A Systematic Stochastic Design Strategy Achieving an Optimal Tradeoff Between Peak BGL and Probability of Hypoglycaemic Events for Individuals Having Type 1 Diabetes Mellitus

Graham C. Goodwin, Maria M. Seron*, Adrian M. Medioli

Priority Research Centre for Complex Dynamic Systems and Control School of Electrical Engineering and Computing, The University of Newcastle, Australia.

Tenele Smith, Bruce R. King, Carmel E. Smart

Department of Paediatric Endocrinology and Diabetes, John Hunter Children's Hospital, Newcastle, Australia.

Abstract

This paper has two key contributions. The first contribution is a systematic procedure for fitting an envelope of models which captures a range of possible blood glucose level (BGL) responses for a particular individual having Type 1 diabetes. An important aspect of the procedure is that it requires minimal testing on the individual. Moreover, the testing can be carried out by the individual at home. The developed envelope of models, termed 'Metabolic Digital Twin Envelope' (MDTE) takes into account the quantification of possible errors including those arising from utilising a simplified model (commonly called "bias" errors) and those arising from unmodelled disturbances and noise (commonly called "variance" errors). The second, and most important, contribution is a methodology that allows convex optimisation to be used to develop an insulin injection policy which minimises mean square peak BGL whilst ensuring that there is a strict lower bound on the probability of hyperglycaemic events. The optimisation methodology is posed as a stochastic design strategy based on using the probabilistic models for each individual afforded by the MDTE.

Keywords: medical control systems, Type 1 diabetes mellitus, stochastic strategies for diabetes management, blood glucose regulation, insulin bolusing, diabetes modelling and estimation, system identification, stochastic embedding

1. Introduction

Stochastic optimisation incorporating chance constraints has been widely discussed in many engineering areas and, specifically, in recent control theory literature, see e.g., [1, 2, 3]. The importance of this idea has also been recognised in the management of Type 1 Diabetes Mellitus (T1DM). For example, work reported in [4, 5, 6, 7, 8, 9] proposed a stochastic modelling framework and a methodology for stochastic targeted glycemic control for critically ill patients. It is a natural choice in the context of T1DM management since targets and constraints are often stated in a probabilistic setting (e.g., percentage of time in normoglycaemic or hypoglycaemic

*Corresponding author

range, etc.) A key contribution of the current paper is to develop a systematic design procedure which trades off average peak BGL versus the probability of having a hypoglycaemic event.

We use the notion of 'digital twins' to formalise and answer this problem. In this context, a digital twin associated with a physical system is a mathematical model that represents the system. We note that the idea of digital twins is ubiquitous in engineering. For example all engineering designs utilise computer models that represent the object to be designed. Models appear in essentially every area, e.g., aircraft design, automobile engine design, bridge design, etc. The idea of digital twins is also gaining traction in medicine [10], [11].

Our use of the term *"metabolic digital twin"* is quite specific. Our goal is to utilise the notion of a digital "twin" to corroborate and extend insights arising from contemporary research in the area of T1DM management.

We pay particular attention to the impact of potential errors in the digital twins. In particular, we describe an envelope of possible twins, which we term the Metabolic Digital Twin Envelope (MDTE). We view these sets of digital twins as allowing one to predict a range of behaviour for an individual.

There exists a vast literature on the development of predictive models for blood glucose response. For example, a recent paper [12] surveys 140 articles related to personalised blood glucose prediction strategies. Many of these strategies have been developed to aid the design of a, so-called, "artificial pancreas." The goal of an artificial pancreas is to automate the delivery of insulin in response to measured variables, in particular the measured blood glucose response. The models used for this purpose are usually "control-oriented" models with a small number of adjustable parameters. Examples of models utilised for the design of artificial pancreas strategies are described in [13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25].

A related research area focuses on models that describe the behaviour of a physiologically similar cohort of individuals. These models are typically more complex and are useful to simulate groups of persons with diabetes using, so-called, "in-silico" trials. A well-known example of a model of this type is the UVA-Padova simulator [26, 27, 28]. Other models that have been used for in-silico trials were presented in [29, 30, 31, 32, 33, 12].

It is acknowledged in the research referred to above that real models for individuals with diabetes exhibit a high degree of uncertainty. This uncertainty is taken into account when evaluating different strategies. However, this is after the control strategy has been specified. We propose here a different approach in which the model uncertainty is explicitly accounted for at the design stage. Indeed, of importance in the description of the MDTE for an individual is that the formalism emphasises that a multitude of output events are typically observed for an individual in clinical practice under seemingly similar conditions. To address this problem, we provide a statistical description of an envelope of possible behaviours. This is in the spirit of work on stochastic embedding described in [34, 35, 36]. The availability of this statistical description allows one to answer more sophisticated questions than could be answered using deterministic models. For example, given the MDTE for an individual one can ask, "what is the best management policy for a particular meal subject to the constraint that the probability of experiencing a hypoglycaemic event is less than, say, 15%?" Such questions are impossible without the availability of a statistical description of the associated uncertainty.

Our work is aimed broadly at personalised diabetes management. This includes, but is not limited to, the development of closed-loop strategies. Recent examples of target management strategies by our medical team include [37] and [38], where different open-loop insulin dosing strategies are investigated for specific types of meals. One of the goals of the current paper is to

complement and give insights into the results of these clinical trials. In parallel work described in [39], a personalised open-loop technique to optimise insulin bolusing for meals having high fat and protein content was presented. The model used in [39], however, does not account for disturbances and uncertainties and hence does not allow for the design of insulin strategies that account for probabilistic outcomes, as we propose here.

Other works that have proposed models of the general type presented here are [31, 27, 40]. These works use Bayesian inference and Monte Carlo techniques to derive posterior distributions of model parameters given data collected for a particular individual. The procedure yields models that have been referred to as virtual 'clones' [27] or 'stochastic virtual subjects' [31]. In [27] the background model is the model behind the UVA-Padova simulator, whereas the model used in [31] is the Willinska-Hovorka model [29]. In that regard they offer valuable improved versions of the virtual populations coded in their respective simulators. An important difference with the work reported here is that [31, 27] base their results on plasma glucose and plasma insulin measurements. By way of contrast, the proposed MDTE used here is derived from interstitial blood glucose sensor data and insulin pump information. Hence the models can be estimated from minimally invasive free-living tests and used to directly compute insulin dosages for use with pumps or injections. Also, the proposed envelopes are based on simple linear models and uncertainty descriptions and have a minimal number of adjustable parameters. A model with similar characteristics was proposed in [41], and was used to identify individualised models to data collected under free-living conditions. The model has an error component that is modelled as filtered white noise. The individualised model was shown to predict hypoglycaemic events more accurately than a single average model for the whole population. Stochastic models of limited complexity, such as our proposed MDTE framework and the individualised model of [41], are well suited to the design of optimal management algorithms that specifically consider uncertainty in the formulation. In this regard, a key contribution of the current paper is a methodology to design optimal probabilistic open-loop insulin dosing strategies. The methodology incorporates a mechanism to adjust the probability that the predicted BGL will violate a 'safety' constraint imposed on hypoglycaemic events. The strategy can thus be used to investigate optimal tradeoffs between performance and robustness of policies for a particular individual.

In summary there are two main contributions in the paper. The first contribution is a systematic procedure for fitting an envelope of models which captures a range of possible BGL responses for a particular individual having Type 1 diabetes. This contribution is explained in detail in Sections 3 to 5. The second contribution is a methodology that employs convex optimisation to develop an insulin injection policy which minimises mean square peak BGL whilst ensuring that there is a strict lower bound on the probability of hypoglycaemic events. This contribution is explained in detail in Sections 6 to 10.

2. Phases of the Procedure

The main steps involved in developing and using an MDTE for an individual are shown schematically in Figure 1. The procedure shown in the figure is organised into 6 phases which are briefly described below. (More details are presented in subsequent sections of the paper.)

Phase 1 (Individualisation): The process begins with a specific individual.

Phase 2 (Trial design and data collection): A suitable set of clinical trials is planned based on minimal disruption to the individual. Data is collected from trials conducted at home by the individual.

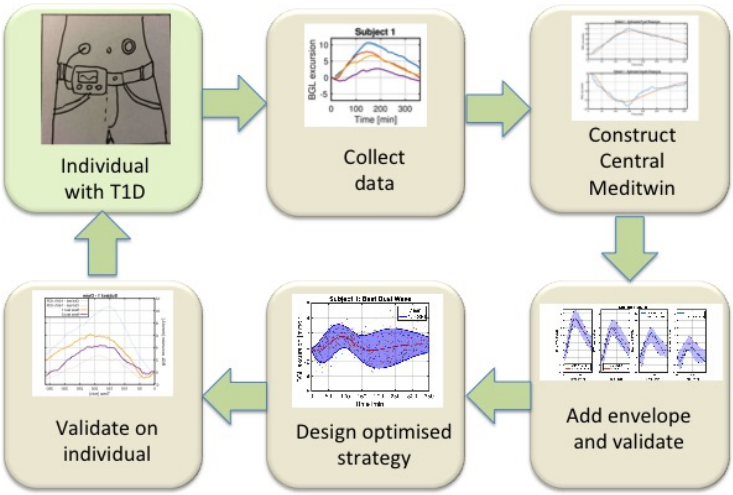

Figure 1: The main stages of fitting and utilising a MDTE.

Phase 3 (Estimation): A "central" Metabolic Digital Twin (Meditwin) is estimated for the individual.
Phase 4 (Envelope quantification and validation): The impact of disturbances is quantified and added to the model. The envelope responses are then compared with multiple data records.
Phase 5 (Optimise strategies): Designs of "optimal" management strategies are performed for the individual using the MDTE leading to suggested improved management strategies.
Phase 6 (Validate strategies): The suggested designs are tested and validated on the real individual.

This paper reports results of clinical trials on 12 individuals. Phase 6 was restricted to 2 of the original individuals.

3. Individualisation, Trial Design and Data Collection

This section describes phases 1 and 2 shown in Figure 1.

3.1. Individualisation

One or more individuals are chosen for which it is desired to build a personalised MDTE and to formulate a diabetes management strategy.

3.2. Trial Design

Data is collected at home by the individual. To maximise the value of data, planned experiments are used to focus on the behaviour of interest [42]. The testing needs to be minimally intrusive to the individual yet reveal as much information as possible. It is thus important that the tests span the behaviour of interest and allow the impact of different inputs to be separated. For example, to identify the response to a single meal and insulin, one needs at least two tests with different meal size or insulin magnitude. Alternatively, if one wishes to separate the impact of 3 meal macronutrients (e.g., carbohydrates, fat, protein), two types of exercise (aerobic, anaerobic)

and insulin, then 6 separate tests would be needed having different meal composition, exercise patterns and insulin profiles.

Caution is needed in carefully selecting the pattern of tests. One reason for this is that the human body is a complex mechanism. This means that a single model that describes all possible scenarios would need to be very complex. Indeed, it might reasonably be argued that such models are unrealistic. Even if one were optimistic about the existence of such a "universal" model then the model would contain a large number of adjustable parameters making it extremely difficult to fit to an individual. On the other hand, restricted complexity models will contain "bias errors".

"Bias" errors [42, 43, 44, 45, 46, 47] are particularly problematic since they are operating condition dependent. However, the impact of such errors can be mitigated by ensuring that the conditions under which models are fitted are similar to the conditions under which the model will be ultimately used.

In the case of the MDTEs, tests comprising a single meal consumed over a short period of time combined with a dual wave insulin pattern comprising an initial pulse of insulin (bolus) followed by a constant flow over an extended period (extended bolus) will be used. This choices is motivated by two facts:

1. Meals are typically consumed over relatively short periods of time, and
2. clinical experience suggests that a good insulin strategy for most meal types comprises either a single bolus or a single bolus followed by some other extended insulin pattern.

Remark 3.1. *Note that "bias error" distribution is not usually addressed in the statistics literature which, almost exclusively, assumes that the true system lies in the chosen model set [43]. There does exist some discussion of the question of bias error distribution in recent control engineering literature—see for example [35, 36, 45, 47].*

3.3. Data Collection

Consistent with the discussion in Section 3.2, we use data obtained during a recent medical trial conducted at home and aimed at evaluating the efficacy of an open-loop 'dual wave' insulin strategy, i.e., the use of a standard bolus followed by an 'extended bolus' (square wave), for a high-fat, high-protein meal. The target meal contained 30 g of carbohydrates (CHO), 40 g of fat and 50 g of protein. A summary of the trial protocol is as follows: In the two weeks prior to study commencement participants were contacted every second day to review BGLs and adjust basal insulin rates and insulin-to-carbohydrate ratio (ICR) reference to meet pre-meal and post-meal BGL targets of 4 to 8 mmol/L. Test meals were given at breakfast and participants were restricted from doing strong exercise during the postprandial monitoring period. No restrictions were placed on exercise the day prior. Participants were asked to fast overnight and were excluded from the test day if they had an episode of hypoglycaemia (capillary BGL < 3.5 mmol/L or symptomatic, resulting in food on board) or administered a correction insulin dose after 3 am (resulting in extra insulin on board). Participants were required to meet a pre-meal blood glucose target of 4 to 10.9 mmol/L and consume the test meal in 20 minutes. The data was collected by individuals at home.[1] Twelve individuals from this trial were studied, identified as Subjects 1, 2, 4, 5, 6, 7, 9, 10, 11, 12, 13 and 16. (The missing numbers, i.e., 3 and 8 correspond to individuals for whom the trials had to be aborted due to equipment malfunction, sickness, etc.) Four tests

[1] This research was approved by the Hunter New England Human Research Ethics Committee, reference number 16/12/14/4.01 of the Hunter New England Local Health District.

were carried out on each individual (with a small difference[2] in the case of Subject 2). These tests were performed on separate days for each individual. Each test comprised a dual wave with a 60:40 split of the insulin dose between the standard bolus and the extended bolus. Insulin delivery was commenced 15 min before food ingestion. The extended bolus was applied over 180 minutes following the standard bolus. Each test utilised a different quantity of insulin as a function of the individual's ICR. The quantities used in the tests were:

Test 1: the amount suggested by the person's ICR for the particular meal's CHO content (the 'ICR dose'),

Test 2: 20% more than the ICR dose,

Test 3: 40% more than the ICR dose and

Test 4: 60% more than the ICR dose.

The data for 47 trials (4 tests on 11 individuals and 3 tests on Subject 2) is shown in Figure 2. Time zero in this figure corresponds to the time meal ingestion was commenced. The time of application of the bolus is 15 min earlier.

We note that the BGL unit used in this paper is mmol/L and the insulin unit is a value, V, relative to the individual's ICR. (For example, 0.7V is 70% of the individual's ICR dose multiplied by the CHO content of the meal.)

4. MDTE Estimation

Here we describe phases 3 and 4 of the procedure shown in Figure 1.

We chose a simple linear model having a small number of parameters to balance the requirement of excessive testing to fit many parameters in a complex model versus inconvenience to the individual. We note that other works, e.g., [32, 33], have also adopted linear models. In this context, we stress the point that using a more complex model is not guaranteed to yield better predictability on new data due to the potential for over-fitting when a small amount of data is used to fit the model and when many adjustable parameters are contained in the model [42]. The simple model was chosen as a trade-off between good predictability and the number of paramers to fit with limited data.

The following Subsections 4.1 to 4.3 describe the steps for creating a "Central" Meditwin for a particular individual. Subsection 4.5 explains how this is extended to a set of models as used in the MDTE.

4.1. Pulse Response Fitting

In the sequel we subtract the entry BGL and basal insulin flow so that we deal with deviations around steady state.

We denote the measured BGL data for each test (as a deviation from the entry value) by

$$Y^{\mu} \doteq (y_1^{\mu}, \ldots, y_N^{\mu}), \quad \mu \in \{1, 1.2, 1.4, 1.6\}. \tag{1}$$

[2] The 100% test was not performed on Subject 2, but enough data was collected through the other tests to allow for model fitting.

