Integral of Error forms for Blood Clotting Speed Control using Warfarin when Data are Missing

C. J. Taylor and D. G. Millevithanatchy
School of Engineering, Lancaster University, UK
Email: c.taylor@lancaster.ac.uk

E.D. Wilson
SCC, Lancaster University, UK
Email: e.d.wilson1@lancaster.ac.uk

***Abstract*—A control theory approach to the management of blood clotting speed using the anticoagulant warfarin is investigated. Proportional Integral (PI) and Proportional-Integral-Plus (PIP) controllers are developed for models identified from patient data. These are used to estimate treatment decisions subject to stochastic disturbances, model uncertainty and missed observations, the latter representing missed clinic appointments. The focus is on the relative performance of various integral-of-error (IOE) forms, which are used to track the target International Normalised Ratio (INR) at steady state. These are adapted in novel ways to handle the missed observations. Preliminary Monte Carlo simulations suggest that forward difference, mean and trapezoidal IOE forms with modified pole assignment, could lead to more desirable outcomes than the standard case, but the differences are relatively small and more research is required into scenarios leading to set point deviations.**

***Index Terms*—Adaptive Treatment; Anticoagulation; Missing Measurements; Integral-of-Error; Proportional Integral (PI).**

I. Introduction

Using control theory to determine medical treatments for individual patients 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 problem [2–4]. However, the biostatistical literature in this area generally focuses on modelling, estimation and inference, as opposed to control.

Algorithms from control theory have been applied to various medical treatments. These include, for example, proportional-integral (PI) and proportional-integral-derivative (PID) control of blood glucose levels by adjusting insulin inputs [5] and for anaesthetic drug delivery [6], among other areas [1] .

In this article, we consider control of blood clotting speed using warfarin. Simplified pharmacodynamic models for the dose response have been developed using methods from both control theory [7] and biostatistics [8], whilst recent work involving two of the present authors has used these to develop new control theory approaches to dose guidance [9–11]. More generally, there is extensive research into anticoagulation dosage, with selected examples including [12–17].

The controlled variable is the International Normalised Ratio (INR), a standardised measure with high values indicating long clotting times [18]. Warfarin is a commonly prescribed anticoagulant globally, with a relatively narrow therapeutic range (2–3 INR) and high degree of inter-patient dose variability, making it a difficult drug to manage, despite being approved for human use since 1954 [19]. In fact, warfarin ranks third on the list of drugs for hospital admission due to adverse effects, with the main side effect being bleeding, which can result in fatalities [20]. Current treatment regimens utilise dose guiding algorithms, however, there is still large inter-patient variability due to genetic and clinical factors, with roughly 30% to 60% of variability being unexplainable. This is further complicated by irregular clinic recordings and missed appointments to measure INR and adjust dosage.

In control theory, it is often assumed that observations and inputs occur at fixed intervals, i.e. uniform sampling intervals. However, in practice, medical treatments are rarely measured and applied at such fixed times. Appointment times may be irregular, planned visits may not be adhered to, and individuals may drop out from trials [21]. Indeed, the warfarin data alluded to below are not uniformly sampled. Wilson *et al.* [9] used proportional-integral-plus (PIP) and model-predictive control (MPC) to investigate warfarin control when data are missing. For PIP control, an integral-of-error (IOE) term [22] was used to ensure steady state tracking of the target INR in the presence of model uncertainty and disturbances, whilst a new trapezoidal IOE form was introduced to handle missed appointments but not fully investigated.

The present work builds on [9] by investigating a wider range of possible IOE states, including forward and backward difference forms, and shows that the trapezoidal case requires an update to the control gain calculation that was not previously addressed. The simulations are based on straightforward linear difference equation models, with the focus on preliminary investigations of the new IOE forms, rather than model inference. Whilst [9] was limited to the simplest first order model, the present research is expanded to consider a second order model that might be a better representation of real patients (however, model identification is not directly considered in this article). PIP control is selected for analysis, since this provides a convenient framework to investigate the proposed IOE forms [23]. Control performance is evaluated for Monte Carlo simulations subject to stochastic disturbances, model mismatch and missed clinic appointments.

