

Committed Moving Horizon Estimation for Meal Detection and Estimation in Type 1 Diabetes

Hongkai Chen[†], Nicola Paoletti[‡], Scott A. Smolka^{*} and Shan Lin[†]

Abstract—We introduce a model-based meal detection and estimation method for the treatment of type 1 diabetes that automatically detects the occurrence and estimates the amount of carbohydrate (CHO) intake from continuous glucose monitor (CGM) data. Meal detection and estimation play a critical role in closed-loop insulin control by enabling automatic regulation of post-meal insulin dosing in artificial pancreas systems without manual meal announcements by the patient. Our approach to meal detection is based on a novel technique we call *Committed Moving Horizon Estimation* (CMHE), an extension of *Moving Horizon Estimation* (MHE). While MHE alone is not well-suited for disturbance estimation and meal detection, CMHE aggregates the meal disturbances estimated by multiple MHE instances to balance future and past information at decision time, thus providing timely detection and accurate estimation. We evaluated our CMHE-based meal detection and estimation method in-silico, using a nonlinear ODE gluco-regulatory model and random meal profiles to generate blood glucose and CGM signals. CGM data is used to detect meal occurrences and to estimate their onset, duration, and CHO amount. At the optimal operating point of the detector, we achieve an 88.5% daily detection rate and, more importantly, a 100% detection rate, with an average of 18.86 minutes onset deviation, and 70.50% CHO amount estimation accuracy for the main meals (i.e., excluding snacks).

I. INTRODUCTION

Type 1 Diabetes Mellitus (T1DM) is an autoimmune disease in which the human pancreas is unable to produce a sufficient amount of insulin to regulate *blood glucose* (BG) levels. In healthy subjects, insulin is released in amounts commensurate with current BG levels, and circulating insulin promotes glucose uptake in muscle and adipose (fatty) tissue. This process maintains BG within a healthy, safe range (70–180 mg/dL) [1]. In T1DM, insufficient insulin causes *hyperglycemia* (high BG), a condition that if untreated, can lead to health issues such as cardiovascular disease, kidney damage, and blindness.

In the U.S., approximately 30.3 million people have diabetes, about which 5–10% is T1DM [2]. T1DM patients require everyday insulin therapy to maintain healthy BG levels. The *artificial pancreas* (AP) [3] is a system for automated, closed-loop insulin delivery, consisting of a *continuous glucose monitor* (CGM) that provides readings of

subcutaneous glucose levels, a wearable insulin pump for the infusion of insulin, and control algorithms for computing the insulin amount that best keeps the BG in range.

Designing a fully closed-loop AP is, however, difficult: meals are a major (and unknown) disturbance, as they increase BG rapidly unless a timely and adequate insulin dose is administered. Indeed, no commercial AP device exists that is able to automatically regulate post-meal glucose. This means that every T1DM patient has to announce, at every meal, the amount of ingested carbohydrates (CHO) so that the appropriate insulin dose can be delivered. This manual procedure is not just a burden on the patient but also inherently dangerous, as inaccurate or delayed CHO information can lead to an incorrect insulin dosage. In particular, insulin overdose causes *hypoglycemia* (low BG), which may result in coma, brain damage, or even death.

This paper focuses on *automated meal detection and estimation* (MDE), i.e., the problem of detecting, only from CGM measurements, the occurrence of meals in the recent past, as well as estimating their time of occurrence and amount of ingested CHO. Accurate and timely MDE has the potential to replace meal announcements, thus enabling fully closed-loop insulin therapy.

Our approach to MDE is based on a novel technique that we call *Committed Moving Horizon Estimation* (CMHE), which allows us to estimate both the timing and size of unannounced meals, information that is vital to achieve closed-loop insulin control. As the name implies, CMHE is based on *Moving Horizon Estimation* (MHE) [4], a constrained optimization technique for state estimation that given a plant model and a bounded history of observed outputs (or measurements), determines the sequence of system states and disturbances that minimizes the discrepancy between predicted and observed outputs. In the MDE context, the model describes the human gluco-regulatory system and the disturbances are the meal amounts that we seek to estimate. MHE has been previously used for state estimation in insulin control for T1DM [5]–[7], but not as a disturbance estimation technique.

MHE provides a principled way to deal with meal disturbance estimation. First, MHE has an important probabilistic interpretation, as it (roughly) corresponds to maximizing the likelihood of state and disturbance estimations given the observed outputs [8]. In contrast, Kalman filters are optimal in this sense only if disturbances are normally distributed, which is not the case for meals in T1DM. Second, MHE allows one to incorporate additional constraints on the size, onset, duration, and shape of the meal disturbances, in such

[†]Hongkai Chen and Shan Lin are with the Electrical and Computer Engineering Department, Stony Brook University, Stony Brook, NY, 11794, USA {hongkai.chen, shan.x.lin}@stonybrook.edu

[‡]Nicola Paoletti is with Department of Computer Science, Royal Holloway, University of London, UK Nicola.Paoletti@rhul.ac.uk

^{*}Scott Smolka is with the Department of Computer Science, Stony Brook University, Stony Brook, NY, 11794, USA sas@cs.stonybrook.edu

a way that only realistic meals are considered. For example, one can incorporate bounds on the meal-intake profile based on patient data [5]. Third, MHE provides a unique solution to both state and meal estimation in insulin control, even though the present paper focuses on meal estimation only.

MHE alone, however, is not sufficient for accurate meal estimation. An individual MHE instance (i.e., a standard MHE optimization problem) executed at time t returns N disturbance estimates, one for each time point in the MHE window $[t-N, t-1]$. Here N is the MHE window size; i.e., the number of time steps in the past (and thus the number of past measurements) we consider for estimation. This implies that one has N admissible estimates for the disturbance value at any given time point t , because t is included in the time windows of all the MHE instances executed at times $t+1, \dots, t+N$. Thus, the problem arises of how to select a value for the disturbance at time t out of these N candidates, a problem that MHE alone cannot solve. Below we argue that we cannot commit to only one of the N estimates, and discuss how our method addresses this issue.

