

RESEARCH REPORT

Multivariate learning framework for long-term adaptation in the artificial pancreas

Dawei Shi | Eyal Dassau | Francis J. Doyle III

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

Correspondence

Francis J. Doyle III, Harvard John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138. Email: frank_doyle@seas.harvard.edu

Funding information

National Institute of Diabetes and Digestive and Kidney Diseases, Grant/Award Numbers: DP3DK104057 and UC4DK108483

Abstract

The long-term use of the artificial pancreas (AP) requires an automated insulin delivery algorithm that can learn and adapt with the growth, development, and lifestyle changes of patients. In this work, we introduce a data-driven AP adaptation method for improved glucose management in a home environment. A two-phase Bayesian optimization assisted parameter learning algorithm is proposed to adapt basal and carbohydrate-ratio profile, and key feedback control parameters. The method is evaluated on the basis of the 111-adult cohort of the FDA-accepted UVA/Padova type 1 diabetes mellitus simulator through three scenarios with lifestyle disturbances and incorrect initial parameters. For all the scenarios, the proposed method is able to robustly adapt AP parameters for improved glycemic regulation performance in terms of percent time in the euglycemic range [70, 180] mg/dl without causing risk of hypoglycemia in terms of percent time below 70 mg/dl.

KEYWORDS

artificial pancreas, data-driven algorithms, parameter adaptation, safety-critical control

1 | INTRODUCTION

For several decades, engineers and clinicians have been working on artificial pancreas (AP) systems to achieve automated blood glucose (BG) regulation for patients with Type 1 diabetes mellitus (T1DM). With the encouraging outcome of recent large-scale/long-duration outpatient clinical trials,¹⁻⁴ current AP systems require user intervention around meals and exercise. Since clinical factors such as basal rate (BR) and carbohydrate ratio (CR) are not fixed but do change for patients with Type 1 diabetes, adaptation and learning are needed and expected from AP systems to improve outcome and enable better control as insulin sensitivity changes. In addition, the effects of physical activity on glucose utilization, production, insulin sensitivity, and absorption strongly depend on its timing, type, and intensity,^{5,6} which vary with lifestyle changes. This forms another motivation to incorporate adaptation in AP design.

To achieve this goal, different approaches to AP adaptation have been investigated in the literature. Automated BR modulation was considered in recent outpatient clinical trials^{1,2} and the first U.S. Food

and Drug Administration (FDA) approved commercial hybrid AP system in the United States.³ With different degrees of automation and clinician effort, pump setting adaptation was also considered in multiple clinical studies,⁷⁻⁹ although the detailed techniques were not explicitly reported. In a 12-week outpatient study,⁴ an automated algorithm monitored by study physicians was employed to perform joint adaptation of BR and CR. In earlier investigations,^{10,11} a run-to-run approach was proposed to update BR and meal bolus sizes based on sparse BG measurements; this approach was recently revisited in Reference 12 to adapt CR profile during the day and BR at night based on continuous glucose monitor (CGM) measurements, and was verified through multiweek simulations that took account of insulin sensitivity circadian rhythm. A closely related approach that can exploit CGM measurements and continuous insulin delivery is iterative learning control, which can be regarded as a two-timescale enhancement of run-to-run methods.¹³ In Reference 14, iterative learning model predictive control (MPC) was proposed to adapt the reference trajectory of the closed-loop controller used for glucose regulation, which was tested in a pilot study recently.¹⁵

through long-term AP adaptation algorithms, based on historical performance metric records from the lower layer. With the successful deployment of large and long-term out-patient clinical studies, long-term AP adaptation has gained in importance only recently, and is the main focus of this article.

With the dual-layer control scheme, we are now ready to introduce the parameter adaptation problem to be investigated. Specifically, the parameters to be adapted include

1. segments $\{\beta_n\}$ that compose the BR profile $\beta := [\beta_1, \beta_2, \dots, \beta_N]$ used in the BR calculator. Note that each β_n is associated with a time period T_n^β during which it is active. In this work, we assume that T_n^β is fixed and predetermined by clinicians.
2. segments $\{\gamma_n\}$ that compose the CR profile $\gamma := [\gamma_1, \gamma_2, \dots, \gamma_M]$ used in the meal controller. Similar to β , each segment γ_n is used for a particular time period T_n^γ , which is assumed to be fixed by clinicians as well.
3. parameters in the closed-loop controller. Three parameters in the zone MPC developed in Reference 20 are considered, including \hat{R} that represents the control penalty parameter for insulin delivery above the BR, \bar{D} that determines the upper bound of a glucose zone for which the velocity penalty is active, and a coefficient γ_{IOB} that determines the responsiveness of the insulin-on-board (IOB) constraint.

The goal of this work is to develop an automatic parameter learning algorithm that can correctly adapt the parameters in the lower-layer control algorithms and is robust to lifestyle disturbances, with minimal patient/clinician involvement and without causing risks of hypoglycemia during the adaptation procedure.

To achieve this goal, a data-driven multivariate parameter learning framework is proposed. Considering the different roles of the lower-layer control algorithms, the adaptation procedure is divided into two phases. In Phase I, the parameters in the feedforward control algorithms (namely, the BR and CR profiles) are optimized. Based on the obtained/updated BR and CR profiles, parameters in the feedback control algorithms are adjusted in Phase II. In both phases, we consider the challenging but realistic case that the patient is under closed-loop control, such that the “open-loop” parameters are adjusted in Phase I based on closed-loop data. The detailed adaptation algorithms are presented in Sections 2.1 and 2.2, respectively.

2.1 | Learning feedforward control parameters

In this subsection, we focus on adapting the parameters in the BR and CR profiles, which is Phase I of the adaptation procedure. The aim here is to obtain reasonable rather than optimal profiles, such that an appropriate operating point is provided for the feedback controller, which helps enhance the safety of closed-loop glucose control. More importantly, considering the fact that the feedback control may not be available for some periods (e.g., when the controller runs out of battery or lost CGM connection), the obtained parameter needs to be “safe” in the sense that no hypoglycemia would be caused when the

patient loses closed-loop control. In the following, we first separately introduce the methods for BR and CR adaptation, based on which a hybrid time- and event-triggered method is introduced to iterate the procedures that adapt BR and CR.

2.1.1 | Adaptation of BR profile

Intuitively, the BR profile β is designed to manage “healthy” fasting glucose levels without considering meal-induced glucose excursions. However, when the fasting glucose levels (measured as glucose levels at night or before meals) are not satisfactory, it is difficult to diagnose which segment β_n in β is the root cause. This problem becomes more complicated when the effects of meal boluses are considered, due to a further loss of “identifiability.”²¹ A helpful observation, however, is that the “effective” real-time BR is ultimately determined by the feedback controller, which is able to adjust a potentially inappropriate BR provided by the BR profile to a certain extent. This observation allows us to perform BR profile adaptation using a controller-led approach, without the need of explicitly performing root cause diagnosis. To reduce the risk of controller-induced hypoglycemia, safety constraints are designed to eliminate the adoption of aggressive BR profiles. Through controller-led BR adaptation, the problem of lack of identifiability/diagnosability can be automatically solved, which leads to a direct separation of the adaptation of BR and CR profiles. The difficulty of algorithm design and implementation, however, is independent of the number of segments in the BR profile.

Controller-led BR profile update

BR adaptation determines the values of BR segments $\{\beta_n^{k+1}\}$ at iteration $k + 1$ based on $\{\beta_n^k\}$ and available glucose and insulin delivery information. The idea is to update the BR segment β_n with the averaged nonmeal-related insulin microboluses commanded by the feedback controller during the same time interval, namely, T_n^β . Although different methods can be used to distinguish meal/nonmeal-related insulin, we consider a simple approach in this work and define nonmeal-related insulin as the insulin deliveries that happen τ_m hours after the previous meal. This defines a set of admissible meal-related insulin deliveries \mathcal{I}_n^d for each BR segment β_n :

$$\mathcal{I}_n^d := \{(t, I_t^d) | t \in T_n^\beta, t \geq T_m^d(t) + \tau_m\}, \quad (1)$$

where I_t^d denotes an insulin delivery at time instant t on day d , T_n^β denotes the time period for which β_n is active, and $T_m^d(t)$ denotes the time of the previous meal that happens before time t on day d . We take $\tau_m = 3$ hr in our implementation. An initial BR estimate $\bar{\beta}_n^k$ for β_n^k can be written as

$$\bar{\beta}_n^k = \frac{\sum_{d=1}^{n_d} \sum_{t=1}^{n_w} I_t^d \cdot \mathbf{1}((t, I_t^d) \in \mathcal{I}_n^d)}{\sum_{d=1}^{n_d} \sum_{t=1}^{n_w} \mathbf{1}((t, I_t^d) \in \mathcal{I}_n^d)}, \quad (2)$$

where n_d denotes the number of insulin microboluses in a day, and n_w denotes the number of days in iteration k of the adaptation process, and $\mathbf{1}(\cdot)$ denotes the indicator function. Although the actual real-time

BR will be decided by the feedback control algorithm (using the knowledge of β_n), the key consideration on the obtained $\bar{\beta}_n$ is that it should not overestimate the “true” BR and cause hypoglycemic events, which is particularly important when the feedback controller is unavailable. To address this concern, two types of safety constraints are incorporated, as will be introduced below.