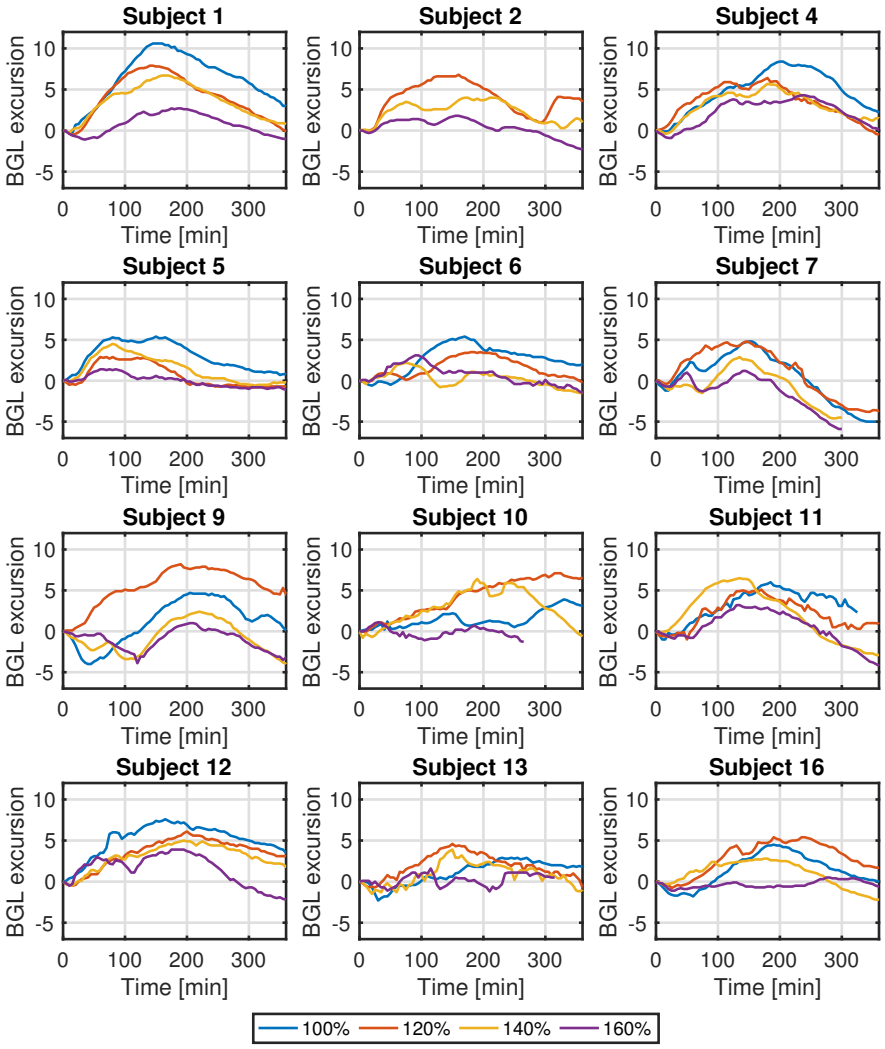

Figure 2: Raw data for all subjects. BGL excursions are relative to the initial value at the start of the test.

In (1), N is the number of samples and the values of μ correspond to the insulin doses used in Test 1 to Test 4 described in Section 3.3. In the linear modelling framework, superposition applies. Thus the BGL deviations can be modelled as

$$Y^{\mu} = R_F + \mu\, R_I, \tag{2}$$

where R_F is the BGL response to the given meal (represented as a pulse) and R_I is the BGL response to the ICR dose administered as the dual-wave input described in Section 3.3. Since we are interested in investigating alternative insulin patterns beyond the dual wave scheme tested in the trials, it is useful to extract the response to a single insulin bolus, i.e., a pulse of insulin. This is done by deconvolving the insulin response, as explained below.

Let $R_I \doteq (r_1^I, \ldots, r_N^I)$ be the response to a dual wave comprising a pulse of size α (the standard bolus) applied at time $t = 0$ and a step of size β (the extended bolus) applied from sample 0 to $M - 1$. Denote the unit pulse response by $H_I \doteq (h_1^I, \ldots, h_N^I)$.

By superposition, we have that the response R_I is given by

$$\begin{bmatrix} r_1^I \\ r_2^I \\ r_3^I \\ \vdots \\ r_M^I \\ r_{M+1}^I \\ \vdots \\ r_N^I \end{bmatrix} = \underbrace{\begin{bmatrix} \alpha+\beta & & & & & & & \\ \beta & \alpha+\beta & & & & & & \\ \beta & \beta & \alpha+\beta & & & & & \\ \vdots & \vdots & \ddots & \ddots & & & & \\ \beta & \beta & \dots & \beta & \alpha+\beta & & & \\ 0 & \beta & \dots & \dots & \beta & \alpha+\beta & & \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \\ 0 & \dots & 0 & \beta & \dots & \dots & \beta & \alpha+\beta \end{bmatrix}}_{\Phi^I} \begin{bmatrix} h_1^I \\ h_2^I \\ h_3^I \\ \vdots \\ h_M^I \\ h_{M+1}^I \\ \vdots \\ h_N^I \end{bmatrix},$$

where Φ^I is a lower-triangular matrix. The above equation can be expressed succinctly as $R_I = \Phi^I H_I$. Combining this with (2) yields

$$Y^\mu = R_F + \mu \Phi^I H_I. \tag{3}$$

The data from the trials for each individual will be considered *"biologically consistent"* if the areas under the positive BGL excursions (i.e., effectively a quantity proportional to the sum of all positive BGL excursions) are 'ordered' so that larger areas correspond to lower amounts of insulin injected. Only biologically consistent data will be used to estimate the MDTE. To describe the deconvolution step, suppose that for a given individual all data is biologically consistent so that all 4 tests can be used for model estimation. As outlined in Section 3.3, our tests corresponded to $\mu_1 = 1$, $\mu_2 = 1.2$, $\mu_3 = 1.4$ and $\mu_4 = 1.6$. The following set of equations then apply:

$$\begin{bmatrix} Y^{\mu_1} \\ Y^{\mu_2} \\ Y^{\mu_3} \\ Y^{\mu_4} \end{bmatrix} = \underbrace{\begin{bmatrix} I_N & \mu_1 \Phi^I \\ I_N & \mu_2 \Phi^I \\ I_N & \mu_3 \Phi^I \\ I_N & \mu_4 \Phi^I \end{bmatrix}}_{\Gamma} \begin{bmatrix} R_F \\ H_I \end{bmatrix} \tag{4}$$

(I_N is the $N \times N$ identity matrix.) The food and insulin pulse responses can be obtained from (4) by multiplying both sides from the left by the pseudoinverse of the matrix Γ as long as there are at least 2 tests that yield consistent data.

To illustrate, we will use Subject 1 from the previously described trial.[3] For this individual all collected data is biologically consistent so the 4 tests were used to estimate the responses. The food and insulin to BGL pulse responses, R_F and H_I, extracted from (4) are shown by blue solid lines in Figure 3.

4.2. Transfer Function Fitting for the Central Meditwin

The next step is to fit transfer function models. Transfer function models represent a dynamic system's frequency response, and have an equivalent representation as a set of differential equations [48]. The latter are commonly used to describe BGL responses, see e.g., the models surveyed in [12].

We fit a 3-real-pole transfer function model to the food and insulin pulse responses, denoted by G_F and G_I, respectively. The food model contains 4 parameters i.e., the gain and the location

[3] The results for other individuals are presented in Section 4.3.

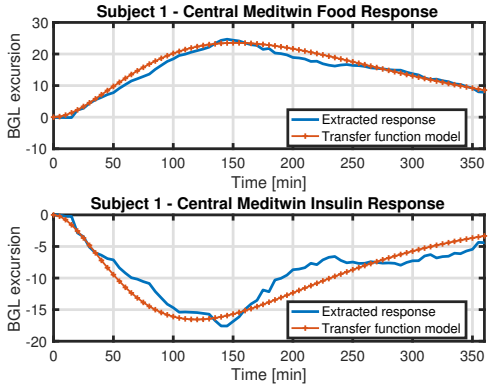

Figure 3: Extracted food and insulin pulse responses and fitted transfer function model responses for the Central Meditwin for Subject 1.

of the 3 real poles:

$$G_F(s) = \frac{K_F}{(a_1 s + 1)(a_2 s + 1)(a_3 s + 1)} \doteq \mathcal{G}_F(\theta_F), \tag{5}$$

where s is the Laplace transform variable, and the parameter vector is $\theta_F \doteq (K_F, a_1, a_2, a_3)$, where K_F [(mmol/L)/(g/min)] is the gain and a_1, a_2 and a_3 [min] are the time constants associated with the food response. The insulin model contains five parameters:

$$G_I(s) = \frac{K_I \, \mathrm{e}^{-s\tau}}{(b_1 s + 1)(b_2 s + 1)(b_3 s + 1)} \doteq \mathcal{G}_I(\theta_I). \tag{6}$$

The parameter vector is $\theta_I \doteq (K_I, b_1, b_2, b_3, \tau)$, where K_I [(mmol/L)/(V/min)] is the gain, b_1, b_2 and b_3 [min] are the time constants associated with the insulin response and τ is a pure delay. The latter is fixed at 15 minutes based on extensive prior clinical experience.

The total BGL response model for the central Meditwin is then

$$y = G_F \, f + G_I \, u, \tag{7}$$

where y [mmol/L] denotes BGL, and u [V/min], f [g/min] are the insulin and food inputs, respectively.

Each transfer function in (5) and (6) was fitted separately using the extracted responses R_F and H_I shown in Figure 3. The following fitting cost function was employed (illustrated here for the fitting of $H_I \doteq (h_1^I, \ldots, h_N^I)$):

$$\theta_I^* = \arg\min_{\theta_I} \Big\{ \sum_{i=1}^{N} [h_i^I - \hat{h}_i(\mathcal{G}_I(\theta_I))]^2 + 50(h_{max} - \hat{h}_{max})^2 + (t_{max} - \hat{t}_{max})^2 \Big\}, \tag{8}$$

where $\hat{h}_i(\mathcal{G}_I(\theta_I))$ is the pulse response of the insulin to BGL transfer function (6), $(\hat{t}_{max}, \hat{h}_{max})$ are the peak time and peak value of this response and (t_{max}, h_{max}) are the peak time and peak value of the response H_I extracted from the data. The weights 50 and 1 in the last two terms were found by trial and error. The motivation behind the above choice of cost function is that, from clinical experience, the peak response time and value are key characteristics in determining the resulting BGL response and obtaining a realistic insulin management strategy. The optimisation (8) was performed using Matlab's 'fmincon' function initialised from different starting points to circumvent local minimum difficulties.

The pulse responses of the food and insulin to BGL transfer functions of the Central Meditwin for Subject 1, fitted as described above, are shown by the solid-'plus' red lines in Figure 3. The Central Meditwin so obtained was next used to predict the responses for each of the tests of the trial. These predictions are shown in Figure 4 for Subject 1 together with the original data.

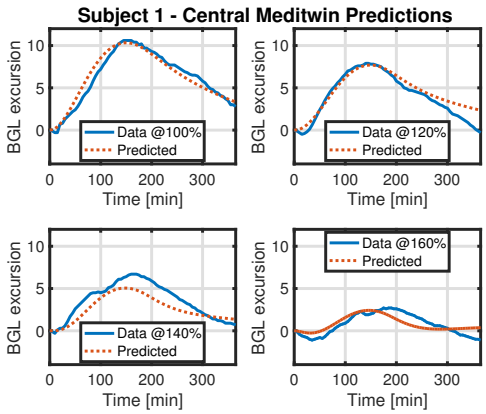

Figure 4: Raw data and Central Meditwin predictions for Subject 1.

4.3. Central Meditwin Fitting for the Full Set of Individuals

The Central Meditwin fitting procedure described in the previous subsections for Subject 1 was repeated for all remaining individuals in the trial using the biologically consistent data sets for those individuals. The results are reported in Appendix A. For each individual, the 2 plots on the left show the estimated food response (upper plot) and the estimated insulin response (lower plot). The responses obtained by the procedure explained in Section 4.1 are shown by solid blue lines. The responses of the transfer function models, fitted as explained in Section 4.2, are shown by solid-'plus' red lines. The 4 subplots on the right of the figures show the predicted responses to the dual-wave inputs used in the trials, together with the raw data collected during the trials.

We observe that some of the extracted responses have large sharp variations. These variations are difficult to explain given the trial conditions where the only recorded 'external' inputs over the trial duration were a food pulse and an insulin dual wave at the start of the test. Thus there are clearly other factors (e.g., other disturbances, sensor or actuation malfunction) affecting the data for those individuals. The disturbance model presented in Section 4.4 and the construction of the MDTE as explained in Section 4.5 are intended to encompass some of these errors.

4.4. Disturbance model

In determining an appropriate disturbance model, it is relevant to examine possible sources of error. These include, but are not limited to:

(i) sensor noise and drift,
(ii) actuator (insulin pump) not delivering exactly what is requested,
(iii) intra-individual variability[4] (on a daily basis) in food and insulin responses.

Points (i) and (ii) are evolving as technology improves [49, 50].

We fit the disturbance model to the collected data. We make the simplifying assumption that, although food and insulin models need to be personalised, the disturbance response is the same,

[4]Note that inter-individual variability is not relevant here since the modelling is directed at each individual.

or at least similar, for all individuals. We hypothesise that superposition applies and include in (7) the disturbance induced errors, e, in additive form as follows:

$$y = G_F\, f + G_I\, u + e. \tag{9}$$

Given the Central Meditwin estimates $\hat{G}_F^i$ and $\hat{G}_I^i$ of each transfer function G_F^i and G_I^i (where the index $i \in \{1, \ldots, N_i\}$ denotes each individual and N_i is the number of individuals in the trial) and the collected data $\{y_k^{i,j}\}$, for $k = 1, \ldots, N$, $i = 1, \ldots, N_i$, $j = 1, \ldots, N_T$ (where $N_T = 4$ is the number of tests), a total of $N_i N_T$ realisations of the time series $\{e_k^{i,j}\}$ are available. An estimate of these time series is obtained by subtracting the response to the central model, i.e.,

$$\hat{e}_k^{i,j} = y_k^{i,j} - \hat{G}_F^i f_k + \hat{G}_I^i u_k^j. \tag{10}$$

The estimates $\{e_k^{i,j}\}$ are shown in Figure 5. To remove the effect of the initial conditions and to have sequences of equal length (matching the shortest data sequence), the interval $k \in [10, 54]$, corresponding to time between 50 and 270 minutes, was used to fit the disturbance model.

We chose to model e_k as the output of a second-order filter, $T_e(z)$, driven by zero-mean, unity variance, white noise, ν_k. Error models of this type have been used elsewhere in the literature [51, 52, 53] to describe sensor errors. Here our goal is to capture these errors plus errors due to other sources including stress, mild exercise, daily model parameter variability, etc. This disturbance model amounts to 'low-pass' filtered white noise. It is a discrete-time counterpart of the diffusion process generated by a continuous-time representation of disturbances.

To find the filter parameters, the averaged periodogram of the estimates (10) was fitted to the averaged periodogram of the outputs of the filter driven by $N_i N_T$ realisations of the noise ν_k. The data and fitted averaged periodograms are shown in Figure 6. This method resulted in the discrete-time transfer function $T_e(z)$ in the model $e_k = T_e(z)\, \nu_k$ given by

$$T_e(z) = \frac{2\sqrt{0.02}}{(z/0.8 - 1)^2}. \tag{11}$$

The variance of the resulting output noise e_k is approximately 1.15, resulting in a standard deviation of approximately 1.07 mmol/L.

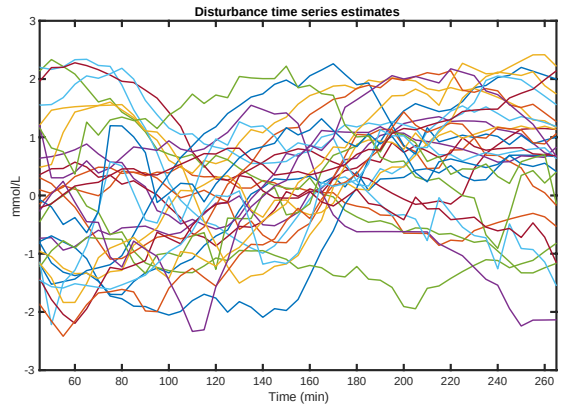

Figure 5: Estimates of the disturbance time series obtained using (10).

