Research and Applications

Automated meal detection from continuous glucose monitor data through simulation and explanation

Min Zheng,[1] Baohua Ni,[2] and Samantha Kleinberg[1]

[1]Computer Science, Stevens Institute of Technology, Hoboken, New Jersey, USA, and [2]Electronic Engineering, Tsinghua University, Beijing, China

Corresponding Author: Samantha Kleinberg, PhD, Computer Science, Stevens Institute of Technology, 1 Castle Point on Hudson, Hoboken, NJ 07030, USA; samantha.kleinberg@stevens.edu

ABSTRACT

Background: Artificial pancreas systems aim to reduce the burden of type 1 diabetes by automating insulin dosing. These systems link a continuous glucose monitor (CGM) and insulin pump with a control algorithm, but require users to announce meals, without which the system can only react to the rise in blood glucose.
Objective: We investigate whether CGM data can be used to automatically infer meals in daily life even in the presence of physical activity, which can raise or lower blood glucose.
Materials and Methods: We propose a novel meal detection algorithm that combines simulations with CGM, insulin pump, and heart rate monitor data. When observed and predicted glucose differ, our algorithm uses simulations to test whether a meal may explain this difference. We evaluated our method on simulated data and real-world data from individuals with type 1 diabetes.
Results: In simulated data, we detected meals earlier and with higher accuracy than was found in prior work (25.7 minutes, 1.2 g error; compared with 48.3 minutes, 17.2 g error). In real-world data, we discovered a larger number of plausible meals than was found in prior work (30 meals, 76.7% accepted; compared with 33 meals, 39.4% accepted).
Discussion: Prior research attempted meal detection from CGM, but had delays and lower accuracy in real data or did not allow for physical activity. Our approach can be used to improve insulin dosing in an artificial pancreas and trigger reminders for missed meal boluses.
Conclusions: We demonstrate that meal information can be robustly inferred from CGM and body-worn sensor data, even in challenging environments of daily life.

Key words: diabetes mellitus type 1; meal detection; continuous glucose monitoring; pancreas; artificial

INTRODUCTION

Type 1 diabetes (T1D) is a chronic, lifelong disease that affects millions of individuals.[1] T1D is characterized by the inability to produce insulin, so individuals must manage their blood glucose (BG) using an insulin pump or injections. While meals are a key cause of BG changes, other factors such as stress and activity also affect BG and insulin sensitivity, making T1D complex to manage. Many complications of diabetes such as chronic kidney disease[2] are a result of long-term hyperglycemia (dangerously high BG).[3] Tight glycemic

control can prevent microvascular complications of diabetes and yet tighter control increases the risk of hypoglycemia, which can lead to seizure, coma, and death.[4]

Researchers have aimed to reduce the burden of T1D through artificial pancreas (AP) systems that measure glucose with continuous glucose monitors (CGMs), determine BG trajectory, and deliver insulin through a pump to automatically correct high values. CGMs measure glucose in fluid between cells as a proxy for BG, but because this is delayed relative to BG a rise in glucose may appear

well after a meal has begun. AP systems either require users to provide input about their meals (making this no longer a closed loop) or infer meals from CGM data so that the system can correct for them. Most approaches to detect meals track glucose rate of change (ROC),[5] but in the real world, this is susceptible to false positives due to physical activity (which can increase BG) and noise. While the Glucose Rate Increase Detector (GRID)[6] addresses those challenges, this comes at the expense of delays in recognizing meals and GRID does not yet identify meal amount, as is needed for insulin dosing. Meal detection can also help identify missed meal insulin boluses,[7] which are associated with poor glycemic control.[8]

To address this, we propose a method to detect meals and amount of carbohydrates consumed by combining CGM, insulin, and activity monitor data with a simulation of glucose. We build on the Food and Drug Administration–approved glucose-insulin model (GIM),[9,10] which incorporates physical activity. When simulated and observed glucose differ significantly, that triggers a search for a meal that may explain the difference.

Background

We aim to automatically infer meals from data collected by body-worn sensors (specifically, CGM and activity sensors), as this allows for a fully automated solution. This is similar to finding latent variables (as meals are not directly observed), but prior work has aimed to identify causal structures with latent variables,[11] rather than to recover when latent events happen, as is needed for meal detection. A meal's effect on glucose could be viewed as a changepoint, but methods such as Bayesian online changepoint detection[12] only identify that a change has occurred (not why) and cannot provide the meal information needed to adjust insulin.

Thus, we focus on related work in detecting meals for T1D management. Many approaches use the ROC of CGM measurements,[6] as meals can lead to a rapid increase in BG. One approach tracking the speed of glucose changes combines a set of methods (eg, Kalman filters) to predict glucose, and identifies meals when they are supported by three-quarters of the methods.[5] Lee and Bequette[13] used Kalman filters to find meals and identified meal size by feeding the filtered glucose trajectory into a finite impulse response filter, where a meal is an impulse. The variable state dimension (VSD) approach[14] also builds on Kalman filters and identifies meal time and size simultaneously. However, these methods have high false positives in the presence of exercise[15] and stress,[16] as both can lead to an increase in glucose that could be mistaken for a meal. Further, meal detection is delayed, as the methods must wait until a meal affects BG enough to be distinguished from other factors that also increase BG, such as intense exercise.