Statistical IOB constraint

Similar to the case with lower-layer control algorithms, the IOB information can be exploited to prevent aggressive behavior of the parameter adaptation algorithm. On the basis of the IOB calculation used in zone MPC,²⁰ a statistical IOB constraint is proposed for BR adaptation. Specifically, a dynamic database $\mathcal{D}(\beta_n)$ is built to eliminate overestimated values of β_n , which stores the information of a triplet $\{\text{IOB}_n^i, \beta_n^i, \gamma_{m\sim n}^i\}$ that have led to low-glucose events in history for $i \in \{1, 2, \dots, k\}$, with IOB_n^i denoting the averaged IOB value immediately before β_n becomes active in adaptation iteration i , and $\gamma_{m\sim n}^k$ denotes the values CR used within T_n^β in iteration k . When $\bar{\beta}_n^{k+1}$ is computed for iteration $k+1$ and $\bar{\beta}_n^{k+1} > \beta_n^k$, the IOB constraint is enforced and implemented as

$$\left[\text{IOB}_n^k, \bar{\beta}_n^{k+1}, 1/\gamma_{m\sim n}^k \right] \leq \left[\text{IOB}_n^i, 0.95\beta_n^i, 1/\gamma_{m\sim n}^i \right], \quad (3)$$

for all $\{\text{IOB}_n^i, \beta_n^i, \gamma_{m\sim n}^i\} \in \mathcal{D}(\beta_n)$, where “ \leq ” denotes an element-wise partial order holds for the two vectors compared. Here, IOB_n^k is used to estimate the average IOB value immediately before β_n becomes active in adaptation iteration $k+1$. If any of the IOB constraints are violated, $\bar{\beta}_n^{k+1} = 0.95\beta_n^k$ to avoid causing hypoglycemia risk.

Smoothness constraint

This constraint ensures that the BR segment β_n does not change too much compared with its neighboring segments β_{n^-} and β_{n^+} , with $n^- = N - \text{mod}(N - n + 1, N)$ and $n^+ = \text{mod}(n, N) + 1$, respectively. The underlying intuition is that the insulin requirement of the metabolic system to maintain a healthy fasting glucose level does not change abruptly in time. Mathematically, this constraint is implemented as

$$\beta_n^{k+1} = \min \left(\beta_n^{k+1}, \lambda_s^\beta \min \left(\bar{\beta}_{n^-}^{k+1}, \bar{\beta}_{n^+}^{k+1}, \beta_{n^-}^{k+1}, \beta_{n^+}^{k+1} \right) \right), \quad (4)$$

where λ_s^β denotes the “smoothness coefficient” and is selected as 1.3 in our work to compromise smoothness with performance.

2.1.2 | Adaptation of CR profile

Different from BR adaptation, the effect of different CR segments on glucose profiles are relatively isolated, as meals are usually several hours apart. This helps decouple the adaptation problems for different γ_n values, for which a data-driven optimization approach can be developed.

Zone objectives

Instead of optimizing the CR profile for a specific glycemic metric, our goal for CR adaptation is to only ensure that the average BG levels $\{\gamma_n^i | n=1, 2, \dots, M\}$ after meals are taken for τ_γ hours should settle

within a certain zone $[\underline{\gamma}^\gamma, \bar{\gamma}^\gamma]$; in our implementation, this zone is meal-dependent and is selected as [125, 155], [135, 165], and [125, 155] mg/dl for breakfast, lunch, and dinner, respectively, based on simulation data from the UVA/Padova simulator. On one hand, this choice of a control-to-range objective helps efficiently handle the uncertainties caused by lifestyle disturbances (e.g., sizes and timing of meals). Conversely, as the final value of meal boluses are normally decided by the patient, the meal bolus sizes calculated using the CR profile are used as references, thus a robust estimate of CR would be more practical from an application perspective. In this work, τ_γ is selected to be 4 or the length of time elapsed before the next meal is taken if it is less than 4 hr, based on the observations of the UVA/Padova Simulator.

Parameter optimization with safety constraints

To adapt the CR profile toward the target zone, a virtual optimization problem is formulated for each γ_n . The cost function is selected as the average postprandial BG levels

$$y_n^\gamma := f_n^\gamma(\gamma_n, \theta_n^\gamma), \quad (5)$$

where $f_n^\gamma(\gamma_n, \theta_n^\gamma)$ is used to represent the underlying unknown dependency of y_n^γ on γ_n and other parameters θ_n^γ . The adaptation process is then performed by solving a sequence of constrained optimization problems with this unknown cost function:

$$\min_{\gamma_n^{k+1}} f_n^\gamma(\gamma_n^{k+1}, \theta_n^\gamma) \quad (6)$$

$$\text{s.t. } |\gamma_n^{k+1} - \gamma_n^k| \leq \lambda^\gamma \gamma_n^k, \quad (7)$$

$$|f_n^\gamma(\gamma_n^{k+1}, \theta_n^\gamma) - f_n^\gamma(\gamma_n^k, \theta_n^\gamma)| \leq \delta^{y_n^\gamma} \quad (8)$$

$$\gamma_n^{k+1} \geq \underline{\gamma}_n^{k+1}. \quad (9)$$

To ensure the smoothness of the adaptation procedure, three safety constraints are considered in the optimization problem above. The first and second constraints, respectively, restrict the rate of change of γ_n and $f_n^\gamma(\gamma_n, \theta_n^\gamma)$. In our implementation, λ^γ is chosen as 30%, and $\delta^{y_n^\gamma}$ is selected as 12 mg/dl. The third constraint directly bounds γ_n from below to avoid hypoglycemia risks caused by an underestimated CR; the lower bound $\underline{\gamma}_n^{k+1}$ is updated dynamically by taking the maximum value of γ_n^i that has caused risk of hypoglycemia during the adaptation process. Note that γ_n^k and $\underline{\gamma}_n^{k+1}$ are known and $f_n^\gamma(\gamma_n, \theta_n^\gamma)$ can be calculated based on historical CGM measurements when $f_n^\gamma(\gamma_n^{k+1}, \theta_n^\gamma)$ is optimized for γ_n^{k+1} , but the explicit expression of $f_n^\gamma(\gamma_n^{k+1}, \theta_n^\gamma)$ is generally not known. To solve this problem, a data-driven BO-assisted algorithm will be provided in Section 2.3.

2.1.3 | Alternating BR and CR adaptation

Considering the coupling effects of BR and CR on glucose management, the adaptation of $\{\beta_n | n = 1, 2, \dots, N\}$ and $\{\gamma_n | n = 1, 2, \dots, M\}$ are performed in an interactive fashion. For safety consideration, only one profile (BR or CR) is adapted in each iteration. To do this, BR adaptation is performed in a time-driven manner such that the values of $\{\beta_n | n = 1, 2, \dots, N\}$ are adapted for N^β consecutive iterations, which is set to 2 in our implementation. CR adaptation is performed in a

combined time- and event-triggered fashion; the adaptation procedure ends either when the zone objectives are achieved or when a maximum number of iterations N^Y is reached, which is set to 5 in our implementation. As meal boluses can be adjusted by the patients based on their knowledge and preference, the overall feedforward parameter adaptation phase begins by adapting the BR profile and iterates according to the alternating procedure described. An illustration of the procedure is provided in Figure 2. The phase terminates when the zone objective criterion for CR adaptation remains valid after performing the BR adaptation.

2.1.4 | Choice of parameters

A few parameters are adjustable in the proposed BR and CR adaptation procedure, including λ^Y , δ^{Y_n} , N^β , and N^Y , the values of which used in the in silico studies of this work have been provided in the above subsections. Specifically, λ^Y and δ^{Y_n} are safety factors that limit the rates of change of the elements in the CR profile, while N^β and N^Y are the maximum allowable inner iterations to adapt BR and CR profiles (see also Figure 2) and thus jointly determine the duration of the adaptation process. With these physical interpretations, other values of these parameters can be selected based on clinical experience and knowledge of particular patients in clinical studies.

2.2 | Learning feedback control parameters

In this subsection, we focus on Phase II of the adaptation procedure, which deals with parameter adaptation for feedback control. As the appropriately adjusted BR and CR profiles in Phase I provide optimized operating conditions for the closed-loop controller, only moderate changes on the key parameters are needed to achieve satisfactory glucose regulation. Based on this observation, the approach here is to first determine the bottle-neck parameter that limits the performance of closed-loop control for a specific patient, and then to dynamically learn the appropriate value of the selected parameter. Ideally, improved performance could be potentially obtained by considering combined dynamic parameter selection and adaptation, but the improvement comes with compromised risk of hypoglycemia and time needed to complete the adaptation procedure.

