. C. Palerm, H. Zisser, L. Jovanovic, D. E. Seborg, and F. J. Doyle. Glucose estimation and prediction through meal responses using ambulatory subject data for advisory mode model predictive control. *Journal of diabetes science and technology*, 1(6):825–833, 2007.

[16] R. Ivanov, J. Weimer, A. Simpao, M. Rehman, and I. Lee. Early detection of critical pulmonary shunts in infants. In *Proceedings of the 6th ACM/IEEE International Conference on Cyber-Physical Systems*, ICCPS '15, 2015.

[17] T. Kobayashi, S. Sawano, T. Itoh, K. Kosaka, H. Hirayama, and Y. Kasuya. The pharmacokinetics of insulin after continuous subcutaneous infusion or bolus subcutaneous injection in diabetic patients. *Diabetes*, 32(4):331–336, 1983.

[18] B. P. Kovatchev, M. Breton, C. D. Man, and C. Cobelli. In silico preclinical trials: A proof of concept in closed-loop control of type 1 diabetes. *Diabetes Sci Technol*, 3(1):44–55, 2009.

[19] H. Lee, B. A. Buckingham, D. M. Wilson, and B. W. Bequette. A closed-loop artificial pancreas using model predictive control and a sliding meal size estimator. *Journal of diabetes science and technology*, 3(5):1082–1090, 2009.

[20] I. Lee, O. Sokolsky, S. Chen, et al. Challenges and research directions in medical cyber–physical systems. *Proceedings of the IEEE*, 100(1):75–90, 2012.

[21] L. Magni, D. Raimondo, C. D. Man, G. D. Nicolao, B. Kovatchev, and C. Cobelli. Model predictive control of glucose concentration in type i diabetic patients: An in silico trial. *Biomedical Signal Processing and Control*, 4(4):338 – 346, 2009.

[22] C. D. Man, R. A. Rizza, and C. Cobelli. Meal simulation model of the glucose-insulin system. *Biomedical Engineering and IEEE Transactions on*, 54(10):1740–1749, oct. 2007.

[23] T. M. Mitchell, R. Hutchinson, M. A. Just, et al. Classifying instantaneous cognitive states from fmri data. In *AMIA Annual Symposium Proceedings*, pages 465–469, 2003.

[24] G. Nucci and C. Cobelli. Models of subcutaneous insulin kinetics. a critical review. *Computer methods and programs in biomedicine*, 62(3):249–257, 2000.

[25] K. Ogata. *Discrete-time control systems*, volume 2. Prentice Hall Englewood Cliffs, NJ, 1995.

[26] A. Pantelopoulos and N. G. Bourbakis. A survey on wearable sensor-based systems for health monitoring and prognosis. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 40(1):1–12, 2010.

[27] S. Saria, D. Koller, and A. Penn. Learning individual and population level traits from clinical temporal data. In *Proceedings of Neural Information Processing Systems*, pages 1–9, 2010.

[28] S. Saria, A. K. Rajani, J. Gould, et al. Integration of early physiological responses predicts later illness severity in preterm infants. *Science translational medicine*, 2(48):48ra65, Sept. 2010.

[29] L. L. Scharf. *Statistical signal processing*, volume 98. Addison-Wesley Reading, MA, 1991.

[30] L. L. Scharf and C. Demeure. *Statistical Signal Processing*. Addison-Wesley Publishing Company, 1991.

[31] J. Weimer, S. A. Ahmadi, J. Araujo, et al. Active actuator fault detection and diagnostics in hvac systems. In *Proceedings of the Fourth Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 107–114, 2012.

[32] J. Weimer, R. Ivanov, A. Roederer, S. Chen, and I. Lee. Parameter invariant design of medical alarms. *IEEE Design & Test*, accepted.

[33] J. Weimer, S. Kar, and K. H. Johansson. Distributed detection and isolation of topology attacks in power networks. In *International Conference on High Confidence Networked Systems*, pages 65–72, 2012.

[34] J. Weimer, D. Varagnolo, and K. H. Johansson. Distributed model-invariant detection of unknown inputs in networked systems. In *International Conference on High Confidence Networked Systems*, pages 127–134, 2013.

[35] J. Weimer, D. Varagnolo, M. Stankovic, and K. Johansson. Parameter-invariant detection of unknown inputs in networked systems. In *Conference on Decision and Control*, pages 4379–4384, 2013.

[36] A. S. Willsky. A survey of design methods for failure detection in dynamic systems. *Automatica*, 12(6):601–611, 1976.

[37] T. Wolever and C. Bolognesi. Prediction of glucose and insulin responses of normal subjects after consuming mixed meals varying in energy, protein, fat, carbohydrate and glycemic index. *The Journal of nutrition*, 126(11):2807–2812, 1996.

Table 2: Physical meanings, units, and a typical Type 1 diabetic value of the parameters in the T1DMS maximal model

Parameters	Physical Meaning	Units	Typical Type 1 value
m_1	Rate parameter	min^{-1}	0.20
m_2	Rate parameter	min^{-1}	0.28
m_3	Rate parameter	min^{-1}	0.30
m_4	Rate parameter	min^{-1}	0.11
k_i	Rate parameter	min^{-1}	0.0084
P_{2u}	Rate parameter	min^{-1}	0.0034
V_i	Insulin volume	liter/kg	0.053
I_b	Basal insulin level	pmol/liter	108
BW	Body weight	kg	86
k_1	Rate parameter	min^{-1}	0.063
k_2	Rate parameter	min^{-1}	0.13
k_{p1}	Extrapolated EGP	mg/kg/min	4.6
k_{p2}	Liver glucose effectiveness	min^{-1}	0.0042
k_{p3}	Insulin action on liver	mg/kg/min per pmol/liter	0.0095
V_{m0}	Model parameter	mg/kg/min	4.9
V_{mx}	Model parameter	mg/kg/min per pmol/liter	0.054
K_{m0}	Model parameter	mg/kg	237
V_g	Glucose volume	dL/kg	1.8
k_{a1}	Rate parameter	min^{-1}	0.0034
k_{a2}	Rate parameter	min^{-1}	0.016
k_d	Rate parameter	min^{-1}	0.016
k_{sc}	Rate parameter	min^{-1}	0.095
k_{gri}	Rate of grinding	min^{-1}	0.039
b	Meal related parameter	unitless	0.79
d	Meal related parameter	unitless	0.18
k_{max}	Rate parameter	min^{-1}	0.039
k_{min}	Rate parameter	min^{-1}	0.074
f	Meal related parameter	unitless	0.9
k_{abs}	Rate parameter	min^{-1}	0.21
k_{e1}	Glomerular filtration rate	min^{-1}	0.0005
k_{e2}	Renal threshold of glucose	mg/kg	339
F_{snc}	Insulin-independent glucose utilization	mg/kg/min	1