Chemotherapy Optimization using Moving Horizon Estimation based Nonlinear Model Predictive Control ⋆

Bence Czakó * **Máté Siket** * **Dániel András Drexler** *
Levente Kovács *

* *Physiological Controls Research Center, Research and Innovation Center of Óbuda University, Óbuda University, Hungary (e-mails:czako.bence, siket.mate, drexler.daniel, levente.kovacs@uni-obuda.hu).*

Abstract: A Moving Horizon Estimator (MHE) based Nonlinear Model Predictive Controller (NMPC) was designed for an impulsive minimal tumor growth model. The estimator computes the time-varying model parameters using mean square error with parameter deviation penalization and provides state estimations for the controller. The controller computes optimal doses for non-equidistant, fixed time instants while constraining the administered drug dose. Tuning of the MHE was based on experimental time series measurements, while for the NMPC a virtual population was generated. The robustness of the combined approach was tested in silico on a virtual population, where the simulation was tailored to a real experimental scenario.

Keywords: Physiological Model; Parameter and state estimation; Nonlinear predictive control; Robustness analysis

1. INTRODUCTION

Chemotherapy is one of the most efficient ways to deal with cancerous diseases in everyday clinical practice, especially for metastatic tumors. Healthcare professionals use empirical protocols to determine the course of therapy for a given patient. It was conjectured early on that by employing mathematical models these treatment protocols can be optimized, which can alleviate the side effects of chemotherapy. Therefore, in the past decades, a vast effort was put into the development of mathematical models and algorithms that can accurately compute these optimal protocols, as it was indicated in the reviews by Sbeity and Younes (2015) and Shi et al. (2011).

A remarkable aspect of the optimized therapies could be the prolonged effectiveness of a chemotherapeutic agent before drug resistance occurs in the patient. During conventional therapy, a maximum tolerable dose (MTD) is given to the patient with a couple of weeks of duration in between therapeutic sessions. One issue is that the therapy frequently causes cell mutations in the tumor such that they render the drug ineffective. It is theorized that by applying a low dose therapy with more frequent administration sessions, which is called metronomic therapy, the onset of drug resistance could be delayed (Carrère (2017)).

Our goal is to implement a model-based algorithm in the context of metronomic therapy which will be verified using mice experiments afterward. The algorithm is based on a simple reaction kinetics model which was established in Drexler et al. (2020) and was fitted to experimental data using nonlinear mixed-effects modeling. One limitation of the model is that it does not incorporate the time-varying nature of the tumor cell population in its parameters and assumes them to be constants. In our work, we aim to overcome this issue by online estimating a subset of the model parameters using Moving Horizon Estimation (MHE), based on Siket et al. (2020). This estimator is combined with a Nonlinear Model Predictive Controller (NMPC) which computes the optimal treatment protocol, adapted from Czakó et al. (2020). In our setting, we assume that the therapy days are given in advance and the administrations are carried out impulsively. Similar approaches can be found in Chen et al. (2012) and Belfo and Lemos (2021). We also constrain the maximum amount both for a single dose and for a given time frame so that cumulative toxicity can be minimized. The parameters of the algorithm were tuned using the experimental data from Füredi et al. (2017), which is also utilized for the generation of a virtual population for our simulation study.

In Section 2, we recall the tumor model used in our study and explain its terms briefly. In Section 3, the algorithm is introduced with a separate exposition on the MHE and the NMPC. In Section 4, we describe the

⋆ Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 679681). Project no. 2019-1.3.1-KK-2019-00007. has been implemented with the support provided from the National Research, Development and Innovation Fund of Hungary, financed under the 2019-1.3.1-KK funding scheme. Bence G. Czakó and Máté Siket were supported by the ÚNKP-20-3 New National Excellence Program of the Ministry for Innovation and Technology from the source of the National Research, Development, and Innovation Fund and the Eötvös Loránd Research Network Secretariat under grant agreement no. ELKH KÖ-40/2020 ('Development of cyber-medical systems based on AI and hybrid cloud methods').

tuning of the algorithms and show the stability of the combined framework on virtual populations. Conclusions and directions for further research can be found in Section 5.

