Data-Driven Robust Control for a Closed-Loop Artificial Pancreas

Nicola Paoletti, Kin Sum Liu, Hongkai Chen, Scott A. Smolka, and Shan Lin

Abstract—We present a fully closed-loop design for an artificial pancreas (AP) that regulates the delivery of insulin for the control of Type I diabetes. Our AP controller operates in a fully automated fashion, without requiring any manual interaction with the patient (e.g. in the form of meal announcements). A major obstacle to achieving closed-loop insulin control are the "unknown disturbances" related to various aspects of a patient's daily behavior, especially meals and physical activity. Such disturbances can significantly affect the patient's blood glucose levels. To handle such uncertainties, we present a data-driven, robust, model-predictive control framework in which we capture a wide range of individual meal and exercise patterns using uncertainty sets learned from historical data. These uncertainty sets are then used in the insulin controller to achieve automated, precise, and personalized insulin therapy. We provide an extensive *in silico* evaluation of our robust AP design, demonstrating the potential of the approach. In particular, without the benefit of explicit meal announcements, our approach can regulate glucose levels for large clusters of meal profiles learned from population-wide survey data and cohorts of virtual patients, even in the presence of high carbohydrate disturbances.

1 Introduction

Type 1 diabetes (T1D) is a chronic autoimmune disease in which the pancreas is unable to produce a sufficient amount of insulin to regulate blood glucose (BG) levels. In healthy subjects, pancreatic β cells are responsible for the release of insulin in amounts commensurate with current BG levels. Circulating insulin promotes glucose uptake in muscle and adipose (fatty) tissue. This regulatory process maintains BG within healthy ranges, normally between 70-200 mg/dL. In T1D, T-cell mediated autoimmune destruction of β cells occurs, leading to high BG levels.

In the U.S. alone, more than 29 million people suffer from diabetes, among which approximately 5% have T1D [1]. An estimated 1 million T1D patients worldwide wear an insulin pump to keep their BG levels under control. Insulin pumps are devices for the continuous infusion of insulin, in a manner which provides more accurate therapy and superior BG profiles compared to traditional insulin injections. Insulin pumps deliver *basal insulin*, a low and continuous dose that covers insulin needs outside of meals, and *bolus insulin*, a single high dose for covering meals.

The concept of closed-loop control of insulin, a.k.a. the artificial pancreas (AP), involves a continuous glucose monitor (CGM) that provides glucose measurements (with a typical period of 5 minutes) to a control algorithm running inside the insulin pump or on a peripheral device (e.g. smartphone or tablet) connected to the pump. The controller adjusts the insulin therapy to maintain healthy BG levels, thereby avoiding *hyperglycemia* (BG levels above the healthy range) as well as *hypoglycemia* (BG levels below the healthy range).

While some temporary postprandial (i.e., after-meal) hyperglycemia is admissible, untreated and prolonged hyperglycemia can lead to critical health issues, including cardiovascular disease, kidney damage, and blindness. Hypoglycemia can cause even more serious effects, including coma and death [2].

AP systems have been extensively studied in the last 20 years [3], but only recently cleared for clinical trials [4] and commercialization. The recently FDA-approved MINIMED 670G by Medtronic [5] is the first commercial AP system, and can automatically regulate the basal insulin rate. It is referred to as a "hybrid closed-loop" device as patients need to enter the amount (in grams) of mealtime carbohydrates (CHO) to receive the appropriate bolus insulin dose. This manual procedure is a burden to the patient and inherently dangerous as incorrect information can lead to incorrect insulin dosing and, in turn, harmful BG levels.

Besides meals, another important source of uncertainty in BG control is physical activity (exercise), which accelerates glucose absorption and thus requires a reduced insulin dosage [6]. To build fully automated *closed-loop* AP systems, it is essential to design insulin control algorithms that are *robust* to the patient's eating and exercising behavior.

In this paper, we propose a *data-driven, robust model-predictive control* (robust MPC) framework for the closed-loop control of insulin administration, both basal and bolus, for T1D patients under uncertain meal and exercise events. Such a framework seeks to eliminate the need for meal announcements by the patient, thereby fully automating insulin regulation. We capture the wide range of individual meal and exercise patterns through *uncertainty sets* learned from historical data.

N. Paoletti is with the Department of Computer Science, Royal Holloway, University of London, Egham TW20 0EX, UK. Email: nicola.paoletti@rhul.ac.uk

K. Liu and S. Smolka are with the Department of Computer Science, Stony Brook University, Stony Brook, NY 11794-2424, USA. Email: kiliu@cs.stonybrook.edu *(K. Liu),* sas@cs.stonybrook.edu *(S. Smolka).*

H. Chen and S. Lin are with the Department of Electrical and Computer Engineering, Stony Brook University, Stony Brook, NY 11794-2350, USA. Email: hongkai.chen@stonybrook.edu *(H. Chen),* shan.x.lin@stonybrook.edu *(S. Lin)*

Following [7], we construct uncertainty sets from data so that they cover the underlying (unknown) distribution of the data with *a priori* probabilistic guarantees. Leveraging such information, our robust MPC system computes the insulin-administration profile that minimizes the worst-case performance with respect to these uncertainty sets, thereby providing a principled way to deal with uncertainty.

Besides uncertainty, another challenging aspect of closed-loop control is *state estimation*, which is needed to recover the full physiological state of the patient model (used in MPC) from CGM measurements. Not only are these measurements noisy, but there is also a delay between glucose appearance in the blood and in the sensor (the CGM detects glucose in the interstitial fluid) [8]. Moreover, we need to estimate, along with the state, the current unknown meal and exercise inputs. For this purpose, we designed a moving-horizon state estimator (MHE) [9] that, similar to MPC, exploits a prediction model to find the most likely estimate given the observations. Crucially, MHE enables one to estimate both state and disturbances, unlike more traditional methods based on Kalman filters that only support restricted classes of disturbances [10].

To the best of our knowledge, our robust MPC design for an AP is the first to combine data-driven techniques with a robust insulin controller, and to support both meal and exercise uncertainties. In summary, our main contributions are the following.

- We formulate a closed-loop AP design based on robust MPC to optimize BG levels under meal and exercise uncertainties.
- We apply data-driven techniques to construct uncertainty sets that provide probabilistic guarantees on the robust MPC solution.
- We design an MHE that is able to estimate the patient state as well as meal and exercises disturbances.
- We provide an extensive *in-silico* evaluation of our design, including one-meal simulations, two-day high carbohydrate intake scenarios, and one-day simulations of meal disturbance models learned from population-wide survey data sets (CDC NHANES). We further evaluate our design on a cohort of virtual patients obtained from a clinically validated T1D simulator [11]. Overall, our robust closed-loop AP is able to keep BG within safe levels between 77% and 97% of the time.

This paper extends our previous conference paper [12] by including the following.

- A more realistic model of the CGM sensor (Eqs. 2–5).
- Several new experiments (see Sections 5.2, 5.3, 5.5, and 5.7).
- A detailed description of the differential-equation model of glucose regulation (Section 3).
- A self-contained explanation of the method for constructing uncertainty sets from both patient data and arbitrary distributions (Sections 4.2 and 4.3).
- The procedure used to process meal information from the CDC NHANES database (Section 5.3).

The paper is structured as follows. Section 2 introduces our robust AP design. Section 3 illustrates the gluco-regulatory model and how this is affected by meal and exercise disturbances. Section 4 formalizes the robust MPC and MHE problems, including the construction of uncertainty sets. Section 5 presents the results of our experimental evaluation. Section 6 considers related work. Section 7 offers our concluding remarks.

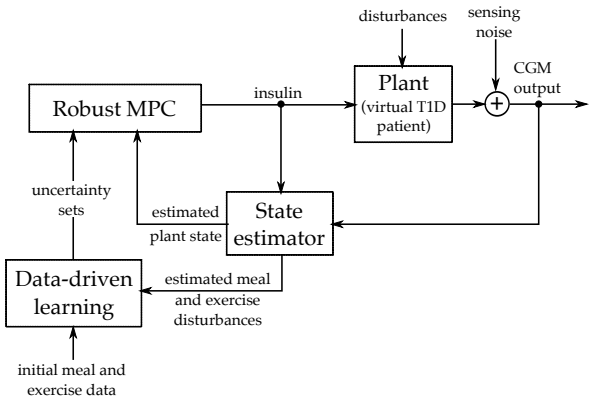

Fig. 1. Robust artificial pancreas design.

2 System Overview

The design of our proposed data-driven robust artificial pancreas is illustrated in Figure 1. The *robust MPC* component (described in Section 4) is responsible for computing the insulin administration strategy (both basal and bolus) that optimizes, over a finite time horizon, the predicted BG profile against worst-case realizations of the meal and exercise disturbances.

Uncertainty sets describe the domains of the disturbances and are derived by the *data-driven learning* component (see Section 4.2), starting from a dataset about the patient's meal and exercise schedules. Uncertainty sets can be also updated online as new data (estimated or announced) come along, in this way enabling the continuous learning of the patient's behavior.

In this work, we analyze our robust artificial pancreas design *in silico*. Thus, the *plant* is given by a system of differential equations (see Section 3) describing the gluco-regulatory dynamics of a virtual T1D patient, as well as the effects of insulin and random disturbances.

In order to approximate real-life settings, we assume that the state of the plant (BG) cannot be observed by the controller, but that we can only access (noisy) CGM measurements. We designed a *moving-horizon state estimator* (described in Section 4.1) that, based on a bounded history of CGM measurements and estimations, computes the most likely plant state. Importantly, this component also provides estimates for the unknown disturbances, which can be used to update the uncertainty sets, even though in this paper we will evaluate our AP design on fixed uncertainty sets.

3 Plant Model

3.1 Disturbances

To account for uncertainty in meal consumption, we consider the input D_G^t, which describes the *rate of CHO ingestion* at time t. Following [13], physical activity is represented by

inputs MM^t, the *percentage of active muscular mass* at time t, and $O2^t$, the *percentage of maximum oxygen consumption* which can be combined to reproduce arbitrary kinds of physical activity.

MM^t corresponds to the ratio between the active muscular mass and the total muscular mass, with typical values being $MM^t = 0\%$ at rest and $MM^t = 25\%$ for a two-legged exercise. $O2^t$ describes the oxygen consumed relative to the maximum oxygen consumption of the subject, and thus, represents a subject-independent measure of exercise workload. Typical values are 8% at rest, 30% for light activity, 60% for moderate activity, and 90% for intense activity [14]. In our scenario, these meal and exercise inputs are not observed or measured, and are thus represented by a disturbance vector $\mathbf{d}^t = (D_G^t, MM^t, O2^t)$. The effects of these inputs on blood glucose are described in the patient's gluco-regulatory model, presented below.

