

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 asym-

metry 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* pre-clinical 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-

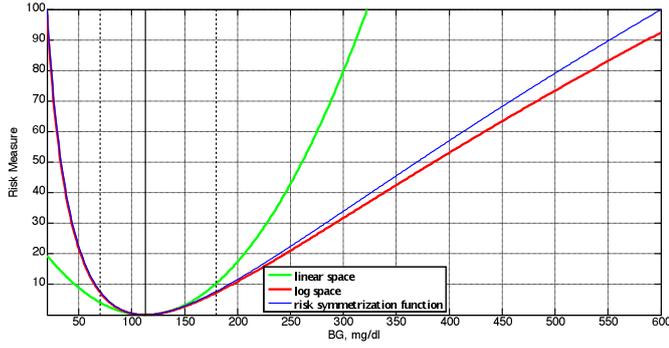


Fig. 1. Comparison of Risk Coordinates

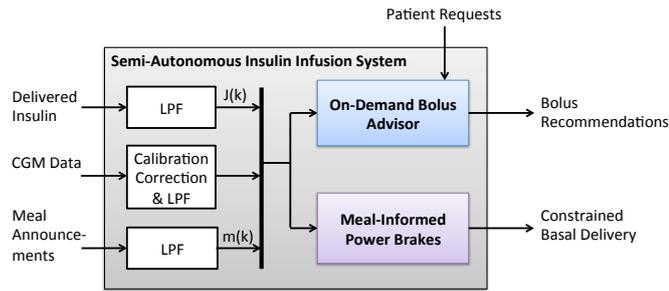


Fig. 2. System Algorithm Schematic Diagram

Informed Power Brakes function by continuously constraining basal insulin delivery based on the predicted risk of hypoglycemia. Since the Power Brakes method has been described elsewhere (cf. Hughes et al. [2010], Patek et al. [2012]), the following sections serve to describe state estimation and model-predictive insulin advice functions of the On-Demand Bolus Advisor.

4. STATE ESTIMATION

State estimation within the On-Demand Bolus Advisor is accomplished through a combination of feedforward estimation and Kalman filtering, as in Grosman et al. [2010]. Here, insulin transport and gut states, expressed by Eqs. (4)-(9), are “estimated” in an open loop fashion, the results of which are fed into a Kalman filter, which uses CGM data to estimate the “core” states of the risk space control model.

To set the stage for this discussion, we define $q(t) = (Q_0(t), Q_1(t), Q_2(t))'$. Then, discretizing Eqs. (7)-(9), we may express the gut model as a discrete-time LTI system:

$$q(k+1) = A_Q q(k) + B_Q m(k), \quad (10)$$

where $q(k)$ refers to the k -th sample of the state vector, $m(k)$ is carbohydrate ingestion held constant for the entire sampling interval, and

$$Q_2(k) = C_{Q2} q(k) + D_{Q2} m(k). \quad (11)$$

Similarly, defining $i(t) = (I_{SC1}(t), I_{SC2}(t), I_P(t))'$ and discretizing Eqs. (4)-(6), we may express the insulin transport model as a discrete-time LTI system:

$$i(k+1) = A_I i(k) + B_I J(k), \quad (12)$$

where $i(k)$ refers to the k -th sample of the state vector, $J(k)$ is insulin delivery held constant for the entire sampling interval, and

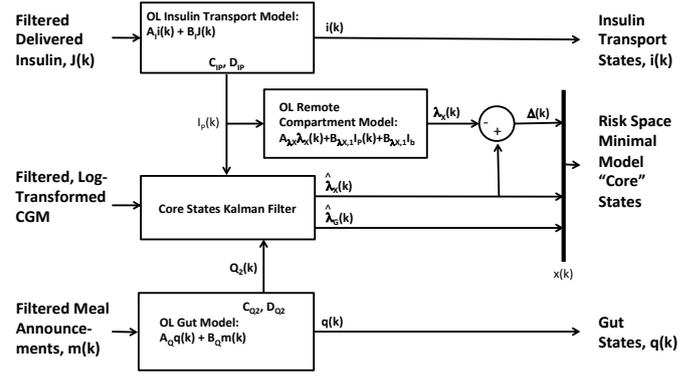


Fig. 3. State Estimation Module

$$I_P(k) = C_{IP} i(k) + D_{IP} J(k). \quad (13)$$

Finally, defining $\eta(t) = (\lambda_G(t), \lambda_X(t))'$ and discretizing Eqs. (1)-(2), we may describe the “core” states of the risk space control model as a discrete time LTI system:

$$\eta(k+1) = A_L \eta(k) + B_{L,1} I_P(k) + B_{L,2} I_b + G_L Q_2(k), \quad (14)$$

where $\eta(k)$ refers to the k -th sample of the core state vector. For convenience later, we introduce a notation for the two rows of the state spaces matrices:

$$A_L = \begin{bmatrix} A_{\lambda G} \\ A_{\lambda X} \end{bmatrix}, \quad B_{L,1} = \begin{bmatrix} B_{\lambda G,1} \\ B_{\lambda X,1} \end{bmatrix}, \quad (15)$$

$$B_{L,2} = \begin{bmatrix} B_{\lambda G,2} \\ B_{\lambda X,2} \end{bmatrix}, \quad G_L = \begin{bmatrix} G_{\lambda G} \\ G_{\lambda X} \end{bmatrix}.$$

4.1 Feedforward “Estimation” of $i(k)$, $q(k)$, and $\lambda_X(k)$

Estimates of $i(k)$ and $q(k)$ are computed in an open loop fashion from knowledge of the inputs: insulin injections $J(k)$ and meals $m(k)$, as depicted in Fig. (3). Also, as a point of reference, we maintain an open loop estimate of remote insulin state using the open loop estimate of $I_P(k)$ and the equation $\lambda_X(k+1) = A_{\lambda X} \lambda_X(k) + B_{\lambda X,1} I_P(k) + B_{\lambda X,2} I_b(k)$.

4.2 Kalman Estimation of $\eta(k)$

CGM-driven estimates ($\hat{\lambda}_G(k), \hat{\lambda}_X(k)$) of the risk model states are computed using a Kalman filter for the dynamic system of Eq. (14) with an additive disturbance process $\omega(k)$:

$$\eta(k+1) = A_L \eta(k) + B_L U_L(k) + \omega(k) \quad (16)$$

where $B_L = [B_{L,1} \ B_{L,2} \ G_L]$ and $U_L = [I_P'(k) \ I_b' \ Q_2'(k)]'$ is treated as a known input (from the open loop estimates above). The innovation process is computed from CGM measurements $y(k) = \ln(\text{CGM}(k)/G_b)$ and is modeled as having additive measurement noise $\nu(k)$:

$$y(k) = C_L \eta(k) + D_L U_L(k) + \nu(k), \quad (17)$$

where C_L and D_L come from the discretization of the risk space control model. For the Kalman filter used in the *in silico* experiments of Section 6, we assume that

$$E(w(k)w(k)') = \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \quad \text{and} \quad E(\nu(k)\nu(k)') = 0.05. \quad (18)$$

(We assume that $\omega(k)$ and $\eta(k)$ are uncorrelated.) In Section 5, we use the difference between the Kalman filter estimate $\hat{\lambda}_X(k)$ and the feedforward estimate $\lambda_X(k)$

$$\Delta(k) = \hat{\lambda}_X(k) - \lambda_X(k). \quad (19)$$

as a representation of model error.

5. CORRECTION BOLUS ADVICE ON DEMAND

Here, we describe how correction bolus advice is computed using the estimated state of the risk space control model. As discussed in Section 3, the system is designed to produce correction bolus advice any time that the patient asks for it. While the effect of *previous* meals is accounted for in the computation (through the estimate of the gut states q), if the bolus advice is being requested *at the time of a meal*, then the content of the meal is ignored in the computation. Meal-related insulin is computed at meal times using the patient's carbohydrate ratio.