As an alternative, Turksoy et al[17] introduced a multivariate adaptive model to detect meals using an unscented Kalman filter (UKF) for state estimations. However, due to the use of the UKF, parameter estimation may not converge to the true value, and instead can become stuck in a local minimum. Ramkissoon et al[18] also built on a UKF, but this approach not yet been evaluated on real-world data. Other approaches[19] use the minimal glucose physiological model,[20] modifying it to capture physiological trends and detect meals. This does not require tuning to individual parameters, but does not identify the meal size, as is needed to accurately estimate insulin needs.

Table 1 summarizes the state of the art. Only 2 methods incorporate physical activity, and half were tested on real-world data. GRID[6] does both, but detects meals ∼42 minutes after they start,

and did not evaluate meal size error. While Lee and Bequette[13] reported low meal size error, the approach was tested in simulations and without exercise, so it is unknown whether similar results will be obtained in the real world. VSD identified meal size[14] and has been applied to real-world data,[21] but it has significant delays and the meal size errors can lead to overestimates of insulin doses. Samadi et al[23] also identified meal size in real-world data, but with a high standard deviation in errors (28 g, roughly the amount of carbohydrates in a medium potato or banana). Further, that work evaluated meal size using the set of carbohydrate estimates within 2 hours of the start of the meal, making it difficult to compare results to methods that identify the meal and its quantity at the same time. We are not aware of any evaluations in real-world data in the presence of exercise. Our simulation-based explanation (SBE) approach instead finds meal times and amounts earlier and more accurately, enabling integration into insulin reminders and fully automated systems.

Significance

1. *Earlier and more accurate meal detection*: Unlike ROC-based methods, which cannot detect a meal until there is a large BG change, we leverage simulation to rapidly identify meals that explain changes from expected glucose. This allows insulin to be delivered earlier, preventing unhealthy glucose excursions.
2. *Robustness to real-world challenges*: Real-world CGM data are noisy, and many factors other than meals affect glucose. Other methods prioritize BG prediction, but this comes at the expense of meal detection, as errors compound throughout a day. We provide more robust results by allowing periodic restarts, and optimizing for meal detection by focusing on how well meals *explain* observed glucose.
3. *Realistic evaluation*: We evaluate accuracy for detecting meal time and size on realistic full-day simulations that include meals and physical activity, and on challenging real-world data.

MATERIALS AND METHODS

We first formalize the problem and notation, then introduce the model we build on, and last, describe our approach for meal detection: SBE.

Problem description and notation

Our primary motivation is improving AP systems, which use a CGM and insulin pump. Thus, we assume that these data sources are available. The pump records continuous basal insulin (pmol/kg/min) and discrete insulin boluses (pmol/kg). CGMs measure glucose (mg/dL) at 5-minute intervals. We assume that individual information is available (body weight, resting heart rate [HR]), individuals wear an activity tracker capturing heart rate, and meals are the only unmeasured variable affecting BG. However, in experimental results, we demonstrate that our approach succeeds even when this does not hold.

We denote the data by $D = [\boldsymbol{G}, \boldsymbol{HR}, \boldsymbol{B}, \boldsymbol{I}]$, where each bold variable is a time series. $\boldsymbol{G} = [G(1), \ldots, G(n)]$ is a subject's glucose from time 1 to time n, $\boldsymbol{HR}$ denotes heart rate, $\boldsymbol{B}$ is insulin boluses, and $\boldsymbol{I}$ the insulin infusion rate at each time point. Because these devices record at different frequencies, we synchronize all to the 5-minute rate of the CGM.

Our task is using observations to detect meals, meaning their start time (m_{st}), duration (m_{du}), and size (m_c, grams of carbohydrates).

Table 1. Related work in meal detection from continuous glucose monitoring

Method	Real-World Data	Exercise	Mean Delay (min)	Mean Meal Size Error (g)
MDA[5]	Yes	No	30	NR
Lee and Bequette[13]	No	No	35 ± 8.3	−0.75 ± 9.3
GRID[6]	Yes	Yes	42	NR
VSD (2015)[14]	No	No	34 ± 19	4.95-10.4
Ramkissoon et al[18]	No	Yes	28 ± 3	NR
VSD (2017)[21]	Yes	No	45 ± 14	5.72-10.4
Kölle et al[22]	No	No	35 (max)	NR
Samadi et al[23]	Yes	No	34.8 ± 22.8	1.7 ± 28.1
SBE (our method)	Yes	Yes	25.7 ± 5	1.2 ± 3.6

Values are mean ± SD or range, unless otherwise indicated. Meal size error is given in grams of carbohydrates.

GRID: Glucose Rate Increase Detector; MDA: meal detection algorithm; NR: not reported; SBE: simulation-based explanation; VSD: variable state dimension.

Background: GIM

Our approach relies on simulation, rather than fitting to the observed glucose as in other approaches, such as VSD[14] and GRID,[6] as this lets us overcome measurement errors. We build on GIM,[9,10,24] which is well validated and models the effect of physical activity on glucose.