2. TUMOR GROWTH MODEL

The tumor growth model was proposed in Drexler et al. (2020) and consists of four differential equations

$$\begin{aligned}\dot{x}_1 &= (a-n)x_1 - b\frac{x_1 x_3}{ED_{50}+x_3}\\ \dot{x}_2 &= nx_1 + b\frac{x_1 x_3}{ED_{50}+x_3} - wx_2\\ \dot{x}_3 &= -(c+k_1)x_3 + k_2 x_4\\ \dot{x}_4 &= k_1 x_3 - k_2 x_4\end{aligned} \tag{1}$$

where x_1 is the living tumor volume [mm^3], x_2 is the dead tumor volume [mm^3], x_3 is the drug level in the central compartment [mg/kg], and x_4 is the drug level in the peripheral compartment [mg/kg]. The output of the system, which is the complete tumor volume, is given by $y = x_1 + x_2$. The underlying model represents a set of impulsive differential equations, where the drug is administered under an infinitesimal time instant into the subject. We denote the time of administrations with t_i, $i \in \mathbb{N}_0$ with $t_0 < t_1 < \cdots < t_i < t_{i+1} < \ldots$, which constraints the system dynamics on the intervals $(t_0, t_1], (t_1, t_2], \ldots, (t_i, t_{i+1}], \ldots$ with the rule

$$\boldsymbol{x}(t_i^+) = \boldsymbol{x}(t_i^-) + (0\ 0\ 1\ 0)^\top u_i \tag{2}$$

where u_i is the amount of drug injected into the subject at time t_i. The initial condition of the system is often given as $x(t_0) = (x_0, 0, u_0, 0)^\top$ with a strictly positive initial living tumor volume $x_0 > 0$ and arbitrary $u_0 \geq 0$. One can intuitively think about this impulsive system as solving an initial value problem at the beginning of each interval t_i^+, where the initial values are the state of the system at the end of the previous interval, t_i^-, modified by adding u_i amount of drug to $x_3(t_i^-)$.

3. CONTROL ALGORITHM

As it was pointed out in the introduction, our algorithm consists of an estimator and a controller. The MHE can online estimate the parameters of the model and provide a full state estimate to the controller. The controller then uses this information to predict the future evolution of the system from which, an optimal dosage sequence is calculated.

3.1 Moving Horizon Estimation

MHE is an optimization based state and parameter estimation algorithm. During each measurement update, the algorithm minimizes a cost function on a sliding window, which contains a number of previous measurements. It can be thought of as a generalization of the Kalman filter with the additional benefit that it can incorporate various constraints during the optimization.

In this particular case, the cost function is defined to be the sum of two terms which are trajectory error and parameter difference. The former one is the mean-squared error between the measured and estimated volumes in the current time step. The latter one is the difference between the optimized model parameters and a nominal parameter set, which is the fixed effect ($\boldsymbol{\theta}$) of the model, contained in Table 1. This term is beneficial concerning the identifiability properties of the system since it limits the number of possible solutions. Only the most sensitive parameters, $\hat{\boldsymbol{p}} = (a, b, n, w)^\top$, of the model (1) are estimated; for the remaining parameters their nominal value is used. The cost function at time t_i is defined as

$$\begin{aligned}\min_{\hat{\boldsymbol{p}} \in \mathbb{R}^4} \quad & J_M(\hat{\boldsymbol{p}}; t_i) = \sum_{k=i-M}^{i} \Delta\hat{y}_k^2 + d\sum_{l=1}^{4}\left(\frac{\Delta p_l}{p_l}\right)^2\\ \text{s.t.} \quad & \boldsymbol{p} \in [\underline{\boldsymbol{p}}, \bar{\boldsymbol{p}}],\\ & \boldsymbol{A}\boldsymbol{p} \leq c_i,\end{aligned} \tag{3}$$

where $M > 0$ is the length of the horizon, $d > 0$ is a scalar tuning parameter, $\Delta\hat{y}_k = \tilde{y}_k - \hat{y}_k$ where $\tilde{y}_k$ is the measured tumor volume, $\hat{y}_k$ is the estimated volume, which is the output from the solution $(\hat{\boldsymbol{x}}(t), t \in (t_{i-M}, t_i])$ of model (1) at time t_k, $\Delta p_l = p_l - \hat{p}_l$ where $p_l \in \boldsymbol{p}$ is the l-th element of the nominal parameter vector $\boldsymbol{p}$, and $\hat{p}_l \in \hat{\boldsymbol{p}}$ is the parameter vector which we optimize over. In the simplest case where no measurement error is introduced $\tilde{y}_k = y_k$. To obtain the solution $\hat{\boldsymbol{x}}(t)$, the model is integrated from the initial condition $\hat{\boldsymbol{x}}(t_{i-M}^+)$ obtained from rule (2) with $\hat{\boldsymbol{x}}(t_0^+) = (x_0, 0, u_0, 0)^\top$. As such, we integrate the model numerically with the optimal parameter set at each measurement time instant on the full horizon and save the endpoint value of the estimate for the next optimization task at the new measurement.

