Lecture Notes in Bioengineering

Harald Kirchsteiger
John Bagterp Jørgensen
Eric Renard
Luigi del Re *Editors*

Prediction Methods for Blood Glucose Concentration

Design, Use and Evaluation

Springer

Lecture Notes in Bioengineering

More information about this series at http://www.springer.com/series/11564

Harald Kirchsteiger · John Bagterp Jørgensen
Eric Renard · Luigi del Re
Editors

Prediction Methods for Blood Glucose Concentration

Design, Use and Evaluation

Springer

Editors
Harald Kirchsteiger
Institute for Design and Control of Mechatronical Systems
Johannes Kepler University Linz
Linz
Austria

John Bagterp Jørgensen
Department of Applied Mathematics
Technical University of Denmark
Kongens Lyngby
Denmark

Eric Renard
Institut de Génomique Fonctionnelle de Montpellier
Montpellier
France

Luigi del Re
Institute for Design and Control of Mechatronical Systems
Johannes Kepler University Linz
Linz
Austria

ISSN 2195-271X ISSN 2195-2728 (electronic)
Lecture Notes in Bioengineering
ISBN 978-3-319-25911-6 ISBN 978-3-319-25913-0 (eBook)
DOI 10.1007/978-3-319-25913-0

Library of Congress Control Number: 2015953261

Springer Cham Heidelberg New York Dordrecht London

Printed on acid-free paper

Springer International Publishing AG Switzerland is part of Springer Science+Business Media (www.springer.com)

Preface

Standard diabetes insulin therapy for type 1 diabetes and late stages of type 2 is based on the expected development of blood glucose (BG) both as a consequence of the metabolic glucose consumption as well as of meals and exogenous insulin intake. Traditionally, this is not done explicitly, but the insulin amount is chosen using factors that account for this expectation.

The increasing availability of more accurate continuous blood glucose measurement (CGM) systems is attracting much interest to the possibilities of explicit prediction of future BG values. Against this background, in 2014 a two-day workshop on the design, use and evaluation of prediction methods for blood glucose concentration was held at the Johannes Kepler University Linz, Austria. One intention of the workshop was to bring together experts working in various fields on the same topic, in order to shed light from different angles on the underlying problem of modeling the glucose insulin dynamics of type 1 diabetes patients. Among the international participants were continuous glucose monitoring developers, diabetologists, mathematicians and control engineers, both, from academia and industry. In total 18 talks were given followed by panel discussions which allowed to receive direct feedback from the point of view of different disciplines.

This book is based on the contributions of that workshop and is intended to convey an overview of the different aspects involved in the prediction. The individual chapters are based on the presentations given by the authors at the workshop but were written afterward which allowed to include the findings and conclusions of the various discussions and of course updates.

The chapter "Alternative Frameworks for Personalized Insulin–Glucose Models" by Harald Kirchsteiger et al. asks the question whether more and more detailed physiological descriptions of the glucose metabolism with an ever-increasing degree of sophistication and number of modeled phenomena are really what is needed for pushing the boundaries in glucose prediction for control. As an alternative, the chapter introduces two data-based approaches that focus not on the prediction of exact future blood glucose values, but rather on the prediction of changes in the patients' blood glucose range.

The chapter "Accuracy of BG Meters and CGM Systems: Possible Influence Factors for the Glucose Prediction Based on Tissue Glucose Concentrations" by Guido Freckmann et al. discusses performance metrics used to characterize the accuracy of continuous glucose measurement devices. This topic is highly relevant for prediction models since many of them rely on the data given by the continuous sensors which are previously calibrated with blood glucose meter measurements which are also subject to measurement errors. Inaccurate measurements will directly affect the performance of the corresponding predictions.

The chapter "CGM—How Good Is Good Enough?" by Michael Schoemaker and Christopher G. Parkin also tackles the problem of continuous glucose monitor performance evaluation. Several performance metrics used in different published studies are compared and their individual characteristics analyzed. The chapter reveals why the comparison of a sensor evaluated in two different clinical studies is not always straightforward.

The chapter "Can We Use Measurements to Classify Patients Suffering from Type 1 Diabetes into Subcategories and Does It Make Sense?" by Florian Reiterer et al. makes use of continuous time prediction models to describe the interaction between ingested carbohydrates, subcutaneously injected insulin, and continuously measured glucose concentration. The identified model parameters of 12 subjects were analyzed and statistically significant correlations between the parameters and patient characteristics such as weight and age could be found.

The chapter "Prevention of Severe Hypoglycemia by Continuous EEG Monitoring" by Claus Borg Juhl et al. shows how to use EEG signals to predict upcoming hypoglycemic situations in real-time by employing artificial neural networks. The results of a 30-day long clinical study with the implanted device and the developed algorithm are presented.

The chapter "Meta-Learning Based Blood Glucose Predictor for Diabetic Smartphone App" by Valeriya Naumova et al. demonstrates how a highly sophisticated glucose prediction model can be ported from a development language running on a PC to a format such that it can be used conveniently by the patients. A unique feature of the algorithm is its independence of any user input other than historic CGM data which is automatically transmitted from a CGM device. No parameter estimation nor prediction model individualization is required.

The chapter "Predicting Glycemia in Type 1 Diabetes Mellitus with Subspace-Based Linear Multistep Predictors" by Marzia Cescon et al. uses data-based methods to develop individualized prediction models. The model can be considered as a combination of physiological models to precompute the rate of appearance of injected insulin and ingested carbohydrates in the bloodstream and of data-based models to combine this information and compute predictions up to 120 min in the future. The results show the performance on data from 14 type 1 diabetes patients in a clinical trial.