The remainder of the paper is structured as follows: section II provides the methodology; sections III and IV develop the various new IOE forms and revised pole assignment algorithm, and show how these are adapted to handle missing data; finally, the simulation results and discussion/conclusions are presented in sections V and VI respectively.

II. Methods

Data from 152 patients with chronic conditions under warfarin anticoagulation treatment were used to develop the models. 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. The gaps between visits were not regular: 67% of observations were within 15% of either 7, 14, 21 or 28 days, with 27% of intervals in excess of 32 days.

These data were used to estimate a model as detailed by [1]. The first order structure was initially chosen as it is simple, but provides a reasonable description of the data,

$$y(k) = -a_1 y(k-1) + b_1 u(k-1) + \varepsilon \tag{1}$$

where the output $y(k)$ is *log*(INR) and the input $u(k)$ is dose (mg). Based on a typical patient and obtained using the `lm` function in R, $a_1 = -0.4$, $b_1 = 0.25$ and $\varepsilon \sim \mathrm{N}(0, \sigma^2)$ with $\sigma = 0.25$. By contrast, Avery *et al.* [10] estimate the following second order model from similar data,

$$y(k) = -a_1 y(k-1) - a_2 y(k-2) + b_1 u(k-1) + b_2 u(k-2) + \varepsilon \tag{2}$$

For a typical patient, $a_1 = 0.2608$, $a_2 = -0.0901$, $b_1 = 0.1917$, $b_2 = 0.0158$ and $\sigma = 0.21$. Model identification is not the focus here, hence both models are utilised to investigate the relative performance of the new IOE forms (with patient variability investigated via Monte Carlo simulation).

A. Open loop control

The theoretical dose required to achieve the set point is obtained from the steady state gain of the model (1) or (2), i.e. $u(k) = ((1 + a_1)/b_1)d(k)$ or $u(k) = ((1 + a_1 + a_2)/(b_1 + b_2))d(k)$. Here, $d(k)$ is the target value for INR. Such open-loop control design is sensitive to modelling errors and disturbances, but provides a baseline for comparison purposes.

B. Integral of error

The integral action in PI, PID or PIP design ensures steady state tracking of the set point d (type 1 servomechanism). An integral-of-error (IOE) term is introduced, conventionally defined as $q = \int (d - y)dt$ [22] or, in discrete-time,

$$q(k) = q(k-1) + d(k) - y(k) \tag{3}$$

PIP control is usually developed within a non-minimum state space framework, in which $q(k)$ represents a state [23]. Since pole assignment is utilised in this article, for brevity the control designs are instead obtained by straightforward algebra.

C. Closed-loop PI and PIP control

For the model (1), the PIP controller reduces to PI form,

$$u(k) = -f_0 y(k) + k_I q(k) \tag{4}$$

where f_0 and k_I are the proportional and integral gains. Substituting (3) and (4) into (1), and equating with the desired response $y(k) = -p_1 y(k-1) - p_2 y(k-2) + k_I b_1 d(k-1)$ where p_1 and p_2 are coefficients chosen by the designer (e.g. via closed-loop poles on the unit circle), yields,

$$f_0 = (-a_1 - p_2)/b_1 \quad ; \quad k_I = (1 + p_1 + p_2)/b_1 \tag{5}$$

The incremental form, used for updating the dose when there is a new observation, is obtained by substituting (3) into (4),

$$u(k) = u(k-1) - f_0(y(k) - y(k-1)) + k_I(d(k) - y(k)) \tag{6}$$

Equivalent results for the second order model (2) are omitted for brevity, but yield the PIP incremental form [23],

$$\begin{aligned} u(k) = u(k-1) &- f_0(y(k) - y(k-1)) - f_1(y(k-1) - y(k-2)) \\ &- g_1(u(k-1) - u(k-2)) + k_I(d(k) - y(k)) \end{aligned} \tag{7}$$

where the control gains $\mathbf{k} = [f_0, f_1, g_1, k_I]$ are obtained by equating the closed-loop system with $y(k) = -p_1 y(k-1) - p_2 y(k-2) - p_3 y(k-3) - p_4 y(k-4) + k_I(b_1 d(k-1) + b_2 d(k-2))$ in which $\mathbf{p} = [p_1, p_2, p_3, p_4]$ are coefficients. For example, a deadbeat response, analogous to the model inversion used in regret-regression methods for adaptive treatment [8], is obtained using $p_i = 0, \forall i$. This is equivalent to setting all the poles to the origin of the complex z–plane.

