



On the Role of the Objective in the Optimization of Compartmental Models for Biomedical Therapies

Urszula Ledzewicz^{1,2}  · Heinz Schättler³

Received: 25 July 2020 / Accepted: 9 September 2020 / Published online: 30 September 2020
© Springer Science+Business Media, LLC, part of Springer Nature 2020

Abstract

We review and discuss results obtained through an application of tools of nonlinear optimal control to biomedical problems. We discuss various aspects of the modeling of the dynamics (such as growth and interaction terms), modeling of treatment (including pharmacometrics of the drugs), and give special attention to the choice of the objective functional to be minimized. Indeed, many properties of optimal solutions are predestined by this choice which often is only made casually using some simple ad hoc heuristics. We discuss means to improve this choice by taking into account the underlying biology of the problem.

Keywords Optimal control · L_1 - and L_2 -type objectives · Mathematical modeling · Cancer therapies · Pharmacometrics

Mathematics Subject Classification 49K15 · 92C50

1 Introduction

Optimal control is a methodology to obtain well thought through solutions to practical problems for dynamical processes. It has found numerous applications in the sciences and has become a staple of engineering design, but applications to biomedical problems are less well known. Yet, optimal control as a tool in the analysis of mathematical models of cancer chemotherapy dates back to the 1970s and 1980s (e.g., see [1–6]),

✉ Urszula Ledzewicz
uledzew@siue.edu

Heinz Schättler
hms@wustl.edu

¹ Lodz University of Technology, 90-924 Lodz, Poland

² Southern Illinois University Edwardsville, Edwardsville, IL 62026-1653, USA

³ Washington University, St. Louis, MO 63130, USA

and these efforts have continued actively throughout the years to the present day (e.g., see [7–14] and the many references in [15]).

In the modeling, considerable attention is given to the formulation of the underlying dynamics. Based on the tremendous progress that has been made over the years in understanding biomedical problems, and cancer in particular, even minimally parameterized models are generally well thought through and realistic. As a matter of fact, however, considerably less effort is made when the objective is formulated that will be imposed on the process. Omnipresent are the papers in which the formulation of the criterion at best is a mere after-thought rested in questionable heuristics based on the notion of a “systemic cost”—whatever that may be. Naturally, the choice of the objective in an optimal control problem is a very important aspect of the overall approach and quite often it is here where what will be eventual properties of the optimal solution are already determined. One should not confuse this with obtaining meaningful information on the underlying biological problem. While, on the one hand, it is not that straightforward to come up with a proper objective criterion in biological problems, on the other hand, choosing the objective based on mathematical tractability with disregard of its meaning will not pass muster.

In this paper, which is intended as a review and discussion, using mathematical models for cancer treatments as the main vehicle of presentation, we discuss the main aspects that enter into the modeling of biomedical problems as optimal control problems. We give some attention to the formulation of the dynamics (such as growth and interaction terms for the population) and treatment aspects, but our emphasis will be on the formulation of the objective to be minimized. We stress the need for a meaningful interpretation of the objective. This should be related to side effects of the treatment in models for cancer therapies and/or to the actual cost. Yet, such costs are not always easily quantifiable. Just think of ‘social distancing’ during an epidemic and its effects on the economy. Using examples from our work, we illustrate how information from the underlying biology or medically relevant limitations can be used to formulate the optimization criterion so that optimal solutions then give meaningful information for the underlying biomedical problem. We point out some pitfalls to be avoided, but also make useful simplifications in the modeling, both in the dynamics and the objective.

2 Typical Structure of the Dynamics for Biomedical Therapies

In this section, a general model is formulated which represents a class of dynamical systems common in the modeling of biomedical therapies. These models are affine in the controls $u \in \mathbb{R}^m$ with states $x \in \mathbb{R}_+^n$. The components x_i , $i = 1, \dots, n$, of the state x are nonnegative numbers which, in a typical example, represent volumes or cell counts of various populations (compartments) of cells relevant to the underlying biomedical condition or, in epidemiological models, population counts in various stages of the disease. The controls u_j , $j = 1, \dots, m$, stand for possible external interventions and strategies such as drug therapy in cancer treatments or vaccination efforts in epidemiological models. Mathematically, the general structure we consider in this paper is of the form

$$\dot{x} = F(x) + G(x)u, \quad x \in \mathbb{R}_+^n, \quad u \in U \subset \mathbb{R}^m, \quad (1)$$

with F and G typically nonlinear vector fields called the *drift* and *control vector fields*, respectively, and U the control set which represents restrictions on the size of the controls.

2.1 The Uncontrolled Dynamics: Model of the Underlying System

The drift vector field F describes the underlying dynamics of the system without any outside interventions. It generally consists of various terms which should include the following aspects: *growth terms* of the populations as well as *transition* and *interaction terms* between the various compartments.

2.1.1 Growth Terms

These terms describe the mechanisms behind the increase in all or some of the populations in the compartments that define the state. These mechanisms can be modeled in various ways taking into account whether or not such growth saturates or not. The simplest model for a single population $x \in \mathbb{R}$ is given by *exponential growth*, $\dot{x} = ax$, with a a constant representing the growth rate. The same expression with $a < 0$ describes a basic death term of the population and, more generally, these two processes can be combined with a describing the net effect. While this is an adequate model for some phases of the dynamics, often used in short-term models or in epidemiology, over longer time periods saturation effects need to be taken into account. Here, the most common models are *logistic growth*, $\dot{x} = ax \left(1 - \frac{x}{K}\right)$, with K denoting a fixed carrying capacity of the population. Another saturation type model, which has been verified empirically as the model for late-stages in breast cancer [16,17], is *Gompertzian growth* of the form $\dot{x} = -ax \ln\left(\frac{x}{K}\right)$. Logistic growth is often preferred because of its simpler mathematical structure. Other growth models exist, but are less common (e.g., see [18]).

