

An Adaptive Model Predictive Controller to Address the Biovariability in Blood Clotting Response During Therapy With Warfarin

Edoardo Faggionato ¹, Alessandro Guazzo ¹, Elena Pegolo ¹, Ruggero Carli ¹, *Member, IEEE*,
Mattia Bruschetta ¹, and Simone Del Favero ¹

Abstract—Objective: Effective dosing of anticoagulants aims to prevent blood clot formation while avoiding hemorrhages. This complex task is challenged by several disturbing factors and drug-effect uncertainties, requesting frequent monitoring and adjustment. Biovariability in drug absorption and action further complicates titration and calls for individualized strategies. In this paper, we propose an adaptive closed-loop control algorithm to assist in warfarin therapy management. **Methods:** The controller was designed and tested *in silico* using an established pharmacometrics model of warfarin, which accounts for inter-subject variability. The control algorithm is an adaptive Model Predictive Control (a-MPC) that leverages a simplified patient model, whose parameters are updated with a Bayesian strategy. Performance was quantitatively evaluated in simulations performed on a population of virtual subjects against an algorithm reproducing medical guidelines (MG) and an MPC controller available in the literature (I-MPC). **Results:** The proposed a-MPC significantly ($p < 0.05$) lowers rising time (2.8 vs. 4.4 and 11.2 days) and time out of range (3.3 vs. 7.2 and 12.9 days) with respect to both MG and I-MPC, respectively. Adaptivity grants a significantly ($p < 0.05$) lower number of subjects reaching unsafe INR values compared to when this feature is not present (8.9% vs. 15% of subjects presenting an overshoot outside the target range and 0.08% vs. 0.28% of subjects reaching dangerous INR values). **Conclusion:** The a-MPC algorithm improves warfarin therapy compared to the benchmark therapies. **Significance:** This *in-silico* validation proves effectiveness of the a-MPC algorithm for anticoagulant administration, paving the way for clinical testing.

Index Terms—Adaptive control, anticoagulant, in-silico simulation, model-based control, precision medicine.

Manuscript received 11 October 2022; revised 17 February 2023; accepted 18 March 2023. Date of publication 27 March 2023; date of current version 30 August 2023. This work was supported by the MIUR, Italian Minister for Education under the initiative Departments of Excellence (Law 232/2016). (Edoardo Faggionato, Alessandro Guazzo, and Elena Pegolo contributed equally to this work.) (Corresponding author: Simone Del Favero.)

Edoardo Faggionato, Alessandro Guazzo, Elena Pegolo, Ruggero Carli, and Mattia Bruschetta are with the Department of Information Engineering, University of Padova, Italy.

Simone Del Favero is with the Department of Information Engineering, University of Padova, 35131 Padova, Italy (e-mail: sdfavero@dei.unipd.it).

Digital Object Identifier 10.1109/TBME.2023.3261962

I. INTRODUCTION

COAGULATION is the physiological process through which blood changes from a liquid to a gel state to prevent the leakage of blood from a damaged vessel. This is accomplished by the formation of fibrin threads that aggregate platelets creating the so-called blood clot. Thrombotic disorders are the pathological formation of blood clots within the vessels that might obstruct blood flow. The most severe complications associated with thrombosis include stroke, heart attack, and serious breathing problems.

Anticoagulants are drugs that act to inhibit specific pathways associated with the coagulation cascade. To date, warfarin is the most widespread anticoagulant for the prevention and treatment of thrombotic disorders. Its mechanism of action consists in the inhibition of the hepatic synthesis of vitamin K-dependent coagulation factors preventing the formation of new blood clots. Despite the great efficacy of this drug, dosing of warfarin proves to be challenging; in fact, an insufficient amount of anticoagulant might fail to prevent clot formation, while excessive dosing might cause bleeding and hemorrhages. Challenges in anticoagulant dosing include the uncertainty and fluctuation in drug sensitivity caused by the intrinsic complexity of underlying pharmacological mechanisms and the presence of several disturbing factors, such as diet, short-term physical exercise, concomitant drugs, and comorbidities. This calls for monitoring of the drug effects, possibly on a daily basis, and for consequent drug dose adjustment.

In order to monitor the effects of anticoagulant therapy, especially during its initiation, the International Normalized Ratio (INR) is frequently monitored. INR is the standardized measure of the prothrombin time that is the time required for blood to coagulate *in vitro* [1]. The typical value of INR for healthy subjects is 1 (with a range of 0.8-1.3), whereas, in patients on warfarin medication, the most common target INR is 2.0-3.0 [2]. Home monitoring of INR can be performed by the patient itself using an INR-meter, placing a drop of blood in a disposable test strip. The blood is collected through a finger-prick, making the process less painful than a blood test. Thereafter, based on the measured INR, the drug dose can be adjusted to provide the most effective treatment. To date, this adjustment is typically performed by the attending physician or by the patient following standardized medical guidelines. This process introduces a

manual feedback control loop in the therapy that contributes to mitigating the impact of fluctuations and disturbing factors. Nevertheless, manual closed-loop actions are often suboptimal, prone to errors, and burdensome for the clinician and the patient.

Recently, computer-based closed-loop drug delivery has emerged as a viable, safe, and effective therapeutic alternative to manual closed-loop in several medical applications. A notable example is blood-glucose control via subcutaneous insulin infusion in patients with type 1 diabetes [3], [4] and type 2 diabetes [5], including hospitalized patients [6] and pregnant women [7]. Another noteworthy example is automated closed-loop drugs delivery for total intravenous anesthesia [8], [9], where up to three drugs (hypnotic, analgesic, and neuromuscular blocker) are co-administered. In both the applications mentioned above, model predictive control (MPC) [10], was one of the most adopted and successful control techniques [11], [12], [13], [14], [15]. The list of closed-loop drug delivery applications based on MPC extends further, including chemotherapy for cancer treatment [16] and naltrexone automatic administration to treat fibromyalgia, a chronic pain condition [17]. These recent applications show that the MPC framework is promising for computer-based drug delivery, mainly due to its ability to account for constraints on drug dose, leveraging on a (possibly adaptive) model of the plant.

In this paper, we explored the development of a computer-based closed-loop warfarin dose delivery algorithm that leverages MPC to assist both clinicians and patients in the management of thrombotic disorders. A major challenge to effective dosing of warfarin is the high inter-individual variability response to the drug [1]. In fact, biovariability in both pharmacokinetics (PK) and pharmacodynamics (PD) of warfarin contributes up to a 20-fold difference in the drug dose requirements [18]. To cope with this large variability, clinical, demographic, anthropometric, and genetic factors are commonly employed as predictors of warfarin dose. The most used ones are age and weight; however, genetic variations are known to play a key role in warfarin dosing. Indeed, subjects with CYP2C9 and/or VKORC1 variant alleles have shown an increased risk of over-coagulation, and thus, require lower drug intakes, and more time to achieve the desired INR target [19], [20]. To date, physicians follow established medical guidelines to determine an individualized initial warfarin dose suitable for the patient and then adapt the therapy based on the subject response, thus, personalizing the manual feedback strategy [21]. This work addresses this large biovariability in the absorption and action of the anticoagulant by proposing an adaptive MPC algorithm that learns from the measurements, adapts a patient-specific model, and then uses it to formulate a patient-specific drug dose adjustment. The MPC-based control strategy is developed, tested, and compared with state-of-the-art alternatives on a simulator of the physiology of the patients. The simulator is based on a nonlinear mixed-effects model of warfarin PK and PD that the authors validated on data coming from a large dataset [18]. A key feature of this model is that it accounts for inter-patient variability thanks to a mixed-effects formulation. In a mixed-effects model framework, subjects (i.e., patient-specific model parameters) are assumed to be random realizations of a

population whose distribution is shaped around some parameters (the so-called fixed effects). Subjects are described using the deviation of each individual parameter from the population value (random effects) and possible covariates introduced in the model to explain the fraction of variability directly related to that macroscopic feature. Emulating the covariates of the original population and generating the random effects, we produced a new set of covariates-parameters, each one describing a possible *in-silico* subject. The availability of this *in-silico* population is key to realistically test control algorithms, including the variability among subjects in drug response. The simulation environment was then enriched with the addition of instrument error noise, as well as with quantization of doses in input and of INR measurements in output. Finally, the impact of poor data on the developed controllers was assessed by testing different measurement noise scenarios.

