

# Model Predictive and Proportional Integral Control of Blood Clotting Speed using Warfarin when Data are Missing

Emma D. Wilson\*, Quentin Clairon<sup>†</sup>, Robin Henderson<sup>†</sup> and C. James Taylor\*

\* *Engineering Department, Lancaster University, Lancaster, UK, LA1 4YW*

*Email: e.d.wilson1@lancaster.ac.uk, c.taylor@lancaster.ac.uk*

<sup>†</sup>*School of Mathematics, Statistics and Physics, Newcastle University, Newcastle upon Tyne, UK, NE1 7RU*

*Email: quentin.clairon@newcastle.ac.uk, robin.henderson@newcastle.ac.uk*

**Abstract**—A control theory approach to the management of the blood clotting speed using the anticoagulant Warfarin is investigated. Controllers are developed and analysed using hospital data from patients with chronic conditions under Warfarin anticoagulation treatment. Proportional Integral (PI) and Model Predictive (MPC) controllers are used to estimate treatment decisions. These controllers are adapted in a novel manner, to enable their use with missing or irregularly sampled data. The performance of the controllers is evaluated both using a simulation of the system and by retrospectively comparing actual decisions in the data to those suggested by the control algorithms. It is shown that when the blood clotting speed is within a target range, the decisions suggested by the control algorithms are similar to those actually made (by medical staff), so would likely have led to similar desirable outcomes. When the blood clotting speed is outside the desirable range and too high or too low, the control algorithms on average suggest lower, or higher inputs respectively. These suggestions are likely to lead to improved outcomes.

**Keywords**—Adaptive Treatment; Anticoagulation; Missing Measurements; Proportional Integral (PI); Model Predictive Control (MPC); Non-Minimum State Space (NMSS)

## I. INTRODUCTION

The problem of determining effective medical treatments for individual patients can be treated as a control problem. Using control theory to determine medical treatments is appealing as it provides a systematic way of achieving desired performance in the presence of uncertainty, external disturbances and noise [1]. The question of determining individually tailored treatments has also been considered in the biostatistical literature, where the problem is cast as an optimal dynamic treatment (ODT) problem [2], [3], [4]. The ODT biostatistical literature focuses on modelling, estimation and inference, as opposed to control.

Algorithms from control theory have been applied to various medical treatments. These include: i) the control of blood glucose levels by adjusting insulin inputs using proportional-integral-derivative (PID) [5], [6] and model predictive control (MPC) [7], [8]; ii) anaesthetic drug delivery using PI, PID and robust control [9], [10], [11], [12]; and iii) behavioural research to develop interventions using MPC [13], [14], [15].

In this article, we consider the control of the blood clotting speed using the anticoagulant Warfarin. Anticoagulants are one of the most commonly prescribed group of drugs and Warfarin is probably the most used anticoagulant. Simplified pharmacodynamic models for the Warfarin dose response have been developed using methods from control theory [16]. However, with the exception of recent work by the present authors [1], regulating the blood clotting speed has not been considered using a control theory approach.

We use observational data from patients undergoing Warfarin anticoagulation treatment to design and evaluate control algorithms based on both MPC and PI control. A PI controller was selected for this research, since it is ubiquitous and remains one of the main controllers used for industrial applications. MPC was selected because of its ability to handle constraints and, like PI control, has a track record of use for various medical control applications.

In control theory it is often assumed that observations and inputs occur at fixed intervals, with uniform sampling intervals pre-determined by the sampling period  $T$ . However, in practice, medical treatments are rarely measured and applied at such fixed times. Appointment times (or measurements) may be scheduled without constant time gaps, planned visits may not be adhered to, and individuals may drop out from trials [17]. Indeed, the Warfarin data evaluated in this contribution is not uniformly sampled.

We consider a non-minimal state space (NMSS) framework for both PI and MPC design. For further details on NMSS methods, see e.g. [18], [19] and the references therein. We show how these controllers can be adapted to enable their application to blood clotting speed, where output data are missing or irregularly sampled. Here, the conventional integral-of-error state used for NMSS design, is replaced by a new form that facilitates application to irregular data. Control performance is evaluated with missing measurements on a simulated system, as well as retrospectively using observational data. In the latter case, we compare decisions made by the healthcare provider to those suggested by the algorithms. This provides a useful comparison since we can reasonably expect the healthcare practitioners to have made good decisions on input doses.