2.1.2 Transitions Between the Compartments

Transition terms model the exchanges between populations. For example, in *cell-cycle specific models for cancer chemotherapy* [8], quiescent and proliferating cells are distinguished and proliferating cells move through growth phases, synthesis and mitosis, all represented by different coordinates of the state vector x . Similarly, in models including *drug resistance*, random changes from sensitive to resistant cells occur which, in cases of cancer chemotherapy, also may be reversible. In *epidemiological models*, transitions between sensitive, exposed, infectious and recovered individuals are crucial to understand the spread of the disease. All such transitions are generally modeled by linear transition rates of the form

$$\dot{x}_i = -\alpha_i x_i + \sum_{j \neq i} \beta_{ij} x_j,$$

with $\alpha_i \geq 0$ representing the outflow from the i th compartment and $\beta_{ij} \geq 0$ for $i \neq j$ representing the inflow from the j th into the i th compartment.

2.1.3 Interactions Between the Compartments

Cells in various compartments stimulate and/or inhibit cells in other compartments. Often these interactions are modeled by the law of mass actions leading to products of terms representing the cell counts or volumes in various compartments. This generates expressions of the form $x_i x_j$ or, more generally, expressions which include powers of the variables. The simplest structure shows up in *epidemiological models* when susceptible and infectious populations S and I interact generating an outflow of the form βSI from the susceptible compartment which becomes an inflow to the infectious population I with β accounting for the probability of such interactions. Powers arise, for example, in a model for *angiogenic signaling* by Hahnfeldt et al. [19] where the interactions are between the tumor surface through which angiogenic inhibitors need to be released and the carrying capacity of the vasculature. If both tumor volume and the carrying capacity are modeled by volumes p and q , respectively, this generates an interaction term of the form $p^{\frac{2}{3}}q$. In a classical model by Stepanova of *tumor immune system interactions* [20], the fact that the immune system is able to control small tumors, but increasingly becomes overwhelmed by large tumors is modeled by an interaction term of the form $\alpha(x - \beta x^2)y$ where x represents the tumor volume and y a non-dimensional, order of magnitude variable related to the actions of the immune system called the immunocompetent cell density. Here, $x = \frac{1}{\beta}$ corresponds to a threshold beyond which the immunological system becomes depressed by the growing tumor. The coefficients α and β are used to calibrate these interactions and collectively describe state-dependent interactions between the cancer cells and the immune system. Limiting effects of the interactions between various compartments can be taken into account by so-called Michaelis–Menten terms. As an example, in a model for CML (chronic myeloid leukemia) [21] which includes proliferating cancer cells P and effector cells E of the immune system, this leads to terms of the type

$$\frac{P_{\max} \cdot P}{PC_{50} + P} \cdot E$$

with P_{\max} denoting the maximum possible effect and PC_{50} the concentration of proliferating cancer cells for which half of the maximum effect is realized. Similar terms describe stimulating or inhibiting effects of such interactions depending on the nature of the variables.

2.1.4 Multi-stability and Regions of Attractions

Important aspects of the uncontrolled dynamics are **multi-stability** and the geometry of the **regions of attraction** of locally stable equilibria. We illustrate these features using two classical models for tumor growth and its interactions with the tumor micro-environment: the model for angiogenic signaling by Hahnfeldt et al. [19] and Stepanova's model for tumor immune-system interactions [20,22,23]. These

models also serve as illustrations for the general elements of the dynamics described above.

(a) *angiogenic signaling* In this model by Hahnfeldt et al. [19], the spatial aspects of the underlying consumption-diffusion process that stimulate and inhibit angiogenesis are incorporated into a non-spatial 2-compartment model with the *primary tumor volume*, p , and the *carrying capacity* of the vasculature, q , as its principal variables. The dynamics consists of two ODEs that describe the evolution of the tumor volume and its carrying capacity which are given by

$$\dot{p} = -\xi p \ln\left(\frac{p}{q}\right), \quad (2)$$

$$\dot{q} = bp - dp^{\frac{2}{3}}q - \mu q. \quad (3)$$

Thus, a Gompertzian growth term is used to model tumor growth. The carrying capacity q and tumor volume p are balanced for $p = q$, while the tumor volume shrinks for inadequate vascular support ($q < p$) and increases if this support is plentiful ($q > p$). Different from traditional approaches, the carrying capacity is not considered constant, but becomes a state variable whose evolution is governed by a balance of stimulatory and inhibitory effects. The term bp represents the stimulation of the vasculature by the tumor and is taken proportional to the tumor volume. The term $dp^{\frac{2}{3}}q$ models the inhibition of the vasculature by the tumor which, as mentioned above, is represented as an interaction term between the tumor surface and the vasculature. The coefficients b and d are mnemonically labeled for birth and death, respectively. The last term $-\mu q$ simply represents a natural death term related to the vasculature.

This model has a unique equilibrium at

$$\bar{p} = \left(\frac{b - \mu}{d}\right)^{\frac{3}{2}} = \bar{q} \quad (4)$$

which is globally asymptotically stable (relative to the natural state-space $\mathcal{D} = \{(p, q) : p > 0, q > 0\}$) [24] (see Fig. 1). The dynamics has a strong differential-algebraic character with the q -dynamics of the vasculature fast and the p -dynamics for the tumor volume slow. Essentially, the system follows the corresponding null-cline (defined by the equation $\dot{q} = 0$) into the unique equilibrium point. For realistic values of the parameters, this equilibrium is not biologically viable and therapy is needed to lower or even eradicate this equilibrium point. In the later case, the objective then becomes to drive the state of the system toward the singular point $(p, q) = (0, 0)$.

