HHS Public Access

Author manuscript

Proc IEEE Conf Decis Control. Author manuscript; available in PMC 2017 May 05.

Published in final edited form as:

Proc IEEE Conf Decis Control. 2014 December ; 2014: 310–315. doi:10.1109/CDC.2014.7039399.

Moving-horizon-like state estimation via continuous glucose monitor feedback in MPC of an artificial pancreas for type 1 diabetes

Ravi Gondhalekar, **Eyal Dassau**, and **Francis J. Doyle III**

Department of Chemical Engineering and Institute for Collaborative Biotechnologies, University of California Santa Barbara (UCSB), USA.

Abstract

An extension of a novel state estimation scheme is presented. The proposed method is developed for model predictive control (MPC) of an artificial pancreas for automatic insulin delivery to people with type 1 diabetes mellitus; specifically, glycemia control based on feedback by a continuous glucose monitor. The state estimation strategy is akin to moving-horizon estimation, but effectively exploits knowledge of sensor recalibrations, ameliorates the effects of delays between measurements and the controller call, and accommodates irregularly sampled output measurements. The method performs a function fit and a sampling action to synthesize a mock output trajectory for constructing the state. In this paper the structure of the fitted function prototype is divorced from the structure of the function that is sampled, facilitating the strategic elimination of prediction artifacts that are not observed in the actual plant. The proposed estimation strategy is demonstrated using clinical data collected by a Dexcom G4 Platinum continuous glucose monitor.

I. Introduction

Model Predictive Control (MPC) finds increased favor due to its flexibility in handling, e.g., constraints, complex dynamics, and sophisticated cost objectives [1,2]. In MPC the control input is characterized by optimizing a future control input trajectory with respect to the predicted evolution of the plant, based on a dynamic model. A crucial element of every MPC implementation is a mechanism to characterize an initial condition from which to perform predictions. This initial condition is typically a function of the model's past outputs and inputs. For regressor models, e.g., an Auto-Regressive model with eXogenous inputs (ARX), the initial condition may be constructed trivially from past measurements and control commands. For state-space models, used more typically in MPC, the initial *state* is provided by a *state estimator*, of which various flavors exist. Simple recursive linear estimators, e.g., the Luenberger observer and Kalman filter, are straightforward to implement and frequently highly effective. Moving-Horizon Estimation (MHE) has been gaining traction as an alternative to recursive strategies and provides, to estimation, the benefits that MPC provides, to control synthesis, over its alternatives [2,3].

For correspondence contact Francis J. Doyle III. doyle@engineering.ucsb.edu. gondhalekar@engineering.ucsb.edu, dassau@engineering.ucsb.edu

In [4] the authors proposed a novel state estimator akin to MHE, but conceptually distinct, that has three advantages over above-mentioned alternatives: **1**) Sensor recalibrations – step-changes in sensor output not related to changes in plant output – are accommodated gracefully. Without special treatment recalibrations cause highly dynamic estimator outputs that can cause undesirable control action. **2**) Measurement times and controller call-times are exploited, facilitating superior state estimates when controller and sensor are not synchronized. **3**) Irregularly timed sensor output can be accommodated; alternative estimators are based on fixed sample-periods, and naively *feeding-in* irregular measurements causes inaccuracies. These advantages are pertinent to the authors' application. The proposed estimator functions by fitting a continuous-time function to recent measurements, and sampling it at the controller call time, and exact integer multiples of the MPC model's sample-period, to synthesize a mock output trajectory for constructing the state estimate. Effects of delays between sensor and controller, and of irregular measurements, are ameliorated by the function fit and sampling. Sensor recalibrations are accommodated by permitting a discrete step in the value, but not the derivatives, of the fitted function following sensor recalibrations.

A contribution of this paper over [4] is that the fitted function is not sampled directly. Instead, a second function is derived from it, and the mock outputs for constructing the state estimate are synthesized by sampling this new function. For example, to obtain a good fit and exploit a long history of measurements, a high order polynomial may serve as the prototype function. However, higher orders may introduce prediction artifacts not observed during plant operation. To yield superior predictive power it may be beneficial to eradicate higher order effects by sampling a straight line. A minor contribution of this paper is that an equality constraint is introduced in the function fit at measurement times following a recalibration. This explicitly forces the fitted function to equal the data obtained at a sensor recalibration. The paper makes two pedagogic contributions. First, the action and benefits of the proposed approach are demonstrated on Continuous Glucose Monitor (CGM) data, not numerical examples, in contrast to [4]. Second, the amelioration of plant-model mismatches is demonstrated. This was partially done in [4] with respect to steady-state errors, but in this paper we demonstrate a crucial advantage with respect to dynamic errors.