Hence, the novel contributions of the article are the application of control methods to Warfarin data, the adaptation of these methods for irregularly sampled data, and the use of a retrospective data analysis. The remainder of the paper is structured as follows: Section 2 provides the methods, including details on the Warfarin data and model, the PI and MPC control algorithms, and how these are adapted to deal with missing data, as well as information about the simulations and retrospective data analysis. Subsequently, in Section 3 we present the results, followed by the discussion in Section 4 and conclusions in Section 5.

## II. METHODS

In this section, the data, model, control algorithms and simulation methods are all described. The system model was developed in R and control simulations were implemented using Matlab.

### A. Patient Data

Data from 152 patients with chronic conditions under Warfarin anticoagulation treatment were used. The controlled variable is the blood clotting speed, which is measured through the International Normalised Ratio (INR). The INR is a standardised measure, with high values indicating long clotting times. Usually the aim is to achieve INR values in the range 2–3 units [20]. The input used to achieve control is the Warfarin dose, which is measured in mg. Available covariates include age and sex, with the mean age of patients 84 and range of ages 75–97. The data were recorded in Newcastle upon Tyne, UK between 1995–2013. Treatment periods for each patient varied from 186 to 5925 days and involved 6 to 215 clinic visits. The gaps between visits were not regular. For all patients, 67% of observations were within 15% of either 7, 14, 21 or 28 days, with 27% of intervals in excess of 32 days. Measurements recorded as INR=0 were assumed to be missing and so are excluded from the analysis.

### B. System Model

The Warfarin data was used to estimate a simulated model as detailed in [1]. The model structure was chosen as it is simple, but provides a reasonable description of the data. Note that, in contrast to [1], the present article develops PI and MPC designs that address the missing data problem. The output  $y(k)$  is the  $\log(\text{INR})$  and the input  $u(k)$  the dose (mg) of Warfarin. The output is modelled as,

$$y(k) = a_1 y(k-1) + b_1 u(k-1) + \varepsilon \quad (1)$$

where, based on a typical patient and obtained using the `lm` function in R,  $a_1 = 0.4$  and  $b_1 = 0.25$ ,  $\varepsilon \sim N(0, \sigma^2)$  with  $\sigma = 0.25$ .

### C. PI Control Algorithm for Missing Data

The PI control algorithm was based on state variable feedback methods, with a non-minimal state vector  $x(k) = [y(k) \quad q(k)]^T$ , where  $q(k)$  is an integral-of-error state used to provide steady state tracking of the target or set point (see definition below). Proportional ( $k_p$ ) and integral ( $k_I$ ) gains were calculated using standard linear quadratic optimisation and weights  $w_e = w_x = w_u = 1$ ; for further details on these weights and NMSS control design, see Section 5.4 in [18]. The control input or dose is determined as follows,

$$u(k) = -k_p y(k) + k_I q(k) \quad (2)$$

Previous research into NMSS design has used a simple backwards rectangular integration to define the integral-of-error state [18], [19]. By contrast, the present work uses a trapezoidal integration to improve accuracy when data are missing. Here the integral-of-error is,

$$q(k) = q(k-1) + 0.5T(e(k) + e(k-1)) \quad (3)$$

where  $e(k) = y_d(k) - y(k)$  and  $y_d(k)$  is the desired set point and  $T$  is the sampling time. In the nominal case,  $T = 1$  is assumed. In the case of missing data,  $e(k-1)$  and  $q(k-1)$  may not be available. Hence, we can write,

$$q(k) = q(k-N) + 0.5NT(e(k) + e(k-N)) \quad (4)$$

where  $y(k-N)$  is the previous measurement and  $N$  is the number of samples between the current and previous measurement. When data are not missing  $N = 1$ . The following notation is introduced:

$$\begin{aligned} z^{-1}q(k) &= q(k-1) \quad ; \quad z^{-N}q(k) = q(k-N) \\ q(k) - q(k-N) &= q(k)(1 - z^{-N}) \\ (1 - z^{-N}) &= \Delta_N \end{aligned}$$