3.2 Patient Model

We consider the nonlinear ODE gluco-regulatory model of Jacobs et al. [13], which extends Hovorka's well-established model [15] to capture the effect of exercise on BG. The model describes the dynamics of glucose and insulin in the human body, i.e., their absorption, metabolism, excretion and transport between compartments (tissues and organs), and is regarded as one of the most physiologically realistic models. In addition to insulin, Jacobs' model also allows for the automated control of glucagon, i.e. the hormone antagonistic to insulin that protects against hypoglycemia. In our work, however, we leave aside glucagon because the single-hormone insulin therapy is the most common one.

Model parameters (available in [16]) represent the physiological characteristics (e.g. transport or consumption rates) of a single virtual subject, even though we also simulate virtual cohorts of patients as described at the end of this section.

At time t, the inputs to the system are the subcutaneous insulin infusion rate u^t (mU/min), and disturbances $\mathbf{d}^t = (D_G^t, MM^t, O2^t)$. The output corresponds to the CGM measurement. The state-space representation of the system is as follows:

$$\dot{\mathbf{x}}(t) = \mathbf{F}\left(\mathbf{x}(t), u^t, \mathbf{d}^t\right),\ t \in \mathbb{R}^{\geq 0} \quad (1)$$

$$y(k) = h\left(\mathbf{x}(k)\right) + \epsilon^k,\ k \in \mathbb{Z}^{\geq 0} \quad (2)$$

where $\mathbf{x}$ is the 14-dimensional continuous state vector that evolves according to the ODE system $\mathbf{F}$, given below. Eq. 2 describes the CGM measurement y, which is collected at discrete time instants k and derived from $\mathbf{x}$ using the measurement model h. Also, y is subject to an additive measurement noise ε^k, which is non-white and non-Gaussian, but follows a Johnson's S_U-distribution [17]. Specifically, ε^k is generated by an autoregressive moving average (ARMA) process which was estimated in [17] using real patient data collected by the FreeStyle Navigator™ CGM and YSI 2300 Stat Plus™ analyzer. The model of the sensor noise is given by:

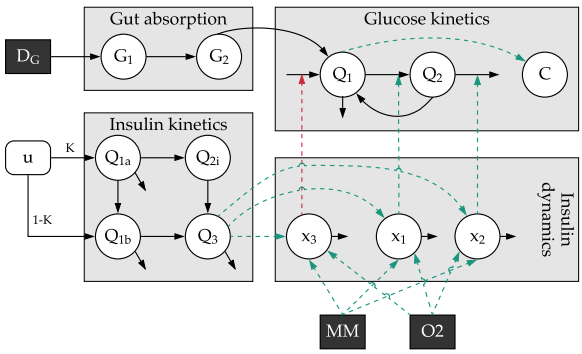

Fig. 2. Schema of the gluco-regulatory ODE system and its four main subsystems. White circles: ODE variables; black boxes: disturbances; white rounded box: insulin input; solid black arrows: flows of glucose or insulin; dashed green/red arrows: positive/negative interactions between variables.

$$e^0 = w^0 \quad (3)$$

$$e^k = 0.7 \cdot (e^{k-1} + w^k) \quad (4)$$

$$\varepsilon^k = \xi + \lambda \sinh \frac{e^k - \gamma}{\delta} \quad (5)$$

where each w^k is sampled i.i.d. from $\mathcal{N}(0,1)$, and $\xi, \lambda, \gamma, \delta$ are the S_U-distribution parameters estimated from data.

Figure 2 illustrates a high-level schema of the ODE system $\mathbf{F}$. The ODE system $\mathbf{F}$, illustrated in Figure 2, consists of the following subsystems:

Gut absorption: the model uses a chain of two compartments, G_1 and G_2 (mmol), to describe digestion of ingested CHO, given by the disturbance D_G^t [11]:

$$\dot{G}_1(t) = \frac{-G_1(t)}{T_{\max}} + A_g \cdot D_G^t, \quad \dot{G}_2(t) = \frac{G_1(t) - G_2(t)}{T_{\max}} \quad (6)$$

where A_g (unitless) is the CHO bio-availability, and $T_{\max}$ (min) is the time of maximum appearance rate of glucose. The gut absorption rate $U_g(t) = G_2(t)/T_{\max}$ characterizes the flow of glucose into the blood compartment Q_1 (given below).

Glucose kinetics: describes the glucose masses in the accessible (where BG measurements are made) and non-accessible compartments, respectively through variables Q_1 and Q_2 (mmol) as follows:

$$\begin{aligned}\dot{Q}_1(t) =& -F_{01c} - x_1 \cdot Q_1(t) + k_{12} \cdot Q_2(t) - F_R \\ &+ U_g(t) + EGP_0 \cdot (1 - x_3(t)) \\ \dot{Q}_2(t) =& x_1(t) \cdot Q_1(t) - k_{12} \cdot Q_2(t) - x_2(t) \cdot Q_2(t)\end{aligned} \quad (7)$$

where F_{01c} and F_R (mmol/min) are parameters for the non-insulin mediated glucose uptake and renal glucose clearance, respectively; variables x_1, x_2, x_3 describe the effect of insulin on glucose (see the *insulin dynamics* subsystem); and parameter EGP_0 (mmol/min) is the glucose production at a theoretical zero-insulin concentration. BG concentration, G (mmol/L), is the main variable we aim to control, and is derived from Q_1 as $G(t) = Q_1(t)/V_G$, where V_G is the glucose distribution volume.

C (mmol/L) corresponds to the glucose detected by the CGM sensor and thus, the measurement function h of Eq. 2

maps the state vector $\mathbf{x}(t)$ to $C(t)$. C has a delayed response w.r.t. the blood concentration G, and is given by:

$$\dot{C}(t) = k_{a\_int} \cdot (G(t) - C(t)) \tag{8}$$

where k_{a_int} is the transport rate parameter.

Insulin kinetics: it models the absorption of the fast-acting insulin u^t, i.e. our control input, and its transport through compartments Q_{1a}, Q_{1b}, Q_{2i} and Q_3 (mU) [18]:

$$\begin{aligned}
\dot{Q}_{1a}(t) =& K \cdot u^t - k_{ia1} \cdot Q_{1a}(t) - \frac{V_{\max,LD} \cdot Q_{1a}(t)}{k_{m,LD} + Q_{1a}(t)} \\
\dot{Q}_{1b}(t) =& (1-K) \cdot u^t - k_{ia2} \cdot Q_{1b}(t) - \frac{V_{\max,LD} \cdot Q_{1b}(t)}{k_{m,LD} + Q_{1b}(t)} \\
\dot{Q}_{2i}(t) =& k_{ia1} \cdot Q_{1a}(t) - k_{ia1} \cdot Q_{2i}(t) \\
\dot{Q}_3(t) =& k_{ia1} \cdot Q_{2i}(t) + k_{ia2} \cdot Q_{1b}(t) - k_e \cdot Q_3(t)
\end{aligned} \tag{9}$$

This model assumes a slow insulin absorption pathway consisting of compartments Q_{1a} (subcutaneous insulin mass) and Q_{2i} (non-accessible insulin), and a fast pathway that includes only Q_{1b} (subcutaneous). Parameter K is the proportion in which the input insulin u^t is distributed into the two pathways. Q_3 is the plasma insulin mass, from which we derive the plasma insulin concentration I (mU/L) as $I(t) = Q_3(t)/V_I$, where V_I is the insulin distribution volume. $V_{\max,LD}$ (mU/min) and $k_{m,LD}$ (mU) are the Michaelis-Menten constants characterizing local insulin degradation.

Insulin dynamics: it defines the effects of plasma insulin levels on blood glucose through variables x_1, x_2, x_3. Variable x_1 (min^{-1}) promotes glucose distribution; x_2 (min^{-1}) promotes glucose disposal; and x_3 (unitless) inhibits endogenous glucose production.

$$\begin{aligned}
\dot{x}_1(t) =& k_{a1} \cdot (-x_1(t) + M_{PGU}(t) \cdot M_{PIU}(t) \cdot S_{IT} \cdot I(t)) \\
\dot{x}_2(t) =& k_{a2} \cdot (-x_2(t) + M_{PGU}(t) \cdot M_{PIU}(t) \cdot S_{ID} \cdot I(t)) \\
\dot{x}_3(t) =& k_{a3} \cdot (-x_3(t) + M_{HGP}(t) \cdot S_{IE} \cdot I(t))
\end{aligned} \tag{10}$$

where M_{PGU}, M_{PIU} and M_{HGP} (unitless) are factors depending on the patient's physical activity (described below); S_{IT}, S_{ID}, and S_{IE} are parameters for the insulin sensitivity on glucose distribution, disposal and production inhibition, respectively. The overall subsystem promotes removal of BG mass from Q_1 and Q_2 and thus, decreases the BG concentration G.

Physical activity: this subsystem consists of two state variables (not shown in Figure 2): the glucose uptake due to active muscular tissue UA (mg/min), and the actual percentage of maximum oxygen consumption $O2_m$ (unitless):

$$\begin{aligned}
\dot{UA}(t) =& k_{UA} \cdot (\overline{UA}(t) - UA(t)) \\
\dot{O2}_m(t) =& k_{O2} \cdot (O2_m(t) - O2^t)
\end{aligned} \tag{11}$$

where $O2^t$ is the disturbance describing the target workload, and $\overline{UA}(t) = f(O2_m(t))$ is the target value of UA, computed as a function of $O2_m$, where f is estimated in [19] using quadratic regression.

The effects of exercise on peripheral glucose uptake (M_{PGU}), on peripheral insulin uptake (M_{PGU}), and on hepatic glucose production (M_{HPG}) depend on UA and $O2_m$ as follows:

$$\begin{aligned}
M_{PGU}(t) =& 1 + k_{PGU} \cdot UA(t) \cdot MM^t \\
M_{PIU}(t) =& 1 + k_{PIU} \cdot MM^t \\
M_{HPG}(t) =& 1 + k_{HPG} \cdot UA(t) \cdot MM^t.
\end{aligned} \tag{12}$$

Initial conditions: The initial state of the system is derived at a steady-state BG level of 7.8 mmol/L [20], assuming no meal and exercise. We use a nonlinear equation solver (MATLAB's `fsolve`) to find $\mathbf{x}(0)$ and basal insulin level $\bar{u}$ such that $\dot{\mathbf{x}}(0) = F(\mathbf{x}(0), \bar{u}, \mathbf{d}^0) = \mathbf{0}$ (see Eq. 1), where the disturbances $\mathbf{d}^0$ are given by $D_G^0 = 0$, $MM^0 = 0$ and $O2^0 = 8$ (oxygen consumption at rest).