D. Initialisation

In a clinical setting, the initial dose is usually based on age and other factors [18]. However, in many engineering applications, the controller (6) or (7) is initialised for $k = 1$ by assuming $y(0) = u(0) = 0$. When $y(0) \neq 0$, an alternative approach is required to avoid an initial 'kick' in the response. The present work adapts II-A e.g. $u(0) = (1 + a_1)d(0)/b_1$. For $k > 0$, the controller determines adjustments to the initial dose on the basis of the observations in the usual manner. In the following, it is assumed that $k = 0$ is not a missed observation (conceptually, this is the first clinic visit and hence first measurement, data point, and prescribed warfarin dose).

E. Simulation study

For each Monte-Carlo realisation, the plant is given by (1) or (2), with the model coefficients varied via standard deviations of 0.01. Data are simulated as missing completely at random (MCAR) [1], such that for each sample the probability that the output was missing was a defined value (e.g. 50%). The effect of deterministic load, measurement noise, and random walk disturbances, with various standard deviations, were investigated. The desired INR $y_d(k)$ was log(2.5) then log(3.5). These are the normal targets for patients with occasional and recurrent deep vein thrombosis respectively [18], while the increase mimics a change in the diagnosis. The metrics were based on 1000 simulations for each scenario.

III. Integral-of-error forms

Various IOE forms are proposed for handling missing data. In the present section, these all reduce to the standard IOE (3) when data are available. To illustrate the concept, consider the evolution of $q(k)$ from initial $q(0)$ over samples $k = 1 \to 4$ with $k = 2$ missing. The simplest approach is to ignore the missed sample in the control loop and, for the next attended clinic appointment, update $q(3)$ in the standard manner,

$$\begin{aligned} q(1) &= q(0) + d(1) - y(1) \\ q(3) &= q(1) + d(3) - y(3) \\ q(4) &= q(3) + d(4) - y(4) \end{aligned}$$

Generalising for multiple missed samples,

$$q(k) = q(k-N) + d(k) - y(k) \tag{8}$$

where N is the number of samples since the previous observation. For notational brevity, a time index for N is omitted, with its value at each sample being determined by context. In the example above, for which there is no observation at $k = 2$, then $N = 2$ for the calculation of $q(3)$. When data are not missing, as for $q(4)$, then $N = 1$ and equation (8) reduces to (3). For the first order model, the controller (6) becomes,

$$u(k) = u(k-N) - f_0(y(k) - y(k-N)) + k_I(d(k) - y(k)) \tag{9}$$

Equation (7) is revised in a similar manner.

A. Backward difference IOE

For the same $k = 1 \to 4$ example, a second approach is to update $q(2)$ using the previous output value as a surrogate for the missed observation, i.e. $q(2) = q(1) + d(2) - y(1)$ and $q(3) = q(2) + d(3) - y(3)$. We will call this the backward difference approximation. Instead of updating the control loop for the missed sample, an equivalent is to express $q(3)$ as a direct update to $q(1)$ at the next clinic appointment, i.e.,

$$q(3) = q(1) + d(2) - y(1) + d(3) - y(3) \tag{10}$$

Generalising for multiple missed samples yields,

$$q(k) = q(k-N) + \sum_{i=k-N+1}^{k} d(i) - (N-1)y(k-N) - y(k) \tag{11}$$

To illustrate for the first order model, the incremental form is,

$$\begin{aligned} u(k) &= u(k-N) - f_0(y(k) - y(k-N)) \\ &+ k_I \left(\sum_{i=k-N+1}^{k} d(i) - (N-1)y(k-N) - y(k) \right) \end{aligned} \tag{12}$$

Although $q(2)$ is stated above, this is only to derive the general form. The controller (12) is updated when there is a clinic visit. For the present example, at samples 1 ($N = 1$), 3 ($N = 2$) and 4 ($N = 1$). Conceptually, the new form provides an estimate of the total sum of errors, hence might be an improved approximation of the IOE when there are missing observations.