The chapter "Empirical Representation of Blood Glucose Variability in a Compartmental Model" by Stephen D. Patek et al. shows a modeling technique designed to extract the information on the net effect of meals on the blood glucose concentration. By assuming that all major unexplained glycemic excursions can be

attributed to oral glucose ingestion, a meal vector is estimated which significantly improves the mathematical model. Results are shown on three patients during a clinical trial and on virtual patients where it is shown how the method can be used for adjustments of the basal insulin rate.

The chapter "Physiology-Based Interval Models: A Framework for Glucose Prediction Under Intra-patient Variability" by Jorge Bondia and Josep Vehi tries to cope with the large intrasubject variability by using the concept of interval predictions. Instead of predicting a single blood glucose value in the future, a whole solution envelope is determined. With the presented theory it can be guaranteed that the real value is always inside of the envelope and moreover the envelope is not conservative. The method is evaluated on a physiological diabetes model.

The chapter "Modeling and Prediction Using Stochastic Differential Equations" by Rune Juhl et al. considers uncertainty in the dynamics between different patients as well as within a patient by making use of stochastic differential equations. It is shown how the mixed effects modeling methodology can be applied such that the underlying information of several datasets from different patients is extracted to form the model.

The chapter "Uncertainties and Modeling Errors of Type 1 Diabetes Models" by Levente Kovács and Péter Szalay analyzes the effect of prediction model uncertainties on the control system during a design procedure involving the steps model reduction by elimination of state variables, state estimation using extended Kalman Filters and Sigma Point filters and linear parameter-varying control synthesis.

The chapter "Recent Results on Glucose–Insulin Predictions by Means of a State Observer for Time-Delay Systems" by Pasquale Palumbo et al. introduces a prediction model which in real time predicts the insulin concentration in blood which in turn is used in a control system. The method is tested in simulation on a time-delay system representing the glucose–insulin system.

The chapter "Performance Assessment of Model-Based Artificial Pancreas Control Systems" by Jianyuan Feng et al. makes use of prediction models to compute treatment advices. The novelty of the proposed algorithm consists in explicitly considering (among others) the model prediction error and model error elimination speed. A retuning of the advisory system is done in case the prediction model does not perform well. Results on 30 virtual patients show the performance of the control system.

We would like to thank all people involved in the process of writing this book: All authors for their individual contributions, all reviewers of the book chapters, Daniela Hummer for the entire organization of the workshop, Boris Tasevski for helping with the typesetting, Florian Reiterer for his help editing the book, as well as Oliver Jackson and Karin de Bie for the good cooperation with Springer.

Linz
August 2015

Harald Kirchsteiger
John Bagterp Jørgensen
Eric Renard
Luigi del Re

Contents

Alternative Frameworks for Personalized Insulin–Glucose Models

Harald Kirchsteiger, Hajrudin Efendic, Florian Reiterer and Luigi del Re

Abstract The description of the insulin–glucose metabolism has attracted much attention in the past decades, and several models based on physiology have been proposed. While these models provide a precious insight in the involved processes, they are seldom able to replicate and much less to predict the blood glucose (BG) value arising as a reaction of the metabolism of a specific patient to a given amount of insulin or food at a given time. Data-based models have proven to work better for prediction, but predicted and measured values tend to diverge strongly with increasing prediction horizon. Different approaches, for instance the use of vital signs, have been proposed to reduce the uncertainty, albeit with limited success. The key assumption hidden behind these methods is the existence of a single "correct" model disturbed by some stochastic phenomena. In this chapter, instead, we suggest using a different paradigm and to interpret uncertainty as an unknown part of the process. As a consequence, we are interested in models which yield a similar prediction performance for all measured data of a single patient, even if they do not yield a precise representation of any of them. This chapter summarizes two possible approaches to this end: interval models, which provide a suitable range; and probabilistic models, which provide the probability that the BG lies in predetermined ranges. Both approaches can be used in the framework of automated personalized insulin delivery, e.g., artificial pancreas or adaptive bolus calculators.

1 Introduction

Glucose is the main energy source for the human body, but in excessive amounts it can be detrimental as well. In healthy persons, the correct glucose level is regulated by the insulin–glucose metabolism. In diabetes patients, this control system does not work properly any more, due, roughly speaking, either to lack of endogenous insulin

H. Kirchsteiger · H. Efendic · F. Reiterer (✉) · L. del Re
Institute for Design and Control of Mechatronical Systems,
Johannes Kepler University Linz, 4040 Linz, Austria
e-mail: florian.reiterer@jku.at
URL: http://desreg.jku.at/

H. Kirchsteiger et al. (eds.), *Prediction Methods for Blood Glucose Concentration*,
Lecture Notes in Bioengineering, DOI 10.1007/978-3-319-25913-0_1

production or to a resistance to insulin action [23]. In such cases, the therapy of choice consists in providing the necessary insulin externally. However, the delay times associated with this procedure and the danger of removing too much glucose from the blood circulation—with possibly lethal consequences—make the determination of the correct insulin quantity quite difficult and delicate.

Against this background, there has been a sustained interest in describing the glucose–insulin metabolism by physiological models of increasing complexity [2, 3, 25, 44] . A model resulting from the cooperation between the Universities of Padova and Virginia [12] has been accepted in 2008 by the Food and Drug Administration (FDA) as a substitute for preclinical trials for certain insulin treatments. The model was recently extended with new features [11].

