Preprints of the 21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

Toxicity-Centric Cancer Chemotherapy Treatment Design

Joseph T. Liparulo [*] **Timothy D. Knab PhD** [*,**]
Robert S. Parker PhD [*,***,****,†]

[*] *Department of Chemical and Petroleum Engineering, Swanson School of Engineering, University of Pittsburgh, Pittsburgh, PA 15261. (email: jtl19@pitt.edu, knab.tim@gmail.com, rparker@pitt.edu)*
[**] *Metrum Research Group, Tariffville, CT 06081*
[***] *Department of Critical Care Medicine, University of Pittsburgh, Pittsburgh, PA 15261*
[****] *Department of Bioengineering, Swanson School of Engineering, University of Pittsburgh, Pittsburgh, PA 15261*
[†] *McGowan Institute for Regenerative Medicine, University of Pittsburgh Medical Center, Pittsburgh, PA 15213*

Abstract: Cancer chemotherapy scheduling in the mathematics and engineering literature has generally focused on optimal control formulations and tumor kill, using constraints on dose magnitude and duration to implicitly mitigate toxicity. We introduce a framework for scheduling that focuses on clinically-relevant toxicity mitigation allowing clinicians to specify toxicity limits in terms they understand. Building from the model predictive control framework, we explicitly use the pharmacokinetic model of drug distribution as well as pharmacodynamic models of both antitumor effect and drug toxicity in the optimization problem. Clinical and logistical constraints round out the treatment design problem. Rather than direct inversion, we synthesize the optimization problem in an input-discretized form and solve via graphical processing unit (GPU) calculation. The resulting suboptimal solution is shown to be clinically indistinguishable from an optimal solution (calculated via nonlinear least squares (NLS) from a relaxation of the input and logistical constraints to continuous variables). Using a docetaxel administration case study, the algorithm controlled neutropenia within user-specified toxicity constraints while maintaining tumor eradication rates equivalent to, or better than, clinically-implemented dosing schedules. Changes in patient response – both antitumor efficacy and toxic drug sensitivity are captured via a nonlinear least squares (NLS) calculation at the end of each treatment cycle and updated in the next cycle design. By explicitly controlling treatment toxicity, this algorithm has the potential to improve patient quality-of-life.

Keywords: Systems Medicine, Nonlinear Dynamical Modeling, Model Predictive Control, Optimal Treatment Design, Cancer Chemotherapy, Docetaxel, Pharmacokinetic-Pharmacodynamic Modeling, Toxicity

1. INTRODUCTION

More than 1.7 million new cases of cancer, a disease characterized by the uncontrolled growth and spread of malignant cells, are expected to be diagnosed in 2019; roughly 607,000 Americans are expected to die in 2019 as a direct result of the disease (Street, 2019). Untreated cancer leads to organ failure and death of the host organism as a result of either the primary tumor or metastasis. Metastasis occurs when cancer cells translocate to a tissue distant from the original tumor site; the exact metastasis process remains undetermined (Gupta and Massagué, 2006). Treatment modalities are selected depending on the type and location of the cancerous tumor. When possible, the tumor is excised surgically. However, due to the location of certain tumors (*e.g.*, in the brain) or the potential resulting loss-of-function from organ removal(*e.g.*, in the pancreas), this is not always an option; non-localized cancers (*e.g.*, hematologic cancers like leukemia) also cannot be surgically removed. Once the primary tumor is detected, undetectable distant metastases likely exist, thus necessitating full-body treatment (*e.g.*, chemotherapy).

Docetaxel is a commonly employed anticancer drug in both mono-agent and combination regimens (Pollard et al., 2017; US Food and Drug Administration, 2012). Chemotherapy, a commonly-used systemic cancer treatment, takes advantage of the quick rapid proliferation of cancer cells by preferentially attacking rapidly dividing cells. Selectivity of chemotherapeutics toward diseased cells is desired but not always achieved, leading to patient toxicities. White blood cells are one of the most commonly affected healthy cells during chemotherapy. The balance between minimizing chemotoxic side-effects and maximizing elimination of cancerous cells creates a dichotomy for clinicians, who have to select a chemotherapy schedule

(defined as dose magnitude and dosing frequency) that balances antitumor efficacy against tolerable toxic side-effects. These schedules are guided by empirical evidence from preclinical trials during the drug discovery phase and refined through clinical testing. Formal model-based optimization methods provide for a rigorous analysis of this treatment trade-off, to the degree that mathematical models of toxicity and treatment are available.