4.5. Envelope Quantification

The Central Meditwin fitting procedure described in Section 4 delivered a single model for each individual. However, this model is too optimistic to be useful in practice. We thus propose to

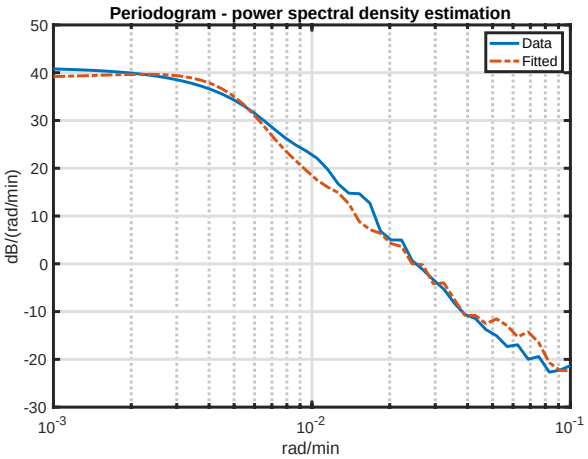

Figure 6: Disturbance time series fitting.

quantify the impact of possible realisations of e_k on the system identification procedure. This leads to an envelope of models around the Central Meditwin. This envelope captures the impact of uncertainty in the model parameters. We proceed by producing pL realisations of the process e_k via simulation, where p is the number of tests whose data is used to fit the food and insulin models. Adding the resulting vectors E_ℓ^j, $\ell = 1, \ldots, L$, $j = 1, \ldots, p$ to the measured time responses Y^{μ_j} (of the form (1)) used to fit the model, the procedure described in Sections 4.1 and 4.2 can be repeated L times to generate L distinct estimates of the joint parameter vectors $(\hat{\theta}_F^\ell, \hat{\theta}_I^\ell)$, $\ell = 1, \ldots, L$.

Remark 4.1.

1. *The model estimation was performed using Matlab's 'fmincon' optimisation function. The original 'central' estimates, $(\hat{\theta}_F^0, \hat{\theta}_I^0)$, obtained in Sections 4.1 and 4.2 were used as the initial guess for the optimisation.*
2. *The cost function used for the MDTE fitting is of a similar form to* (8) *but the two last terms were replaced by a 'regularisation' term [45] penalising large departures from the Central Meditwin parameters, i.e., via the 2-norm of $(\hat{\theta}_F^\ell - \hat{\theta}_F^0, \hat{\theta}_I^\ell - \hat{\theta}_I^0)$.*
3. *The estimates $(\hat{\theta}_F^\ell, \hat{\theta}_I^\ell)$ are correlated*[5] *and hence the parameter estimates for the food and insulin models must be considered as a pair for $\ell = 1, \ldots, L$.*

To illustrate the above procedure, we show in Figure 7 the estimates obtained for the gains K_F and K_I for $L = 1000$.

Note that the gains on food and insulin have opposite signs but are positively correlated in magnitude, i.e., an increase in the magnitude of one gain typically corresponds to an increase in the magnitude of the other. Hence, if the two models are used to evaluate an insulin dose that compensates food, then the increased gain magnitudes on both will tend to cancel each other.

The uncertainty envelope for the model parameters is next embedded in the model via a description of the form:

$$y^s = y_F^s + y_I^s, \qquad y_F^s = \mathcal{G}_F(\theta_F^s)\, f, \quad y_I^s = \mathcal{G}_I(\theta_I^s)\, u, \quad \theta^s \doteq (\theta_F^s, \theta_I^s) \in \mathcal{E}, \tag{12}$$

where $\mathcal{G}_F(\theta_F^s)$, $\mathcal{G}_I(\theta_I^s)$ are transfer functions of the form (5), (6). The delay is not considered

[5] This correlation is natural since reduced insulin sensitivity is correlated with increased sensitivity to food.

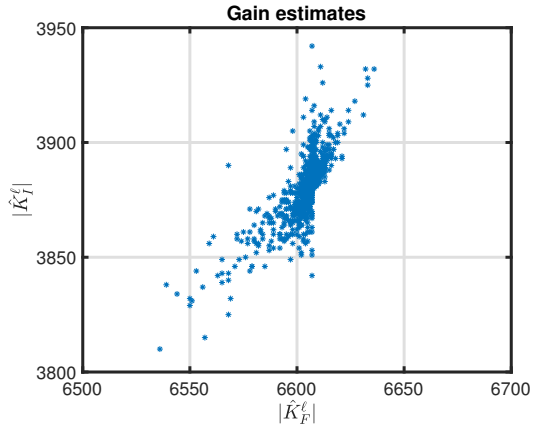

Figure 7: Estimates for the gains K_F and K_I.

in (6) since all designs will assume that insulin administration starts 15 min prior to food consumption, as was done in the clinical trials. In (12), f is the food input, u is the insulin input and θ^s is a stochastic parameter vector with support in $\mathcal{E} \subset \mathbb{R}^8$ and probability measure $\mathbb{P}_\theta$. Since $\mathbb{P}_\theta$ is unknown, when testing different insulin administering strategies, a probabilistic analysis of BGL predictions will be performed by utilising a large number of parameter samples (or scenarios) obtained as described above.

5. MDTE Validation

To validate the MDTE we compare all the collected raw data with the MDTE predictions in response to the dual-wave inputs used in the trials. To do this, we add to the responses $\hat{y}$ given by (12) a set of independent realisations, $e^s(t)$, of the process $e(t)$. This yields a set of possible predicted responses of the form

$$y^s = \mathcal{G}_F(\theta_F^s)\, f + \mathcal{G}_I(\theta_I^s)\, u + e^s. \tag{13}$$

As an illustration, Figure 8 shows the resulting MDTE predictions based on an envelope comprising $L = 100$ model estimates for Subject 1. The blue shaded areas are the ± 1.5 standard-deviation (SD) envelopes around the MDTE mean prediction (plotted in red). The results for the remaining individuals are shown in Appendix B.

It can be seen that for Subject 1 the envelopes do indeed capture (within ± 1.5 SD bounds) the responses obtained in all the tests. For all individuals, the ± 1.5 SD envelopes around the MDTE mean predictions are within the order of ± 2 mmol/L. The criteria to consider the MDTE predictions reliable for an individual are: (i) the data is "biologically consistent" as explained in Section 4.1; and (ii) the measured BGL traces of all tests including those not used in the fitting step are within the ± 1.5 SD envelope of the MDTE mean predictions for the relevant curve's conditions for at least 70% of the time.

6. Stochastic Optimisation Using the MDTE

Here we describe phase 5 of the procedure shown in Figure 1.

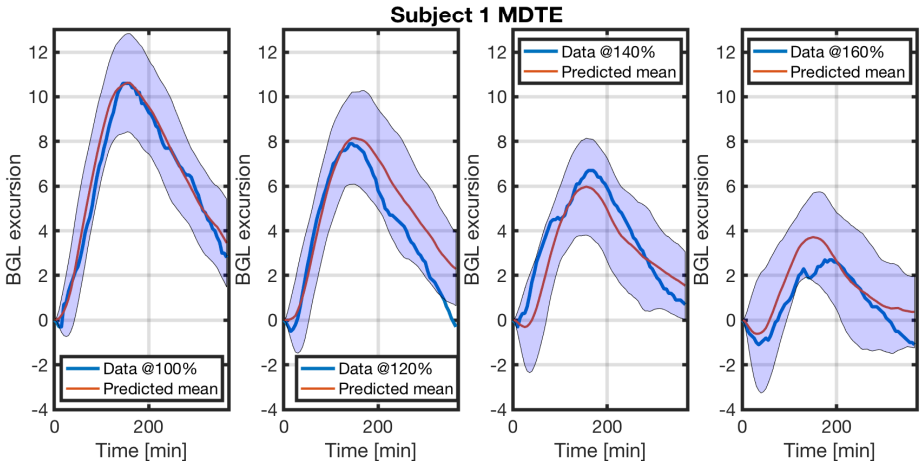

Figure 8: MDTE validation for Subject 1. The blue shaded areas are the ± 1.5 standard-deviation envelopes around the MDTE mean prediction (plotted in red).

6.1. Overview

This section describes one of the key results of the paper, namely how we use the digital twin envelopes in a systematic stochastic design procedure which trades off peak BGL minimisation and predicted probability of hypoglycaemic events based on the uncertainty inherent in θ^s and e^s.

As an illustration of the power of utilising MDTEs to complement the clinical trials on real individuals with T1DM, we ask the following questions for the given meal:

(i) What is the best dual wave strategy when the duration of the extended bolus can be chosen?
(ii) What is the best dual bolus (or 'split bolus') strategy?
(iii) What is the best dual wave strategy when the duration of the extended bolus is set to 180 minutes?
(iv) What is the best single bolus strategy?

In all cases, we specifically account for model uncertainty in the design.

To find the best strategy in cases (i) to (iv), we refer to equation (12) and consider L model realisations corresponding to parameter estimates $(\theta_F^\ell, \theta_I^\ell)^s$, $\ell = 1, \ldots, L$. For each of these realisations we consider a discrete-time model equivalent via equation (12). We let $\mathcal{G}_F^d(\hat{\theta}_F^\ell)$ and $\mathcal{G}_I^d(\hat{\theta}_I^\ell)$ be zero-order hold discretisations, with a chosen sampling period T_s, of the continuous-time transfer functions in (12). The sample instants obtained from the sampling period T_s are denoted by $k = 0, 1, \ldots$. Unless otherwise stated, the sampling period is taken as $T_s = 10$ min.

We then seek an 'optimal' insulin input $u(k) = u^*(k)$ using the following relaxed optimisation:

$$u^*(\cdot) = \arg\min_{u(\cdot)} \Big\{ J \doteq \sum_{\ell=1}^{L} \sum_{k=1}^{N} \hat{y}_\ell^2(k) \mid \hat{y}_\ell(k) = \mathcal{G}_F^d(\hat{\theta}_F^\ell) f(k) + \mathcal{G}_I^d(\hat{\theta}_I^\ell) u(k), \hat{y}_\ell(0) = 0, \\ u(k) \text{ of desired form}, \hat{y}_\ell(k) \geq y_\#, \forall k \geq 0, f(k) = F\delta(k), F > 0 \Big\}. \quad (14)$$

F in (14) is the amount of CHO in the meal and $\delta(k)$ denotes a unit-area pulse. In the sequel we take $F = 1$, since the amount of CHO ingested was implicitly identified as part of the food transfer function gains. A very important aspect of the cost function (14) is that the constraint parameter $y_\#$ is *not* the target minimum BGL. Instead, it represents a *constraint relaxation parameter* used in the design to control the probability of occurrence of hypoglycaemic events.

The 'desired form' for the insulin input u varies in each of the above cases, namely dual wave, split bolus, dual wave with fixed duration and single bolus. We take the number of samples as $N = 36$, that is an optimisation horizon of $N\,T_s$ =360 min.

6.2. Choice of Constraint Relaxation Parameter

A key element of our proposed design strategy is the use of the relaxation parameter $y_{\#}$. We use this parameter as a mechanism to 'convexify' the problem of trading off average peak BGL reduction and the probability of violation of the constraint $y(k) \geq y_{min}$. We note that there are two sources of error in our predictive model, namely (i) uncertainty in the model parameters resulting from different values of $e(t)$ experienced during data collection and (ii) uncertainty in the values of $e(t)$ that are encountered when the model is used. To simplify the explanation, say that the range of possible $e(t)$ is bounded in the interval $[-a,\ a]$, for some positive number a. Our design ensures that no possible value of $\hat{\theta}$ can result in a response such that $\min_t \hat{y}(t) \leq y_{\#}$. Figure 9 shows possible values of $\min_t \hat{y}(t)$ and $\hat{e}(t^*)$ where $t^* \doteq \arg\min_t \hat{y}(t)$. We note that some values of $\hat{e}(t^*)$ play a beneficial role whilst others have a negative impact.

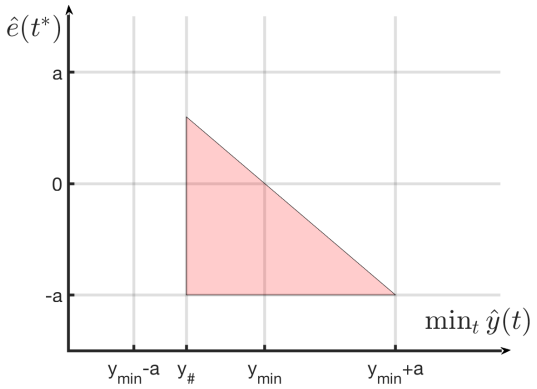

Figure 9: Illustration of the use of $y_{\#}$ to adjust the probability of violation of the hypoglycaemic constraint.

The responses to the left of $y_{\#}$ in Figure 9 are ruled out by the design of the insulin input using the relaxed cost function. On the other hand, any combination of $\min_t \hat{y}(t)$ and $\hat{e}(t^*)$ lying in the shaded area will result in $\min_t y(t) < y_{min}$, where y is as in (13), i.e., violation of the hypoglycaemic constraint. It is clear from Figure 9 that when $y_{\#}$ is increased then the probability of constraint violation decreases since the set of offending values of $\min_t \hat{y}(t)$ and $\hat{e}(t^*)$ is a strict subset of the previous values. Conversely, if $y_{\#}$ is decreased then the probability of constraint violation necessarily increases. In addition, as $y_{\#}$ is increased the cost function value is non-decreasing, and vice versa. Hence the use of $y_{\#}$ allows one to convexify the problem of achieving a trade-off between the average peak BGL response (as measured by the cost function) and robustness (as measured by the probability of constraint violation).

We have extensively tested the above idea for Subject 1. For example,[6] using $y_{\#} = -5$ mmol/L and $y_{min} = -3$ mmol/L for the optimal dual wave design gives a 5% chance of constraint violation due to uncertainty in the model parameters alone and 18% chance of constraint violation due to the combined uncertainty in both the model parameters and the future error process realisation.

[6]Note that y_{min} is the desired constraint whereas $y_{\#} < y_{min}$ is the constraint relaxation parameter.

In the sequel, we will exclusively use parameter uncertainty but note that if uncertainty in the future error process were to be also considered then this would increase the probability of constraint violation by about 10%. As discussed above, both probabilities can be adjusted via the constraint relaxation parameter $y_{\#}$. This tradeoff is further examined in Section 8.1.

Remark 6.1. *We recall that here we deal only with deviations from the entry BGL. Thus y_{min} is actually a (negative) deviation from the entry BGL. To illustrate the ideas we will initially assume an entry BGL of 7 mmol/L. Later in Section 8.2 we will explore the impact of different entry BGL values on the design strategy.*

7. The Optimised Insulin Policies

Tha availability of the MDTE for an individual allows us to examine and optimise many different strategies without needing to further bother the individual.

We first illustrate the above ideas for Subject 1. We then repeat for the remaining individuals.

7.1. Best Dual Wave Strategy

We perform the optimisation (14), where the form of the input is constrained to be a bolus (whose size is to be determined) followed by an extended bolus (square wave) whose amplitude and duration are also to be found by the optimisation. The envelope size is taken initially as $L = 100$ and the results will then be compared with those using more models in the envelope.

We observe that, for each fixed duration of the square wave, say w samples, the problem (14) is a quadratic programme in the decision variables $u(0), \ldots, u(N-1)$, where $u(0)$ represents the size of the initial bolus, and $u(1) = \cdots = u(w)$, $u(w+1) = \cdots = u(N-1) = 0$ represent the additional input constraints required to yield the desired square wave. Since quadratic programmes are convex and computationally efficient, we will solve problem (14) for each duration w in the integer interval $[6, 18]$, which for the sampling period $T_s = 10$ min correspond to actual extended bolus durations of 60 min, 70 min, etc. up to 180 min. Each value of w yields a cost $J = J(w)$ resulting from the optimisation (14). The optimal duration, w^*, is then the one that yields the smallest cost, that is $w^* = \arg\min_w J(w)$.

Proceeding as described above we obtain the following results:
– Optimal duration: $w^* = 11$, i.e., 110 min.
– Total insulin: 1.7V.
– Split: 39:61.
– Max predicted mean BGL deviation: 1.5 mmol/L.
– Probability $\min_t \hat{y}(t) < y_{min}$ of 5%.

The predicted BGL excursions, $\hat{y}(t)$, resulting from the best dual wave strategy for the MDTE associated with Subject 1 are shown in the first subplot of Figure 10. The mean response is shown by the solid red line, the ± 1.5 standard deviation bounds are shown by blue lines, and all $L = 100$ MDTE responses are shown by dashed lines.

Comparing with the results obtained in the clinical trial, the above values suggest that the duration of 180 min used in the testing phase for the extended bolus is longer than is optimal for Subject 1.[7] Indeed, the optimal dual wave responses for this individual, as seen in the first sub-

[7]The clinical members of our team have taken note of this discrepancy and are hence currently testing dual wave policies having reduced time for the extended bolus.

