Preprints of the 19th World Congress
The International Federation of Automatic Control
Cape Town, South Africa. August 24-29, 2014

Virtual Trials with *b-spline* Basis Functions and Stochastic Differential Equations

Liam M. Fisk * **Paul D. Docherty** * **Christopher G. Pretty** *
Jennifer L. Dickson * **Geoffrey M. Shaw** ** **J. Geoffrey Chase** *

* *Department of Mechanical Engineering, University of Canterbury, Christchurch, New Zealand (e-mail: liam.fisk@pg.canterbury.ac.nz, geoff.chase@canterbury.ac.nz).*
** *Department of Intensive Care, Christchurch School of Medicine and Health Sciences, New Zealand.*

Abstract: Virtual trials have proved useful in developing safe and efficacious glycaemic control protocols. However, these trials rely on lumping all changes in patient condition into the insulin sensitivity parameter. As electronic data collection provides higher temporal resolution than paper-based charts, irregular timings of both therapies and measurements clash with a regular, hourly insulin sensitivity profile. Additionally, unobservable endogenous changes are a factor for hour-to-hour variability. This research extends the virtual trial protocol to natively handle irregular data by regularising the insulin sensitivity profile, and utilising a simple stochastic differential equation. The insulin sensitivity profile was re-interpreted as a *b-spline* basis, allowing a higher order description with greater local support. The fitting error resulting from this regularisation was absorbed by a stochastic element in the glucose compartment, representing the hour-to-hour changes that cannot be attributed to changes in insulin sensitivity. The resulting virtual patients were demonstrated to be equivalent to the originals when a 0^{th} order basis was used. Inclusion of the stochastic element in this case simply ensured the model still fitted during periods of unmodelled high endogenous glucose production, while a 2^{nd} order basis uses this element to natively control the balance between changes in patient state and hour-to-hour unmodelled changes due to noise and endogenous processes. The resulting virtual trials are thus better able to preserve information in irregular data sets, and regulate the balance between controllable and uncontrollable glycaemic changes.

Keywords: Identification and validation; chronic care and/or diabetes; decision support and control.

1. INTRODUCTION

Virtual trials [Chase et al., 2010] are a valuable tool for *in-silico* development of control algorithms. However, validity of this methodology relies heavily on how model parameters are identified. Chase et al. describe a rigid parameter identification procedure that is not readily generalisable to more irregular data. This research describes a robust parameter identification and simulation procedure that features continuous parameter variation and native inclusion of internal noise.

Virtual trials were a key technique used to develop the STAR (Stochastic Targeted [Fisk et al., 2012]) and SPRINT (Specialised Relative Insulin and Nutrition Tables [Chase et al., 2008]) protocols *in-silico*. Avoiding physical trials during initial development enabled pre-informed protocols to be implemented in pilot studies, with the virtual trial results giving a high degree of confidence in safety and efficacy. Virtual trials also provide context analysing clinical trial results, as shown in [Chase et al., 2010]. An indicative measure of compliance and performance can be provided by the comparison between virtual trial results and true observations. Virtual trials are thus an valuable tool for model-based control design.

The key steps in a virtual trial are: A) fitting the underlying "true" parameter profile, B) using a protocol to choose the new treatment, and C) using the "true" parameter profile to solve for the resulting virtual blood glucose (BG) measurement. Correctly fitting the data is the critical step for a representative virtual trial, as failure to fit translates directly to loss of information and lesser virtual trial confidence. Due to low data density in a real-world glycaemic control setting only one parameter, insulin sensitivity, can be reliably identified.

Prior research developed the virtual trial methodology using the ICING (Intensive Control Insulin-Nutrition-Glucose [Lin et al., 2011]) model and integral-based fitting [Hann et al., 2005]. The available data was from SPRINT [Chase et al., 2008] in a summary spreadsheet with hourly slots for measurements leading to hourly-binned data. Consequently, fitting was carried out with an hourly piecewise-constant insulin sensitivity (S_I) profile, and linear interpolation to create intermediate BG estimates for 2-hourly intervals.

Implementation of the STAR protocol on a computerised tablet led to more precise timestamps recorded for measurements and therapy. With BG measurement times often

offset from the hour, and nutrition changes happening between measurements, the use of piecewise-constant S_I and interpolated measurements introduced new difficulties in creating virtual trials. A more physiological description of S_I was deemed possible and desirable.