In-silico design and testing of control algorithms, already broadly used in other fields of engineering [22], is emerging as a powerful and effective tool also in biomedical engineering. In both the notable examples of closed-loop drug delivery systems mentioned above, *in-silico* testing has been widely employed: accurate simulators of subjects with type 1 diabetes were used for insulin delivery to control blood glucose [23], [24], while, an open-source simulator of patients' physiology is available and largely employed to test closed-loop control strategies in computer-assisted drug delivery for anesthesia [25].

A. State of Art

At the state of the art, some of the most used and validated algorithms that employ pharmacogenetic factors are those developed by Gage and colleagues [26], and by the International Warfarin Pharmacogenetics Consortium [27], [28]. These methods were individualized with linear regression approaches; however, in recent years, more advanced techniques have also been investigated, including data-driven algorithms based on multiple linear regression, support vector machines, and artificial neural networks [21]. Moreover, closed-loop approaches were proposed, using deadbeat, proportional integral [29], and proportional integral plus [30] controllers. From these works, the need for a flexible but still robust technique has emerged as one of the critical requirements for automated dosing of warfarin. Avery and colleagues suggested the use of MPC as an attractive solution to address this problem, both because of its ability to automatically handle constraints, and because it allows long-term strategies while keeping the flexibility to react to unexpected responses [30]. However, to the best of our knowledge, only two works employed this technique for computer-assisted warfarin administration. Pannocchia and Brambilla [31] proposed an MPC controller based on a second-order critically damped system to predict the INR response. Model parameters were preliminarily initialized, and then, estimated when at least 3-4 samples were available; after these steps, the model was no longer updated. In a more recent work, Wilson and colleagues [29] proposed an MPC controller developed using a linear time-discrete model. The controller was specifically developed to address the presence of missing data.

This second more recent algorithm is considered a baseline comparator for the performance of the adaptive MPC that we propose. Moreover, the algorithms were compared against a manual feedback strategy based on the medical guidelines.

One limitation of the literature works is associated with the validation setup. In fact, both contributions employ the same linear model used in the MPC to test their control algorithm rather than using a physiological nonlinear one. Given that no model-plant mismatch was introduced, the resulting performance evaluation is likely to be optimistic. In fact, in a real scenario, a considerable model mismatch is unavoidable due to the complexity of the underlying biological processes and the unavailable model parameters. Such model-mismatch is known to be a potentially critical aspect for model-based control applications and should be considered in the evaluation. Moreover, both papers tested the considered algorithms against only a set of possible model parameters value, i.e., one virtual patient, not accounting for intra-patient variability. The validation proposed in this paper aims at improving the realism of the performance evaluation by resorting to a validated physiological model and a population of subjects.

II. METHODS

A. Patient Simulator

1) Pharmacometric Model of Warfarin: A nonlinear mixed effects model of warfarin PK and PD [18] was employed in this work to simulate the INR response to an oral administration of warfarin. The model was developed by Hamberg and colleagues employing data collected from two studies: an Italian study (57 subjects) and a larger Swedish study (1,426 subjects). The former and a subset of the latter were merged to create a dataset of 196 subjects for model building, while the remaining 1,287 subjects were used for its validation. The developed dose-response model was able to describe the INR time course by taking into account the PK and PD biovariability that affects treatment with warfarin, together with three macroscopic features of the subjects, i.e., age, CYP2C9, and VKORC1 genotypes, to be used as predictors for model parameters. In fact, warfarin clearance is known to reduce with age, enzyme CYP2C9 is the main responsible for the metabolism of warfarin, while enzyme VKORC1 is the main target enzyme of warfarin action.

Warfarin is administered to patients in tablets containing a racemic mixture of R- and S-warfarin enantiomers, with the S-form being 3 to 5 times more efficacious [1]. The pharmacometric model is based on the assumption that only S-warfarin (and not R-warfarin) contributes to the dynamic effect of the drug; thus, it takes as input half of the administered dose. The disposition of S-warfarin is described through a PK one-compartment model with first-order absorption (parameter k_a) and first-order elimination (k_e). The amount of S-warfarin in this compartment is employed to compute the dose ratio, $DR(t)$, which is then input in a PD model. Here, $DR(t)$ is employed as the only predictor of drug-exposure, $E(t)$, in an inhibitory E_{max} model with parameters EC_{50} and γ . $E(t)$ expresses the drug-induced inhibition of the VKORC1 enzyme, which leads to the depletion of the reduced form of Vitamin K, blocking the two coagulation

pathways. Thereafter, the delay between drug exposure and INR response is modeled through two transit three-compartment chains, one with a mean transition time (MTT) of about 29 h, and the other with an MTT of about 118 h. Finally, the quantities in the last compartment of each chain determine the INR value. The compartmental structure of the model and all the formulas are reported in Fig. 1.

Inter-subject variability is described through the following stochastic models of the parameters CL , k_e , and EC_{50} :

$$\begin{cases} CL = [1 - \theta_{AGE,CL}(AGE - 71)] \cdot \theta_{CYP2C9,CL} \\ k_e = \frac{CL}{V} \cdot \exp(\eta_{k_e}) \\ EC_{50} = CL \cdot \theta_{VKORC1,EC_{50}} \cdot \exp(\eta_{EC_{50}}) \end{cases} \quad (1)$$

where CL is the warfarin clearance rate, $\theta_{AGE,CL}$ is a fixed coefficient that describes the decrease of warfarin clearance due to aging, AGE is the age of the subject, $\theta_{CYP2C9,CL}$ and $\theta_{VKORC1,EC_{50}}$ are two parameters that vary depending respectively on CYP2C9 and VKORC1 genotypes of the subject, V represents the distribution volume of the compartment Q_1 , and η_{k_e} and $\eta_{EC_{50}}$ are two independent Gaussian random variables with zero mean and standard deviation ω_{k_e} and $\omega_{EC_{50}}$, respectively. Of note, since the random effects η_{k_e} and $\eta_{EC_{50}}$ are normally distributed, then, from (1), the model parameters k_e and EC_{50} belong to a log-normal distribution, which is a common assumption for non-negative physiological parameters. In this model, the age and the two genotypes are employed to describe a portion of the inter-subject variability in a deterministic fashion, while the residual parameter variability is modeled through random effects. Except for CL , k_e , and EC_{50} , whose values depend on the realization of random variables, all the other model parameters are fixed to the value estimated in [18]. More information about the employed datasets, model identification, and model validation are available in [18].

2) Virtual Subject Generation: A virtual subject is a set of model parameters, that can in turn be seen as a realization of random effects and covariates sampled from a known distribution. Random effects were sampled from a zero-mean Gaussian distribution: $\eta_{k_e} \sim \mathcal{N}(0, \omega_{k_e}^2)$ and $\eta_{EC_{50}} \sim \mathcal{N}(0, \omega_{EC_{50}}^2)$; while covariates (i.e., AGE , $CYP2C9$, and $VKORC1$) were sampled trying to reproduce as close as possible the distribution of the original dataset where the model was identified on [18]. Therefore, AGE was drawn from a truncated Gaussian distribution with mean and standard deviation reflecting those of the original dataset. Truncation was performed by discarding samples smaller than 18 or greater than 92 years and then re-sampling a new subject. Conversely, genotypes CYP2C9, and VKORC1 were obtained with a discrete random sampling weighted by the same frequencies as those of the original dataset [18]. The result of the sampling process was a vector $[\eta_{k_e}, \eta_{EC_{50}}, AGE, CYP2C9, VKORC1]$ that, in turn, characterizes a specific subject.