2.2.1 | Parameter selection based on sensitivity analysis

Generally, parameter selection determines the parameter $\varphi \in \{\hat{R}, \bar{D}, \gamma_{IOB}\}$ to be adapted in Phase II. In this work, a sensitivity analysis approach is utilized to achieve automatic parameter selection, by

rerunning the closed-loop control algorithm with different parameter settings using the most recent historical glucose measurements for a specific patient, which is the so-called advisory-mode analysis.²² To do this, for each $\varphi \in \{\hat{R}, \bar{D}, \gamma_{IOB}\}$, a pair of upper and lower bounds $\{\varphi^+, \varphi^-\}$ that limit the range of feasible choices of φ in the adaptation procedure are specified; in our implementation, $\{\varphi^+, \varphi^-\}$ is chosen as {150%, 50%}, {122%, 89%}, and {120%, 70%} of the corresponding nominal values for \hat{R} , \bar{D} , and γ_{IOB} , respectively. Advisory mode comparisons are performed by separately choosing $\varphi = \bar{\varphi}$ and $\varphi = \underline{\varphi}$ and keeping parameters in $\{\hat{R}, \bar{D}, \gamma_{IOB}\} \setminus \{\varphi\}$ unchanged. The sensitivity analysis is then performed by solving

$$\varphi^* = \operatorname{argmin}_{\varphi \in \{\hat{R}, \bar{D}, \gamma_{IOB}\}} |\mathcal{I}(\varphi^+) - \mathcal{I}(\varphi^-)|, \quad (10)$$

where $\mathcal{I}(\varphi^+)$ and $\mathcal{I}(\varphi^-)$ denote the total amount of insulin calculated in the advisory mode analysis for $\varphi = \varphi^+$ and $\varphi = \varphi^-$, respectively. This procedure determines the parameter that has the "strongest" controllability of insulin delivery, hence the closed-loop glycemic regulation can be adjusted most efficiently.

2.2.2 | Parameter optimization

The goal of adapting the selected parameter φ is to achieve satisfactory average glucose level without having risk of hypoglycemia, which implies improved percentage time in euglycemic range [70, 180] mg/dl as the glucose profile is restricted below by considering constraints on hypoglycemia. To achieve this goal, we again formulate the adaptation procedure as a sequential optimization problem:

$$\min_{\varphi^{k+1}, \theta^\varphi} f^\varphi(\varphi^{k+1}, \theta^\varphi) \quad (11)$$

$$\text{s.t. } |\varphi^{k+1} - \varphi^k| \leq \lambda^\varphi \varphi^k, \quad (12)$$

$$|f^\varphi(\varphi^{k+1}, \theta^\varphi) - f^\varphi(\varphi^k, \theta^\varphi)| \leq \delta^{Y^\varphi}, \quad (13)$$

$$\max\{\varphi^-, \underline{\varphi}^{k+1}\} \leq \varphi^{k+1} \leq \min\{\varphi^+, \bar{\varphi}^{k+1}\}. \quad (14)$$

Here, φ^{k+1} denotes the value of φ at the $(k+1)$ th iteration, $f^\varphi(\varphi^{k+1}, \theta^\varphi)$ denotes the average glucose value obtained using φ^{k+1} with θ^φ being potential dependent parameters. The constraints in Equations 12 and 13 restrict the rate of change of φ and $f^\varphi(\varphi, \theta^\varphi)$, respectively. In our implementation, λ^φ is selected as 30%, and δ^{Y^φ} is chosen as 6 mg/dl; the roles of these two parameters are identical to those of λ^Y and δ^{Y_n} discussed in Section 2.1.4. The constraints in Equation 14 bound the feasible region of φ ; in addition to φ^- and φ^+ , $\underline{\varphi}^{k+1}$ and $\bar{\varphi}^{k+1}$ provide additional dynamic bounds based on historical values of φ , namely, $\{\varphi^1, \dots, \varphi^k\}$, to help avoid hypoglycemia risks. The sequential optimization procedure ends either when average glucose level becomes satisfactory (less than 135 mg/dl in our implementation) or when the bounds in Equation 14 become active, the latter of which means the controller has achieved its performance limitation. Similar to the case of adapting γ_n , the analytical expression of the cost function $f^\varphi(\varphi^{k+1}, \theta^\varphi)$ is not known. A BO-assisted optimization algorithm will be introduced to solve this problem.

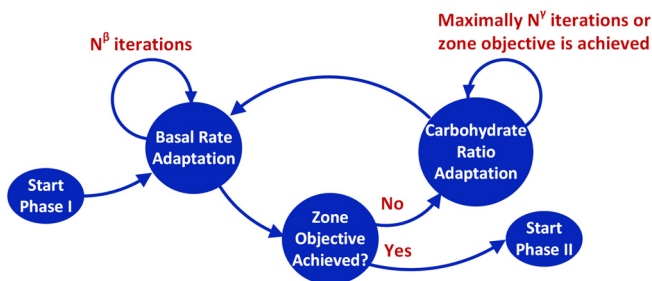


FIGURE 2 Adaptation of feedforward control parameters

2.3 | BO-assisted parameter adaptation

In Sections 2.1.2 and 2.2.2, optimization problems with unknown objective functions and constraints need to be solved. Specifically, the two problems in Equations 6-9 and 11-14 share the form

$$\min_{\psi^k} f(\psi^k, \theta^\psi) \quad (15)$$

$$\text{s.t. } |f(\psi^k, \theta^\psi) - y_{k-1}| \leq \delta, \quad (16)$$

$$g(\psi^k) \leq 0, \quad (17)$$

where $f(\psi^k, \theta^\psi)$ is an unknown function of ψ^k parameterized by θ^ψ , y_{k-1} is a noisy measurement of $f(\psi^{k-1}, \theta^\psi)$, δ is a known parameter, and $g(\psi^k)$ is a known (linear) function of ψ^k . Noticing the fact that a noisy measurement/estimate y_k of $f(\psi^k, \theta^\psi)$, which is average glucose level for a fixed ψ^k , is available, we solve this problem through designing a BO-assisted optimization algorithm, to make efficient use of historical measurements of $\{f(\psi^k, \theta^\psi) | k \in \pi_k\}$ for an admissible set of iteration indexes π_k ; here, π_k is defined as the set of iteration indexes for which the same set of parameters are adapted while other adaptation parameters are kept constant. The main idea is to obtain a data-driven estimate $\hat{f}(\psi_k, \theta_k | \mathcal{D}_k)$ of $f(\psi^k, \theta^\psi)$ based on available historical data \mathcal{D}_k . In the BO literature, different forms of $\hat{f}(\psi_k, \theta_k | \mathcal{D}_k)$ have been proposed.²³ As the size of data set for adaptation is small compared with those available for machine learning problems, a linear kernel is adopted in this work:

$$\hat{f}(\psi^k, \theta_k | \mathcal{D}_k) := \theta_{k,1} \psi^k + \theta_{k,2} \quad (18)$$

with $\theta_k := [\theta_{k,1}, \theta_{k,2}]^\top$. Note that to make efficient use of data, the value of θ_k is made ψ -dependent and is obtained based on data segment $\bar{\mathcal{D}}_k \subseteq \mathcal{D}_k$, which is the subset of data \mathcal{D}_k corresponding to iteration indexes in π_k . This further explains the rationale for using a linear kernel: the cost function in Equation 18 simply represents a local linearization of the unknown cost function $f(\psi^k, \theta^\psi)$ around the adapted values of ψ^k .

To solve Equations 15-17, a BO-based algorithm is proposed (see Algorithm 1), which iteratively adapts tuning parameter ψ_k until the problem-dependent terminal conditions are satisfied. For each iteration, the algorithm starts with updating the admissible set π_k , and obtains $\bar{\mathcal{D}}_k \subseteq \mathcal{D}_k$ based on π_k (line 2 of the algorithm). For the problems in Equations 6-9 and 11-14, the effect of the adaptation parameter ψ on $f(\psi^k, \theta^\psi)$ (which can be understood as the sign of the partial derivative) is known, and can be exploited to determine the search direction \mathcal{S}_k (line 3); this helps ensure that the algorithm can evolve along the correct direction with noisy measurements $\{y_k\}$. If there is not enough data to perform data fitting (line 4), heuristic adjustment of φ_k will be performed along \mathcal{S}_k for safety concerns, where δ^S is selected as 10% in our implementation. For all other scenarios, φ_k is adjusted through the proposed BO procedure (lines 5-11). The BO first estimates θ_k based on $\bar{\mathcal{D}}_k$ (line 6). With the obtained θ_k , ψ_k is calculated by solving a constrained optimization problem (line 7). As $\hat{f}(\varphi_k, \theta_k | \mathcal{D}_k)$ is linear with respect to φ_k , the optimization problem is actually a linear programming problem. To enhance the robustness of the algorithm against noises in $\{y_k\}$, the deviation of ψ_k from ψ_{k-1} is compared with \mathcal{S}_k ; the value of ψ_k will be recalculated if inconsistency is observed (line 8). Safety checks are further performed for ψ_k (line 10), where

the value of ψ_k is truncated if the constraint is violated. The obtained φ_k is then implemented to obtain y_k , which is collected to update \mathcal{D}_k (line 12).