A variety of model-based methods have been used to design optimal treatment algorithms targeting tumor growth and the subsequent elimination of cancerous cells, with model parameters often fitted to patient data Norton (1988); Florian, Jr. (2008). The design tools have focused on optimal control formulations using a fixed final time-point for treatment; most aim to minimize tumor volume at a final time, with constraints on input magnitude (Martin and Teo, 1994; Swan, 1986; Ledzewicz and Schättler, 2007). Patient treatment, however, does not have a final time defined *a priori*. Instead, real-time or cycle-wise treatment decisions are driven by toxicity and efficacy, which may change over time in a single patient and/or across patients, such that the endpoint of disease remission is not easily predicted (Cella et al., 2003). Thus, treatment is given in cycles to allow clinicians to evaluate patient response and use feedback to adjust treatment accordingly. Herein, we synthesize a model-based framework that allows clinicians to easily assess trade-offs between treatment efficacy and toxic side-effects.

Our first treatment design algorithm (Harrold and Parker, 2009) was based on a preclinical animal model and solved the treatment design problem using mixed-integer linear programming. Here we employ physiologically-based nonlinear pharmacokinetic (PK), pharmacodynamic (PD), and toxicity models in a receding-horizon-based treatment design framework akin to model predictive control (MPC). Toxicity limits and clinical logistic concerns are incorporated explicitly as constraints. Tumor volume is minimized over the treatment horizon of one or more therapy cycles, subject to the constraints. To address changes in patient response over time, tumor volume and toxicity measurements at the end of each treatment cycle are used to update the patient-specific model. The resulting algorithm provides a rigorous model-based approach to balance toxicity and efficacy, which could yield superior patient quality-of-life.

2. METHODS

2.1 Treatment Design Algorithm

Overall, the chemotherapy scheduling problem can be posed as follows, drawing upon the MPC framework:

$$\min_{D_d(q)} \sum_{d=1}^{N_d}(N(d)) + \Gamma_u \sum_{q=1}^{m_q} D_d(q)^2 + \Gamma_{cyc} N(N_d) \quad (1)$$

$$s.t. \text{ Drug Pharmacokinetics} \quad (2)$$

$$\text{PD: Tumor Kill} \quad (3)$$

$$\text{PD: Toxicity} \quad (4)$$

$$\sum_{k=1}^{5} b_d(k) \leq 1 \forall d \in \{1,2,3\} \quad (5)$$

$$D_d^{min} \leq D_d(q) \leq D_d^{max} \quad (6)$$

$$\sum_{q=1}^{m_q} \leq D_d^{total} \quad (7)$$

Here tumor volume ($N(d)$) is minimized weekly over N_d weeks, with the problem typically solved over 1 or more multi-week cycles. Drug doses ($D_d(q)$, with maximum number m_q) can be penalized (via weight Γ_u), leading to small doses not being administered (they contribute toxicity, but no cancer kill), while large doses kill sufficiently to overcome the penalty. A terminal penalty, Γ_{cyc} allows preferential weighting of the end-of-cycle tumor volume, if desired. Equation (5) is a logistics constraint limiting chemotherapy administration to only once per calendar work week (Monday - Friday), and no more than 3 weeks in a row in a given administration cycle. This satisfies patient quality-of-life constraints (not visiting the doctor more often than weekly for a drug infusion) and hospital/health insurer cost and staffing concerns (dosing only during the typical work week and workday). The final two constraints, (6) and (7), bound the magnitude of each individual dose and limit the amount of drug administered per cycle, respectively.

2.2 Case Study and Models

This work focuses on docetaxel administration to treat solid tumors, such as non-small cell lung, head-and-neck, and androgen-independent (castrate-resistant) prostate cancer (Clarke and Rivory, 1999). Common clinical schedules are 100 $\frac{\text{mg}}{\text{m}^2}$ every 21 days or 35 $\frac{\text{mg}}{\text{m}^2}$ weekly, 3 weeks of 4. The primary associated toxicity is neutropenia (low absolute neutrophil count (ANC)), where toxicity grades are shown in Table 1.