Using the above notation Eq. (4) can be written as,

$$q(k) = \frac{NT}{2\Delta_N}(e(k) + e(k-N)) \quad (5)$$

Substituting this into Eq. (2) yields an expression for the input in terms of available data,

$$u(k) = -k_p y(k) + k_I \frac{NT}{2\Delta_N}(e(k) + e(k-N)) \quad (6)$$

To implement control, the input was expressed in incremental form, to avoid wind-up problems, i.e.,

$$u(k) = u(k-N) - k_p(y(k) - y(k-N)) + \frac{k_I NT}{2}(e(k) + e(k-N)) \quad (7)$$

but with the constraint that  $u(k) \geq 0$ . This approach yields an input in terms of the available signals, and is used to appropriately update  $u(k)$  when data are missing. The input dose is only updated if data are available, otherwise the input is held at the last observed value, i.e.  $u(k) = u(k-N)$

#### D. MPC Algorithm for Missing Data

The NMSS-based MPC algorithm described by [21] was used, rather than alternatives such as [22]. This algorithm was selected since it provides integral action without requiring an explicit estimation of the integral-of-error state (which is unavailable when analysing retrospective data). For the model (1), the state is  $\mathbf{x}_m(k) = [\Delta_1 y(k) \quad y(k)]^T$  where  $\Delta_1 y(k) = y(k) - y(k-1)$ . The MPC cost function is,

$$J = \gamma \sum_{i=1}^{N_p} (\mathbf{C} \mathbf{x}_m(k+i|k) - y_d(k+i))^2 + \lambda \sum_{i=0}^{N_c-1} (\Delta_1 u(k+i|k))^2 \quad (8)$$

where  $\mathbf{C} = [0 \quad 1]$ . For the results presented below, the control and input horizon were equal, with  $N_p = N_c = 10$ , while  $\gamma = 1$  and  $\lambda = 1$ . This MPC problem can be cast as a quadratic program at each sampling instant  $k$ , i.e. to find the vector  $\Delta_1 \mathbf{U} = [\Delta_1 u(k) \quad \Delta_1 u(k+1|k) \quad \dots \quad \Delta_1 u(k+N_p-1|k)]$  that minimises the cost function in Eq. (8) subject to constraints. Constraints on the change in input  $\Delta_1 \mathbf{U}$  were imposed such that the input dose could never be zero, and (for illustrative purposes) the change in dose was always less than unity. The constrained quadratic program was solved using the Matlab function `quadprog`. This yields a suggested dose sequence that was calculated as a function of the previous input and the current state  $\mathbf{x}_m(k)$ .

In MPC, generally only the first element of the sequence of suggested inputs is applied, while subsequent inputs are recalculated by re-estimating the optimal input sequence at each sampling instant. For MPC with missing data, one approach is to hold the control input constant and equal to the first element, which is updated only when data become available. This is a similar approach to that used for PI control above, and is termed ‘MPC hold’ in the results below. Alternatively, the sequence of MPC inputs determined via  $\Delta_1 \mathbf{U}$  can be used until more data becomes available. This is termed ‘MPC tail’ in the results. When data are missing  $y(k-1)$  was unavailable. We therefore made the approximation that  $y(k-1) \approx y(k-N)$  and so  $\Delta_1 y(k) \approx \Delta_N y(k) \approx y(k) - y(k-N)$ . The approximation of  $\Delta_1 y(k)$  is used in the estimate of the current state.

#### E. Control Simulations

The performance of the PI and MPC control algorithms was evaluated in simulation. Here, the plant was modelled using Eq. (1) and the estimated model parameters were varied from the data-based estimated values using a Monte-Carlo simulation, with standard deviation of 0.01 for both  $a_1$  and  $b_1$ . Missing data were simulated as missing completely at random (MCAR) [1], such that for each sample the probability that the output was missing was 50%.

The desired INR  $y_d(k)$  was 2.5 for 50 samples, then 3.5 for a further 50 samples. These are the normal targets for patients with occasional and recurrent deep vein thrombosis respectively [20], while the increase mimics a change in the diagnosis. The percentage of measurements within  $\pm 0.5$  of the desired INR was used as a performance metric. Performance metrics were based on 1000 different simulations, each with 100 data points.

#### F. Retrospective Data Analysis

At each visit, the inputs suggested by the control algorithms (based on previous measured output and input) were compared to those made by healthcare providers in the data set. Control estimates were made corresponding to each visit using the PI and MPC controllers described above. The PI input was calculated in terms of the data measured at the current and previous visit using Eq. (7), and the MPC input calculated by solving the quadratic programme Eq. (8).