The main parts of GIM are endogenous glucose production, glucose utilization, and how exercise influences glucose. Glucose utilization includes insulin-independent utilization $U_{ii}(t)$ (constant glucose uptake by the brain and erythrocytes) and insulin-dependent utilization $U_{id}(t)$ (the effect of endogenous and exogenous insulin). Total glucose utilization U at time t is

$$U(t) = U_{ii}(t) + U_{id}(t). \tag{1}$$

To incorporate physical activity, GIM builds on Breton,[25] whose approach linked insulin action and glucose effectiveness to current HR ($HR(t)$) and basal heart rate (HR_b). The approach models changes in glucose uptake due to energy consumption (Y) and insulin action (Z):

$$\dot{Y}(t) = -\frac{1}{T_{HR}}\,[Y(t) - (HR(t) - HR_b)],\ \ Y(0) = 0 \tag{2}$$

$$\dot{Z}(t) = -\left[\frac{f(Y(t))}{T_{in}} - \frac{1}{T_{ex}}\right] Z(t) + f(Y(t)),\ \ Z(0) = 0 \tag{3}$$

$$f(Y) = \frac{\left(\frac{Y}{aHR_b}\right)^n}{1 + \left(\frac{Y}{aHR_b}\right)^n}. \tag{4}$$

The variables T_{HR}, T_{in}, and T_{ex} are parameters to model the physiological reaction to energy expenditure and exercise. To model the effect of exercise on insulin-dependent glucose utilization, $U_{id}(t)$, Dalla Man et al[10] added a scale variable $W(t)$ defined as

$$W(t) = \begin{cases} \int_0^t (HR(t) - HR_b)dt \text{ for } t < t_z \\ 0 \text{ otherwise} \end{cases} \tag{5}$$

where t_z is the end of the activity's effect on insulin action, Z. Exercise (reflected in HR) has a cumulative effect, but is modeled with a step increase in HR. As we discuss in the following section, we use HR measurements at each time to capture exercise intensity.

Meal detection using SBE

Our method has 3 parts: (1) find when observed glucose diverges from predicted glucose, (2) generate predicted glucose for potential meals using GIM, and (3) test if potential meals explain observed glucose. Figure 1 shows an overview.

Find diverging point

First, we must determine when to search for a meal. To do this efficiently, we identify times when observed and predicted glucose differ, and then determine whether this difference can be explained by a meal. We continually update the insulin infusion rate $IIR(t)$ and boluses using the insulin pump data at each t and update $HR(t)$ in equation 5 using the HR data. Thus, GIM is updated with observed variables (HR, B, I) (all except G), to obtain the predicted glucose time series G'. We aim to determine when the difference between G and G' is large enough that a meal may explain it. As CGM data can be noisy, we first smooth G using Savitzky-Golay filter,[26] then compare the means of G' and G over a window ζ. This handles the noise and outlying values of CGM data better than comparison at a single time point. When the difference between G' and G is greater than a threshold ϕ, we call this time point t the diverging point.

More formally, we test

$$|\overline{G'[t - \zeta : t]} - \overline{G[t - \zeta : t]}| > \phi \tag{6}$$

where $\overline{G'[t - \zeta : t]}$ denotes the mean of G' from time $t - \zeta$ to time t. We set ϕ as a percent of the mean observed glucose (5% in this work).

Generate predicted glucose sequence

When t is a diverging point, the algorithm then searches for a meal that may have begun before t. A meal is defined as $m = [m_{st}, m_c, m_{du}]$, with m_{st} being the start time, m_c being the meal size (grams of carbohydrates), and m_{du} being the meal duration. To constrain the search space, we set a maximum meal duration Γ, maximum meal size m_{max}, and maximum lag τ between meal start and when glucose diverges. Thus, $t - \tau \leq m_{st} \leq t$.

We then perform a grid search over combinations of meal timing (start, duration) and size, generating predicted glucose G' for each meal, and comparing it against the observed glucose. For efficiency, we begin with a resolution of 1 g for meal size and 2 minutes for meal time until $d_{Euc} < 2\epsilon$. Then we perform a finer search, using 0.1 g and 1 minutes. These values were chosen as a tradeoff between efficiency and precision.

For each potential meal m, we update the parameters in GIM beginning at m_{st}. In GIM, glucose at time t is proportional to the carbohydrate ingestion rate $d_{\delta(t)}$ (in mg/min):

$$G'(t) \propto d_{\delta(t)} \tag{7}$$

We assume that meals are consumed at a constant rate, as this is more realistic than assuming all carbohydrates are ingested at a single time as in other approaches. Thus, the ingestion rate at time t is:
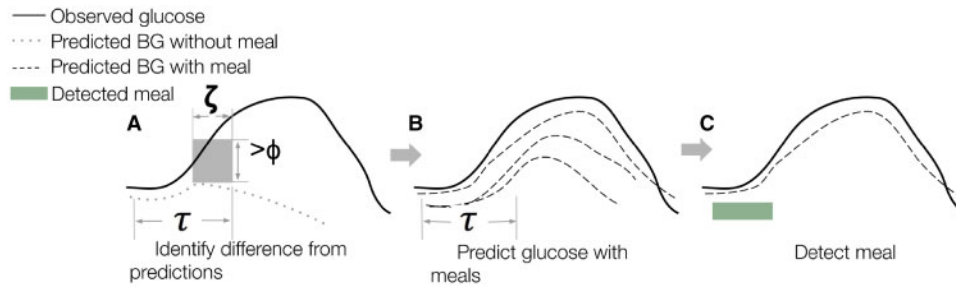
Figure 1. Overview of method: (A) identify times in which observed glucose differs significantly from predicted glucose, (B) simulate glucose trajectories with varying meal sizes and times, and (C) determine which, if any, inferred meals account best for observed glucose. BG: blood glucose.

$$d_{\delta(t)} = \frac{m_c}{m_{du}}, \ t \in [m_{st}, \ m_{st} + m_{du}] \tag{8}$$

Updating $d_{\delta(t)}$ at each time t during the meal yields the predicted glucose G' for each potential meal m.