Each parameter is constrained to be in the box defined by $\underline{\boldsymbol{p}} = (0\ 0.01\ 0.1\ 0)^\top$ and $\bar{\boldsymbol{p}} = (2\ 1\ 1\ 1)^\top$. An additional adaptive linear constraint is used to avoid unstable parameter sets if possible. The theoretical possibility of tumor volume decrease upon administration can be guaranteed (see Drexler et al. (2018)) if the parameters fulfil that $a - b - n < 0$ which entails $\boldsymbol{A} = (1\ {-1}\ {-1}\ 0)$. Under ideal conditions $c_i = 0$, however, if the measured volume is more than 20% larger than the estimated, the constraint is loosened in an iterative way by incrementing c_i with 0.1 and restarting the optimization. The adaptive constraint guarantees stable parameter set under normal conditions, and provides better estimation when abrupt changes occur in the physiology of the tumor.

3.2 Nonlinear Model Predictive Control

The NMPC takes the solution $\hat{\boldsymbol{p}}$ of problem (3) and the estimated states $\hat{\mathbf{x}}(t)$ of system (1) and predicts its evolution on the (not necessary equidistant) intervals, $(t_i, t_{i+1}], i \in \{0, \ldots, N\} \subset \mathbb{N}_0$ with N being the number of prediction intervals. To each starting point of these intervals an input u_i is assigned, according to rule (2), which will be the subject of the optimization. One can think about this prediction as the continuation of the solution $\hat{\mathbf{x}}(t)$ from the MHE with the computed vector $\hat{\boldsymbol{p}}$, and parametrized by each input signal. The stage cost is defined as

$$\ell(y_i, u_i) = q\int_{t_i}^{t_{i+1}} \left(\frac{y_i - y_{\text{ref}}}{\tilde{y}_0}\right)^2 \mathrm{d}t + r\left(\frac{u_i}{\bar{u}_i}\right)^2, \tag{4}$$

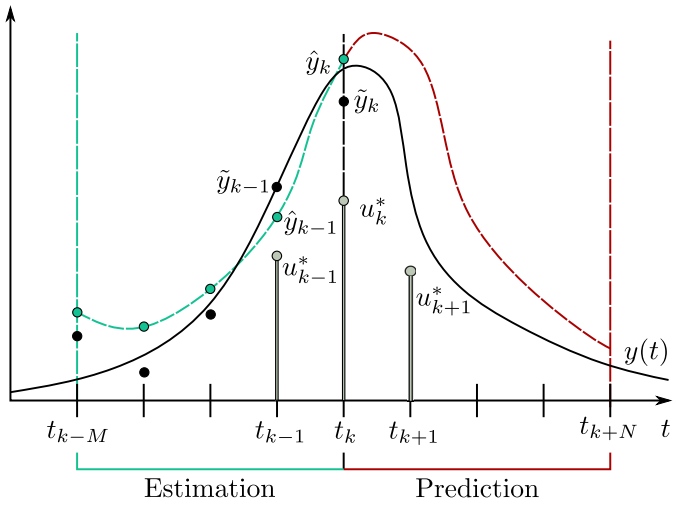

Fig. 1. Schematic diagram of the proposed approach. The evolution of the tumor volume (black line) also represents one growth-shrink cycle in Section 4.1

where q and r are scalar control parameters, the output y_i is defined on the interval $(t_i, t_{i+1}]$ by the solution of model (1), y_{ref} is the reference tumor volume, $\tilde{y}_0$ is the measured tumor volume at the beginning of the treatment, u_i is the term in (2), and $\bar{u}_i$ is the maximum admissible drug in a single administration. In practice, the model is solved numerically while the integral is approximated by the trapezoidal rule. The optimization problem at time t_k can now be formulated as

$$\begin{aligned} \min_{\boldsymbol{u} \in \mathbb{R}^N} \quad & J_N(\boldsymbol{u}; t_i) = \sum_{k=0}^{N-1} \ell(y_{i+k}, u_{i+k}) \\ \text{s.t.} \quad & \boldsymbol{u} \in [\boldsymbol{0}, \bar{\boldsymbol{u}}], \\ & \boldsymbol{u^c} \in [\boldsymbol{0}, \bar{\boldsymbol{u}}^{\boldsymbol{c}}], \end{aligned} \tag{5}$$

where $\boldsymbol{u} = (u_i, u_{i+1}, \dots u_{i+N-1})$ is the optimal input sequence where each element is constrained to lie in the interval $0 \leq u_{i+k} \leq \bar{u}$, which entails $\bar{\boldsymbol{u}} = \bar{u}\boldsymbol{1}$ (where $\boldsymbol{1}$ is an N dimensional column vector with all of its entries being equal to 1), and the second group of constraints $0 \leq u^c_{i+k} \leq \bar{u}^c$ denote the cumulative dosage associated with each u_{i+k}, from which $\boldsymbol{u^c} = \bar{u}^c\boldsymbol{1}$. The cumulative dosage u_i^c is defined here as the sum of doses in the past 10 days,

$$u_i^c = \sum_J u_j, \;\; J := \{\, j \mid t_i - 10 \leq t_j \leq t_i \,\}. \tag{6}$$