Let us denote with $\delta_t^{t+1}, \dots, \delta_t^{t+N}$ the estimates for the disturbance at time t obtained through the $t+1, \dots, t+N$ -th MHE instances, respectively. Committing to δ_t^{t+1} only (i.e., estimate of the $t+1$ -th MHE instance) is not a suitable strategy because the corresponding MHE instance does not consider measurements after time t , and thus there is no evidence of the effects of the disturbance on the dynamics. In glucose metabolism, like many other control systems, the effect on the output is not immediate, but rather depends on delays due to digestion and transport from the blood to the subcutaneous compartment where measurements are taken. In other words, this strategy fails because it does not have enough look-ahead.

On the other hand, committing to the disturbance with the most look-ahead, δ_t^{t+N} , is not ideal either because this strategy ignores past information; i.e., the series of past measurements and disturbances outside the prediction window that led to the state at t .¹ Indeed, on observing a glucose increase at time t it is impossible, without past information, to decide whether the increase was due to a meal at time t or at an earlier time. Even more importantly, using the estimate of the $t+N$ -th MHE instance, this strategy prevents timely MDE as it requires a time delay of N in order to produce the final estimate for the current time. In our context, the MHE window size N is of the order of hours, a delay that is unacceptable for meal-detection purposes, whereas choosing a smaller N would affect the performance of MHE and the quality of its estimations.

Our proposed *Committed MHE* (CMHE) is an extension of MHE that solves the above commitment problem. CMHE considers a so-called *commitment level* $V < N$, and instead of committing to only one estimate, obtains the final disturbance value at time t by combining (e.g. via a weighted average) the estimates of the V MHE instances from time

$t+1$ to $t+V$. Thus, CMHE overcomes the above problems because it accounts for both future and past information, as it considers MHE estimations having up to time V of look-ahead (δ_t^{t+V}), and up to time N of history (δ_t^{t+1}). See also Figure 1 for an illustration of CMHE. Another key advantage of CMHE is that it provides timely detection, with *guaranteed and adjustable delay* V .

The idea of CHME is inspired by [9], where the authors apply the the principle of commitment level to finite-horizon control problems. To the best of our knowledge, we are the first to apply this principle to the estimation problem.

In summary, the main contributions of this paper are the following:

- We introduce a model-based meal detection method for artificial pancreas systems that automatically announces carbohydrate intake and estimates the amount. Compared to existing solutions that focus on the detection of the occurrence of a meal, our method also estimates the size of a meal, which is essential for the safe and effective operation of a closed-loop insulin pump.
- We propose a novel technique called Committed Moving Horizon Estimation that derives the final meal estimate from the candidates produced by the MHE instances. CMHE crucially ensures that the final detection decision occurs within a fixed and adjustable delay.
- We design an online detection algorithm that filters out potential noise from CMHE estimates to derive the start time, duration, and CHO amount of meals.
- We evaluate our approach *in silico* using synthetic CGM measurements generated from a high-fidelity gluco-regulatory model [10] and randomly generated meal profiles. We achieve a 88.5% overall detection rate and most importantly, for large meals (i.e., excluding snacks), a 100% detection rate, with an average of 18.86 minutes of onset deviation, and 70.50% CHO amount estimation accuracy.

The structure of the rest of the paper is the following. Section II provides relevant background information. Section III introduces our Committed Moving Horizon Estimation technique. Section IV presents our experimental results. Section V considers related work. Section VI offers our concluding remarks and directions for future work.

II. BACKGROUND

A. CGM measurements from virtual patient

CGM measurements for our experiments are generated from Hovorka’s well-established virtual patient model [10], a system of non-linear ODEs describing the human glucose-insulin metabolism across several physiological compartments (e.g., gut, blood, subcutaneous). The model is executed in closed-loop with the MPC-based insulin controller of [5]. The CGM is derived by adding Gaussian noise to the subcutaneous glucose variable of the model. CGM measurements are produced with period of 5 minutes. Since, as we will see, our MDE method works with a 1-minute resolution, we interpolate the CGM signal using a Savitzky-Golay filter.

¹The arrival cost term in the MHE objective function summarizes past information about the measurements, but does not capture past disturbances; see Eq. 9.

B. Estimation model

For efficiency reasons, linear models are often used to predict and estimate the physiological state of the patient for insulin control and state estimation algorithms in the AP. For MHE, we also employ a linear model, an extension of the model presented in [11], which in turn builds on the well-known Bergman's minimal model [12]. The model equations are given below:

$$\dot{G}(t) = -p_1 \cdot G(t) - p_2 \cdot I(t) + p_3 + p_4 \frac{m(t)}{t_G V_G} \quad (1)$$

$$\dot{C}(t) = \frac{G(t) - C(t)}{t_{G,int}} \quad (2)$$

$$\dot{g}(t) = A_G \cdot D_G(t) - \frac{g(t)}{t_{\max,G}} \quad (3)$$

$$\dot{m}(t) = \frac{g(t) - m(t)}{t_{\max,G}} \quad (4)$$

$$\dot{I}(t) = -k_e \cdot I(t) + \frac{k_a}{V_I} x(t) \quad (5)$$

$$\dot{x}(t) = -k_a \cdot x(t) + u(t) \quad (6)$$

where $G(t)$ is the BG concentration (mmol/L); $C(t)$ is the subcutaneous glucose concentration; $m(t)$ (mmol/min) is the rate of BG appearance; $g(t)$ (mmol/min) is the glucose in the gut compartment; $D_G(t)$ (mmol/min) is the CHO input (i.e., the disturbance); $u(t)$ (mU/min) is the insulin input (i.e., the control input); $x(t)$ (mU) is the insulin mass in the subcutaneous compartment; and $I(t)$ is the plasma insulin concentration (mU/L). V_G and V_I (L) are the glucose and insulin distribution volumes, known parameters that depend on the body weight. $t_G = 30$ min and $t_{G,int} = 8$ min are, respectively, the delays for BG appearance in the blood and for glucose transport from blood to the subcutaneous compartment.