Table 1. Toxicity grades and neutrophil counts

Grade	Cell Count
0 (Normal)	ANC $\geq 2.0 \times 10^6$/ mL
1	1.5×10^6/ mL $\leq$ ANC $< 2.0 \times 10^6$/ mL
2	1.0×10^6/ mL $\leq$ ANC $< 1.5 \times 10^6$/ mL
3	0.5×10^6/ mL $\leq$ ANC $< 1.0 \times 10^6$/ mL
4	ANC $< 0.5 \times 10^9$ / L

Physiology-based Pharmacokinetic Model Physiologically-based pharmacokinetic (PBPK) modeling data was obtained from a PK case study that administered docetaxel to severe combined immunodeficient (SCID) mice bearing SOV-3 human ovarian cancer xenografts (Strychor et al., 2005; Zamboni et al., 2008; Florian, Jr., 2008). By taking mass balances around all tissues, a compartmental description of docetaxel PBPK can be constructed as shown in Figure 1 Florian, Jr. (2008). The model was scaled to humans by altering tissue flowrates and volumes Snyder et al. (1975), and allowing the "Other" compartment intercompartment transport rates to increase (docetaxel is lipophilic, and humans carry a higher percentage of body fat than mice). All other rate parameters were held constant. Docetaxel is delivered via Intravenous (IV) administration. Tissue compartments (except the gut and bone marrow) were composed of vascular and extravascular compartments, where extravascular compartments included a protein-bound docetaxel concentration state (necessitated by the significant lipophilicity and protein

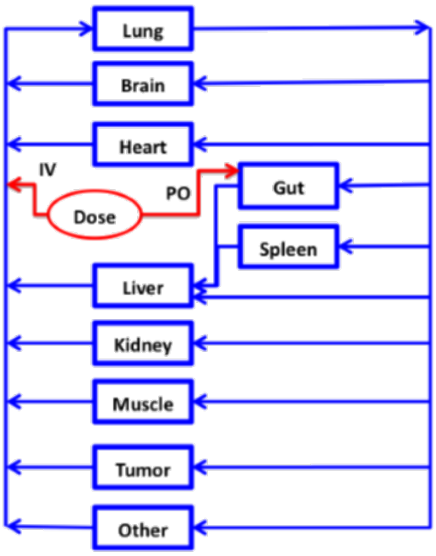

Fig. 1. Physiologically-based pharmacokinetic model for docetaxel, as developed in Florian, Jr. (2008).

binding characteristics of docetaxel). Protein binding was also accounted for in the bloodstream. Drug metabolism in the liver was modeled via first-order kinetics based on extravascular concentration. The result was a 35-state linear dynamical model of docetaxel built in mice and scaled to humans. This PBPK model creates challenges in computation time, particularly during optimization, where the model is included in a set of dynamic constraints. We therefore reduce the linear ODE model using balanced truncation (via the *balmr* command in MATLAB©2019, The Mathworks, Natick, MA). The 4-state reduced model outputs are plasma concentration (for comparison to patient data), tumor concentration (to model drug-induced kill), and marrow concentration (to capture drug-induced toxicity to neutrophils).

2.3 Neutrophil toxicity Model

We employ a mechanism-based model of neutrophil trafficking originally developed in (Ho et al., 2013), and composed of 16 ODEs and 55 parameters. This model captures responses to both inflammation and chemotherapy challenges. For the present work, the inflammatory components of the model were removed, and model reduction was performed to reduce state dimension and to focus model sensitivity into a smaller number of more identifiable parameters (the model is not *a priori* identifiable), as outlined in (Ho, 2014). The 9-state model shown in Figure 2 is mathematically represented as follows:

$$\frac{dPr(t)}{dt} = \frac{B_{min}k_G + B_{max}G_{CSF}(t)}{k_G + G_{CSF}(t)} - k_{tox}\frac{E_{max}C_{BM}(t)}{EC_{50} + C_{BM}(t)}Pr(t) - k_{tr}Pr(t) \quad (8)$$