According to Füredi et al. (2017) the maximal tolerable dose (MTD) of PLD in mice is 8 mg/kg which could be repeated in every 10 days without triggering an irreversible weight loss. To define a safe cummulated dose threshold in our method, the maximum given PLD of 16 [mg/kg] was lowered to $\bar{u}_i^c = 14$ [mg/kg] in 10 days to minimize the possibility of severe systemic toxicity. The upper bound for each dose was set to $\bar{u}_i = 6$ [mg/kg] and is smaller than the 8 [mg/kg] which is the MTD that is given to a subject during a conventional course of therapy.

In Figure 1, one can see the combined approach graphically. The black line is the evolution of the tumor volume ($y(t)$), black circles correspond to noisy measurements ($\tilde{y}_k$), turquoise circles are estimates by the MHE ($\hat{y}_k$), grey lines are the optimal administrations (u_k^*) and the carmine line is the prediction of the tumor evolution. Note that the prediction is a continuation of the MHE estimate $\hat{y}_k$ and its shape is determined by the estimated parameters $\hat{\boldsymbol{p}}$ and the optimal input sequence $\boldsymbol{u}$.

4. NUMERICAL EXPERIMENTS

In order to tune the parameters of our algorithm and assess its performance, we conducted several numerical simulations. While there exists no separation principle for nonlinear systems, we chose to tune the observer and the controller separately to mitigate the computational burden of the tuning procedure. To obtain realistic parameters, we used the experimental data from Füredi et al. (2017). The experimental data contain 10 time series of Murine breast cancer evolution, treated with Pegylated Liposomal Doxorubicin (PLD). In the case of the MHE, it was previously shown in Siket et al. (2020), that satisfying performance can be achieved with a window $M = 14$. Since in this paper we used a different nominal parameter set, a value of $d = 20$ leads to similar results. For the NMPC, a virtual population was generated by fitting the time series to the model (1), which we describe here in detail. The combined algorithm is then tested on a different virtual population with added noise to incorporate the effect of real measurements.

4.1 Virtual population generation

For tuning the control parameters q and r in (4), a virtual population was generated, which represents a number of different parameter combinations for the model (1). To obtain a realistic population, we fitted the model using the Stochastic Approximation Expectation Maximization (SAEM) algorithm to experimental data, similar to Drexler et al. (2020). In this paper, instead of fitting the parameters of the model to the whole time series, we cut the time series into multiple intervals empirically, each containing one growth-shrink cycle (similar to the black curve in Figure 1). This means that each cycle begins with tumor growth which is then subject to injection treatment and then followed by a shrink to a quasi-stationary value. An example can be seen in Figure 2, where the grey vertical lines are the cuts. We also excluded resistant artifacts, where the injection does not lead to remission. The reason behind cutting the time series is that the tumor cell population is time-varying in nature, which is not reflected by constant model parameters. By cutting the time series, we implicitly assume that the model parameters are piecewise constant during each cycle, which leads to better fit and computationally more tractable than fitting time-varying functions for each full time series. In order to perform the fit, we used **sbiofitmixed** routine with **nlmefitsa** option in MATLAB 2021a and the **scattersearch** algorithm of **sbiofit** to calculate a proper initial value for the mixed-effects fit. We have also used the identified pharmacokinetic parameters from Drexler et al. (2020). The tumor volume was approximated in this article with

$$y = \frac{\pi}{3}(lw)^{(3/2)}, \tag{7}$$

where l is the length, w is the width of the tumor. This approximation is required since in the experiment they measure only the length and the width of the tumor with a

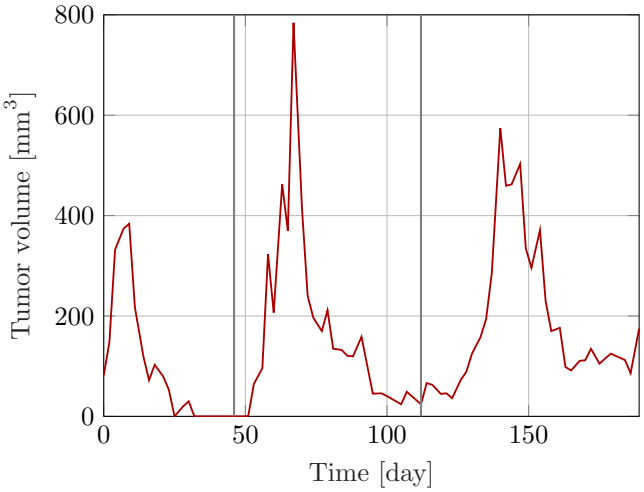

Fig. 2. An example cut of the time series **PLD 6** from Füredi et al. (2017). Between each grey line a single growth-shrink cycle is contained which is assumed to be independent from the rest of the time series.