Following [13], we further assess the physiologic feasibility of the initial conditions by checking that: 1) in absence of insulin, the steady-state BG is above 300 mg/dL, and 2) delivery of a high insulin dose (15 U/h) results in a steady-state BG below 100 mg/dL.

Virtual patient cohort: In Section 5.5, we validate our robust AP design on a virtual cohort of T1D patients, where each patient has a different parameterization of the ODE model. Specifically, we sample ODE parameters from the distributions of [11, Table 2], which are part of a T1D simulator validated with clinical data and based on Hovorka's model [15].

4 Robust MPC

Since we want to optimize the BG profile against worst-case realizations of the disturbances (constrained by the uncertainty sets), at each time step t, the robust MPC computes the insulin infusion u^t as the solution of the following nonlinear min-max optimization problem:

$$\min_{u^t,\ldots,u^{t+N_c-1}} \max_{\mathbf{d}^t,\ldots,\mathbf{d}^{t+N_p-1}} \sum_{k=1}^{N_p} d(\tilde{\mathbf{x}}(t+k)) + \beta \cdot \sum_{k=0}^{N_c-1} (\Delta u^{t+k})^2 \tag{13}$$

subject to

$$u^{t+k} \in D_u, \quad k = 0, \ldots, N_c - 1 \tag{14}$$

$$u^{t+k} = \bar{u}, \quad k = N_c, \ldots, N_p - 1 \tag{15}$$

$$\mathbf{d}^{t+k} \in \mathcal{U}^{t+k}, \quad k = 0, \ldots, N_p - 1 \tag{16}$$

$$\tilde{\mathbf{x}}(t) = \hat{\mathbf{x}}(t) \tag{17}$$

$$\dot{\tilde{\mathbf{x}}}(t+k) = F(\tilde{\mathbf{x}}(t+k), u^{t+k}, \mathbf{d}^{t+k}),\ k = 0, \ldots, N_p - 1 \tag{18}$$

where N_c and N_p are the control and prediction horizon (in minutes), respectively; (14) states that the control input u must belong to some set D_u of admissible insulin inputs; (15) imposes that u is fixed to the basal insulin rate $\bar{u}$ outside the control horizon; (16) states that, at any time point $t+k$ in the prediction horizon, disturbances $\mathbf{d}^{t+k}$ must belong to the corresponding uncertainty sets $\mathcal{U}^{t+k}$; constraints (17) and (18) restrict how the robust MPC computes the predicted state vector $\tilde{\mathbf{x}}$: for the initial state, it uses the estimated plant state at time t, $\hat{\mathbf{x}}(t)$, while subsequent states are predicted using the plant model (1). We set control and prediction horizons to $N_c = 100$ min and $N_p = 180$ min, respectively.

We design the cost function so as to optimize the following two objectives:

1: Minimize the sum of squared distances between the predicted BG level $\tilde{\mathbf{x}}_G(t+k)$ and a target trajectory $R(t+k)$:

$$d(\tilde{\mathbf{x}}(t+k)) = \gamma(t+k) \cdot (\tilde{\mathbf{x}}_G(t+k) - R(t+k))^2 \quad (19)$$

where $\gamma(t+k) = \gamma$ if $\tilde{\mathbf{x}}_G(t+k) < R(t+k)$ and 1 otherwise. Recall that $\mathbf{x}_G(t) = G(t)$ in the glucose kinetics subsystem. Parameter $\gamma \geq 1$ allows defining asymmetric cost functions where predicted BG values below the target are penalized more than those above the target. Glucose control is naturally asymmetric given that hypoglycemia leads to more severe consequences than (temporary) hyperglycemia, and, as shown in [21], asymmetric costs effectively contribute avoiding hypoglycemia.

2: Minimize step-wise changes in the control input $(\Delta u^{t+k})^2$ to penalize abrupt therapy changes, where $\Delta u^{t+k} = u^{t+k} - u^{t+k-1}$, and u^{t-1} corresponds to the control input in the previous iteration, or to the basal insulin rate $\bar{u}$ if $t = 0$.

In our setup, we fix the target trajectory to $R(t+k) = 7.8$ mmol/L for all time instants and set penalty β to $1/50$. We set the asymmetric cost penalty to $\gamma = 2$, after experimenting with different values (see Section 5.7)

Optimization method

We solve (13) using nonlinear optimization techniques. We reduce the problem into two nested optimization problems, where the objective function value of the outer minimization problem for a control sequence $u^t, \ldots, u^{t+N_c-1}$ is the result of maximizing w.r.t. disturbances $\mathbf{d}^t, \ldots, \mathbf{d}^{t+N_p-1}$ the cost function of (13) with fixed $u^t, \ldots, u^{t+N_c-1}$. To solve both minimization and maximization problems, we employ the interior point method of MATLAB's `fmincon`.

Since we use a nonlinear predictive model and nonlinear solvers, our method is not guaranteed to find a solution (saddle point) of the min-max problem. Nonlinear min-max MPC problems are in general intractable and thus require one to weaken the problem statement (e.g., by approximating the set of model states reachable under all possible disturbances) or the model dynamics (e.g., by linearizing the dynamics) [22]. Our approach addresses a weakened version of the problem in that it finds sub-optimal solutions to both the inner maximization and outer minimization problems, but does not approximate the nonlinear predictive model. Importantly, using nonlinear solvers, we can bound the maximum number of candidate solutions to be evaluated, and thus we can arbitrarily reduce the runtime of our method (sacrificing the quality of the final solution).

To further improve efficiency, we reduce the number of decision variables by assuming that, in the prediction model, control inputs change with period 10 min, and disturbances with period 30 min. We employ a heuristic to select the initial point x_0 for the search of the sub-optimal insulin therapy, which considerably affects the performance of the nonlinear solver. Specifically, we choose x_0 as the control strategy with best performance among 1) the strategy constantly equals to the basal rate; 2) the strategy that linearly transitions (for the control horizon duration) from the previous control input to the basal rate; 3) strategies that linearly transition to the basal rate starting from an insulin rate such that the total insulin dosage (for the prediction horizon duration) is 1 U, 2 U, 3 U, and 4 U, respectively.

Non-robust variant

We introduce a non-robust variant of the controller that will serve as the baseline controller in our experimental evaluation of Section 5. This controller has no knowledge of meals and exercise, and thus is equivalent to fixing the disturbances to their default values at rest in the prediction model. Such a controller mimics the behavior of hybrid closed-loop artificial pancreas systems where only basal insulin is automatically regulated and the patient is responsible for bolus insulin. The optimization problem of the non-robust controller reduces to:

$$\min_{u^t,\ldots,u^{t+N_c-1}} \sum_{k=1}^{N_p} d(\tilde{\mathbf{x}}(t+k)) + \beta \cdot \sum_{k=0}^{N_c-1} (\Delta u^{t+k})^2 \quad (20)$$

s.t. $(14, 15, 17, 18)$ and $\mathbf{d}^{t+k} = (0, 0, 8), k = 0, \ldots, N_p - 1$.

Note that the constraints on the insulin therapy are the same of the robust controller (14-15) meaning that the non-robust controller is free to synthesize bolus-like therapy profiles.

4.1 State Estimation

This component allows to recover an estimate of the current state, which is used in the following iteration by the robust MPC as the initial state for its predictions (see Eq. 17). We designed a moving-horizon state estimator (MHE) [9] that works in a finite-horizon fashion similar to MPC, and allows estimating the current state by reconstructing a state trajectory from previous estimations and a bounded history of observed CGM measurements.

To distinguish from the system variables $\mathbf{x}$ and $\mathbf{d}$, we denote the corresponding estimator variables with $\boldsymbol{\chi}$ and $\boldsymbol{\delta}$, respectively.

For an estimation window of size N, MHE is based on simulating a model of the plant from time $t - N$ to t and aims at finding the model trajectory $\boldsymbol{\chi}(t-N), \ldots \boldsymbol{\chi}(t)$ that minimizes the discrepancies between simulated and estimated states, and between simulated and measured outputs (CGM). Then, the current state estimate $\hat{\mathbf{x}}(t)$ is chosen as the final state of the optimal trajectory, i.e., $\hat{\mathbf{x}}(t) = \boldsymbol{\chi}(t)$.

Crucially, our estimator also works as a meal and physical activity detector: in addition to the plant state, we compute the most likely sequence of disturbances $\boldsymbol{\delta}^{t-N}, \ldots, \boldsymbol{\delta}^t$, corresponding to decision variables in our optimization problem as they are unknown inputs of the model. The MHE problem boils down to the following non-linear optimization problem:

$$\min_{\substack{\boldsymbol{\chi}(t-N),\ldots\boldsymbol{\chi}(t),\\ \boldsymbol{\delta}^{t-N},\ldots,\boldsymbol{\delta}^{t-1}}} \mu \cdot \|\boldsymbol{\chi}(t-N) - \hat{\mathbf{x}}(t-N)\|^2 + \rho \cdot \sum_{k=0}^{N-1} \|v^{t-k}\|^2 \quad (21)$$

subject to

$$\dot{\boldsymbol{\chi}}(t-k) = F(\boldsymbol{\chi}(t-k), u^{t-k}, \boldsymbol{\delta}^{t-k}), \quad k = N, \ldots, 1 \quad (22)$$

$$v^{t-k} = \Big(y(t-k) - \mathrm{E}[\varepsilon^{t-k}]\Big) - h(\boldsymbol{\chi}(t-k)), k = N-1, \ldots, 0 \quad (23)$$

where (22) states that $\boldsymbol{\chi}$ evolves according to the same ODE model of the plant, with u^{t-k} being the insulin input previously computed by the robust MPC; and (23) defines the measurement discrepancy at time $t-k$, v^{t-k}, as the difference between the measured output, $y(t-k)$, and the simulated output, $h(\boldsymbol{\chi}(t-k))$. We take into account the fact that our sensor model has a non-zero-mean noise (see Eq. (2)) by subtracting from $y(t-k)$ the expected value of the sensor noise $\mathrm{E}[\varepsilon^{t-k}]$, which we obtained as the sample mean of multiple realizations of the stochastic noise model of Eq. (3–5). The problem is solved using the interior point method of MATLAB's `fmincon` non-linear solver.