(b) *Tumor immune system interactions* In this by now classical mathematical model of Stepanova [20], two ordinary differential equations formulate the aggregated interactions between cancer cell growth and the activities of the immune system during the development of cancer. Precisely because of its simplicity—a few parameters incorporate many medically important features—the underlying equations have been widely accepted as a basic model. The main features of tumor immune system interactions are aggregated into just two variables, the tumor volume, p , and the immunocompetent cell densities, r , a non-dimensional, order of magnitude quantity related to various

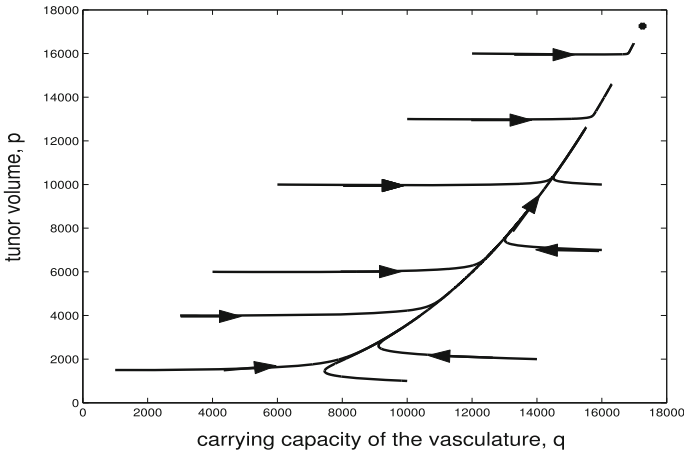


Fig. 1 Phase portrait for the dynamical systems (2)–(3) for the parameter values $\xi = 0.192$, $\mu = 0.02$, $b = 5.85$ and $d = 0.00873$. The parameter values are based on biologically validated data for Lewis lung carcinoma in mice [19], a fast growing cancer. The almost horizontal segments indicate a strong differential-algebraic character with trajectories eventually following the \dot{q} -nullcline defined by $\dot{q} = 0$

types of immune cells (T-cells) activated during the immune reaction. The resulting dynamics takes the following form:

$$\dot{p} = \xi p F(p) - \theta pr, \quad (5)$$

$$\dot{r} = \alpha(p - \beta p^2)r + \gamma - \delta r, \quad (6)$$

with F a tumor growth function. Greek letters denote constant coefficients which model the main interactions of the tumor with the immune system. Various organs such as the spleen, thymus, lymph nodes and bone marrow, each contribute to the development of immune cells in the body and the parameter γ in Eq. (6) models a combined rate of influx of T-cells generated through these primary organs; δ is simply the rate of natural death of the T-cells. The first term in this equation models the proliferation of lymphocytes. For small tumors, it is stimulated by tumor antigen and this effect is taken to be proportional to the tumor volume p . Large tumors suppress the activity of the immune system, mostly because of a general suppression of immune lymphocytes. This feature is expressed in the model through the inclusion of the term $-\beta p^2$. Thus, $1/\beta$ corresponds to a threshold beyond which the immunological system becomes depressed by the growing tumor. The coefficients α and β are used to calibrate these interactions and collectively describe state-dependent influence of the cancer cells on the stimulation of the immune system. The first equation, (5), models tumor growth with ξ a tumor growth coefficient that has been separated from the qualitative behavior of the growth function F . The second term, $-\theta pr$, models the beneficial effects of the immune system reaction on the cancer volume and θ denotes the rate at which cancer cells are eliminated through the activity of T-cells.

The dynamical behavior of this system is fundamentally different from the model for angiogenic signaling. For a typical set of parameter values, the dynamics is multi-stable

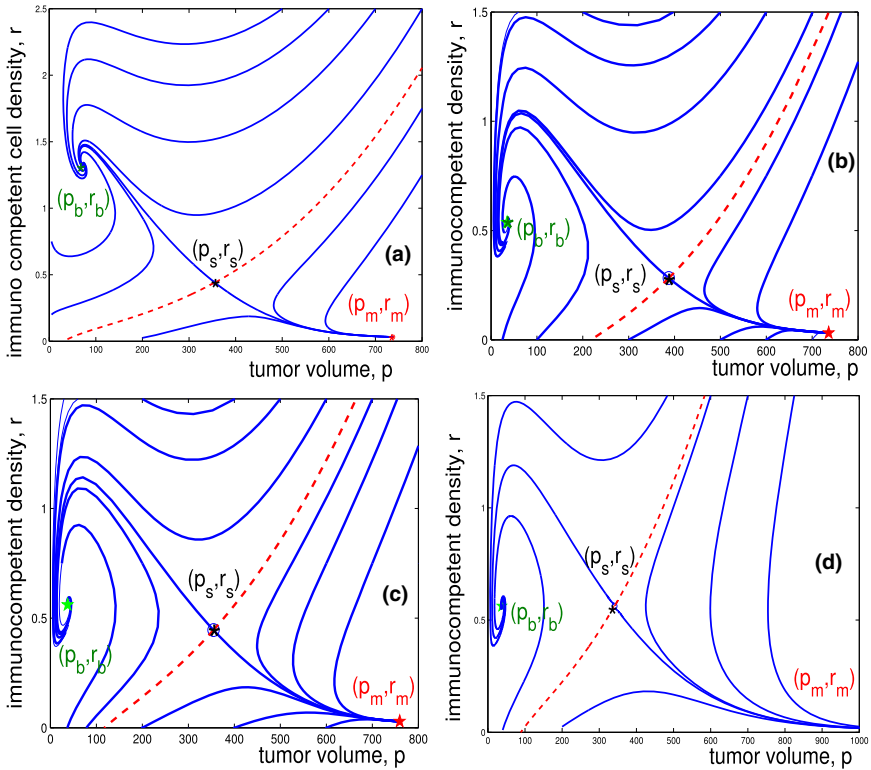


Fig. 2 Phase portraits for the system (5)–(6) for the parameter values given in Table 1 and varying growth functions F . The following growth functions were used in the respective panels: **a** Gompertzian growth: $F_G(p) = -\ln\left(\frac{p}{p_\infty}\right)$, **b** logistic growth: $F_L(p) = 1 - \frac{p}{p_\infty}$, **c** generalized logistic growth with exponent 2: $F_{GL}(p) = 1 - \left(\frac{p}{p_\infty}\right)^2$, and **d** exponential growth: $F_E(p) \equiv 1$. The points (p_b, r_b) and (p_m, r_m) are the locally stable benign and malignant equilibria, respectively, and (p_s, r_s) denotes the unstable saddle point. The dashed red curve is the stable manifold of this saddle

and the system has both locally asymptotically stable microscopic and macroscopic equilibrium points as well as an unstable saddle point. Figure 2 shows phase portraits of the system for the parameter values listed in Table 1 for four different growth models which qualitatively all exhibit the same multi-stability. At the microscopic equilibrium point, the tumor volumes are small and immuno-competent cell densities are upregulated. This corresponds to a situation when the immune system is controlling the tumor. We denote the corresponding equilibrium point by (p_b, r_b) and call it *benign*. The macroscopic equilibrium point is characterized by more than tenfold higher tumor volumes and depressed immunocompetent cell densities. In these solutions, the tumor has suppressed the immune system and it almost reached its carrying capacity. We denote the corresponding equilibrium point by (p_m, r_m) and call it *malignant*. Both of these equilibria are locally asymptotically stable, (p_b, r_b) a stable focus and (p_m, r_m) a stable node. We call the regions of attraction associated with the benign and malignant equilibria the *benign*, respectively, the *malignant region*. For the model with an