caliper. The form of the equation is a proper choice for the approximation of caliper measurements, as it was shown in Sápi et al. (2015).

The estimated fixed effects (θ^*) are contained in Table 1 with their standard errors, and random effect variance coefficients. For c and k_1, the low values of the random effects can be attributed to the fact that we initialized their values from the results of the previous pharmacokinetic parameter estimation.

We assumed that one parameter sample is from the normal distribution $\theta_i^* \sim \mathcal{N}(\theta^*, \Sigma^*)$, where θ_i^* is the generated parameter set, θ^* is the estimated fixed effects in Table 1., and Σ^* is the random effect covariance matrix, which is in our case a diagonal matrix with elements from the last column in Table 1. Because the values in Table 1 are log-transformed, one must back-transform the calculated parameters θ_i^*-s to obtain their values in the original space i.e. $\theta_i = \exp(\theta_i^*)$. For each sample we draw the corresponding initial value from the uniform distribution $x_0 \sim \mathcal{U}(l_x, u_x)$, where the parameter l_x, u_x corresponds to the smallest and largest tumor volume where the first dose was applied in the growth-shrink cycles. We further generated statistics about the cycles which we used to impose certain conditions during the generation such that the generated tumors reflects real dynamics. Our first statistics was the mean duration of the cycles $\bar{t}_r = 26.08$ [day] with their standard deviation $\sigma_r = 11.13$ [day] and the mean value of the difference between the peak tumor volume and the starting volume where the dose was applied, denoted by $\bar{y}_p = 1057.42$ [mm^3] with standard deviation $\sigma_p = 614.45$ [mm^3].

During the generation, we imposed three conditions that each θ_i-s must obey. The first condition is that for each θ_i the untreated tumor should grow, i.e., $a - n > 0$. Our second condition filters those parameter sets for which the tumor is resistant and grows continuously for an initial 8 [mg/kg] dose, i.e., it filters out the cases for which $a - n - b \geq 0$. The last condition restricts the maximum deviation of the tumor volume between the volume at the injection

Table 1. Log-transformed parameter values of the identification.

Parameter	**Fixed effect**	**Standard errors**	**Random effect covariance**
a	-0.84	0.21	0.087
b	-0.2	0.17	0.251
n	-2.03	0.65	0.133
w	-2.44	0.17	0.417
ED_{50}	-6.71	32.7	0.008
c	0.46	3.19	$5.9 \cdot 10^{-8}$
$k1$	1.46	58.8	$2.38 \cdot 10^{-8}$
$k2$	1.42	49.96	0.228
x_0	4.08	0.35	1.78

($u_0 = 8$ [mg/kg]) and the peak tumor volume to $y_p \pm \sigma_p$. This entails that the generated parameter sets lead to non-resistant tumors and they also share similar traits to their real-life counterparts. Using these restrictions we generate a virtual population $\mathcal{V}_r$ with 100 elements that will be used for robustness analysis. We have also generated a different population $\mathcal{V}_o$ with a condition, that each tumor should shrink under 10 [mm^3] between $t_r \pm \sigma_r$, from a generated initial condition x_0 with initial dose $u_0 = 8$ [mg/kg]. The role of $\mathcal{V}_o$ is to provide species for which a single dose can reduce the tumor completely so that we can easily compare the effectiveness of therapies attributed to different control parameters with the MTD injection.

4.2 NMPC tuning

For the tuning of the controller, we assumed that the administrations take place on Mondays and Thursdays. Because PLD treatment is given through the tail vein of the animal, a minimum of 3 days recovery is required between drug injections. Technically, the frequency of tail vein administration should be limited to a minimum to avoid unnecessary stress (Hedrich and Bullock (2004)) and, additionally, the PLD treatment could cause inflammation and necrosis if administered more frequently. In practice, this means that the intervals of the optimization are defined recursively as

$$\begin{aligned} (t_k, t_k + d],\ d = \begin{cases} 3, & \text{if } k \text{ is even} \\ 4, & \text{otherwise} \end{cases} \\ t_{k+1} = t_k + d, \\ t_0 = t_i, k \in \{0, \dots, N\} \subset \mathbb{N}_0 \end{aligned} \tag{8}$$