The first addend of the cost function penalizes the discrepancy between the initial state of the simulated trajectory and the corresponding state estimation, where $\mu > 0$ is a weighting factor. The second addend penalizes measurement discrepancies, weighted by a factor ρ. In the original formulation of the MHE [9], the cost function includes discrepancies for all the states in the trajectory. Our simplification comes from the fact that we do not consider random noise in the model (but only in the measurements), and thus, the trajectory $\boldsymbol{\chi}(t-N),\ldots,\boldsymbol{\chi}(t)$ is fully determined by the initial state $\boldsymbol{\chi}(t-N)$ and by the estimated disturbances sequence $\boldsymbol{\delta}^{t-N},\ldots,\boldsymbol{\delta}^{t}$. Further, this greatly improves computational efficiency because variables $\boldsymbol{\chi}(t-N+1),\ldots,\boldsymbol{\chi}(t)$ are strictly constrained by the ODE in Eq. (22). In practice, this means that the decision variables reduce to $\boldsymbol{\chi}(t-N),\boldsymbol{\delta}^{t-N},\ldots,\boldsymbol{\delta}^{t}$.

The MHE has an important probabilistic interpretation: when $N = t$ (unbounded horizon), the MHE problem corresponds to maximizing the joint probability for the trajectory of states $\boldsymbol{\chi}(t-N),\ldots,\boldsymbol{\chi}(t)$ given the measurements $y(t-N),\ldots,y(t)$ [9].

4.2 Building Data-Driven Uncertainty Sets

In this section, we describe how to build the uncertainty sets used within the robust MPC and the state estimator to restrict the domain of the admissible meal and exercise inputs. We apply the approach of [7] where the authors present a general schema for designing uncertainty sets from data for robust optimization (of which robust MPC is an instance). The key idea is to define an uncertainty set that captures random realizations of the disturbance with desired coverage probability, and then optimize against worst-case realizations within this set. In this approach, we see a disturbance as a random variable distributed according to $\mathbb{P}^*$, the true disturbance distribution, and we build the corresponding uncertainty set using data, i.e., a random sample $\mathcal{S}$ drawn from $\mathbb{P}^*$. For instance, in our context, elements of $\mathcal{S}$ are the CHO intakes at some time of the day for a sample of the population. $\mathbb{P}^*$ is the true population distribution of CHO intakes at that time. Importantly, this method requires no information about $\mathbb{P}^*$ and provides a probabilistic guarantee (an upper bound) on the likelihood that the true realized cost — i.e., the cost assuming the true distribution $\mathbb{P}^*$ — is higher than the optimal worst-case cost computed by the robust controller — i.e., the worst-case cost w.r.t. the uncertainty set.

Let us characterize an uncertainty set $\mathcal{U}$ by means of a so-called robust constraint $f(\mathbf{d},\mathbf{x}) \leq 0$, where $\mathbf{d}$ is the unknown disturbances and $\mathbf{x}$ is the optimization variable, corresponding in our case to the state vector plus insulin input. The analytical form for the true distribution $\mathbb{P}^*$ of $\mathbf{d}$ is unknown. Given a confidence level $\epsilon > 0$, $\mathcal{U}$ should satisfy two conditions: (1) the robust constraint f is computationally tractable. (2) $\mathcal{U}$ implies a probabilistic guarantee for $\mathbb{P}^*$ at (significance) level ϵ, that is, for any solution $\mathbf{x}^* \in \mathbb{R}^k$ and for any function $f(\mathbf{d},\mathbf{x})$ concave in $\mathbf{d}$ for all $\mathbf{x}$,

$$\text{if } f(\mathbf{d},\mathbf{x}^*) \leq 0\ \forall \mathbf{d} \in \mathcal{U}, \text{then } \mathbb{P}^*(f(\mathbf{d},\mathbf{x}^*) \leq 0) \geq 1-\epsilon.$$

The data-driven schema that we follow uses a finite set of data points $\mathcal{S}$, which are seen as i.i.d. samples of the true distribution $\mathbb{P}^*$, and applies hypothesis testing to construct uncertainty sets with the above guarantees. In particular, for confidence level $\alpha < 1$, the schema employs the corresponding $(1-\alpha)$ confidence region to build $\mathcal{U}$. With the proper construction, the following theorem from [7, Sect. 3.2] holds.

Theorem 1. *With probability at least $1-\alpha$ with respect to the sampling, the resulting set $\mathcal{U}(\mathcal{S},\epsilon,\alpha)$ implies a probabilistic guarantee at least ϵ for $\mathbb{P}^*$.*

In [7, Table 1], the authors show how different uncertainty sets are built depending on the assumptions about $\mathbb{P}^*$, and, in turn, on the suitable statistical test. In this work we consider box sets (i.e. multi-dimensional intervals), whose construction need no assumptions on $\mathbb{P}^*$. In particular, we derive one box-type uncertainty sets for each time step of the simulation (see Eq. (16)), thereby capturing the typical daylong variations in meal and exercise behavior. For the sake of simplicity, below we omit time dependence and work with uncertainty sets for a generic time step.

Recall that $\mathbf{d} \in \mathbb{R}^d$ denotes the random disturbance vector and d_i denotes its components. Let $\mathcal{S} = \{\hat{\mathbf{d}}^1,\ldots,\hat{\mathbf{d}}^S\}$ be the set of samples (data points). The construction of the box-type uncertainty sets is based on the hypothesis test for the $(1-\epsilon/d)$-quantile of distribution $\mathbb{P}^*$ [7]. For $0 < p < 1$, denote with $Q_p^{\mathbb{P}^*}(d_i) = \inf\{v : \mathbb{P}^*(d_i \leq v) \geq 1-p\}$ the $(1-p)$-quantile of $\mathbb{P}^*$ marginalized over d_i. Given $\bar{q}_{i,0}, \underline{q}_{i,0} \in \mathbb{R}, \forall i = 1,\ldots,d$, we consider the multi-variate hypothesis

$$H_0 : Q_{\epsilon/d}^{\mathbb{P}^*}(d_i) \geq \bar{q}_{i,0} \text{ and } Q_{\epsilon/d}^{\mathbb{P}^*}(-d_i) \geq \underline{q}_{i,0},\ \forall i = 1,\ldots,d.$$

In other words, H_0 describes the hypothesis that the last (resp. first) ϵ/d-quantile of component d_i of the disturbance vector is not below $\bar{q}_{i,0}$ (resp. not above $\underline{q}_{i,0}$). We proceed by identifying values $\underline{q}_i$ and $\bar{q}_i$ such that H_0 is rejected at level α (see Theorem 1), i.e., such that $d_i \in [\underline{q}_i, \bar{q}_i]$, for all components $i = 1,\ldots,d$, with probability at least $1-\alpha$.

Assuming that we have S random samples, define the index s as

$$s = \min\left\{k \in \mathbb{N} : \sum_{j=k}^{S} \binom{S}{j} \left(\frac{\epsilon}{d}\right)^{S-j} \left(1-\frac{\epsilon}{d}\right)^{j} \leq \frac{\alpha}{2d}\right\}. \tag{24}$$

For each component d_i of $\mathbf{d}$, denote with $d_i^{(1)}, d_i^{(2)}, \ldots, d_i^{(S)}$ the sequence of d_i components in the sample set arranged in increasing order. We have that H_0 is rejected at level α if

$$\hat{d}_i^{(s)} < \bar{q}_i \text{ or } -\hat{d}_i^{(S-s+1)} < \underline{q}_i,\ \forall i = 1,\ldots,d,$$

which follows from the fact that, for each i, such defined $\bar{q}_i$ and $\underline{q}_i$ respectively reject hypotheses $Q^{\mathbb{P}^*}_{\epsilon/d}(d_i) \geq \bar{q}_{i,0}$ and $Q^{\mathbb{P}^*}_{\epsilon/d}(-d_i) \geq \underline{q}_{i,0}$ at level $\alpha/2d$ [23]. By this construction, Theorem 8 in [7] shows that if s defined by (24) satisfies $S - s + 1 < s$, then, with probability at least $1 - \alpha$ over the sample, the box set

$$\mathcal{U} = \left\{ \mathbf{d} \in \mathbb{R}^d : \hat{d}_i^{(S-s+1)} \leq d_i \leq \hat{d}_i^{(s)},\ \forall i = 1, \ldots, d \right\}$$

implies a probabilistic guarantee for $\mathbb{P}^*$ at level ϵ.

Note that, even though the above sets are constructed component-wise for each $i = 1, \ldots, n$, i.e., by observing samples from the marginal distributions of $\mathbb{P}^*$, we do not need to assume that $\mathbb{P}^*$ has independent marginals.

To further shrink the size of $\mathcal{U}$ and make it less conservative, we employ the following two strategies. First, prior to set construction, we classify the input data and partition it into a number of clusters so as to obtain tighter sets and more customized, patient-specific control strategies. Second, based on Algorithm 1 of [7], we use bootstrapping [24] to approximate the threshold of the test statistics, by estimating the sampling distribution of the statistics through re-sampling with replacement.

Our uncertainty set construction is still valid if we consider a cluster $\mathcal{S}' \subseteq \mathcal{S}$, with probabilistic guarantees w.r.t. the distribution of the sub-population to which $\mathcal{S}'$ belongs. Suppose, for instance, that we are interested in building an uncertainty set for a sub-population with low CHO intake, e.g. $D_G < 5$ mmol/min. Let $\mathcal{S}' = \{\mathbf{d} \in \mathcal{S} \mid D_G < 5\}$ be the corresponding subset of samples. Then, the uncertainty set for $\mathcal{S}'$ has probabilistic guarantees for the conditional distribution $\mathbb{P}^*(\mathbf{d} \mid D_G < 5)$. Further details about the clustering are provided in Section 5.3.

We remark that the construction of uncertainty sets is performed off-line and thus has no computational footprint on the robust controller.

4.3 Building Uncertain Sets from Probabilistic Models

We show how to build uncertainty sets when meal and exercise disturbances follow arbitrary distributions, a scenario that we evaluate in some of our experiments (see Sections 5.1, 5.2, and 5.4).

In this model, we assume that the start time of a meal, t_m, and the total amount of ingested carbohydrates, CHO, are random. Meal duration d_m is fixed, during which carbohydrate ingestion happens at a constant rate. Similarly, each exercise episode has random start time t_e, percentage of muscular mass MM, and percentage of maximum oxygen consumption $O2$. Exercise duration d_e is random too.