3) Simulation of Warfarin Treatments: The different treatment strategies were tested on a population of virtual subjects created by sampling 10,000 sets of parameters-covariates to be employed in the pharmacometric model of warfarin described above. Multiple-day simulations were realized by integrating the




Fig. 1. Compartmental representation of the warfarin PK/PD model. $u(t)$ represents the milligrams of S-warfarin dose which is administered to the subject. The PK model is a one-compartment model with first-order absorption and first-order elimination. The amount of drug in this compartment determines the dose ratio, $DR(t)$, which is input in the PD model. The PD model consists of an inhibitory E_{max} model to compute the drug exposure $E(t)$, and two transit three-compartment chains to account for an exposure-response delay. Finally, the quantities in the last two compartments of the chains are employed to compute the $INR(t)$ value.

differential equations of the model with the MATLAB built-in solver *ode45()* [32]. Initial conditions on the first day were set to $\mathbf{Q}(0) = [0, 1, 1, 1, 1, 1, 1]$, thus, $INR(0) = INR_{base}$. INR was sampled at the end of each day and provided to the controller, which calculated the dose $u(t)$ to be administered (i.e., added to the initial condition of the first compartment Q_1 for the next day). Measurement error was added to the samples of INR assuming a Gaussian white noise with a standard deviation $\sigma_v = 0.11$ as if measures were taken from capillary blood with a portable INR-meter [33]. In addition, to set up a more realistic simulation environment, both measures of INR and warfarin inputs were quantized. In particular, $INR(t)$ was rounded up to multiples of 0.1 to recreate instrument quantization [34], while $u(t)$ was round up to multiples of 0.5 mg (i.e., 0.25 mg of S-warfarin) since, in the US, warfarin is available in 1 mg tablets that can be divided into two halves [35].

B. MPC Algorithms

Two types of MPC approaches were developed in this study: a *Standard MPC* and an *Adaptive MPC*. The models used by the two algorithms have the same structure; however, in the *Standard MPC* model parameters were fixed to population values, while in the *Adaptive MPC* model parameters were adaptively estimated in real-time on patient data as they became available. This was performed by an online linear Bayesian estimation using the previously estimated distributions as prior information.

1) Model Used by the Control Algorithms: As in [29], both controllers employed the following linear time-discrete model to predict INR response:

$$\begin{cases} x(k+1) = ax(k) + bu(k) \\ y(k) = x(k) \end{cases} \quad (2)$$

Input $u(k)$ is the dose in milligrams of S-warfarin only at the k^{th} day, whereas model output $y(k)$ is the $\log(INR(k))$. Model parameters a and b were determined by averaging the parameters estimated on data obtained from 10,000 virtual subjects, resulting in $a = 0.84$ and $b = 0.05$. The estimated distributions of a and b are reported in Fig. 2. Note that the estimated standard




Fig. 2. Boxplots of the 10,000 estimates of parameters a , Panel (a), and b , Panel (b). Final values of the two parameters were obtained by averaging these results, while mean and standard deviation were employed to construct a prior to be used by the *Adaptive MPC*.

deviations for a ($\sigma_a = 0.07$) and b ($\sigma_b = 0.03$) are considerably higher than those introduced in [29] ($\sigma_{a,wilson} = \sigma_{b,wilson} = 0.01$) to perform the simulations, showing that, in this work, algorithms were tested in a more challenging environment.

2) Standard MPC: In the *Standard MPC*, model parameters a and b were always fixed to the mean value (i.e., $a = 0.84$ and $b = 0.05$). To obtain offset-free reference tracking, an integral action was included in the controller, by resorting to a full-increment velocity form [36]. The model (2) was augmented as follows:

$$\begin{cases} \begin{bmatrix} \Delta x(k+1) \\ x(k+1) \end{bmatrix} = \begin{bmatrix} a & 0 \\ a & 1 \end{bmatrix} \begin{bmatrix} \Delta x(k) \\ x(k) \end{bmatrix} + \begin{bmatrix} b \\ b \end{bmatrix} \Delta u(k) \\ y(k) = \begin{bmatrix} 0 & 1 \end{bmatrix} \begin{bmatrix} \Delta x(k) \\ x(k) \end{bmatrix} \end{cases} \quad (3)$$

where $\Delta x(k) = x(k) - x(k-1)$ and $\Delta u(k) = u(k) - u(k-1)$.




Fig. 3. Reference INR profile, $y_r(t)$ for the controller. The day of therapy and the INR are reported on the x and y axes, respectively. On the first day, the profile reaches the target range (2-3). On the second day, the target value (2.5) is achieved. The image is zoomed on the first ten days to better appreciate the initial behavior; once the target value of 2.5 is reached it is maintained for the entire prediction horizon.

At each time instant k , the MPC controller computes the sequence of N_p control input variations $\Delta u(k|k), \dots, \Delta u(k + N_p - 1|k)$ that minimizes a cost function $J(k)$, and then, applies the first element of the sequence, $\Delta u(k|k)$, in compliance with the receding horizon principle [10]. The cost function $J(k)$ is chosen as

$$J(k) = q \left\{ \sum_{i=1}^{N_p-1} [y(k+i|k) - y_r(k+i)]^2 + x(k + N_p|k)^T S x(k + N_p|k) \right\} + \sum_{i=0}^{N_p-1} \Delta u(k+i|k)^2 \quad (4)$$

where the first term of the cost penalizes the deviation of the predicted output sequence $y(k+i|k)$ from the reference INR profile $y_r(k+i)$, while the last term penalizes large therapy adjustments $\Delta u(k+i)$. The hyperparameter q weights these two components and determines the controller aggressiveness. The prediction horizon N_p was set to 20 days. The term $x(k + N_p|k)^T S x(k + N_p|k)$ is the terminal cost introduced to promote stability, with S that is calculated by solving the Riccati equation associated with the infinite horizon unconstrained optimal control problem for the augmented system. The reference INR profile $y_r(k)$ that the controller was designed to follow is shown in Fig. 3. To optimize the performance during the transient phase, a gradual transition to the target INR = 2.5 is used as reference. The reference profile enters the target range (INR = 2) within one day and on the second day achieves the desired value (INR = 2.5).

Input dose $u(k)$ was constrained to allow a maximum warfarin intake of 20 mg (i.e., 10 mg of S-warfarin) per day and avoid

the suggestion of negative doses (apparently not administrable) resulting in the constraints:

$$0 \text{ mg} \leq u(k) \leq 20 \text{ mg} \quad \forall k.$$

These constraints are then reformulated as constraints on the optimization variable Δu . The constrained quadratic problem was implemented in MATLAB and solved using the built-in function `quadprog()`.

3) Adaptive MPC: The *Adaptive MPC* is derived from that previously described and differs from it only for the model used. The model used here shared the same structure of (2), but model parameters a and b were updated at every iteration according to a Bayesian linear regression.

In particular, from (2), the following equation can be written at the k^{th} day as

$$Y^+(k) = \begin{bmatrix} y(2) \\ y(3) \\ \dots \\ y(k) \end{bmatrix} = a \begin{bmatrix} y(1) \\ y(2) \\ \dots \\ y(k-1) \end{bmatrix} + b \begin{bmatrix} u(1) \\ u(2) \\ \dots \\ u(k-1) \end{bmatrix} = aY(k) + bU(k), \quad (5)$$

then, grouping $Y(k)$ and $U(k)$ in the same matrix $G(k)$, the previous equation becomes

$$Y^+(k) = \begin{bmatrix} y(2) \\ y(3) \\ \dots \\ y(k) \end{bmatrix} = \begin{bmatrix} y(1) & u(1) \\ y(2) & u(2) \\ \dots & \dots \\ y(k-1) & u(k-1) \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = G(k) \theta. \quad (6)$$

Then, the vector of parameters $\hat{\theta}(k) = \begin{bmatrix} \hat{a}(k) \\ \hat{b}(k) \end{bmatrix}$ can be estimated at each time step using a Bayesian strategy and specifically by minimizing the cost function

$$J_p(k) = \frac{1}{\sigma_v^2} [G(k)\hat{\theta}(k) - Y^+(k)]^T [G(k)\hat{\theta}(k) - Y^+(k)] + \lambda [\hat{\theta}(k) - \mu_p]^T \Sigma_p^{-1} [\hat{\theta}(k) - \mu_p] \quad (7)$$

with

$$\mu_p = \begin{bmatrix} 0.84 \\ 0.05 \end{bmatrix} \quad \text{and} \quad \Sigma_p = \begin{bmatrix} 0.07^2 & -0.0007 \\ -0.0007 & 0.03^2 \end{bmatrix},$$

and σ_v the previously introduced measurement noise error variance. The first addend of $J_p(k)$ accounts for the model fit of the data, while the second term introduces the available *a priori* knowledge of the parameter values in the estimation process by penalizing the deviation of the estimated parameters from their typical value. The prior distribution of the parameters θ , described by their mean and variance, μ_p and Σ_p , is estimated on the population as previously described (see Fig. 2). The impact of the two terms on the cost $J_p(k)$ is weighted by the hyperparameter λ , which should be tuned to reflect the reliability of the prior with respect to the accuracy of the measurements.