3 | RESULTS

In this section, we evaluate the proposed adaptation algorithm through multiple-month simulations on the 111-patient cohort of the US FDA accepted UVA/Padova simulator.¹⁶ To mimic intra-day variations in insulin sensitivity and the effect of dawn phenomenon, the diurnal patterns of the parameters in the endogenous glucose production and glucose utilization models are implemented according to Reference, 24 Inter-day variation in insulin sensitivity is further introduced by adding random noise obeying a uniform distribution within $\pm 5\%$ of the nominal values of the parameters that determine insulin sensitivity and dawn phenomenon. To consider the effect of illness on insulin sensitivity, a healthy in silico patient has a 5% chance of entering a sick state that can last up to five consecutive days; when an illness event happens, it can either increase by 50% or decrease by 100% the magnitudes of the insulin sensitivity and dawn phenomenon parameters with probabilities of 0.5, throughout the illness period. We assume that the state of sickness is not announced to the adaptation algorithm. To mimic lifestyle disturbances, the in silico subjects take breakfast, lunch, and dinner with normally distributed meal sizes (with means and standard deviations equal to [50, 75, 75] g and [3, 4, 4] g carbohydrate [CHO]) and meal times uniformly distributed in [07:00, 09:00], [11:00, 13:00], and [18:00, 20:00], respectively; in addition, each meal can be skipped with probability 0.1. The CGM measurement noise is generated according to a random noise seed on each day. We assume that the meals are all fully announced but the meal boluses are calculated with potentially inappropriate CR profiles. The updated parameters obtained in each iteration are used for 1 week (7 days)¹⁰ before the next iteration, so that enough data can be collected for performance evaluation. For implementation purpose, we assume the BR profile contains five segments, the effective period of which are [02:00, 05:00], [05:00, 10:00], [10:00, 16:00], [16:00, 21:00], and [21:00, 02:00], respectively; the CR profile are composed of four segments, the effective period of which are [05:00, 10:00], [10:00, 16:00], [16:00, 21:00], and [21:00, 05:00], respectively; only the first three segments that are responsible to breakfast (B), lunch (L), and dinner (D) are adapted and the segment that effective overnight is set to the default value from the simulator.

To evaluate the safety, effectiveness, and robustness of the adaptation algorithm, three scenarios are considered. In the first scenario (Scenario I), the patients are assumed to have doubled CR and halved BR profile segments compared with the default values in the simulator; both of these settings will lead to increased hyperglycemia due to conservative insulin delivery. In the second scenario (Scenario II), the patients are initiated with doubled CR and doubled BR profiles; the former would cause conservative meal boluses but the latter would counteract with relatively larger insulin microboluses, which makes it challenging for the adaptation algorithm to identify the appropriate tuning parameters. The third scenario (Scenario III) mimics real-life situations, in which different segments in the BR profile and CR profile

TABLE 1 Glycemic metrics obtained before and after the proposed adaptation algorithm

Metrics	Scenario I			Scenario II			Scenario III		
	Week 1	Week 24	p value	Week 1	Week 24	p value	Week 1	Week 24	p value
Percentage time < 54 mg/dl	0.0 (0.0)	0.0 (0.0)	.017*	7.0 (7.6)	0.0 (0.0)	<.001*	0.0 (0.0)	0.0 (0.0)	.763
Percentage time < 70 mg/dl	0.0 (0.0)	0.0 (0.4)	<.001*	15.7 (8.0)	0.0 (0.1)	<.001*	0.1 (1.0)	0.0 (0.2)	.077
Percentage time 70–180 mg/dl	41.1 (18.2)	88.1 (9.5)	<.001*	75.6 (9.9)	88.8 (10.2)	<.001*	81.6 (12.4)	88.7 (13.8)	<.001*
Mean BG (mg/dl)	194.3 (15.9)	142.3 (8.8)	<.001*	110.5 (10.8)	142.0 (8.9)	<.001*	144.3 (14.2)	142.4 (9.1)	.756

Data in this table are shown as median (inter quartile range). Statistical significance is assessed by Wilcoxon signed-rank test. Statistically significant ($p < .05$) changes are highlighted in bold with asterisks.

may deviate within a reasonable range from either above or below the corresponding appropriate values; to do this, the in silico subjects are initialized with randomized BR and CR profiles, in which the segments are generated according to uniform distributions within [50%, 150%] of the corresponding default values. For all in silico subjects, we assume the zone MPC developed in Reference 20 with default parameters is used. The parameters in the adaptation algorithm are specified in Sections 2.1–2.3. For all three scenarios, the first week is utilized to collect data for initialization and thus no parameter is adjusted; the adaptation process starts from the second week. All simulations are run for 24 weeks. The key glycemic metrics obtained before and after the proposed adaptation algorithm are provided in Table 1.

3.1 | Results for Scenario I

The results for Scenario I are shown in Figures 3 and 4. Figure 3 provides the trends of key glycemic metrics during the adaptation

procedure, including average BG levels, percentage time in the euglycemic range [70, 180] mg/dl, percentage time below 70 mg/dl, and percentage time below 54 mg/dl. Due to the joint effect of underestimated BR and overestimated CR profiles, the in silico subjects have elevated glucose levels on Week 1. From Week 2, monotonic and steady improvements are achieved by the proposed adaptation algorithm, and a trend of convergence in the performance metrics is observed around Week 12 as no significant changes happened in the rest of the simulations. For the 111-subject cohort, the adaptation algorithm manages to improve glycemic control performance dramatically in terms of both average glucose levels (from 194.3 mg/dl [Week 1] to 142.3 mg/dl [Week 24]; $p < .001$) and average percent time in euglycemia range [70, 180] mg/dl (from 41.0% to 88.1%; $p < .001$). The adaptation algorithm is safe in the sense that no hypoglycemia risk is caused during the process.

The trends of parameter changes during the adaptation procedure are provided in Figure 4. The important observation here is that

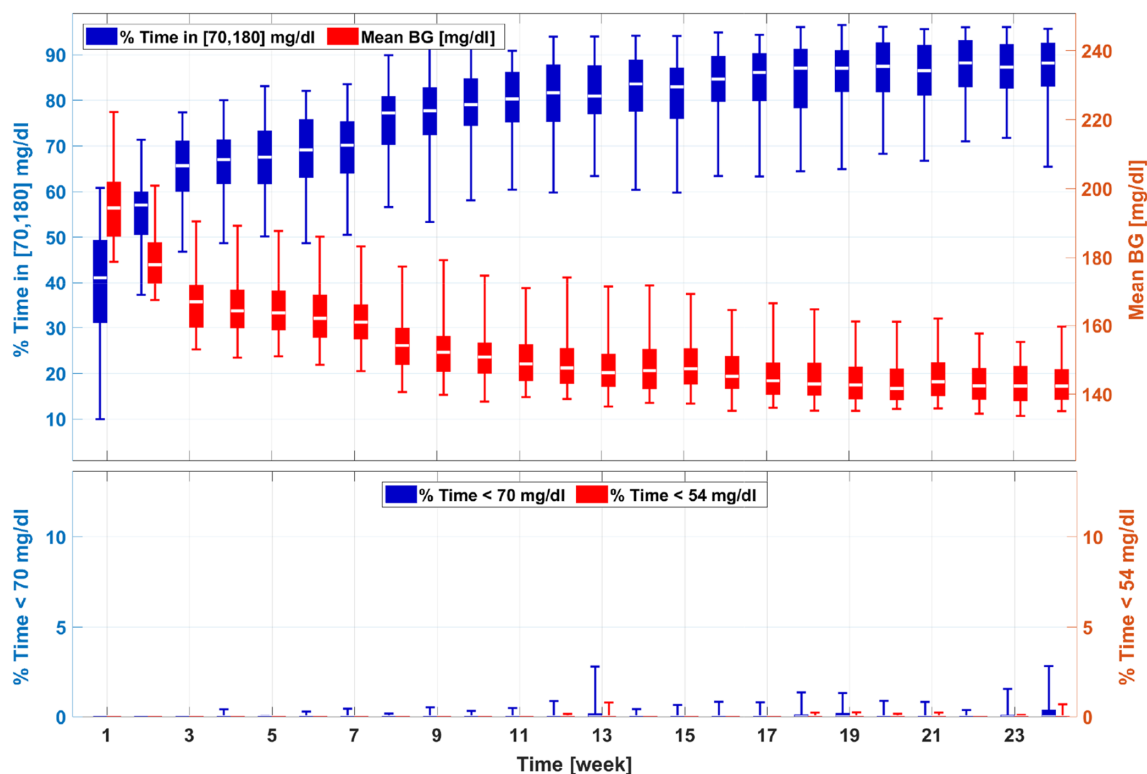


FIGURE 3 Trends of the glycemic management metrics in the adaptation procedures for the 111-patient cohort (Scenario I). A box-and-whisker approach is used to plot the data, where on each box, the central white line is the median, the edges of a box denote the 25% and 75% percentiles, and the whiskers denote the 5% and 95% percentiles

