Proceedings of the 19th World Congress
The International Federation of Automatic Control
Cape Town, South Africa. August 24-29, 2014

Model-Based Control of Type 1 Diabetes in "Risk Space" ⋆

Stephen D. Patek * **Marc Breton** * **Pavel Vereshchetin** *
Boyi Jiang * **Boris P. Kovatchev** *

* *Systems and Information Engineering, University of Virginia; and University of Virginia Center for Diabetes Technology*

Abstract: The clinical significance of glycemic variability in Type 1 Diabetes is asymmetric: a 40 mg/dl deviation below a nominal 110 mg/dl would represent a significant risk of hypoglycemia, while the same deviation above would not cause major concern. The Blood Glucose (BG) risk function of Kovatchev et al. [1997], which is widely used in retrospective analysis of BG data, reflects this asymmetry as a disutility function that is quadratic in the logarithm of BG. Interestingly, the prospective use of the same risk function in model-predictive control can be complicated by the requirement for on-line numerical methods in computing insulin doses that minimize risk over a given prediction horizon. In this work we propose an empirical *linear* model that expresses the dynamic relationship between plasma glucose and remote-compartment insulin in logarithmic coordinates, a model that (i) provides a natural representation of the multiplicative effect of insulin action on glucose clearance and (ii) is such that linear-quadratic methods applied to the model naturally reflect the BG risk function *with closed-form solutions*. We demonstrate the potential of this approach through the design of a Semi-Automated Insulin Advisor that uses continuous glucose monitoring to continuously estimate the patient's metabolic state, informing both episodic correction advice prompted by the patient (for the treatment of hyperglycemia) and automated basal insulin attenuation (for prevention of hypoglycemia). *In silico* pre-clinical trials show favorable performance with respect to idealized "optimal" open-loop treatment, even in scenarios involving miscalibrated carbohydrate ratios and misestimated carbohydrate content in meals.

Keywords: Stochastic Control, Biomedical Systems, Behavior

1. INTRODUCTION

Type 1 Diabetes (T1D) is a lifelong condition characterized by the auto-immune destruction of pancreatic beta cells, destroying the body's ability to produce insulin which is necessary for glucose homeostasis. Insulin replacement therapy is the only proven treatment of T1D, addressing both short- and long-term complications of the disease. Unfortunately, insulin self-treatment represents a significant cognitive burden for the patient, even with the use of an insulin pump. This, along with the opportunity for significantly improved control of Blood Glucose (BG), has given rise to the current wave of interest in Artificial Pancreas (AP) technology. Encouraging results have been reported recently for proportional-integral-derivative control (cf. Weinzimer et al. [2012]), Model Predictive Control (MPC) (cf. Hovorka et al. [2010], Cobelli et al. [2012], Breton et al. [2012], Russell et al. [2012], Dassau et al. [2013]), and fuzzy logic-based strategies (cf. Phillip et al. [2013], Mauseth et al. [2013]).

One of the persistent challenges of designing closed-loop algorithms for the control of T1D is the inherent asymmetry of risk associated with Blood Glucose excursions away from euglycemia. For example, a 40 mg/dl excursion below a euglycemic target of 110 mg/dl presents a significant risk of dangerous hypoglycemia, while a 40 mg/dl excursion above 110 mg/dl lies well within the ADA recommended range of 70-180 mg/dl and is not particularly alarming. Acknowledging this, Parker et al. [2000] and Dua et al. [2009] have proposed the use of an objective function for model-predictive control that penalizes BG deviations asymmetrically so as to emphasize the importance of avoiding hypoglycemia, and Hernjak and Doyle III [2005], et al. [2013] have demonstrated the benefits of also including an asymmetric control penalty term. While these methods have proven to be effective in avoiding hypoglycemia in MPC settings, they have the significant drawback of requiring on-line numerical solvers for computing insulin doses at each stage, even when the underlying plant model is linear.

The BG risk function of Kovatchev et al. [1997] reflects the asymmetry of risk by (i) equating the risks of severe hypoglycemia (20 mg/dl) and severe hyperglycemia (600 mg/dl) and (ii) similarly equating the risks associated with endpoints of the clinically recommended [70, 180] mg/dl target range. The BG risk function is central to the "risk space" computational framework for retrospective analysis of BG data (cf. Kovatchev et al. [2001]), encompassing the

⋆ This work was sponsored in part by the National Science Foundation (NSF/CNS 0931633), the National Institutes of Health (NIH/NIDDK, RO1 DK 08562). This content is solely the responsibility of the authors and does not necessarily represent the official views of the NSF or the NIH.

Low and High Blood Glucose Indices (LBGI and HBGI) and the Average Daily Risk Range (ADRR), which have proven to be predictive of future significant hypoglycemia, hyperglycemia, and extreme glycemic variability, respectively (see Cobelli et al. [2009] for a review). However, the prospective use of the existing risk symmetrization function as a criterion for model-based control presents a challenge since online numerical methods are generally required to compute optimal actions, cf. Magni et al. [2009].