Meal detection

After simulating G' for all potential meals, m we must determine which meals (if any) explain the observed glucose G. We accept a meal when the average Euclidean distance, d_{Euc}, between G' and G is smaller than a threshold ϵ. For a diverging point t and potential meal $m = [m_{st}, \ m_{st} + m_{du}]$, d_{Euc} is:

$$d_{Euc} = \frac{\sqrt{\sum_{i=m_{st}}^{t} |G'(i) - G(i)|^2}}{t - m_{st}} \tag{9}$$

If multiple meals have $d_{Euc} < \epsilon$, we select the one with the lowest d_{Euc}. The value for ϵ was chosen after tuning (see the Supplementary Appendix for details on this and algorithmic complexity).

Figure 1 shows an overview of SBE. We start with observed glucose (solid line) and find when it diverges from predicted glucose (dotted line). Then, we search for meals that may explain the difference. After generating predicted glucose trajectories (dashed lines), we detect the meal (green bar) that accounts best for the observed glucose.

Experimental approach

We evaluate our approach using simulated data with ground truth and real-world data in which meal information is missing and must be inferred. While the real-world data do not have ground truth, using operational criteria, we demonstrate that our approach works in realistic cases. We compare our approach with VSD, following the details and parameter settings of Xie and Wang.[21] We selected this method for comparison because it detects meals and their size simultaneously and has been previously applied to real-world data with high accuracy.

Simulated data

We first evaluate whether the algorithms can detect a full day of meals in the presence of exercise. All simulations are created using our implementation of the GIM. Experiments for parameter estimation are described in the Supplementary Appendix. We first generate a representative day, before generating a cohort of 100 individuals. For the single day, we use the parameters: body weight 60 kg, basal HR 60 beats/min, and basal insulin rate 2.5 pmol/kg/min. The sequence of events is shown in Figure 2 and Table 2.

We then simulate 100 more individuals for 24 hours each with 1 episode of exercise, varying the parameters for body weight (40-90 kg),

exercise duration (30-50 minutes), meal duration (10-40 minutes), meal size (10-50 g), and meal start time.

Our primary evaluation metrics are delay and meal size error. Delay is the difference between when a meal is identified and the actual meal start time. This metric is critical because identifying a meal more quickly allows earlier bolus reminders or insulin adjustment via an AP. Meal size error is the difference in grams of carbohydrates between the detected meal and ground truth. This information is necessary for accurate insulin dosing. We additionally evaluate recall and precision of eating minutes. Recall is how much eating time is correctly recovered, and precision is how much identified eating time is actually during a meal.

Real-world data

To evaluate our approach on real-world data, we use the Diabetes Management Integrated Technology Research Initiative (DMITRI) dataset.[28] DMITRI includes data for 17 individuals (10 men, 7 women) in daily life for 3 days each (around 72 hours). Participants were active, exercising 2-3 times per week (13 of 17, >4 times per week), and during the study activities included running, cycling, and yoga. The data used here are glucose (Dexcom 7+ CGM), insulin basal and bolus rates (insulin pumps), and HR (Polar chest strap). Participants used their own insulin pumps. The HR and insulin data were synchronized to the 5-minute intervals of the CGM. At the beginning of data collection, body weight and resting HR were recorded. To handle missing CGM data (mean 24.6% missing), we use the imputation approach of Rahman et al,[27] which had high accuracy on this dataset. We excluded 2 individuals with >40% missing data. After imputation, we had 1437 hours of data for 15 people (8 men, 7 women).

The key difference from simulation is in evaluation. While participants photographed their meals, these data were not synchronized to the rest and is not part of the publicly available DMITRI dataset, so there were no ground truth meal times. Thus, we use operational criteria based on prior work to determine if inferences are plausible. We accept a meal if (1) it happens before an increase in BG and (2) BG increases by 4 mg/dL within 30 minutes after the meal starts. This threshold is commonly used in works on identifying meal onset from CGM data.[13,14] This approach allows quantitative evaluation of the algorithms, but it remains possible that there are true meals that do not meet this criteria or nonmeal factors that do. We also examine whether there is a nearby insulin bolus, as boluses are generally administered in conjunction with a meal.