$$\frac{dT_1(t)}{dt} = k_{tr}Pr(t) - k_{tr}T_1(t) \quad (9)$$

$$\frac{dT_2(t)}{dt} = k_{tr}T_1(t) - k_{tr}T_2(t) \quad (10)$$

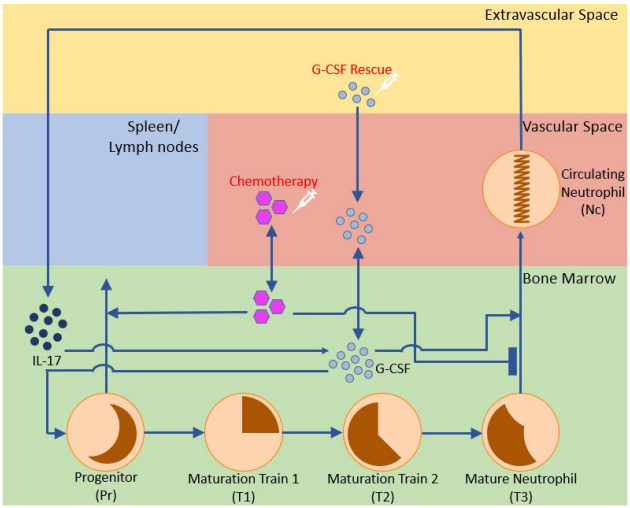

Fig. 2. Reduced neutrophil toxicity model (Ho, 2014).

$$\frac{dT_3(t)}{dt} = k_{tr}T_2(t) - k_{tr}T_3(t) \quad (11)$$

$$\frac{dN_c(t)}{dt} = k_{tr}T_3(t) - k_dN_c(t) \quad (12)$$

$$\frac{dIL_{17}(t)}{dt} = \frac{B_{G_{max}} * k_N}{k_N + N_c(t)} - k_{d_{IL_{17}}}IL_{17}(t) \quad (13)$$

$$\frac{dC_{BM}(t)}{dt} = -k_{bmv}C_{BM}(t) + k_{vbm}C_{plasma}(t) \quad (14)$$

$$\frac{dG_{CSF_T}(t)}{dt} = -k_{sc_{bl}}G_{CSF_T}(t) - k_{d_{GCSF_T}}G_{CSF_t}(t) \quad (15)$$

$$\frac{dG_{CSF}(t)}{dt} = k_{IL_{17}}IL_{17}(t) - k_{d_{GCSF}}G_{CSF}(t) + k_{sc_{bl}}G_{CSF}(t) \quad (16)$$

The stem-cell-like progenitor cell population ($Pr(t)$) generates cells that mature through the maturation train ($T_1(t)$, $T_2(t)$) to become mature neutrophils ($T_3(t)$) in the bone marrow. These cells migrate to the vascular space to become circulating neutrophils ($N_c(t)$), and are measurable as ANC. Decreases in ANC lead to production of $IL-17$, through a series of cell types and signaling molecules not explicitly represented in the reduced model. Elevated $IL-17$ drives production of G-CSF, a powerful stimulator of progenitor cells that also speeds migration of mature neutrophils to the bloodstream. Exogenously administered G-CSF, often used to rescue neutrophil toxicity, is released from a subcutaneous injection site to the bloodstream, which rapidly equilibrates with the bone marrow compartment. Systemic chemotherapy kills progenitor cells and slows the release of mature neutrophils from the marrow into the bloodstream, causing toxic side-effects.

2.4 Docetaxel PD Efficacy Model

Tumor growth models of varying complexity have been developed, and the seminal work of Norton (Norton, 1988) describing breast cancer tumors offered a robust tumor growth model that grows exponentially for small tumors volumes, but then slows its growth rate as the tumor expands. This is characteristic of solid tumors whose irregular blood vessel network causes resource limitation and challenges in delivering chemotherapeutics. A more

recent model from Norton *et al.* (Norton and Massagué, 2006) explores a different hypothesis (self-seeding) for growth that follows a similar temporal trajectory:

$$\frac{dN}{dt} = k_{gr}N(t)^{\frac{a}{c}} - k_{die}N(t)^{\frac{b}{c}} - k_{eff}N(t)D(t) \quad (17)$$

Here $N(t)$ is the total number of cancer cells, k_{gr} and k_{die} are the tumor cell growth and death rates, respectively, and a, b, and c are constants. By setting $3 \geq b \geq c \geq a \geq 2$, proliferation occurs in tumor regions with lower fractal dimension than the region where apoptosis occurs (Norton and Massagué, 2006). Tumor kill is size-dependent and implemented using a bilinear term where $D(t)$ is the drug concentration inside the tumor and k_{eff} is the drug effectiveness kill constant.