B. Backward difference IOE variation

A variation of the above approach is obtained by using $d(1)$ instead of $d(2)$ in the estimation of $q(2)$, hence,

$$q(3) = q(1) + d(1) + d(3) - y(1) - y(3) \tag{13}$$

Equations (10) and (13) are similar, but use different set points to handle the missed sample. Although not normally the case for warfarin, this might be significant when the set point changes. Generalising for multiple missed samples,

$$\begin{aligned} q(k) = q(k-N) + (N-1)\left(d(k-N) - y(k-N)\right) \\ + d(k) - y(k) \end{aligned} \tag{14}$$

C. Forward difference IOE and variation

Here, $q(2)$ is updated using the next available observation $y(3)$ as a surrogate for the missing $y(2)$,

$$q(2) = q(1) + d(2) - y(3) \tag{15}$$

We call this the forward difference approximation. In practice, the IOE cannot be updated until the measurement is taken at the next clinic appointment. Hence, substituting $q(2)$ into $q(3)$, yields an equivalent practical implementation form,

$$q(3) = q(1) + d(2) - y(3) + d(3) - y(3) \tag{16}$$

Generalising for multiple missed samples,

$$q(k) = q(k-N) + \sum_{i=k-N+1}^{k} d(i) - Ny(k) \tag{17}$$

A variation of equation (15) is to utilise $d(3)$ instead of $d(2)$, hence $q(3) = q(1) + 2\left(d(3) - y(3)\right)$. Generalising yields,

$$q(k) = q(k-N) + N\left(d(k) - y(k)\right) \tag{18}$$

D. Initialisation based on set point or output

This approach is based on the controller initialisation step discussed in section II-D. It does not involve the IOE, but is included here for comparison purposes. The control input at the first clinic appointment following missed observations is given by e.g. for the first order case $u(k) = (1 + a_1)d(k)/b_1$, otherwise (6) is used. Instead of $d(k)$, an alternative is to use the latest observation, i.e. $u(k) = (1 + a_1)y(k)/b_1$.

E. Initialisation of IOE

At the first clinic appointment following any missed observations, the IOE is initialised by rearranging the control algorithm. Using (1), (4) and the earlier $k = 1 \to 4$ illustration, $q(2) = (u(k) + f_0 d(2))/k_I$, and substituting $q(2)$ into $q(3)$ as usual yields $q(3) = (u(1) + f_0 d(2))/k_I + d(3) - y(3)$. A switch is required to determine $q(k)$ in the general case,

$$q(k) = \begin{cases} p(k) & \text{if } N > 1 \\ q(k-1) + d(k) - y(k) & \text{otherwise} \end{cases} \tag{19}$$

where $p(k) = (u(k-N) + f_0 d(k-1))/k_I + d(k) - y(k)$.

F. Mean IOE

For the $k = 1 \rightarrow 4$ example with $y(2)$ missing, $q(2) = q(1)+d(2)-0.5\left(y(1)+y(3)\right)$ i.e. using the mean of the next available $y(3)$ and previous $y(1)$ measurements. Generalising for multiple missed samples, $q(k) =$

$$q(k-N)+\sum_{i=k-N+1}^{k} d(i)-0.5(N-1)y(k-N)-0.5(N+1)y(k) \tag{20}$$

Note that equation (20) correctly reduces to the conventional IOE (3) for not missing samples (i.e. with $N = 1$).

IV. Trapezoidal with pole assignment

Here, the trapezoidal IOE given by equation (3) of Wilson et al. [9] is utilised at every control sample k, i.e.,

$$q(k) = q(k-1)+0.5(d(k)-y(k)+d(k-1)-y(k-1)) \tag{21}$$

With, for example, $y(2)$ missing, both $q(2)$ and $q(3)$ now require substitutions for the missed observation. One option is to adapt the trapezoidal IOE (21) with forward and backward approximations for $q(2)$ and $q(3)$ respectively, i.e.,