For our approach, we set body weight and resting HR using the measured value for each subject. We set $\zeta = 30$ minutes (the length of time series used to trigger meal search). As CGM data are every 5 minutes, this allows us to use 6 data points rather than only 2.
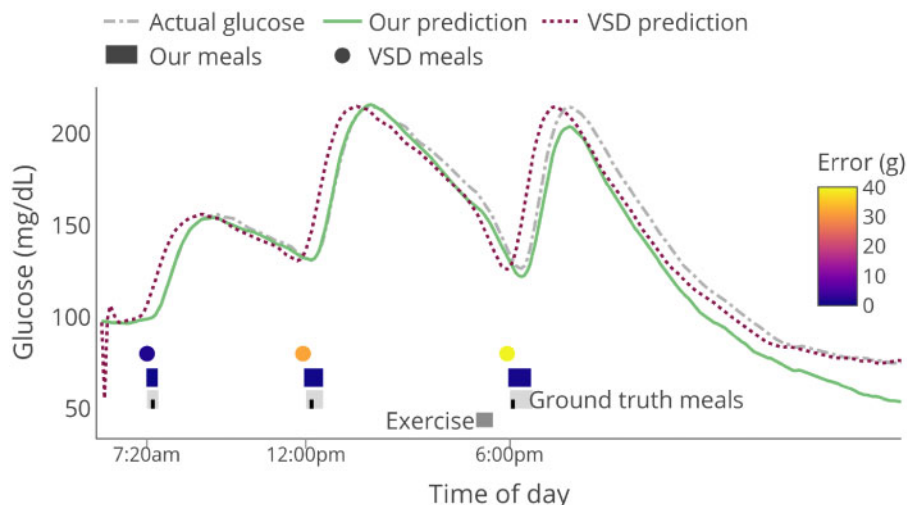
Figure 2. Twenty-four hours of simulated data with ground truth meals (gray boxes), insulin doses (vertical black bars in meals), and exercise (light gray box). Our approach finds all meals with low error (bars show meal duration), while variable state dimension (VSD) has significant error in meal size estimate (circles show meal time). Meals for both methods are shaded to indicate meal size error (dark blue = 0, yellow = 40-g error).

Table 2. Detailed series of events for the whole-day simulation depicted in Figure 2

Time	Event
6:00 AM	Simulation begins
7:20-7:40 AM	Breakfast, 14-g carbohydrates
7:30 AM	Insulin bolus, 70 pmol/kg
12:00-12:30 PM	Lunch, 24.7-g carbohydrates
12:10 PM	Insulin bolus, 80 pmol/kg
5:00-5:30 PM	Run, average HR 140 beats/min
6:00-6:40 PM	Dinner, 40.8-g carbohydrates
6:05 PM	Insulin bolus, 80 pmol/kg

HR: heart rate.

Similarly, we set $\phi = 10\%$ (threshold for identifying a diverging point) and $\epsilon = 5$ mg/dL (threshold for average Euclidean distance) due to the noise in real data. For the maximum delay between the start of a meal and the diverging point, we set $\tau = 60$ minutes. For VSD we use the parameter settings specified by Xie and Wang.[21]

RESULTS

Results on simulated data

We begin by discussing the case study, then the results for a simulated cohort of 100 individuals. Both use a 24-hour day with 3 meals at varying times plus exercise of varied duration and intensity. This is challenging, as the effect of one event (exercise, meal) may persist past the start of another.

Figure 2 shows results for both methods on the case study. First, we detect breakfast of the exact duration and size as the ground truth 25 minutes after it begins, and starting 1 minute earlier (recall 95%, precision 94.7%). VSD also identifies the meal start time with a 25-minute delay, but overestimates size by 1.5 g. We detect lunch with a 30-minute delay, with high accuracy for meal size (24.7 g ground truth, 25 g inferred), but VSD identifies a meal of 56.3 g with a 45-minute delay, which is a substantial overestimate. Finally, we detect dinner 26 minutes after it begins, again with high accuracy (40 g detected, 40.8 g ground truth). In contrast, with a delay of

50 minutes, VSD finds a meal of 80.3 g, which is nearly double the actual meal size, and would lead to a dramatic overestimate of insulin if used for dosing. Across the whole day, our approach has an average delay of 27 minutes vs 40 minutes for VSD. We estimate all meals to within 1 g, while VSD has an average meal size error of 23.3 g. This suggests that our approach can be accurately used for estimating insulin.

Results for the 100 simulated individuals for the full day are shown in Table 3. Our approach detects meals sooner (mean 25.7 ± 5 minutes) than VSD (48.3 ± 9 minutes) and with lower meal size error (1.2 ± 3.6 g) than VSD (17.2 ± 8 g). This difference would have a substantial impact on insulin dosing. Note that VSD assumes that a meal is ingested at a single time point (ie, it does not have a duration), which is unrealistic and may contribute to errors.

Meal detection on real-world data

We now show that it is possible to apply our approach to real-world data. This is challenging, as data can be missing, sensors are noisy, and many unobservable factors affect BG. Nevertheless, we show that in this challenging environment, our approach recovers a larger number of plausible meals than VSD does.

On the DMITRI dataset, our approach detected 30 meals, with 23 (76.7%) accepted as plausible and 7 (23.3%) rejected using our operational criteria. VSD detected 33 meals, but only 13 (39.4%) were accepted and 20 (61.6%) were rejected. Thus our approach has both more true positives and fewer false positives. We discuss 4 representative cases (breakfast, lunch, dinner, snacks). Full results for all subjects are in the Supplementary Appendix.