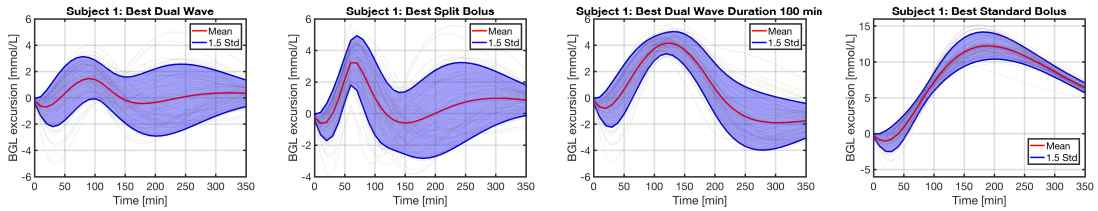

Figure 10: Best dual wave, split bolus, dual wave of duration 180 min and standard bolus strategies for Subject 1.

plot of Figure 10, achieve peak BGL values similar to those obtained for the $+60\%$ test (cf. the fourth subplot of Figure 8) but applying slightly more insulin, i.e., 70% more than the standard bolus for this individual over the shorter duration of 110 min.

The optimisation was repeated for different envelope sizes. The resulting optimal mean responses are shown in Figure 11 for $L \in \{100, 200, 300, 400\}$. We can see that the results are consistent for all sizes, with a slightly larger peak for $L \geq 200$. Hence, unless otherwise stated, we will henceforth consider $L = 100$.

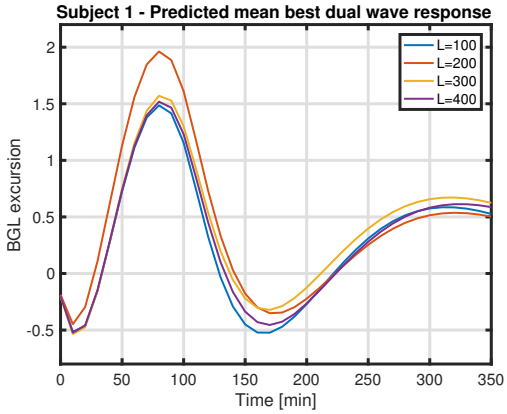

Figure 11: Effect of different envelope sizes. L: number of different model realisations included in the MDTE.

7.2. Best Split Bolus Strategy

We perform the optimisation (14), where the input is now constrained to consist of an initial bolus of size to be determined and a second bolus whose size and application time are also to be found by the optimisation. The application time is considered with respect to the application of the first bolus, which is 15 min before food consumption. Here again we observe that for fixed application time, w, of the second bolus, the problem (14) is a quadratic programme in the decision variables $u(0), \ldots, u(N-1)$, where $u(0)$, $u(w)$ represent the bolus sizes to be found, and all other input values are constrained to be zero. As for the dual-wave case, we solve problem (14) for a range of values w for the application time of the second bolus. Each value of w yields a cost $J = J(w)$ resulting from the optimisation (14). The optimal application time of the second bolus, w^*, is then the one that yields the smallest cost, that is $w^* = \arg\min_w J(w)$.

The predicted BGL excursions, $\hat{y}(t)$, resulting from the best split bolus strategy for the MDTE associated with Subject 1 are shown in the second subplot of Figure 10.

The corresponding results are:
– Second bolus application time: 60 min.
– Total insulin: 1.65V.

– Split: 41:59.
– Max predicted mean BGL deviation: 3.22 mmol/L.
– Probability $\min_t \hat{y}(t) < y_{min}$ of 5%.

7.3. Best Dual Wave Strategy Having Fixed Duration

Here we perform the optimisation as described in Section 7.1, but we fix the duration of the extended bolus to be 180 min, i.e., the square wave is applied over $w = 18$ samples. This duration corresponds to that used in the clinical trial described above. The predicted BGL responses for the best dual wave strategy of duration 180 min for the MDTE associated with Subject 1 are shown in the third subplot of Figure 10.

The corresponding results are:
– Total insulin: 1.76V.
– Split: 41:59.
– Max predicted mean BGL deviation: 4.16 mmol/L.
– Probability $\min_t \hat{y}(t) < y_{min}$ of 34%.

Comparing with the results obtained in the clinical trial, the above values are similar to those obtained for the $+60\%$ test, cf. the fourth subplot of Figure 8. Here, however, the recommendation is to apply approximately 15% more insulin than the latter test, and reverse the split to 40:60 instead of 60:40 as used in the trial.

Note that the probability of constraint violation is very high, 34%. Any attempt to reduce this value by increasing $y_{\#}$ (cf. Section 6.2) changed the nature of the optimal solution leading to the standard (single) bolus strategy.

7.4. Best Standard Bolus Strategy

Here we design a 'standard' bolus strategy consisting of a single bolus applied 15 min prior to food ingestion. The best strategy provides the size of the bolus for which the optimal input (14) yields the minimum cost. The predicted BGL responses for the best standard bolus strategy for the MDTE associated with Subject 1 are shown in the last subplot of Figure 10.

The corresponding results are:
– Total insulin: 0.8V.
– Max predicted BGL deviation: 12.25 mmol/L.
– Probability $\min_t \hat{y}(t) < y_{min}$ of 5%.

From the above values and the traces in the last subplot of Figure 10 (note the different vertical axis scale in this subplot), it is clear that a (single) standard bolus gives an inferior result for this individual (Subject 1) and for the given food.

7.5. Digital Trials for All Individuals

The optimisation procedures reported above were repeated for the MDTEs associated with all individuals. The numerical results are summarised in Table 1. The MDTE BGL predictions, $\hat{y}(t)$, resulting from the best dual wave, split bolus, dual wave with fixed duration, and single bolus strategies are given in Appendix C. All tests were aimed at achieving a probability of violation of the hypoglycaemic constraint $\hat{y}(t) \geq y_{min} = -3$ mmol/L of approximately 5% or less and assume an entry BGL of 7 mmol/L. For the best standard bolus and Subjects 1, 2, 7 and 13, an additional test was performed where this probability was allowed to be higher (see

the pink cells under 'Best Standard Bolus' in Table 1), to illustrate the performance-robustness tradeoff discussed in Section 8.1 below.

Table 1: Digital trial results for all 12 individuals. Subject data: # is the subject identification (ID) number in the trial and 'Tests' displays the numbers of the tests whose data were used to fit the models. The other variables are: w: duration of the extended bolus; t_{app}: time of application of the second bolus. 'Peak' is the peak of the average response. '$< y_{min}$' refers to the percentage of violation of the hypoglycaemic constraint BGL $\geq y_{min} = -3$ mmol/L.

Subject Data		**Best Dual Wave**					**Best Split Bolus**				
#	Tests	Insulin	Split	w	Peak	$< y_{min}$	Insulin	Split	t_{app}	Peak	$< y_{min}$
1	(1,2,3,4)	1.7080	39:61	110	1.4649	5%	1.6502	41:59	60	3.2249	5%
2	(2,3,4)	1.5212	51:49	110	1.3197	5%	1.4038	52:48	50	2.7382	4%
4	(1,2,3,4)	2.0181	57:43	170	1.7178	5%	1.9050	66:34	100	2.3938	5%
5	(1,2,4)	1.5009	63:37	120	0.5860	0%	1.4803	70:30	70	0.8588	2%
6	(1 2,4)	1.3524	86:14	110	0.1739	0%	1.3502	89:11	70	0.1615	0%
7	(2,3,4)	1.2782	55:45	80	0.8401	4%	1.2553	55:45	40	1.1514	5%
9	(1,2,3,4)	1.4844	54:46	190	0.3464	4%	1.4449	64:36	110	0.6253	4%
10	(1,2,3,4)	3.3515	4:96	350	0.9062	5%	1.7970	37:63	150	4.6320	5%
11	(1,2,4)	1.2318	100:0	0	2.8128	5%	1.2318	100:0	0	2.8128	5%
12	(1,2,3,4)	1.7613	89:11	350	0.6685	5%	6.9726	22:78	350	0.8188	5%
13	(2,3,4)	1.3766	51:49	160	1.4514	5%	1.3006	63:37	100	2.4906	5%
16	(2,3,4)	1.3783	29:71	90	2.2675	5%	1.3720	43:57	60	2.5798	5%

Subject Data		**Best Dual Wave w** $= 180$				**Best Standard Bolus**					
#	Tests	Insulin	Split	Peak	$< y_{min}$	Insulin		Peak		$< y_{min}$	
1	(1,2,3,4)	1.7610	41:59	4.1602	34%	0.8067	1.1187	12.2485	8.8866	5%	36%
2	(2,3,4)	1.5027	49:51	4.4145	5%	0.9751	1.1417	8.2008	6.7485	5%	48%
4	(1,2,3,4)	2.0678	56:44	1.7200	5%	1.2939		4.7452		5%	
5	(1,2,4)	1.5494	69:31	0.8133	3%	1.1301		3.100		5%	
6	(1 2,4)	1.3672	89:11	0.2369	0%	1.3213		0.3730		0%	
7	(2,3,4)	1.2955	53:47	3.8876	6%	0.7987	1.1410	6.5635	3.1693	5%	61%
9	(1,2,3,4)	1.4721	53:47	0.3507	4%	1.0794		5.0914		5%	
10	(1,2,3,4)	1.7950	8:92	4.9064	5%	0.6881		7.2635		5%	
11	(1,2,4)	1.2318	100:0	2.8128	5%	1.2318		2.8128		5%	
12	(1,2,3,4)	1.6940	93:7	0.6126	5%	1.5569		0.8181		0%	
13	(2,3,4)	1.3934	53:47	1.5756	5%	0.8156	1.1345	5.4574	3.6523	5%	31%
16	(2,3,4)	1.4735	34:76	3.9446	6%	1.0332		6.0744		5%	

7.6. Discussion

- Subjects with ID# in a green cell have Central Meditwins and MDTEs that reproduce/contain the data reasonably well (i.e., they satisfy the validation criteria given in Section 5). Subjects with ID# in an orange cell have Central Meditwins and MDTEs that fail to reproduce/contain larger parts of the data. Subjects with ID# in a red cell have Central Meditwins and MDTEs that fail to reproduce/contain large parts of the data in addition to having highly inconsistent data (i.e., several of the tests do not respect the 'area ordering').
- The yellow cells flag abnormal results. The probability of constraint violation of 34% for Subject 1 was explained in Section 7.3. The yellow cells in the best split bolus results for Subject 12 are not realistic since they advocate a large amount of insulin applied at the end of the optimisation horizon (see also the second subplot in Figure C.39). In fact, the results in Table 1 and Figure C.39 suggest that the standard bolus achieves the best compromise between performance (low BGL peak) and robustness for Subject 12.
- All strategies for Subject 11 coincide with the best standard bolus strategy.

- For Subjects 1, 2, 7 and 13, the insulin amount for the best standard bolus is less than 1, i.e., less than the individual's ICR (dark blue cells). Note that in those cases the split in the other policies tends to be around 50% (or less) for the initial bolus (lighter blue cells for the same individuals). These results suggest that for those subjects (ID# 1, 2, 7 and 13) and this particular meal, it may be advisable to be cautious with the amount of insulin applied upfront to avoid hypoglycaemia.
- For Subjects 1, 2, 7 and 13, the pink cells in the best standard bolus results correspond to a larger bolus (between 1.12 and 1.14 times the person's ICR). Note that these amounts naturally yield lower peak BGL values. However, this comes at the cost of a much larger probability of violation of the hypoglycaemia constraint (cf. Section 8.1).

8. Design Tradeoffs

Here we examine the inherent design tradeoffs in more detail.

8.1. Tradeoff Between Performance and Robustness

Allowing a larger probability of constraint violation makes a strategy 'less robust', in the sense of increasing the chance of leaving the 'safe' BGL zone. On the other hand, it improves the associated performance by reducing the BGL peaks. This tradeoff is illustrated for Subjects 1 and 4 in Figure 12, where the average maximum BGL excursion achieved by the best standard bolus strategy (right vertical axis) and the corresponding amount of insulin (left vertical axis) are plotted against the percentage of violation by $\hat{y}(t)$ of the hypoglycaemic constraint BGL $\geq y_{min} = -3$ mmol/L (assuming an entry BGL of 7 mmol/L).

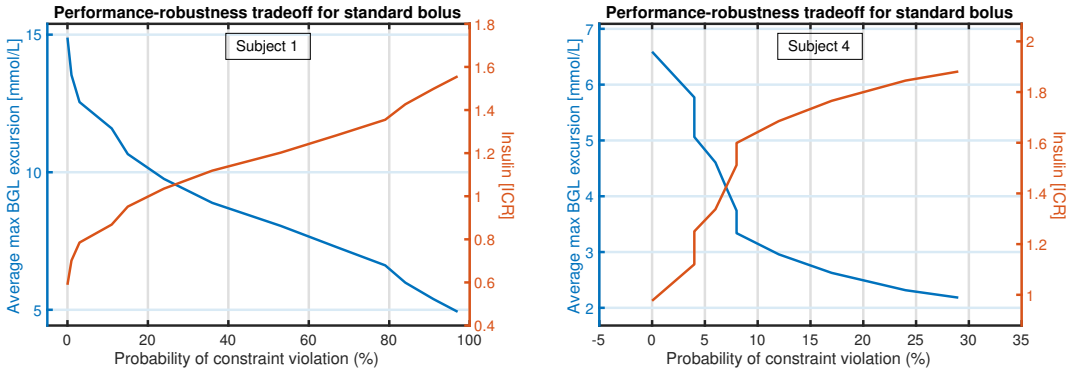

Figure 12: Performance-robustness tradeoff for Subjects 1, 4, and the best standard bolus strategy.

For example, for Subject 1, if one is prepared to accept a prediction that there is a 20% chance that $\hat{y}$ will fall below y_{min}, then an amount of insulin equal to 1 ICR would achieve an average peak excursion of approximately 10 mmol/L. Note also that, for this individual and this meal, the results suggest that it is not possible to reduce the average maximum BGL excursion below (approximately) 5 mmol/L. This can be achieved with (approximately) 1.6 ICR and at the (clinically unacceptable) cost of a 97% chance of violating the hypoglycaemia constraint. This reinforces the claim, made in Section 7.4, that other strategies are preferable for Subject 1 and this particular meal. On the other hand, the performance-robustness tradeoff is very different for Subject 4. Indeed, from the right plot of Figure 12 we see that a very low value of (approximately) 2.2 mmol/L for the average maximum BGL excursion can be achieved with (approximately) 30% probability of having $\hat{y} < y_{min}$ if a single bolus of (approximately) 1.9 ICR is

applied upfront. In fact, as seen in Section 7.5, for Subject 4 and this particular meal a standard bolus is likely to yield similar results to those achieved for the other tested strategies where insulin is split between a bolus and some form of extended (or second) bolus. Similar conclusions are shown for other individuals in Section 7.5.

8.2. Impact of Entry BGL

In the above analysis, we have assumed an entry BGL of 7 mmol/L by way of illustration. We note that the average entry BGL for all tests and individuals in the trial was 8.8 mmol/L. The minimum entry BGL was 4.1 mmol/L and the maximum 9.8 mmol/L. For cases where the entry BGL is larger than 7 mmol/L then one can safely use a more aggressive policy. Similarly, if the entry BGL is less than 7 mmol/L then one should use a more cautious policy. That is, the methodology can take into account any entry BGL value by simply redefining the minimum BGL deviation constraint. Thus, the implicit assumption of equal entry BGL is not a limitation of the methodology but rather a choice to simplify the above discussion. The final tests reported in Section 10, where the optimised strategies were tested on a subset of the individuals, necessarily account for entry BGL.

For a fixed probability (i.e., 5%) of having $\hat{y} < y_{min}$, Figure 13 shows the impact of having a different entry BGL value for Subjects 1 and 4.