with an optimization variable u_i at each t_k. The considered time span of the simulations was 40 days for each sample, starting with $t_0 = 0$. During the tuning, we have presumed that the model parameters are known precisely, and we have access to full state measurement so that we can omit the use of MHE, which alleviates the computational burden of the tuning procedure. Our initial approach was to utilize a simple grid search on the domain $q \times r$ where $q, r \in 10^i, i \in \{-7, \dots, 3\}$ with $N = 3$ which is two weeks of prediction. The setpoint was set to $y_{\text{ref}} = 1$, because zero volume can not be attained with the model. During our initial trials, we used **fmincon** to solve (5), however the discontinuous nature of the impulse treatment lead to poor convergence properties. By using the derivative-free, global optimizer **patternsearch**, the problem could be tackled efficiently.

The initial grid search revealed two important information. Using smaller values for q, r leads to better convergence

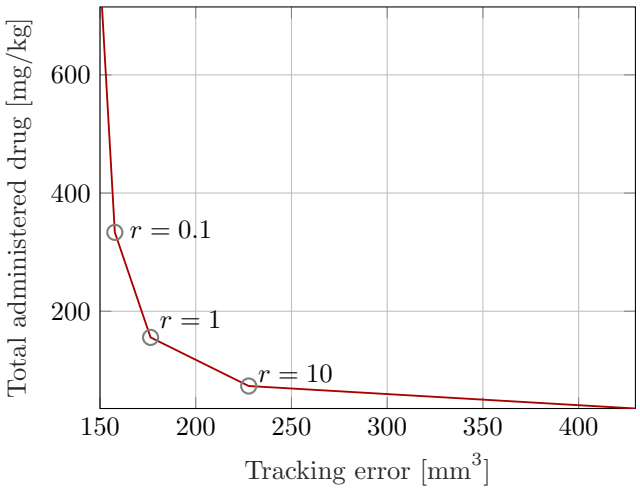

Fig. 3. Pareto front of the grid search for $q = 1$ and varying r values

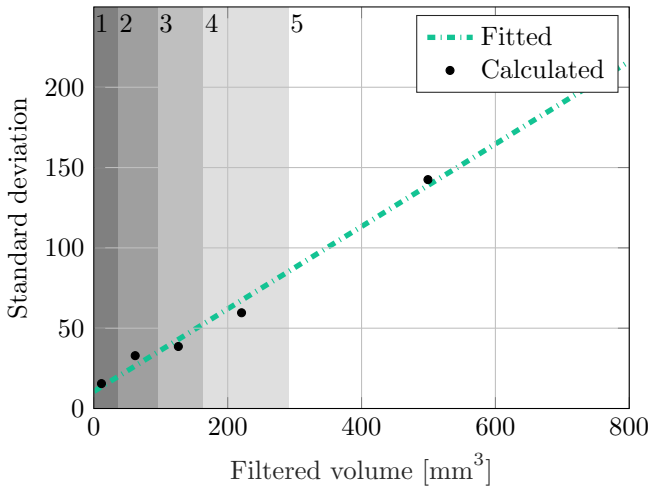

Fig. 4. Fitted noise model and the comparison of the distributions. Numbers (1-5) corresponds to the shaded regions, in which the standard deviation is calculated.

properties in each of the virtual patients. It could also be seen from the simulations, that the distance between q, r on logarithmic scale determines the quality of the solutions. We calculated for each q, r value the total amount of drug administered for the 100 patients in $\mathcal{V}_o$ and the sum of the tracking error at the end of the simulation interval (denoted by u_{sum} and e_{sum} respectively). Surprisingly, for each q, r which has the same relative distance lead to essentially the same u_{sum} and e_{sum} values. As such, we show the Pareto front of these two metrics in Figure 3. for $q = 1$. One can see that by decreasing the value of r, the total administered drug increases while the tracking error shrinks. It can be seen that the optimal choice, which results in the best trade off between the two metrics, is $r = 1$, for which $u_{\text{sum}} = 155$, $e_{\text{sum}} = 176$ which is significantly better than the single 8 [mg/kg] dose case with $u_{\text{sum}} = 800, e_{\text{sum}} = 414$.