It is possible to address the concerns created by the nature of the irregularly dataset by modifying the virtual trial approach. In this work, the S_I profile was made into a continuous function, and the error introduced regularising S_I was captured in an additional term appended to the ICING model. The glucose model was thus converted to a novel, simplistic stochastic differential equation (SDE [Van Kampen, 1976]).

2. METHODS

2.1 ICING model

The ICING model defines glucose and insulin kinetics and dynamics in critically ill patients:

$$\dot{I} = \frac{U_x + (1-x_L)U_n}{V_I} - n_I(I-Q) - n_K I - n_L \frac{I}{1+\alpha_G I} \quad (1)$$

$$\dot{Q} = n_I(I-Q) - n_C \frac{Q}{1-\alpha_G Q} \quad (2)$$

$$\dot{G}(t) = -p_G G - S_I(t)\frac{G(t)Q(t)}{1-\alpha_G Q(t)} + \frac{P_{EGP} - P_{CNS} + P_N(t) + \max\left(P_{max}, d_2 P_2(t)\right)}{V_G} \quad (3)$$

$$P_2(t) = d_1 P_1(t) - \max\left(P_{max}, d_2 P_2(t)\right) \quad (4)$$

$$P_1(t) = P(t) - d_1 P_1(t) \quad (5)$$

where $G(t)$ [mmol.L^{-1}] is the total plasma glucose, $I(t)$ [mU.L^{-1}] is the plasma insulin, and interstitial insulin is represented by Q(t) [mU.L^{-1}]. Exogenous insulin input is represented by $U_{ex}(t)$ [mU.min^{-1}], and glucose-dependent endogenous insulin production is estimated with U_{en} [mU.min^{-1}] [Pretty, 2012]. $S_I(t)$ [L.mU^{-1}.min^{-1}] is the identified insulin sensitivity profile, $P_1(t)$ [mmol] represents the glucose in the stomach and $P_2(t)$ [mmol] represents glucose in the gut. Enteral glucose input is denoted $P(t)$ [mmol.min^{-1}], while parenteral glucose input is denoted $P_N(t)$ [mmol.min^{-1}]. All model constants are shown in Table 1.

2.2 Parameter Identification

A simple version of the integral-based fitting method will be introduced via a differential equation with an unknown linear parameter, θ:

$$\dot{x} = f_0(t,x) + \theta f_1(t,x) \quad (6)$$

where x is the conserved quantity, θ the unknown constant parameter, $f_1(t,x)$ is the function corresponding to this parameter, and $f_0(t,x)$ contains the remaining known parameters and functions. The model estimate of x at time t is:

$$x_{mod}(t) = x_0 + \int_{t_0}^{t} f_0 \mathrm{d}t + \theta \int_{t_0}^{t} f_1 \mathrm{d}t \quad (7)$$

Assuming multiple data points, a residual error (ϵ) will occur. The error at the i^{th} measurement, $\epsilon_i = x_{mod}(t_i) - x(t_i)$, is:

$$\epsilon_i = -x(t_i) + x_0 + \int_{t_0}^{t_i} f_0 \mathrm{d}t + \theta \int_{t_0}^{t_i} f_1 \mathrm{d}t \quad (8)$$

Table 1. ICING constant model parameters

Variable	Description	Value
p_G	Non-insulin mediated uptake	0.006 min^{-1}
n_I	Insulin transport rate	0.006 min^{-1}
n_K	Renal clearance	0.0542 min^{-1}
n_L	Hepatic clearance	0.1578 min^{-1}
n_C	Interstitial clearance	0.006 min^{-1}
d_1	Stomach clearance	0.0151 min^{-1}
d_1	Gut clearance	0.00301 min^{-1}
P_{max}	Maximal gut cl.	6.11 mmol.min^{-1}
P_{EGP}	Endogenous glucose production	1.16 mmol.min^{-1}
P_{CNS}	Nervous system glucose disposal	0.3 mmol.min^{-1}
x_L	First-pass hepatic extraction	0.67
V_I	Insulin volume of distribution	4.0 L
V_G	Glucose volume of distribution	13.3 L
α_I	Saturation of hepatic insulin clearance	0.0017 L.mU^{-1}
α_G	Saturation of insulin-mediated glucose uptake	0.01538 L.mU^{-1}

which, provided the integrals can be numerically estimated, can be minimised using least squares for n measurements:

$$\begin{bmatrix} \int_{t_0}^{t_1} f_1 \mathrm{d}t \\ \vdots \\ \int_{t_0}^{t_n} f_1 \mathrm{d}t \end{bmatrix} \theta = \begin{bmatrix} x(t_1) - x_0 - \int_{t_0}^{t_1} f_0 \mathrm{d}t \\ \vdots \\ x(t_n) - x_0 - \int_{t_0}^{t_n} f_0 \mathrm{d}t \end{bmatrix} \quad (9)$$