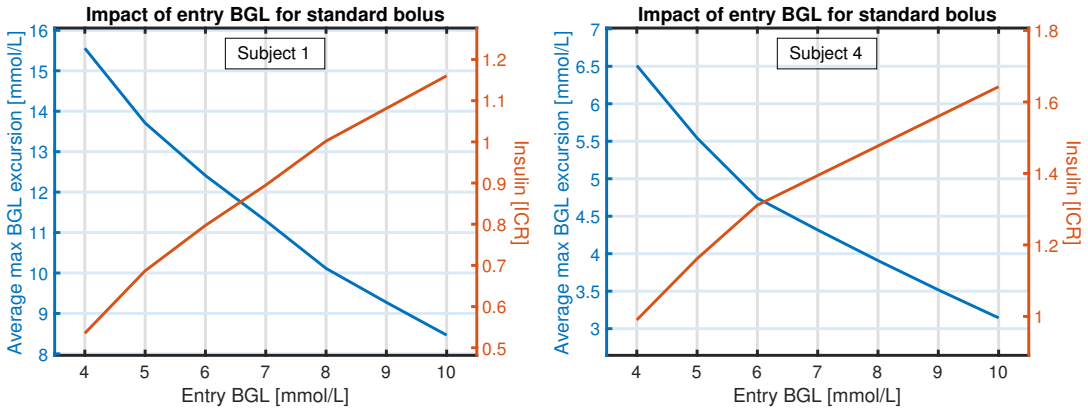

Figure 13: Impact of entry BGL for Subjects 1, 4, and the best standard bolus strategy.

9. Hypotheses Arising from the MDTE Trials

Two main hypotheses arise from the MDTE trials reported above:

Hypothesis A: For the given meal (high-fat, high-protein) there is a significant advantage to be obtained from using some form of extended bolus rather than a single bolus.

Hypothesis B: There appears to be no significant difference between dual wave and split bolus strategies provided they are "tuned" to each individual.

Clinical data is available relating to Hypothesis A above. Specifically, work reported in [38] confirms that for a high-fat, high-protein meal, a dual-wave strategy consistently outperforms a single bolus. Hypothesis B is currently being studied in clinical trials by the medical team associated with the current paper.

10. Final Evaluation on a Subset of the Individuals

This section relates to phase 6 of the procedure shown in Figure 1.

Obviously the ultimate test of the efficacy of the procedure described here is to return to the same individuals to test the proposed "optimised" strategies. It would have been desirable to return to all individuals but this was not achievable due to the passage of time since the original tests were performed (some individuals had moved away, others were not available for further participation in clinical trials). However, it was possible to return to two individuals (Subjects 1 and 4) and test the optimised strategies for the same meal utilised in the original tests. Note that the specific entry BGL now had to be considered in the optimisation. To deal with this issue, a table computed for different entry BGLs varying between 4.5 to 8.5 mmol/L, at intervals of 0.5 mmol/L, and the associated optimised dual wave strategies was prepared for each individual. On the day of the test the individual (or their guardian) called the trial coordinator to report the entry BGL. The coordinator then communicated the corresponding insulin dose, split and duration.

The validation results for the two subjects are briefly described below.

10.1. Subject 4

The "optimal" policy obtained by the procedure discussed in Sections 7.1 and 8.2 was applied twice on two successive days. The details are: *Day 1 (entry BGL 4.5 mmol/L)*: Dose: 1.54, Split: 44:56, Duration: 180 min; *Day 2 (entry BGL 5 mmol/L)*: Dose 1.73, Split 47:53, Duration: 200 min).

The results are shown in Figure 14. The plots correspond to Day 1 on the left and Day 2 on the right. The data is shown by thick blue lines and the green shaded areas are the ± 1.5 standard-deviation envelopes around the MDTE mean predictions.

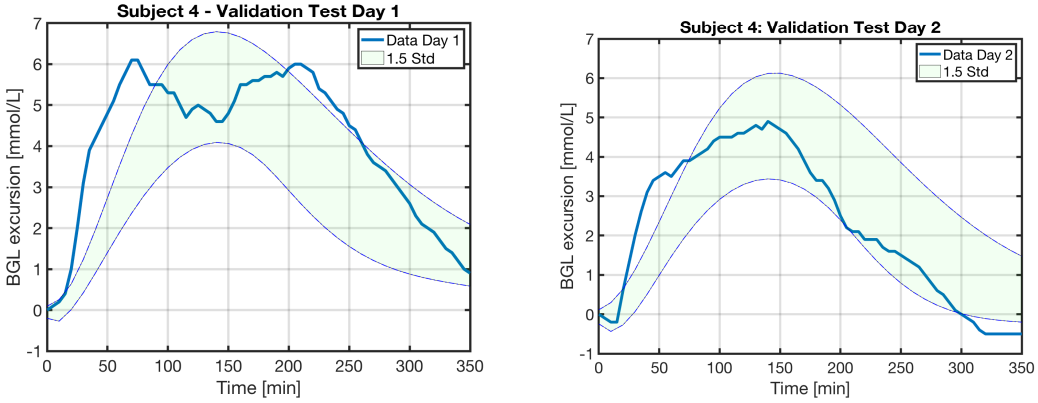

Figure 14: Results of validation test on Subject 4.

Note that the measured BGL traces lie substantially within the prediction envelope provided by the MDTE. Indeed, on Day 1, 65% of the duration of the BGL response lay within the ± 1.5 SD envelope and 82% within the total MDTE envelope (not shown in the figures for clarity) and on Day 2, 74% of the duration of the BGL response lay within the ± 1.5 SD envelope and 81% within the total MDTE envelope. Also, neither test resulted in a hypoglycaemic event (BGL less than 3.9 mmol/L).

10.2. Subject 1

The optimal policy obtained by the procedure discussed in Sections 7.1 and 8.2 was again applied twice. The details are: *Day 1 (entry BGL 7.5 mmol/L)*: Dose 1.71, Split: 43:57, Duration: 110 min; *Day 2 (entry BGL 8 mmol/L)*: Dose: 1.73, Split: 46:54, Duration: 120 min.

The results are shown in Figure 15. The plots correspond to Day 1 on the left and Day 2 on the right. The data is shown by thick blue lines and the green shaded areas are the ± 1.5 standard-deviation envelopes around the MDTE mean predictions.

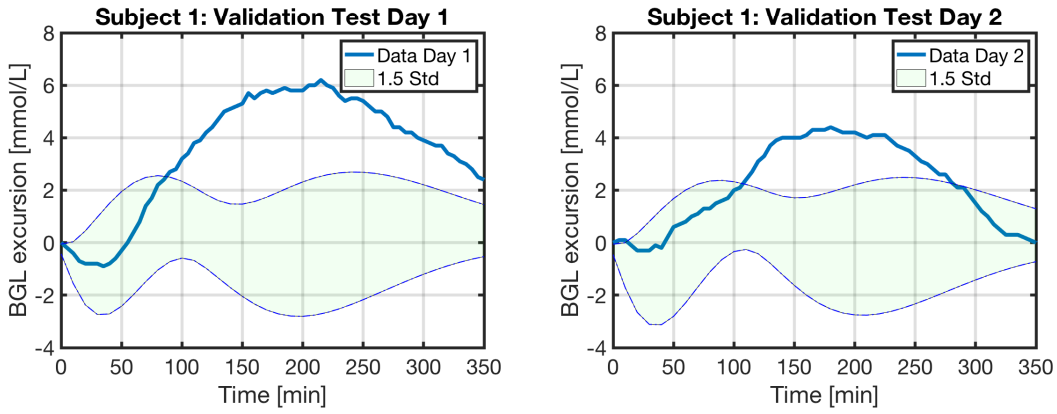

Figure 15: Results of validation test on Subject 1.

Note that a hypoglycaemic event did not occur on either day. However, the responses deviate from the prediction envelope: on Day 1, 24% of the duration of the BGL response lay within the ± 1.5 SD envelope and 54% within the total MDTE envelope (not shown in the figures for clarity); the results improved on Day 2, for which the BGL response duration was 54% within the ± 1.5 SD envelope and 81% within the total MDTE envelope.

The results suggest that the BGL responses for this individual may have changed over the 2 years between the original trials and the new tests. Hence refitting the models and new MDTE strategy testing seems desirable for this individual and is planned in the future.

Even though the full trajectories are not well captured by the MDTE predictions, we note that the responses are consistent in terms of peak values with the original trial results, see the first plot in Figure 2.

11. Extensions and Embellishments

The work reported in this paper has been based on a single meal comprising high fat and high protein. The procedure, as described, does not allow the individual impact of fat, protein or CHO to be separated. Also, other disturbances such as exercise have been excluded. The basic principle described here can be extended to cover other scenarios. In particular, we have so far used two (or more) tests for the fixed high-fat, high-protein meal so as to separate the response to two inputs, namely the given meal and injected insulin. As an extension, say that one wished to distinguish the response to four inputs (e.g., CHO, fat, protein and insulin). Then four (or more) tests on each individual would be needed where the amount of each variable of interest were varied in each test. This represents a modest increase in the testing needed on each individual. Such testing is planned over the next few years and extra funding to support this extension has been sought.

12. Conclusions

This paper has proposed a systematic stochastic optimisation design strategy which explicitly trades off peak BGL minimisation versus predicted probability of experiencing a hypoglycaemic event. The design uses a metabolic digital twin envelope (MDTE), which is a probabilistic model description that takes into account model uncertainty. The MDTE can be fitted using data from minimally intrusive clinical trials conducted at home. It has been emphasised that the data collected from the individual needs to contain as many independent tests (with different input levels) as there are variables of interest, e.g., at least 4 tests are needed to distinguish CHO, protein, fat and insulin responses. The new strategy has been employed to obtain personalised optimal insulin policies of different form, i.e., dual wave, split bolus and standard bolus, for the MDTEs associated with 12 individuals from a recent digital trial performed by the medical team. Finally, the methodology has been tested by returning to (a subset of) the individuals to apply the optimised insulin strategies suggested by the MDTE-based stochastic design.

Acknowledgments

This research was supported by a Diabetes Australia Research Program, Millennium Type 1 Award grant.

M.M. Seron, A.M. Medioli and G.C. Goodwin also acknowledge financial support from the Priority Research Centre for Complex Dynamic Systems and Control, The University of Newcastle.

The authors thank the children and families who participated in this research; the clinical team at John Hunter Children's Hospital, Newcastle, Australia, including Dr Patricia Crock, Dr Don Anderson, Megan Paterson, Michelle Neylan, Marcelle Pappas and Emma Fisher; and the clinical team at Queensland Children's Hospital, Brisbane, Australia, including Dr. Mark Harris and Michelle Fuery.

The authors have no potential conflicts of interest. Unless explicitly attributed, the opinions expressed in this work are those of the authors only and do not represent the official view of Hunter New England Local Health District nor the New South Wales Government.

Appendix A. Central Meditwin Fitting and Predictions for All Individuals

The following figures display the results for the Central Meditwin model fitting and corresponding predictions for all individuals other than Subject 1.

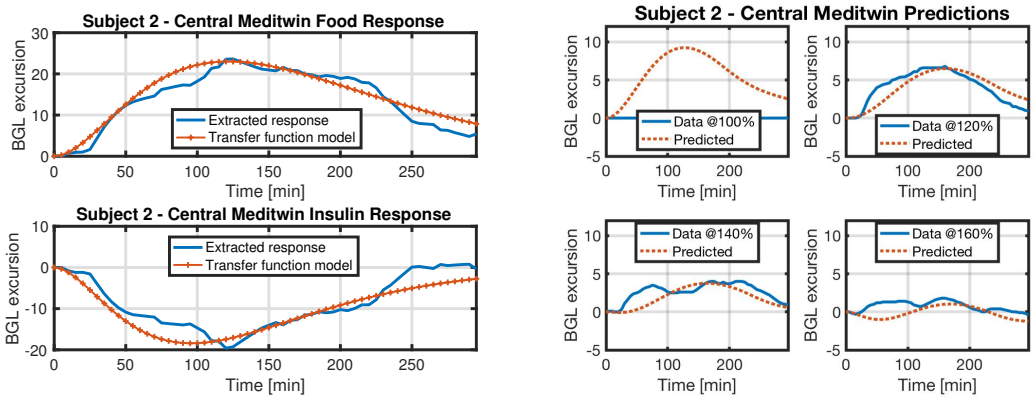

Figure A.16: Central Meditwin fitting and predictions for Subject 2. Tests having consistent data used for fitting: (2,3,4). Note that Test 1 @100% was not performed for this subject, which is indicated by a zero line in the corresponding figure.

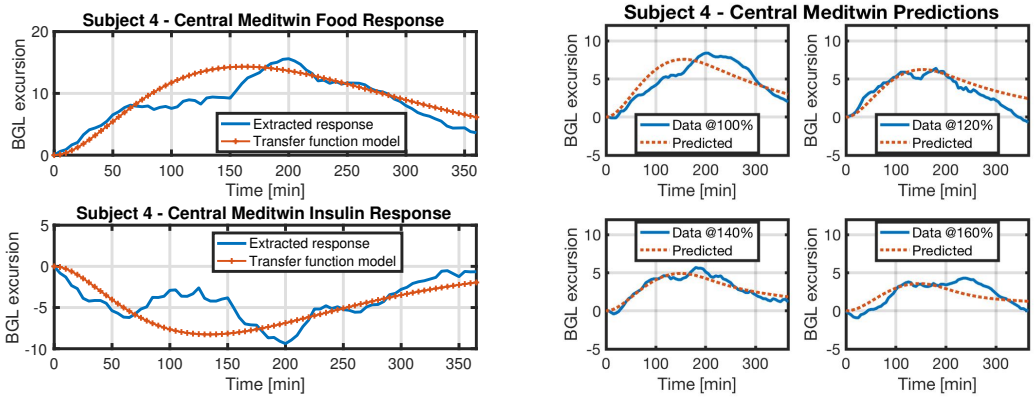

Figure A.17: Central Meditwin fitting and predictions for Subject 4. Tests having consistent data used for fitting: (1,2,3,4).

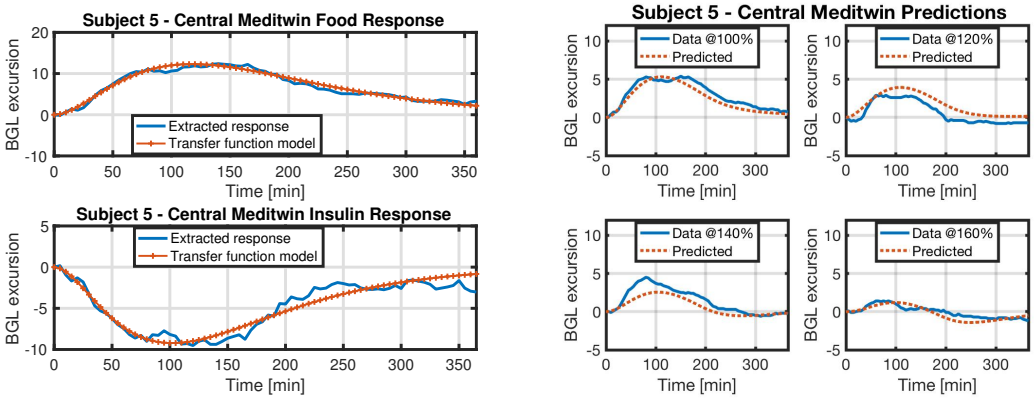

Figure A.18: Central Meditwin fitting and predictions for Subject 5. Tests having consistent data used for fitting: (1,2,4).

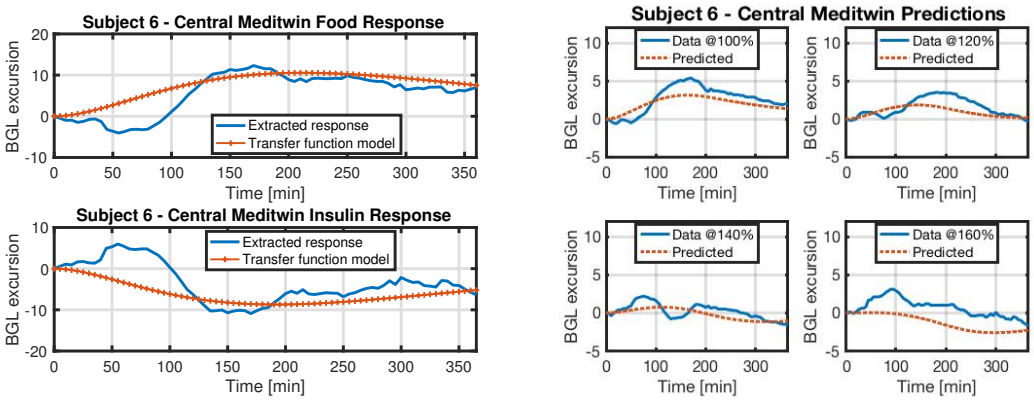

Figure A.19: Central Meditwin fitting and predictions for Subject 6. Tests having consistent data used for fitting: (1,2,4).