We used the robust dataset $\mathcal{V}_r$ to validate the obtained controller parameters with a fixed $q = 1$ and $r \in 10^i, i \in \{-1, \dots, 2\}$. The uncontrolled dataset with a single 8 [mg/kg] dose at the beginning can be seen in Figure 5. Here, we simulated each patient to 350 days to ensure that they can be controlled to the setpoint. While the controllers were able to shrink each tumor, the best tracking error was obtained by using $r = 0.1$, which is in accordance with our previous result. As such, we chose $q = 1$ with $r = 0.1$ to achieve a more aggressive response from the controller to eradicate the tumor as fast as possible from the subject while retaining the economic drug administration. One last issue is that the tracking error at the end of the treatment is still quite large with $e_{\text{sum}} = 67$ because the controller is not tracking the setpoint correctly in some cases. By increasing the prediction horizon to $N = 8$, the stability of the scheme was improved to $e_{\text{sum}} \approx 0$.

4.3 Simulated measurement noise

Measurement noise has been characterized based on the complete time series. During the experiments, the cross-section of the tumor was measured using calipers which combined with (7) approximates the tumor volume on a given day. Since the approximation is a rough estimate, because the tumor, in reality, is not a perfect ellipsoid, one must include measurement noise in the model. To quantify the noise present in the time-series data, a lowpass Butterworth filter was applied in a zero-phase setting. The noise is approximated as the difference between the raw, and the filtered time series. We created a histogram from these error terms with five bins (containing an approximately equal number of measurements), where we calculated their standard deviations. We found that the standard deviation of the noise can be accurately approximated by an affine function of the volume, which can be seen in Figure 4. The generated noise is assumed to be an additive Gaussian process in the form of

$$\begin{aligned} \sigma(y) &= 0.1 + 12.4y, \\ \nu(y) &\sim \mathcal{N}(0,\, \sigma^2(y)), \\ \tilde{y} &= y + \nu(y), \end{aligned} \tag{9}$$

where $\tilde{y}$ is the measured tumor volume and y is the output of the model given in (1).

4.4 Verification of the algorithm

The combined MHE-NMPC algorithm is verified on a virtual population generated with the same conditions as $\mathcal{V}_r$ but containing different subjects. We assumed that during a week five measurements are taken, from Monday to Friday, and the administrations are applied as described in (8). The measurements are also corrupted with the noise as described earlier and each parameter of the algorithm remains the same as above mentioned. For the MHE, problem (3) is solved using **fmincon** from two initial points to improve convergence. One of them is the nominal parameter set while the other is the resulting parameters of the previous optimization. Problem (5), corresponding to the NMPC, is solved with the global optimizer as mentioned earlier.

The results of the simulation can be seen in Figure 6. Each tumor shrinks down to the neighborhood of the setpoint with varying speed. The total drug used was $u_{\max} = 1992$ which means that on average a single subject received 20 [mg/kg] PLD during the simulation interval. Because measurement errors are introduced, it is not obvious to determine whether a true setpoint tracking was achieved in

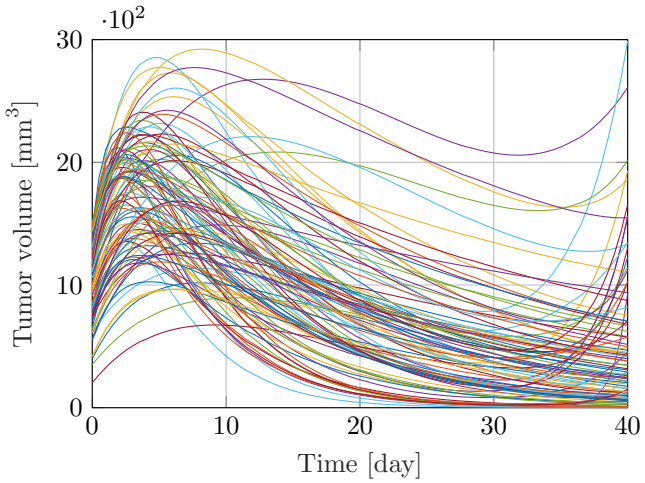

Fig. 5. Robust virtual population $\mathcal{V}_r$ subject to a single 8 [mg/kg] administration on the first day.

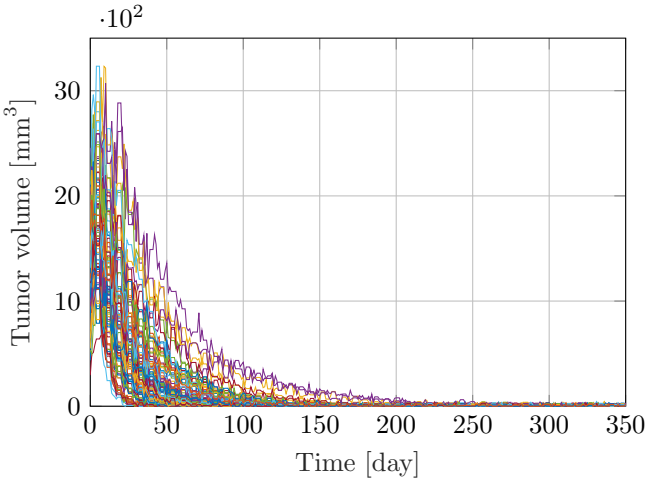

Fig. 6. The measured outputs of each virtual patients subject to the algorithmic treatment