$$q(2) = q(1) + 0.5\left(d(1) - y(1) + d(3) - y(3)\right) \tag{22}$$

$$q(3) = q(2) + 0.5\left(d(1) - y(1) + d(3) - y(3)\right) \tag{23}$$

In a similar manner to the earlier examples, $q(2)$ cannot be updated in the control loop until the measurement $y(3)$ is taken, hence $q(3)$ is equivalently obtained by substituting (22) into (23). The same result is obtained directly using (21) and, for the same $k = 1 \rightarrow 4$ example considered earlier, yields the following IOE forms at each clinic visit,

$$\begin{aligned} q(1) &= q(0) + 0.5\left(d(0) - y(0) + d(1) - y(1)\right) \\ q(3) &= q(1) + d(1) - y(1) + d(3) - y(3) \\ q(4) &= q(3) + 0.5\left(d(3) - y(3) + d(4) - y(4)\right) \end{aligned} \tag{24}$$

In other words, the PI controller investigated by Wilson *et al.* [9], implicitly utilised the trapezoidal definition (21) for most control samples and the backward difference variation (13) for any missed samples. Generalising for multiple missing observations, the general form of the IOE is,

$$q(k) = q(k-N)+0.5N(d(k)-y(k)+d(k-N)-y(k-N)) \tag{25}$$

Wilson *et al.* [9] only considered the simpler model (1) and utilised existing computational methods in the CAPTAIN toolbox [24] to solve the linear quadratic control problem. Although this pragmatic approach worked well in practice, it ignores the different state space model associated with using (21) instead of (3), and the different closed-loop difference equation (or equivalent transfer function) subsequently obtained. By contrast, the present work algebraically solves pole assignment to obtain revised control gains.

For the case of the first order plant (1), the new proportional gain $f_0 = (-a_1 + 0.5(1 + d_1 + d_2) - d_2)/b_1$, which differs from that shown by equations (5) in section II-C, whilst k_I is unchanged. For the plant (2), the solution is concisely expressed in matrix form, i.e. $\mathbf{k} = \mathbf{\Sigma}(\mathbf{p}-\mathbf{r})$, where $\mathbf{r} = [a_1 - 1, a_2 - a_1, -a_2, 0]^T$ and,

$$\mathbf{\Sigma}^{-1} = \begin{bmatrix} b_1 & 0 & 1 & 0.5b_1 \\ b_2 - b_1 & b_1 & a_1 - 1 & 0.5(b_1 + b_2) \\ -b_2 & b_2 - b_1 & a_2 - a_1 & 0.5b_2 \\ 0 & -b_2 & -a_2 & 0 \end{bmatrix} \tag{26}$$

which is obtained in a similar manner to Appendix E of reference [23], adapted here for the new trapezoidal IOE state. Again, f_0 differs from that obtained in section II-C.

V. Simulation results

Fig. 1 uses the example of deadbeat to show that pole assignment based on (26) yields the correct design response, $y(k) = 0.5k_I(d(k-1) - d(k-2))$, whereas the uncorrected algorithm yields an oscillatory response in this extreme case. By contrast, the following examples are based poles of $(0.6579, 0.4042)$ for the model (1) and $(0.72, 0.47, -0.1955, 0)$ for (2). These are the closed-loop poles obtained from a linear quadratic optimal design approach as described by [9], solved here by pole assignment.

The IOE states presented in section III all reduce to the standard form (3) for the case that $N = 1$, hence when there are no missed observations, the response of the associated PI and PIP controllers are identical. By contrast, the approach proposed in section IV utilises (21) $\forall k$ hence the response differs slightly from (3). In practice, simulations for a range of scenarios suggest that (3) and (21) yield very similar performance in terms of the robustness to modelling errors and disturbances, when there are no missing data.

Hence, returning to the motivation for the introduction both of the trapezoidal form (21), and the variations of (3) effective in the case of missed observations, Table I and Fig. 2 show the results from an illustrative simulation scenario, for selected controllers applied to the model (1), with a high level of missed observations, subject to a random walk load disturbance and additive white noise with variance $\sigma^2 = 0.25$. The metrics in Table I are: percentage of samples for which the output was within 0.5 INR of the target; mean absolute error between target and output; and variances of the error and input.