The optimal parameters θ at time k , $\hat{\theta}(k)$, that minimize $J_p(k)$, can be computed in closed-form as

$$\hat{\theta}(k) = \left[\frac{1}{\sigma_v^2} G(k)^T G(k) + \lambda \Sigma_p^{-1} \right]^{-1} \left[\frac{1}{\sigma_v^2} G(k)^T Y^+(k) + \lambda \Sigma_p^{-1} \mu_p \right]. \quad (8)$$

This Bayesian estimator has the important role of permitting a robust estimation of the two model parameters, preventing possible large deviations from the typical values caused by measurement noise. This is particularly important in the initial phase when only a few INR measurements are available.

The adaptation process starts when at least two INR measures were available. Before that, the estimates of a and b collapsed on the prior (no adaptation). Therefore, for the first two days, $\hat{\theta}(k)$ was equal to μ_p , and the *Adaptive MPC* behaved like the *Standard MPC*. Then, from the third day on, $\hat{a}(k)$ and $\hat{b}(k)$ started adapting according to the past INR response of the subject. Moreover, by rewriting the cost function J_p as

$$J_p = \frac{1}{\sigma_v^2} \sum_{i=1}^k [\hat{a}(k)y_i(k) + \hat{b}(k)u_i(k) - y_i^+(k)]^2 + \frac{\lambda}{\sigma_a^2} [\hat{a}(k) - \mu_a]^2 + \frac{\lambda}{\sigma_b^2} [\hat{b}(k) - \mu_b]^2 \quad (9)$$

it is easy to see that the weight of the first term, accounting for model fit, increases in the cost function as more and more data come in since more and more addends are included in the summation. This means that the Bayesian prior becomes less important with the reduction of the risk of large estimation errors, caused by large measurement noise corrupting a few data.

4) Hyperparameter Optimization: The proposed MPC approaches have different hyperparameters that need to be properly tuned. The *Standard MPC* approach has a single hyperparameter that needs to be optimized, namely the controller aggressiveness q that weights the prediction-reference distance $[y(k+i|k) - y_r(k+i)]^2$ and the variation of the control action $\Delta u(k+i|k)^2$ in the cost function (4). Instead, the *Adaptive MPC* approach has two hyperparameters to be optimized, namely the controller aggressiveness q , already discussed, and the prior strength λ that defines the reliability of the prior μ_p with respect to the measures $Y(k)$ in the cost function (7) used to estimate the MPC model parameters a and b .

Hyperparameters were optimized with the aim of reaching an acceptable INR value (i.e., $\text{INR} > 2$) in the shortest time window, while avoiding overshooting out of the acceptable and safe range (i.e., $\text{INR} < 3$). In order to allow for an automatic and quantitative evaluation of the INR profile obtained using a certain MPC hyperparameters configuration, a custom cost function was developed: the Weighted Out Of Target Area (WOOTA) cost function, defined as follows:

$$\text{WOOTA} = A_{\text{INR}<2} + w A_{\text{INR}>3} \quad (10)$$

where $A_{\text{INR}<2}$ is the area between the INR profile and the lower bound of the target range ($\text{INR} = 2$), and $A_{\text{INR}>3}$ is the area between the upper bound of the target range ($\text{INR} = 3$) and

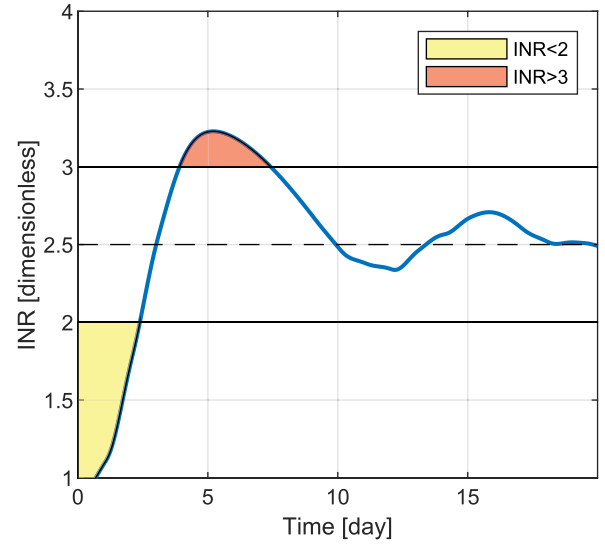


Fig. 4. Out of target area considered by the WOOTA cost function for a virtual subject. $A_{\text{INR}<2}$ is depicted in yellow meanwhile $A_{\text{INR}>3}$ is depicted in red. The total value of the WOOTA cost function is obtained by the weighted sum of the two highlighted areas.

the INR profile. By minimizing the WOOTA, it is possible to simultaneously obtain a fast rising time (small area under $\text{INR} = 2$) and limit the overshoot (small area above $\text{INR} = 3$). Finally, w is a design parameter that weights the two terms in the cost function (10). It was set to 100 to penalize the overshoot more than the raising time, favoring a less aggressive tuning of the MPC.

An example of the areas considered by the WOOTA cost function for a virtual subject INR profile is shown in Fig. 4.

The optimization step was performed by simulating 5,000 subjects, optimizing the hyperparameters for each one, and then obtaining the population hyperparameters by averaging those obtained for each subject. The best hyperparameters were selected using a random search approach after a grid search approach was used to set the extreme values of the random distributions of the hyperparameters [37]. Hyperparameters values were randomly extracted from log-uniform distributions: 1,000 possible values of q ranging from 10^2 to 10^3 were considered for the *Standard MPC*, meanwhile, 1,000 possible combinations of (q, λ) , with q ranging from 10^2 to 10^3 and λ ranging from 10^{-3} to 10, were examined for the *Adaptive MPC*. The best hyperparameters were selected as those that led to the minimum value of the WOOTA cost function.

C. Benchmark Therapies

In order to assess the different MPC algorithms developed in this work, two benchmark therapies were re-implemented and compared against the proposed control strategies.

The first benchmark therapy was the *Medical Guidelines* [38]. Usually, clinicians follow standard guidelines to tailor warfarin therapy for a specific patient. In particular, for treatment initiation, an age-adjusted protocol is used, which recommends higher dosages for younger patients. The initial phase of the

therapy generally lasts four to five days; dosage is tuned based on both the age of the patient and the morning INR measurement. Thereafter, a maintenance therapy is established as percentage-based changes in warfarin intake aiming at maintaining the INR in the target range. INR measurements are performed less frequently, usually once a week, when the desired value is achieved.

The second benchmark therapy was the MPC approach developed by Wilson and colleagues (*Wilson MPC*) [29]. Their work aimed at developing an MPC algorithm for the control of INR response focusing on not-uniformly sampled data. The *Wilson MPC* algorithm was reimplemented in MATLAB by solving the constrained quadratic problem with the built-in function *quadprog()*. The controller parameters reported in the paper were used without any adjustments.

In this work, *Medical Guidelines* from [38] and *Wilson MPC* from [29] were implemented in MATLAB reproducing as closely as possible the description available in their respective publications, and then simulated according to the procedure previously described in Section II-A3 (i.e., with a daily monitoring and therapy adjustment) in order to compare the tested strategies in a fair scenario.