this case. At first glance, the MHE provides close estimates to the actual measurements and leads to less variation in each upcoming measurement. However, the fact that the parameter estimates from the MHE vary significantly in many subjects could easily mean that the computed therapy, with these parameters, just aggressive enough to shrink the tumor size completely, but unable to provide a realistic setpoint tracking. This could be attributed to the fact that the model introduces several discontinuities because of its impulsive nature which can not be properly handled by the gradient-based solver. Moreover, the use of the global optimizer, as in the case of the NMPC, did not improve the convergence properties of the MHE.

5. CONCLUSION

A combined MHE-NMPC algorithm was presented which was tuned based on a generated virtual population. These populations, in conjunction with the measurement noise model, could provide a useful basis for further numerical studies on the improvements of the current algorithm. It can be noted that the algorithm in simulation can shrink tumors with a high degree of variability in their parameters. Further research should focus on improving the stability of the MHE estimates and understanding the setpoint tracking capabilities of the algorithm in more detail.

REFERENCES

Belfo, J.P. and Lemos, J.M. (2021). *Optimal Impulsive Control for Cancer Therapy*. Springer International Publishing. doi:10.1007/978-3-030-50488-5.

Carrère, C. (2017). Optimization of an in vitro chemotherapy to avoid resistant tumours. *Journal of Theoretical Biology*, 413, 24–33. doi:10.1016/j.jtbi.2016.11.009.

Chen, T., Kirkby, N.F., and Jena, R. (2012). Optimal dosing of cancer chemotherapy using model predictive control and moving horizon state/parameter estimation. *Computer Methods and Programs in Biomedicine*, 108(3), 973–983. doi:10.1016/j.cmpb.2012.05.011.

Czakó, B.G., Drexler, D.A., and Kovács, L. (2020). Impulsive control of tumor growth via nonlinear model predictive control using direct multiple shooting. In *2020 European Control Conference (ECC)*. IEEE. doi:10.23919/ecc51009.2020.9143755.

Drexler, D.A., Ferenci, T., Füredi, A., Szakács, G., and Kovács, L. (2020). Experimental data-driven tumor modeling for chemotherapy. *21st IFAC World Congress 2020*.

Drexler, D.A., Nagy, I., Romanovski, V., Tóth, J., and Kovács, L. (2018). Qualitative analysis of a closed-loop model of tumor growth control. In *2018 IEEE 18th International Symposium on Computational Intelligence and Informatics (CINTI)*. IEEE. doi:10.1109/cinti.2018.8928208.

Füredi, A., Szebényi, K., Tóth, S., Cserepes, M., Hámori, L., Nagy, V., Karai, E., Vajdovich, P., Imre, T., Szabó, P., Szüts, D., Tóvári, J., and Szakács, G. (2017). Pegylated liposomal formulation of doxorubicin overcomes drug resistance in a genetically engineered mouse model of breast cancer. *Journal of Controlled Release*, 261, 287–296. doi:10.1016/j.jconrel.2017.07.010.

Hedrich, H.J. and Bullock, G. (2004). *The Laboratory Mouse*. Elsevier. doi:10.1016/b978-0-12-336425-8.x5051-1.

Sápi, J., Kovács, L., Drexler, D.A., Kocsis, P., Gajári, D., and Sápi, Z. (2015). Tumor volume estimation and quasi-continuous administration for most effective bevacizumab therapy. *PLOS ONE*, 10(11), e0142190. doi:10.1371/journal.pone.0142190.

Sbeity, H. and Younes, R. (2015). Review of optimization methods for cancer chemotherapy treatment planning. *Journal of Computer Science & Systems Biology*, 8(2). doi:10.4172/jcsb.1000173.

Shi, J., Alagoz, O., Erenay, F.S., and Su, Q. (2011). A survey of optimization models on cancer chemotherapy treatment planning. *Annals of Operations Research*, 221(1), 331–356. doi:10.1007/s10479-011-0869-4.

Siket, M., Eigner, G., and Kovács, L. (2020). Sensitivity and identifiability analysis of a third-order tumor growth model. In *2020 IEEE 15th International Conference of System of Systems Engineering (SoSE)*, 417–422. doi:10.1109/SoSE50414.2020.9130530.