If measured data is dense enough, or appropriate assumptions used, these integrals can be fully formed without further computation, and the linear system solved directly [Hann et al., 2005]. If sparse data or noise causes these integrals to be poorly approximated by the available data, this approach can be applied iteratively, where new parameters produce a modified solution, which is then used to update the integrals [Docherty et al., 2012].

As SPRINT data consisted of intermittent 1 and 2 hour measurement intervals, a distinct parameter value was fitted for each hour interval. Thus, θ was described as m piecewise-constant functions, where:

$$\theta(t) = \sum_{j=1}^{m} \gamma_j g_j(t) \quad (10)$$

where γ_j is the j^{th} θ value and $g_j(t)$ is a rectangle function, non-zero on a single hour interval. This description of θ is equivalent to a 0^{th} order uniform *b-spline* basis [De Boor, 1972] with $m+1$ knots, each an hour apart. Equation (7) thus expands to:

$$x_{mod}(t) - x_0 = \int_{t_0}^{t} f_0 \mathrm{d}t + \gamma_1 \int_{t_0}^{t} g_1 f_1 \mathrm{d}t + \ldots + \gamma_m \int_{t_0}^{t} g_m f_1 \mathrm{d}t \quad (11)$$

Accordingly, Equation (9) becomes:

$$\hat{A}_{1,(n,m)} \hat{\Gamma}_{(m,1)} = \hat{b}_{1,(n,1)} \quad (12)$$

where $\hat{\Gamma}_{(m,1)} = [\gamma_1, \ldots, \gamma_m]^\top$, and:

$$\hat{A}_{1,(n,m)} = \begin{bmatrix} \int_{t_0}^{t_1} g_1 f_1 \mathrm{d}t & \dots & \int_{t_0}^{t_1} g_m f_1 \mathrm{d}t \\ \vdots & \ddots & \vdots \\ \int_{t_0}^{t_n} g_1 f_1 \mathrm{d}t & \dots & \int_{t_0}^{t_n} g_m f_1 \mathrm{d}t \end{bmatrix} \quad (13)$$

$$\hat{b}_{1,(n,1)} = \begin{bmatrix} x(t_1) - x_0 - \int_{t_0}^{t_1} f_0 \mathrm{d}t \\ \vdots \\ x(t_n) - x_0 - \int_{t_0}^{t_n} f_0 \mathrm{d}t \end{bmatrix} \quad (14)$$

If $m + 1 > n$, the linear system is clearly indeterminate. Such a case almost always occurs with SPRINT data, as 2 hour BG measurements were common [Chase et al., 2008]. To circumvent this issue, SPRINT data was re-sampled hourly, an assumption that introduces fitting error if measurements are offset and can force unusual parameter spikes if additional glucose is added parenterally close to a re-sampled measurement. A continuous profile would therefore be beneficial, and controlling the order of the basis and knot locations provides a natural method for regularising the shape of the S_I profile. However, using a knot at each measurement forces the shape of the S_I profile to reach a maxima/minima in the middle of the measurement period, as well as forcing all changes to be the direct result of changes in S_I.

2.3 S_I profile

A distinction should be made in the use of this model for control and for virtual trials. In control, the model is refitted over a 6 hour period to ensure the initial conditions for the past hour are insensitive to modified data, and thus can be relied on for prediction. Previously, all dynamic changes were lumped into the S_I profile. This raw S_I profile showed a number of clinically important trends [Chase et al., 2011, Pretty et al., 2012, Ferenci et al., 2013] highlighting the importance of S_I. Clearly, the model-based metric evolves over time, and sudden rises greatly increase the risk of hypoglycaemia when patients are undergoing insulin therapy. Sudden increases in S_I predispose patients to hypoglycaemia, as injected insulin has an amplified effect on BG, and insulin dose is selected based on the prior (reduced) effect of insulin.