Figure 3A shows a breakfast detected by both methods, around 6:00 AM However, VSD detects a second meal at 8:20 AM around the time an insulin bolus is given, even though glucose has a downward trend at the time. In Figure 3B, there is a smaller afternoon meal (3:00 PM) and then what we expect is dinner based on the timing. Both methods detect the snack, but VSD detects the second meal too early, at a time when glucose is still decreasing and far from the insulin bolus. Our approach detects a more reasonable meal time. For lunch (Figure 3C), both methods detect a meal starting around 12:30 PM, but VSD has a false positive at 4:20 PM, while glucose is

Table 3. Experimental results for 100 simulated people

Experiment	Our Method				VSD			
	Recall	Precision	Delay (min)	Meal Size Error (g)	Recall	Precision	Delay (min)	Meal Size Error (g)
Breakfast	0.942[a]	0.947	24.9 ± 3[a]	0.8 ± 2.1[a]	0.920	0.951[a]	25.4 ± 5	5.1 ± 4.2
Lunch	0.917[a]	0.919	28.9 ± 6[a]	1.0 ± 3.2[a]	0.901	0.930[a]	43.5 ± 8	16.1 ± 6.3
Dinner	0.871[a]	0.881	27.8 ± 7[a]	1.4 ± 4.3[a]	0.845	0.893[a]	53.4 ± 11	23.6 ± 9.1
Overall	0.880[a]	0.933	25.7 ± 5[a]	1.2 ± 3.6[a]	0.860	0.936[a]	48.3 ± 9	17.2 ± 8.0

Delay is time between meal start and its detection and size error is in grams of carbohydrates. Recall and precision for VSD are at the meal level, rather than at the minute level, as it does not infer duration.

VSD: variable state dimension.

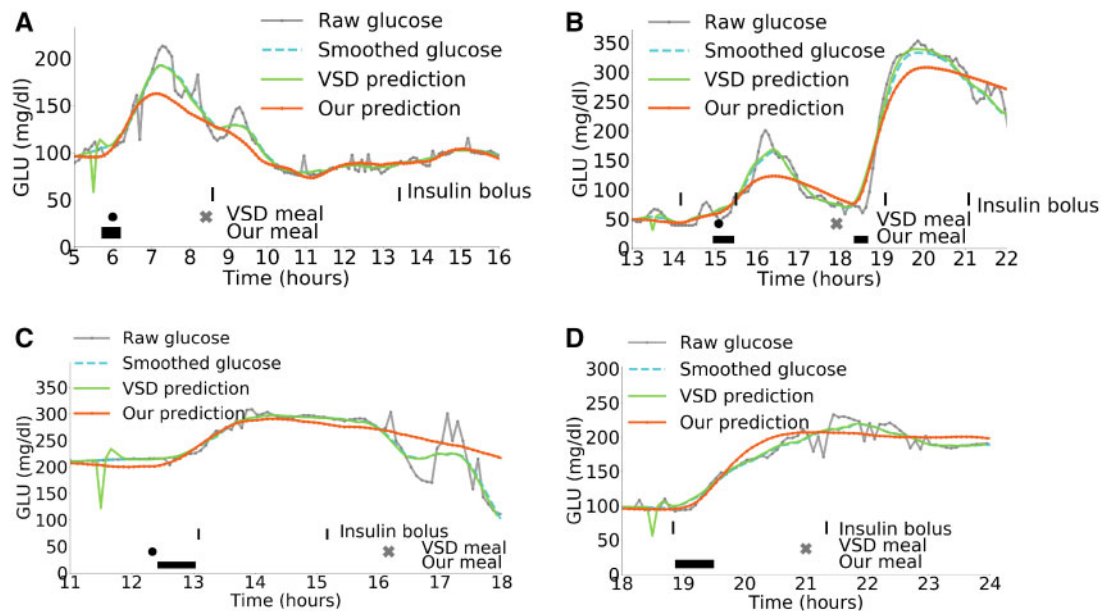[a]Best result for the metric and experiment.

Figure 3. Real-world continuous glucose (GLU) monitor (gray line) and insulin bolus data (vertical black bars) with accepted meals (black bar for our method or black circle for variable state dimension [VSD]) and rejected meals (gray bar for our method or gray X for VSD), smoothed GLU (dashed blue), our predicted GLU (orange line), and VSD-predicted GLU (green line). (A) Subject A (breakfast). (B) Subject B (snacks and dinner). (C) Subject C (lunch). (D) Subject D (dinner).

decreasing and almost an hour after an insulin bolus. For the dinner in Figure 3D, we detect a meal beginning at 7:00 PM This time is close to an insulin bolus and glucose increases right after the identified meal. In contrast, VSD detected a meal starting at 9:00 PM, which is well after the initial rise in glucose and just before a second insulin bolus.

Overall, while VSD detected 13 plausible meals, it also often (5 times) found a second false meal afterward, detected a meal with a significant delay, or detected meals while BG was decreasing (20 total rejected meals). Because VSD uses the error between predicted and observed glucose to detect the onset of a meal, when the error is negative and larger than the threshold, it erroneously detects a meal. However, detecting extra meals, especially if an insulin bolus has already been given, could lead to dangerously low BG if used in an AP.

DISCUSSION

Implications for diabetes treatment

Identifying meals automatically can have a substantial impact on treatment of T1D, by allowing automated insulin dosing in an AP,

or reminding individuals about missed insulin boluses with meals. While prior work has attempted to identify meals from CGM data, meals and their quantity have not been inferred together in real-world data that also include physical activity. Our approach outperformed VSD on both simulated and real-world data. We further showed that meal times and their quantities can be reliably inferred from CGM data even in the presence of physical activity in noisy real-world environments. The reduced delay and higher accuracy of our approach yields more actionable information for insulin dosing with fewer potentially dangerous false positives and better estimation of insulin needs. Integrating meal detection with an AP may allow a fully closed-loop system that does not have to wait for glycemic excursions to adjust insulin—reducing variation in BG. Our identification of meals can be used to prompt both reminders for missed meal boluses as well as suggestions for the size of the bolus needed. Beyond fully automated BG control via an AP, our work can be extended to help individuals learn to better manage their BG. One potentially significant area of future work is identifying not only missed, but also inaccurate (eg, too small) meal boluses. This can be done by simulating different sized boluses than the actual one given to identify the optimal quantity and by providing this feedback

to patients. Owing to the lifelong nature of T1D, a deployed system in the real world will have significantly more than 3 days of data on each individual. A promising avenue for future work is investigating personalization, which may be able to further reduce the time needed to detect a meal by adapting to each individual's glycemic response.