Table 1 Variables and parameters used for the phase portraits shown in Fig. 2

Variable parameters	Interpretation	Numerical value	References
p	Tumor volume		[20]
p_∞	Fixed tumor carrying capacity	780	[22]
r	Immuno-competent cell density		[20]
α	Tumor stimulated proliferation rate	0.00484	[22]
β	Inverse threshold for tumor suppression	0.00264	[22]
γ	Rate of influx of T-cells	0.1181	[22]
δ	Death rate	0.37451	[22]
θ	Tumor–immune interaction rate	1	[22]
ξ	Tumor growth parameter	0.5618	

exponential growth function, the tumor size is not limited and in this case no malignant equilibrium exists, but the malignant region is characterized by the fact that the p -component of the trajectories diverges to ∞ , while the r -component converges to 0. The benign and malignant regions are separated by the stable manifold (shown as a dashed red curve in each diagram) of a saddle point which forms the common boundary of these regions.

These phase portraits represent a snapshot of the true dynamics: a stationary situation for fixed parameter values. Because of unmodeled aspects of the dynamics and/or random events, however, this situation will constantly change in time. Thus, it is intuitively clear that a benign equilibrium solution can be disrupted by events affecting the immune system which were not included in the mathematical modeling which then may lead to a transition of the state into the malignant region. How likely this becomes is related to the size of the region of attraction of this equilibrium: If the benign region is small, even minor events may bring up the disease, whereas the immune system may be able to control the disease if this region is large. While there exist good reasons to believe that the immune system is able to control some tumors initially, over a longer period of time, the neoplasm will develop various strategies to evade the actions of the immune system (*immuno-editing*) and this allows the tumor to recommence growing into clinically apparent tumors and eventually reach its carrying capacity [25]. Tumor–immune system interactions thus exhibit a multitude of dynamic properties that include persistence of both benign and malignant scenarios. From a practical point of view, the question then becomes how to move an initial condition that lies in the malignant region into the benign region. This requires therapy and can naturally be formulated and analyzed as an optimal control problem.

2.2 The Controlled Dynamics: Mathematical Models for Therapy and Other Interventions

Besides the underlying dynamics represented by the drift vector field F , the other essential ingredient in biomedical problems is a mathematical model for outside inter-

ventions. These comprise therapies aimed at curing a patient in cancer treatment or actions to limit the spread of an epidemic. Within the general structure (1) of the model considered in this paper, the effects of these actions are modeled linearly through the *control vector fields*

$$G(x)u = \sum_{i=1}^m G_i(x)u_i. \quad (7)$$

Here the control variables u_i , $i = 1, \dots, m$, represent various therapy options such as chemotherapy, radiotherapy, immunotherapy in mathematical models for cancer treatments and/or other available actions to influence the state of the system such as, for example, social distancing measures in epidemiological models. The components g_{ij} , $j = 1, \dots, n$, of the vector field G_i , $i = 1, \dots, m$, model the effects which the j th intervention modality u_j has on the i th compartment x_i of the state.

Generally, the number m of different modalities for outside interventions in a biomedical problem setting is *small* and m typically lies in the range of $m = 1, 2, 3$. Even for the current COVID-19 pandemic, in essence, all efforts to limit the spread of the pandemic are related to social distancing measures in an attempt to lower the probabilities of infection. If available, vaccination would be a different second intervention strategy, but, clearly, the overall number of options is not very large. Similarly, in cancer therapy, the number of possible qualitatively different treatments options is small. If one ignores surgery, as this is an option that would be pursued at the onset prior to any time frame that typically would be modeled mathematically, the main modalities are chemotherapy, radiotherapy and immunotherapy. If chemotherapy is given through a cocktail of drugs (and this is the common procedure), then it is still typically advantageous to model the effects of the combination of all the drugs together and one may have $m = 1$, while the case $m > 1$ often is reserved for true *combination therapies* that combine, for example, chemo- or radiotherapy with antiangiogenic treatment. Thus, the typical scenario is that the number m of controls in the system is small.

2.2.1 Strengths and Limitations of Control Affine Models

The particular mathematical structure considered in Eq. (7)—linear in u —clearly does not cover the full range of possibilities, but it is worth emphasizing that this is by far the most common, most important, and, in many aspects, most natural structure.

Essentially, whether or not a control affine model (7) is valid depends on the choice of what the controls represent. If the controls are tied in with the effects of treatment (e.g., the control represents the concentration of an agent rather than its dose rate), then this model may not be appropriate. If, however, the controls u_i are related to the administration of the therapies (dose rates) and are separated from their effects, these models will be linear in the controls. Below we shall discuss in more detail mathematical models for drug therapies, but we want to mention at least in passing the main model for radiotherapy, the so-called *linear quadratic model*.

In this model, the probability of cell survival is represented in the form $\exp(-\alpha D - \beta D^2)$ where D denotes the total radiation dose and α and β are radiosensitive parameters. Briefly, the rationale for the model is as follows [15,26]: Radiation

causes abnormalities (lesions) that correspond to ruptures of the molecular bonds on the double-stranded DNA. This damage is made up of two components: (i) simultaneous breaks in both DNA strands that are caused by a single particle and (ii) adjacent (both in location and time) breaks on separate strands. While a double-strand break leads to a loss of proliferative abilities, a single-strand breakage is not considered lethal since DNA has the ability to repair it. Only if a second adjacent break occurs on the other strand before the first one can be repaired, this will lead to loss of proliferation properties. Both scenarios (i) and (ii) are considered lethal in the sense that the cell is no longer able to proliferate. The first situation leads to a linear term in the cell survival model, while the second one contributes quadratic terms. This is accounted for by the parameters α and β which are called the LQ parameters.

