Automatica 71 (2016) 237–246

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Periodic zone-MPC with asymmetric costs for outpatient-ready safety of an artificial pancreas to treat type 1 diabetes☆

Ravi Gondhalekar, Eyal Dassau, Francis J. Doyle III

Harvard John A. Paulson School of Engineering & Applied Sciences, Harvard University, Cambridge, MA 02138, USA

ARTICLE INFO

Article history:
Received 12 May 2015
Received in revised form
30 November 2015
Accepted 29 March 2016
Available online 1 June 2016

Keywords:
Model predictive control
Periodic control
Safety-critical control
Artificial pancreas
Type 1 diabetes mellitus

ABSTRACT

A novel Model Predictive Control (MPC) law for an Artificial Pancreas (AP) to automatically deliver insulin to people with type 1 diabetes is proposed. The MPC law is an enhancement of the authors' zone-MPC approach that has successfully been trialled in-clinic, and targets the safe outpatient deployment of an AP. The MPC law controls blood-glucose levels to a diurnally time-dependent zone, and enforces diurnal, hard input constraints. The main algorithmic novelty is the use of asymmetric input costs in the MPC problem's objective function. This improves safety by facilitating the independent design of the controller's responses to hyperglycemia and hypoglycemia. The proposed controller performs predictive pump-suspension in the face of impending hypoglycemia, and subsequent predictive pump-resumption, based only on clinical needs and feedback. The proposed MPC strategy's benefits are demonstrated by *in-silico* studies as well as highlights from a US Food and Drug Administration approved clinical trial in which 32 subjects each completed two 25 h closed-loop sessions employing the proposed MPC law.

1. Introduction

Type 1 Diabetes Mellitus (T1DM) is an auto-immune disease that destroys the pancreas' β-cells, rendering people with T1DM incapable of producing insulin, a hormone that facilitates absorption of glucose from the blood-stream into various types of cell, and that plays a crucial role in the endocrine feedback mechanisms that lead to glucose homeostasis in healthy people. People with T1DM tend to suffer chronic hyperglycemia and a lack of glucose homeostasis, causing severe and incurable health problems in later life, e.g., premature cardiovascular diseases, nephropathy, retinopathy, and neuropathy (Centers for Disease Control and Prevention, 2014; The Diabetes Control and Complications Trial Research Group, 1993). The number of people with T1DM in the United States is estimated to be about 1.46 million, nearly 0.5% of the population (Centers for Disease Control and Prevention, 2014). Treating T1DM using an external source of insulin is effective, albeit burdensome, but determining the required dosage is difficult, or impossible, even for experienced and diligent patients. Insulin over-delivery causes hypoglycemia, which may quickly lead to seizures, coma, and death. This work is motivated by the enormous potential for automatic feedback control of insulin delivery to improve the clinical outcomes, and alleviate the burden, resulting from the treatment of T1DM.

Research into a so-called Artificial Pancreas (AP), a device that performs automatic insulin dosing and delivery to people with T1DM, started in the 1970s (Clemens, 1979; Clemens, Chang, & Myers, 1977), but through the development of the CGM (Hovorka, 2006) only became feasible beyond intensive care units much later (Cobelli et al., 2009; Cobelli, Renard, & Kovatchev, 2011; Doyle III, Huyett, Lee, Zisser, & Dassau, 2014; Harvey et al., 2010; Zisser, 2011). AP control laws based on Model Predictive Control (MPC) (Breton et al., 2012; Hovorka et al., 2004; Magni et al., 2009; Parker, Doyle III, & Peppas, 1999; Turksoy, Bayrak, Quinn, Littlejohn, & Cinar, 2013), proportional–integral–derivative control (Marchetti, Barolo, Jovanovič, Zisser, & Seborg, 2008; Steil, Rebrin, Darwin, Hariri, & Saad, 2006), or MD/fuzzy logic (Mauseth et al., 2013; Nimri et al., 2014) have been deployed in human trials. Other control schemes have been proposed and tested *in-silico*, e.g., $\mathcal{H}_\infty$ (Colmegna, Sánchez Peña, Gondhalekar, Dassau, & Doyle III, 2014; Parker, Doyle III, Ward, & Peppas, 2000) and linear parameter-varying (Colmegna, Sánchez-Peña, Gondhalekar, Dassau, & Doyle III, 2015) control. The authors' group has been focusing increasingly on developing zone-MPC strategies (Gondhalekar,

☆ The authors thank the National Institutes of Health (NIH) for funding: DP3DK094331, DP3DK104057, R01DK085628. The material in this paper was partially presented at the 2014 American Control Conference, June 4–6, 2014, Portland, OR, USA. This paper was recommended for publication in revised form by Associate Editor Shun-ichi Azuma under the direction of Editor Toshiharu Sugie.