Fig. 2 shows one closed-loop realisation associated with the results in Table I, demonstrating the high level of noise involved in these simulations (consistent with patient data and the associated noise variance estimates). The trapezoidal case is highlighted in blue for illustrative purposes. Finally, Table II and Fig. 3 show similar results for the model (2), albeit for a different simulation scenario.

For the results presented in this article, the differences between the various IOE forms are rather small, and might not be significant in practical terms, whilst the open-loop and initialise approaches generally yield much poorer performance. Given the challenging nature of warfarin control, any improvement might be welcomed and, in many cases, the forward difference (16), mean (20) and trapezoidal (21) IOE forms appear to yield the most promising results.

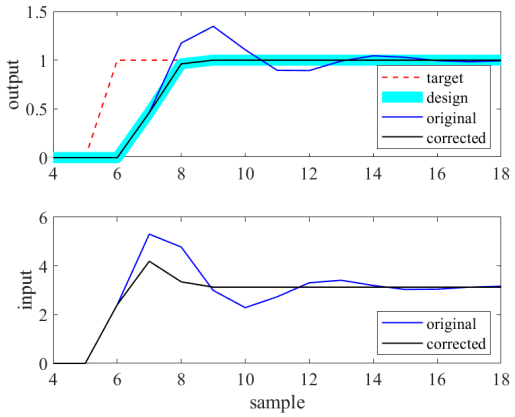

Fig. 1. Response of the second order plant (2) with trapezoidal IOE (21) for the original and corrected values of f_0. Upper: unit step set point, deadbeat design theoretical response, and outputs. Lower: control inputs.

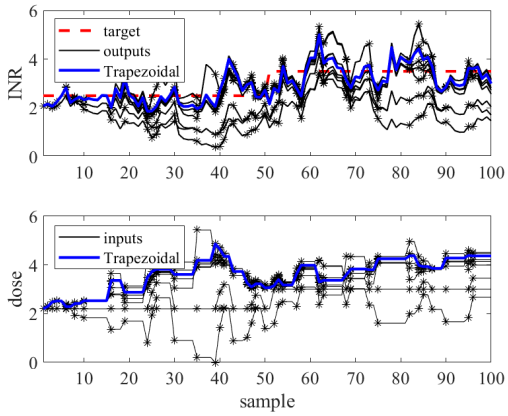

Fig. 2. Illustrative realisation from the results summarised in Table I. Upper: set point (step 2-3 INR at $k = 50$) and outputs from each controller. Lower: control inputs. Legend highlights response of one particular controller, i.e. based on trapezoidal IOE (21) with corrected f_0. Control samples with observations are shown with * (50% MCAR in this scenario).

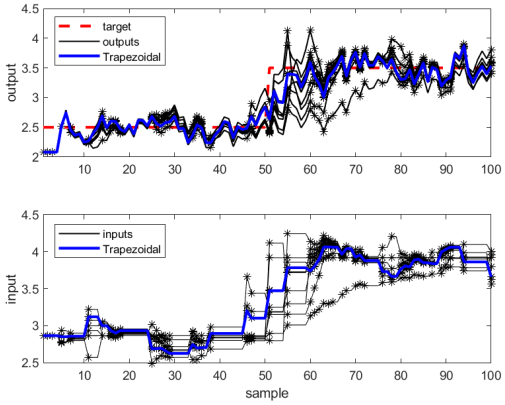

Fig. 3. Illustrative realisation from the results summarised in Table II. Similar legend to Fig. 2, except 20% MCAR in this scenario.

TABLE I
MONTE CARLO SIMULATION RESULTS FOR THE MODEL (1).

Controller	% ± 0.5 target	MAE	$var(e)$	$var(u)$
Open (II-A)	27.53	1.82	5.07	0.16
Standard (3)	63.29	0.47	0.35	0.74
Backward (11)	62.44	0.51	0.51	0.96
Forward (16)	65.50	0.49	0.92	1.02
Mean (20)	66.77	0.46	0.42	0.94
Initialise (III-D)	35.21	1.33	2.96	0.39
Trapezoidal (21)	64.31	0.49	0.46	0.96
Trapezoidal new f_0	66.61	0.46	0.42	0.93

TABLE II
MONTE CARLO SIMULATION RESULTS FOR THE MODEL (2).