To set the stage for this discussion, let us assume that the patient asks for advice at stage k . Section 5.1 below shows how the effect of a proposed bolus is predicted using the individualized risk space control model. Section 5.2 presents the optimization model used to compute the optimal correction.

5.1 Predicting the Effect of a Bolus at Stage k

The model-based advice of this section is informed by the predicted trajectory that results from the correction bolus at stage k . Predictions are initialized by the feedforward and CGM-driven estimates of Section 4. The error variable $\Delta(k) = \hat{\lambda}_X(k) - \lambda_X(k)$ of Eq. (19) is what links the open-loop and feedback elements of the state estimate. In particular, for the purposes of predicting the impact of an insulin bolus u_{bolus} (mU/min) at stage k , we predict future log-glucose states as

$$\begin{aligned} \check{\lambda}_G(k + \tau + 1) = & A_{\lambda G} \begin{pmatrix} \check{\lambda}_G(k + \tau) \\ \check{\lambda}_X(k + \tau) + \check{\Delta}(k + \tau) \end{pmatrix} \\ & + B_{\lambda G,1} \check{I}_P(k + \tau) + B_{\lambda G,2} I_b \\ & + G_{\lambda G} \check{Q}_2(k + \tau) \end{aligned} \quad (20)$$

for $\tau = 0, 1, \dots, N$ (the prediction-horizon of the bolus advisor optimization model), where $\check{\lambda}_G(k) = \hat{\lambda}_G(k)$ (the estimate from the Kalman filter at stage k) and where $\check{I}_P(k) = C_{IP} \check{i}(k) + D_{IP} \check{J}(k)$, and $\check{Q}_2 = C_{Q2} \check{q}(k) + D_{Q2} \check{m}(k)$ are computed from

$$\begin{aligned} \check{\lambda}_X(k + \tau + 1) = & A_{\lambda X} \begin{pmatrix} \check{\lambda}_G(k + \tau) \\ \check{\lambda}_X(k + \tau) \end{pmatrix} \\ & + B_{\lambda X,1} \check{I}_P(k + \tau) + B_{\lambda X,2} I_b \\ & + G_{\lambda X} \check{Q}_2(k + \tau) \end{aligned} \quad (21)$$

$$\check{\Delta}(k + \tau + 1) = \alpha \check{\Delta}(k + \tau) \quad (22)$$

$$\check{i}(k + \tau + 1) = A_I \check{i}(k + \tau) + B_I \check{J}(k + \tau) \quad (23)$$

$$\check{q}(k + \tau + 1) = A_Q \check{q}(k + \tau) + B_Q \check{m}(k + \tau) \quad (24)$$

where $\check{\lambda}_X(k) = \hat{\lambda}_X(k)$, $\check{\Delta}(k) = \Delta(k)$, $\check{i}(k) = i(k)$, and $\check{q}(k) = q(k)$ (all from the state estimation module), $\check{J}(k +$

$\tau)$ reflects the effect of the insulin bolus u_{bolus} at stage k (specifically, $\check{J}(k) = J(k) + u_{bolus}$ and is otherwise just $J(k)$), and $\check{m}(k + \tau) = m(k + \tau)$ reflects the filtered carb input from prior meals (but not including the meal that might be arriving at stage k). Note that per Eq. (24) the carbs that may be ingested at stage k are *ignored* in the prediction, since the goal is to compute an optimal correction bolus. (If there happens to be a meal at stage k , then meal related insulin will be computed by the usual method using the patient's carbohydrate ratio.) Finally, note that the parameter $\alpha = \exp(-5/720)$ causes the initial error $\Delta(k)$ to be "forgotten" over the prediction horizon of optimal bolus computation.