The overall goal of this work is an artificial pancreas for the automated delivery of insulin to people with type 1 diabetes mellitus [5,6,7,8]. To achieve this the authors and their collaborators have been developing MPC strategies [9,10,11,12,13]. The feedback signal is an estimate of the blood glucose concentration, provided by a CGM [14]. A CGM's output drifts and thus it must periodically be recalibrated with respect to an accurate reference measurement. The CGMs used by the authors typically update their output every 5 minutes; the sample-period of the MPC model [11]. However, current CGMs are not intended for automatic control purposes, but are designed to produce human readable output, and withhold update information at times of high uncertainly. Furthermore, the CGM sensor communicates wirelessly with the receiver, and communications interference may introduce delays. This leads to irregularly spaced measurements and delays between the receipt of a measurement update and the next controller update, which should occur every 5 minutes as precisely as possible. In the currently employed recursive state estimators these factors

induce output predictions that, sometimes very clearly, are not extensions of the measured output trajectory, do not correspond to the actual output trajectory measured subsequently, and may produce insulin infusion commands that are inappropriate or dangerous. The purpose of the proposed estimator is to facilitate effective and safe control.

II. Background

A. Insulin-glucose model summary

The insulin-glucose model used for performing MPC predictions was described in [11]; it is a discrete-time, third-order ARX model, with $T = 5$ min sample-period, insulin-infusion input, and blood glucose level output. The linearization point is the patient's basal-rate input, and an output concentration of 110 mg/dL. The model can be written as the single-input single-output linear time-invariant (LTI) system

$$x_{i+1}=Ax_i+Bu_i, \quad y_i=Cx_i \quad (1)$$

with state $x \in {}^{n}$, $n = 3$, and (A, C) observable (see [4]).

B. Model predictive control outline

Let denote the set of integers, $\mathbb{Z}_+$ the set of positive integers, and ${}_{a}^{b}$ the set $\{a, \ldots, b\}$ of consecutive integers from a to b. Let $N \in \mathbb{Z}_+$ denote the prediction horizon, and u and x the predicted values of input u and state x. Then, MPC performs closed-loop control by applying, at each step i, the first control input u_0^* of the predicted, optimal control input trajectory $\left\{\mathsf{u}_0^*, \ldots, \mathsf{u}_{N-1}^*\right\}$, characterized by the minimization

$$\left\{\mathsf{u}_0^*, \ldots, \mathsf{u}_{N-1}^*\right\} := \underset{\{\mathsf{u}_0, \ldots, \mathsf{u}_{N-1}\}}{arg\,min} \; J\left(x_i, \{\mathsf{u}_0, \ldots, \mathsf{u}_{N-1}\}\right)$$

of a suitable cost function $J(\cdot)$ (details omitted for brevity), subject to suitable constraints, and furthermore subject to the predictions performed employing model (1):

$$\mathsf{x}_0 := x_i, \quad \mathsf{x}_{k+1} := A\mathsf{x}_k + B\mathsf{u}_k \quad \forall k \in {}_{0}^{N-1}. \quad (2)$$

The predicted state trajectory is initialized in (2) to the estimated model state x_i, the value of which profoundly affects the performance of the resulting MPC control law. No notational distinction between the actual and estimated state is made, because state x of (1) can *only* be estimated.

C. Recursive linear estimator

At each step i let $\tilde{y}_i \in$ denote the most recent sensor measurement. The state estimator employed currently is the recursive linear estimator of (3) with gain L obtained by solving a

Riccati equation (details omitted). Crucially, this estimator is based on model (1) used for MPC predictions.

$$\hat{x}_i = A x_{i-1} + B u_{i-1}, \quad \hat{y}_i = C \hat{x}_i \quad (3a)$$

$$x_i = \hat{x}_i + L\,(\tilde{y}_i - \hat{y}_i) \quad (3b)$$

Model (1) is a state-space realization of an ARX model, as is stated in Section II-A, and the MPC predictions could trivially be initialized using the history of sensor measurements and control commands. However, the recursive estimator of (3) is favored because it provides a convenient framework for tuning the state estimator's noise rejection capabilities.

D. State-reconstruction given exact input-output history

At each step i, given the *exact* sequence $\{y_k\}_{k=i-n+1}^{i}$ of past outputs (and present), and further given the *exact* sequence $\{u_k\}_{k=i-n+1}^{i-1}$ of past control inputs, the current state x_I of model (1) may be reconstructed, given the above-stated observability of the (A, C) pair. For brevity the mechanics of the reconstruction procedure are omitted; the interested reader is referred to Section 2.5 of [4].

E. Comments

During a recalibration the sensor output $\tilde{y}$ undergoes a step change not present in the plant output. After a recalibration all three above-mentioned initialization procedures – the recursive estimator (3), the regressor initialization of model (1) mentioned in Section II-C, and the reconstruction procedure outlined in Section II-D – produce MPC predictions oriented strongly in the direction of the recalibration, not in accordance with the actual CGM output *trend* (see Figure 1).