In the actual data, the majority of measurements were separated by 7, 14, 21 or 28 days, hence we define  $dsep = [7 \quad 14 \quad 21 \quad 28]$ . In calculating the control input, the sampling period  $T$  was set to  $1/\text{mean}(dsep)$ . This is because on average the sampling period of the control model (i.e. that the controllers were optimised for) was  $\approx 1/\text{mean}(dsep)$  times that of the actual data.

We first considered instances where the healthcare provider made a ‘good’ decision – this was defined as a dose decision that resulted in a INR in the range 2–3 at the subsequent visit. For such good decisions, a percentage difference ( $pdu$ ) was estimated as,

$$pdu(k) = \frac{|u_a(k) - u_s(k)|}{(u_a(k) + u_s(k))/2} \quad (9)$$

where  $u_a(k)$  is the actual input from the data, and  $u_s(k)$  is the suggested control input from the algorithms.

Instances where the healthcare provider made a ‘bad’ decision were then considered. These were split into decisions that led to too high an outcome (e.g. dose decision that resulted in an INR above 3), and decisions that led to too low an outcome (e.g. dose decision that resulted in an INR below 2). In these cases, the proportion of decisions in which the controllers suggested a lower preceding dose for high outcomes, or a higher dose for low outcomes, was calculated. In these cases, we assume that the control input would have improved the outcome; however, as the analysis of data is retrospective, this of course cannot be tested.

### III. RESULTS

#### A. Simulated Control

Three different control algorithms were implemented in simulation. These were i) a PI controller with inputs held constant when data are missing; ii) MPC with inputs held constant when data are missing (MPC hold); and iii) MPC with the sequence of optimal inputs used when data are

missing (MPC tail). The performance of each is evaluated for 1000 simulations in Table I. An example dose and INR sequence for each control method is shown in Fig. 1.

In the best case scenario, the target range is only achieved in 45% of measurements<sup>1</sup>. This is due to the simulated (but realistic) high levels of random noise. As might be expected, these results show that MPC outperforms the PI controller. Under missing data, MPC hold yields improved results in comparison to MPC tail. This is unexpected and requires further research. However, it is likely to be related to the high noise to signal ratio, and the fact that there is only one change in the desired set point in these simulations. In fact, when no noise was added to the simulations ( $\varepsilon = 0$ ), then MPC tail yields the best performance. In this case, with zero noise and missing data, 90% of the results are in range when holding the input (MPC hold) and 94% when using the optimal sequence (MPC tail). Fig. 2 illustrates a typical example of the latter scenario, and demonstrates that using the tail of the input sequence improves control when the reference signal changes – especially when data are missing around this change in the reference signal.

### B. Comparison with Actual Decisions

The dose inputs in the data set, determined by healthcare practitioners, were compared to those suggested by MPC and PI control algorithms. Table II compares the suggested and actual doses for fifteen patients when the actual doses led to good outcomes (i.e.  $INR=2-3$ ), taken from the full set of 152 patients at random. In general the inputs suggested by the control algorithms are within a similar range to the actual decisions. Figure 3 plots the suggested and actual doses for two patients: one where the percentage change in dose was usually small, and the other where it was often large. The larger percentage differences in Fig. 3b) can be attributed to the lower mean dose level. In general input changes are more frequent in the control decisions. This is because a set point of 2.5 was being tracked by the control algorithms but, in practice, the medical aim is to keep the INR within the range 2–3. Further constraints could be incorporated into the MPC algorithm to prevent small dose changes (other than zero) since the actual doses tend to vary at fixed increments. There are occasionally large variations between the PI suggestion and the actual dose, suggesting MPC control may be more appropriate, due to its ability to handle constraints.

Table III compares the suggested and actual doses for fifteen patients when the actual doses led to poor outcomes ( $INR < 2$  or  $INR > 3$ ). When the dose led to too high a response the proportion of times the controller suggestion is lower (so may have improved the actual medical outcome) is given, and vice-versa when the dose led to too low a response. For all patients, the controllers suggested a

<sup>1</sup>For the actual data, 54% of measurements are within the target range of 2–3 INR; this discrepancy is because the real data do not encompass a change in the set point.

dose that may have improved outcomes over 60% of the time, implying that the controller may have outperformed the medical decision takers. However, as the analysis is retrospective this cannot be confirmed in practice.

Method	No Data Missing	50% of Data Missing
MPC hold	45	42
MPC tail	45	38
PI	37	32

Table I  
PERCENTAGE OF MEASURED INR VALUES WITHIN DESIRED RANGE.