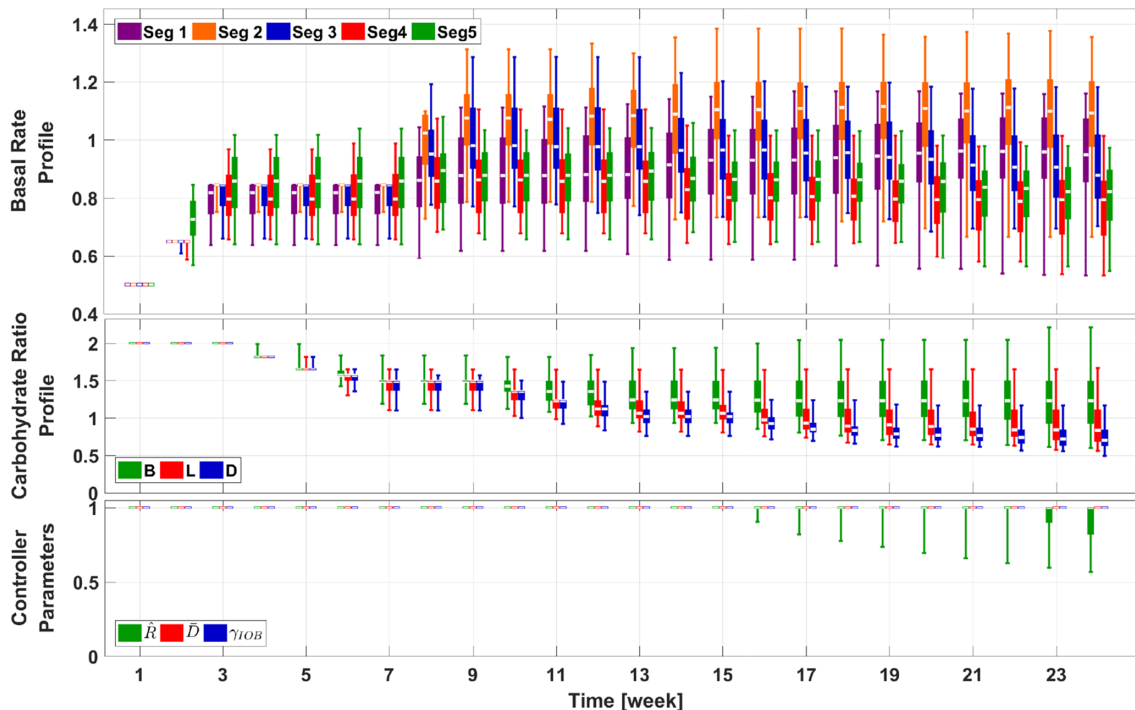


FIGURE 4 Trends of the adaptation parameters in the adaptation procedures for the 111-patient cohort (Scenario I). The same box-and-whisker approach as that in Figure 3 is used to plot the data. For illustration purpose, the relative values of parameters against default values in the UVA/Padova simulator are provided

through the use of statistical IOB constraints and smoothness constraints, strong performance is achieved by the BR profiles with their segments aligning around the reference value (which is 1 in Figure 4),

instead of profiles with extremely large and small neighboring segments. In addition, the CR profile segments are aligned around the reference value; due to the safety constraints, risky (small) values for the

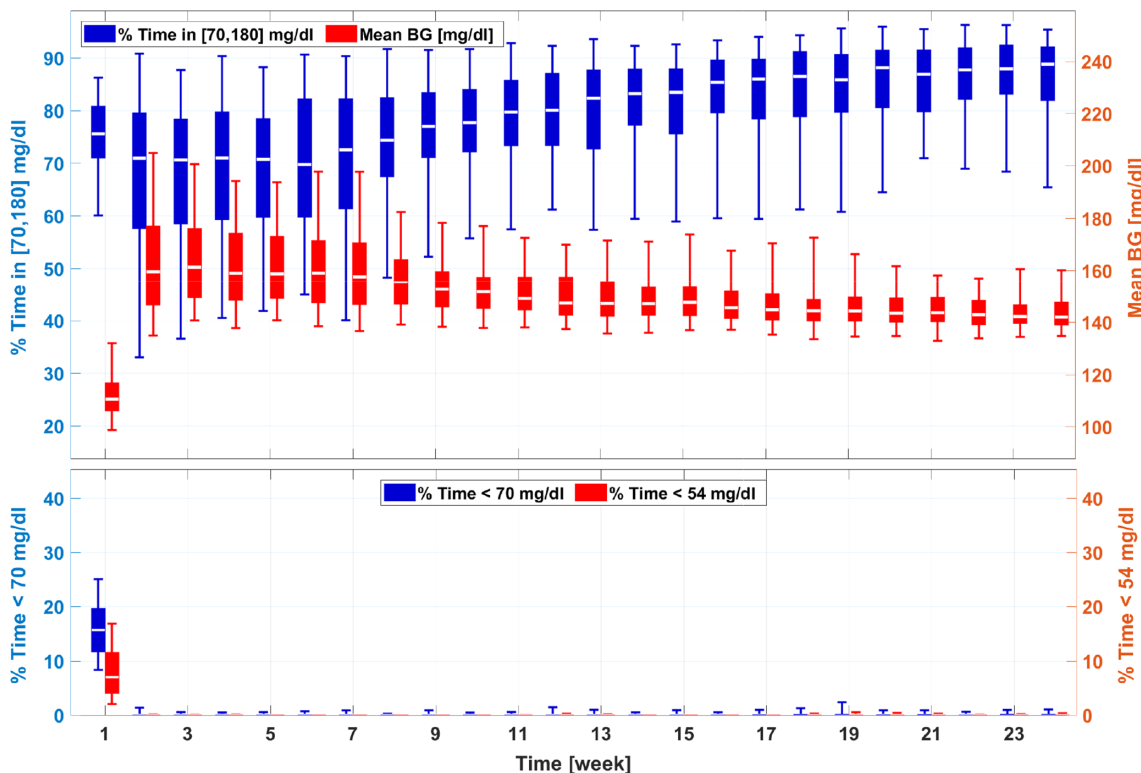


FIGURE 5 Trends of the glycemic management metrics in the adaptation procedures for the 111-patient cohort (Scenario II). The same box-and-whisker approach as that in Figure 3 is used to plot the data




FIGURE 6 Trends of the adaptation parameters in the adaptation procedures for the 111-patient cohort (Scenario II). The same box-and-whisker approach as that in Figure 3 is used to plot the data. Keys are the same as those in Figure 4

CR segments are avoided. In addition, only moderate changes are observed for parameters in the feedback controller, which is expected as the obtained feedforward control parameters (BR and CR profiles) have set up an optimized operating point for the feedback controller,

and thus the default values can achieve satisfactory glucose management. It is interesting to note that the second segment of the BR profile is larger than the other elements, which helps counteract against diurnal insulin sensitivity changes and dawn phenomenon.

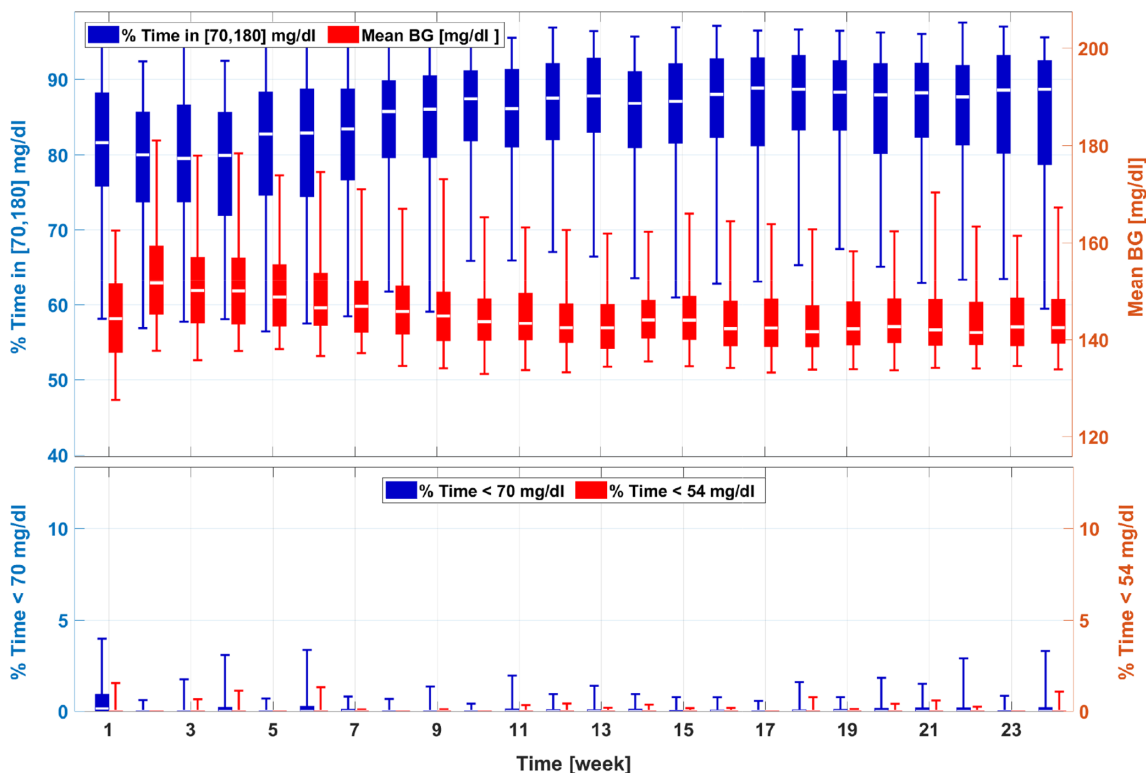


FIGURE 7 Trends of the glycemic management metrics in the adaptation procedures for the 111-patient cohort (Scenario III). The same box-and-whisker approach as that in Figure 3 is used to plot the data