5.2 Optimization Model

The correction bolus advice is computed to minimize the sum of quadratic costs:

$$\min_{u_{bolus} \in \mathfrak{R}} \left[\sum_{\tau=1}^N \omega(k - k_m, \tau) \left(\check{\lambda}_G(k + \tau) \right)^2 \right], \quad (25)$$

where k_m is the stage number of the most recent meal (prior to a meal that might be arriving at stage k), N corresponds to a four-hour prediction horizon ($N = 48$), and $\omega(k - k_m, \tau)$ is a nonnegative weight determined by the elapsed time since the most recent meal. The time-varying weights ω are used to ensure that the optimization model itself is well-adapted to the tactical situation at stage k . In particular, if it has been only a short time since the most recent meal the weights are set to zero for several hours after the meal to prevent the correction advice from targeting the post-prandial excursion. The weights are set to non-zero values sooner if it has already been a long time since the most recent meal. The exact mathematical expression for $\omega(k - k_m, \tau)$ is too complex to relate here, but we provide some details as follows. The weights can be nonzero for τ as small as 18 (corresponding to 90 minutes), and this is in the case where the most recent meal is more than 200 minutes in the past. The weights can remain zero for as long as $\tau = 42$ (corresponding to 210 minutes), and this is in the case where the most recent meal was within 25 minutes of the time of the bolus request. The numerical values of the weights vary between zero and one. Finally, note that weighting function itself is not specific to the characteristics of the patient being treated.

Applying the Lagrange formula to Eqs. (20)-(24), we can compute the predicted trajectory of λ_G as a linear function of the current state estimate and u_{bolus} :

$$\begin{aligned} \tilde{\lambda}_G(k + 1) = & \mathcal{A}_{\lambda G} \hat{\lambda}_G(k) + \mathcal{A}_{\lambda X} \hat{\lambda}_X(k) + \mathcal{A}_{\Delta} \Delta(k) \\ & + \mathcal{A}_I i(k) + \mathcal{A}_{Ib} I_b + \mathcal{A}_Q q(k) + \mathcal{B} u_{bolus} \end{aligned} \quad (26)$$

where $\tilde{\lambda}_G(k + 1) = (\check{\lambda}_G(k + 1), \dots, \check{\lambda}_G(k + N))'$, with the matrices $\mathcal{A}_{\lambda G}$, $\mathcal{A}_{\lambda X}$, \mathcal{A}_{Δ} , \mathcal{A}_I , \mathcal{A}_{Ib} , \mathcal{A}_Q , and \mathcal{B} all derived from the state space equations of Section 5.1. The optimization model of Eq. (25) can be expressed as

$$\min_{u_{bolus} \in \mathfrak{R}} \tilde{\lambda}_G(k + 1)' \Omega(k, k_m) \tilde{\lambda}_G(k + 1), \quad (27)$$

where $\Omega(k, k_m) = \text{Diag}(\omega(k - k_m, 1), \dots, \omega(k - k_m, N))$.

The optimal correction bolus can now be expressed in closed-form as $u_{bolus}^* =$

$$\begin{aligned}
 & -\Phi^{-1}\Theta_{\lambda_G}\hat{\lambda}_G(k) - \Phi^{-1}\Theta_{\lambda_X}\hat{\lambda}_X(k) - \Phi^{-1}\Theta_{\Delta}\Delta(k) \\
 & - \Phi^{-1}\Theta_I i(k) - \Phi^{-1}\Theta_{I_b} I_b \\
 & - \Phi^{-1}\Theta_Q q(k)
 \end{aligned} \tag{28}$$

where $\Phi = \mathcal{B}'\Omega(k, k_m)\mathcal{B}$ and $\Theta_{\lambda_G} = \mathcal{B}'\Omega(k, k_m)\mathcal{A}_{\lambda_G}$, $\Theta_{\lambda_X} = \mathcal{B}'\Omega(k, k_m)\mathcal{A}_{\lambda_X}$, and similar expressions hold for Θ_{Δ} , Θ_I , Θ_{I_b} , and Θ_Q .