D. Performance Assessment Criteria

Differences in performance between the proposed MPC approaches and the benchmark therapies were assessed considering four metrics. The first metric was the rising time which is the number of days needed for the INR profile to reach the lower acceptable INR threshold ($\text{INR} = 2$). A small rising time is preferable since the patient will spend fewer days with a too-low INR, which is related to increased blood clotting risk. The second considered metric was the out of range time which is the total number of days where the measured INR value is out of the target range ($2 \leq \text{INR} \leq 3$). This metric condenses both the effect of rising time and possible overshoots and undershoots out of the target range. Two other metrics related to patient safety were considered: the number of subjects reaching a dangerous INR value above 4 at any time during the therapy, and the number of subjects whose first maximum INR value exceeded the therapy window (i.e., above 3). Avoiding overshoot above the maximum INR threshold is key since an elevated INR value is associated with increased bleeding risk and a good INR control approach should be able to induce the INR profile to reach the target as soon as possible without any overshoot.

Performance metrics were computed for a population of 10,000 randomly-generated subjects (different than the 5,000 used for hyperparameters optimization) whose INR profiles were controlled with both the proposed MPC controllers and the benchmark therapies within a 40-day simulation. A one-tailed t-test was then used to assess whether the average of the continuous metrics (i.e., rising and out of range times) obtained using the *Adaptive MPC* was smaller than the average continuous metrics obtained using other control techniques without assuming equal variance [39]. A chi-squared test was used to assess the statistical significance of differences in dichotomous metrics (i.e., the number of subjects reaching unsafe INR values).

E. Robustness to Poor Data and Sensor Malfunctioning

In order to test the robustness of the developed control techniques against poor data and sensor malfunctioning, two scenarios including abnormally inaccurate measurements were considered. In the first scenario, a positive offset of 1 INR unit was added to the measured INR with a probability of 5%, thus producing a bimodal INR noise with peaks in 0 and 1. In the second scenario, the faults of the measurement instrument were modeled as follows: the instrument enters a faulty state with a probability of 5% and a positive offset of 1 is added to every measurement. At every new sampling time, the system can either return to the normal state with a probability of 50% or remain in the faulty one. This mechanism produces sequences of consecutive, abnormally inaccurate, measurements of random duration. The two scenarios were simulated on $N = 10,000$ subjects for all the dosing strategies and compared against the previous case of standard Gaussian white noise.

III. RESULTS

A. Adaptive Estimation

Adaptive estimation was implemented via a linear Bayesian estimator that updates the model parameters according to the past INR response of the patient. In the various panels of Fig. 5, the evolution of model parameters is reported by comparing three different values of λ , the hyperparameter that controls the adherence to the prior. As it can be noted, a low adherence ($\lambda = 0.001$) led to significant fluctuations in the behavior of a and b , undermining the model reliability in predicting the INR response; a strong adherence ($\lambda = 10$), instead, could prevent the adaptation process as the estimator struggled in deviating from the prior values. This shows the need for accurate tuning of this hyperparameter in order to grant an effective balance in the robustness and flexibility of the controller.

B. Hyperparameter Optimization Results

The procedure for hyperparameter optimization (Section II-B4) led to the following results: a comparable value of control aggressiveness q was selected as optimal for the *Standard MPC* ($q = 491$) and the *Adaptive MPC* ($q = 525$), while the optimal value of the prior confidence λ for the *Adaptive MPC* was 2.

C. Comparison Among Different Therapies

Table I shows the results for the performance metrics (rising time, out of range time, number of subjects reaching an $\text{INR} > 4$, and number of subjects with first INR maximum > 3) obtained out of the values computed from INR profiles of 10,000 simulated subjects. The proposed *Standard MPC* and *Adaptive MPC* approaches achieved lower rising times with respect to *Medical Guidelines* and *Wilson MPC* (2.7 and 2.8 vs. 4.4 and 11.2 days, respectively), and the proposed *Adaptive MPC* approach achieved the overall lowest out of range time (3.3 days), as well as the lowest number of subjects reaching dangerous values of $\text{INR} > 4$ (only 8 subjects over the 10,000

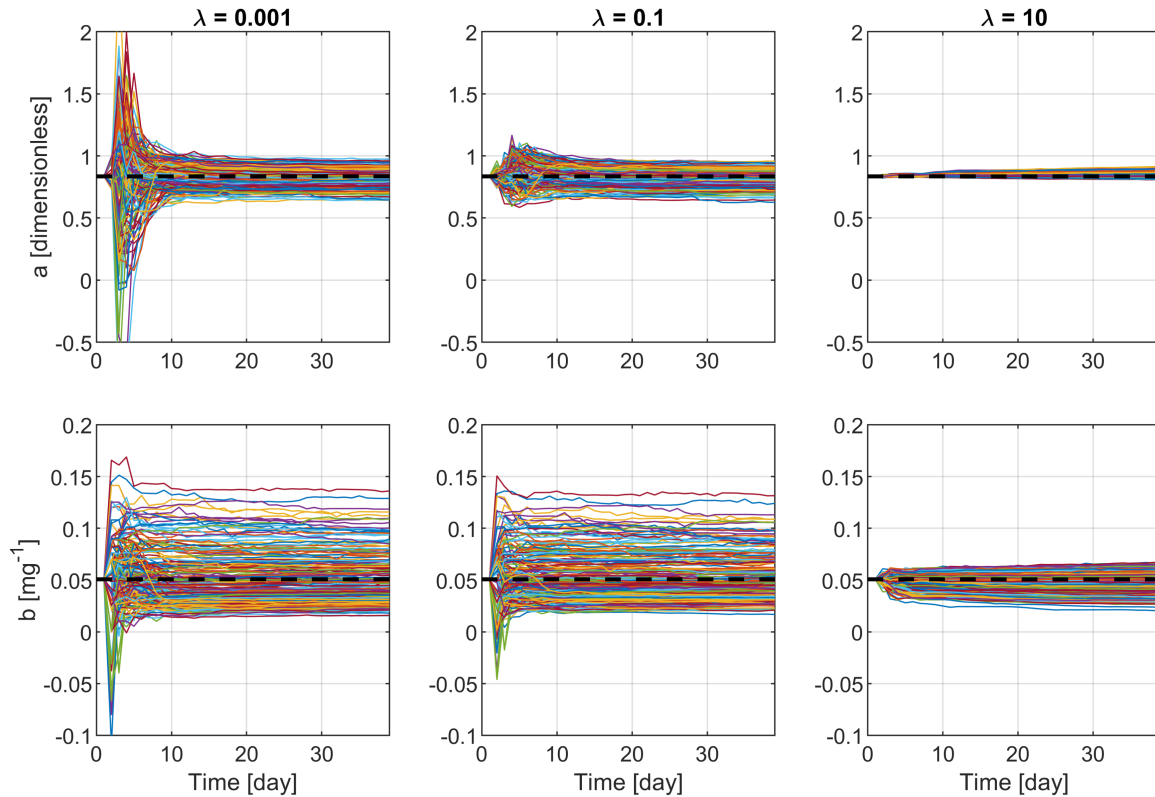


Fig. 5. Evolution of the estimates of a and b inside the *Adaptive MPC* applied on 200 virtual subjects during a 40-day simulation using different values of λ : from the left-hand panels, where the adherence to the prior is low ($\lambda = 0.001$), to the right-hand panels, where the adherence to the prior is high ($\lambda = 10$). The values of the prior are marked with a black dashed line.