Repair of DNA damage generally is incomplete. Once this is taken into account as well, then the surviving fraction of cancer cells can be represented in the form

$$\exp\left(-\alpha \int_0^T u(t)dt - 2\beta \int_0^T \int_0^t u(t)u(s)e^{-\rho(t-s)}dsdt\right)$$

with u denoting an arbitrary, time-varying dose rate, $u = u(t) \geq 0$, over the therapy interval $[0, T]$ and ρ the tissue repair rate. The total dose D is given by $D = \int_0^T u(t)dt$. This model can be written in a control-affine form [15,27] as follows:

$$\dot{p} = pF(p) - (\alpha + 2\beta r)pu, \quad p(0) = p_0, \quad (8)$$

$$\dot{r} = -\rho r + u, \quad r(0) = 0. \quad (9)$$

Here, p once more represents the tumor volume and the equation for r models the temporal effects of tissue repair; α and β are the LQ parameters. The specific parameter values depend on the tissue treated, and there will be separate equations modeling the effects on cancerous and healthy tissue. The control u represents a time-varying fractionated radio-therapy dose rate. The linear model in u represents the way therapy is administered, and this almost always is adequately modeled by a linear term. The quadratic effects of radiotherapy are subsumed at the expense of introducing the extra state variable r .

Similarly, while the actions of different treatment modalities clearly are interrelated, this concerns their effects, not the way the therapies are administered. If chemotherapy is considered with multiple drugs, then synergistic or antagonistic effects which cannot be modeled linearly become important. Aside from the fact that such interactions often are far from being understood, it also matters what the control u represents. If the control represents drug concentrations, then such interactions enter into the model through the controls and a linear model is inadequate. On the other hand, if the controls represent the dose rates of the various agents, then the concentrations of these agents merely become additional states and these interactions enter into the drift vector field F of a generally already nonlinear dynamics. This then preserves the control linear structure (7).

In essence, if the controls u_i , $i = 1, \dots, m$, represent the administration of the therapies (dose rates) and as variables are separated from the effects of the actions

(which, for example, depend on the concentrations), then a model which is linear in the controls is not only adequate, but is the correct one.

We expand further on this for the case of chemotherapy and, as a matter of simplicity of the presentation, focus on the case of a single control u .

2.2.2 Pharmacodynamics: Modeling Drug Therapies Through Drug Concentrations

Pharmacodynamics is the branch of pharmacology which quantifies the effects which a drug has at concentration c . According to the law of mass action, the speed of a chemical reaction is proportional to the product of the active masses (concentrations) of the reactants. If a drug is administered, in an ideal situation it has been postulated that cell death under cancer drugs follows first-order kinetics [28], i.e., the decrease in the number of cancer cells per unit of time is proportional to the number of cancer cells with the rate depending on the concentration of the anti-cancer drug. Any functional relation of the form $e = s(c)x$ with x denoting the cells in a compartment on which the drug is acting (e.g., tumor cells, immune system, healthy cells, etc.) and the coefficient $s(c)$ representing the concentration dependent killing rate of the drug is considered a pharmacodynamic model. A linear model of the form $s(c) = \varphi c$ with φ a positive constant, the so-called linear *log-kill hypothesis* based on Skipper's paper [28], is the most commonly used form. Incorporating this structure into a standard growth law results in the following growth model under chemotherapy,

$$\dot{p} = pR(p) - \varphi cp. \quad (10)$$

If the concentration c is considered the control u in the system, $u = c$, then the linear log-kill model fits the general linear form (7). The product represents an interaction term as discussed in Sect. 2.1.3. Such terms can be both inhibiting (minus sign) and stimulating (plus sign). Clearly, in the example considered above, the action is inhibiting as the agent is cytotoxic and kills the cells. On the other hand, in Stepanova's model the action of an immune stimulatory agent (such as interleukin-2, IL-2, representing immunotherapy) would be included through the addition of a positive term $+\rho ru$ into Eq. (6).

While a linear model $s(c) = \varphi c$ is applicable over a specified range of the concentration c , it is not valid for low or high drug concentrations. This simply stems from the fact that most drugs have little to no effect if their concentrations are too low and, on the other hand, their effectiveness saturates at high concentrations. If side effects allow, drugs are preferably given in the upper range of this dose-effect curve and an E_{\max} model of the form

$$s(c) = \frac{E_{\max}c}{EC_{50} + c}, \quad (11)$$

describes the effect of "fast" acting drugs more accurately for high concentrations. Here, EC_{50} denotes the concentration at which half of the maximum effect E_{\max} is realized. Sigmoidal functions also capture the behavior at lower concentrations. In these cases, if the drug concentrations are used as the control, after normalization this leads to terms of the form $\frac{u}{1+u}$ which no longer fit the proposed structure (7). This,

however, is merely caused by using the concentrations as the controls. It can easily be removed by changing the model to treat the dose rates of the drug administration as the control.

2.2.3 Pharmacokinetics: Modeling Drug Therapies Through Drug Dose Rates

The second major aspect of drug delivery is pharmacokinetics. Pharmacokinetic equations model the time evolution of a drug's concentration in the blood plasma. The standard model that describes the concentration $c = c(t)$ of the drug in the bloodstream in response to a continuous dose rate $u = u(t)$ is one of exponential growth and decay given by

$$\dot{c} = -\rho c + u, \quad c(0) = 0, \quad (12)$$

with ρ the *clearance rate* of the drug. More generally, if the concentrations in the bloodstream are separated from the concentrations in a peripheral compartment (such as tissue or an organ of interest) and the interactions between those compartments are taken into account, a two-dimensional linear system with positive eigenvalues is used. Higher-dimensional compartmental models are used less frequently, but arise, for example, if the peripheral compartment is divided further into a 'shallow' and 'deep' compartment. For example, a 3-compartment model is often used for insulin pumps [29].