And, of course, there has been a large number of attempts to use the model information to determine the right quantities of external insulin needed, so to say to replace the missing control action, to "close the loop." Positive results have been reported for overnight operation [26], but also a reduction of variability for the daily use have been reported [33]. New and improved sensors have been introduced, faster insulin is entering the market, and all this nurtures the hope that more progress will come. Unfortunately, in spite of over 40 years of research [1], the results of this "artificial" or "virtual pancreas" are still not there where they should be and simple safety rules—e.g., avoiding insulin infusion during or near to hypoglycemia—seem to be able to offer the largest part of the benefits of closed-loop control in a much simpler way as well.

Against this background, it is natural to lean back for a moment and wonder whether we are asking the right questions. In this chapter, we are addressing this very issue by presenting two promising alternatives to modeling, namely interval models and probabilistic modeling techniques.

2 Alternatives for Modeling

The key assumption for modeling is the reproducibility of results. To some extent, of course, this is always true, for instance insulin does reduce the blood glucose (BG) concentration and carbohydrates (Carb) intake increases it. Many other effects are known but hard to quantify—e.g., some hormones increase it as well, muscular work reduces it—but there are many control loops in the human body which can lead to a complex response, e.g., in the case of prolonged physical work [21]. Other effects, like the circadian variation of sensitivities, are known as well [46].

Physiological models describe all these phenomena relying on deep understanding of physiology and frequently on very specific measurements, e.g., tracers [24, 42]. It is usually impossible to determine all parameters of such models with simple "external" measurements, e.g., taking only Carb and insulin administration and BG values. Indeed, it has been shown that the "external" behavior—the relation between insulin and Carb intake and BG—of these complex models can be approximated very well with very simple ones [17, 45] (see Fig. 1). This is the one

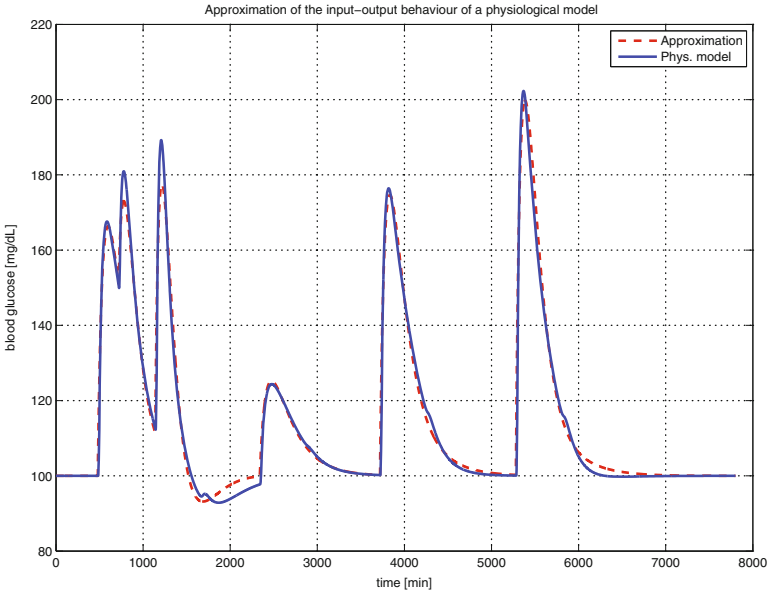

Fig. 1 Approximation of the input–output relationship of model [12] as shown in [45]

reason why comparisons between the values predicted by physiological models and measurements are very seldom, exceptions being for instance [8, 47].

Indeed, in general, physiological models are not able to provide a personalized description.

As the comparison makes clear, the measurements have a high degree of complexity not reflected in the physiological model, the main reason being the many unmodeled effects, e.g., related to the emotional state, which may affect very strongly the BG values, and which cannot be captured by the model because the critical quantities, in this case the concentration of some hormones, are not known.

There are several ways to cope with this problem. On one side, the attempt can be made to find additional measurements to extend the model, e.g., vital signs—acceleration, heart frequency, body temperature, and so on. Similar techniques have proved very useful in the industrial framework, e.g., to detect changes in machines [13], but have never really succeeded in the case of diabetes treatment.

Another approach, related to another chapter of this book (see "Empirical Representation of Blood Glucose Variability in a Compartmental Model" by S. Patek et al.) consists essentially in estimating a "corrective" Carb input to explain the difference between measured and computed values. While this method cannot be used in real time, it allows to study the effect of some changes in therapy, e.g., different amounts of insulin.

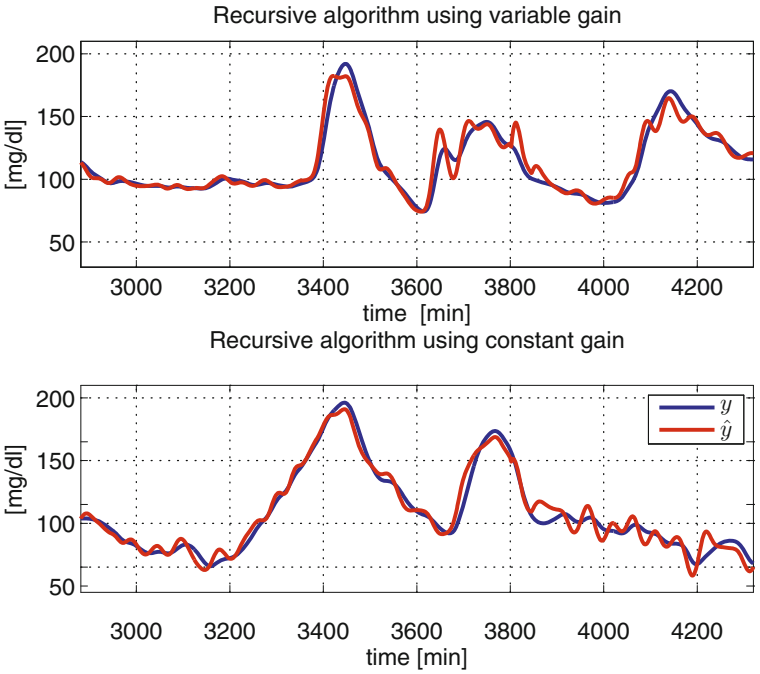

Fig. 2 45 min ahead prediction of glucose for two different patients from [6]