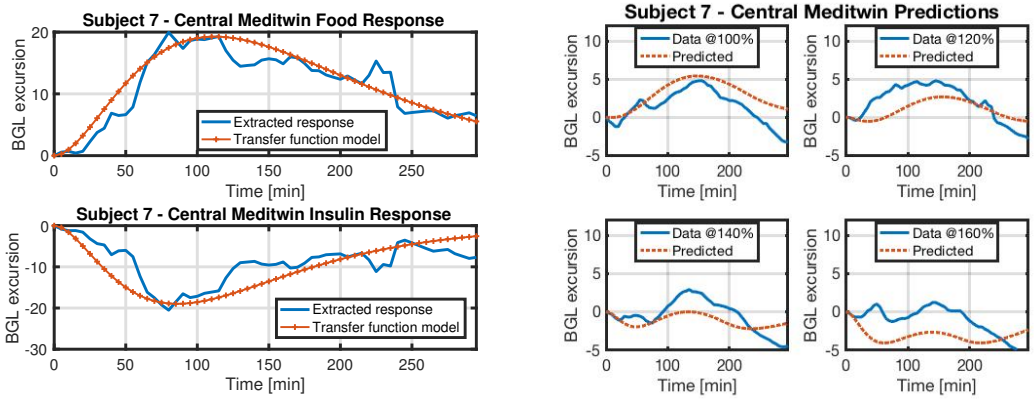

Figure A.20: Central Meditwin fitting and predictions for Subject 7. Tests having consistent data used for fitting: (2,3,4).

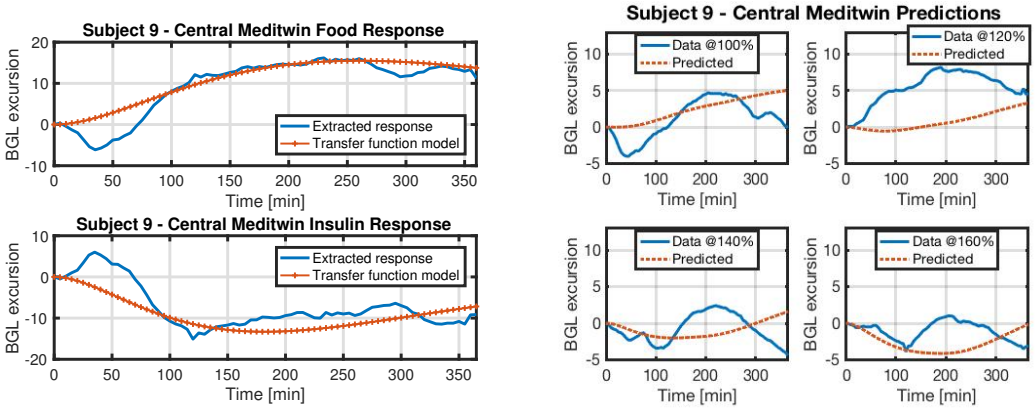

Figure A.21: Central Meditwin fitting and predictions for Subject 9. Tests used for fitting: (1,2,3,4). Note that this data set is not consistent, however, no better fitting was obtained by trying different subsets of this set.

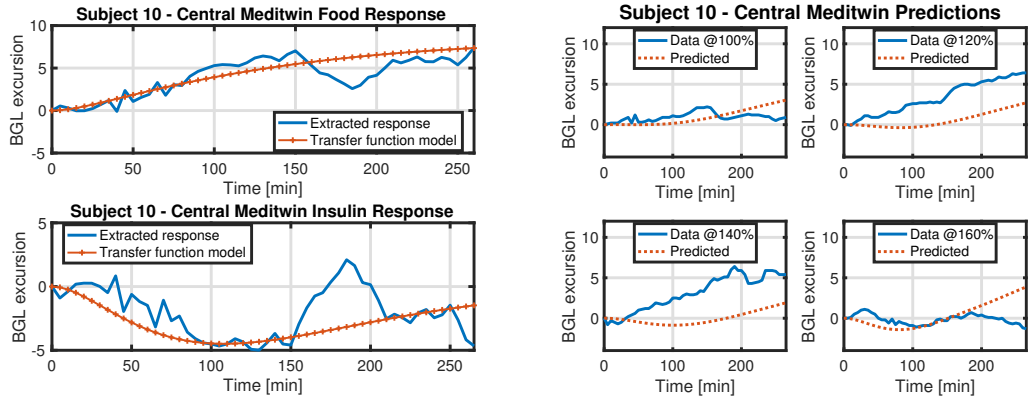

Figure A.22: Central Meditwin fitting and predictions for Subject 10. Tests used for fitting: (1,2,3,4). Note that this data set is not consistent, however, no better fitting was obtained by trying different subsets of this set.

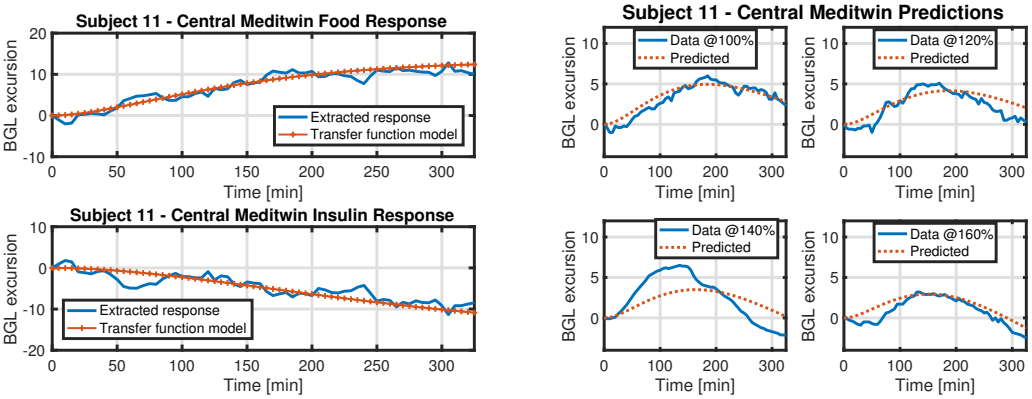

Figure A.23: Central Meditwin fitting and predictions for Subject 11. Tests having consistent data used for fitting: (1,2,4).

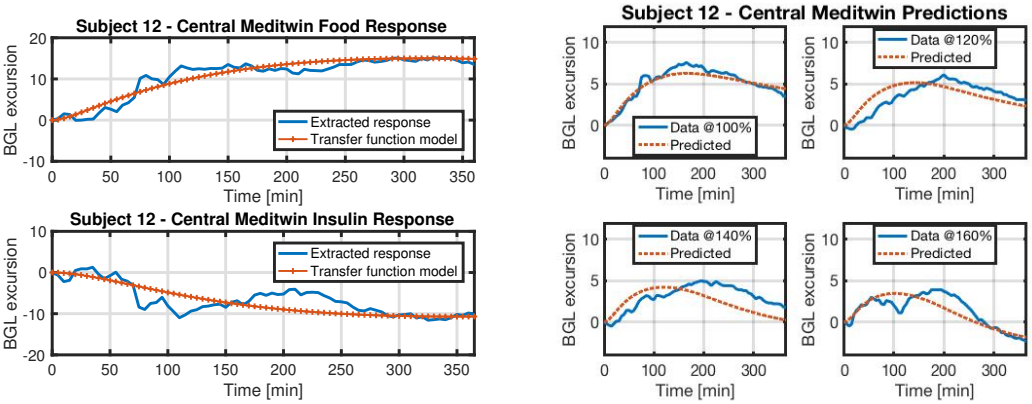

Figure A.24: Central Meditwin fitting and predictions for Subject 12. Tests having consistent data used for fitting: (1,2,3,4).

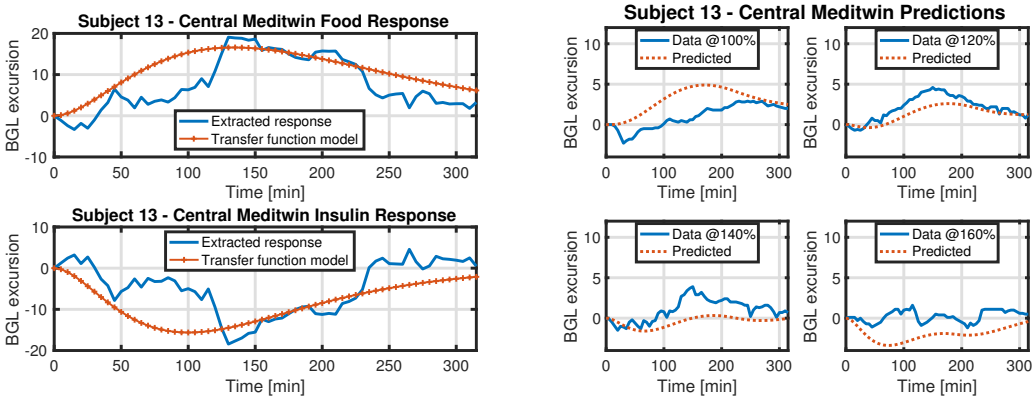

Figure A.25: Central Meditwin fitting and predictions for Subject 13. Tests having consistent data used for fitting: (2,3,4).

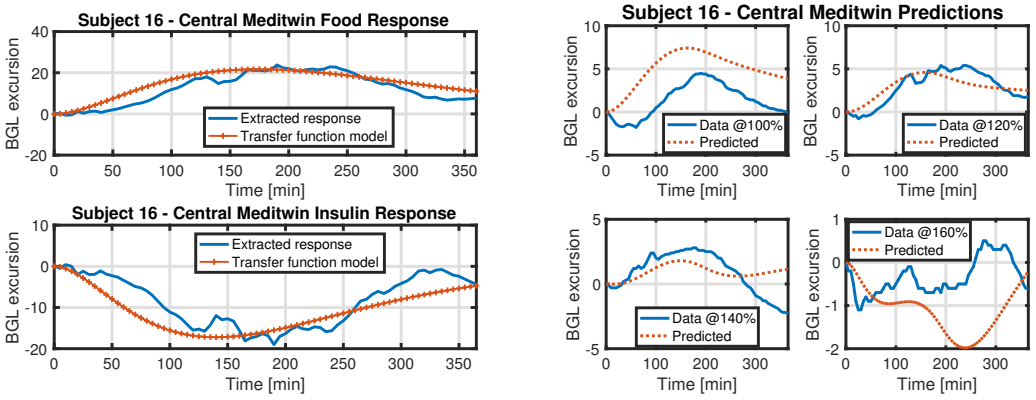

Figure A.26: Central Meditwin fitting and predictions for Subject 16. Tests having consistent data used for fitting: (2,3,4).

Appendix B. MDTE Validation for All Individuals

The following figures display the MDTE validation results for all individuals other than Subject 1.

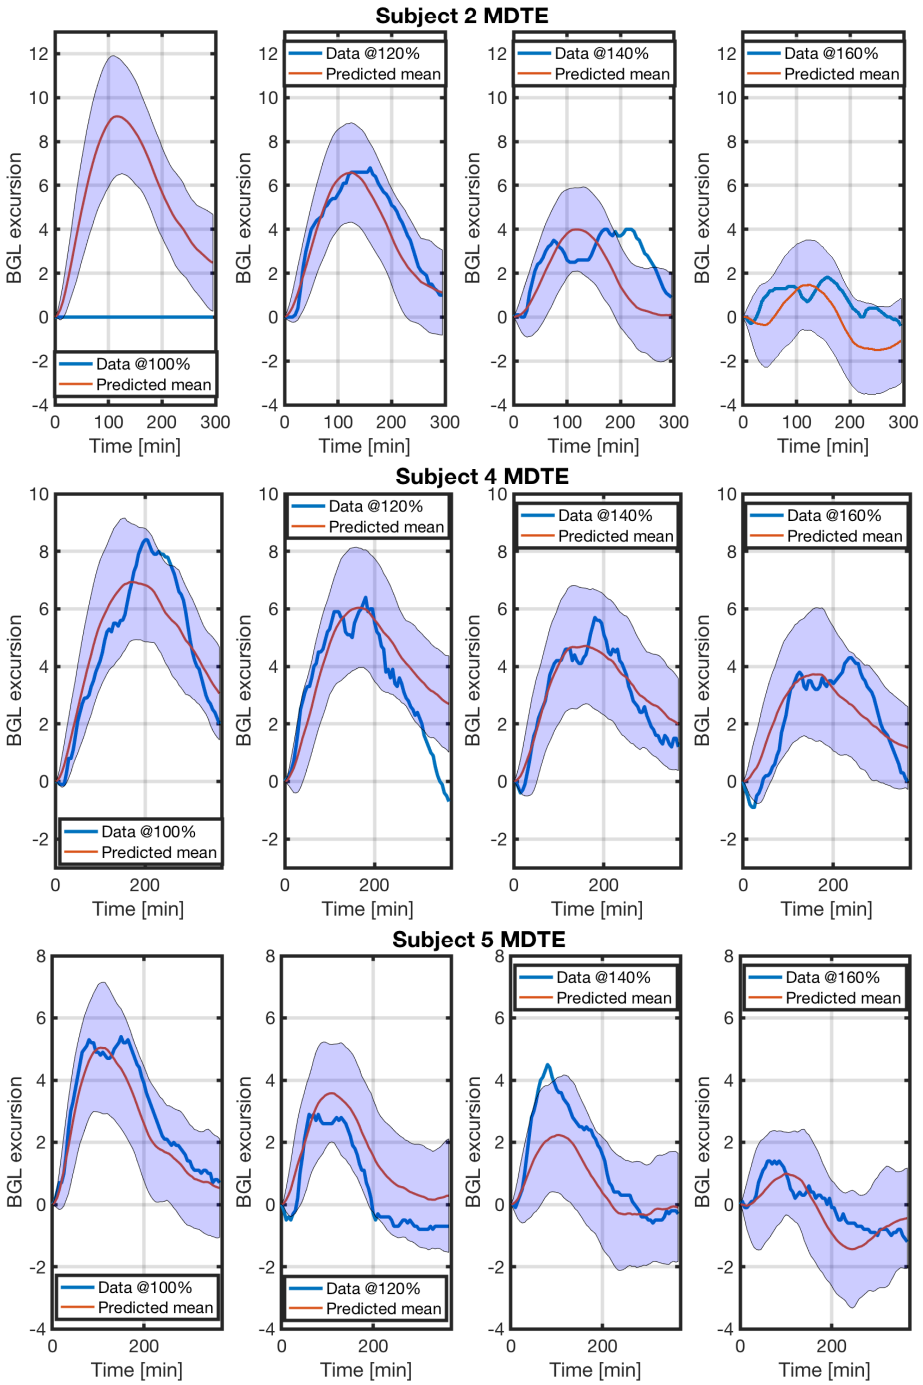

Figure B.27: MTDE validation for Subjects 2, 4 and 5. The blue shaded areas are the ± 1.5 standard-deviation envelopes around the MDTE mean prediction (plotted in red).