Incorporating other factors affecting glycemia

One of the reasons diabetes is so difficult to manage is that many factors, including stress,[28,29] sleep,[30] and hormonal cycles,[31,32] may influence glycemia. These factors are not accounted for by any existing meal detection algorithm. However, there are now wearable sensors and apps that can automatically measure or track these factors. For example, wristbands such as the Empatica E4 (Empatica, Milan, Italy) track stress with electrodermal activity, consumer devices such as the Fitbit (Fitbit, Inc, San Francisco, CA) track sleep, and there are numerous apps for tracking menstrual cycles.[33] Ultimately, we believe that these signals could be automatically integrated with our work to more accurately detect meals and provide better guidance on insulin dosage. While patients are mainly taught to count carbohydrates, meals with the same carbohydrate content but different protein and fat can lead to different glycemic responses.[34–36] Dosing suggestions and feedback can ultimately be improved by integrating our meal detection module and information on foods consumed such as from recently developed methods for automated dietary monitoring.[37,38]

CONCLUSION

We introduced SBE, a new approach to automated meal detection that combines simulation with observations of glucose and activity to robustly identify meals even in noisy data. On simulated data, we detect meals on average 10 minutes sooner than VSD, with an average error of only 1.2 g compared with 17.2 g in prior work. On real-world data, our approach finds a larger number of plausible meals while making fewer false positives. In future work, we aim to develop an online version of the approach that could be run efficiently as part of an AP and to incorporate data from other body-worn sensors that capture other factors that influence glucose (sleep, stress). Ultimately, our approach may be used to automatically provide meal information to an AP and remind individuals managing their own diabetes about missed insulin doses.

FUNDING

AUTHOR CONTRIBUTIONS

MZ and NB implemented the methods and conducted experiments. MZ and SK wrote the article.

SUPPLEMENTARY MATERIAL

Supplementary material is available at *Journal of the American Medical Informatics Association* online.

ACKNOWLEDGMENTS

CONFLICT OF INTEREST STATEMENT

None declared.

REFERENCES

1. Gale EA. The rise of childhood type 1 diabetes in the 20th century. *Diabetes* 2002; 51 (12): 3353–61.
2. Levey AS, Coresh J, Balk E, *et al*. National Kidney Foundation practice guidelines for chronic kidney disease: evaluation, classification, and stratification. *Ann Intern Med* 2003; 139 (2): 137–47.
3. Nathan DM, Group F. The diabetes control and complications trial/epidemiology of diabetes interventions and complications study at 30 years: overview. *Diabetes Care* 2014; 37 (1): 9–16.
4. Frier BM. Hypoglycaemia in diabetes mellitus: epidemiology and clinical implications. *Nat Rev Endocrinol* 2014; 10 (12): 711–722.
5. Dassau E, Bequette BW, Buckingham BA, Doyle FJ. Detection of a meal using continuous glucose monitoring. *Diabetes Care* 2008; 31 (2): 295–300.
6. Harvey RA, Dassau E, Zisser H, Seborg DE, Doyle FJ. Design of the glucose rate increase detector: a meal detection module for the health monitoring system. *J Diabetes Sci Technol* 2014; 8 (2): 307–20.
7. Peyrot M, Barnett AH, Meneghini LF, Schumm-Draeger P-M. Insulin adherence behaviours and barriers in the multinational global attitudes of patients and physicians in insulin therapy study. *Diabet Med* 2012; 29 (5): 682–9.
8. Simmons JH, Chen V, Miller KM, *et al*. Differences in the management of type 1 diabetes among adults under excellent control compared with those under poor control in the T1D exchange clinic registry. *Diabetes Care* 2013; 36 (11): 3573–7.
9. Dalla Man C, Raimondo DM, Rizza RA, Cobelli C. GIM, simulation software of meal glucose-insulin model. *J Diabetes Sci Technol* 2007; 1 (3): 323–330.
10. Dalla Man C, Breton MD, Cobelli C. Physical activity into the meal glucose-insulin model of type 1 diabetes: in silico studies. *J Diabetes Sci Technol* 2009; 3 (1): 56–67.
11. Entner D, Hoyer PO. On causal discovery from time series data using FCI. In: *Proceedings of the 5th European Workshop on Probabilistic Graphical Models*; 2010.
12. Adams RP, MacKay DJ. Bayesian online changepoint detection. *arXiv* 2007 Oct 19 [E-pub ahead of print].
13. Lee H, Bequette BW. A closed-loop artificial pancreas based on model predictive control: Human-friendly identification and automatic meal disturbance rejection. *Biomed Signal Process Control* 2009; 4 (4): 347–54.
14. Xie J, Wang Q. Meal detection and meal size estimation for type 1 diabetes treatment: a variable state dimension approach. In: *ASME 2015 Dynamic Systems and Control Conference*; 2015.
15. Grimm J. Exercise in type 1 diabetes. In: Nagi D, ed. *Exercise and Sport in Diabetes*. New York, NY: Wiley; 2005: 25–43.
16. Capes SE, Hunt D, Malmberg K, Pathak P, Gerstein HC. Stress hyperglycemia and prognosis of stroke in nondiabetic and diabetic patients: a systematic overview. *Stroke* 2001; 32 (10): 2426–32.
17. Turksoy K, Samadi S, Feng J, Littlejohn E, Quinn L, Cinar A. Meal detection in patients with type 1 diabetes: a new module for the multivariable adaptive artificial pancreas control system. *IEEE J Biomed Health Inform* 2016; 20 (1): 47–54.
18. Ramkissoon CM, Herrero P, Bondia J, Vehi J. Meal detection in the artificial pancreas: implications during exercise. *IFAC-Papers OnLine* 2017; 50 (1): 5462–7.
19. Weimer J, Chen S, Peleckis A, Rickels MR, Lee I. Physiology-invariant meal detection for type 1 diabetes. *Diabetes Technol Ther* 2016; 18 (10): 616–24.
20. Bergman RN. Toward physiological understanding of glucose tolerance: minimal-model approach. *Diabetes* 1989; 38 (12): 1512–27.
21. Xie J, Wang Q. A variable state dimension approach to meal detection and meal size estimation: in silico evaluation through basal-bolus insulin therapy for type 1 diabetes. *IEEE Trans Biomed Eng* 2017; 64 (6): 1249–60.