If we are interested in obtaining models which are sufficiently simple to allow their use and parameter estimation in real time, it might be better to look for other approaches. A very efficient adaptive model was developed by [5] which relies on a simple hypothesis, a so-called ARX model, and determines the parameters continuously, concentrating on the most critical BG ranges. Figure 2 shows the performance of such a model as predictor.

In this chapter, however, we suggest two different approaches. The key idea is not to get rid of uncertainty, assuming one particular value to be true, but to design models valid for the whole region, implicitly assuming that a full range of values are possible and in some sense true. One possible approach to this end are interval models, i.e., models which compute an output range and not a single value. The other alternative is using a probabilistic approach. Indeed, the exact BG value is not really important in itself, the clinician is more interested in keeping it inside the usual ("euglycemic") range and preventing to reach a dangerous one, e.g., hypoglycemia. Markov jump model can help in describing the physiology the way it really is—i.e., to some extent random.

Both models can be used for automated insulin delivery as well. In the case of interval models, the problem can be stated in terms of a min/max problem such that the carbohydrate amount optimizes the cost functions for all cases of uncertainty. In

the probabilistic case, it becomes the minimization of the probability that, under the action of insulin and meals, the BG leaves the good range to reach a dangerous one.

The further sections of this chapter are organized as follows: Before coming to the topic of interval models itself, a review about empirical continuous time transfer function models is given in Sect. 3 and the differences between possible model structures are discussed. The models presented in that section are a special form of control-oriented data-based models for describing the input–output relationship of the glucose metabolism that have proven quite powerful in the recent past (see e.g., [37, 38]). One of those presented model structures is then further used for the interval modeling introduced in Sect. 4. After a quick overview about the topic of interval modeling in general, some details about the methods used here for deriving interval models from data are given in Sect. 4.1. In following Sect. 4.1, results for the interval modeling are shown, both for simulated and for real patient data. The subsequent Sect. 5 then describes a probabilistic framework that can be used for predicting changes from one BG range to another. It starts in Sect. 5.1 with an overview about Gaussian mixture models that have been used in this context. Section 5.2 gives some details about the used model structure and the methodology of predicting transitions in the BG range, whereas actual prediction results for real patient data are presented in the following Sect. 5.3. The chapter finished with some final conclusions and discussion given in Sect. 6.

3 Model Structures

A model consists of a mathematical structure and of parameters. Thus the first step in modeling consists of fixing the model structure, and thereafter the parameters have to be tuned to get the best correspondence between measured and computed values. Of course, every model represents a simplification of the real system, and not every model structure is able to capture the behavior of the system under observation in a sufficiently general way. This is especially true in the case of a simplified model we are interested in.

Table 1 summarizes previously proposed model structures to describe the blood glucose dynamics where $BG(s)$, $Carb(s)$, and $I(s)$ correspond to the blood glucose concentration, ingested meal carbohydrates, and subcutaneously injected insulin bolus, respectively, all transformed into the Laplace domain. The table also lists the number of parameters which need to be estimated from data. All models thus use the same amount of information. A typical dataset which could be used for parameter estimation is shown in Fig. 3 (data from the DIAdvisor project [14]).

It is immediately visible that the model inputs are impulse-shaped quantities which are zero most of the time. That is because subcutaneous insulin injections are discrete events and meal ingestions, regardless of the quantity and time it takes to actually consume the food, are commonly treated as discrete events also. Therefore, for model analysis, the *impulse* responses are of great interest whereas *step* responses, commonly used in control engineering, do not provide sufficient information. A step

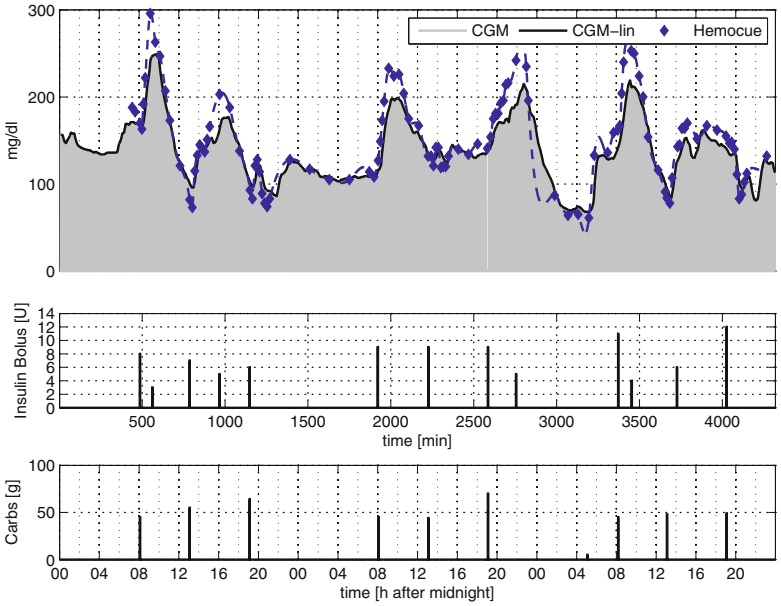

Fig. 3 Illustration of the relevant measurement data for a specific patient, CHU0102 (data from the DIAdvisor project [14])