E-mail addresses: gondhalekar@seas.harvard.edu (R. Gondhalekar), dassau@seas.harvard.edu (E. Dassau), frank_doyle@seas.harvard.edu (F.J. Doyle III).

http://dx.doi.org/10.1016/j.automatica.2016.04.015
0005-1098/

Dassau, & Doyle III, 2014; Gondhalekar, Dassau, Zisser, & Doyle III, 2013; Grosman, Dassau, Zisser, Jovanovič, & Doyle III, 2010; van Heusden, Dassau, Zisser, Seborg, & Doyle III, 2012), whereby the controller reduces insulin delivery from, or supplements insulin in addition to, subjects' basal-insulin only when blood-glucose levels are predicted to make an excursion from a target zone, rather than deviate from a singular setpoint. This was motivated by clinical intuition; there is not one optimal glucose level, instead all glucose levels considered safe form an interval. Furthermore, zone-MPC has proven effective in real-life operation of an AP, yielding control laws that exhibit limited intervention. Use of a zone induces robustness to plant-model mismatch, model bias, and CGM sensor errors; the controller does not respond to small deviations from the setpoint, instead intervenes only when there is a strong indication that intervention is required.

Only recently has an AP been considered feasible in outpatient settings (Hovorka et al., 2014; Kovatchev et al., 2014; Phillip et al., 2013; Russell et al., 2014), facilitated by improvements in CGM accuracy and the availability of consumer-oriented Continuous Subcutaneous Insulin Infusion (CSII) pumps. Safety concerns for outpatient AP deployment are different than for in-clinic use, and it is a contribution of the MPC strategy proposed in this paper to explicitly address these. A primary concern is that while asleep patients cannot monitor themselves or their equipment, and may not respond to alarms. Thus it is the responsibility of the control system to safeguard patients from hypoglycemia, the main immediate risk when treating with insulin, without requiring user-interaction. The proposed strategy improves safety by employing a glucose target zone that is diurnal, i.e., periodic based on the time of day (see Section 2.3). At night, assumed (and enforced in trials) to be the time of sleep, the target zone is raised, encouraging elevated glucose levels and thereby reduced hypoglycemia risk. Additionally, the proposed strategy enforces a diurnal input constraint that limits nighttime insulin infusion to 1.8 times the subjects' basal-rate, limiting the controller's leeway to correct hyperglycemia (see Section 2.4). Diurnal zones and input constraints were first described in Gondhalekar et al. (2013).

Protection from hypoglycemia is more critical at home than in-clinic where, e.g., rescue by intra-venous glucose infusion is feasible. Thus, even when patients are awake and aware of their current state, the control system must prevent hypoglycemia suitably *before* glucose concentrations descend to levels at which patients experience symptoms. What is required are predictive insulin delivery suspensions. The proposed MPC strategy (see Section 2.6) performs appropriate predictive pump-suspensions, and subsequent predictive pump-resumptions, promoted by the use of novel asymmetric input cost functions in the MPC formulation, first described in Gondhalekar et al. (2014) (see Section 4).

Elements of the proposed MPC strategy were proposed previously, but are brought together in this work, and with settings tuned over multiple previous clinical trials. The MPC algorithm proposed in this paper was trialled in the first clinical deployment of the authors' zone-MPC approach in an outpatient setting; for details consult the clinical companion paper (Dassau et al., 2015). It is a contribution of this paper to describe the proposed MPC strategy in reproducible detail, and to demonstrate its efficacy and features using data obtained from real-life testing during US Food and Drug Administration (FDA) approved trials (Dassau et al., 2015). The paper is organized as follows: The feedback MPC strategy is described in Section 2. A feed-forward method to announcing meal intake to the control system is presented in Section 3. The novel asymmetric input cost functions are discussed in Section 4. In Section 5 an outline of the simulation test procedures and results to obtain FDA approval are provided. Highlights of clinical trials using the presented approach are discussed in Section 6.