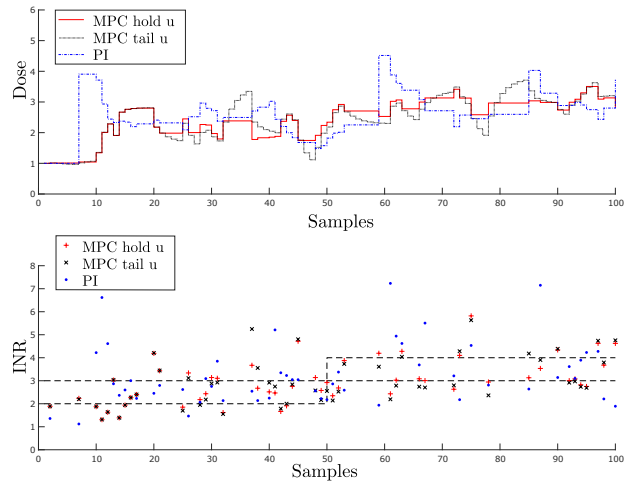


Figure 1. Example of INR and warfarin inputs for each control method. Only the available data are plotted.

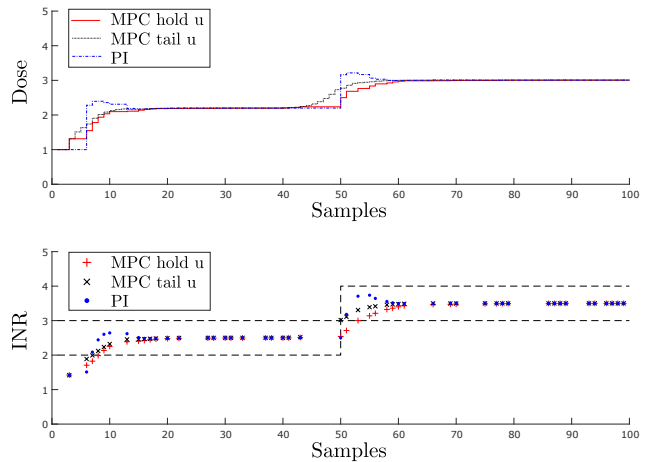


Figure 2. Example of INR and warfarin inputs for each control method. In these simulations, the noise  $\varepsilon = 0$ . Only the available data are plotted.

## IV. DISCUSSION

In applying control to medical treatments a key issues is ensuring patient safety. Hence, in many areas the use of

Patient No.	No. of obs.	No. of good decisions	Input difference (%)	
			MPC	PI
8	34	22	4.64	6.75
14	125	93	2.77	4.51
38	25	16	8.44	5.98
45	41	28	4.59	4.34
47	26	14	10.98	11.32
54	29	17	4.44	8.82
57	22	13	3.85	8.89
79	113	59	8.92	6.83
82	174	119	5.27	6.11
91	97	40	8.26	5.57
100	64	50	2.79	4.27
107	61	29	13.52	13.19
110	50	31	10.70	9.11
130	109	65	19.30	21.21
139	14	11	13.73	14.59
Mean	60	38	7.93	8.83
Weighted Mean	-	-	7.48	8.04

Table II

COMPARISON OF ACTUAL AND SUGGESTED DOSES (CALCULATED USING EQ. (9)) FOR FIFTEEN EXAMPLE PATIENTS WHEN THE ACTUAL DOSES LED TO GOOD OUTCOMES. AVERAGE MEANS ARE GIVEN FOR THE FULL SET OF 152 PATIENTS, THE WEIGHTED MEAN IS WEIGHTED TO TAKE INTO ACCOUNT THE NUMBER OF OBSERVATIONS.

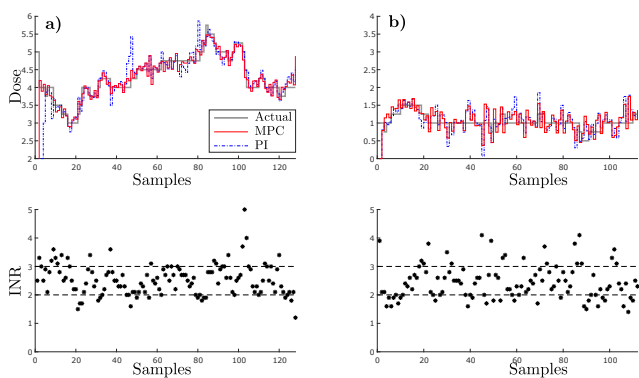


Figure 3. INR and dose, with comparison to doses suggested by MPC and PI control methods at each visit: a) Patient 14 and b) Patient 130.

automatic controllers is not standard practice and decisions are made by health practitioners. However, control theory could help to guide decisions. The controllers evaluated in this paper led to similar or improved decisions when compared to the decisions made by the healthcare practitioner. Given that we can expect the health practitioners to make good, experienced decisions, this motivates further investigation into the potential use of control theory in determining individual medical treatments.