Table 1 Selected model structures previously published

No.	Model structure	Parameters	Reference
1	$BG(s) = \frac{K_1}{(1+sT_1)^2 s} Carb(s) + \frac{K_2}{(1+sT_2)^2 s} I(s)$	4	[29]
2	$BG(s) = \frac{K_1 \exp(-\tau_1 s)}{(1+sT_1)s} Carb(s) + \frac{K_2 \exp(-\tau_2 s)}{(1+sT_2)s} I(s)$	6	[36]
3	$BG(s) = \frac{K_1}{(1+sT_1)s} Carb(s) + \frac{K_2}{(1+sT_2)s} I(s)$	4	[9]
4	$BG(s) = \frac{K_1}{(1+sT_1)^2} Carb(s) + \frac{K_2}{(1+sT_2)^2} I(s)$	4	[4][a]
5	$BG(s) = \frac{K_1 \exp(-\tau_1 s)}{(1+sT_1)(1+sT_2)} Carb(s) + \frac{K_2 \exp(-\tau_2 s)}{(1+sT_3)s} I(s)$	7	[10]

[a]Extended with a dynamic model for carbohydrates

response would actually mean that food or insulin is added to the metabolism in a continuous way over an extended period of time. While this might be true for the case of insulin, it is definitely an unrealistic case for food. In all given references, continuous insulin delivery is not considered for modeling.

From a physiological point of view, the impulse response gives precise information on the effect of one gram of carbohydrate and one unit of insulin, respectively, on the blood glucose concentration.

The impulse responses for both inputs of model structure 1 in Table 1 are shown in Fig. 4 for various selections of the parameters T_1 and T_2. It is immediately clear that the parameters K_1 and K_2 correspond to the steady-state change in BG. Furthermore, the parameters T_1 and T_2 are time constants which determine the time it takes until

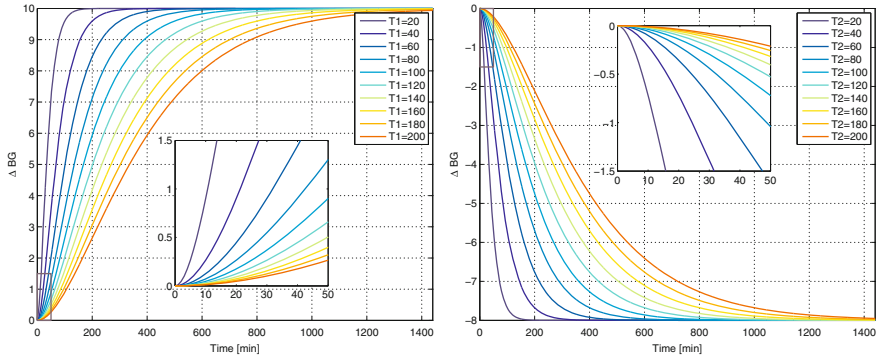

Fig. 4 Impulse response of model structure 1 from Table 1 using $K_1 = 10$, $K_2 = -8$

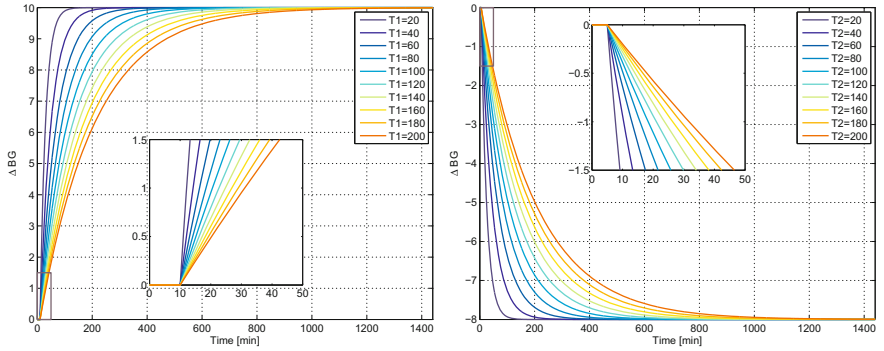

Fig. 5 Impulse response of model structure 2 from Table 1 using $K_1 = 10$, $\tau_1 = 10$, $K_2 = -8$, $\tau_2 = 5$

this steady state is reached. From this perspective, the physiological interpretation of the model parameters is straightforward.

The impulse responses for both inputs of model structure 2 in Table 1 are shown in Fig. 5 for various selections of the parameters T_1 and T_2. Compared to the impulse responses in Fig. 4, model structure 2 involves a time delay determined by the parameters τ_1 and τ_2. Furthermore, after this delay, the impulse response shows a discontinuity (see the magnified plots in Fig. 5) which hardly appears in real subjects. However, the overall responses of model structures 1 and 2 are similar.

The impulse responses for model structure 3 in Table 1 are not shown explicitly since they are very similar to those of model structure 2 shown in Fig. 5, except for the time delay which is zero in this case ($\tau_1 = 0$, $\tau_2 = 0$).