6. IN SILICO PRECLINICAL TRIALS

To evaluate the semi-automated insulin advisor developed in Sections 3-5 relative to conventional CSII therapy (mealtime boluses only without low-glucose insulin attenuation) we have conducted *in silico* preclinical trials using the 100 adult subjects that accompany the U. Virginia / U. Padova FDA-accepted Type 1 Simulator. For each experimental setting, we present summary results in terms of (i) percentage time in the range of [70 mg/dl 180 mg/dl] and (ii) percentage time under 70 mg/dl (hypoglycemia).

Miscalibrated Carb Ratio- Three Meals a Day: In this experiment we study the case where the patient's simulated carbohydrate ratio is miscalibrated (within a range of values), so that the meal is underinsulinized to varying degrees, and the On-Demand Bolus Advisor is triggered one hour after the meal to compensate for the inadequate bolus. Each of the *in silico* subjects is challenged with three meals in a 24 hour period, with a breakfast of 0.7 g/kg at 08:00, lunch of 1 g/kg at 13:00, and a dinner of 1 g/kg at 20:00. Mealtime corrections are also computed by the bolus advisor, but this is done without knowledge of the meal amount. The Meal-Informed Power Brakes are enabled for the duration of the experiment. As can be seen in the Fig. 4, the risk-space correction advice serves to reduce the time in the hyperglycemia, especially in the case of heavily underinsulinized meals. The incidence of hypoglycemia in either case is negligible: 0.01 percent time on average below 70 for the entire *in silico* population, with and without the optimal correction.

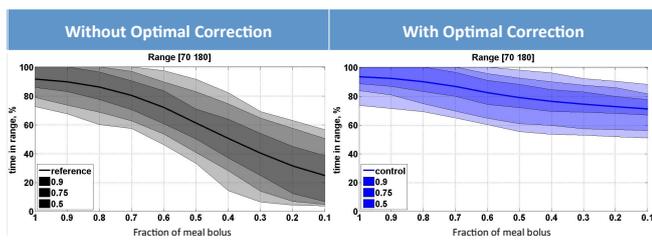


Fig. 4. Graphical summary of the simulation results across the adult subjects ran under a three-meal scenario. The X-axis of each plot shows the fraction of the nominal meal boluses being delivered, ranging from 1 to .1. Average percentage time in range is shown with thick solid lines, accompanied by 50, 75, and 90% envelopes.

Misestimated Carb Content Here, we study the more challenging scenario where the actual carbohydrate content of a meal ranges from -50% to +50% of the true value. (With the incorrect estimate of the size of the meal, the estimate of the patient's state will be thrown off for the

timeframe after the meal.) Each *in silico* subject experiences a 0.8 g/kg meal, and the meal-related insulin dose is computed using the patient's carbohydrate ratio. The On-Demand Bolus Advisor is called one hour after each meal. Mealtime corrections are also computed by the bolus advisor, but this is done without knowledge of the meal amount. The Meal-Informed Power Brakes are enabled for the duration of the experiment. As can be seen in Fig. 5 the bolus advisor manages to improve upon conventional therapy with up to 50% under- and over-estimation of carbohydrates in meals. It is worth noting that the improvement in the case of 50% underestimation is relatively small, due probably to the fact that the Kalman filter has to "catch up" to the truth that a large meal was taken. The improvement in the case of overestimation is due mostly to hypo-mitigating effect of the meal-informed power brakes, which indeed manage to prevent hypoglycemia.

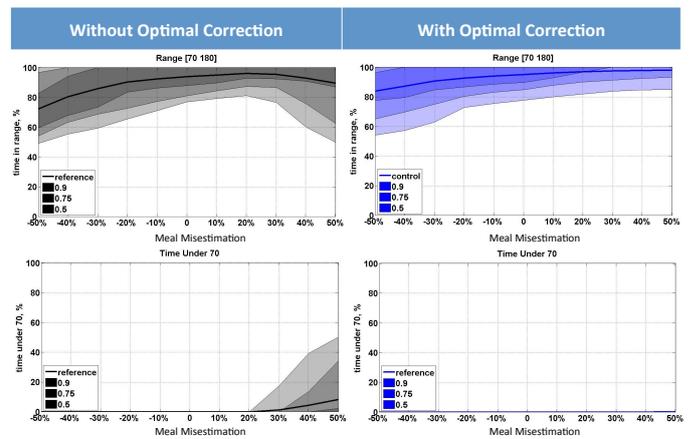


Fig. 5. Graphical summary of the simulation results across the adult subjects ran under a one-meal scenario where the carb content is misestimated from -50% to +50% (X axis).