The initialization schemes above are appropriate when measurements are synchronous with controller call times, and temporally evenly distributed at exactly the sample-period of model (1); $T = 5$ min. In a clinical application these assumptions are not always true, and timing discrepancies may cause badly initialized predictions (see [4]).

III. Estimator design

A. Measurements and recalibrations

Throughout Section III we suppose that the control system is called at time $t = 0$, and thus each output measurement has a negative time-stamp associated with it. Each measurement is defined by a triple $(\tilde{y}_j, t_j, r_j)$ where $j \in {}_1^M$ denotes the measurement index, $M \in \mathbb{Z}_+$ denotes a history horizon that is a design parameter, $\tilde{y}_j$ is the measured output value obtained at time t_j, satisfying $t_{j+1} < t_j < 0 \quad \forall j \in {}_1^{M-1}$, and $r_j \in \{0, 1\}$ denotes a recalibration flag defined such that $r_j := 0$ if no sensor recalibration occurred in the period $[t_j, t_1]$, and $r_j := 1$ if a recalibration was performed within that period. For simplicity it is assumed that exactly one

sensor recalibration has occurred over the history horizon. If no recalibration was performed the method as stated below can be straightforwardly simplified (details omitted). The case of multiple recalibrations during the history horizon is not described for clarity of exposition, but can be accommodated.

B. Estimation procedure

We first select a continuous-time function prototype $\times \Theta \rightarrow$, where Θ is a set of admissible parameters. For example, in this paper we employ the p-order polynomial (4) where p is a design parameter. Functions other than polynomials may be employed or preferable, depending on the application and the properties of the plant's output.

$$\begin{aligned} f(t,\theta) &:= \sum_{k=0}^{p} a_k t^k = \begin{bmatrix} 1 & t & \cdots & t^p \end{bmatrix} \theta \\ a_k &\in \forall k \in {}_0^p \\ \theta &:= \begin{bmatrix} a_0 & \cdots & a_p \end{bmatrix}^{\mathsf{T}} \in {}^{p+1} \end{aligned} \tag{4}$$

Let $\delta \in$ denote the size of the sensor recalibration, that is identified from the measurement data, and further let

$$\bar{\theta} := \begin{bmatrix} \delta & \theta^{\mathsf{T}} \end{bmatrix}^{\mathsf{T}} \in {}^{p+2}.$$

For each $j \in {}_1^M$ define the error, between the discontinuous fitted function $f(\cdot)$ and the data, as follows:

$$e_j := \tilde{y}_j - \begin{bmatrix} r_j & 1 & t_j & \cdots & t_j^p \end{bmatrix} \bar{\theta}\,.$$

The use of r_j facilitates the inclusion or omission of the step δ for uncalibrated or recalibrated measurements, respectively.

Let θ_{i-1}^* denote the optimal parameter as characterized in the previous controller call, and further let $\hat{j} \in {}_1^M$ denote the index of the first measurement since the recalibration, i.e., $\hat{j}$ is the only value of $j \in {}_1^M$ such that $r_{\hat{j}+1} - r_{\hat{j}} = 1$.

The optimal augmented parameter vector $\bar{\theta}^*$ is characterized by solving the quadratic program

$$\bar{\theta}^* := \underset{\bar{\theta} \in \mathbb{R}^{(p+2)}}{arg\,min} (\theta - \theta_{i=1}^*)^{\mathsf{T}} Q (\theta - \theta_{i-1}^*) + \sum_{k=1}^{M} R_k e_k^2$$

subject to the equality constraint

$$\tilde{y}_{\hat{j}} = \begin{bmatrix} r_{\hat{j}} & 1 & t_j & \cdots & t_j^p \end{bmatrix} \bar{\theta}, \quad (5)$$

where $Q \in {}^{(p+1)\times(p+1)}$, $Q \succeq 0$ denotes a cost for penalizing parameter deviations $\theta - \theta_{i-1}$, and $R_k \in \mathbb{R}_{>0}\ \forall k \in {}_1^M$ denote costs to penalize errors e_j, time-dependent with respect to relative time the measurement was taken, but not time-dependent with respect to actual time.

We next derive a further continuous-time function $g: \to$ from $f(t, \theta^*)$. Again, depending on the application and properties of the plant's output trajectory, the reader's preference may vary. In this work we keep it simple and define the affine function

$$\begin{aligned} g(t) &:= f(0, \theta^*) + \dot{f}(0, \theta^*)\, t \\ &= \begin{bmatrix} 1 & t & 0 & \cdots & 0 \end{bmatrix} \theta^*. \end{aligned} \quad (6)$$

We next synthesize a mock output trajectory $\{y_i, y_{i-1}, y_{i-2}, \ldots\}$ by sampling the function $g(t)$ at times $t \in \{0, -T, -2T, \ldots\}$. In combination with the known (assumed) trajectory of past control inputs, this synthesized output trajectory is employed in the procedure outlined in Section II-D to construct the current state x_i.