Parameters $p_1, \dots, p_4, k_a, k_e$ are typically estimated from patient data. In our case, they are estimated from open-loop trajectories of the virtual patient model of Section II-A.

C. Moving Horizon Estimation

Consider a dynamical system described by the equations

$$x^+ = f(x, u, d) \quad (7)$$

$$y = h(x) + v \quad (8)$$

where $x \in \mathbb{R}^n$ is the state of the system; x^+ denotes x at the next sample time; $y \in \mathbb{R}^m$ is the (noisy) measurement; u is a known system input; $d \in \mathbb{R}^g$ is the process disturbance; and v is the measurement noise. In our case, f is the AP model of (1)-(6), y is the CGM value, h returns the subcutaneous glucose variable C (2), u is the insulin input, and d is the meal disturbance. In the MHE context, d , v , and the initial system state, x_0 , are unknown.

MHE [13] seeks to estimate the trajectory of states x using only a finite set of measurements y by solving an optimization problem that minimizes the error between measured outputs and model-predicted outputs. In estimating the state trajectory, MHE also estimates the sequence of unknown

disturbances that yield the best fit between measurements and model predictions. This allows us to detect meals and estimate their size.

In the following, we will use the notation $x_{i,\dots,i+j}$ to denote the indexed sequence x_i, \dots, x_{i+j} . Analogous notation applies to all variables. We will use Greek letters χ, ν, δ for the MHE variables in order to distinguish them from the system variables x, v, d .

The MHE problem solved at time t , also called the t -th MHE (instance), seeks to find the sequence of states $\chi_{t-N:t}$ and sequence of disturbances $\delta_{t-N:t-1}$ that minimize the following cost function:

$$C_{\text{MHE}}^t(\chi_{t-N}, \delta_{t-N,\dots,t-1}) = \mu \cdot \|\chi_{t-N} - \hat{x}_{t-N}\|^2 + \sum_{k=1}^N \frac{\|\nu_{t-k}\|^2}{q_{t-k}} \quad (9)$$

subject to

$$\nu_{t-k} = y_{t-k} - h(\chi_{t-k}), \quad k = N, \dots, 0 \quad (10)$$

$$\chi_{t-k+1} = f(\chi_{t-k}, u_{t-k}, \delta_{t-k}), \quad k = N, \dots, 1 \quad (11)$$

where N is the MHE window size; δ_{t-k} is the estimated disturbance at time $t-k$; (10) defines the measurement discrepancy ν_{t-k} at time $t-k$ as the difference between the measured output y_{t-k} and predicted output $h(\chi_{t-k})$; and (11) states that the MHE state variable χ evolves according to the system dynamics f (1)-(6), where u_{t-k} is the (known) insulin input.

The first summand of the cost function C_{MHE}^t given in (9) is the so-called arrival cost, which penalizes the discrepancy between the MHE state at the start of the window, χ_{t-N} , and the optimal estimate of the state at time $t-N$, \hat{x}_{t-N} , obtained by the $t-N$ -th MHE. This term is designed to summarize the information about the past, out-of-window, measurements $y_{0,\dots,t-N+1}$. $\mu \in \mathbb{R}^+$ is a weighting factor. The second summand of (9) accounts for the discrepancy between measured and predicted outputs. $q_{t-k} \in \mathbb{R}^+$ is a weighting factor, which is typically chosen to reflect the variance (known or estimated) of the measurement noise v at time $t-k$.

Note that the cost function depends only on the first state of the MHE window because the subsequent states $\chi_{t-N+1,\dots,t}$ are automatically determined by the deterministic dynamics f .

III. METHOD

A. Committed MHE

Let us denote the t -th MHE decision variables by $\chi_{t-N,\dots,t}^t$, $\delta_{t-N,\dots,t-1}^t$, and $\nu_{t-N,\dots,t}^t$, where the superscript index indicates the corresponding MHE instance.

In MHE-based state estimation, the decision variable χ_t^t is selected as the estimated state at time t . One could be tempted to apply the same strategy for disturbance estimation, and select the most recently estimated disturbance value, δ_t^{t+1} , as the *final disturbance estimate* at time t . We denote the latter with Δ_t . This strategy, however, is

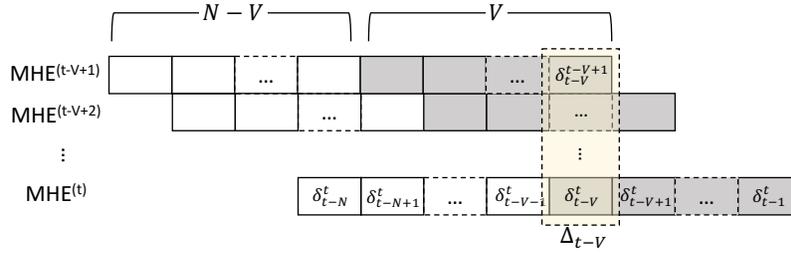


Fig. 1. Illustration of Committed MHE (CMHE), with MHE window size N and commitment level V . At each step t , CMHE computes the final disturbance estimate at time $t - V$, Δ_{t-V} , by aggregating the V estimates at time $t - V$, $\delta_{t-V}^t, \dots, \delta_{t-V+1}^t$, of the last V MHE instances.

not adequate because it does not have look-ahead, i.e., it ignores the delay in glucose appearance due to digestion and transport from the blood to the subcutaneous compartment where measurements are taken. Then, one could follow the opposite strategy, and wait for the $t + N$ -th MHE in order to set Δ_t to the estimate with the longest look-ahead, δ_t^{t+N} . This choice is not ideal either, as δ_t^{t+N} lies at the start of the $t + N$ -th MHE prediction window, and thus ignores information about the BG trajectory and disturbances before time t ; i.e., it does not account for history. This might raise false alarms, when MHE wrongly detects a meal on observing a BG increase at time t , which is instead caused by a meal happened before t . Most importantly, this strategy prevents timely MDE, because there is an estimation delay of N minutes, which is normally hours. In summary, MHE does not provide a clear mechanism to select a value for Δ_t out of the N disturbances at time t , $\delta_t^{t+1}, \dots, \delta_t^{t+N}$, respectively estimated by the $t + 1, \dots, t + N$ -th MHE.