Advice Timing Relative to Meals Here, we explore the ability of the system to provide correction bolus advice at different times after meals. Returning to a 24-hour simulation scenario, each subject experiences three meals: breakfast of 0.7 g/kg at 8:00, lunch of 1 g/kg at 13:00, and dinner of 1 g/kg at 20:00. In each case the patient receives 50% of his/her meal bolus due to miscalibrated carbohydrate ratio. In separate runs we provide advice at different times after the underbolused meal, ranging from 15 to 240 minutes. Again, the advisor is also invoked at meal times, and the Meal-Informed Power Brakes are continuously enabled. From Fig. 6 we see that the On-Demand Bolus Advisor manages to represent an improvement over the no-advice condition, even when the advisor is invoked 15 minutes after the underbolused meal. Again, the incidence of hypoglycemia in either case is negligible: 0.04 percent time on average below 70 for the entire *in silico* population, with and without the optimal correction.

7. CONCLUSIONS

In this paper we have introduced a framework for *control in risk space*, where, with the use of an empirical model (the risk space control model) expressed in a logarithmic

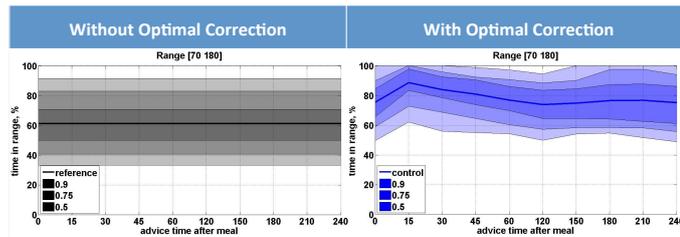


Fig. 6. Graphical summary of the simulation results across the adult subjects ran under one-meal scenarios where the advice is requested at different times relative to the meal.

coordinate system, we obtain the effect of an asymmetric penalty function for model-predictive control but with very simple linear-quadratic computations. Control in risk space provides a natural mechanism for encouraging safe controller action by heavily penalizing BG deviations below 112.5 mg/dl, more so than with comparable deviations above. As demonstrated by our *Semi-Automated Insulin Advisor*, which involves a combination of (i) an On-Demand (risk-based) Bolus Advisor and (ii) Meal-Informed Power Brakes, the risk-space approach tremendously simplifies the design of model-predictive control systems for type 1 diabetes. Specifically, the risk space control model (individualized with the patient's open-loop treatment parameters) facilitated the design of a model-predictive correction bolus advisor that provides safe and effective correction bolus advise for under-bolused meals. Our *in silico* preclinical trial results show that the system is highly robust to misestimated carb content in meals, and this is achieved without requiring patient-specific objective function weights. In addition, we have found that the risk-space control model admits accurate predictions of the effect of post-meal corrections Jiang et al. [2013] (paper forthcoming). In future work we plan to quantify the safety and performance improvements afforded by the risk space approach relative to other methods of representing asymmetric BG risk. In addition, we plan to explore the role and use of the risk space approach in other automated insulin dosing configurations including fully closed-loop artificial pancreas systems.