2.5 Generation of the Feasible Schedule Set

Clinically optimal chemotherapy schedules are obtained by first generating a set of feasible schedules which meet or conform to all logistic constraints. These constraints limit a number of aspects such as the number of allowable doses per specified time period, the allowable treatment administration times and the maximum total dose. The feasible schedules are generated by choosing the resolutions for the discretization of time and dose. These resolutions set the minimum difference in time between administered treatments and the minimum difference between a dose and the next largest (or smallest) dose. The chosen values may result from clinical concerns such as the usual length of a clinical visit (*e.g.*, one hour) or the common clinical increment in dose (*e.g.*, 5 mg steps). Coarser discretizations will result in significantly fewer feasible schedules, thereby leading to much faster simulation of the schedule set. Care should be taken to ensure that the selected discretizations are clinically relevant, as there is no need to discretize the schedule to clinically irrelevant precision (*e.g.*, nearest second for treatment or nearest μg for dose).

A recursive Python function was implemented to generate feasible schedules based on timing constraints, and a second recursive function was used to find all possible dose regimens within a defined maximum number of doses. Regardless of the chosen discretization for dose, a zero dose is permitted enabling schedules with fewer than the maximum allowable number of doses to be considered. The combinatorial set of feasible administration times and doses is then generated and saved as a matrix of feasible regimens in a binary format for use in the GPU simulation step.

2.6 GPU Simulation of the Feasible Schedule Set

Utilizing C++ along with VEXCL (`https://github.com/ddemidov/vexcl`), a function was written to compile a GPU kernel. This kernel takes in the matrix of all feasible schedules and solves, in parallel at each time step (δt), the PKPD ODEs comprising the PK model from Section 2.2.1 and the PD Equations (8) to (17). The ODE solver used in the GPU kernel is from the Boost odeint library (`https://www.boost.org/doc/libs/1_66_0/libs/numeric/odeint/doc/html/index.html`) and employs a Runge-Kutta DOPRI5 method. Larger time steps lead to significantly faster solutions to the ensemble of ODE models at the cost of accuracy, and care must be taken to choose a δt that ensures all feasible chemotherapy administrations fall exactly on a time step. During runtime, after each time step, the GPU kernel updates the objective function, which for the Docetaxel case study presented herein, is given by Equation (1).

2.7 Dose Schedule Design

Starting from the optimization problem in equations (1) to (7), clinicians can apply additional constraints to limit toxicity. Relevant to the Docetaxel case study, the following constraints are added to the problem for all days i:

$$N_c(i) \geq 1.0 \times 10^6/mL \quad (18)$$

$$\max\left([N_c(i), N_c(i+1) \ldots N_c(i+6)]\right) \geq 1.5 \times 10^6/mL \quad (19)$$

The first constraint guarantees no grade 3 toxicity. The second forces any grade 2 toxicity to last no more than 7 consecutive days. Constraints are tested for violation at the end of each GPU simulation. When a particular schedule induces a toxicity constraint violation, all further simulations that use the selected schedule and higher doses are not simulated. The schedule is then incremented to the next entry in the feasible set, and simulation begins again. This process is repeated until the full set of acceptably toxic schedules is evaluated. The optimal schedule is the one with the lowest objective function value from Equation (1).

2.8 Updating the Model During Treatment

Patients return to the clinic for each drug dose, as the drugs are administered intravenously (IV). At the end of each cycle, absolute neutrophil count is measured from a patient blood sample. Antitumor efficacy is typically assessed every other cycle via magnetic resonance imaging (or other clinical imaging method). Since patient response may vary from the model predictions, two model parameters are updated as measurements are collected. Drug sensitivity (k_{tox} in Equation (8)) is updated each cycle, and antitumor effect (k_{eff} in Equation (17)) is updated every other cycle, corresponding with their most informative clinical measurements at their respective rates of collection. Parameter updates are computed via nonlinear least squares according to the following objective function:

$$\min_{k_{tox},k_{eff}} \quad (ANC_{meas} - N_c(N_d))^2 + w_e * (V_{meas} - N(N_d))^2 \quad (20)$$

Then end-of-cycle (N_d) measurements of ANC and tumor volume are compared to model simulation values. k_{tox} affects $N_c(t)$, while k_{eff} drives $N(t)$. The binary variable w_e is 1 for cycles when tumor volume is measured, and 0 otherwise.