In this paper we propose an alternative "risk space" approach to control that starts with the adoption of an empirical model, which we refer to as a "risk space control model," that describes the relationship between the logarithm of plasma glucose and the logarithm of remote-compartment insulin. This representation of the model has two major benefits: (i) it expresses the multiplicative dependence on remote-compartment insulin in glucose clearance in a linear fashion and (ii) it enables a close approximation of the risk symmetrization as a quadratic function of the state vector in the new coordinate system. Using this framework we have designed a Semi-Automated Insulin Advisor (SAIA) that uses CGM to frequently estimate the patient's metabolic state, informing both episodic correction advice prompted by the patient (for the treatment of hyperglycemia) and automated basal insulin attenuation (for prevention of hypoglycemia). *In silico* preclinical trials show favorable performance with respect to idealized "optimal" open-loop treatment, even in challenging scenarios involving miscalibrated carbohydrate ratios and misestimated carbohydrate content in meals.

2. RISK SPACE CONTROL MODEL

We capture the dynamic interaction of plasma glucose and remote insulin in a logarithmic coordinate system through the following model, whose parameters are fitted (below) from transient responses to glucose challenges.

$$\dot{\lambda}_G(t) = -p_1\lambda_G(t) - p_2\lambda_X(t) + p_3 Q_2(t)/BW \quad (1)$$

$$\dot{\lambda}_X(t) = -p_4\lambda_X(t) + p_4[I_P(t)/(V_I BW) - I_b] \quad (2)$$

where

$$\lambda_G(t) = \ln(G(t)/G_b) \quad \text{and} \quad \lambda_X(t) = \ln(X(t)), \quad (3)$$

with $G(t)$ [mg/dl] representing plasma glucose and $X(t)$ [mU/l] representing insulin acting in the remote compartment. Plasma insulin $I_P(t)$ [mU] is modeled as:

$$\dot{I}_{SC1}(t) = -k_d I_{SC1}(t) + J(t) \quad (4)$$

$$\dot{I}_{SC2}(t) = -k_d I_{SC2}(t) + k_d I_{SC1}(t) \quad (5)$$

$$\dot{I}_P(t) = -k_{cl} I_P(t) + k_d I_{SC2}(t) \quad (6)$$

where $J(t)$ [mU/min] is injected insulin. Gut glucose $Q_2(t)$ [mg] is modeled as follows:

$$\dot{Q}_0(t) = -k_1 (Q_0(t) - m(t)) \quad (7)$$

$$\dot{Q}_1(t) = -k_2 (Q_1(t) - Q_0(t)) \quad (8)$$

$$\dot{Q}_2(t) = -k_3 (Q_2(t) - Q_1(t)) \quad (9)$$

where $m(t)$ [mg/min] is ingested carbohydrates. The parameters p_1, p_2, p_3, p_4, V_I, and BW, some of which are patient-specific, have interpretations similar to those in the standard minimal model of glucose kinetics, cf. Bergman et al. [1979]. The gut and insulin transport parameters $k_1, k_2, k_3, k_d, k_{cl}$ are also patient-specific. Basal glucose concentration G_b is set to 112.5 [mg/dl] as a fixed reference. I_b is calculated from the steady state value of $I_P(t)/(V_I BW)$ with $J(t)$ fixed at the patient's average basal rate.

The method for estimating the parameters of the risk space control model is described in Jiang et al. [2013] (paper forthcoming). A subset of the parameters of the model (p_2, p_3, k_d, and the gut transport model parameters) are adjusted to represent the specific characteristics of an individual patient, and the rest are held fixed as "population" values. Generally, the parameters of the model are chosen to maximize four-hour prediction accuracy. As a first step, using the Virginia/Padova Type 1 Simulator as a reference, we have tuned all of the parameters of a "population average" model designed to maximize average prediction accuracy across all of the adult *in silico* subjects. Next, after fixing the population-average parameters, we computed optimal multiplier values for the individualized parameters. The tuning process uses 2x2 design, with (i) two meal scenarios (first, a meal with a mealtime bolus, and second, the same meal/bolus followed by a correction bolus one hour later) and (ii) two prediction windows (first, between 1 and 5 hours following a meal, and second, between 4 and 8 hours following a meal). The optimization criterion that we used for individualization (i) rewards prediction accuracy within each setting of the 2x2 design but (ii) also heavily penalizes mismatch in accuracy across settings. After computing the optimal multiplier values for each *in silico* subject, we fitted the optimized multipliers to a nonlinear functions of CHO:I and ISF values.

Fig. 1 illustrates that with a cost function that is quadratic in the state vector $(\lambda_G(t), \lambda_X(t))^*$ (specifically, $4405.6 \cdot [\ln(G(t)/G_b)]^2$), cf. red trace, we can closely approximate the BG risk function of Kovatchev et al. [1997], cf. blue trace. (The green trace in the figure illustrates the difficulty of approximating the BG risk function with a quadratic function of $G(t)$, in this case $[G(t)-G_b]]^2$.) Thus, the risk space control model provides a linear-quadratic framework that retains the benefits of the risk space framework, with a computationally tractable model. While we believe that this framework has broad applicability in both advisory and closed-loop algorithms for the treatment of diabetes, we illustrate the use of the framework in the design and *in silico* evaluation of the Semi-Automated Insulin Advisor in which the risk space control model informs model-predictive bolus advice *on demand*.

3. SEMI-AUTOMATED INSULIN ADVISOR

As an illustrative use of the risk space control model, we present a Semi-Automated Insulin Advisor (SAIA), which as shown in Fig. 2 consists in two main modules: an On-Demand Bolus Advisor and Meal-Informed Power Brakes, both of which continuously process insulin history, CGM data, and meal information. The On-Demand Bolus Advisor is invoked episodically by the patient and provides correction bolus advice using a model-predictive approach (using the risk space control model). The Meal-