TABLE I
PERFORMANCE METRICS

	Rising time [days]	Out of range time [days]	Subjects with INR>4	Subjects with 1 st max INR>3
<i>Medical Guidelines</i>	4.4 ± 2.2 *	7.2 ± 5.7 *	76 [$<1\%$] *	2451 [25%] *
<i>Wilson MPC</i>	11.2 ± 4.2 *	12.9 ± 5.5 *	19 [$<1\%$] *	872 [8.7%]
<i>Standard MPC</i>	2.7 ± 0.8	3.7 ± 2.0	28 [$<1\%$] *	1503 [15%] *
<i>Adaptive MPC</i>	2.8 ± 1.0	3.3 ± 1.5	8 [$<0.1\%$]	886 [8.9%]

Results were obtained over a population of $N = 10,000$ virtual patients. Results for rising and out of range times are reported as mean ± standard deviation. Results for reaching an unsafe INR value are reported as the number and the percentage of subjects that experienced the event. (*) highlights statistically significant increases according to a one-tailed t-test without assuming equal variance at the 0.05 confidence level with respect to the *Adaptive MPC*.

simulated) and the second lowest number of subjects with overshoot above the target window (886 over 10,000). The proposed MPC approaches outperformed both the *Wilson MPC* and the *Medical Guidelines*, with the *Adaptive MPC* being the overall best-performing control technique according to the considered performance metrics. According to a one-tailed t-test without assuming equal variance, rising and out of range times were statistically significantly ($p < 0.05$) higher for *Medical Guidelines* and *Wilson MPC*, when compared to *Adaptive MPC*. According to a chi-squared test, instead, the number of subjects reaching dangerous INR values above 4 was statistically significantly ($p < 0.05$) higher for *Medical Guidelines*, *Wilson MPC*, and *Standard MPC* when compared against *Adaptive MPC*, and the number of subjects overshooting above the target range (INR = 3) was statistically significantly higher for *Medical*

Guidelines and *Standard MPC* when compared against *Adaptive MPC*.

Fig. 6 shows a population boxplot of the paired differences obtained considering rising time, Panel (a), and out of range time, Panel (b), for *Medical Guidelines*, *Wilson MPC*, and *Standard MPC* with respect to the *Adaptive MPC*. This figure illustrates that the *Adaptive MPC* systematically outperforms both *Medical Guidelines* and *Wilson MPC*, which exhibit higher rising and out of range times in the vast majority of the subjects. The *Adaptive MPC* is comparable to the *Standard MPC* in terms of rising time, but it led to shorter out of range times in most patients.

Fig. 7 shows the INR evolution over time in the population and compares the INR achieved using the *Adaptive MPC* with respect to the comparators: the *Medical Guidelines* and the *Wilson MPC*. Median INR profiles in the population are highlighted




Fig. 6. Boxplots of performance difference of *Medical guidelines* (MG), *Wilson MPC* (WM), and *Standard MPC* (SM) with respect to *Adaptive MPC*. Rising time is shown in Panel (a), out of range time in Panel (b). The red area indicates worse performance than *Adaptive MPC* and the green area better results.

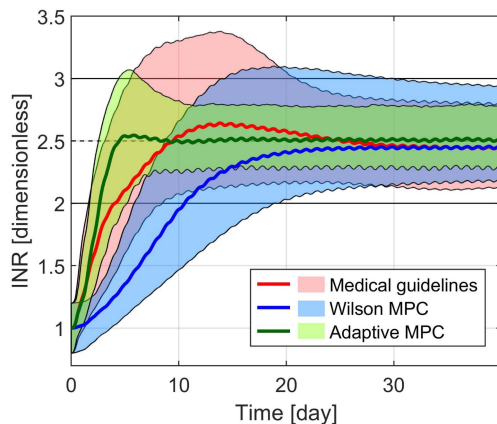


Fig. 7. Comparison of 90% prediction intervals of the simulated INR responses generated with the *Adaptive MPC* (green shaded area), the *Wilson MPC* (blue shaded area) and the *Medical Guidelines* (red shaded area). Simulations were performed on $N = 10,000$ virtual subjects. Medians are marked with a bold solid line.

with a bold line with different colors for each treatment. The shaded area is delimited by the 5-th and 95-th percentile of the INR in the population. This representation shows that the *Adaptive MPC* proposed in this study produced a superior INR profile: the first to enter the target range and leading to the smallest and shortest overshoot. *Medical Guidelines* reached the target INR soon, but at the cost of a considerable overshoot while *Wilson MPC* led to a lower but prolonged overshoot at the cost of later entrance in the INR target.

Fig. 8, compares the INR time evolution achieved by the *Adaptive MPC* with the *Standard MPC*. The *Adaptive MPC* performed better than the *Standard MPC* despite having a similar control aggressiveness q . The difference in the results was thus entirely imputable to the adaptation. In the *Adaptive MPC*, the parameters of the MPC model, a and b , were estimated from

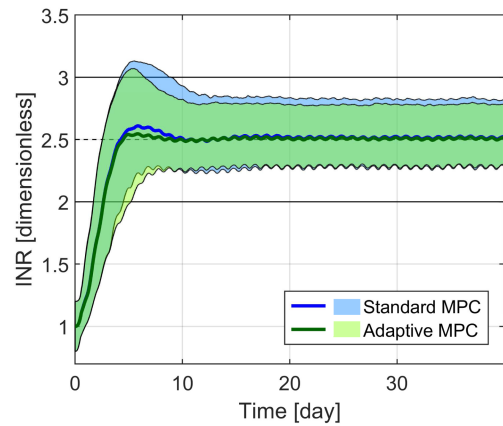


Fig. 8. Comparison of 90% prediction intervals of the simulated INR responses generated with the *Adaptive MPC* (green shaded area) and the *Standard MPC* (blue shaded area). Simulations were performed on $N = 10,000$ virtual subjects. Medians are marked with a bold solid line.

past noisy data points at each step of the simulation by iterating a Bayesian estimation procedure. This was key in the first days of simulation to avoid any overshoot while reaching the acceptable INR target as fast as possible.

D. Robustness to Poor Data and Sensor Malfunctioning

In Table II, the results of the robustness analysis are reported. The considered performance metrics (i.e., rising time, out of range time, number of subjects with $\text{INR} > 4$, and number of subjects with first max $\text{INR} > 3$) were obtained over a population of 10,000 simulated subjects for each of the four dosing strategies. Results in the considered poor data scenarios are consistent with the ones previously presented: the proposed MPC approaches outperform the benchmark therapies, reaching overall preferable performances. The adaptive MPC achieves the best performance and, even in presence of poor data, ensures the smallest number of subjects spending time in the unsafe INR ranges.

No divergent or unstable behavior was observed in none of the numerous simulated scenarios. Importantly, this holds true also for the adaptive strategy, in spite of its higher complexity.

IV. DISCUSSION

In this study, we investigated computer-assisted warfarin dosing for feedback regulation of blood coagulation. INR measurements are provided to two newly developed model-based control algorithms to suggest dose adjustments. Two different MPC approaches have been developed: a *Standard MPC* and an *Adaptive MPC*.

The first approach is an MPC implementation inspired by [29] with an optimized control aggressiveness. The second is an evolution of the former, in which the model parameters are adapted online based on the observed patient response. The adaptive updated model was then used by the MPC algorithm to address the high biovariability present in the INR response after oral warfarin administration. One of the strengths of this

TABLE II
PERFORMANCE METRICS IN THE POOR DATA SCENARIOS

		Rising time [days]	Out of range time [days]	Subjects with INR>4	Subjects with 1 st max INR>3
Bimodal noise	<i>Medical guidelines</i>	4.6 ± 2.4	7.4 ± 5.6	72 [<1%]	2199 [22%]
	<i>Wilson MPC</i>	11.7 ± 4.6	12.9 ± 5.3	13 [<1%]	716 [7.2%]
	<i>Standard MPC</i>	2.8 ± 0.9	4.6 ± 2.6	134 [1.3%]	1439 [14%]
	<i>Adaptive MPC</i>	2.9 ± 1.1	4.4 ± 2.5	61 [<1%]	988 [9.9%]
Persistent Fault	<i>Medical guidelines</i>	4.7 ± 2.7	8.9 ± 6.3	70 [<1%]	2062 [21%]
	<i>Wilson MPC</i>	12.1 ± 5.0	13.3 ± 5.5	13 [<1%]	661 [6.6%]
	<i>Standard MPC</i>	2.9 ± 1.2	5.9 ± 3.3	180 [1.8%]	1360 [14%]
	<i>Adaptive MPC</i>	3.0 ± 1.3	5.7 ± 3.2	115 [1.2%]	917 [9.2%]

Results were obtained over a population of N = 10,000 virtual patients. Results for rising and out of range times are reported as mean ± standard deviation. Results for reaching an unsafe INR value are reported as the number and the percentage of subjects that experienced the event.

work was the use of a realistic simulation environment for both tuning and validating the proposed control strategies: a nonlinear mixed-effects model of warfarin identified and validated on a large dataset was adopted to simulate the virtual subjects as realistically as possible. The developed simulation environment allowed the optimization of the controller hyperparameters, which was critical to meet satisfactory control performance. Moreover, the simulated population was then used to compare different closed-loop treatments.