In the biostatistical literature, research into ODTs has progressed rapidly from the papers of Murphy [2] and Robins [3], with methods including Q-learning and A-learning [4]. However, there have been very few attempts to make use of control theory in optimal dynamic selection. Hence, our overarching aim is to combine ideas from well established approaches in control with the statistical theory of ODT regimes. Progress has been made towards this by

using robust control theory in ODT selection [23].

One of the challenges in developing a medical control strategy is the fact that data are not uniformly sampled and patients may miss appointments. As a first step, in the present article we have considered the case where data are missing completely at random (MCAR). However, the missingness mechanism is important when dealing with observational data, as discussed in more depth by [1]. In the present study, the model order was low and so only the current and previous measurements were needed to calculate the control inputs. The proposed method needs extending to the case when more history is required by the control algorithm to calculate the inputs, in which case there is more potential for useful information to be missing. Hence, future work also aims to look at developing a more general method to implement control when measurements are missing from the history vector.

## V. CONCLUSIONS

In this paper PI and MPC controllers were applied to the problem of controlling blood clotting speed using Warfarin. These controllers were adapted to enable application to data which was not uniformly sampled. Retrospective data analysis suggests that control theory could be applied to the problem of dose assignment to achieve as good, or better decisions than those made by the health care practitioner. However, this is only determined from retrospective analysis. Due to patient safety issues it is difficult to analyse online performance in a real system. The underlying model used to develop the controllers was a relatively simple first order model. Future work aims to extend the control algorithms and analysis to more complex system models.

## ACKNOWLEDGMENT

This work is supported by the UK Engineering & Physical Sciences Research Council (EPSRC) grant EP/M015637/1.

## REFERENCES

- [1] E. D. Wilson, Q. Clarion, R. Henderson, and C. J. Taylor, "Dealing with observational data in control," *Annual Reviews in Control*, 2018 (<https://doi.org/10.1016/j.arcontrol.2018.05.001>).
- [2] S. A. Murphy, "Optimal dynamic treatment regimes," *Journal of the Royal Statistical Society Series B*, vol. 65, pp. 331–355, 2003.
- [3] J. M. Robins, "Optimal structural nested models for optimal sequential decisions," in *Proceedings of the Second Symposium on Biostatistics*, D. Y. Lin and P. Heagerty, Eds. New York: Springer, 2004, pp. 189–326.
- [4] B. Chakraborty and E. Moodie, *Statistical Methods for Dynamic Treatment Regimes*. New York: Springer, 2013.
- [5] M. Hoekstra, M. Vogelzang, E. Verbitskiy, and M. W. Nijsten, "Health technology assessment review: Computerized glucose regulation in the intensive care unit— how to create artificial control," *Critical Care*, vol. 13, no. 5, p. 223, 2009.

Patient No.	No. of obs.	No. of decisions giving too high an outcome	Proportion of high decisions where controller suggestion is lower (%)		No. of decisions giving too low an outcome	Proportion of low decisions where controller suggestion is higher (%)	
			MPC	PI		MPC	PI
8	34	4	75	50	8	88	75
14	125	17	82	65	15	73	80
38	25	2	0	0	7	86	71
45	41	6	67	83	7	71	57
47	26	7	57	57	5	80	40
54	29	3	33	33	9	89	78
57	22	2	50	0	7	100	86
79	113	28	57	57	26	65	50
82	174	16	50	44	39	79	85
91	97	23	39	30	34	74	68
100	64	3	67	67	11	55	55
107	61	15	60	53	17	76	59
110	50	10	60	40	9	33	56
130	109	19	68	74	25	80	76
139	14	1	100	100	2	0	0
Mean	60	10	59	56	12	68	65
Weighted Mean	-	-	57	56	-	68	66

Table III

COMPARISON OF ACTUAL AND SUGGESTED DOSES FOR FIFTEEN EXAMPLE PATIENTS WHEN THE ACTUAL DOSES LED TO POOR OUTCOMES. WHEN THE DOSE LED TO TOO HIGH A RESPONSE THE PROPORTION OF TIMES THE CONTROLLER SUGGESTION IS LOWER (SO MAY HAVE IMPROVED THE OUTCOME) IS GIVEN, AND VICE-VERSA WHEN THE DOSE LED TO TOO LOW A RESPONSE.