Variability in S_I exists due to changes in patient state, measurement noise, and mismodelled dynamics. Some literature suggests the major effect of intensive insulin therapy is the suppression of hepatic glucose production [Thorell et al., 2004], though the relationship between plasma insulin and endogenous glucose production is poorly understood. Pulsatile delivery, intraportal concentrations, and arteriovenous concentration gradients are all thought to have an effect, amongst other factors. However, quantifying the relationship between the liver and insulin is impossible using the data available at the bedside in critical care. It is sufficient to say that the relationship changes, and thus not all changes in patient state can truly be labelled "insulin-dependent".

This research regularises the S_I profile to restrict the frequency of "insulin-dependent" state changes by utilising

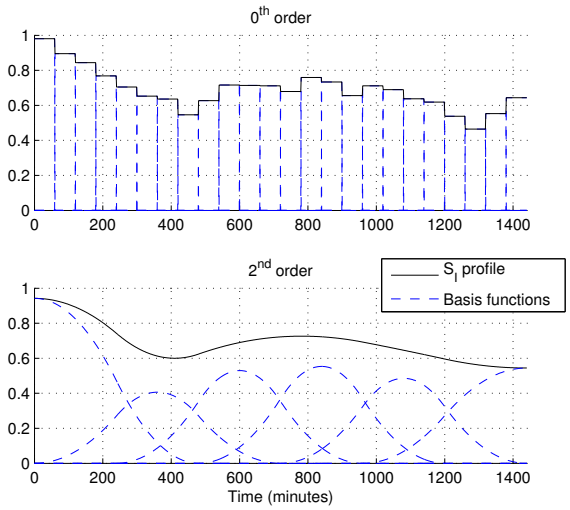

Fig. 1. *Comparison of the current 0^{th} order basis (60 minute knot widths) with a proposed 2^{nd} order basis (240 minute knot widths).*

the generalisable description of a b-spline [De Boor, 1972] basis. Figure 1 compares the current 0^{th} order basis with 60 minute knot widths with a proposed 2^{nd} order basis with 240 minute knot widths. A knot width of 240 minutes was chosen for the proposed description as 180 minutes is a typical measurement interval for STAR, and timing errors up to 60 minutes can occur. Thus, BG changes for between 1 to 3 measurement intervals are described by a single function. The local support provided by a 2^{nd} order description further regularises the profile.

2.4 Noise profile

Regularising the S_I profile highlights a bias vs. variance tradeoff, as multiple measures per function introduces fitting error with real data. As the S_I profile becomes smoother, greater error is introduced. Thus, additional fitting is required to prevent information being lost in the generation of virtual patients. Such information loss would result in virtual trials showing an artificial improvement. Returning to Equation (6), a zero-mean internal noise (process noise, $\phi(t)$) can be added:

$$\dot{x} = \phi(t) + f_0(t, x) + \theta f_1(t, x) \quad (15)$$

Typically, $\phi(t)$ would take the form of a wiener process. However, the added computational intensity associated with the increased resolution, and non-deterministic forward simulation, is not necessary in this application. In this simple SDE, $\phi(t)$ becomes the integral of a wiener process between two measurements, and captures unmodelled dynamics and measurement noise that cannot be incorporated by the now-regularised S_I profile. Thus, $\phi(t)$ is a piecewise-constant function, individual values of which can be fitted using:

$$\begin{bmatrix} \hat{A}_{1,(n,m)} & \hat{A}_{2,(n,n)} \end{bmatrix} \begin{bmatrix} \hat{\Gamma}_{(m,1)} \\ \hat{\Phi}_{(n,1)} \end{bmatrix} = \hat{b}_{1,(n,1)} \quad (16)$$

where $\hat{\Phi}_{(n,1)} = [\phi_1, \dots, \phi_n]^\top$ (ϕ_i is the i^{th} value of the piecewise constant $\phi(t)$), and:

$$\hat{A}_{2,(n,n)} = \begin{bmatrix} -(t_1-t_0) & 0 & \dots & 0 \\ -(t_1-t_0) & -(t_2-t_1) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ -(t_1-t_0) & -(t_2-t_1) & \dots & -(t_n-t_{n-1}) \end{bmatrix} \tag{17}$$

As $n+m > n$ when $m > 0$ this system is always rank-deficient. However, the definition of $\phi(t)$ as zero-mean noise can be utilised to fill the rank of the system. As zero-mean noise, $\int_{t_0}^{t_n} \phi(t)\mathrm{d}t = 0$, and by the definition of the basis functions, $\sum_{i=1}^{m} g_i(t) = 1 \ \forall$ t. Thus, $\phi(t) \equiv \sum_{i=1}^{m} \phi(t) g_i(t)$. If zero-mean is enforced for each component $\phi(t)g_i(t)$:

$$\int_{t_0}^{t_n} \phi(t) g_i(t)\mathrm{d}t = 0 \ \ \forall \ \ i = 1, \dots, m \tag{18}$$