2. Control law design

2.1. Insulin–glucose dynamics: control-relevant model

Control law design is based on the discrete-time, linear time-invariant (LTI) model of insulin–glucose dynamics proposed in van Heusden et al. (2012), with sample-period $T := 5$ [min]. The time step index is denoted by i. The scalar plant input is the administered insulin bolus $u_{\mathrm{IN},i}$ [U] delivered per sample-period, and the scalar plant output is the subject's blood-glucose value $y_{\mathrm{BG},i}$ [mg/dL]. The model is linearized around the steady-state of the subject-specific, time-dependent basal input rate $u_{\mathrm{BASAL},i}$ [U/h], achieving a blood-glucose output $y_{\mathrm{s}} := 110$ [mg/dL]. The LTI model's input u_i and output y_i are defined as:

$$u_i := u_{\mathrm{IN},i} - \frac{u_{\mathrm{BASAL},i}\, T}{60\ \mathrm{min/h}}, \qquad y_i := y_{\mathrm{BG},i} - y_{\mathrm{s}}.$$

We denote by $\mathcal{Y}(z^{-1})$ and $\mathcal{U}(z^{-1})$ the z-transform of the signals of input u_i and output y_i, respectively. The transfer characteristics from u to y are described by

$$\frac{\mathcal{Y}\left(z^{-1}\right)}{\mathcal{U}\left(z^{-1}\right)} = \frac{1800\, Fc}{u_{\mathrm{TDI}}} \cdot \frac{z^{-3}}{\left(1-p_1 z^{-1}\right)\left(1-p_2 z^{-1}\right)^2} \tag{1}$$

with poles $p_1 := 0.98$, $p_2 := 0.965$, the subject specific *total daily insulin* amount $u_{\mathrm{TDI}} \in \mathbb{R}_{>0}$ [U], and where $c := -60\,(1-p_1)\,(1-p_2)^2$ is used to set the correct gain, and for unit conversion. The so-called *safety factor* F is unitless and provides a mechanism to personalize the model gain to the subject; however, $F := 1.5$ is fixed throughout this paper. The 1800 term stems from the "1800 rule" to estimate blood-glucose decrease with respect to (w.r.t.) the delivery of rapid-acting insulin (Walsh & Roberts, 2006). The state-space realization of (1) used in this work is

$$x_{i+1} = Ax_i + Bu_i, \qquad y_i = Cx_i \tag{2}$$

$$A := \begin{bmatrix} p_1+2p_2 & -2p_1p_2-p_2^2 & p_1p_2^2 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \in \mathbb{R}^{3\times 3}$$

$$B := \frac{1800\, F\, c}{u_{\mathrm{TDI}}}\,[1 \quad 0 \quad 0]^\top \in \mathbb{R}^3, \qquad C := [0 \quad 0 \quad 1] \in \mathbb{R}^{1\times 3}.$$

The structures of A and C indicate that, in the absence of noise, at time-step i the three state elements $x_{[3]}$, $x_{[2]}$, and $x_{[1]}$ correspond to y_i, y_{i+1}, and y_{i+2}, respectively.

2.2. State-estimation

The proposed MPC strategy is designed for use with a CGM that updates its glucose measurement output $\tilde{y}_i$ at the controller's update period $T = 5$ [min]. At each step i let $\tilde{y}_i \in \mathbb{R}$ denote the most recent CGM measurement. An estimate x_i of the state of model (2) is provided at each step i by linear recursive state-estimator (3) (Luenberger observer, see, e.g., Levine, 2011). No notational distinction between the actual and estimated state is made, because state x of (2) can *only* be estimated. Adjusting penalization term R allows tuning the estimator's noise rejection capabilities; the stated value was arrived at by experimentation using the University of Virginia/Padova (UVA/Padova) FDA accepted metabolic simulator (Dalla Man et al., 2014; Kovatchev, Breton, Dalla Man, & Cobelli, 2009).

$$x_i = \hat{x}_i + L\left(\left(\tilde{y}_i - y_{\mathrm{s}}\right) - \hat{y}_i\right) \tag{3a}$$

$$\hat{y}_i = C\hat{x}_i, \qquad \hat{x}_i = Ax_{i-1} + Bu_{i-1} \tag{3b}$$