For X random variable, the intuition is to derive an interval $[X^\perp, X^\top]$ that covers realizations of X with arbitrarily high probability p. Let F be the corresponding cumulative distribution function of X and $Q(p) = \inf\{x \in \mathbb{R} \mid F(x) \geq p\}$ be its p-quantile. Then, we choose $X^\perp = Q\left((1-p)/2\right)$ and $X^\top = Q\left((p+1)/2\right)$, which leads to the desired coverage probability p. An example in our experiments is when X is normally distributed with mean μ and standard deviation σ, written $X \sim \mathcal{N}(\mu, \sigma)$. In this case we set $X^\perp = \mu - 3\sigma$ and $X^\top = \mu + 3\sigma$, which gives a coverage probability of $\approx 99.74\%$. When X is instead uniformly distributed in $[a, b]$, written $X \sim \text{Unif}(a, b)$, then we choose the interval defined by $X^\perp = a$ and $X^\top = b$, which covers all realizations with probability 100%.

$$\mathcal{U}^t_{|D_G} = [{D_G^t}^\perp, {D_G^t}^\top],\ \begin{array}{l} {D_G^t}^\perp = CHO^\perp/d \text{ if } t_m^\top \leq t \leq t_m^\perp + d_m, 0 \text{ o/w} \\ {D_G^t}^\top = CHO^\top/d \text{ if } t_m^\perp \leq t \leq t_m^\top + d_m, 0 \text{ o/w} \end{array}$$

$$\mathcal{U}^t_{|MM} = [{MM^t}^\perp, {MM^t}^\top],\ \begin{array}{l} {MM^t}^\perp = MM^\perp \text{ if } t_e^\top \leq t \leq t_e^\perp + d_e^\perp, 0 \text{ o/w} \\ {MM^t}^\top = MM^\top \text{ if } t_e^\perp \leq t \leq t_e^\top + d_e^\top, 0 \text{ o/w} \end{array}$$

$$\mathcal{U}^t_{|O2} = [{O2^t}^\perp, {O2^t}^\top],\ \begin{array}{l} {O2^t}^\perp = O2^\perp \text{ if } t_e^\top \leq t \leq t_e^\perp + d_e^\perp, 8 \text{ o/w} \\ {O2^t}^\top = O2^\top \text{ if } t_e^\perp \leq t \leq t_e^\top + d_e^\top, 8 \text{ o/w} \end{array}$$

TABLE 1
Uncertain sets at time t for CHO ingestion rate D_G, active muscular mass MM and oxygen consumption $O2$.

With this method, we can thus define intervals $[t_m^\perp, t_m^\top]$, $[CHO^\perp, CHO^\top]$, $[t_e^\perp, t_e^\top]$, $[d_e^\perp, d_e^\top]$, $[MM^\perp, MM^\top]$, $[O2^\perp, O2^\top]$ bounding the realization the above variables describing start time, duration and intensity of the disturbance. Then, the box-type uncertainty sets at time t relative to inputs D_G (CHO ingestion), MM and $O2$ are defined as per Table 1. In addition, when a meal or exercise episode happens with probability lower than 1 (as in the experiment at Section 5.4), we set the lower bounds of the corresponding uncertainty sets to their values at rest (0 for D_G and MM, 8 for $O2$).

5 Results and Discussion

Hardware and runtime

We ran the experiments on a Windows 8 machine with an Intel Core i7 processor and 32GB of DDR3 memory. We used MATLAB version 2016b. With this configuration, the average time to compute the insulin therapy over all the experiments ranged from 4 to 18 seconds, which is well within the CGM measurement period of 5 minutes. Given the significant performance improvement of modern embedded and mobile devices, we expect our algorithm to perform similarly once deployed on such hardware platforms. In this work we do not focus on optimizing the performance on embedded platforms. For this purpose, methods for learning function approximations of MPC trajectories could be used [25].

Performance indicators

To measure the efficacy of our robust controller design over multiple runs, we consider the following indicators (the means below are intended as averages over multiples runs):

$t_{<3.9}$, $t_{3.9-11.1}$, $t_{>11.1}$*:* mean percentage of time spent in, respectively, hypoglycemia (BG < 3.9 mmol/L), normal ranges (BG between 3.9 and 11.1), and hyperglycemia (BG > 11.1). Clearly, we wish to maximize $t_{3.9-11.1}$ and minimize the other two indicators, keeping in mind that we can tolerate some temporary postpandrial hyperglycemia while hypoglycemia should be avoided as much as possible.

$BG_{\min}$, $BG_{\max}$*:* mean low and peak BG level, respectively, in mmol/L. The controller should keep $BG_{\min}$ and $BG_{\max}$ as close as possible to the target BG level.

$\sum u$*:* mean total non-basal insulin (in U). It measures the amount of insulin injected by the controller in order to cover meals, and thus excludes the contribution of basal insulin.

TABLE 2
Performance indicators for the one-meal experiments of Figure 3. The best value of each indicator between the robust and the non-robust controllers is highlighted in bold.

	$t_{<3.9}$	$t_{3.9-11.1}$	$t_{>11.1}$	$BG_{\min}$	$BG_{\max}$	$\sum u$
1), Perfect	0%	99.69%	0.31%	7.15	9.91	4.38
1), Non-rob.	**1.72%**	75.47%	22.81%	**5.56**	12.13	8.5
1), Robust	1.99%	**94.91%**	**3.11%**	5.43	**10.07**	6.87
2), Perfect	0%	100%	0%	7.03	8.84	4.67
2), Non-rob.	**0.39%**	69.04%	30.57%	**5.81**	13.58	7.09
2), Robust	0.47%	**77.54%**	**21.99%**	5.11	**12.15**	7.1
3), Perfect	0%	100%	0%	7.22	9.3	5.06
3), Non-rob.	**0%**	76.52%	23.48%	**6.7**	12.48	5.86
3), Robust	1.01%	**97.75%**	**1.24%**	5.11	**9**	6.78

5.1 One-Meal Experiments

We consider 300-minute simulations comprising a single meal, and three different synthetic scenarios (illustrated in Figure 3 (a-c)), i.e. where meal disturbances follow arbitrary distributions (see Section 4.3). For each scenario and controller, we collect results for 50 repetitions.

Scenario 1, meals as expected: we assume a uniformly distributed meal with start time $t_m = \text{unif}(30, 90)$, total amount of CHO (grams) $\mathsf{CHO} = \text{unif}(42, 78)$ and meal duration fixed to 20 minutes, during which CHO ingestion happens at a constant rate. This scenario allows us evaluating the adequacy of the controller when the plant behaves according to a known distribution, in other words, when we have accurate information for building uncertainty sets.

Scenario 2, outliers: in this case, random meals behave as statistical outliers, i.e. they are constantly distant from the expected value of the underlying distribution. To this purpose, we build the uncertainty sets under the assumption that meals are normally distributed with parameters $t_m = \mathcal{N}(60, 15)$ and $\mathsf{CHO} = \mathcal{N}(60, 9)$. The uncertainty sets are built so as to cover all possible realizations with z-score between -3 and 3 (i.e. between -3 S.D. and +3 S.D. around the mean). However, to reproduce outliers, meal disturbances are sampled from the tails of the distributions (z-scores in $[-4, -3]$ and $[3, 4]$).

Scenario 3, late meals: here we consider the same settings as in Scenario 1, but with each random meal delayed of one hour. This models the situation where the controller has wrong information about the meal schedule, since it expects the meal to start, on the average, one hour earlier.

Results in Figure 3 show that our robust controller attains very good performance, closely following the ideal behavior of the perfect controller in the first and third scenarios, where the virtual patient stays in normal ranges for >94% of the time. In the outliers scenario, we register some postprandial hyperglycemia, because this scenario is characterized by frequent high CHO intake. Overall, the robust controller is able to limit the time spent in hypoglycemia below 2% and consistently outperforms the non-robust controller, staying in normal BG ranges for 8.5% to 21.23% more. Performance indicators are given in Table 2.

5.2 Regulation during Exercise

We evaluate the behavior of the robust controller when the virtual patient is involved in physical activity, which, contrarily to meals, contributes to decreasing BG levels. Specifically, we simulate a scenario where a meal is followed by physical activity, a situation that typically leads to hypoglycemia because the BG decrease due to the meal-time insulin dose is exacerbated by exercise. After a meal constructed based on the *meals as expected* scenario, we simulate a two-legged exercise consisting of two phases:

1) **Moderate activity**, with start time $t_e = \text{unif}(40, 80)$, duration $d_e = \text{unif}(24,\ 36)$, active muscular mass $MM = \text{unif}(0.15, 0.35)$, and oxygen consumption $O2 = \text{unif}(45, 75)$; followed by
2) **Light activity**, where parameters stay as in the previous phase except for $O2 = \text{unif}(15, 45)$.

Results, reported in Figure 4, show that the robust controller effectively prevents hypoglycemia ($t_{<3.9} = 0.57\%$), maintaining BG in range for 86% of the time. In contrast, without any knowledge about plausible ranges for future disturbances, the non-robust controller fails to predict the occurrence of exercise after the meal, resulting in insulin overshooting and critically low BG levels ($t_{<3.9} = 11.87\%$).

We also performed a variant of this experiment that does not include the initial meal, obtaining 100% of the time in range for all controllers. Our results outperforms those of Resalat et al. [13], that experimented with a similar scenario for a dual-hormone MPC (300-minute simulation with a 45-minute exercise at fixed $O2 = 60$ and $MM = 0.8$). While we use their same plant model, their MPC design is different in two ways: it can regulate both insulin and glucagon (to prevent hypoglycemia) and is not robust, in that the controller knows the true exercise onset time and intensity. Despite that, however, their evaluation resulted into some episodes of hypoglycemia and hyperglycemia, while our controller is able to keep BG for 100% of the time in healthy ranges without meal announcements.

5.3 One-Day Experiments using NHANES Survey Data

We test our robust controller with meal disturbances extracted from real population data in the CDC's National Health and Nutrition Examination Survey (NHANES) database [26].

We consider the 2013 survey, comprising 8,611 participants, and retrieve meal information from the dietary interview, where each participant reports the timings, types and amounts of each meal during a typical day. Through a moving average filter, we transform the meal events of each participant into a one-day trajectory describing the CHO intake rate, so that it can be mapped into the input disturbance D_G. To avoid building a single uncertainty set built from the whole database, which would result in a unrealistic and overly conservative sets, we classify the trajectories into 10 groups using k-means clustering. Grouping meal data into clusters aligns with the common pre-pump assessment questionnaire conducted by medical personnel.