C. Comments

The function $g(\cdot)$ is an addition to the proposal of [4], and is considered useful because dynamic artifacts can strategically be eradicated from the state estimate by eliminating elements of the structure of $f(\cdot)$. For example, longer history horizons M may facilitate identifying a fit $f(\cdot)$ that better captures the output trend of the measurements. However, long histories require higher order polynomials, including second, and maybe higher, order terms, to capture any curvature in the output trajectory. Performing a straight-line fit would possibly not permit capturing the trend well. However, a state estimate reconstructed from an output trajectory with strong curvature results in predictions that have high curvature (see Section IV-F), but such predictions were not observed in the authors' application. By capturing the trend in $f(\cdot)$, but synthesizing the mock output trajectory via the tangent $g(\cdot)$, all curvature effects of second, and higher, order are expressly rendered non-existent in the state estimate.

The equality constraint (5) is a further addition to the proposal described in [4] and forces the fitted function $f(t, \theta^*)$ to equal the sensor measurement after a recalibration.

IV. Demonstrative example

A. The CGM data

In this section four state estimation schemes are contrasted using CGM data depicted by blue dots in Figure 1; each subplot depicts the same CGM trajectory. The data were obtained from a type 1 diabetic adult wearing the Dexcom G4 Platinum CGM. To facilitate a demonstration of the effects of recalibrations the data were manipulated as follows: The

data-points over the period 18:31–01:06 were raised by 30 mg/dL, increasing the size of a pre-existing recalibration at 18:31, and introducing a recalibration at 01:11 that was not originally present. This is an arbitrary change made for demonstrative purposes for *this paper*, but is not unrealistic.

B. MPC strategy

Depicted by red dots and lines in Figure 1 are the output trajectories predicted by the MPC algorithm. In each subfigure of Figure 1 the MPC algorithm employed was identical and used a prediction horizon $N = 9$ ($NT = 45$ min). Other MPC details are omitted for brevity and are not relevant.

C. Insulin delivery history

The CGM data are fixed; there is no "closed loop" from insulin input to CGM output. Thus, for simplicity and consistency between subfigures, the insulin delivery history employed when characterizing the state estimate is constant and the subject's basal rate, i.e., $u_i = 0$ in terms of LTI model (1). This issue is somewhat revisited later in Section V.

D. High gain recursive linear state estimator

Depicted in Figure 1 (**A**) is the result when using a recursive linear state estimator. The estimator has a high gain, and the predicted output trajectories' starting values are generally in close accordance with the CGM values. The estimator outputs Cx_i are connected by a green solid line that is mostly not visible, because the CGM data are neatly superimposed, as is desired. However, at times of high-speed transients, e.g., after 05:00, the estimator output trails behind the CGM trajectory due to a delay of roughly 4 minutes between the CGM data and subsequent controller call-times.

After sensor recalibrations (18:31, 01:11) the estimator output *value* Cx_i converges quickly to the recalibrated CGM output trajectory. However, the state estimate experiences a highly dynamic response for a prolonged period of time following the recalibration, because knowledge of a recalibration cannot be taken into account. The MPC predictions are obviously initialized at an appropriate *value*, but inappropriate *velocity* and acceleration. Improving the MPC predictions after recalibrations was the primary motivation for developing the proposed state estimation strategy [4].

The reader is cautioned that model (1) characterized in [11] is a so-called *control-relevant model*, designed for use in control synthesis, and suffers large plant-model mismatches that are difficult to eradicate due to the high inter-subject and intra-subject variability of human physiology. The three poles are purposefully sluggish, and the inability to quickly manipulate the rate of change of the predicted outputs, despite a predicted insulin delivery over the predictions, is evidence of this. At times of high acceleration in sensor output the predictions diverge from the output trajectory that is traversed in reality. To a degree such divergence is acceptable, and not correctable by design of the state estimator.

The predictions starting at roughly 19:00, 02:10, and 08:40 are initialized such that they veer away from the CGM trajectory, despite a lack of such trend components in the CGM data

preceding the controller call. This phenomenon is demonstrated and discussed in more detail later in Section V.

E. Proposed estimator: General comments

Figures 1 (**B,C,D**) depict the results obtained when using the proposed state estimator with three different tunings, described in Sections IV-F through IV-H. Although different, the three tunings' results share two important commonalities. First, the value and velocity of the initial condition of the MPC predictions appears to appropriately accommodate the CGM recalibrations. Instead of wild deviations, the predictions capture the trend of the recent CGM trajectory despite the step change in value of the sensor output. The second commonality is that during the rapid transients after 05:00, the estimator output trajectory Cx_i is nearly not discernible, because the CGM trajectory is closely superimposed. This indicates that CGM time-stamps and controller call-times are exploited effectively to reject delays between sensor outputs and controller updates, in contrast to the recursive estimator.