3. RESULTS

3.1 Model Validation

A sample patient fit of the PBPK model, using data from UPCI study 01-150, is shown in Figure 3. The

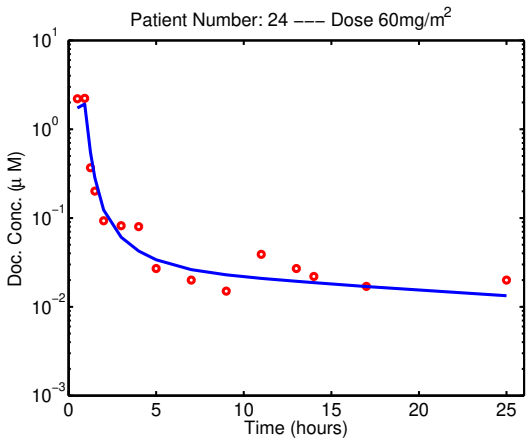

Fig. 3. Sample model fit to human patient data from UPCI study 01-150. A single docetaxel dose of 60 $\frac{\text{mg}}{\text{m}^2}$ was administered, and plasma concentrations were measured at intervals after IV dosing.

model captures well the single-dose concentration profile of docetaxel in plasma, demonstrating that the animal-derived model can fit patient data adjusting only flows, volumes, and the intercompartment transfer rates for the "Other" compartment.

3.2 Optimal Treatment vs. Clinical Practice

Typical treatment schedules for docetaxel are 100 mg/m^2 every 21 days (q21d) or 35 mg/m^2 every week for three weeks followed by one week off (qw 3 of 4)(US Food and Drug Administration, 2012). The decision support system (DSS) algorithm predicted an optimal dosing schedule of 55 mg/m^2, 45 mg/m^2, and 5 mg/m^2 on days 0, 11, and 18, respectively, of a 28 day dosing cycle. Although this schedule may be mathematically optimal, a more clinically relevant (sub)optimal schedule is 55$\frac{\text{mg}}{\text{m}^2}$ and 50$\frac{\text{mg}}{\text{m}^2}$ on days 0 and 11 (DSS), respectively. The three schedules are compared in Figure 4. The q21d schedule demonstrates superior tumor kill compared to the qw 3 of 4 and DSS over the course of 84 days. It should be noted, however that the q21d schedule has administered 4 cycles of treatment, or 25% more drug, compared to three cycles for DSS and qw 3 of 4. The q21d schedule also induces significant grade 3 neutropenia, which is clinically undesirable. The toxicity and tumor-killing effect is similar for the qw 3 of 4 and DSS treatment schedules, but the DSS schedule only requires two clinical visits per cycle, compared to three visits for qw 3 of 4. This would improve patient quality-of-life and reduce treatment costs for the hospital and insurer with no increase in undesirable toxicity.

3.3 Managing Interpatient Variability

As patient sensitivities (efficacy, toxicity) change, the DSS patient model can be updated to provide patient-tailored optimal treatment. Figure 5 shows the ability of the DSS to alter treatment due to changing sensitivity of E_{max}. A new patient ($k_{tox} = 1.6$) arrives for their first cycle of therapy. The DSS nominal value for drug sensitivity is $k_{tox} = 1.0$. The first cycle of therapy is highly toxic – a grade 3 toxicity is observed, which would have been far worse if