- [6] R. Rattan and S. A. Nasraway, "The future is now: software-guided intensive insulin therapy in the critically ill," *Journal of Diabetes Science and Technology*, vol. 7, no. 2, pp. 548–554, 2013.
- [7] 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*, vol. 25, no. 4, p. 905, 2004.
- [8] A. Chakrabarty, S. Zavitsanou, F. Doyle III, and E. Dassau, "Event-triggered model predictive control for embedded artificial pancreas systems," *IEEE Transactions on Biomedical Engineering*, 2017.
- [9] G. A. Dumont, "Closed-loop control of anesthesia— a review," *IFAC Proceedings Volumes*, vol. 45, no. 18, pp. 373–378, 2012.
- [10] G. A. Dumont, A. Martinez, and J. M. Ansermino, "Robust control of depth of anesthesia," *International Journal of Adaptive Control and Signal Processing*, vol. 23, no. 5, pp. 435–454, 2009.
- [11] S. Bibian, C. R. Ries, M. Huzmezan, and G. Dumont, "Introduction to automated drug delivery in clinical anesthesia," *European Journal of Control*, vol. 11, no. 6, pp. 535–557, 2005.
- [12] A. Gentilini, C. W. Frei, A. H. Glatfledler, M. Morari, T. J. Sieber, R. Wymann, T. W. Schnider, and A. M. Zbinden, "Multitasked closed-loop control in anesthesia," *IEEE Engineering in Medicine and Biology Magazine*, vol. 20, no. 1, pp. 39–53, 2001.
- [13] D. E. Rivera, M. D. Pew, and L. M. Collins, "Using engineering control principles to inform the design of adaptive interventions: A conceptual introduction," *Drug & Alcohol Dependence*, vol. 88, pp. S31–S40, 2007.
- [14] S. Deshpande, N. N. Nandola, D. E. Rivera, and J. W. Younger, "Optimized treatment of fibromyalgia using system identification and hybrid model predictive control," *Control Engineering Practice*, vol. 33, pp. 161–173, 2014.
- [15] K. Bekiroglu, C. Lagoa, S. A. Murphy, and S. T. Lanza, "Control engineering methods for the design of robust behavioral treatments," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 3, pp. 979–990, 2017.
- [16] S. Kim, A. E. Gaweda, D. Wu, L. Li, S. N. Rai, and M. E. Brier, "Simplified warfarin dose-response pharmacodynamic models," *Biomedical Engineering (Singapore)*, vol. 27, no. 1, 2015.
- [17] P. Diggle, *Analysis of longitudinal data*. Oxford University Press, 2002.
- [18] C. J. Taylor, A. Chotai, and P. C. Young, "State space control system design based on non-minimal state-variable feedback: further generalization and unification results," *International Journal of Control*, vol. 73, no. 14, pp. 1329–1345, 2000.
- [19] C. J. Taylor, P. C. Young, and A. Chotai, *True digital control: statistical modelling and non-minimal state space design*. John Wiley & Sons, 2013.
- [20] T. P. Baglin, D. M. Keeling, H. G. Watson, and the British Committee for Standards in Haematology, "Guidelines on oral anticoagulation (warfarin): third edition – 2005 update," *British Journal of Haematology*, vol. 132, pp. 277–285, 2006.
- [21] L. Wang and P. C. Young, "An improved structure for model predictive control using non-minimal state space realisation," *Journal of Process Control*, vol. 16, no. 4, pp. 355–371, 2006.
- [22] V. Exadaktylos and C. J. Taylor, "Multi-objective performance optimisation for model predictive control by goal attainment," *International Journal of Control*, vol. 83, no. 7, pp. 1374–1386, 2010.
- [23] Q. Clairon, E. D. Wilson, R. Henderson, and C. J. Taylor, "Adaptive biomedical treatment and robust control," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 12 191–12 196, 2017.