In this experiment, we focused on and evaluated a cluster consisting of 274 people whose meal patterns are characterized by a CHO-rich breakfast at around 9am, as visible in the uncertainty set of Figure 5 (a). Research [27] has shown that hyperglycemia occurs after breakfast and timing of pre-breakfast insulin injection is important to

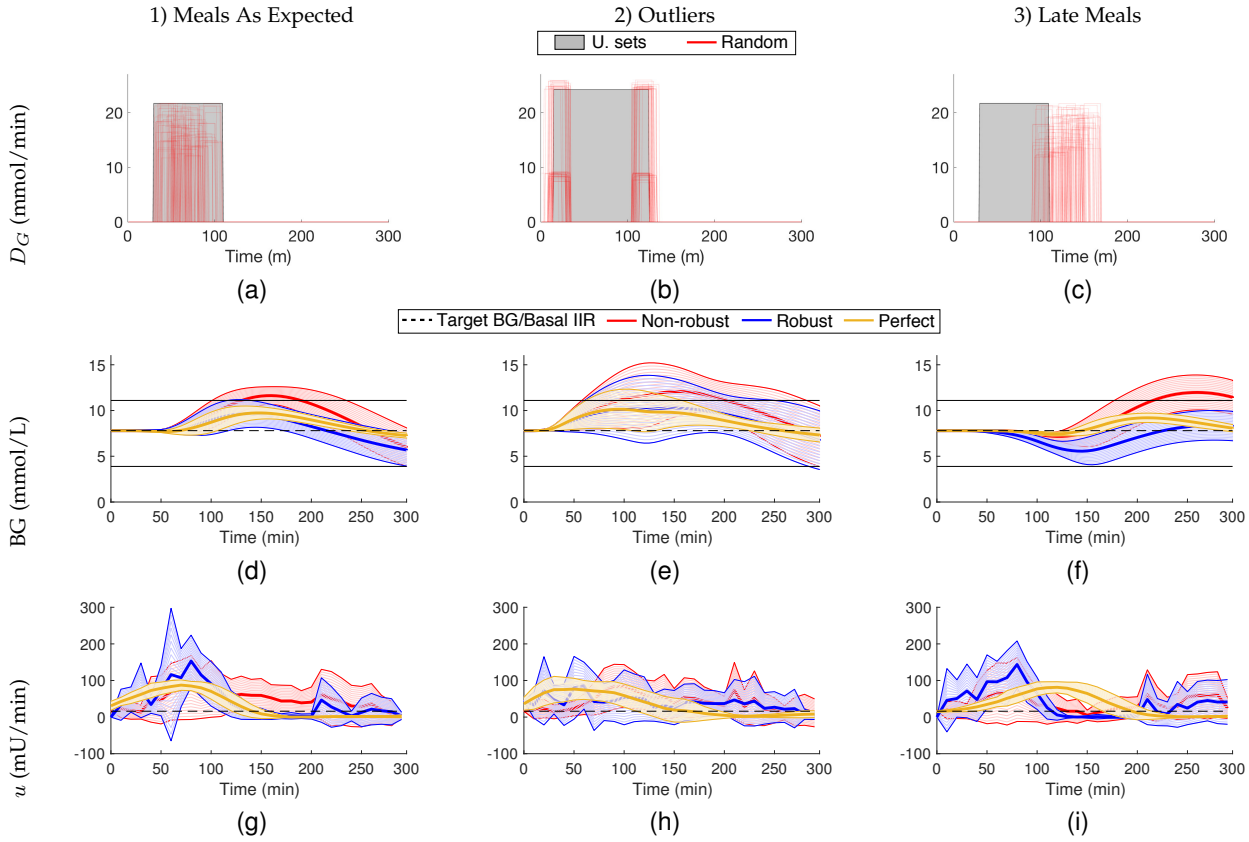

Fig. 3. One-meal, 300-minute experiments (50 repetitions). **Top:** uncertainty sets and random realizations of disturbance D_G (rate of CHO ingestion). **Middle:** BG profiles (with solid black lines indicating the normal BG range). **Bottom:** synthetized insulin therapies. Thick solid lines indicate average BG/insulin values, and are surrounded by an area spanning $\pm$ 1 S.D. Performance indicators are reported in Table 2.

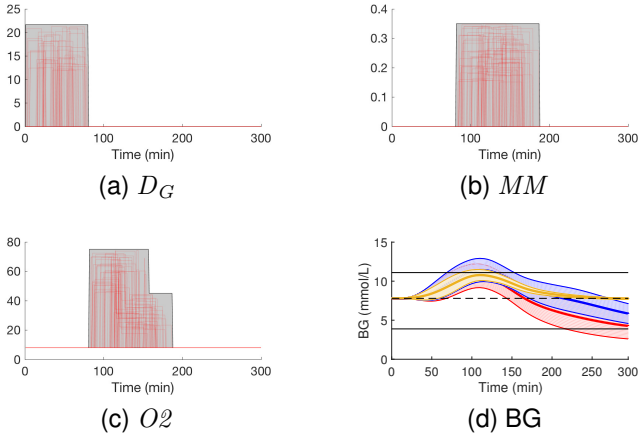

	$t_{<3.9}$	$t_{3.9-11.1}$	$t_{>11.1}$	$BG_{\min}$	$BG_{\max}$	$\sum u$
Perfect	0%	94.83%	5.17 %	7.73	10.93	3.12
Non-rob.	11.87%	80.77%	**7.37%**	4.15	**11.02**	6.16
Robust	**0.57%**	**86.29%**	13.15%	**5.75**	11.65	5.4

Fig. 4. Regulation during random exercise following a meal (50 repetitions). a), b), and c) show uncertainty sets and realizations for rate of CHO ingestion (D_G), active muscular mass (MM), and oxygen consumption ($O2$). Legend is as in Figure 3.

maintain normal BG level. We want to investigate how our data-driven approach can alleviate such negative effects automatically without patient's involvement. From this cluster, we build the corresponding uncertainty sets as explained in Section 4.2 (choosing $\alpha = 0.2$ and $\epsilon = 0.2$).

Due to the poor quality of physical activity data in NHANES, we generated synthetic exercise data consisting of one random exercise episode for each patient, with random start time (outside meal times) and one hour duration. The intensity is randomly chosen among light, moderate, and intense activity, with oxygen consumption and active muscular mass sampled according to the ranges below:

- light: $MM \sim \text{unif}(0.1, 0.25)$, $O2 \sim \text{unif}(15, 45)$;
- moderate: $MM \sim \text{unif}(0.2, 0.35)$, $O2 \sim \text{unif}(45, 75)$;
- intense: $MM \sim \text{unif}(0.3, 0.5)$, $O2 \sim \text{unif}(75, 100)$.

Results for this cluster were obtained with 20 random meal profiles and are reported in Figure 5 (d) and Table 3. Our robust controller has a close-to-ideal performance, with >95% of time spent in normal BG ranges. It outperforms the non-robust controller, which fails to predict the correct BG levels during sleep (time $<$ 500 min), leading to excessive insulin therapy and to dangerous overnight hypoglycemia.

We further experiment with uncertainty sets built from two sub-sets of the selected cluster, shown in Figure 5 (b) and (c), comprising 92 and 91 records, respectively. In this way, we can evaluate how the choice of tighter and less conservative uncertainty sets (induced by the smaller number of records) affect glucose control. We do not register significant differences with the results of the full cluster: our robust AP performs only slightly worse on the two sub-clusters, yielding respectively 90.06% and 92.66% in range.

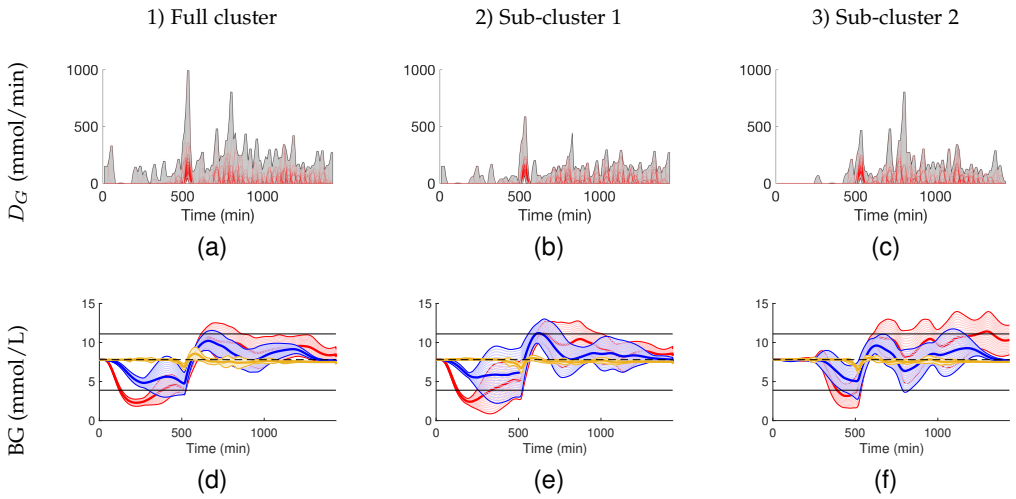

Fig. 5. BG regulation for virtual patient learned from NHANES database (20 repetitions). Uncertainty sets in experiments 2) and 3) are derived from a subset of the data used for 1). Legend is as in Figure 3. Performance indicators are reported in Table 3.

TABLE 3
Performance indicators for the experiments of Figure 5 with NHANES data.

	$t_{<3.9}$	$t_{3.9-11.1}$	$t_{>11.1}$	$BG_{\min}$	$BG_{\max}$	$\sum u$
1), Perfect	0%	100%	0%	6.37	8.92	33.65
1), Non-rob.	20.31%	71.29%	8.4%	2.28	11.71	32.84
1), Robust	**1.18%**	**95.7%**	**3.12%**	**4.2**	**11.18**	32.35
2), Perfect	0%	100%	0%	6.32	8.39	34.2
2), Non-rob.	20.91%	69.99%	9.1%	2.28	**12.38**	33.99
2), Robust	**1.14%**	**90.06%**	**8.81%**	**4.15**	12.48	32.55
3), Perfect	0%	100%	0%	6.3	8.55	33.63
3), Non-rob.	10.3%	64.18%	25.51%	2.77	13.61	27.2
3), Robust	**2.57%**	**92.66%**	**4.76%**	**3.53**	**11.19**	31.71

TABLE 4
High carbohydrate intake simulation parameters of [20]. Meals in the plant are sampled uniformly based on the above intervals and probabilities.