REFERENCES

- R. N. Bergman, Y. Z. Ider, C.R. Bowden, and C. Cobelli. Quantitative estimation of insulin sensitivity. *Am. Journal of Physiology*, 236:E667–E677, 1979.
- M. Breton, A. Farret, et al. Fully integrated artificial pancreas in type 1 diabetes: modular closed-loop glucose control maintains near normoglycemia. *Diabetes*, 61(9): 2230–7, 2012.
- C. Cobelli, C. Dalla Man, G. Sparacino, L. Magni, G. De Nicolao, and B. P. Kovatchev. Diabetes: Models, signals, and control. *IEEE Reviews in Biomedical Engineering*, 2:54–96, 2009.
- C. Cobelli, E. Renard, et al. Pilot studies of wearable outpatient artificial pancreas in type 1 diabetes. *Diabetes Care*, 35(9):e65–7, 2012.
- E. Dassau, H. Zisser, et al. Clinical evaluation of a personalized artificial pancreas. *Diabetes Care*, 36(4): 801–9, 2013.
- P. Dua, F. J. Doyle III, and E. N. Pistikopoulos. Multi-objective blood glucose control for type 1 diabetes. *Med Biol Eng Comput*, 47:343–352, 2009.
- R. Gondhalekar et al. Asymmetric objective functions for model predictive control of an artificial pancreas. In *Poster and Abstract: Diabetes Technology Meeting*, 2013.
- B. Grosman, E. Dassau, H. C. Zisser, L. Jovanovic, and F. J. Doyle III. Zone model predictive control: a strategy to minimize hyper- and hypoglycemic events. *J Diabetes Sci Technol*, 4(4):961–75, 2010.
- N. Hernjak and F. J. Doyle III. Glucose control design using nonlinearity assessment techniques. *AIChE J.*, 51(2):544–554, 2005.
- R. Hovorka, J. M. Allen, et al. Manual closed-loop insulin delivery in children and adolescents with type 1 diabetes: a phase 2 randomised crossover trial. *Lancet*, 375:743–751, 2010.
- C. S. Hughes, S. D. Patek, M. Breton, and B. P. Kovatchev. Hypoglycemia Prevention via Pump Attenuation and Red-Yellow-Green “Traffic” Lights using CGM and Insulin Pump Data. *J Diab Sci and Tech*, 4:1146–1155, 2010.
- B. Jiang et al. Dynamic model of glycemic risk: application to t1dm bolus advisory system in artificial pancreas platforms. In *Poster and Abstract: Diabetes Technology Meeting*, 2013.
- B. Kovatchev, D. Cox, L. Gonder-Frederick, and W. Clarke. Symmetrization of the blood glucose measurement scale and its applications. *Diabetes Care*, 20: 1655–1658, 1997.
- B. Kovatchev, M. Straume, D.J. Cox, and L.S. Farhy. Risk analysis of blood glucose data: a quantitative approach to optimizing the control of insulin dependent diabetes. *J Theor Med*, 3:1–10, 2001.
- L. Magni, D. M. Raimondo, C. Dalla Man, G. De Nicolao, B. Kovatchev, and C. Cobelli. Model predictive control of glucose concentration in type 1 diabetic patients: An in silico trial. *Biomed. Signal Processing Contr.*, 4: 338–346, 2009.
- R. Mauseth, I. B. Hirsch, et al. Use of a “fuzzy logic” controller in a closed-loop artificial pancreas. *Diabetes Technology and Therapeutics*, 15(8):628–633, 2013.
- R. S. Parker, F. J. Doyle III, J. H. Ward, and N. A. Peppas. Robust h_∞ control in diabetes using a physiological model. *Bioengineering, Food, and Natural Products AIChE Journal*, 46:2537–2549, 2000.
- S. D. Patek, L. Magni, E. Dassau, et al. Modular closed-loop control of diabetes. *IEEE Trans Biomed Eng*, 59(11):2986–2999, 2012.
- M. Phillip, T. Battelino, et al. Nocturnal glucose control with an artificial pancreas at a diabetes camp. *N Engl J Med*, 368(9):824–33, 2013.
- S. J. Russell, F. H. El-Khatib, et al. Blood glucose control in type 1 diabetes with a bihormonal bionic endocrine pancreas. *Diabetes Care*, 35(11):2148–55, 2012.
- S. A. Weinzimer, J. L. Sherr, et al. Effect of pramlintide on prandial glycemic excursions during closed-loop control in adolescents and young adults with type 1 diabetes. *Diabetes Care*, 35(10):1994–9, 2012.