FIGURE 8 Trends of the adaptation parameters in the adaptation procedures for the 111-patient cohort (Scenario III). The same box-and-whisker approach as that in Figure 3 is used to plot the data. Keys are the same as those in Figure 4

3.2 | Results for Scenario II

The results for Scenario II are provided in Figure 5 and 6. Due to the increased BRs, severe hypoglycemia is observed on Week 1 of the simulation. The controller-led BR adaptation algorithm, however, is able to recognize safe BR profiles for the in silico cohort from Week 2 and perform safe fine tunes for the rest of the adaptation procedure. The CR adaptation procedure works properly as well, which manages to decrease the segments in the CR profiles for improved glycemic metrics. By eliminating risks of hypoglycemia (average percent time below 70 mg/dl, from 15.7% [Week 1] to 0.0% [Week 24], $p < .001$), the adaptation algorithm improves average percent time in [70, 180] mg/dl (from 75.6% [Week 1] to 88.8% [Week 24], $p < .001$). The obtained patterns of the BR profile and CR profile are similar to that of Scenario I.

3.3 | Results for Scenario III

This scenario is devoted to simulate a 6-month study of AP parameter adaptation, for patients with mixed and reasonably over- or underestimated segments in their BR and CR profiles. The results are provided in Figures 7 and 8. As expected, with slightly incorrect BR and CR profiles, the in silico patients generally have acceptable glucose management performance, which challenges the proposed adaptation method on the ability of recognizing and adjusting small but realistic mismatches in the adaptation parameters. From both figures, we observe that the algorithm is able to discern these parameter mismatches and achieve improved average percentage time in [70, 180] mg/dl (from 81.6% on Week 1 to 88.7% on Week 24), $p < .001$, without causing risk of hypoglycemia. It appears that average percentage time below 70 mg/dl is decreased as well, but the significance is not sufficiently small (from 0.1% to 0.0%, $p = .077$). Similar patterns in the BR and CR profiles to that obtained in the first two scenarios are realized.

3.4 | Intra-patient changes

To further illustrate the effect of the adaptation process, glucose and insulin profiles of a particular patient simulated using the adaptation parameters obtained on Weeks 1, 8, 16, and 24 are provided in Figure 9 (Scenario I), Figure 10 (Scenario II), and Figure 11 (Scenario III), in which a 24-hr protocol composed of a 50 g CHO breakfast at 08:00, a 75 g CHO lunch at 12:00, and a 75 g CHO dinner at 19:00 with the sensor noise seed is used. The results obtained are consistent with the population-level analysis in this section.

4 | DISCUSSION

Several challenges exist in approaching the AP parameter adaptation problem. First, the effect of different AP parameters on glucose management is usually strongly coupled. For instance, the size of a real-time insulin microbolus is jointly determined by the BR profile and the closed-loop feedback controller, while BR, feedback controller, and user-requested meal boluses can all contribute to meal-related insulin delivery. This causes loss of identifiability to determine the correct root cause of poor glucose management; adjusting an incorrect parameter, however, may cause disastrous consequences. Second, lifestyle disturbances can also mask issues in glucose management. For instance, if a patient had a low BR for some time period while he/she happened to be jogging (which would increase insulin sensitivity) at the same time, the glucose traces might still appear to be acceptable, which would prevent an adaptation algorithm from identifying the real problems in glucose regulation. Finally, a considerable number of parameters in AP need to be considered as candidate variables in the adaptation procedure (which is 12 in our in silico tests in this work), while from a user experience perspective, the time of adaptation is limited (e.g., 24 weeks). As an iteration of the adaptation

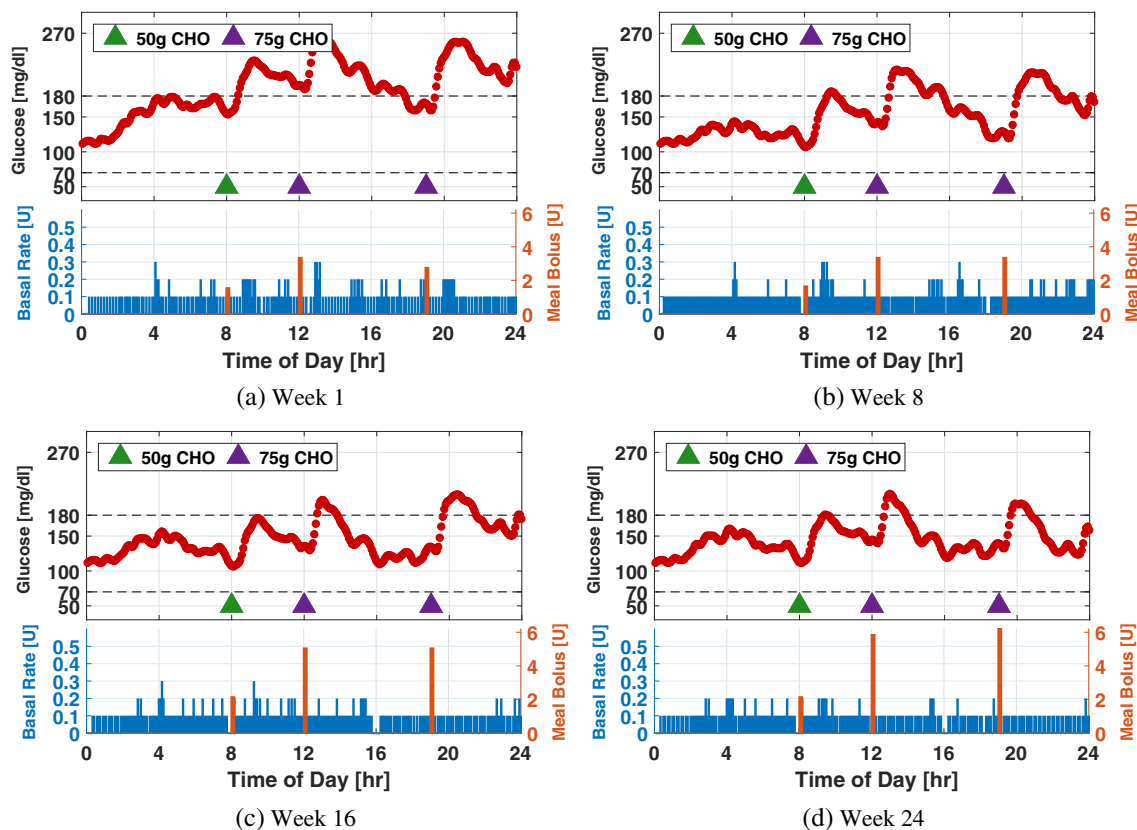


FIGURE 9 Twenty-four-hour glucose and insulin profiles simulated using the system parameters obtained on Weeks 1, 8, 16, and 24 of Scenario I for a particular patient. For comparison purpose, the same meal protocol and measurement noise sequence are applied to generate the data in the four subplots. Green and purple triangles denote meals of 50 g and 75 g CHO, respectively

procedure may take 1 week to average out lifestyle disturbances, this basically implies that the task of AP adaptation through adjusting 12 coupled parameters need to be achieved within 24 iterations.

To overcome these challenges, a data-driven multivariate learning framework is developed in this work, on the basis of a dual-layer control scheme for long-term home use of AP. In this scheme, feedback/feedforward control algorithms operate in the lower layer to achieve real-time glucose regulation, while the adaptation algorithm is implemented in the upper layer based on data from the lower layer. The proposed adaptation procedure is composed of two phases. The first phase focuses on adaptation of feedforward control parameters, including BR and CR profiles. To bypass the nonidentifiability issue, a controller-led approach is proposed for BR adaptation; the key idea is to exploit the intelligence from lower-layer feedback control algorithm to achieve autonomous decision of BR profiles. IOB and smoothness constraints are proposed to avoid hypoglycemia risks in the controller-led decision procedure. The CR profile is adapted through optimizing the postprandial BG levels toward a user-specified target zone while considering dynamically updated data-driven safety constraints. A hybrid time- and event-triggered iterating procedure is proposed to achieve joint adaptation of BR and CR profiles. The second phase is devoted to adapting feedback control behavior. To improve the efficiency of adaptation, a sensitivity analysis is proposed through performing advisory mode comparison based on historical glucose data,²² so that the bottleneck control parameter can be selected for adaptation. The selected

parameter is then updated by optimizing average glucose levels while restricting risks of hypoglycemia.