All these models, however, are described by linear time-invariant systems with real eigenvalues. If one therefore introduces the drug *dose rates* as the controls, then the associated drug concentrations c become state variables and thus become an integral part of the modeling of the uncontrolled dynamics that enters the form of the drift vector field F . In particular, all the nonlinearities that arise because of saturation effects or because of synergistic properties of drug administration are subsumed in the drift vector field F , which, as the example for angiogenic signaling shows, may already be a highly nonlinear vector field to begin with. Actual drug administration, however, is described by linear models.

3 Significance of the Choice of the Objective

Once medical professionals have decided on a course of action, a number of natural questions need to be answered: (1) HOW MUCH of the agent should be given? (2) WHEN and HOW OFTEN should drugs be given? (3) IN WHAT ORDER should the treatments be applied? All these questions can naturally be answered within the framework of an optimal control problem. Clearly, we are only dealing with a mathematical model and no claims are made that the real problem could be solved in this way. But interesting and useful information can be obtained. To give a concrete example, chemotherapy needs the tumor vasculature to deliver cytotoxic agents that kill the tumor cells. It therefore was conventional wisdom to first give chemotherapy when antiangiogenic treatments and chemotherapy were combined in the treatment of cancer. This notion, however, was challenged by Jain et al. [30,31] who argued that a normalization of the vasculature ("pruning") prior to the administration of chemotherapy is

beneficial for the delivery of the cytotoxic agents and that it gives better results. These findings can be strongly supported through a mathematical analysis of the Hahnfeldt model with both chemotherapy and antiangiogenic therapy [32]. Thus, while optimal control cannot provide the ultimate answers to the underlying problems, it can become a valuable tool in understanding these processes.

3.1 A General Form for the Objective

The ultimate aim in biomedical problem formulations is to find a safe and reasonable strategy to solve the underlying medical problem. By adopting an optimal control formulation, we are aiming for what in a certain way can be considered the “best possible” solution. What this best possible solution should be, however, is often up to wide interpretations. It is here that the role of the objective functional chosen to minimize matters greatly—sometimes to the extent that this choice already preprograms the eventual solution. This choice then is the third essential aspect of the problem formulation.

The following general form encompasses a wide range of possibilities:

$$\text{minimize } J(u) = \int_0^T \left(L(x(t)) + \sum_{j=1}^m \theta_j u_j^k(t) \right) dt + \varphi(x(T)), \quad k = 1 \text{ or } 2. \quad (13)$$

Here, L is a sufficiently smooth function of the state called the Lagrangian, the coefficients θ_j , $j = 1, \dots, m$, are nonnegative weights and φ is a penalty function on the terminal state. The terminal time T can be fixed or free. The minimization is over a class \mathcal{U} of *admissible controls* u which are locally bounded Lebesgue measurable functions such that the associated controlled trajectory (x, u) satisfies all other constraints imposed on the system. These generally include control-constraints and terminal point constraints, but could also come from a wide variety of other restrictions imposed possibly by state-space constraints or other limitations inherent to the problem formulation. We discuss the significance of each of the terms starting with the controls.

3.2 Modeling Toxicities as a Soft Constraint: L_1 Versus L_2 Approaches

Generally, and without loss of generality, we make the assumption that the controls represent nonnegative quantities (e.g., dose rates or concentrations). Thus, the control term $\sum_{j=1}^m \theta_j u_j^k(t)$ in the objective (13) for $k = 1$ or $k = 2$ represents a weighted L_1 - or L_2 -norm, respectively. Including such a positive and increasing term in the objective represents what in the engineering literature would be called a *soft modeling of control constraints*. Clearly, medical treatments, especially the strong treatment options that need to be pursued in cancer treatment such as chemo- or radiotherapy, have significant side effects and these need to be taken into account in the modeling. Including the control term in the objective to be minimized inherently limits the use of the controls. This represents an attempt to make a reasonable compromise between

maximizing the tumor kill which requires to give a large amount of drugs and limiting the side effects of treatment (these are measured indirectly through the integral of the control) which requires to limit this amount of drugs. Whether the optimal control does this adequately, then becomes an after-thought that always needs to be checked and, if the optimal solution is still inadequate, the weights need to be adjusted. If the weights on the drugs are too small, optimal solutions will simply give drugs all the time at full dose; if the weights are too high, no drugs will be given. It is clearly necessary to calibrate the weights to obtain meaningful results. Such back-and-forth approaches are quite common in engineering and they simply reflect the fact that the form of the objective, and especially the weights that appear, are *variables of choice* in such problems. Generally, and this is the typical application of optimal control in engineering, various choices of the weights will be tested until a solution has been found which appears adequate, possibly also with respect to other criteria which were not included in the mathematical model.

The choice of taking an L_1 - or L_2 -norm, however, is not irrelevant. Generally speaking, an L_1 -formulation is to be preferred from a modeling perspective while the L_2 -formulation offers significant mathematical and numerical advantages.

In chemotherapy, there exists a clear interpretation for the integral $\int_0^T u(t)dt$ related to the total dose given regardless of whether u represents the dose rate of the drug or its concentration. This is an essential pharmacological quantity also referred to as the AUC (area under the curve), and it plays a major role in interpreting the overall efficacy and side effects of a drug. Thus, making the choice of an L_1 -formulation as the penalty of the control leads to a realistic modeling. The handicap of an L_1 -type approach lies with the fact that the mathematical analysis generally becomes difficult. Optimal solutions need to be synthesized from concatenations of bang and singular controls (and we shall discuss this more in the next section) and reliable numerical methods to solve such problems do not exist.

It is for this reason—and in our opinion, for this reason only—that the use of an L_2 -criterion in the literature is wide-spread. For, in this case, the application of the necessary conditions for optimality of the Pontryagin maximum principle [33–36] leads to a closed form formula for the control as a function of state and multiplier that can easily be solved numerically. This, however, does not guarantee that a numerically computed solution is optimal as only extremals (pairs of controlled trajectories and multipliers which satisfy the first-order necessary conditions for optimality) are obtained, a fact often understated by many researchers. Indeed, the projection from the space of extremals into the state space (i.e., projecting out the multiplier) can have—and often does have—singularities, the equivalent of the well-known conjugate points from the calculus of variations. For problems employing an L_2 -type objective, higher-order tests for local optimality are given in [15, Chapter 4] or [37] and are not difficult to apply, but generally this is not done in the literature.