As $\phi(t)$ is constant between two measurements:

$$\phi_1 \int_{t_0}^{t_1} g_i(t)\mathrm{d}t + ... + \phi_n \int_{t_{n-1}}^{t_n} g_i(t)\mathrm{d}t = 0 \tag{19}$$

Thus, the system in Equation (16) can be made full rank, and can be solved for all variables:

$$\begin{bmatrix} \hat{A}_{1,(n,m)} & \hat{A}_{2,(n,n)} \\ \hat{0}_{(m,m)} & \hat{A}_{3,(m,n)} \end{bmatrix} \begin{bmatrix} \hat{\Gamma}_{(m,1)} \\ \hat{\Phi}_{(n,1)} \end{bmatrix} = \begin{bmatrix} \hat{b}_{1,(n,1)} \\ \hat{0}_{(m,1)} \end{bmatrix} \tag{20}$$

where:

$$\hat{A}_{3,(m,n)} = \begin{bmatrix} \int_{t_0}^{t_1} g_1(t)\mathrm{d}t & \dots & \int_{t_{n-1}}^{t_n} g_1(t)\mathrm{d}t \\ \vdots & \ddots & \vdots \\ \int_{t_0}^{t_1} g_m(t)\mathrm{d}t & \dots & \int_{t_{n-1}}^{t_n} g_m(t)\mathrm{d}t \end{bmatrix} \tag{21}$$

The new form of the ICING model is therefore:

$$\begin{aligned} \dot{G}(t) &= G_x(t) - p_G G - S_I(t)\frac{G(t)Q(t)}{1-\alpha_G Q(t)} \\ &+ \frac{P_{EGP} - P_{CNS} + P_N(t) + \max\left(P_{max}, d_2 P_2(t)\right)}{V_G} \end{aligned} \tag{22}$$

where $S_I(t) = \theta(t)$ consists of *b-spline* basis functions, and $G_x(t) = \phi(t)$ is the new stochastic element. During a virtual trial, $G_x(t)$ is treated as the observed realisation of the stochastic process, and used in conjunction with $S_I(t)$ to calculate deterministic outcomes to modified therapeutic inputs.

2.5 STAR Cohort

BG, insulin, and nutrition data was collected as part of routine use of the STAR protocol in Christchurch Hospital medical and surgical ICU between July 2011 and February 2013. For use in this study, datasets were split when a gap greater that 5 hours occurred between consecutive BG measurements. Observational ethics was granted by the National Ethics Advisory Committee (New Zealand). Available cohort details are shown in Table 2.

2.6 Analyses

The intent of this research is to develop and implement a new fitting method to be used for virtual trials. Accordingly, two main categories of analyses were carried out. The first analysis was intended to demonstrate the problem solved by the research, and thus justify the added complexity. The second analysis was intended to give an understanding of the behaviour of the new methodology. For the iterative fits, convergence was assumed if maximum RMS error was within 0.1 mmol.L^{-1} (minimum resolution of a glucometer). To speed iterations, the G_x profile was updated between iterations by using the current S_I profile.

Table 2. STAR cohort details, presented as median [interquartile range] when appropriate.

	Cohort details
Episodes	207
Total hours	11538
Total BG measures	6517
Age (years)	61 [48 - 71]
Sex	66.2% male, 33.8% female
Length of episode (hours)	32 [15 - 68]
BG measures	19 [10 - 39]

Initially, STAR data was fitted using the standard methodology yielding an hourly piecewise-constant S_I profile fitted non-iteratively to linearly-interpolated hourly BG measurements. The fit was analysed with RMS error between BG measurements and modelled plasma glucose, to quantify how well a virtual trial would reproduce the original dataset. Fitting error in any patient indicates a virtual trial run using this patient could not recreate the original observed data, and thus error is an important metric.

The closest analogue of this method was then fitted using the basis function SDE approach. An hourly piecewise-constant profile was not possible to fit without introducing further constraints, and so constant S_I and G_x values were fitted for each measurement interval using the presented methodology. RMS error in BG was presented for comparison with the original, and the G_x distribution was presented as a cumulative density function (CDF). Finally, the regularised S_I profile was fitted and error in BG was calculated.