For optimization problems encountered in CR profile and feedback control adaptation, analytical relationships between optimization parameters and performance metrics are generally not known, which precludes the use of standard optimization methods.²⁵ To overcome this additional challenge, the optimization problems are solved utilizing a BO approach, which was developed to optimize unknown objective functions based on noisy measurements in the machine learning community,²³ and has been recently used in on-line dynamics learning, automatic controller tuning and nonlinear adaptive control.^{26–29} In our work, we integrate BO with safety requirements and clinical experience, so that the “black-box” glucose regulation process can be safely adapted and improved.

The proposed adaptation method is evaluated on the basis of the 111-patient cohort of the U.S. FDA accepted UVA/Padova simulator¹⁶ for three in silico scenarios. The first two scenarios focus on edge cases with (a) underestimated BR and overestimated CR profiles and (b) overestimated BR and CR profiles, while the third scenario considers a real-life situation that different segments in the BR and CR profiles can be either overestimated or underestimated within a moderate extent. To challenge the proposed algorithm, lifestyle disturbances caused by timing and size changes for meals, inter- and intra-day insulin sensitivity changes, and occasionally occurring sickness that can significantly offset insulin sensitivity. We show that for all scenarios, the proposed method is able to correctly identify and

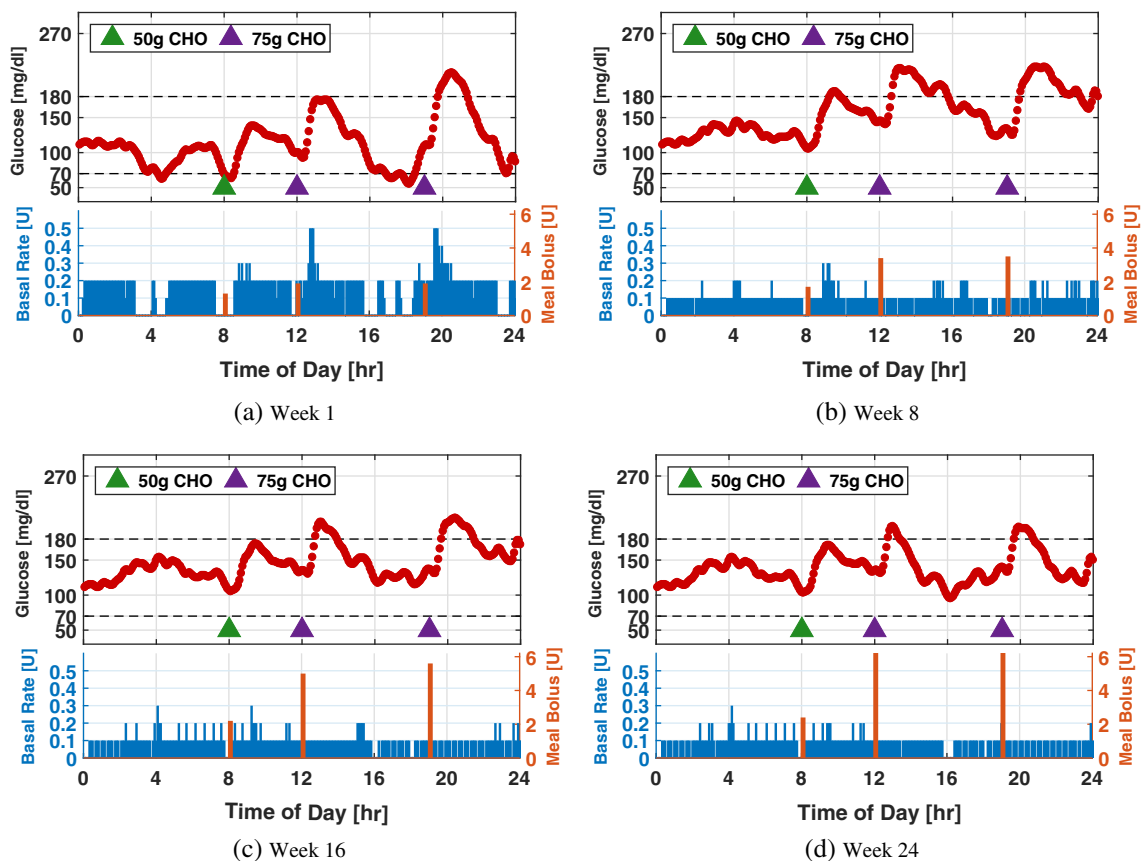


FIGURE 10 Twenty-four-hour glucose and insulin profiles simulated using the system parameters obtained on Weeks 1, 8, 16, and 24 of Scenario II for a particular patient. Meal protocol and sensor noise sequence are the same as those in Figure 9

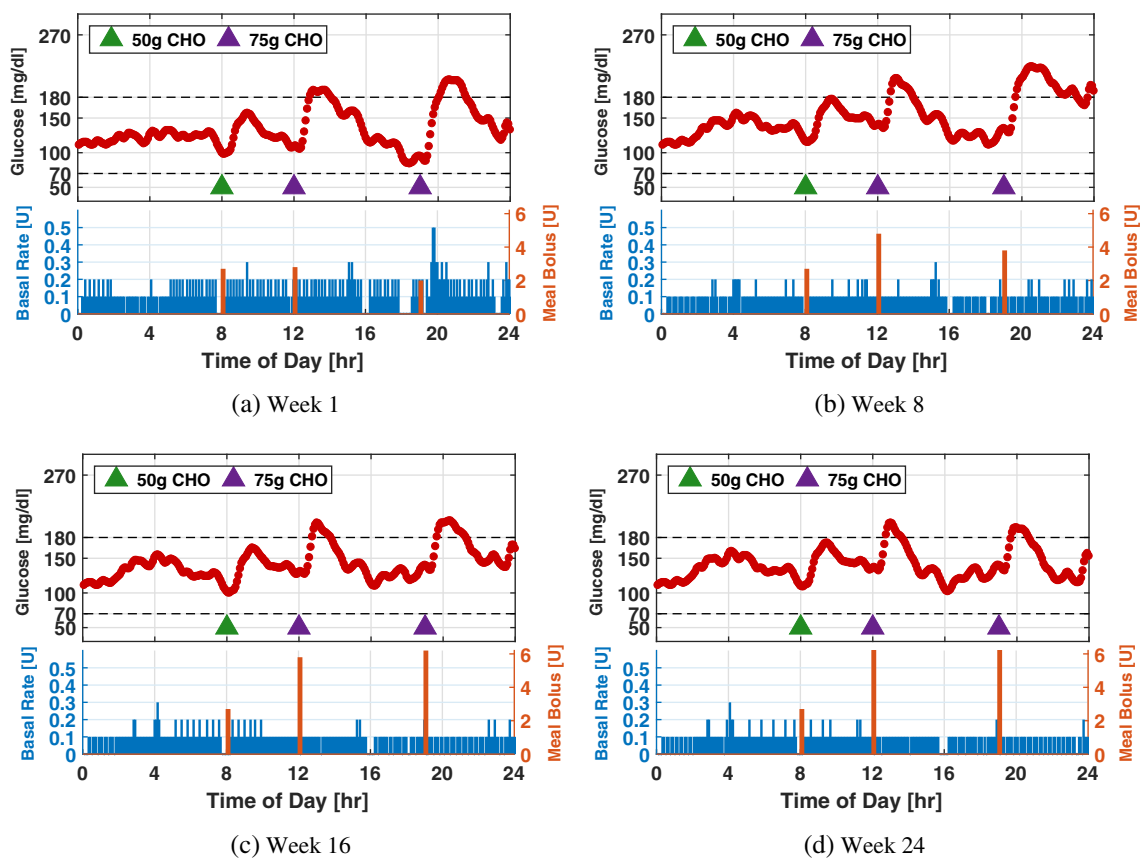


FIGURE 11 Twenty-four-hour glucose and insulin profiles simulated using the system parameters obtained on Weeks 1, 8, 16, and 24 of Scenario III for a particular patient. Meal protocol and sensor noise sequence are the same as those in Figure 9

adaptively adjust the inappropriate parameters to achieve improved and satisfactory glucose regulation, without causing risks of hypoglycemia throughout the adaptation procedure.