Model structure 4 of Table 1 gives a significantly different impulse response than the models discussed above, see Fig. 6. Since there is no integrating behavior, the impulse responses return to the steady state of zero. Physiologically, the assumption is that even in the absence of insulin, glucose will be removed from the circulation after meal ingestion. This is fundamentally different than the assumptions in the

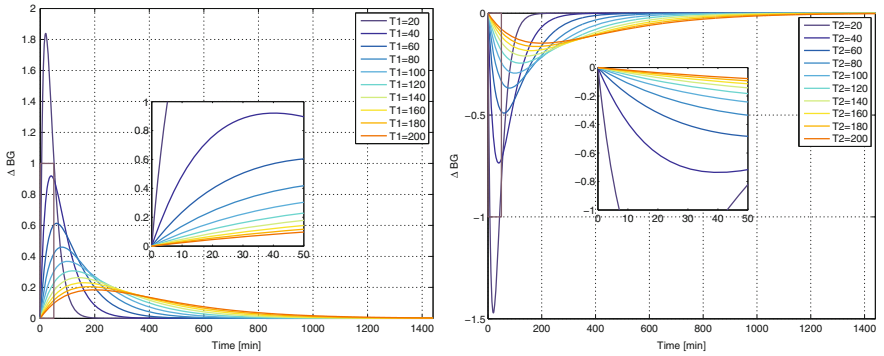

Fig. 6 Impulse response of model structure 4 from Table 1 using $K_1 = 100$, $K_2 = -80$

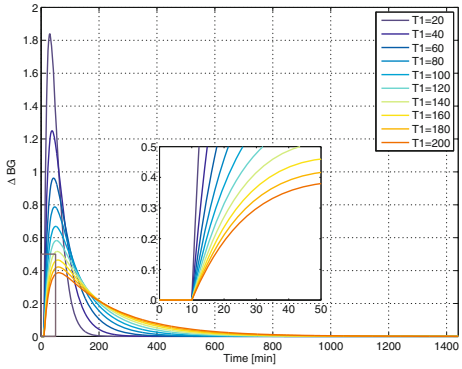

Fig. 7 Impulse response of model structure 5 from Table 1 using $K_1 = 100$, $T_2 = 20$, $\tau_1 = 10$

previous structures, where glucose will not be removed until insulin is supplied. As it can be seen in the experimental results in the corresponding papers, both model assumptions can lead to a good approximation of real data. Note also that model parameters are not directly visible from the impulse response.

Finally, the impulse response related to the carbohydrate input from model structure 5 in Table 1 is shown in Fig. 7. The response related to the insulin input is the same as for model structure 2 and can be seen in Fig. 5. The main difference in this structure is the assumption of two different time constants T_1, T_2, both for the carbohydrate dynamics. Therefore, there is more flexibility in the response compared to the case of using only one value for both constants ($T_1 = T_2$). However, it is questionable whether distinctive values for T_1 and T_2 can really be identified from clinical data [29].

4 Interval Models

There are essentially two ways to characterize interval models: one viewpoint is to design the model such that all possible realizations of the uncertainties will be considered and the true system output will be within the computed model output all the time, minimizing at the same time the conservativeness. The other viewpoint is to allow a certain amount or points to lie outside of the computed interval and thereby obtaining a much smaller interval [35]. The first approach is treated in another chapter of this book (see "Physiology-Based Interval Models: A Framework for Glucose Prediction Under Intra-Patient Variability" by J. Bondia and J. Vehi) while we focus on the latter here.

We will choose the model structure 1 from Table 1

$$BG(s) = \frac{K_1}{(1+sT_1)^2 s} Carb(s) + \frac{K_2}{(1+sT_2)^2 s} I(s) \tag{1}$$

because it contains only four parameters to be estimated and does not lead to a discontinuity in the impulse response. The generalized description of this model (assuming only one input for simplicity of notation) is a continuous time process model of the form

$$G(s) = \frac{B(s)}{A(s)} = \frac{b_0 + b_1 s + \cdots + b_m s^m}{a_0 + a_1 s + \cdots + s^n}. \tag{2}$$

Modeling a specific patient means to assign values to the variables a_i, b_i in (2). This is an identification problem [34, 43] and can be tackled in two ways. By transforming (2) into an equivalent discrete-time formulation, all the available tools of discrete-time system identification can be applied. However, such a transformation might introduce additional parameters and the initial (physiological) interpretation of the continuous time parameters gets lost [20]. On the opposite, there are continuous time identification methods which directly estimate the parameters in (2) without any need for transformation [19]. The benefits of such a direct estimation in the context of the human glucose insulin system have been treated in [10, 30]. In this contribution, we will thus also focus on a direct continuous time estimation method.

4.1 Continuous Time System Identification

Beginning first without taking into account the uncertainty, estimates of the model parameters can be found by minimizing a quadratic criterion of the form

$$J_1(\theta) = \frac{1}{N} \sum_{k=1}^{N} \varepsilon^2(k, \theta) \tag{3}$$

where ε denotes the difference between measurement and model output,

$$\varepsilon(t, \theta) = y(t) - \hat{y}(t, \theta) \tag{4}$$

N is the total number of available measurements and the vector θ contains all parameters. Various optimization techniques exist for actually minimizing the criterion (3) (see e.g., [43]).