F. Second-order f(·), but without proposed g(·)

Depicted in Figure 1 (**B**) is the result when employing the proposed strategy's fitting function $f(\cdot)$ with the following settings: $p = 2$, i.e., $f(\cdot)$ is a quadratic function; $M = 5$, i.e., 5 past CGM samples are used in the function fit; $\{R_1, \ldots, R_M\} = \{1, 1, 0.8, 0.5, 0.3\}$, i.e., CGM measurements prior to the previous two are de-emphasized in the data fitting; $Q = \mathrm{diag}(0, 1, 0)$, i.e., there is a penalty on the rate of change (a "viscosity") of the optimal parameter vector θ_i^* as i progresses, counteracting high-frequency disturbances. However, we do not define the sampling function $g(\cdot)$ according to (6) as proposed, but sample $f(\cdot)$ in accordance with the original proposal of [4], i.e., $g(t) := f(t, \theta^*)$.

This tuning clearly induces predictions that contain excessive second-order components, and the resulting predictions tend to diverge from the observed CGM trajectory, sometimes strongly, even when CGM fluctuations are very mild. During the oscillations after 05:00 the CGM trajectory's curvature is continued, and exacerbated, in the MPC predictions, leading to significantly amplified maximum and minimum predicted values, compared to any of the three alternatives. Interestingly, similar phenomena occur when employing the standard notion of MHE, but this is not discussed or demonstrated further in this paper.

G. Proposed estimator: Second-order f(·)

Figure 1 (**C**) depicts the result when employing the same settings as above for $f(\cdot)$, but defining the sampling function $g(\cdot)$ as in (6), as proposed in this paper. As $g(\cdot)$ is an affine function, second-order dynamics are explicitly eradicated from the state estimate, resulting in improved predictive power in MPC. The authors consider this state estimator superior to that of Section IV-F. However, except during post-recalibration periods where it is superior, it appears comparable to the recursive estimator of Figure 1 (**A**).

H. Proposed estimator: First-order f(·)

Depicted in Figure 1 (**D**) is the result when employing the proposed strategy with the following settings: $p = 1$, $M = 3$, $\{R_1, \ldots, R_M\} = \{1, 1, 0.01\}$, $Q = \mathrm{diag}(0, 5, 0)$. This choice

of settings implies $f(\cdot)$ is a straight-line fit, using three data points, with very high emphasis on the first two when no recalibration has occurred within the last three steps, and an exact fit to all three previous CGM points when one recalibration was performed. The value of Q implies the viscosity imposed on the velocity term is higher than in the previous two tunings of the proposed estimator. As $f(\cdot)$ is a straight line it holds that $g(t) = f(t, \theta^*)$.

The resulting predictions generally look comparable to those of the recursive estimator of Figure 1 (**A**), except that recalibrations and delays are properly accounted for, as described in Section IV-E. One subtle phenomenon appears to have been eliminated by the proposed estimator. This effect is a result of dynamic plant-model mismatch, and is described and discussed further in Section V. The effect manifests itself in the MPC predictions shortly before 06:00; the predicted peak glucose values are 29 mg/dL lower (253 vs. 282 mg/dL) than with the recursive estimator of Figure 1 (**A**). Furthermore, the predictions at 08:20 are a continuation of the CGM trajectory, as opposed to the recursive estimator where predicted glucose values accelerate upwards.

For the case studied in this paper the shorter history horizon $M = 3$ seemed to outperform the longer one $M = 5$; the authors emphasize that none of the presented estimators can be considered “best in class”. A possible reason is that polynomials constitute bad prototype functions for this data. A better choice of function $f(\cdot)$ may permit exploiting a longer history horizon and glean superior state estimates.

V. Plant-model mismatch

In this section we provide a numerical example to demonstrate how the proposed estimator eliminates some dynamic plant-model mismatches, by basing its estimate on the data alone, *for* a specific model's state, not employing the actual model, or its state, when constructing the state estimate, as the recursive linear estimator does. We consider the CGM trajectory depicted by blue dots in Figure 2: Constant at 70 mg/dL from midnight to 01:00, then rising at 2 mg/dL/min until 03:00, then constant at 300 mg/dL until 06:00. CGM measurements and controller call-times are synchronous.