Controller	% ± 0.3 target	MAE	$var(e)$	$var(u)$
Open (II-A)	32.27	0.62	0.16	0.28
Standard (3)	73.57	0.25	0.12	0.41
Backward (11)	73.78	0.22	0.10	0.43
Forward (16)	76.21	0.21	0.08	0.41
Mean (20)	77.18	0.20	0.08	0.41
Initialise (III-D)	39.83	0.50	0.15	0.32
Trapezoidal (21)	75.67	0.21	0.09	0.41
Trapezoidal new f_0	76.75	0.20	0.08	0.40

VI. CONCLUSIONS

Linear models for the response of INR to warfarin dose were used to investigate IOE forms. These were adapted to handle missing data, representing missed clinic appointments. The performance of controllers based on each IOE was evaluated by simulation. The models used were based on a 'typical' patient developed from retrospective data. For individual dose guidance, generalised models could be established for segregated patient groups, with the model subsequently updated for patient specific measurements at each clinic appointment. In combination with feedback control, this could help determine the maintenance dose more quickly and robustly than the trial-and-error approach commonly used. Preliminary results are discussed above but require further research [25]. Beyond warfarin, the novel IOE forms developed in this article may have value in other applications with missing observations.

The trapezoidal IOE (21) implicitly weights the latest and previous measurements. The authors are investigating the optimisation of these weightings e.g. in comparison to equation (24) using $q(3) = q(1) + d(2) + d(3) - \beta_1 y(1) - \beta_2 y(3)$ where β_1 and β_2 are coefficients. There are methods in the literature for using the model or interpolation for missing observations and these will be considered, as will the application of different IOE forms to MPC design [9]. Finally, in addition to the biostatistical approaches cited in section I, machine learning algorithms are recently being investigated for warfarin [26, 27] and should be compared.

The controllers in this article are based on somewhat arbitrary design criteria, i.e. closed-loop pole positions. Following further research, we envision that clinicians could alter the criteria to obtain the desired response speed and robustness depending on the patient's previous response to warfarin. The patient could be classified as 'predictable' (low

mismatch variance, which allows the use of faster controllers) or 'unpredictable' (high mismatch variance, so a slower, more robust controller might be appropriate). Priority can be given to keeping the dose variance low by maintaining the response within a wider range, or aggressively aiming for a narrow range with greater dose adjustments, aligning with existing clinical guidance for slow loading and rapid initiation regimens.

REFERENCES

[1] E. D. Wilson, Q. Clairon, R. Henderson, and C. J. Taylor. "Dealing with observational data in control". In: *Annual Reviews in Control* 46 (2018), pp. 94–106.

[2] S.A. Murphy. "Optimal dynamic treatment regimes". In: *J. R. Statist. Soc. B* 65 (2003), pp. 331–366.

[3] J. M. Robins. "Optimal structural nested models for optimal sequential decisions". In: *Proceedings Second Symposium on Biostatistics*. Ed. by D. Y. Lin and P. Heagerty. New York: Springer. 2004, pp. 189–326.

[4] B. Chakraborty and E.E.M. Moodie. *Statistical Methods for Dynamic Treatment Regimes*. NY: Springer, 2013.

[5] M. Hoekstra, M. Vogelzang, E. Verbitskiy, and M. W. N Nijsten. "Health technology assessment review: Computerized glucose regulation in the intensive care unit– how to create artificial control". In: *Critical Care* 13.5:223 (2009).

[6] G. A. Dumont. "Closed-loop control of anesthesia– a review". In: *IFAC Proc.* 45 (2012), pp. 373–378.

[7] S. Kim, A. E. Gaweda, D. Wu, L. Li, S. N. Rai, and M. E. Brier. "Simplified warfarin dose-response pharmacodynamic models". In: *Biomedical Engineering (Singapore)* 27.1 (2015).

[8] R. Henderson, P. S. Ansell, and D. Alsibani. "Regret-regression for optimal dynamic treatment regimes". In: *Biometrics* 66 (2010), pp. 1192–1201.