The nature of the G_x profile was then investigated graphically. The G_x value for a measurement interval was plotted against the initial BG, time between measurements, mean exogenous glucose delivery (enteral and parenteral), and mean exogenous insulin. The contribution of G_x to glucose disposal was also calculated.

3. RESULTS

Figure 2 shows the dramatic reduction in fitting error when the SDE approach is used. No discernable difference is visible between the 0^{th} and 2^{nd} order basis functions. The difference between the two basis function shapes is shown in Figure 3, where the 0^{th} order SDE forces $G_x \to 0$ except at extreme cases. The regularised basis of the 2^{nd} order SDE forces G_x to absorb the remaining BG changes, hence producing a more variable G_x profile. The method thus permits control over the contribution of noise to hour-to-hour BG changes.

Dependence of the G_x profile was investigated in Figure 4. Graphically, the only factor that affects the variance of G_x is measurement interval, which is congruent with the constraints on G_x. Due to the choice of constraint, longer

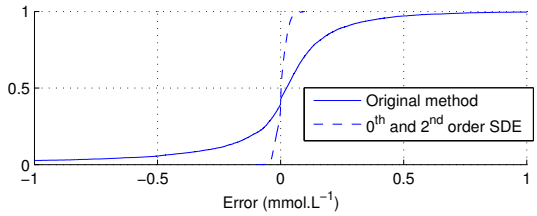

Fig. 2. *Fitting error comparison between the methods. No discernible differences is visible between the two SDE approaches.*

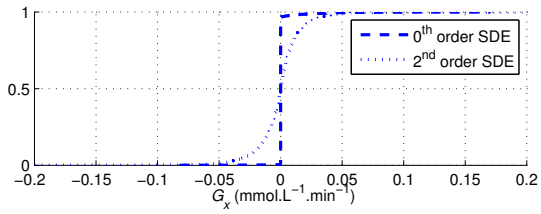

Fig. 3. *Noise signal comparison. The 0^{th} order SDE was fitted with a constant S_I between BG measures, while the 2^{nd} order SDE was fitted with a 2^{nd} order* b-spline *basis with knot widths of 240 min.*

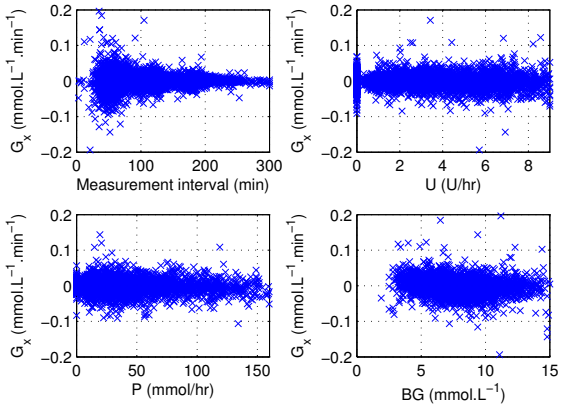

Fig. 4. *Scatter diagram of all fitted G_x vales with a 2^{nd} order basis, plotted against measurement interval, exogenous insulin, exogenous glucose, and initial BG. Hourly insulin is only indicative, as timing errors will affect the value when insulin is delivered in a bolus.*

term changes in BG are mediated by S_I, while short term changes predominantly affect the G_x profile.

Finally, Figures 5 and 6 show the new methods implemented on the same patient and compared to the original. This patient was chosen as the episode includes an initial period of high BG that cannot be fitted by the original method, and has some minor interpolation artefacts near the end of the episode. Figure 5 shows the 0^{th} order basis has almost equivalent S_I, barring some minor timing differences, and G_x is zero everywhere except for the initial high BG period that is impossible to fit using the non-stochastic model. In comparison, the regularised basis used in Figure 6 has a much smoother S_I profile, and where there are disparities between the original and new S_I profiles G_x is forced to be non-zero.