Another way to view this problem is by *commitment*. Both the above strategies commit to only one disturbance estimate at each timestep, and that will be the final decision for Δ_t . To solve this problem, we propose a novel method called *Committed Moving Horizon Estimation (CMHE)*. Our idea, shown in Fig. 1, is to commit not just to one disturbance but to multiple disturbances, in this way balancing between look-ahead and history. In addition to the MHE window size N , CMHE uses another parameter called *commitment level* and denoted by V , such that $V \leq N$. The commitment level describes both how many disturbance estimates CMHE commits to and the delay for computing the final estimation decision. At time t , CMHE performs the following two steps. First, it executes the t -th MHE instance, producing a sequence of estimates for times $t - N, \dots, t - 1$, shown in the last row of Fig. 1. Second, it returns the final estimate for time $t - V$, Δ_{t-V} , by combining the disturbances estimated for the same time point by the $t - V + 1, \dots, t$ -th MHE instances, $\delta_{t-V}^t, \dots, \delta_{t-V+1}^t$, i.e., V estimates in total. The second step is illustrated by the yellow column of Fig. 1.

We introduce a *time-based* aggregation strategy, by which Δ_{t-V} is derived as a weighted average of the V MHE estimates as follows:

$$\Delta_{t-V} = \frac{\left(\sum_{i=t-V+1}^t W(i)^b \cdot \delta_{t-V}^i \right)}{\sum_{i=t-V+1}^t W(i)^b}, \quad (12)$$

where $b \in \mathbb{R}^{\geq 0}$ is a parameter that determines the impor-

tance of the weights. Note that committing only to the last estimated disturbance, i.e., setting $\Delta_{t-1} = \delta_{t-1}^t$, is a special case of CMHE with estimation window N and commitment level $V = 1$. The weights $W(i)$ are defined in such a way to prioritize estimations having a good balance between history (i.e., number of past measurements) and look-ahead (i.e., number of future measurements). For $i = t - V + 1, \dots, t$, the length of history is $N + t - V - i$, and the length of look-ahead is $i + V - t - 1$. Define the absolute deviation between history and look-ahead lengths as $\epsilon(i) = |N + 2(t - i - V) + 1|$. Note that the maximum deviation for all $i = t - V + 1, \dots, t$ is $N - 1$ when $i = t - V + 1$, i.e., when the look-ahead length is 0 and the history length is the highest (which happens in the earliest MHE instance, see Fig. 1); and the minimum deviation is 0 when the two lengths are equal.² To achieve such balance, $W(i)$ should grow as $\epsilon(i)$ decreases, and thus, we define $W(i)$ by the following linear equation:

$$W(i) = N - \epsilon(i),$$

which assigns a maximum weight of N when $\epsilon(i)$ is minimum, and a minimum weight of 1 when $\epsilon(i)$ is maximum.

In CMHE with window size N and commitment level V , at time t , all the timeslots before $t - V$ are already committed, i.e., the final disturbance values Δ_i for $i < t - V$ are already determined. To encourage agreement with these previous estimates, we thus introduce a term in the cost function penalizing the discrepancy between $\delta_{t-N, \dots, t-V-1}^t$ and $\Delta_{t-N, \dots, t-V-1}$. Therefore, the CMHE cost function is defined as:

$$C^t(\chi_{t-N}, \delta_{t-N, \dots, t-1}) = C_{\text{MHE}}^t(\chi_{t-N}, \delta_{t-N, \dots, t-1}) + \eta \cdot \sum_{k=V+1}^N \|\delta_{t-k}^t - \Delta_{t-k}\|^2 \quad (13)$$

where $\eta > 0$ is a weighting factor determining the importance of agreeing with the final disturbance estimates.

In summary, the t -th MHE instance in our CMHE approach is defined as following problem:

$$\min_{\substack{\chi_{t-N, \dots, t}, \\ \delta_{t-N, \dots, t-1}^t}} C^t(\chi_{t-N}, \delta_{t-N, \dots, t-1}) \quad (14)$$

subject to (10) and (11).

²It is easy to see that there exists an i for which $\epsilon(i) = 0$ only if $V \geq (N + 1)/2$. Otherwise, all estimations have indeed strictly more history than lookahead, in which case the minimal $\epsilon(i)$ is obtained for the latest MHE instance. See also Fig. 1.

B. Estimating realistic meal profiles

We further extend the MHE instance (14) to capture realistic meal profiles, assuming that the CHO intake rate during a meal is constant. This leads to a representation of the meal disturbance as a square-wave shape impulse. Another sensible assumption is that, in an estimation window of size N , no more than $m < N$ meals can occur. Under these assumptions, we can formulate the MHE optimization problem as a *Mixed-integer Quadratic Programming* (MIQP) as follows.