A. Dynamic rebound

The recursive estimator results in predictions that are initialized with a rate of change exceeding 2 mg/dL/min for a few steps starting at 01:25, and with a strongly negative rate of change after 03:25, despite *no such signal* being present in the CGM measurements. This is because the CGM trajectory does not constitute an output trajectory that is admissible w.r.t. model (1), as is common in practice, and due to the high gain the state estimate x_I of the recursive estimator must adopt values resulting in a close match between CGM and the estimator *output* Cx_I. In contrast, the proposed state estimator does not suffer this phenomenon. In Figure 1 (**A**) the “dynamic rebounding” of the recursive estimator can be seen at roughly 19:00, 02:10, and 08:40. It may also be the reason why the predictions’ peaks, as predicted around 06:00, are significantly higher in Figure 1 (**A**) than in Figure 1 (**D**).

B. State dependence on input history

Plotted in the top subfigure of Figure 3 are the final predictions, starting at 06:00, resulting from the recursive and proposed estimator as plotted in Figure 2. The predicted output trajectories accelerate downwards due to the predicted insulin delivery. However, the proposed estimator results in predictions *starting* with no rate of change, because the CGM data is constant. In contrast, the recursive estimator induces predictions initialized with a negative rate of change, again, despite *no such signal* being present in the CGM measurements. The reason for this is that Figure 2 is based on a constant basal-rate input ($u_i = 0$), as was Figure 1, and according to the model such an input should lower the blood glucose value. However, in practice it frequently does not. Plotted in the bottom subfigure of Figure 3 are the final predictions, starting at 06:00, of the same simulation scenario as before, but with a constant control input of 4× the basal-rate; this choice of input is arbitrary and made only for demonstrative purposes. Higher insulin delivery lowers glucose values, according to the model, and in both cases the predictions achieve a lower glucose value at their final step, as expected, due to the higher insulin delivery history. Crucially, the recursive estimator *initializes* the predictions with a more negative rate of change than in the top subfigure, due to this higher delivery history. In contrast, the proposed estimator correctly initializes the predictions with a rate of change of zero, in accordance with the CGM trend.

VI. Conclusion

A novel state estimation scheme for use in MPC was extended, and its functionality demonstrated in the context of glycemia control with CGM feedback. The proposed strategy is effective because it exploits knowledge of sensor recalibrations, sensor time-stamps and controller time-stamps, and eliminates certain phenomena due to plant-model mismatches. The modification proposed in this paper permits the strategic eradication of prediction artifacts that are not usually observed in the actual data. The proposed method is underpinned by the use of one function prototype that is fitted to the data, and a sampling function that is determined from the fitting function. In future work the most effective structures for these two functions will be investigated.

Acknowledgments

The authors thank Dr. Justin J. Lee (UCSB) for drawing their attention to prediction phenomena induced when employing the standard notion of MHE, and furthermore gratefully acknowledge funding provided by the National Institutes of Health (NIH): DP3DK094331, R01DK085628.

References

1. Maciejowski, JM. Predictive Control with Constraints. Pearson/Prentice Hall; Harlow, England: 2002.
2. Rawlings, JB., Mayne, DQ. Model Predictive Control: Theory and Design. Nob Hill Publishing; Madison, WI, USA: Aug. 2009
3. Goodwin, GC., Seron, MM., De Doná, JA. Communications and Control Engineering. Springer; 2005. Constrained Control and Estimation: An Optimisation Approach..

4. Gondhalekar, R., Dassau, E., Doyle, FJ, III. State Estimation with Sensor Recalibrations and Asynchronous Measurements for MPC of an Artificial Pancreas to Treat T1DM. Proc. 19th IFAC World Congress; Cape Town, South Africa. Aug 2014; p. 224-230.
5. Cobelli C, Dalla Man C, Sparacino G, Magni L, De Nicolao G, Kovatchev BP. Diabetes: Models, Signals and Control. IEEE Rev Biomed Eng. 2009; 2:54–96. [PubMed: 20936056]
6. Harvey RA, Wang Y, Grosman B, Percival MW, Bevier W, Finan DA, Zisser H, Seborg DE, Jovanovič L, Doyle FJ III, Dassau E. Quest for the Artificial Pancreas: Combining Technology with Treatment. IEEE Eng Med Biol Mag. 2010; 29(2):53–62. [PubMed: 20659841]
7. Cobelli C, Renard E, Kovatchev B. Artificial Pancreas: Past, Present, Future. Diabetes. Nov.2011 60:2672–2682. [PubMed: 22025773]
8. Doyle FJ III, Huyett LM, Lee JB, Zisser HC, Dassau E. Closed Loop Artificial Pancreas Systems: Engineering the Algorithms. Diabetes Care. May.2014 37:1191–1197. [PubMed: 24757226]
9. Parker RS, Doyle FJ III, Peppas NA. A Model-Based Algorithm for Blood Glucose Control in Type I Diabetic Patients. IEEE Trans Biomed Eng. Feb.1999 46:148–157. [PubMed: 9932336]
10. Grosman B, Dassau E, Zisser HC, Jovanovič L, Doyle FJ III. Zone Model Predictive Control: A Strategy to Minimize Hyper- and Hypoglycemic Events. J Diabetes Sci Technol. Jul.2010 4:961–975. [PubMed: 20663463]
11. van Heusden K, Dassau E, Zisser HC, Seborg DE, Doyle FJ III. Control-Relevant Models for Glucose Control Using A Priori Patient Characteristics. IEEE Trans Biomed Eng. Jul.2012 59:1839–1849. [PubMed: 22127988]
12. Gondhalekar R, Dassau E, Zisser HC, Doyle FJ III. Periodic-Zone Model Predictive Control for Diurnal Closed-loop Operation of an Artificial Pancreas. J Diabetes Sci Technol. Nov.2013 7:1446–1460. [PubMed: 24351171]
13. Gondhalekar, R., Dassau, E., Doyle, FJ, III. MPC Design for Rapid Pump-Attenuation and Expedited Hyperglycemia Response to Treat T1DM with an Artificial Pancreas. Proc AACC American Control Conf; Portland, OR, USA. Jun 2014; p. 4224-4230.
14. Hovorka R. Continuous glucose monitoring and closed-loop systems. Diabetic Med. Jan.2006 23:1–12.