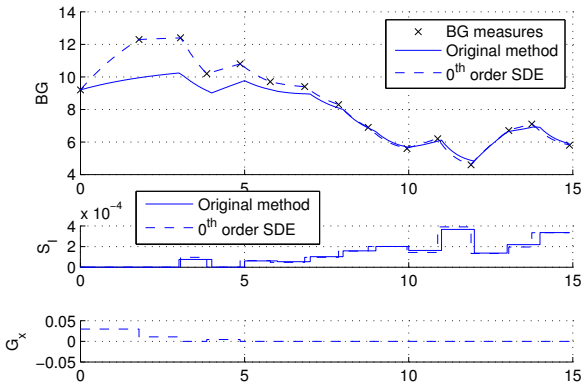

Fig. 5. *0^{th} order basis functions compared to the original methods. S_I is similar, and G_x is only non-zero when S_I is constrained to the lower limit.*

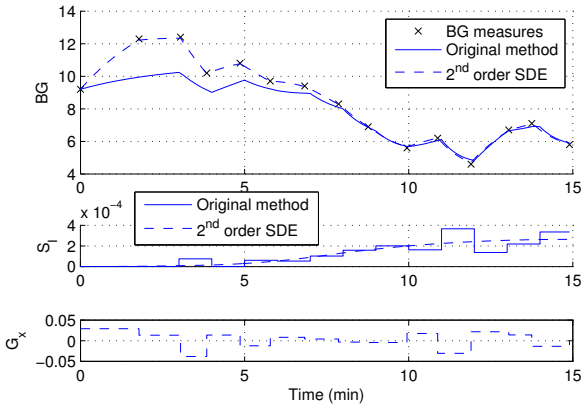

Fig. 6. *2^{nd} order basis functions compared to the original method. S_I is much smoother, and G_x is non-zero both when S_I is constrained and when differences exist between the new and old profiles.*

4. DISCUSSION

An updated fitting methodology is needed when using higher resolution data. An hourly S_I value is useful for control, but timing errors cause information to be lost when creating virtual patients. During a virtual trial, this discrepancy means the original data cannot be created when the original inputs are used. Generalising S_I to a series of b-*spline* basis functions permits timing errors to be natively captured, but introduces fitting error if the basis function local support extends beyond an individual measurement interval. To permit greater local support, and thus regularise the S_I profile, a new time-varying parameter, G_x, was introduced to eliminate fitting error. The methodology presented here is a robust way identify both of these parameters, where forcing G_x to be zero-mean allows the system to be fully defined, and the iterative integral method allows rapid solution.

Figure 2 demonstrates this approach solves the original problem. Fitting error can be forced to zero if enough iterations are carried out, and both the S_I and G_x profiles can be used in a virtual trial that fully recreates the data if the original intervals are used. As S_I is regularised with a higher order basis, and wider knot widths, shorter-term

changes cannot be captured by S_I, instead being forced into the G_x profile. This effect is seen in both Figure 3 and the example patient in Figures 5 and 6. In this way, direct control can be exerted over the balance between changes in S_I and observed noise.

This balance between noise and changes in S_I may also be useful in control. Figure 4 shows a lack of dependence of noise on both exogenous inputs and current BG, showing only the expected dependence on measurement interval. This knowledge may be leveraged to improve model prediction. In particular, a continuous S_I profile and a cohort G_x profile could be used in a non-parametric prediction algorithm, replacing the computationally intensive stochastic model of STAR [Fisk et al., 2012]. Such a non-parametric approach could feasibly be updated in real-time, improving the quality of glycaemic control for longer-stay ICU patients. Safety would be improved for more variable (higher noise) patients, and performance would be improved for more stable patients, neither of whom benefit from a whole-cohort approach. The novel simplicity of this SDE permits stable parameter identification with a relatively computationally light algorithm, which permits use in real-time glycaemic control.

One final possibility created by using this SDE approach is analysis of noise around changes in exogenous inputs. Patterns in the noise profile may possibly be used to identify areas where the model could be improved. For example, changes in nutrition rate may affect BG faster than the glucose absorption submodel permits, which will appear as positive G_x values immediately after nutrition changes. Consistent patterns could be used to update parameter values either for an individual or for the whole-cohort ICING model.

5. CONCLUSIONS

A robust parameter identification method was introduced, permitting identification of a smooth S_I profile, and capturing remaining variation in a simple stochastic element, G_x. This method suits the high resolution data available to STAR. Thus, future work will centre around implementation of a non-parametric prediction algorithm using these two parameters, allowing for direct inclusion into STAR.

REFERENCES

J Geoffrey Chase, Geoffrey Shaw, Aaron Le Compte, Timothy Lonergan, Michael Willacy, Xing-Wei Wong, Jessica Lin, Thomas Lotz, Dominic Lee, and Christopher Hann. Implementation and evaluation of the sprint protocol for tight glycaemic control in critically ill patients: a clinical practice change. *Crit Care*, 12(2):R49, 2008. doi: 10.1186/cc6868. URL `http://dx.doi.org/10.1186/cc6868`.