$$\min_{\substack{\chi_{t-N, \dots, t}, A_{1, \dots, m} \\ \varphi_{t-N, \dots, t-1}^1, \dots, \varphi_{t-N, \dots, t-1}^m}} C^t(\chi_{t-N}, \delta_{t-N, \dots, t-1}) \quad (15)$$

subject to, for $k = N, \dots, 1$ and $i = 1, \dots, m$:

$$v_{t-k} = y_{t-k} - h(\chi_{t-k}) \quad (16)$$

$$\chi_{t-k+1} = f(\chi_{t-k}, u_{t-k}, \delta_{t-k}^t) \quad (17)$$

$$\varphi_{t-k}^i \in \{0, 1\} \quad (18)$$

$$\sum_{i=1}^m \varphi_{t-k}^i \leq 1 \quad (19)$$

$$\lambda_{t-k}^i = I(\varphi_{t-k+1}^i - \varphi_{t-k}^i = 1), \quad k = N, \dots, 2 \quad (20)$$

$$\sum_{k=2}^N \lambda_{t-k}^i \leq 1 \quad (21)$$

$$\delta_{t-k}^t = \sum_{i=1}^m A_i \cdot \varphi_{t-k}^i \quad (22)$$

where $\varphi_{t-k}^i = 1$ if the i -th meal is ongoing at time $t-k$, 0 if not; and A_i is the ingested CHO amount per minute of the i -th meal. (16) and (17) are the MHE constraints on measurement discrepancy and time evolution of χ . (19) indicates that multiple meals cannot exist at the same time. In (20), we define λ_{t-k}^i as a Boolean variable indicating whether or not meal i starts at time $t-k+1$ (I is the indicator function). Equation (21) states that each meal can start at most once. The last constraint, (22), states that, for any meal i , the meal disturbance at time $t-k$ equals to the corresponding ingestion rate A_i if meal i is happening at time $t-k$. Note that the superscript i in φ^i and λ^i represents the i -th potential meal, while superscript t in the cost function C^t and variable δ^t represents the time of the MHE instance.

MIQP is, in general, computationally demanding, and thus, for implementing the method on resource-constrained devices, one could substitute it with a more efficient, albeit sub-optimal solution method. We remark, however, that computational constraints are not the focus of our work. Nevertheless, given the rapid improvement in modern embedded and mobile devices, we expect our algorithm to perform efficiently also on such hardware platform.

C. Online detection algorithm

We designed an online detection algorithm, Algorithm 1, for extracting meal onset, duration, and total CHO amount out of the estimations of meal disturbances produced by CMHE. The state of the algorithm is characterized by the

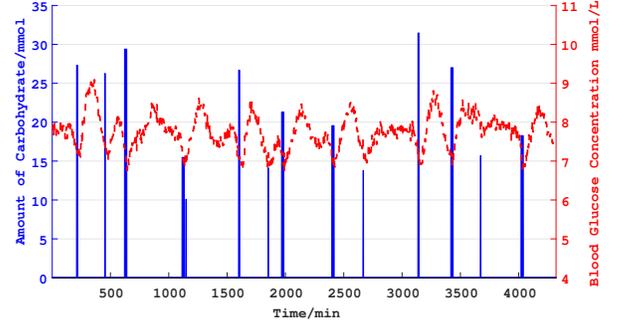


Fig. 2. One of the random meal profiles and corresponding CGM measurements used in our experiments.

variables `onset`, keeping the start time of the most recently detected meal, and `inMeal`, which is true when the beginning of a meal has been detected, but not yet its ending, i.e., the current time is within the meal duration.

When `inMeal` is false, the algorithm detects the onset of a meal if: 1) the current CMHE estimate Δ_t is above a threshold $\bar{\Delta}$ (see line 4); and 2) the sequence of CMHE estimates is non-decreasing for at least 80% of the last w estimates. Variable `num_inc` maintains the number of points in the window $t-w, \dots, t$ where the sequence is non-decreasing (line 5). When both conditions hold, the meal onset is set to $t-w$ and `inMeal` is updated to true. These two conditions are designed to reduce false alarms: the threshold in 1) allows to filter out noise in the estimation signal; 2) ensures that the amount of ingested CHO is actually rising.

Detection of meal offset works in an analogous way. When `inMeal` is true, if Δ_t is below $\bar{\Delta}$ and the sequence of CMHE estimates is non-increasing for at least 80% of the last w estimates, then we mark the current time as the end of the meal and output the corresponding onset, duration and size (i.e., total amount of CHO between onset and offset).

Parameter $\bar{\Delta}$ can be tuned to achieve different trade-offs between false alarms and detection rate. In our experiments, we set w to 5 min.

Our meal detection algorithm receives CMHE estimates with a period of one minute. When the measurement period is longer, as for CGM measurements that have a period of five minutes, we interpolate the missing intermediate measurements.

IV. RESULTS

This section presents an evaluation of the performance of our CMHE-based meal detection and estimation method. We evaluate the method using 10 repetitions of a 3-day experiment, 30 days in total, with randomly generated meal profiles. Our experiments use the CGM measurements described in Section II-A and two types of random meals (large meals and small snacks). See the top half of Table I for details of the random meal profiles and Fig 2 for a meal profile realization and corresponding CGM trajectory. We performed the experiments using MATLAB and YALMIP [15].

Algorithm 1: Online meal detection from CMHE meal disturbance estimates

```

input : threshold  $\bar{\Delta} \in \mathbb{R}$ ,
         meal rise/fall window  $w \in \mathbb{N}^+$ ,
         stream of CMHE estimates  $\Delta$ 
output: stream of meal onset times, durations, and sizes
1 inMeal  $\leftarrow$  false; onset  $\leftarrow$  0;  $t \leftarrow w + 1$ ;
2 while true do
   /* Collect estimated meals at times
       $t - w - 1$  to  $t$  from stream  $\Delta$  */
3  $\Delta_{t-w-1}, \dots, \Delta_t \leftarrow \text{collect}(\Delta, t - w - 1, t)$ ;
4 if  $\neg \text{inMeal} \wedge \Delta_t > \bar{\Delta}$  then
5   num_inc  $\leftarrow \sum_{i=t-w}^t I(\Delta_i - \Delta_{i-1} \geq 0)$ ;
6   if num_inc  $\geq 0.8 \cdot w$  then
7     /* Meal start detected */
8     onset  $\leftarrow t - w$ ; inMeal  $\leftarrow$  true
9   end
10 end
11 if inMeal  $\wedge \Delta_t \leq \bar{\Delta}$  then
12   num_dec  $\leftarrow \sum_{i=t-w}^t I(\Delta_i - \Delta_{i-1} \leq 0)$ ;
13   if num_dec  $\geq 0.8 \cdot w$  then
14     /* Meal end detected */
15     duration  $\leftarrow t - \text{onset}$ ; size  $\leftarrow \sum_{i=\text{onset}}^t \Delta_i$ ;
16     yield (onset, duration, size);
17     inMeal  $\leftarrow$  false;
18   end
19 end