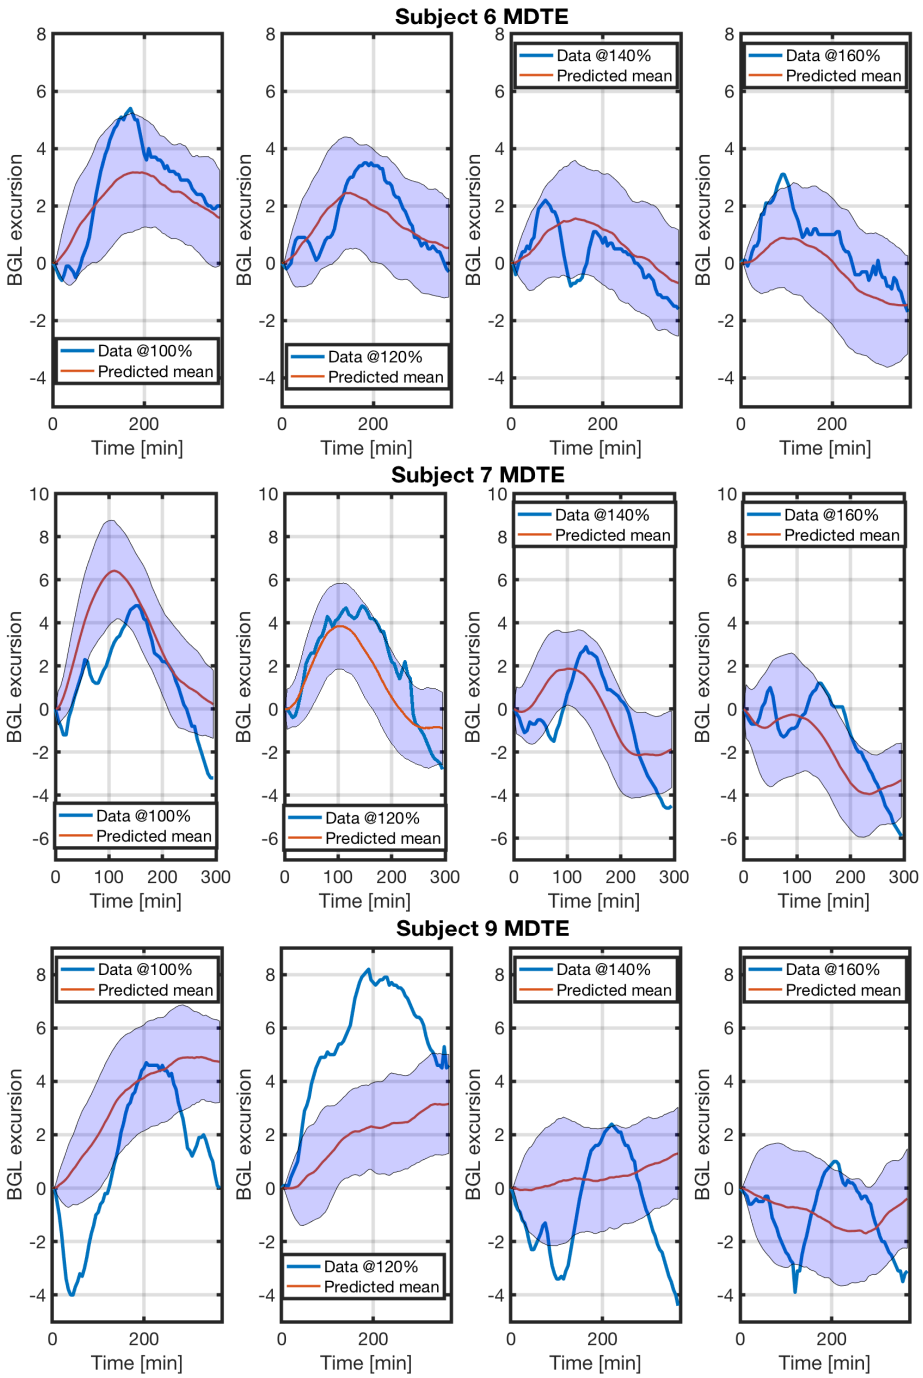

Figure B.28: MTDE validation for Subjects 6, 7 and 9. The blue shaded areas are the ± 1.5 standard-deviation envelopes around the MDTE mean prediction (plotted in red).

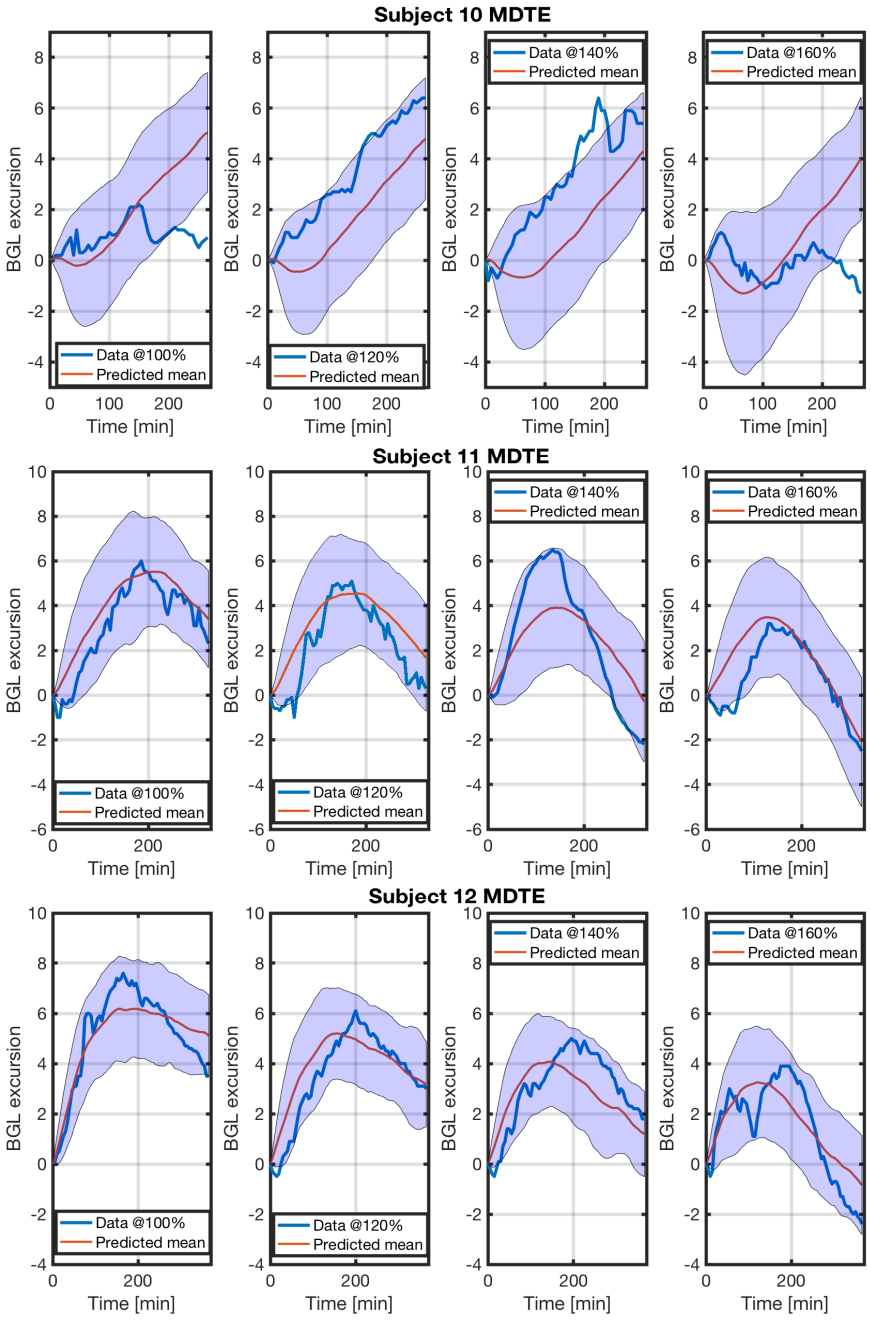

Figure B.29: MTDE validation for Subjects 10, 11 and 12. The blue shaded areas are the ± 1.5 standard-deviation envelopes around the MDTE mean prediction (plotted in red).

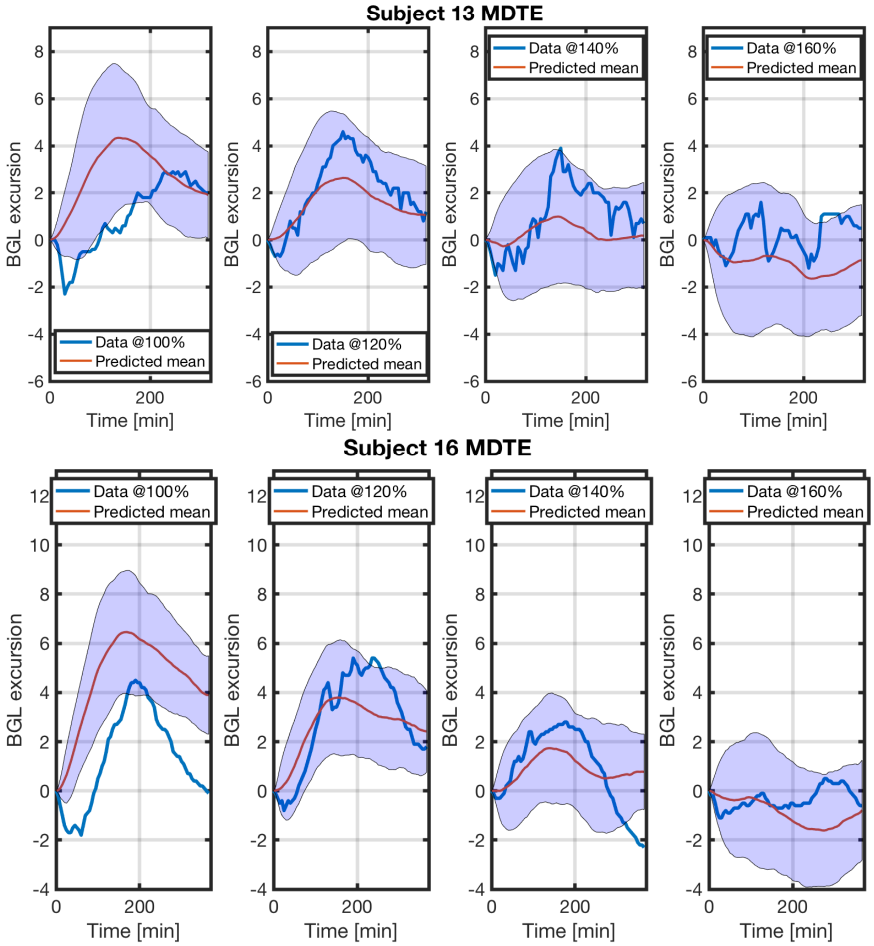

Figure B.30: MTDE validation for Subjects 13 and 63. The blue shaded areas are the ± 1.5 standard-deviation envelopes around the MDTE mean prediction (plotted in red).

Appendix C. Digital Trial Results for All Individuals

Figures C.31–C.40 show the BGL excursions resulting from the best dual wave, dual bolus, dual wave with fixed duration, and single bolus strategies for all remaining individuals in the trial.

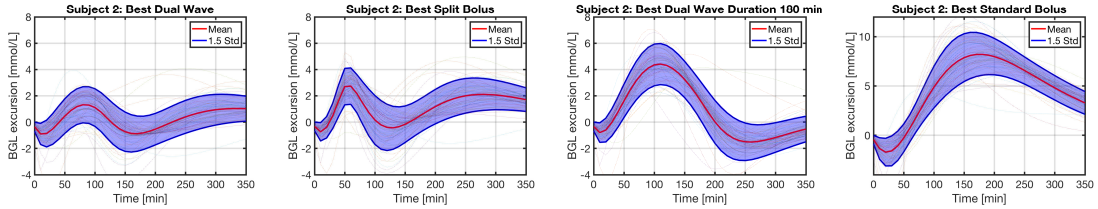

Figure C.31: Best dual wave, split bolus, dual wave of duration 180 min and standard bolus strategies for Subject 2.

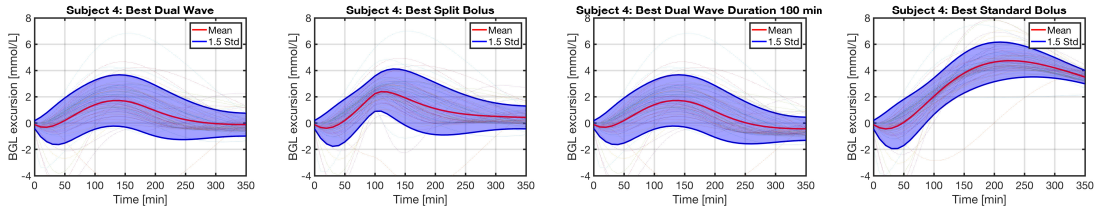

Figure C.32: Best dual wave, split bolus, dual wave of duration 180 min and standard bolus strategies for Subject 4.

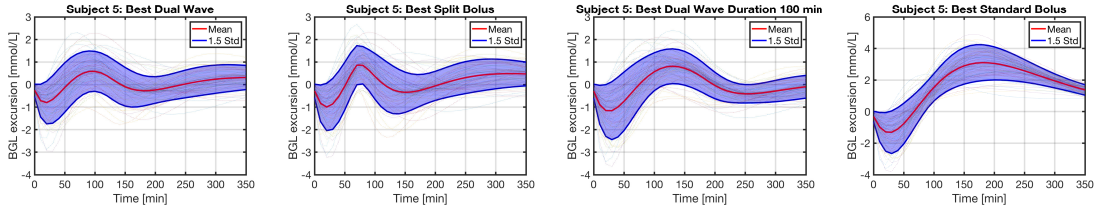

Figure C.33: Best dual wave, split bolus, dual wave of duration 180 min and standard bolus strategies for Subject 5.

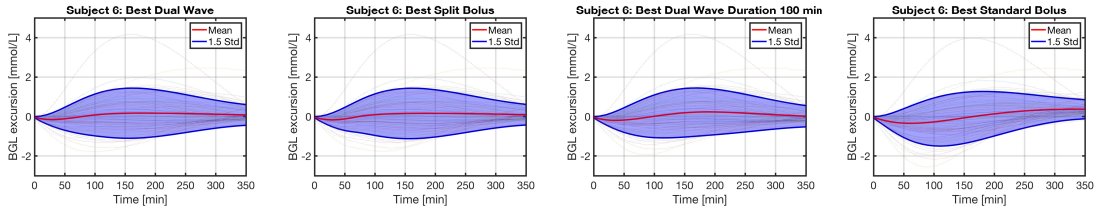

Figure C.34: Best dual wave, split bolus, dual wave of duration 180 min and standard bolus strategies for Subject 6.

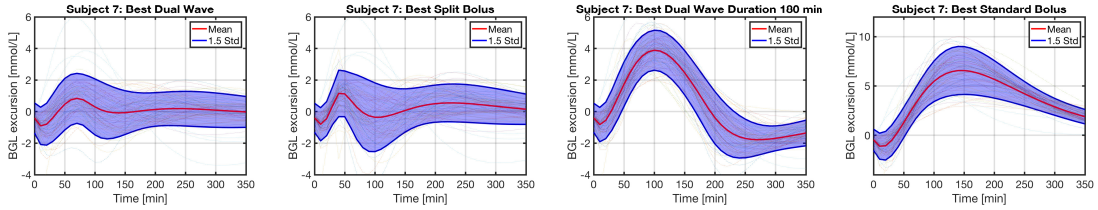

Figure C.35: Best dual wave, split bolus, dual wave of duration 180 min and standard bolus strategies for Subject 7.

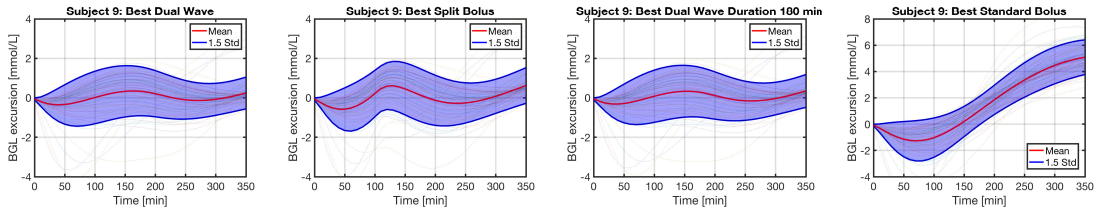

Figure C.36: Best dual wave, split bolus, dual wave of duration 180 min and standard bolus strategies for Subject 9.

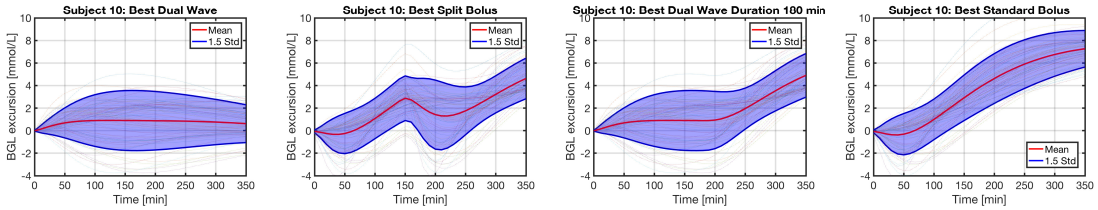

Figure C.37: Best dual wave, split bolus, dual wave of duration 180 min and standard bolus strategies for Subject 10.

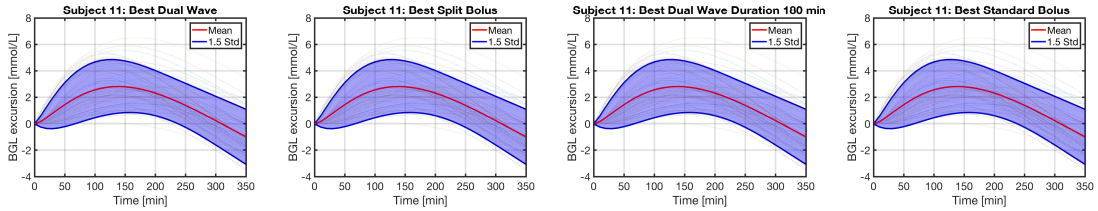

Figure C.38: Best dual wave, split bolus, dual wave of duration 180 min and standard bolus strategies for Subject 11. Note that all results are identical to the standard bolus, which is the best strategy for this subject and meal ingested.