[9] E. D. Wilson, Q. Clairon, R. Henderson, and C. J. Taylor. "Model predictive and proportional integral control of blood clotting speed using warfarin when data are missing". In: *18th IEEE International Conference on Bioinformatics and Bioengineering (BIBE)*. Taichung, Taiwan, Oct. 2018.

[10] P. Avery, Q. Clairon, R. Henderson, C. J. Taylor, and E. Wilson. "Robust and adaptive anticoagulant control". In: *The Journal of the Royal Statistical Society, Series C (Applied Statistics)* 69.3 (2020), pp. 503–524.

[11] Q. Clairon, R. Henderson, N. Young, E. Wilson, and C. J. Taylor. "Adaptive treatment and robust control". In: *Biometrics* 77.1 (2021), pp. 223–236.

[12] C. S. Landefeld and R. J. Beyth. "Anticoagulant-related bleeding: clinical epidemiology, prediction, and prevention". In: *American Journal of Medicine* 95 (1993), pp. 315–328.

[13] P. S. Wells, A. M. Holbrook, N. R. Crowther, and J. Hirsh. "The interaction of warfarin with drugs and food". In: *Annals of Internal Medicine* 121 (2004), pp. 676–683.

[14] S. Rosthøj, C. Fullwood, R. Henderson, and S. Stewart. "Estimation of optimal dynamic anticoagulation regimes from observational data: a regret-based approach". In: *Statistics in Medicine* 25 (2006), pp. 4197–4215.

[15] U. I. Schwarz, M. D. Ritchie, Y. Bradford, C. Li, S. D. Dudek, A. Frye-Anderson, R. B. Kim, D. M. Roden, and M. Stein. "Genetic determinants of response to warfarin during initial anticoagulation". In: *New England Journal of Medicine* 358 (2008), pp. 999–1008.

[16] P. J. Avery, A. Jorgensen, A. K. Hamberg, M. Wadelius, M. Pirmohamed, and F. Kamali. "Proposal of an individualised pharmacogenetic-based warfarin initiation dose regimen for patients commencing anticoagulation therapy". In: *Clinical Pharmacology and Therapeutics* 90 (2011), pp. 701–706.

[17] B. Rich, E. E. M. Moodie, and D. A. Stephens. "Simulating sequential multiple assignment randomized trials to generate optimal personalized warfarin dosing strategies". In: *Clinical Trials* 11 (2014), pp. 435–444.

[18] 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". In: *British Journal of Haematology* 132 (2006), pp. 277–285.

[19] M. Pirmohamed. "Warfarin: almost 60 years old and still causing problems". In: *British Journal of Clinical Pharmacology* 62 (2006), pp. 509–511.

[20] M. Pirmohamed. "Adverse drug reactions as cause of admission to hospital: prospective analysis of 18 820 patients". In: *The BMJ* 329 (2004), pp. 15–19.

[21] P. Diggle. *Analysis of longitudinal data*. Oxford University Press, 2002.

[22] P. C. Young and J. C. Willems. "An approach to the linear multivariable servomechanism problem". In: *International Journal of Control* 15 (1972), pp. 961–979.

[23] C. J. Taylor, P. C. Young, and A. Chotai. *True Digital Control: Statistical Modelling and Non–Minimal State Space Design*. John Wiley and Sons, 2013.

[24] C. J. Taylor, P. C. Young, W. Tych, and E. D. Wilson. "New developments in the CAPTAIN Toolbox for Matlab with case study examples". In: *18th IFAC Symposium on System Identification (SYSID)*. Stockholm, Sweden, July 2018.

[25] D. Millevithanatchy. "PI and PIP control strategies for blood clotting speed using warfarin when data are missing". MSc Dissertation. School of Engineering: Lancaster University, 2023.

[26] S.-U Lee and P. H. Chang. "Development of a system to support warfarin dose decisions using deep neural networks". In: *Scientific Reports* 11.14745 (2021).

[27] V. Jahmunah, S. Chen, S. L. Oh, U. R. Acharya, and B. Chowbay. "Automated warfarin dose prediction for Asian, American, and Caucasian populations using a deep neural network". In: *Computers in Biology and Medicine* 153.106548 (2023).