	Chance of occurrence	CHO (g)	Time of day (h)
Breakfast	100%	40-60	6:00-10:00
Snack 1	50%	5-25	8:00-11:00
Lunch	100%	70-110	11:00-15:00
Snack 2	50%	5-25	15:00-18:00
Dinner	100%	55-75	18:00-22:00
Snack 3	50%	5-15	22:00-00:00

5.4 High Carbohydrate Intake Scenario

We assess the behavior of the controller under irregular meal timing and unusually high CHO intake, following the protocol of [20], reported in Table 4. In this protocol, no physical activity is considered and the closed-loop system is simulated for two days. Uncertainty sets were derived following Section 4.3. Results, obtained with 50 repetitions, are shown in Figure 6.

Our robust controller resulted in 81.02% of time within healthy BG ranges, against the 70.53% of the non-robust controller. Hypoglycemia amounts to 1.95% of the total time, but it consists only of minor episodes, as visible by the standard deviation intervals in the plot and by the average minimum BG ($BG_{\min}$ = 3.82 mmol/L) that falls only slightly below the hypoglycemic level (3.9 mmol/L).

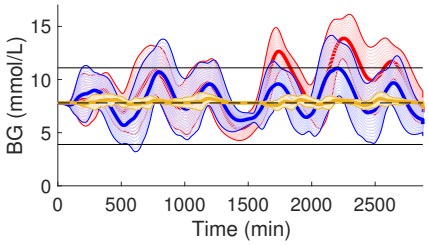

	$t_{<3.9}$	$t_{3.9-11.1}$	$t_{>11.1}$	$BG_{\min}$	$BG_{\max}$	$\sum u$
Perfect	0%	99.52%	0.48%	6.83	10.05	16.19
Non-rob.	**0.7%**	70.53%	28.77%	**4.31**	16	28.82
Robust	1.95%	**81.02%**	**17.04%**	3.82	**14.61**	33.13

Fig. 6. BG profile (top) and performance indicators (bottom) for the high carbohydrate intake scenario (20 repetitions). Legend is as in Fig. 3.

5.5 Validation with Virtual Patient Cohort

In the previous experiments we have evaluated our method using a fixed parametrization of the patient model, that is, a single virtual T1D patient, and random disturbances. In this experiment, we validate our robust AP design on a virtual cohort of 20 T1D patients generated using randomly sampled model parameters from a validated T1D simulator, as described in Section 3.2. Since we want to evaluate the controller under different model parameters, in this experiment we do not consider random disturbances, but a single meal disturbance sequence. The sequence was randomly sampled from the NHANES data used in Section 5.3. We remark that the robust controller knows only the uncertainty sets (shown in Figure 5 a) but not the true disturbances.

Results, reported in Figure 7, show that the robust controller has very good performance on the cohort of virtual subjects, attaining an average time in range of 94.26% and only 0.31% of the time in hypoglycemia. These results align with the performance of the single patient experiment of Section 5.3.

5.6 Evaluation of State Estimator

We chose an MHE scheme for state estimation (see Section 4.1) after having evaluated *extended Kalman filters (EKF)* [28],

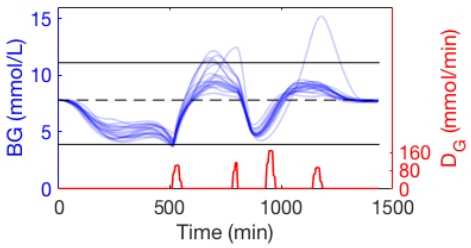

$t_{<3.9}$	$t_{3.9-11.1}$	$t_{>11.1}$	$BG_{\min}$	$BG_{\max}$	$\sum u$
0.31%	94.26%	5.43%	4.01	10.89	23.84

Fig. 7. BG profiles (top) and performance indicators (bottom) obtained by running the robust MPC controller on a cohort of 20 virtual patients. CHO disturbances (red) are fixed and obtained from the NHANES dataset used in Section 5.3.

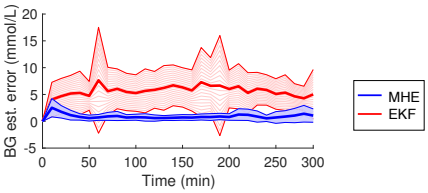

Fig. 8. BG estimation error of Moving Horizon Estimator (MHE) and Extended Kalman Filter (EKF) (50 repetitions).

which are commonly employed for the state estimation of non-linear systems. MHE overcomes some of the typical problems of Kalman filtering, namely, the inability to accurately incorporate state constraints (e.g. non-negative concentrations); poor use of the nonlinear model [29]; and estimations that often diverge, or converge to wrong state predictions [30]. Moreover, Kalman filters are only optimal for linear models and zero-mean normally-distributed disturbances (white Gaussian noise) [10], which cannot capture the random meal and exercise episodes that we are interested in.

We compare the state estimation accuracy between our MHE design and an EKF scheme, according to the *meals as expected* scenario (see Section 5.1). In the EKF, to predict the state estimate at time t, $\hat{\mathbf{x}}(t)$, we use the model of Section 3 as follows: $\dot{\hat{\mathbf{x}}}(t) = F\left(\hat{\mathbf{x}}(t), u^t, \mathrm{E}[\mathbf{d}^t]\right)$, where u^t is the (known) insulin input and disturbances $\mathbf{d}^t$ are replaced with their expected value $\mathrm{E}[\mathbf{d}^t]$. As visible in Figure 8, the MHE outperforms the EKF, with a consistently lower state estimation error. Specifically, the average BG estimation error of EKF is 5.39 mmol/L, compared to an error of 1.09 mmol/L for MHE, meaning that EKF estimations are completely unreliable for insulin control.

5.7 Asymmetric Costs

We evaluate glucose regulation under different asymmetric costs. By choosing $\gamma > 1$ in the controller (Eq. 19), predicted BG trajectories below the target BG level are penalized more than those above the target. As discussed in [21], this strategy contributes to reducing hypoglycemic episodes, motivated by the fact that hypoglycemia leads to more severe consequences than (temporary) hyperglycemia.

We tested the robust controller with $\gamma = 1, 2, 4$ (symmetric, 2x, and 4x penalty, respectively). Simulations were conducted according to the *outliers* scenario (see Section 5.1),

TABLE 5
Performance of different asymmetric cost strategies.

	$t_{<3.9}$	$t_{3.9-11.1}$	$t_{>11.1}$	$BG_{\min}$	$BG_{\max}$
$\gamma = 1$	1.5%	85.35%	13.15%	5.17	11.05
$\gamma = 2$	0%	80.13%	18.87%	5.4	11.38
$\gamma = 4$	0%	76.8%	23.2%	5.6	11.63

which typically generates hypoglycemic episodes and, thus, is an ideal testbed for tuning γ. Table 5 reports the performance indicators obtained by performing 20 repetitions for each value of γ. While for $\gamma = 1$ (symmetric cost) we observe some minor hypoglycemic episodes, hypoglycemia is totally avoided for $\gamma = 2, 4$. Between these two values, we chose $\gamma = 2$ since it yields smaller hyperglycemia. The indicators for average BG peaks and valleys confirm that glucose levels tend to increase with γ.

6 Related Work

Earlier approaches to robust control for the AP (see e.g. [31]) are based on the theory of H_∞ control [32], a technique where the robust controller is synthesized offline as the result of an optimization problem that minimizes the worst-case closed-loop performance of the controlled system. However, H_∞ control only supports linear systems, thus requiring linearization of physiological, non-linear gluco-regulatory models, with inevitable loss of accuracy. Szalay et al. [20] apply robust linear parameter varying (LPV) control, a technique that consists on deriving a piecewise-linear approximation of the non-linear plant and synthesizing a robust H_∞ controller for each linear region.

In contrast to the above techniques, our data-driven robust MPC supports not just meal disturbances, but also physical activity, and is based on non-linear optimization, meaning that it does not require to approximate the system dynamics, leading to more precise predictions. Further, MPC is known to be superior for individualized control strategies [33], even though is computationally more demanding than offline techniques like H_∞ or LPV control, but still feasible within the update periods typical of the artificial pancreas (5-10 minutes). Finally, our data-driven scheme supports continuous learning of the patient's behavior, thus enabling the synthesis of robust and adaptive insulin therapies. On the other hand, H_∞ and LPV controllers are offline and need to be synthesized from scratch in order to adapt to changing patient conditions.

In [34], the authors introduce a data-driven approach for MPC-based insulin control, where the MPC combines multiple linear predictive models, e.g., describing the dynamics under or without disturbances, using prior probabilities of meal occurrence estimated from NHANES data [26]. We also use data from the NHANES database but, in contrast to [34], we learn uncertainty sets that ensure an *a priori* coverage probability of the random meals, and apply these sets in a robust MPC framework. Further, our robust MPC relies on only one (non-linear) predictive model and supports both meal and exercise disturbances.

A simpler strategy employed in a number of AP studies, see e.g. [35], [36], is that of PID control, where the control input results from applying tunable gains to the

error between the system output and a desired setpoint. Synthesizing these gains to obtain robustness guarantees, however, becomes difficult for systems with nonlinear and probabilistic dynamics. To this end, the approach of [37] enables the synthesis of PID gains maximizing the probability that blood-glucose levels stay within a safe range. Being based on SMT solving and probabilistic reachability analysis, the method provides provably correct results, but unlike our work, does not consider data-driven disturbances describing the patient's behavior. Similarly, the work of Kushner and others [38] leverages reachability analysis and data-driven models to find optimal PID gains for the AP. However, this method does not consider meal and exercise disturbances, but rather focuses on the control of gluco-regulatory models estimated from patient data.

7 Conclusions

Thanks to modern wearable devices, patient-specific data about meals and physical activity is becoming readily available, making it possible to offer significantly enhanced personalized medical therapy for type 1 diabetes. Accordingly, we presented a data-driven robust MPC framework for T1D that leverages meal and exercise data to provide enhanced control and state estimation. Our results show that learning a patient's behavior from data is key to achieving fully closed-loop therapy that does not require meal and exercise announcements.

Acknowledgment

Research supported in part by AFOSR Grant FA9550-14-1-0261 and NSF Grants IIS-1447549, CPS-1446832, CNS-1445770, CNS-1445770, CNS-1553273, CNS-1536086, and IIS-1460370.

References

[1] Centers for Disease Control and Prevention, "National diabetes statistics report: estimates of diabetes and its burden in the United States, 2014," *Atlanta, GA: US Department of Health and Human Services*, 2014.

[2] Diabetes Control and Complications Trial Research Group, "The effect of intensive treatment of diabetes on the development and progression of long-term complications in insulin-dependent diabetes mellitus," *New England journal of medicine*, vol. 329, no. 14, pp. 977–986, 1993.