Our main objection to the use of an L_2 -criterion in the controls, however, does not lie with this mathematical aspect, but simply with the fact that for biomedical problems it seems rather impossible to give a satisfactory justification for such a functional form as the meaning of the phrase “systemic cost” employed in the vast literature is hard to capture. In chemotherapy, for example, not only does the quantity

$\int_0^T u(t)^2 dt$ not represent any pharmacologically meaningful quantity, but in addition, if one normalizes the controls to satisfy $0 \leq u_i \leq 1$, as it is often done, then this strongly favors smaller partial doses over the higher full doses as half the maximum dose is only given a quarter of the weight of the maximum dose. Thus, from the very setup, a bias is built into the optimal solutions to be computed which the underlying system does not have. Cost is another questionable justification for an L_2 -type objective as $\2 just is not a very meaningful quantity. It is certainly true that the cost need not be linear in u . This particularly holds in epidemiological models where the cost of vaccinating the first 10% of a population clearly is much smaller than the cost of vaccinating the last 10% of this population as much greater efforts will be needed to reach the tail end of the population. Yet, while this clearly makes a point for some nonlinear relation, a quadratic term does not seem to be the answer to this.

We still remark that it is often the typical scenario that optimal solutions for the L_2 -type term (which necessarily are continuous) exhibit sharp rises from the lower control value $u = 0$ to its maximum value $u = u_{\max}$ mimicking optimal bang–bang controls which would be obtained by using an L_1 -type term. Indeed, from a numerical perspective, it makes sense to consider the problem when the control terms are taken in the form $u + \varepsilon u^2$ and then the limit as $\varepsilon \rightarrow 0$ is taken (if it exists). We note, however, that our interest in solving the optimal control problems is more in trying to obtain a better global understanding of the dynamics, i.e., the underlying problem, than to obtain one particular numerical solution (which is a much easier problem). As a rule, the parameter values in the model are uncertain and the solutions can be quite sensitive to these values (e.g., for Stepanova’s model) which limits the practical use of small numerical studies.

3.3 Example: Cell-Cycle Specific Chemotherapy

We illustrate some of these effects with locally optimal solutions for a cell-cycle specific model for cancer chemotherapy due to Swierniak [8,12]. In such models, the various phases of the cell-cycle—a quiescent phase G_0 , a first growth phase G_1 , synthesis S , a second growth phase G_2 , and mitosis M —are clustered into compartments and the dynamics describes the transitions, i.e., the flow between these compartments. Mathematically, the dynamics becomes a bilinear system of the form

$$\dot{N}(t) = \left(A + \sum_{j=1}^m u_j B_j \right) N(t), \quad N(0) = N_0, \quad (14)$$

with the state N_i representing the average number of cells in the respective compartments and the matrices A and B_j modeling the flow rates between the compartments. The controls represent cytotoxic agents which indiscriminately kill cells, cytostatic agents which block the transitions of cells through the cell-cycle and recruiting agents which entice cells to leave the quiescent phase and become active again so that it is possible to kill them. A typical form of the objective is given by




Fig. 3 The top row shows locally optimal controls for a 3-compartment model with cytotoxic agent u (shown in red) and cytostatic agent v (shown in blue) for the objectives J_2 (left) and J_1 (right) with weights $r = (1, 1, 1)$, $q = (1, 1, 1)$ and $\theta = (2, 0.02)$. The time-evolution of the corresponding trajectories is shown in the bottom row

$$J_k = \sum_{i=1}^n r_i N_i(T) + \int_0^T \left\{ \sum_{i=1}^n q_i N_i(t) + \sum_{j=1}^m \theta_j u_j^k(t) \right\} dt, \quad k = 1 \text{ or } 2, \quad (15)$$

with the coefficients nonnegative weights.

Figure 3 compares locally optimal solutions for a 3-compartment model with the compartments $N_1 \sim G_0/G_1$, $N_2 \sim S$ and $N_3 \sim G_2/M$ for L_1 - and L_2 -type functional terms on the controls which represent a cytotoxic agent u and a cytostatic agent v [11, 15,37]. Optimal controls for the L_1 -type objective are bang–bang, i.e., switch between full dose therapies and rest-periods while optimal controls for the L_2 -type objective are necessarily continuous. For an L_1 -type objective, the optimal administration of the killing agent u is full dose initially followed with a rest period (and this is what medically makes most sense here), whereas for the L_2 -type objective the continuity of the control generates decreasing lower dose administrations of the cytotoxic agent toward the end of the therapy interval. The total doses given in both cases are similar. For the blocking agent v , the optimal L_2 -control clearly mimics the optimal L_1 -control with steep continuous intermediate segments.

The behavior of the controlled trajectories for both objectives is very much the same. The terminal values for the states corresponding to the L_1 - and L_2 -type objectives are

close, and the reduction in the total cancer burden is from the initial normalization of $C(0) = 1$ to $C(21) = 0.8673$ for minimizing J_1 and to $C(21) = 0.8314$ for minimizing J_2 . In both cases, the initial condition was chosen as the normalized steady-state of the uncontrolled dynamics (i.e., assuming no earlier treatment). Computing the total dose of the drugs given for the L_2 -optimal control, one sees, however, that while minimizing the quadratic objective leads to about a 3% improvement in the reduction of the tumor, this is at the expense of increasing the amounts of both the cytotoxic and cytostatic agents by about 20%. Thus, higher side-effects are encountered. Minimizing J_1 would appear to be the preferred strategy.

3.4 Modeling Toxicities as a Hard Constraint: Isoperimetric Constraints

An alternative approach—and the one which seems to be preferred by medical practitioners—is to impose the total amount of drug (or radio therapy dose etc.) given a priori. The argument is that, to begin with, any mathematical model represents only a small view of the underlying problem and that, by sheer necessity of the complexity which, for example, cancer therapy is, many aspects need to be neglected. Thus, the point of view is taken that, based on other assessments such as the side effects of the therapy, the total amount of drug to be given has been determined a priori, $\int_0^T u(t)dt \leq A$. Here, then the question simply becomes the following one: given this amount of drug, how can it be administered in the best possible way? Mathematically, this amounts to adding the trivial equation $\dot{y} = u$, $y(0) = 0$, to the dynamics and the terminal condition $y(T) \leq A$ is imposed as an isoperimetric constraint. This approach, for example, was taken in solving Hahnfeldt's model for antiangiogenic therapy as an optimal control problem [13,15]. In such a case, no control related term will be included in the objective.