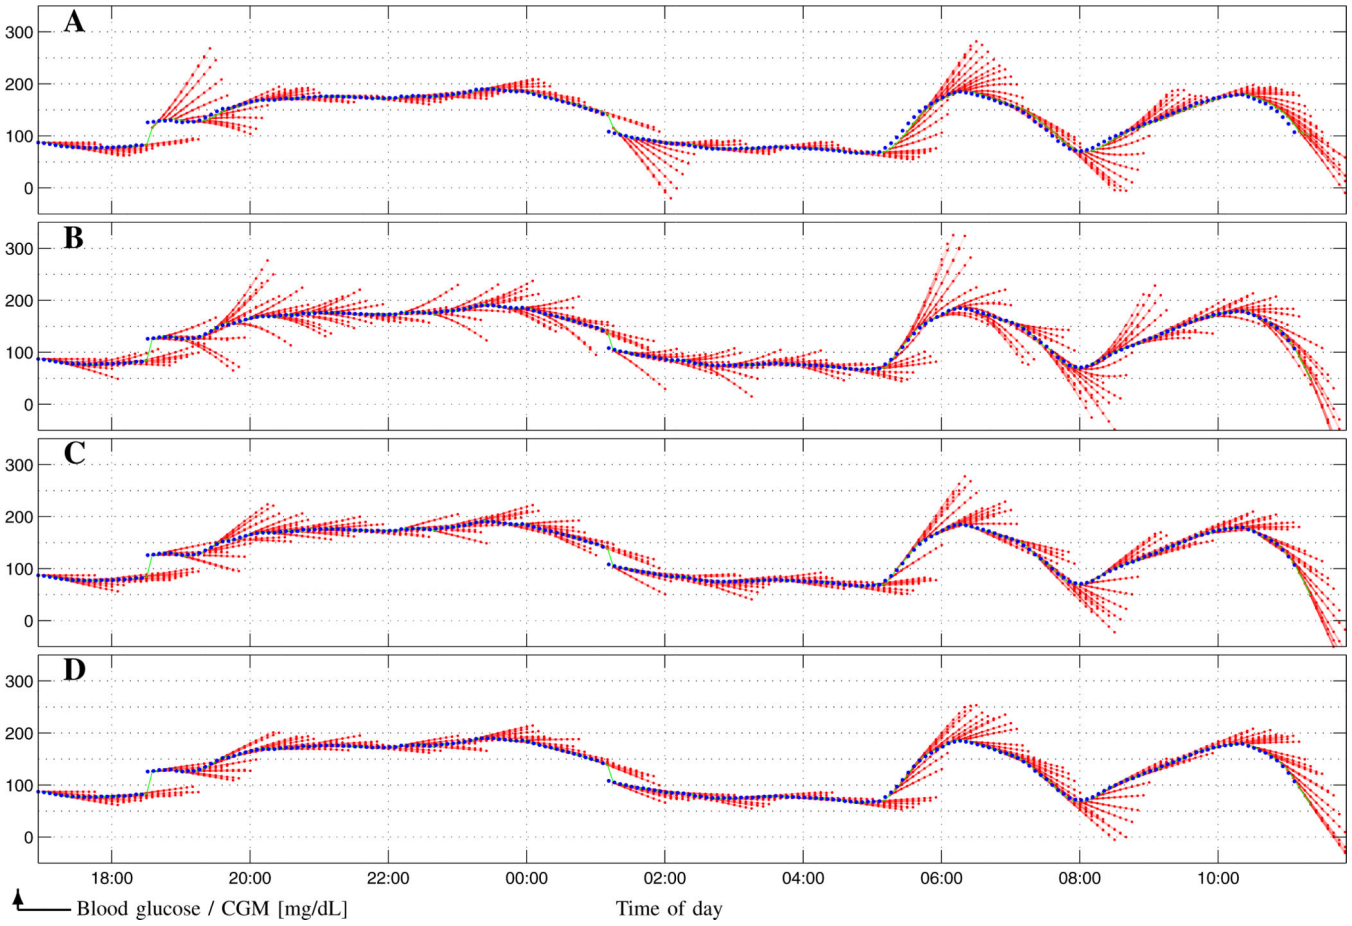

Fig. 1.