The proposed MPC approaches were benchmarked against two control techniques proposed in the literature, namely the *Medical Guidelines* and the *Wilson MPC*. The *Medical Guidelines* approach was based on a MATLAB implementation of a set of rules defined in [38] that are currently used by physicians to plan patient-specific warfarin therapies. Results for the *Medical Guidelines* showed that, on average, the INR profile reached the target range in 4.4 days. However, most simulated subjects had a significant overshoot above the maximum accepted value.

The *Wilson MPC*, recently proposed by Wilson and colleagues [29], was implemented in MATLAB. In the population under study, this approach showed a non-negligible delay (11.2 days) in reaching the INR target range, and several simulated subjects had a long-lasting overshoot. This sub-optimal performance was mostly due to a conservative tuning of the controller that allowed the overshoot to be reduced at the cost of entering the therapeutic window later in time.

The *Standard MPC* approach mainly differs from the *Wilson MPC* for an optimized tuning of the control aggressiveness and a customized INR reference profile (see Section II-B2) instead of a constant value. The introduction of these changes proved to improve the controller performance in terms of rising time at the cost of a slightly increased overshoot. If compared to the *Medical Guidelines* approach, however, a reduction in out of range time is guaranteed.

The *Adaptive MPC* approach was developed to address the high biovariability, providing a personalized treatment tailored to each patient. Adaptation of the MPC model parameters proved to be robust despite the noisy INR measurements even when only a few data points were available. This was achieved thanks to a Bayesian estimation procedure, that introduced a population prior to the model parameters and allowed deviation from this prior as more measurements were collected. This feature was tested also in two challenging scenarios that involved the addition of an offset of 1 to the INR, whose values usually lie between

1 and 4. The use of a patient-tailored model in the MPC led to a more efficient control that was able to maintain the good rising time achieved by the *Standard MPC* while almost completely avoiding overshooting, in its turn, leading to a safer therapy able to avoid the reaching of dangerous INR values for almost all simulated subjects (more than 99.9%). In addition, the adaptation process can be considered completed after the fifth day in most of the subjects. The fact that adaptation does not improve the rising phase is therefore expected, in fact, the rising phase lasted on average only 3 days, less than the time needed to adapt to the patient. In all the tested simulations, including the poor data scenarios, no unstable behavior was observed in the proposed controllers. Adaptivity helped in maintaining the INR within the therapy window and was not observed to produce instability, despite its higher complexity, even during the robustness test. Finally, it is worth stressing that neither the *Standard* nor the *Adaptive MPC* required access to genotype information on the subject. Such information was instead required in the *Medical Guidelines*. Collecting genotype information is time-consuming and costly and thus this information might not be available for all patients.

IV. Limitations and Future Works

This work reports an *in-silico* analysis based on a validated non-linear model. Since all models are approximate descriptions of reality, *in-silico* results need to be confirmed by conclusive clinical investigations. Nevertheless, careful *in-silico* testing can drastically accelerate the control algorithm development pipeline and prevent expensive or dangerous clinical or preclinical tests of suboptimal algorithms. In view of this, the findings of this manuscript can be considered as a first validation step useful to guide the development of new treatments.

Known limitations of the proposed simulated framework include that it does not account for patient variability over time nor it describes the interactions between warfarin and the lifestyle of the patient. For instance, previous studies have shown that certain foods, such as grapefruit, cranberry, and vitamin K-containing vegetables, as well as lifestyle habits such as alcohol and tobacco use may lead to food-drug interactions and influence the effectiveness of warfarin [40], [41]. Even if it should be noticed that patients under warfarin are educated to take the drug between meals and avoid certain foods, future improvements to the simulation environment could include these

aspects, as well as the implementation of a sparse and more realistic INR sampling procedure. It should also be acknowledged that the implementation of *Medical Guidelines* is only partially representative of the actual actions of physicians, who may prefer to manage therapy according to their own professional experience rather than strictly following guidelines. Therefore, the superiority of automatic therapies can be proven only with real clinical trials.

The proposed controller could also be modified to address some of the above-mentioned challenges. A moving window or a forgetting factor in the adaptation process could be used to effectively address the intra-individual variability. Moreover, if the patient is willing to announce the ingestion of food with a known interaction with warfarin, the MPC could provide a natural framework to adjust the drug dosage in response to this event, by accounting for the announcement in the prediction and producing a feed-forward compensative action. Nevertheless, this requires the availability of a model for that specific interaction, either based on physiology or data-driven.

Currently, physicians may be hesitant to let patients adjust their drug dosage on a daily basis without supervision, since this is a complex task and errors could be dangerous. A thoroughly validated decision support system could be of great help in this regard. The automated computation of dose adjustment would leave only the residual risk associated with non-compliance. To minimize this risk, following [42], future work will investigate the inclusion in an MPC framework of quantized suggestion and the reduction in the number of therapy changes.

V. CONCLUSION

In conclusion, two blood coagulation control techniques based on the MPC approach were developed and benchmarked against two other methods proposed in the literature. All the analysis was performed in a simulation environment built from a nonlinear mixed effects model of warfarin PK and PD. The *Standard MPC* reached good control performances thanks to an optimized tuning of the controller aggressiveness. The more advanced *Adaptive MPC* produced the most effective control, leading to an INR profile with a fast rising time and a minimal overshoot above the maximum safe range. Both proposed control techniques outperformed the *Medical Guidelines* and the *Wilson MPC* approaches according to all the considered metrics and a visual inspection of INR profiles of the virtual subjects. This work showed the potential of an adaptive MPC strategy to address biovariability. Future developments include the testing of the proposed strategy on real patients within a randomized controlled trial, to confirm the *in-silico* findings.