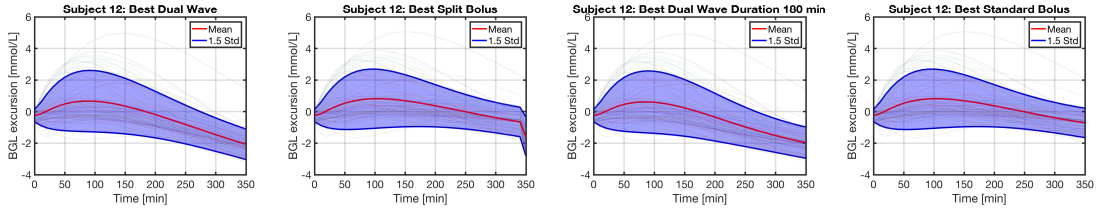

Figure C.39: Best dual wave, split bolus, dual wave of duration 180 min and standard bolus strategies for Subject 12.

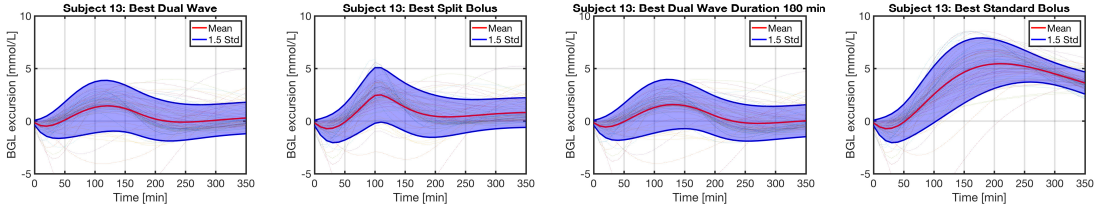

Figure C.40: Best dual wave, split bolus, dual wave of duration 180 min and standard bolus strategies for Subject 13.

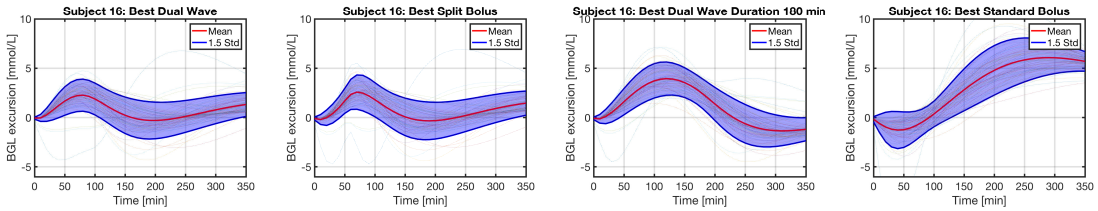

Figure C.41: Best dual wave, split bolus, dual wave of duration 180 min and standard bolus strategies for Subject 16.

References

[1] H. Zhang A. Geletu, M. Klöppel and P. Li. Advances and applications of chance-constrained approaches to systems optimisation under uncertainty. *International Journal of Systems Science*, 44(7):1209–1232, 2013.

[2] J.A. Paulson, E.A. Buehler, R.D. Braatz, and A. Mesbah. Stochastic model predictive control with joint chance constraints. *International Journal of Control*, 2017. Published online: 10 May 2017. https://doi.org/10.1080/00207179.2017.1323351.

[3] A. Shapiro, D. Dentcheva, and A. Ruszczynski. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM - Society for Industrial and Applied Mathematics, second edition, 2014.

[4] J. Lin, D. Lee, J.G. Chase, G.M. Shawand A. Le Compte, T. Lotz, J. Wong, T. Lonergan, and C.E. Hann. Stochastic modelling of insulin sensitivity and adaptive glycemic control for critical care. *Computer Methods and Programs in Biomedicine*, 89:141–152, 2008.

[5] L.M. Fisk, A.J. Le Compte, G.M. Shaw, S. Penning, T. Desaive, and J.G. Chase. STAR Development and Protocol Comparison. *IEEE Transactions on Biomedical Engineering*, 59(12):3357–3364, 2012.

[6] A. Evans, A. Le Compte, C.-S. Tan, L. Ward, J. Steel, C.G. Pretty, S. Penning, F. Suhaimi, G.M. Shaw, T. Desaive, and J.G. Chase. Stochastic targeted (STAR) glycemic control: design, safety, and performance. *Journal of Diabetes Science and Technology*, 6(1):102–115, 2012.

[7] V. Uyttendaele, J.L. Knopp, K.W. Stewart, T. Desaive, B. Benyó, N. Szabó-Némedi, A. Illyés, G.M. Shaw, and J.G. Chase. A 3D insulin sensitivity prediction model enables more patient-specific prediction and model-based glycaemic control. *Biomedical Signal Processing and Control*, 46:192–200, 2018.

[8] J. Lin, N.N. Razak, C.G. Pretty, A. Le Compte, P. Docherty, J.D. Parente, G.M. Shaw, C.E. Hann, and J.G. Chase. A physiological Intensive Control Insulin-Nutrition-Glucose (ICING) model validated in critically ill patients. *Computer Methods and Programs in Biomedicine*, 102(2):192–205, 2011.

[9] B. Paláncz, C.G. Pretty, K. Stewart, J.G. Chase, J. Homlok, and B. Benyó. Estimation of the insulin sensitivity profile for the stochastic variant of the ICING model. In *IEEE 20th Jubilee International Conference on Intelligent Engineering Systems (INES)*, Budapest, Hungary, 30 June – 2 July 2016.

[10] K. Bruynseels, F. Santoni de Sio, and J. van den Hoven. Digital twins in health care: Ethical implications of an emerging engineering paradigm. *Frontiers in Genetics*, 9(31), 2018.

[11] J.G. Chase, J.-C. Preiser, J.L. Dickson, A. Pironet, Y.S. Chiew, C.G. Pretty, G.M. Shaw, B. Benyó, K. Moeller, S. Safaei, M. Tawhai, P. Hunter, and T. Desaive. Next-generation, personalised, model-based critical care medicine: a state-of-the art review of in silico virtual patient models, methods, and cohorts, and how to validation them. *BioMedical Engineering OnLine*, pages 17–24, 2018.

[12] S. Oviedo, J. Vehí, R. Calm, and J. Armengol. A review of personalized blood glucose prediction strategies for T1DM patients. *International Journal for Numerical Methods in Biomedical Engineering*, 33(6), 2017. e2833.

[13] R. Hovorka, V. Canonico, L.J. Chassin, U. Haueter, M. Massi-Benedetti, M.O. Federici, T.R. Pieber, H.C. Schaller, L. Schaupp, T. Vering, et al. Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes. *Physiological Measurement*, 25(4):905, 2004.

[14] S.S. Kanderian, S. Weinzimer, G. Voskanyan, and G.M. Steil. Identification of intraday metabolic profiles during closed-loop glucose control in individuals with type 1 diabetes. *Journal of Diabetes Science and Technology*, 3(5):1047–57, 2009.

[15] P. Soru, G. De Nicolao, C. Toffanin, C. Dalla Man, C. Cobelli, L. Magni, and AP@home Consortium. MPC based artificial pancreas: Strategies for individualization and meal compensation. *Annual Reviews in Control*, 36(1):118–128, 2012.

[16] C. Toffanin, M. Messori, F. Di Palma, G. De Nicolao, C. Cobelli, and L. Magni. Artificial pancreas: Model predictive control design from clinical experience. *Journal of Diabetes Science and Technology*, 7(6):1470–1483, 2013.

[17] P.H. Colmegna and R.S. Sánchez Peña. Analysis of three T1DM simulation models for evaluating robust closed-loop controllers. *Computer Methods and Programs in Biomedicine*, 113:371–382, 2014.

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

[19] P.H. Colmegna, R.S. Sánchez-Peña, R. Gondhalekar, E. Dassau, and F.J. Doyle III. Switched LPV glucose control in Type 1 diabetes. *IEEE Trans. on Biomedical Engineering*, 63(6):1192–1200, 2016.

[20] E. Renard, A. Farret, J. Kropff, D. Bruttomesso, M. Messori, J. Place, R. Visentin, R. Calore, C. Toffanin, F. Di Palma, G. Lanzola, P. Magni, F. Boscari, S. Galasso, A. Avogaro, P. Keith-Hynes, B. Kovatchev, S. Del Favero, C. Cobelli, L. Magni, and J. Hans DeVries. Day-and-night closed-loop glucose control in patients with type 1 diabetes under free-living conditions: Results of a single-arm 1-month experience compared with a previously reported feasibility study of evening and night at home. *Diabetes care*, 39(7):1151–1160, 2016.

[21] L. Kovács, B. Benyó, J. Bokor, and Z. Benyó. Induced L_2-norm minimization of glucose-insulin system for Type I diabetic patients. *Computer Methods and Programs in Biomedicine*, 102:105–118, 2011.

[22] L. Kovács. Linear parameter varying (LPV) based robust control of type-I diabetes driven for real patient data. *Knowledge-Based Systems*, 122:199–213, 2017.

[23] R. Gondhalekar, E. Dassau, and F.J. Doyle III. Velocity-weighting & velocity-penalty MPC of an artificial pancreas: Improved safety & performance. *Automatica*, 91:105–117, 2018.

[24] P.H. Colmegna, R.S. Sánchez-Peña, and R. Gondhalekar. Linear parameter-varying model to design control laws for an artificial pancreas. *Biomedical Signal Processing and Control*, 40:204–213, 2018.

[25] C. Toffanin, R. Visentin, M. Messori, F. Di Palma, L. Magni, and C. Cobelli. Towards a run-to-run adaptive artificial pancreas: in silico results. *IEEE Transactions on Biomedical Engineering*, 65(3):479–488, 2018.

[26] C. Dalla Man, F. Micheletto, D. Lv, M. Breton, B. Kovatchev, and C. Cobelli. The UVA/Padova type 1 diabetes simulator new features. *Journal of Diabetes Science and Technology*, 8(1):26–34, 2014.

[27] R. Visentin, C. Dalla Man, and C. Cobelli. One-day Bayesian cloning of type 1 diabetes subjects: Toward a single-day UVA/Padova type 1 diabetes simulator. *IEEE Transactions on Biomedical Engineering*, 63(11):2416–2424, 2016.

[28] R. Visentin, E. Campos-Náñez, M. Schiavon, D. Lv, M. Vettoretti, M. Breton, B.P. Kovatchev, C. Dalla Man, and C. Cobelli. The UVA/Padova type 1 diabetes simulator goes from single meal to single day. *Journal of Diabetes Science and Technology*, 12(2):273–281, 2018.

[29] M.E. Wilinska and R. Hovorka. Simulation models for in silico testing of closed-loop glucose controllers in type 1 diabetes. *Drug Discovery Today: Disease Models*, 5(4):289–298, 2009.

[30] B. Kovatchev, M. Breton, C. Dalla Man, and C. Cobelli. In Silico Preclinical Trials: A Proof of Concept in Closed-Loop Control of Type 1 Diabetes. *Journal of diabetes science and technology*, 3(1):44–55, 2009.

[31] A. Haidar, M.E. Wilinska, J.A. Graveston, and R. Hovorka. Stochastic virtual population of subjects with type 1 diabetes for the assessment of closed-loop glucose controllers. *IEEE Transactions on Biomedical Engineering*, 60(12):3524–3533, 2013.

[32] N. Magdelaine, L. Chaillous, I. Guilhem, J.-Y. Poirier, M. Krempf, C.H. Moog, and E. Le Carpentier. A long-term model of the glucose-insulin dynamics of type 1 diabetes. *IEEE Transactions on Biomedical Engineering*, 62(6):1546–1552, 2015.

[33] Y. Ruan, M.E. Wilinska, H. Thabit, and R. Hovorka. Modeling day-to-day variability of glucose-insulin regulation over 12-week home use of closed-loop insulin delivery. *IEEE Transactions on Biomedical Engineering*, 64(6):1412–1419, 2017.

[34] G.C. Goodwin, M. Gevers, and B. Ninness. Quantifying the error in estimated transfer functions with application to model order selection. *IEEE Trans. Automatic Control*, 37(7):913–928, 1992.

[35] G.C. Goodwin, J.H. Braslavsky, and M.M. Seron. Non-stationary stochastic embedding for transfer function estimation. *Automatica*, 38(1):47–62, 2002.

[36] L. Ljung, G.C. Goodwin, J.C. Agüero, and T. Chen. Model error modeling and stochastic embedding. *IFAC-PapersOnLine,*, 48(28):75–79, 2015.

[37] P. Lopez, C. Smart, C. Morbey, P. McElduff, M. Paterson, and B.R. King. Extended insulin boluses cannot control postprandial glycemia as well as a standard bolus in children and adults using insulin pump therapy. *BMJ Open Diabetes Research and Care*, 2014. 2:e000050. doi:10.1136/bmjdrc-2014-000050.

[38] P. Lopez, C. Smart, P. McElduff, D.C. Foskett, D.A. Price, M. Paterson, and B. King. Optimizing the combination insulin bolus split for a high-fat, high-protein meal in children and adolescents using insulin pump therapy. *Diabetic Medicine*, 34(10):1380–1384, 2017.

[39] K.J. Bell, E. Toschi, G.M. Steil, and H.A. Wolpert. Optimized mealtime insulin dosing for fat and protein in type 1 diabetes: Application of a model-based approach to derive insulin doses for open-loop diabetes management. *Diabetes Care*, 39:1631–1634, 2016.

[40] H.V. Phan, G.C. Goodwin, D.S. Carrasco, A.M. Medioli, C. Stephen, and A. Feuer. Individualization of stochastic models from population statistics for blood glucose regulation in type 1 diabetes patients. In *Australian Control Conference*, 2016.

[41] C. Toffanin, S. Del Favero, E.M. Aiello, M. Messori, C. Cobelli, and L. Magni. Glucose-insulin model identified in free-living conditions for hypoglycaemia prevention. *Journal of Process Control*, 64:27–36, 2018.

[42] G.C. Goodwin and R.L. Payne. *Dynamic system identification: experiment design and data analysis*. New York: Academic Press, 1977.

[43] L. Ljung. *System Identification – Theory For the User*. Prentice Hall, Upper Saddle River, N.J., 2nd. edition, 1999.

[44] T.A. Johansen. On Tikhonov regularization, bias and variance in nonlinear system identification. *Automatica*, 33(3):441–446, 1997.

[45] T. Chen, H. Ohlsson, and L. Ljung. On the estimation of transfer functions, regularizations and Gaussian processes—revisited. *Automatica*, 48(8):1525–1535, 2012.

[46] Bias variance decomposition. In C. Sammut and G.I. Webb, editors, *Encyclopedia of Machine Learning*, pages 100–101. Springer US, Boston, MA, 2010.

[47] L. Ljung, T. Chen, and B. Mu. A shift in paradigm for system identification. *International Journal of Control*, 2019. DOI: 10.1080/00207179.2019.1578407.

[48] G.C. Goodwin, S.F. Graebe, and M.E. Salgado. *Control system design*. Prentice Hall Upper Saddle River, 2001.

[49] G.E. Umpierrez and D.C. Klonoff. Diabetes technology update: Use of insulin pumps and continuous glucose monitoring in the hospital. *Diabetes Care*, 41(8):1579–1589, 2018.

[50] G. Freckmann, M. Link, U. Kamecke, C. Haug, B. Baumgartner, and R. Weitgasser. Performance and usability of three systems for continuous glucose monitoring in direct comparison. *Journal of Diabetes Science and Technology*, 13(5):890–898, 2019.

[51] M. Breton and B. Kovatchev. Analysis, modeling, and simulation of the accuracy of continuous glucose sensors. *Journal of Diabetes Science and Technology*, 2(5):853–862, 2008.

[52] S. Del Favero, A. Facchinetti, G. Sparacino, and C. Cobelli. Improving accuracy and precision of glucose sensor profiles: Retrospective fitting by constrained deconvolution. *IEEE Transactions on Biomedical Engineering*, 61(4):1044–1053, 2014.

[53] B. Benyó, B. Paláncz, Á. Szlávecz, K. Stewart, J. Homlok, C.G. Pretty, and J.G. Chase. Analysis of stochastic noise of blood-glucose dynamics. *IFAC PapersOnLine*, 50(1):15157–15162, 2017.