22. Kölle K, Fougner AL. Stavdahl Ø. Meal detection based on non-individualized moving horizon estimation and classification. In: *Proceedings of the 2017 IEEE Conference on Control Technology and Applications*; 2017.

23. Samadi S, Rashid M, Turksoy K, *et al*. Automatic detection and estimation of unannounced meals for multivariable artificial pancreas system. *Diabetes Technol Ther* 2018; 20 (3): 235–46.

24. Dalla Man C, Rizza RA, Cobelli C. Meal simulation model of the glucose-insulin system. *IEEE Trans Biomed Eng* 2007; 54 (10): 1740–9.

25. Breton MD. Physical activity—the major unaccounted impediment to closed loop control. *J Diabetes Sci Technol* 2008; 2 (1): 169–74.

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

27. Rahman SA, Huang Y, Claassen J, Heintzman N, Kleinberg S. Combining Fourier and lagged k-nearest neighbor imputation for biomedical time series data. *J Biomed Inform* 2015; 58: 198–207.

28. Riazi A, Pickup J, Bradley C. Daily stress and glycaemic control in Type 1 diabetes: individual differences in magnitude, direction, and timing of stress-reactivity. *Diabetes Res Clin Pract* 2004; 66 (3): 237–44.

29. Lloyd CE, Dyer PH, Lancashire RJ, Harris T, Daniels JE, Barnett AH. Association between stress and glycemic control in adults with type 1 (insulin-dependent) diabetes. *Diabetes Care* 1999; 22 (8): 1278–83.

30. Feupe SF, Frias PF, Mednick SC, McDevitt EA, Heintzman ND. Analysis of nocturnal continuous glucose and sleep stage data in adults with type 1 diabetes in real-world conditions. *J Diabetes Sci Technol* 2013; 7 (5): 1337–45.

31. Trout KK, Rickels MR, Schutta MH, *et al*. Menstrual cycle effects on insulin sensitivity in women with type 1 diabetes: a pilot study. *Diabetes Technol Ther* 2007; 9 (2): 176–82.

32. Goldner WS, Kraus VL, Sivitz WI, Hunter SK, Dillon JS. Cyclic changes in glycemia assessed by continuous glucose monitoring system during multiple complete menstrual cycles in women with type 1 diabetes. *Diabetes Technol Ther* 2004; 6 (4): 473–80.

33. Lee J, Kim J. Can menstrual health apps selected based on users' needs change health-related factors? A double-blind randomized controlled trial. *J Am Med Inform Assoc* 2019; 26 (7): 655–66.

34. Paterson M, Bell KJ, O'Connell SM, Smart CE, Shafat A, King B. The role of dietary protein and fat in glycaemic control in type 1 diabetes: implications for intensive diabetes management. *Curr Diabetes Rep* 2015; 15 (9): 61.

35. Peters AL, Davidson MB. Protein and fat effects on glucose responses and insulin requirements in subjects with insulin-dependent diabetes mellitus. *Am J Clin Nutr* 1993; 58 (4): 555–60.

36. Bell KJ, Smart CE, Steil GM, Brand-Miller JC, King B, Wolpert HA. Impact of fat, protein, and glycemic index on postprandial glucose control in type 1 diabetes: implications for intensive diabetes management in the continuous glucose monitoring era. *Diabetes Care* 2015; 38 (6): 1008–15.

37. Mirtchouk M, Merck C, Kleinberg S. Automated estimation of food type and amount consumed from body-worn audio and motion sensors. In: *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*; 2016.

38. Mirtchouk M, Deierlein AL, Kleinberg S. Automated estimation of food type from body-worn audio and motion sensors in free-living environments. In: *Machine Learning for Healthcare*; 2019.