```

We compare our time-based weighted CMHE described in (12) using $b = 1/2$, $N = 180$ and $V = 40$, with a method that does not combine multiple estimations and that at time t , commits only to the $t - V$ -th estimate, i.e., $\Delta_{t-V} := \delta_{t-V}^t$. That is, the latter method corresponds to using MHE alone and not CMHE.

In Fig. 3, we show the *receiver operating characteristic* (ROC) of the two methods. The ROC curves present the detection rate versus the false alarm rate per day under varying detection thresholds $\bar{\Delta}$ (see Algorithm 1). For each detection method, we report two ROC curves, one for the detection of the large meals only, and one for both large meals and snacks. When the large meals are the only targets, the detector can apply a higher threshold $\bar{\Delta}$ to reduce the false alarms without an attempt to detect snacks with small CHO amount. We observe that CMHE considerably outperforms the simple MHE strategy.

The performance of the CMHE algorithm is a trade-off between the detection rate and the daily false alarm rate. Since the risks associated with hypoglycemia are higher than those associated with hyperglycemia, false alarms of meals and snacks should be considered a more critical event, and their reduction should be an important target. For this purpose, one can always choose another operating point with a higher detection threshold $\bar{\Delta}$, decreasing the detection rate

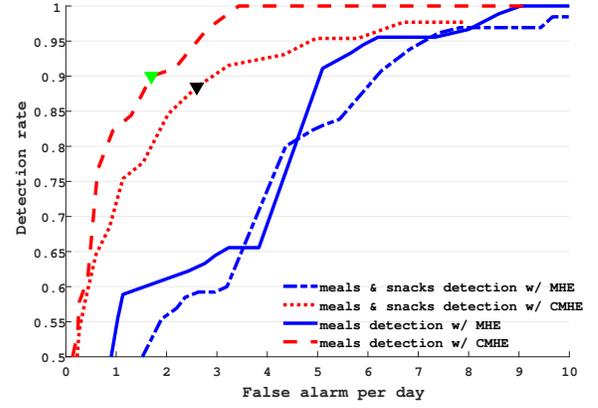


Fig. 3. ROC curve comparison between our CMHE method (red lines) and a simple MHE strategy (blue), with different targeted meal types (main meals only VS main meals and snacks). Triangles represent the optimal operating points for CMHE [14].

TABLE I
CHARACTERISTICS AND CORRESPONDING MDE RESULTS FOR THE RANDOMLY GENERATED MEALS. FOR EACH DAY OF SIMULATION, WE CONSIDER THREE LARGE MEALS (BREAKFAST, LUNCH, DINNER) AND THREE SMALL SNACKS. MEAL CHO AMOUNTS AND STARTING TIMES ARE SAMPLED UNIFORMLY FROM THE REPORTED INTERVALS.

	breakfast	snack1	lunch	snack2	dinner	snack3
Probability of Occurrence	100%	50%	100%	50%	100%	50%
CHO amount (g)	40-60	5-25	70-110	5-25	55-75	5-15
Time of day (h)	1:00-5:00	5:00-8:00	8:00-12:00	12:00-15:00	15:00-19:00	19:00-21:00
MDE performance						
Onset Deviation (min)	22.43	30.67	17.59	25.00	20.68	12.04
CHO Deviation (g)	50.53	20.20	25.54	18.33	22.85	26.93
Detection rate	100%	92.31%	100%	91.67%	100%	27.78%

in exchange for a reduction in false alarms.

We now discuss the performance of CMHE regarding the estimation of both large meals and snacks. Results for each kind of meal are reported on the bottom half of Table I. At the optimal operating point (black triangle in the ROC curve of Figure 3), overall we have 88.5% total detection rate, with an average of 19.99 minutes meal onset deviation and 58.14% CHO amount estimation accuracy. The estimation accuracy tells how close our CHO estimation is to the actual CHO intake of the meal, weighted by the CHO value of that meal. In these settings, we have an average of 2.6 false alarms per day, with 11.08g CHO each.

The reason why the algorithm cannot detect all of the meals and estimate their amount accurately enough is rooted in our complicated random meal profiles. Indeed, for each day we have three CHO-rich meals (breakfast, lunch and dinner) and at most three snacks with small CHO amounts.

In particular, the CHO amount of snacks can be orders of magnitude smaller than that of a regular meal. In order to detect these small snacks, we need to use low values of the detection threshold $\bar{\Delta}$, and thus, the algorithm inevitably incorrectly considers some small detection errors as snacks, leading to an increased number of false alarms. What makes detection even more difficult is the randomized onset times of snacks and meals, which make it possible for them to be very close or even overlapping. In these cases, it is very difficult if not impossible to distinguish a snack from a meal. Under such extreme conditions, however, we are still able to detect 100% of the main meals, and 92.31%, 91.67% and 27.78% of the three snacks, respectively.

In reality, CHO-rich meals lead to much rapid and greater BG increases than smaller meals and snacks, and thus, it is much more dangerous for patients if larger meals are not accurately estimated or even undetected. Therefore, accurate estimation of large meals is of utmost importance for the automation of insulin therapy in diabetes. While we have shown that our algorithm is able to detect multiple types of meals, our focus is on those with larger CHO amount.