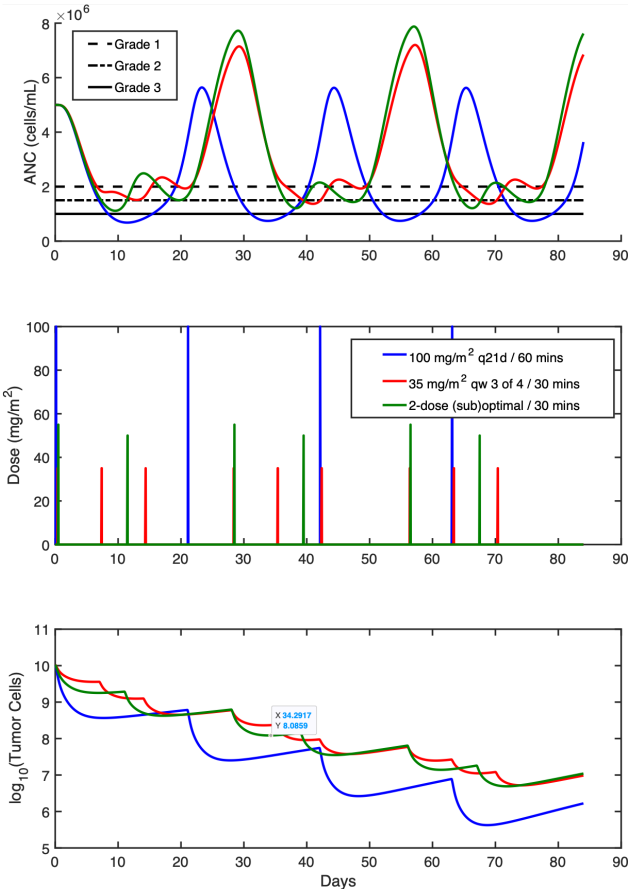

Fig. 4. Comparison of clinical standards of practice (100 $\frac{\text{mg}}{\text{m}^2}$ every 3 weeks, blue; 35 $\frac{\text{mg}}{\text{m}^2}$ 3 weeks of 4, red) and algorithm-calculated docetaxel administration (green). **Top:**. ANC over time, with toxicity grades shown as horizontal lines. **Middle:**. Dosing day and magnitude (in $\frac{\text{mg}}{\text{m}^2}$). **Bottom:** number of tumor cells over time.

this patient had received the standard treatment dose of 100 $\frac{\text{mg}}{\text{m}^2}$ q21d. Clinically, this patient would probably have had dose 2 withheld until ANC recovered to near baseline, followed by a dose reduction – the clinical method for patient tailoring of dose and schedule. The DSS can update k_{tox} each cycle according to Equation (20), which returns a new value of $k_{tox} = 1.6$. For the next two cycles, the patient is controlled within the allowable toxicity limits, using a revised treatment schedule (three smaller doses of 40, 30, and 20 mg/m^2 on days 4, 16, and 24, respectively) rather than the 2 larger doses. This does increase clinical costs (clinician time, patient visits), which could be made a direct trade off by considering the cost of toxicity rescue (∼$10,000/cycle for G-CSF, excluding patient visit needs). Prior to cycle 4, the patient undergoes a toxicity increase ($k_{tox} = 1.9$), while the DSS is unaware (DSS $k_{tox} = 1.6$). After another significant toxicity event, the algorithm again updates to match the patient. As before, doses are further reduced, yielding the optimal treatment schedule of 30, 25, and 25 mg/m^2 on days 7, 17, and 25, respectively. It should be noted that if k_{tox} becomes too high, or k_{eff} goes to zero (meaning the drug is no longer having an antitumor effect), the DSS will return 0 for all dose magnitudes.

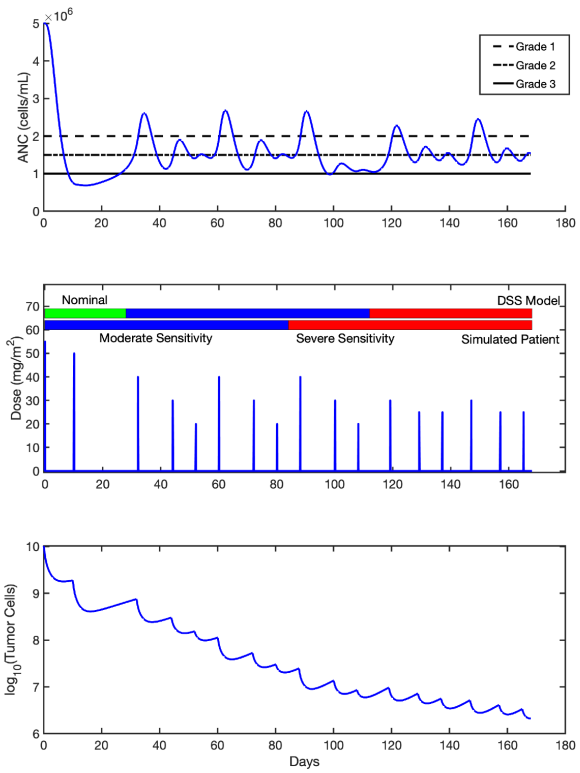

Fig. 5. Simulated patient response to treatment in the presence of model mismatch. **Top:** patient ANC over time. **Middle:** dosing day and quantity (plot) overlayed with toxicity sensitivity for the algorithm-expected patient (top bar) and actual patient (bottom bar). Sensitivities (k_{tox}) are nominal (1.0), moderate (1.6) and severe (1.9). Doses are calculated each cycle (28 days), with algorithm parameter updates after each cycle. **Bottom:**. number of tumor cells over time

4. SUMMARY

A model-based algorithm was synthesized to construct cyclic chemotherapy administration schedules. Pharmacokinetic, pharmacodynamic efficacy, pharmacodynamic toxicity, and tumor growth models are required. The receding-horizon control is suboptimal, but allows more flexible constraint handling and explicit use of nonlinear models that challenge optimal control formulations. Clinically-relevant constraints on toxicity are specified in clinical language and mapped onto the toxicity model, and further incorporation of treatment-relevant logistic constraints impacting patients and insurers / health care providers are easily embedded. The GPU calculations are fast, facilitating real-time use, and the algorithm can be updated to respond to changes in patient response (efficacy and toxicity) to treatment. This treatment design algorithm may provide a framework for improving patient quality-of-life, and it can be extended to other agents for which PK and PD (toxicity) models are available.