J Geoffrey Chase, Fatanah Suhaimi, Sophie Penning, Jean-Charles Preiser, Aaron J Le Compte, Jessica Lin, Christopher G Pretty, Geoffrey M Shaw, Katherine T Moorhead, and Thomas Desaive. Validation of a model-based virtual trials method for tight glycemic control in intensive care. *Biomed Eng Online*, 9:84, 2010.

J Geoffrey Chase, Aaron J Le Compte, Fatanah Suhaimi, Geoffrey M Shaw, Adrienne Lynn, Jessica Lin, Christopher G Pretty, Normy Razak, Jacquelyn D Parente, Christopher E Hann, et al. Tight glycemic control in critical care–the leading role of insulin sensitivity and patient variability: a review and model-based analysis. *Computer methods and programs in biomedicine*, 102(2): 156–171, 2011. URL `http://www.sciencedirect.com/science/article/pii/S0169260710002828`.

Carl De Boor. On calculating with B-splines. *Journal of Approximation Theory*, 6(1):50–62, 1972.

Paul D Docherty, J. Geoffrey Chase, and Timothy David. Characterisation of the iterative integral parameter identification method. *Med Biol Eng Comput*, 50(2):127–134, Feb 2012. doi: 10.1007/s11517-011-0851-y. URL `http://dx.doi.org/10.1007/s11517-011-0851-y`.

Tamás Ferenci, Balazs Benyo, Levente Kovács, Liam Fisk, Geoffrey M Shaw, and J Geoffrey Chase. Daily evolution of insulin sensitivity variability with respect to diagnosis in the critically ill. *PloS one*, 8(2):e57119, 2013.

Liam M. Fisk, Aaron J. Le Compte, Geoffrey M. Shaw, Sophie Penning, Thomas Desaive, and J Geoffrey Chase. Star development and protocol comparison. *IEEE Trans Biomed Eng*, 59(12):3357–3364, Dec 2012. doi: 10.1109/TBME.2012.2214384. URL `http://dx.doi.org/10.1109/TBME.2012.2214384`.

Christopher E. Hann, J Geoffrey Chase, Jessica Lin, Thomas Lotz, Carmen V. Doran, and Geoffrey M. Shaw. Integral-based parameter identification for long-term dynamic verification of a glucose-insulin system model. *Comput Methods Programs Biomed*, 77(3):259–270, Mar 2005. doi: 10.1016/j.cmpb.2004.10.006. URL `http://dx.doi.org/10.1016/j.cmpb.2004.10.006`.

Jessica Lin, Normy N. Razak, Christopher G. Pretty, Aaron Le Compte, Paul Docherty, Jacquelyn D. Parente, Geoffrey M. Shaw, Christopher E. Hann, and J. Geoffrey Chase. A physiological intensive control insulin-nutrition-glucose (icing) model validated in critically ill patients. *Comput Methods Programs Biomed*, 102(2):192–205, May 2011. doi: 10.1016/j.cmpb.2010.12.008. URL `http://dx.doi.org/10.1016/j.cmpb.2010.12.008`.

Christopher G Pretty. *Analysis, classification and management of insulin sensitivity variability in a glucose-insulin system model for critical illness*. PhD thesis, University of Canterbury, 2012. URL `http://hdl.handle.net/10092/6580`.

Christopher G. Pretty, Aaron J. Le Compte, J Geoffrey Chase, Geoffrey M. Shaw, Jean-Charles Preiser, Sophie Penning, and Thomas Desaive. Variability of insulin sensitivity during the first 4days of critical illness: implications for tight glycemic control. *Ann Intensive Care*, 2(1):17, 2012. doi: 10.1186/2110-5820-2-17. URL `http://dx.doi.org/10.1186/2110-5820-2-17`.

Anders Thorell, Olav Rooyackers, Peter Myrenfors, Mattias Soop, Jonas Nygren, and Olle H. Ljungqvist. Intensive insulin treatment in critically ill trauma patients normalizes glucose by reducing endogenous glucose production. *J Clin Endocrinol Metab*, 89(11):5382–5386, Nov 2004. doi: 10.1210/jc.2004-1118. URL `http://dx.doi.org/10.1210/jc.2004-1118`.

Nicolaas G Van Kampen. Stochastic differential equations. *Physics reports*, 24(3):171–228, 1976. URL `http://www.sciencedirect.com/science/article/pii/0370157376900296`.