At the optimal operating point (green triangle in the ROC curve of Figure 3), we achieve 90% overall detection rate and an average of 1.7 false alarms per day. For the detection for three main meals (breakfast, lunch and dinner), we achieve a 100% detection rate, an average of 19.76 minutes onset deviation and an average of 57.40% CHO estimation accuracy. Most importantly, if we further consider the detection for only lunch and dinner (the most prominent meals), we reduce the onset deviation to 18.86 minutes, and increase the CHO estimation accuracy to 70.50%.

Typical detection results of a 3-day experiment are shown in Fig. 4. On the top, we report the results obtained using the optimal threshold for the detection of both main meals and snacks (black triangle in Fig. 3). On the bottom, we use a higher detection threshold corresponding to the optimal operating point for the detection of main meals only (green triangle in Fig. 3). In the former case, we observe that all nine large meals are successfully detected in the 3 days of simulation. Snacks are all detected as well, except for the night snack on the first day, which is too close to dinner on that day, and so the detector mistakes them as a single meal. In the latter case, using a higher detection threshold $\bar{\Delta}$, we miss the night snack on the second day, but most importantly, the number of false alarms decreases dramatically, which alleviates potential risks of hypoglycemia.

V. RELATED WORK

A number of existing meal detection algorithms are based on examining the rate-of-change (RoC) of the CGM signal; see e.g. [16]–[18]), where detection is performed by comparing the first or second derivatives of the CGM with pre-tuned detection thresholds. Model-based approaches have also been proposed. In [19], meals are detected by computing the difference between the CGM signal and the no-meal prediction of a simple insulin/glucose model, and comparing these differences with a set of post-meal glucose

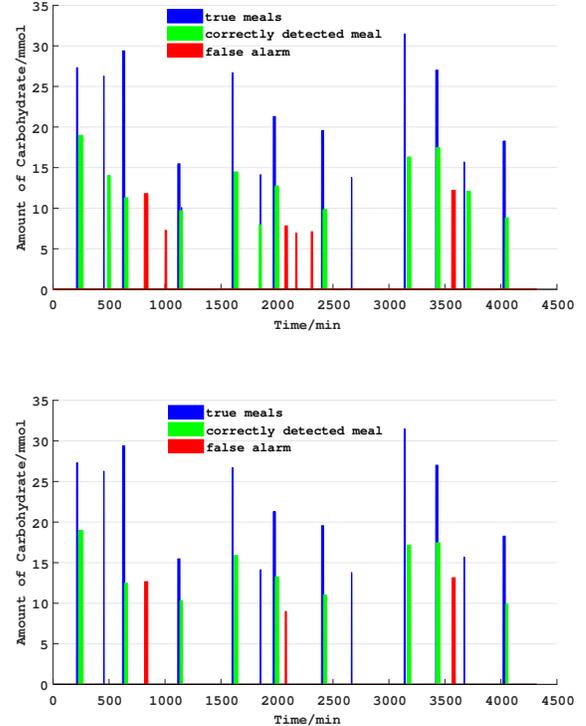


Fig. 4. Meal detection and estimation results of CMHE evaluated on 3-day CGM trajectories for two different detection thresholds $\bar{\Delta}$. Top: results using the optimal operating point for the detection of both main meals and snacks (black triangle in Fig. 3). Bottom: results using a higher detection threshold, corresponding to the optimal operating point for the detection of main meals only (green triangle in Fig. 3).

shapes generated from the model. The work of [20] uses a physiological model to generate test statistics that are invariant to the model parameters and describe the likelihood of meal occurrence in a given time window. All of the above methods, however, only work as meal detectors; i.e., they can detect the occurrence of a meal but they cannot estimate its size (amount of CHO) in a reliable manner.

In [21], time series CGM measurements are converted to a fuzzy qualitative representation whose variable, increase of glucose trend index, is used for meal detection. Meal size is also estimated using a fuzzy system when detection is activated. This method achieves good detection and estimation performance, but it is evaluated only on a fixed meal profile and large meals, unlike our work where we consider random disturbance profiles including both meals and snacks.

In [22], a convolutional neural network is trained to predict the food category and estimate food macronutrients using a nutritional database, food images, and estimated serving size provided by the patients. In contrast to this work, our CMHE approach can provide meal estimates using only CGM data, even though it could be extended to consider additional meal information like food images.

A model-based approach similar to ours is presented in [23], where the authors use an augmented minimal model and an unscented Kalman filter to estimate meal distur-

bances. Meals are detected through three thresholds on, respectively, the difference disturbance signal, the difference CGM signal, and the cross-covariance between the CGM and the difference disturbance signal. However, this work focuses on assessing the detection rate rather than the accuracy of CHO estimates and, unlike our method, does not provide a guaranteed and adjustable detection delay.

VI. CONCLUSION

In this paper, we presented CMHE, a novel model-based online method for meal detection and estimation for type 1 diabetes. CMHE is a new technique for state and disturbance estimation that extends the well-established MHE method. Instead of committing to only one estimate, CMHE aggregates disturbances estimated at different MHE instances in order to achieve a balance between disturbance estimates that are aware of past system measurements and those aware of their effect on future system measurements. Experimental results demonstrate that our method can not only detect multiple types of meals with detection rates close to 100%, but it also estimates the CHO amount of meals with high accuracy for larger meals such as lunch and dinner. As such, our CMHE-based meal detection and estimation method has the potential to replace troublesome manual CHO announcements currently required for insulin therapy, and inform the insulin controller to enable fully closed-loop T1DM therapy. It should be noted, however, that although the in silico experiments show promising performance, the system is still clinically untested.

As future work, we plan to evaluate our method on the UVA/Padova simulator, and eventually on real patient CGM data. To evaluate the performance of our meal detection in a patient cohort, the detection threshold $\bar{\Delta}$ can be optimally personalized to each patient using open-loop glucose-insulin time series collected from the patient.