5 | CONCLUSION

In this article, a data-driven multivariate learning approach has been proposed for long-term parameter adaptation of an AP, on the basis of a dual-layer control scheme. Through a two-phase adaptation procedure, the algorithm gradually adjusts BR profile, CR profile, and parameters in the closed-loop control algorithm. The adaptation of BR profile was performed by based on the behavior of the closed-loop control algorithm while considering IOB and smoothness constraints, and CR profile and feedback control parameters were adjusted through a BO-assisted approach. It was shown that the algorithm was able to adjust inappropriate parameters for improved glucose regulation, despite of lifestyle disturbances caused by skipped meals, sickness, and inter- and intra-day insulin sensitivity changes.

ACKNOWLEDGMENT

Access to the full version of the UVA/Padova metabolic simulator was provided by an agreement with Prof. C. Cobelli (Università degli Studi di Padova) and Prof. B. P. Kovatchev (UVA) for research purposes.

CONFLICT OF INTERESTS

The authors have no conflicts of interest to declare.

Algorithm 1: BO-assisted parameter adaptation

- 1: **while** termination conditions not satisfied **do**.
- 2: update iteration counter $k := k + 1$, update π_k and query \mathcal{D}_k for data $\bar{\mathcal{D}}_k$ used in adaptation of ψ ;
- 3: update S_k based on y_{k-1} ;
- 4: **if** $|\pi_k| \leq n_{\text{BO}}$ **then**.
 - $\psi_k = (1 + \delta_s S_k) \psi_{k-1}$;
- 5: **else**.
- 6: update θ_k through solving.

$$\min_{\theta} \sum_{k \in \pi_k} (y_k - \hat{f}(\psi_k, \theta | \mathcal{D}_k))^2;$$
- 7: obtain ψ_k by solving.

$$\min_{\psi_k} \hat{f}(\psi_k, \theta_k | \mathcal{D}_k)$$

$$\text{s.t., } |\hat{f}(\psi_k, \theta_k | \mathcal{D}_k) - y_{k-1}| \leq \delta;$$
- 8: **if** $\text{sign}(\psi_k - \psi_{k-1})$ is inconsistent with S_k **then**.

$$\psi_k = (1 + \delta_s S_k) \psi_{k-1};$$
- 9: **end if**.
- 10: update ψ_k through safe checks.

$$\hat{g}(\psi_k | \mathcal{D}_k) \leq 0;$$
- 11: **end if**.
- 12: implement ψ_k to obtain y_k , and update data set $\mathcal{D}_{k+1} := \mathcal{D}_k \cup \{\Delta_k\}$ with $\Delta_k := \{y_k, \psi_k\}$;
- 13: **end while**.

REFERENCES

1. Kovatchev BP, Renard E, Cobelli C, et al. Safety of outpatient closed-loop control: first randomized crossover trials of a wearable artificial pancreas. *Diabetes Care*. 2014;37:1789-1796.
2. Thabit H, Tauschmann M, Allen JM, et al. Home use of an artificial beta cell in type 1 diabetes. *N Engl J Med*. 2015;373:2129-2140.
3. Bergenstal RM, Garg SK, Weinzimer SA, et al. Safety of a hybrid closed-loop insulin delivery system in patients with type 1 diabetes. *JAMA*. 2016;316:1407-1408.
4. Dassau E, Pinsky JE, Kudva YC, et al. Twelve week 24/7 ambulatory artificial pancreas with weekly adaptation of insulin delivery settings: effect on hemoglobin A_{1c} and hypoglycemia. *Diabetes Care*. 2017;40:1719-1726.
5. Horton ES. Exercise and physical training: effects on insulin sensitivity and glucose metabolism. *Diabetes Metab Rev*. 1986;2:1-17.
6. Goodyear LJ, Kahn BB. Exercise, glucose transport, and insulin sensitivity. *Annu Rev Med*. 1998;49:235-261.
7. Dassau E, Brown SA, Basu A, et al. Adjustment of open-loop settings to improve closed-loop results in type 1 diabetes: a multicenter randomized trial. *J Clin Endocrinol Metab*. 2015;100:3878-3886.
8. Kovatchev B, Cheng P, Anderson SM, et al. Feasibility of long-term closed-loop control: a multicenter 6-month trial of 24/7 automated insulin delivery. *Diabetes Technol Ther*. 2017;19:18-24.
9. Garg SK, Weinzimer SA, Tamborlane WV, et al. Glucose outcomes with the in-home use of a hybrid closed-loop insulin delivery system in adolescents and adults with type 1 diabetes. *Diabetes Technol Ther*. 2017;19:155-163.
10. Palerm CC, Zisser H, Jovanović L, Doyle FJ III. A run-to-run framework for prandial insulin dosing: handling real-life uncertainty. *Int J Robust Nonlin Control*. 2007;17:1194-1213.
11. Palerm CC, Zisser H, Jovanović L, Doyle FJ III. A run-to-run control strategy to adjust basal insulin infusion rates in type 1 diabetes. *J Process Control*. 2008;18:258-265.
12. Toffanin C, Visentin R, Messori M, Palma FD, Magni L, Cobelli C. Toward a run-to-run adaptive artificial pancreas: in silico results. *IEEE Trans Biomed Eng*. 2018;65:479-488.
13. Wang Y, Gao F, Doyle FJ III. Survey on iterative learning control, repetitive control, and run-to-run control. *J Process Control*. 2009;19:1589-1600.
14. Wang Y, Dassau E, Doyle FJ III. Closed-loop control of artificial pancreatic β -cell in type 1 diabetes mellitus using model predictive iterative learning control. *IEEE Trans Biomed Eng*. 2010;57:211-219.
15. Wang Y, Zhang J, Zeng F, et al. "Learning" can improve the blood glucose control performance for type 1 diabetes mellitus. *Diabetes Technol Ther*. 2017;19:41-48.
16. Dalla Man C, Micheletto F, Lv D, Breton M, Kovatchev B, Cobelli C. The UVA/PADOVA type 1 diabetes simulator: new features. *J Diabetes Sci Technol*. 2014;8:26-34.
17. Messori M, Kropff J, Del Favero S, et al. Individually adaptive artificial pancreas in subjects with type 1 diabetes: a one-month proof-of-concept trial in free-living conditions. *Diabetes Technol Ther*. 2017;19:560-571.
18. Messori M, Paolo Incremona G, Cobelli C, Magni L. Individualized model predictive control for the artificial pancreas: in silico evaluation of closed-loop glucose control. *IEEE Control Syst Mag*. 2018;38:86-104.
19. Cinar A. Artificial pancreas systems: an introduction to the special issue. *IEEE Control Syst*. 2018;38:26-29.
20. Gondhalekar R, Dassau E, Doyle FJ III. Periodic zone-MPC with asymmetric costs for outpatient-ready safety of an artificial pancreas to treat type 1 diabetes. *Automatica*. 2016;71:237-246.
21. Ljung L. *System Identification: Theory for the User*. EngleWood Cliffs, NJ: Prentice Hall PTR; 1999.
22. Gillis R, Palerm CC, Zisser H, Jovanovic L, Seborg DE, Doyle FJ III. Glucose estimation and prediction through meal responses using ambulatory subject data for advisory mode model predictive control. *J Diabetes Sci Technol*. 2007;1:825-833.
23. Shahriari B, Swersky K, Wang Z, Adams RP, Freitas N. Taking the human out of the loop: a review of Bayesian optimization. *Proc IEEE*. 2016;104:148-175.

24. Visentin R, Campos-Náñez E, Schiavon M, et al. The UVA/Padova type 1 diabetes simulator goes from single meal to single day. *J Diabetes Sci Technol*. 2018;12:273-281.
25. Boyd S, Vandenberghe L. *Convex Optimization*. Cambridge, UK: Cambridge University Press; 2008.
26. Trimpe S, Millane A, Doessegger S, D'Andrea R. A self-tuning LQR approach demonstrated on an inverted pendulum. In: Proceedings of the 19th IFAC World Congress, Cape Town, South Africa, August 24-29. 2014;47:11281-11287.
27. Benosman M, Farahmand A.. Bayesian optimization-based modular indirect adaptive control for a class of nonlinear systems. In: 12th IFAC Workshop on Adaptation and Learning in Control and Signal Processing (ALCOSP 2016), Eindhoven, The Netherlands, June 29 - July 01. 2016;49:253-258.
28. Duivendoorn R, Berkenkamp F, Carion N, Krause A, Schoellig AP.. Constrained Bayesian optimization with particle swarms for safe adaptive controller tuning. In: 20th IFAC World Congress, Toulouse, France, July 9-14. 2017;50:11800-11807.
29. Bansal S, Calandra R, Xiao T, Levine S, Tomlin CJ. Goal-driven dynamics learning via Bayesian optimization. In: IEEE 56th Annual Conference on Decision and Control (CDC), Melbourne, Australia, December 12-15. 2017:5168-5173.

How to cite this article: Shi D, Dassau E, Doyle III FJ. Multivariate learning framework for long-term adaptation in the artificial pancreas. *Bioengineering & Translational Medicine*. 2019;4:61-74. <https://doi.org/10.1002/btm2.10119>