[3] R. Hovorka, "Closed-loop insulin delivery: from bench to clinical practice," *Nature Reviews Endocrinology*, vol. 7, no. 7, pp. 385–395, 2011.

[4] B. Kovatchev *et al.*, "Feasibility of long-term closed-loop control: a multicenter 6-month trial of 24/7 automated insulin delivery," *Diabetes Technology & Therapeutics*, 2017.

[5] "MiniMed 670G insulin pump system," 2018. [Online]. Available: https://www.medtronicdiabetes.com/products/minimed-670g-insulin-pump-system

[6] P. G. Jacobs, N. Resalat, J. El Youssef, R. Reddy, D. Branigan, N. Preiser, J. Condon, and J. Castle, "Incorporating an exercise detection, grading, and hormone dosing algorithm into the artificial pancreas using accelerometry and heart rate," *Journal of diabetes science and technology*, vol. 9, no. 6, pp. 1175–1184, 2015.

[7] D. Bertsimas, V. Gupta, and N. Kallus, "Data-driven robust optimization," *Mathematical Programming*, vol. 167, no. 2, pp. 235–292, 2018.

[8] M. E. Wilinska, M. Bodenlenz, L. J. Chassin, H. C. Schaller, L. A. Schaupp, T. R. Pieber, and R. Hovorka, "Interstitial glucose kinetics in subjects with type 1 diabetes under physiologic conditions," *Metabolism*, vol. 53, no. 11, pp. 1484–1491, 2004.

[9] C. V. Rao, J. B. Rawlings, and D. Q. Mayne, "Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations," *IEEE transactions on automatic control*, vol. 48, no. 2, pp. 246–258, 2003.

[10] J. B. Rawlings, "Moving horizon estimation," *Encyclopedia of Systems and Control*, pp. 1–7, 2013.

[11] M. E. Wilinska, L. J. Chassin, C. L. Acerini, J. M. Allen, D. B. Dunger, and R. Hovorka, "Simulation environment to evaluate closed-loop insulin delivery systems in type 1 diabetes," *Journal of diabetes science and technology*, vol. 4, no. 1, pp. 132–144, 2010.

[12] N. Paoletti, K. S. Liu, S. A. Smolka, and S. Lin, "Data-driven robust control for type 1 diabetes under meal and exercise uncertainties," in *International Conference on Computational Methods in Systems Biology*. Springer, 2017, pp. 214–232.

[13] N. Resalat, J. El Youssef, R. Reddy, and P. G. Jacobs, "Design of a dual-hormone model predictive control for artificial pancreas with exercise model," in *Engineering in Medicine and Biology Society (EMBC), 2016 IEEE 38th Annual International Conference of the*. IEEE, 2016, pp. 2270–2273.

[14] M. Hernandez-Ordonez and D. Campos-Delgado, "An extension to the compartmental model of type 1 diabetic patients to reproduce exercise periods with glycogen depletion and replenishment," *Journal of Biomechanics*, vol. 41, no. 4, pp. 744–752, 2008.

[15] R. Hovorka *et al.*, "Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes," *Physiological Measurement*, vol. 25, no. 4, p. 905, 2004.

[16] N. Paoletti, K. S. Liu, S. A. Smolka, and S. Lin, "Data-driven robust control for type 1 diabetes under meal and exercise uncertainties," *CoRR*, vol. 1707.02246, 2017.

[17] M. Breton and B. Kovatchev, "Analysis, modeling, and simulation of the accuracy of continuous glucose sensors," *Journal of Diabetes Science and Technology*, vol. 2, no. 5, pp. 853–862, 2008.

[18] M. E. Wilinska *et al.*, "Insulin kinetics in type-1 diabetes: continuous and bolus delivery of rapid acting insulin," *IEEE Transactions on Biomedical Engineering*, vol. 52, no. 1, pp. 3–12, 2005.

[19] P. J. Lenart and R. S. Parker, "Modeling exercise effects in type I diabetic patients," *IFAC Proceedings Volumes*, vol. 35, no. 1, pp. 247–252, 2002.

[20] P. Szalay, G. Eigner, and L. A. Kovács, "Linear matrix inequality-based robust controller design for type-1 diabetes model," *IFAC Proceedings Volumes*, vol. 47, no. 3, pp. 9247–9252, 2014.

[21] R. Gondhalekar, E. Dassau, and F. J. Doyle, "Periodic zone-MPC with asymmetric costs for outpatient-ready safety of an artificial pancreas to treat type 1 diabetes," *Automatica*, vol. 71, pp. 237–246, 2016.

[22] B. Houska and M. E. Villanueva, "Robust optimization for MPC," in *Handbook of Model Predictive Control*. Springer, 2019, pp. 413–443.

[23] H. A. David and H. N. Nagaraja, "Order statistics," *Encyclopedia of Statistical Sciences*, vol. 9, 2004.

[24] B. Efron and R. J. Tibshirani, *An introduction to the bootstrap*. CRC Press, 1994.

[25] G. Kahn, T. Zhang, S. Levine, and P. Abbeel, "Plato: Policy learning using adaptive trajectory optimization," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 3342–3349.

[26] Centers for Disease Control and Prevention, "National health and nutrition examination survey," 2018. [Online]. Available: https://www.cdc.gov/nchs/nhanes/

[27] A. Kinmonth and J. Baum, "Timing of pre-breakfast insulin injection and postprandial metabolic control in diabetic children," *Br Med J*, vol. 280, no. 6214, pp. 604–606, 1980.

[28] G. Welch and G. Bishop, "An introduction to the Kalman filter," University of North Carolina at Chapel Hill, Chapel Hill, NC, USA, Tech. Rep., 1995.

[29] E. L. Haseltine and J. B. Rawlings, "Critical evaluation of extended Kalman filtering and moving-horizon estimation," *Industrial & Engineering Chemistry Research*, vol. 44, no. 8, pp. 2451–2460, 2005.

[30] L. Perea, J. How, L. Breger, and P. Elosegui, "Nonlinearity in sensor fusion: divergence issues in EKF, modified truncated GSF, and UKF," in *AIAA Guidance, Navigation and Control Conference and Exhibit*, 2007, p. 6514.

[31] R. S. Parker, F. J. Doyle, J. H. Ward, and N. A. Peppas, "Robust H_∞ glucose control in diabetes using a physiological model," *AIChE Journal*, vol. 46, no. 12, pp. 2537–2549, 2000.

[32] A. A. Stoorvogel, *The H_∞ control problem: a state space approach*. Prentice Hall, 1992.

[33] L. Magni *et al.*, "Model predictive control of glucose concentration in type I diabetic patients: An in silico trial," *Biomedical Signal Processing and Control*, vol. 4, no. 4, pp. 338–346, 2009.
[34] F. Cameron, G. Niemeyer, and B. W. Bequette, "Extended multiple model prediction with application to blood glucose regulation," *Journal of Process Control*, vol. 22, no. 8, pp. 1422–1432, 2012.
[35] S. Laxminarayan, J. Reifman, and G. M. Steil, "Use of a food and drug administration-approved type 1 diabetes mellitus simulator to evaluate and optimize a proportional-integral-derivative controller," *Journal of Diabetes Science and Technology*, vol. 6, no. 6, pp. 1401–1412, 2012.
[36] L. M. Huyett, E. Dassau, H. C. Zisser, and F. J. Doyle III, "Design and evaluation of a robust PID controller for a fully implantable artificial pancreas," *Industrial & Engineering Chemistry Research*, vol. 54, no. 42, pp. 10 311–10 321, 2015.
[37] F. Shmarov, N. Paoletti, E. Bartocci, S. Lin, S. A. Smolka, and P. Zuliani, "SMT-based synthesis of safe and robust PID controllers for stochastic hybrid systems," in *Haifa Verification Conference*. Springer, 2017, pp. 131–146.
[38] T. Kushner, D. Bortz, D. M. Maahs, and S. Sankaranarayanan, "A data-driven approach to artificial pancreas verification and synthesis," in *Proceedings of the 9th ACM/IEEE International Conference on Cyber-Physical Systems*. IEEE Press, 2018, pp. 242–252.

Nicola Paoletti is a lecturer at the Department of Computer Science at Royal Holloway, University of London. He obtained his PhD in Information Sciences and Complex Systems in 2014 at the University of Camerino. From 2014 to 2016, he was a post-doctoral researcher at the Department of Computer Science, University of Oxford. From 2016 to 2018, he was postdoctoral associate at the Department of Computer Science at Stony Brook University. His research interests are in the verification, control, and synthesis of stochastic and hybrid systems, with application to biological and biomedical systems.

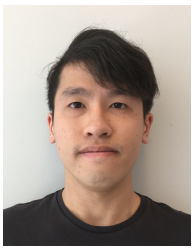

Kin Sum Liu is a Ph.D. candidate at the Computer Science Department of Stony Brook University. His research interests are machine learning and deep learning with applications to graph representation and control system.

Hongkai Chen is a Ph.D. student at the Department of Electrical and Computer Engineering at Stony Brook University. He obtained his MSc degree in Electrical Engineering at Washington University in St. Louis, MO, US in 2016. He obtained his B.S. degree in Electrical Engineering in Nanjing University, China in 2013. His research interests are cyber physical systems, wireless networks, control theory, sensor systems, and activity recognition.

Scott A. Smolka is a SUNY Distinguished Professor of Computer Science at Stony Brook University. He is also a Fellow of the EATCS (European Association on Theoretical Computer Science). His research interests include model checking, runtime verification, and the modeling and analysis of cardiac tissue and other complex systems. He is the Lead PI for the multi- institutional NSF CPS Frontiers project on CyberCardia: Compositional, Approximate, and Quantitative Reasoning for Medical Cyber-Physical Systems. He is perhaps best known for the algorithm he and Paris Kanellakis invented for deciding bisimulation. Smolka's research has resulted in more than 200 publications, generating more than 8,350 citations. He has also been PI/Co-PI on grants totaling more than $23M.

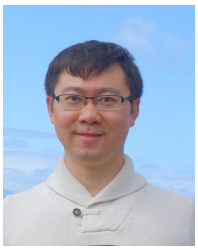

Shan Lin is an assistant professor of the Department of Electrical and Computer Engineering in Stony Brook University. He received his PhD in computer science at the University of Virginia. His research is in the area of networked systems, cyber physical systems, and Internet of Things. He currently works on wireless networking, medical systems and devices, smart buildings, and smart transportation systems. He received the NSF Career award in 2016.