Now, considering uncertainty in the system to be modeled we assume two separate datasets of the same system with different underlying dynamics. The cost function which is minimized now is the sum of two similar terms as above, again taking into account the deviation between model output and measurement. Additionally, we will introduce a third term which penalize the standard deviation of the estimated model parameters for the first dataset compared to the second dataset (5). For example, if this third term is not present, the model parameters for both datasets will be estimated independently and might deviate to a great extent from each other. By introducing and weighting (using the weighting matrix Π which has the tuning parameters in the main diagonal) the third term, a compromise between a good model fit of the individual datasets and compact parameter sets can be obtained.

$$J_2(\theta_1, \theta_2) = \frac{1}{N}\sum_{k=1}^{N} \varepsilon_1^2(k, \theta_1) + \frac{1}{N}\sum_{k=1}^{N} \varepsilon_2^2(k, \theta_2) + \frac{1}{N} \left\| \bar{\sigma}(\theta_1,\ \theta_2) \right\|_{\Pi}^2 \tag{5}$$

The function $\bar{\sigma}$ denotes a vector standard deviation operator which determines component-wise the standard deviation of the model parameters of the parameter vector θ.

Minimization can be done with a Gauss–Newton algorithm [43]. Thanks to the definition of the cost and the model structure, gradient vector and Hessian matrix can be determined analytically [30]. Nevertheless, the optimization is nonlinear and iterative and thus care has to be taken to avoid improper starting values for the parameters. A generalization of the cost function to N_{exp} datasets and the corresponding gradient vectors and the Hessian was derived in [30].

As a result of the continuous time system identification, there is one parameter vector per experiment. The next step is to extract parameter intervals which then define the model output interval. Considering again the model structure 1 from Table 1, there are four model parameters which means the parameter vector has the form

$$\theta = [K_1,\ K_2,\ T_1,\ T_2]^{\top} \tag{6}$$

Denoting with superscripts the corresponding dataset number, at the end of the identification N_{exp} estimates $K_1^1, K_1^2, \ldots K_1^{N_{exp}}$ are available and their maximum value is K_1^{max} and minimum value K_1^{min}. The computation of the interval model output is then

$$\hat{Y}^{\max}(s) = \frac{K_1^{\max}}{(1+T_1^{\min}s)^2 s}U^{(1)}(s) + \frac{K_2^{\max}}{(1+T_2^{\max}s)^2 s}U^{(2)}(s) \tag{7a}$$

$$\hat{Y}^{\min}(s) = \frac{K_1^{\min}}{(1+T_1^{\max}s)^2 s}U^{(1)}(s) + \frac{K_2^{\min}}{(1+T_2^{\min}s)^2 s}U^{(2)}(s) \tag{7b}$$

Note that K_2 describes the effect of insulin and is thus negative. In the following subsection, we will present results when applying the interval model estimation to short data segments representing the breakfast period. The starting point of the data is breakfast time (around 8:00 in the morning) and data end points just before lunch were taken (around 12:30).

4.2 Interval Model Results

Figure 8 shows a typical result obtained with the methodology described above for three independent measurement sets of a single patient. The black stars are the measured glucose levels using a Yellow Spring Instrument *YSI 2300 STAT Plus™ Glucose Analyzer* (YSI) device. When modeling each dataset independently (individual models), the green curves apply which show of course the best performance. Depending on the choice of the tuning matrix Π, the identified parameters then depend stronger or less strong on each other and an output interval according to (7) can be computed, see the red dashed lines in Fig. 8. Finally, also the mean interval model response is shown, which is computed when using as parameter the average values of all the estimated parameters of the interval model.

Remark 1 Again note the dependency of the results on the tuning matrix Π. When choosing $\Pi = 0$ a zero matrix, the corresponding term in the cost function (5) will disappear and the result is an independent estimation of the parameters for each experiment. This also results in the largest possible output interval. In the other extreme, choosing very large values in Π results in very small standard deviations of the model parameters, making them equivalent. This also results in an output interval which degenerates to a single line and is the same result as for a standard multiexperiment identification setup [34].

For an in-depth analysis of the method, it was applied to 10 simulated and 10 clinical datasets. The simulated data was obtained from a time-varying metabolic simulation model [28]. In the following, we will use the term individual model for a model estimated on a single dataset (experiment), the term interval model for a model according to (7), and the term mean interval model when a simulation with the mean value of the interval model parameters is done.

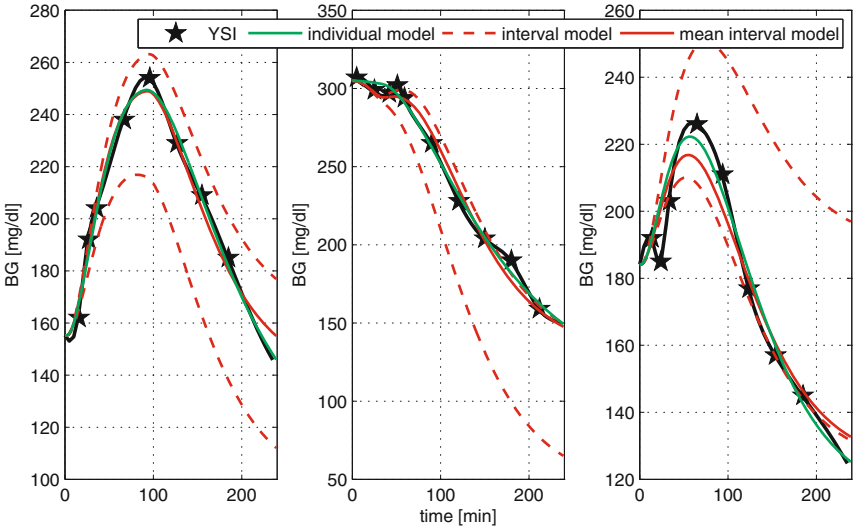

Fig. 8 Output response of the interval model compared to measurements for three datasets