Subfigure A: High gain recursive linear state estimator. **Subfigures B, C, D:** Proposed state estimator with three different tunings (see Sections IV-F, IV-G & IV-H). Blue dots: CGM sensor output. Red dots and line: MPC predictions (9 steps, 45 min) of blood-glucose levels based on state estimate. CGM and MPC law identical in all four cases. The green line connects the output estimate Cx_j that is based on the state estimate, and is mostly not visible. It is clearly visible at recalibration points (18:31, 01:11), and for the recursive linear estimator (Subfigure **A**), when the CGM's rate of change is high (ca. 05:00-08:00).

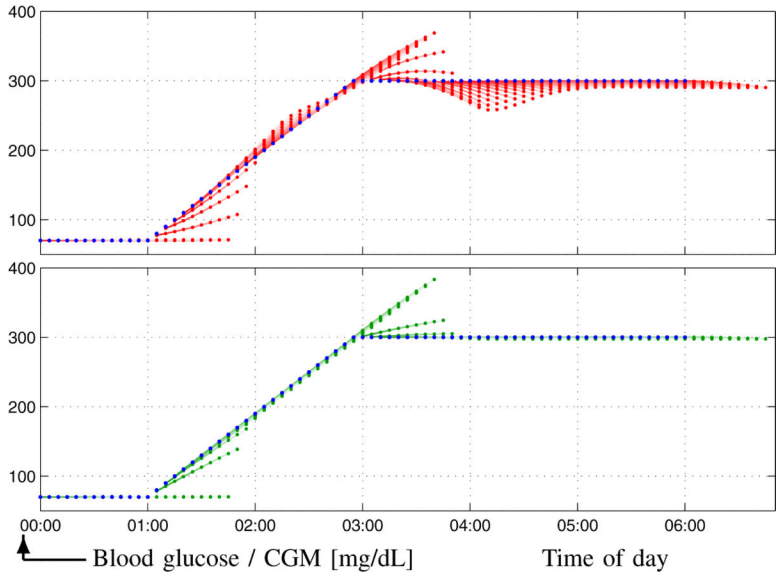

Fig. 2.

Blue dots: CGM. Red/green dots and line: MPC predictions based on state estimate. CGM and MPC law identical in both cases. **Top:** Recursive linear state estimator (Figure 1 (**A**)). **Bottom:** Proposed estimator with parameters of Section IV-H (Figure 1 (**D**)).

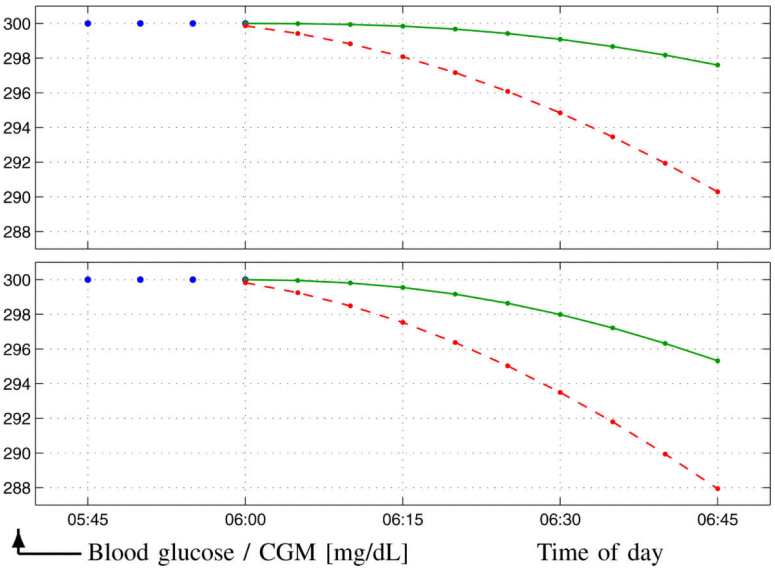

Fig. 3.

Blue dots: CGM. Green dots and solid line: MPC predictions based on proposed state estimator. Red dots and dashed line: MPC predictions based on recursive linear state estimator. CGM and MPC law identical in both cases. **Top:** Basal control input, i.e., $u_i = 0$, corresponding to Figure 2. **Bottom:** Control input elevated at 4× basal-rate, for demonstration.