REFERENCES

Cella, D., Peterman, A., Hudgens, S., Webster, K., and Socinski, M.A. (2003). Measuring the side effects of taxane therapy in oncology: The functional assessment of cancer therapy–taxane (fact-taxane). *Cancer: Interdisciplinary International Journal of the American Cancer Society*, 98(4), 822–831.

Clarke, S.J. and Rivory, L.P. (1999). Clinical pharmacokinetics of docetaxel. *Clinical pharmacokinetics*, 36(2), 99–114.

Florian, Jr., J.A. (2008). *Modeling and Dose Schedule Design for Cycle-Specific Chemotherapeutics*. Ph.d. dissertation, University of Pittsburgh, Pittsburgh, PA.

Gupta, G.P. and Massagué, J. (2006). Cancer metastasis: building a framework. *Cell*, 127(4), 679–695.

Harrold, J.M. and Parker, R.S. (2009). Clinically Relevant Cancer Chemotherapy Dose Scheduling via Mixed-Integer Optimization. *Computers and Chemical Engineering, FOCAPO 2008 Special Issue*, 33, 2042–2054. URL `http://dx.doi.org/10.1016/j.compchemeng.2009.06.005`.

Ho, T. (2014). *A Model-Based Clinically-Relevant Chemotherapy Scheduling Algorithm for Anticancer Agents*. Ph.d. dissertation, University of Pittsburgh, Pittsburgh, PA.

Ho, T., Clermont, G., and S, P.R. (2013). A Model of Neutrophil Dynamics in Response to Inflammatory and Cancer Chemotherapy Challenges. *Comp. Chem. Eng.*, 51, 187–198.

Ledzewicz, U. and Schättler, H. (2007). Optimal Controls for a Model with Pharmacokinetics Maximizing Bone Marrow in Cancer Chemotherapy. *Math. Biosci.*, 206, 320–342.

Martin, R. and Teo, K. (1994). *Optimal control of drug administration in cancer chemotherapy*. World Scientific.

Norton, L. (1988). A Gompertzian model of human breast cancer growth. *Cancer research*, 48(24 Part 1), 7067–7071.

Norton, L. and Massagué, J. (2006). Is cancer a disease of self-seeding? *Nature medicine*, 12(8), 875.

Pollard, M.E., Moskowitz, A.J., Diefenbach, M.A., and Hall, S.J. (2017). Cost-effectiveness analysis of treatments for metastatic castration resistant prostate cancer. *Asian journal of urology*, 4(1), 37–43.

Snyder, W., Cook, M., Nasset, E., Karhausen, L., Howells, G., and Tipton, I. (1975). *Report of the Task Group on Reference Man: A Report Prepared by a Task Group of Committee 2 of the International Commission on Radiological Protection*. Pergamon Press, Oxford.

Street, W. (2019). Cancer facts & figures 2019. *American Cancer Society: Atlanta, GA, USA*.

Strychor, S., Eiseman, J.L., Parise, R.A., Egorin, M.J., Joseph, E., and Zamboni, W.C. (2005). Plasma, tumor, and tissue disposition of docetaxel in scid mice bearing skov-3 human ovarian xenografts.

Swan, G. (1986). Cancer chemotherapy: Optimal control using the verhulst-pearl equation. *Bulletin of mathematical biology*, 48(3-4), 381–404.

US Food and Drug Administration (2012). Highlights of docetaxel prescribing information.

Zamboni, W.C., Strychor, S., Joseph, E., Parise, R.A., Egorin, M.J., and Eiseman, J.L. (2008). Tumor, tissue, and plasma pharmacokinetic studies and antitumor response studies of docetaxel in combination with 9-nitrocamptothecin in mice bearing skov-3 human ovarian xenografts. *Cancer chemotherapy and pharmacology*, 62(3), 417–426.