4.2.1 Performance Metrics

Results will be evaluated based on a fit value (8) and on the error in detecting the peak BG. This is visualized in Fig. 9: ΔBG is the difference between the measured and predicted peak BG and ΔT is the difference in time. Both quantities can be positive or negative, where positive means the measured BG peak is higher and later in time than the simulated. The motivation of using those metrics is their importance in diabetes treatment.

$$\text{fit} = 100\left(1 - \frac{\|\hat{y} - y\|_2^2}{\|y - \bar{y}\|_2^2}\right) \tag{8}$$

4.2.2 Results Using the Simulated Data

First, we will present results from a single simulated patient before presenting statistics of all 10 considered patients. The parameter values and performance metrics reported in Table 2 are from simulated patient number 2. For the interval model, Π (5) was tuned in such a way, that the fit value of the interval model is not significantly less than 90 % of the fit of the individual model, and that the ratio between mean value and standard deviation $(1/CV)$ of each of the four parameters is at least five. In this way, it is ensured that the interval predictions are reasonably tight.

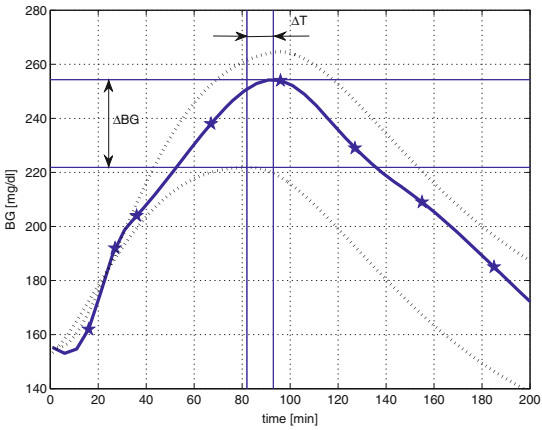

Fig. 9 Performance metrics $(\Delta BG, \Delta T)$ for evaluation of the results. *Blue stars* indicate BG measurements, interpolated with second-order polynomial. *Black dashed* are the min/max responses of the interval model (color figure online)

For the presented results in Table 2 and Fig. 10, the elements in the main diagonal of Π are [15000, 125, 20, 0.3] result in $1/CV = [5.02, -5.90, 5.94, 6.62]$. The parameter estimation was done on the first 3 days (up to $t = 72$ h). For the individual models, the validation is based on mean parameters of the first 3 days. For the interval models, the performance metrics were calculated based on the exact estimations for days 1–3 and validation was done with the mean interval model.

From Table 2, we see that there is a rather small ratio between mean value and standard deviation of the estimated parameters (especially for K_2)—which would mean that the insulin effect is very different—when the experiments are identified independently of each other. This ratio is greatly increased when applying the proposed interval model estimation, without decreasing the performance (on the three training days). Considering validation, the interval model shows better performance compared to the individual model. Note also that in this particular case, the individually estimated model for day 3 is not useful, because $K_2 < 0$, i.e., an insulin injection would result in a glucose increase.

The results for all 10 patients are summarized in Table 3 where the mean values of the performance metrics over all 5 days are shown. The interval model shows the highest fit values in every case, and the average error (for all 10 patients) of correct estimation of the BG peak is only 3.79 mg/dl compared to 10.80 mg/dl of the individual models. Note that this rather high error is to a large extent caused by the specific glucose responses to meals by the simulator for some patients, where the response does not have a single maximal value. In those cases there are two almost equivalent high peaks, but separated in time. Such a shape cannot be reproduced with the chosen model structure (1), and the approximation typically lies in the middle of the two peaks. It is of interest to note that from the 30 individual models, there were 8 that showed to be incorrect from a physiological point of view ($K_1 < 0$ and/or $K_2 > 0$). Considering the interval model, only for patient #6 no physiologically correct models could be found, as this patient's glucose dynamics have a high time

Table 2 Estimated parameters and performance metrics for simulated patient # 2

Individual model							
	Meal #1	Meal #2	Meal #3	Meal #4	Meal #5	Mean	Mean/Std.
Fit	66.54	66.85	87.66	12.73	42.02	55.16	1.92
ΔBG	2.96	0.56	1.35	−10.72	−11.06	5.33	1.03
ΔT	−7.00	−133.00	99.00	−129.00	−132.00	100.00	1.86
K_1	1.28(±0.08)	0.99(±0.05)	0.77(±0.04)	1.01	1.01	1.01	3.97
K_2	−13.83(±0.59)	−5.56(±0.41)	7.37(±0.48)	−4.01	−4.01	−4.01	−0.38
T_1	24.99(±1.33)	20.28(±0.91)	14.69(±0.60)	19.98	19.98	19.98	3.88
T_2	99.77(±12.02)	92.55(±16.48)	61.43(±4.25)	84.58	84.58	84.58	4.15
Interval model							
Fit	72.02	69.68	65.40	33.76	51.93	58.56	3.68
ΔBG	−0.13	1.51	−2.23	−7.20	−6.14	3.44	1.12
ΔT	−9.00	−121.00	−99.00	−114.00	−123.00	93.20	1.94
K_1	1.00	0.96	1.36	1.10	1.10	1.10	5.02
K_2	−16.19	−13.63	−11.53	−13.78	−13.78	−13.78	−5.90
T_1	19.00	21.16	26.26	22.14	22.14	22.14	5.94
T_2	211.86	253.75	287.56	251.05	251.05	251.05	6.62