We also plan to investigate strategies to further improve performance, such as integration of signals from multiple sensors (e.g., microphones and jaw-motion sensors to detect chewing [24]), and mechanisms for the online correction of past CHO estimates when higher-confidence estimates become available.

ACKNOWLEDGMENT

Research supported in part by NSF Grants NSF CPS-1446832, CNS-1445770, IIS-1447549, and CNS-1553273.

REFERENCES

- [1] "Diagnosing Diabetes and Learning About Prediabetes," 2016. [Online]. Available: <http://www.diabetes.org/diabetes-basics/diagnosis/?loc=db-slabnav>
- [2] Centers for Disease Control and Prevention *et al.*, "New CDC report: More than 100 million Americans have diabetes or prediabetes," 2017.
- [3] C. Cobelli, E. Renard, and B. Kovatchev, "Artificial pancreas: past, present, future," *Diabetes*, vol. 60, no. 11, pp. 2672–2682, 2011.
- [4] C. V. Rao, J. B. Rawlings, and D. Q. Mayne, "Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations," *IEEE transactions on automatic control*, vol. 48, no. 2, pp. 246–258, 2003.
- [5] N. Paoletti, K. S. Liu, S. A. Smolka, and S. Lin, "Data-driven robust control for type 1 diabetes under meal and exercise uncertainties," in *International Conference on Computational Methods in Systems Biology*. Springer, 2017, pp. 214–232.
- [6] R. Gondhalekar, E. Dassau, and F. J. Doyle, "Moving-horizon-like state estimation via continuous glucose monitor feedback in MPC of an artificial pancreas for type 1 diabetes," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*. IEEE, 2014, pp. 310–315.
- [7] J. J. Lee, R. Gondhalekar, and F. J. Doyle, "Design of an artificial pancreas using zone model predictive control with a moving horizon state estimator," in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*. IEEE, 2014, pp. 6975–6980.
- [8] E. L. Haseltine and J. B. Rawlings, "Critical evaluation of extended Kalman filtering and moving-horizon estimation," *Industrial & engineering chemistry research*, vol. 44, no. 8, pp. 2451–2460, 2005.
- [9] N. Chen, J. Comden, Z. Liu, A. Gandhi, and A. Wierman, "Using predictions in online optimization: Looking forward with an eye on the past," *ACM SIGMETRICS Performance Evaluation Review*, vol. 44, no. 1, pp. 193–206, 2016.
- [10] R. Hovorka, V. Canonico, L. J. Chassin, U. Haueter, M. Massi-Benedetti, M. O. Federici, T. R. Pieber, H. C. Schaller, L. Schaupp, T. Vering, *et al.*, "Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes," *Physiological measurement*, vol. 25, no. 4, p. 905, 2004.
- [11] S. Chen, J. Weimer, M. R. Rickels, A. Peleckis, and I. Lee, "Towards a model-based meal detector for type I diabetics," *Medical Cyber-physical System Workshop*, 2015.
- [12] R. N. Bergman, Y. Z. Ider, C. R. Bowden, and C. Cobelli, "Quantitative estimation of insulin sensitivity," *American Journal of Physiology-Endocrinology And Metabolism*, vol. 236, no. 6, p. E667, 1979.
- [13] J. B. Rawlings, "Moving horizon estimation," *Encyclopedia of Systems and Control*, pp. 1–7, 2013.
- [14] B. Song, G. Zhang, W. Zhu, and Z. Liang, "ROC operating point selection for classification of imbalanced data with application to computer-aided polyp detection in ct colonography," *International journal of computer assisted radiology and surgery*, vol. 9, no. 1, pp. 79–89, 2014.
- [15] J. Löfberg, "YALMIP : A toolbox for modeling and optimization in MATLAB," in *In Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004.
- [16] E. Dassau, B. W. Bequette, B. A. Buckingham, and F. J. Doyle, "Detection of a meal using continuous glucose monitoring (cgm): implications for an artificial β -cell," *Diabetes care*, 2008.
- [17] 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*, vol. 3, no. 5, pp. 1082–1090, 2009.
- [18] K. Turksoy, S. Samadi, J. Feng, E. Littlejohn, L. Quinn, and A. Cinar, "Meal detection in patients with type 1 diabetes: a new module for the multivariable adaptive artificial pancreas control system," *IEEE journal of biomedical and health informatics*, vol. 20, no. 1, pp. 47–54, 2016.
- [19] F. Cameron, G. Niemyer, and B. A. Buckingham, "Probabilistic evolving meal detection and estimation of meal total glucose appearance," *Journal of Diabetes Science and Technology*, vol. 3, no. 5, pp. 1022–1030, 2009.
- [20] J. Weimer, S. Chen, A. Peleckis, M. R. Rickels, and I. Lee, "Physiology-invariant meal detection for type 1 diabetes," *Diabetes technology & therapeutics*, vol. 18, no. 10, pp. 616–624, 2016.
- [21] S. Samadi, K. Turksoy, I. Hajizadeh, J. Feng, M. Sevil, and A. Cinar, "Meal detection and carbohydrate estimation using continuous glucose sensor data," *IEEE journal of biomedical and health informatics*, vol. 21, no. 3, pp. 619–627, 2017.
- [22] A. Chakrabarty, F. J. Doyle, and E. Dassau, "Deep learning assisted macronutrient estimation for feedforward-feedback control in artificial pancreas systems," in *2018 Annual American Control Conference (ACC)*. IEEE, 2018, pp. 3564–3570.
- [23] C. Ramkissoon, P. Herrero, J. Bondia, and J. Vehi, "Unannounced meals in the artificial pancreas: detection using continuous glucose monitoring," *Sensors*, vol. 18, no. 3, p. 884, 2018.
- [24] E. S. Sazonov and J. M. Fontana, "A sensor system for automatic detection of food intake through non-invasive monitoring of chewing," *IEEE sensors journal*, vol. 12, no. 5, pp. 1340–1348, 2012.