REFERENCES

- [1] N. Eriksson and M. Wadelius, "Prediction of warfarin dose: Why, when and how?," *Pharmacogenomics*, vol. 13, no. 4, pp. 429–440, 2012.
- [2] A.-K. Hamberg, "Pharmacometric models for individualisation of warfarin in adults and children," Ph.D. dissertation, Acta Universitatis Upsaliensis, Uppsala, Sweden, 2013.
- [3] E. Bekiari et al., "Artificial pancreas treatment for outpatients with type 1 diabetes: Systematic review and meta-analysis," *BMJ*, vol. 361, 2018. [Online]. Available: <https://www.bmj.com/content/361/bmj.k1310>
- [4] H. Thabit and R. Hovorka, "Coming of age: The artificial pancreas for type 1 diabetes," *Diabetologia*, vol. 59, no. 9, pp. 1795–1805, 2016.
- [5] H. Thabit et al., "Closed-loop insulin delivery in inpatients with type 2 diabetes: A randomised, parallel-group trial," *Lancet Diabetes Endocrinol.*, vol. 5, no. 2, pp. 117–124, 2017.
- [6] L. Bally et al., "Fully closed-loop insulin delivery improves glucose control of inpatients with type 2 diabetes receiving hemodialysis," *Kidney Int.*, vol. 96, no. 3, pp. 593–596, 2019.
- [7] Z. A. Stewart et al., "Closed-loop insulin delivery during pregnancy in women with type 1 diabetes," *New England J. Med.*, vol. 375, no. 7, pp. 644–654, 2016.
- [8] L. Pasin et al., "Closed-loop delivery systems versus manually controlled administration of total IV anesthesia: A meta-analysis of randomized clinical trials," *Anesth. Analg.*, vol. 124, no. 2, pp. 456–464, 2017.
- [9] M. Ghita, M. Ghita, and D. Copot, "An overview of computer-guided total intravenous anesthesia and monitoring devices—Drug infusion control strategies and analgesia assessment in clinical use and research," *Automated Drug Del. Anesth.*, pp. 7–50, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/B9780128159750000072>
- [10] L. Wang, *Model Predictive Control System Design and Implementation Using MATLAB*. Berlin, Germany: Springer, 2009.
- [11] R. Hovorka et al., "Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes," *Physiol. Meas.*, vol. 25, no. 4, pp. 905–920, 2004.
- [12] C. Toffanin et al., "Artificial pancreas: Model predictive control design from clinical experience," *J. Diabetes Sci. Technol.*, vol. 7, no. 6, pp. 1470–1483, 2013.
- [13] R. Gondhalekar, E. Dassau, and F. J. Doyle III, "Periodic zone-MPC with asymmetric costs for outpatient-ready safety of an artificial pancreas to treat type 1 diabetes," *Automatica*, vol. 71, pp. 237–246, 2016.
- [14] X. Sun et al., "Incorporating prior information in adaptive model predictive control for multivariable artificial pancreas systems," *J. Diabetes Sci. Technol.*, vol. 16, no. 1, pp. 19–28, 2022.
- [15] I. Nascu, R. Oberdieck, and E. N. Pistikopoulos, "Explicit hybrid model predictive control strategies for intravenous anaesthesia," *Comput. Chem. Eng.*, vol. 106, pp. 814–825, 2017.
- [16] T. Chen, N. F. Kirkby, and R. Jena, "Optimal dosing of cancer chemotherapy using model predictive control and moving horizon state/parameter estimation," *Comput. Methods Prog. Biomed.*, vol. 108, no. 3, pp. 973–983, 2012.
- [17] S. Deshpande et al., "Optimized treatment of fibromyalgia using system identification and hybrid model predictive control," *Control Eng. Pract.*, vol. 33, pp. 161–173, 2014.
- [18] A.-K. Hamberg et al., "A pharmacometric model describing the relationship between warfarin dose and INR response with respect to variations in CYP2C9, VKORC1, and age," *Clin. Pharmacol. Therapeutics*, vol. 87, no. 6, pp. 727–734, 2010.
- [19] Z. Ma et al., "Clinical model for predicting warfarin sensitivity," *Sci. Rep.*, vol. 9, no. 1, pp. 1–9, 2019.
- [20] L. Dean, "Warfarin therapy and the genotypes CYP2C9 and VKORC1," *Med. Genet. Summaries*, vol. 2, pp. 257–263, 2012.
- [21] I. G. Asiimwe et al., "Warfarin dosing algorithms: A systematic review," *Brit. J. Clin. Pharmacol.*, vol. 87, no. 4, pp. 1717–1729, 2021.
- [22] B. Kim et al., "Testing autonomous vehicle software in the virtual prototyping environment," *IEEE Embedded Syst. Lett.*, vol. 9, no. 1, pp. 5–8, Mar. 2017.
- [23] M. E. Wilinska et al., "Simulation environment to evaluate closed-loop insulin delivery systems in type 1 diabetes," *J. Diabetes Sci. Technol.*, vol. 4, no. 1, pp. 132–144, 2010.
- [24] R. Visentin et al., "The university of Virginia/Padova type 1 diabetes simulator matches the glucose traces of a clinical trial," *Diabetes Technol. Therapeutics*, vol. 16, no. 7, pp. 428–434, 2014.
- [25] C. M. Ionescu et al., "An open source patient simulator for design and evaluation of computer based multiple drug dosing control for anesthetic and hemodynamic variables," *IEEE Access*, vol. 9, pp. 8680–8694, 2021.
- [26] B. F. Gage et al., "Use of pharmacogenetic and clinical factors to predict the therapeutic dose of warfarin," *Clin. Pharmacol. Therapeutics*, vol. 84, no. 3, pp. 326–331, 2008.
- [27] J. A. Johnson et al., "Clinical pharmacogenetics implementation consortium guidelines for CYP2C9 and VKORC1 genotypes and warfarin dosing," *Clin. Pharmacol. Therapeutics*, vol. 90, no. 4, pp. 625–629, 2011.

- [28] International Warfarin Pharmacogenetics Consortium, "Estimation of the warfarin dose with clinical and pharmacogenetic data," *New England J. Med.*, vol. 360, no. 8, pp. 753–764, 2009.
- [29] E. D. Wilson et al., "Model predictive and proportional integral control of blood clotting speed using warfarin when data are missing," in *Proc. IEEE 18th Int. Conf. Bioinf. Bioeng.*, 2018, pp. 22–27.
- [30] P. Avery et al., "Robust and adaptive anticoagulant control," *J. Roy. Stat. Society: Ser. C. (Appl. Statist.)*, vol. 69, no. 3, pp. 503–524, 2020.
- [31] G. Pannocchia and A. Brambilla, "Model predictive control for optimal oral anticoagulant drug administration," *AICHE J.*, vol. 52, no. 9, pp. 3315–3320, 2006.
- [32] Documentation of MATLAB, version 2020b, Accessed : Jan. 15, 2022. [Online]. Available: <https://it.mathworks.com/help/releases/R2020b>
- [33] *CoaguChek XS System User Manual*, Roche Diagnostics Corporation, Indianapolis, IN, USA, Accessed: Jan. 15, 2022. [Online]. Available: <https://diagnostics.roche.com/us/en/resource-center-pages/coaguheck-xs-system-user-manual-04837991001-07.html>
- [34] J. L. Fuller and E. Miller, "Evaluation of the accuracy and precision of the CoaguChek XS system," *Roche Diagnostics Corporation*, Indianapolis, IN, USA, 2006, Accessed: Jan. 15, 2022. [Online]. Available: <https://diagnostics.roche.com/us/en/resource-center-pages/coaguheck-xs-accuracy-and-precision-white-paper-573-34186-1106.html>
- [35] *Anticoagulant, COUMADIN TABLETS (Warfarin Sodium Tablets, USP) Crystalline, COUMADIN FOR INJECTION (Warfarin Sodium for Injection, USP)*, Bristol-Myers Squibb Company, Princeton, NJ, USA, 2007, Accessed: Mar. 30, 2023. [Online]. Available: https://www.accessdata.fda.gov/drugsatfda_docs/label/2007/009218s1051blv2.pdf
- [36] G. Pannocchia, "Offset-free tracking MPC: A tutorial review and comparison of different formulations," in *Proc. IEEE Eur. Control Conf.*, 2015, pp. 527–532.
- [37] J. Bergstra and Y. Bengio, "Random search for hyper-parameter optimization," *J. Mach. Learn. Res.*, vol. 13, no. 10, pp. 281–305, 2012.
- [38] P. A. Tideman et al., "How to manage warfarin therapy," *Australian Prescriber*, vol. 38, no. 2, pp. 44–48, 2015.
- [39] G. D. Ruxton, "The unequal variance t-test is an underused alternative to Student's t-test and the Mann–Whitney U test," *Behav. Ecol.*, vol. 17, no. 4, pp. 688–690, 2006.
- [40] A. Chock, J. Stading, and E. Sexson, "Food and lifestyle interactions with warfarin: A review," *U. S. Pharmacist*, vol. 34, no. 2, pp. 28–39, 2009.
- [41] M. A. Johnson, "Influence of vitamin K on anticoagulant therapy depends on vitamin K Status and the source and chemical forms of vitamin K," *Nutr. Rev.*, vol. 63, no. 3, pp. 91–97, 2005.
- [42] J. Pavan et al., "Incorporating sparse and quantized carbohydrates suggestions in model predictive control for artificial pancreas in type 1 diabetes," *IEEE Trans. Control Syst. Technol.*, vol. 31, no. 2, pp. 570–586, Mar. 2023.

Open Access provided by 'Università degli Studi di Padova' within the CRUI CARE Agreement