From a modeling perspective, the use of a hard constraint on the controls seems to be more adequate. Mathematically, however, there are close relations between the necessary conditions for optimality for both formulations (either as a soft or hard constraint) and it makes little difference if one or the other formulation is chosen.

3.5 Significance of the Lagrangian and the Penalty Term

The Lagrangian L and the penalty term φ allow to judge the impact of the therapy on the state. Including the Lagrangian L in the objective can also be viewed as incorporating a soft constraint on the state. For example, in cancer therapies, it is clearly important to limit the size of the tumor, or, in another indirect measurement of the side effects, prevent that the level of the bone marrow becomes too small. Clearly, such restrictions could be modeled as state-space constraints imposing upper or lower bounds on some of the variables. For example, in [38] an explicit upper bound in the form

$$N(t) = \sum_{i=1}^n N_i(t) \leq \bar{N}$$

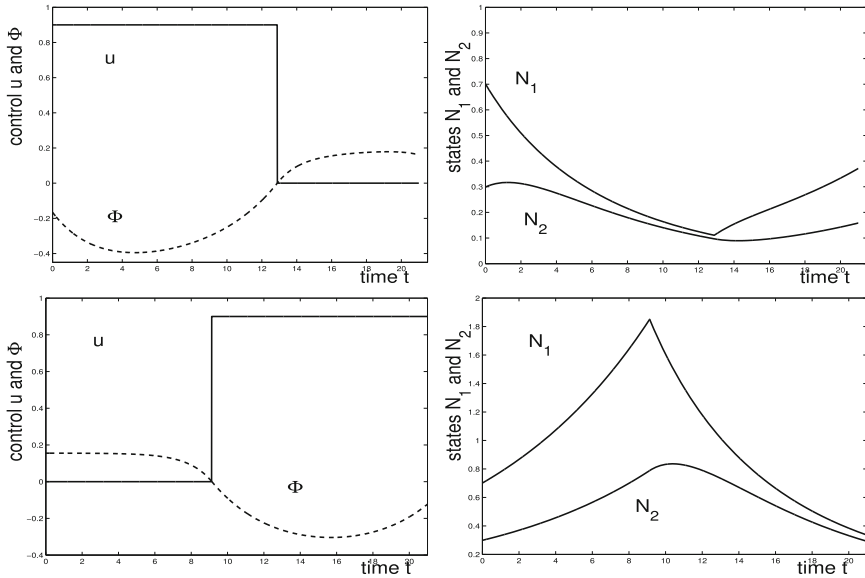


Fig. 4 Examples of locally optimal controls (left) and their corresponding trajectories (right) for a cell-cycle specific 2-compartment model for cancer therapy over a fixed therapy horizon ($T = 21$ days) with Lagrangian term L in the objective (top row, $q = 0.1$) and without Lagrangian term in the objective (bottom row, $q = 0$). The initial condition is the normalized value of the steady-state solution for the uncontrolled system

is imposed on the total number of cancer cells $N(t)$ in a cell-cycle specific model for cancer chemotherapy. Here, however, a modeling approach of using soft constraints is preferred as the necessary conditions for optimality become more involved if state-space constraints are present. If explicit state-space constraints are included in the modeling, then, a priori, the associated multiplier is only known to be a positive Radon-measure. Although it can be shown that this measure is absolutely continuous with respect to Lebesgue measure if the constraint is an embedded hypersurface (e.g., see [39]), the effort of dealing with a state-space constraint is much more involved than having a Lagrangian present in the objective.

A common pitfall in the modeling is not to include such a term. Then, naturally, the emphasis of the minimization is put on the terminal time and quite often control actions will be delayed initially and be concentrated near the terminal time. Once again, however, this is then not a feature of the underlying problem, but it has been imposed, possibly unwillingly, by the particular form of the objective that was chosen. Figure 4 illustrates this feature by comparing the optimal solutions for a 2-compartment cell-cycle specific model for cancer chemotherapy [8,10,15]. The two compartments represent the clustered phase $N_1 \sim G_1/G_0 + S$ and $N_2 \sim G_2/M$ of the cell cycle and the drug is a G_2/M -specific cytotoxic agent. The objective was chosen in the form

$$J(u) = \int_0^{21} \left(q(N_1 + N_2) + \frac{1}{2}u \right) dt + 3(N_1(T) + N_2(T)) \tag{16}$$

for $q = 0.1$ and $q = 0$ over a fixed 3 week therapy period. Optimal controls are bang–bang with one switching. If only a penalty term is used, then the drug administration is all put to the end of therapy with an unacceptable spike in the tumor cells in between.

3.6 Biological Guidance for the Choice of the Objective

The choice of the objective is often made in an heuristic ad hoc manner. Sometimes, however, the underlying biology can be used as a guide. Stepanova’s model provides a beautiful illustration how this can be achieved. Assuming that the parameter values are such that the dynamics is bi-stable with a benign and a malignant region, the practical aim of therapy then can be formulated as moving the state of the system from the malignant region into the region of attraction of the stable, benign equilibrium point while keeping side effects tolerable. Assuming the controls are a strongly targeted cytotoxic agent u and an immuno-stimulatory agent v , this can be achieved using an objective functional of the following form:

$$J = Ap(T) - Br(T) + \int_0^T (Mu(t) + Nv(t) + S) dt. \quad (17)$$

All coefficients are positive and the choice of the weights aims at striking a balance between the benefit at the terminal time and the side effects measured by the total amount of drugs given while it also guarantees the existence of an optimal solution by penalizing the free terminal time T [15,40,41]. We discuss the significance of each penalty term in (17) in detail:

(i) The term $Ap(T) - Br(T)$ at the terminal time is designed to induce a transfer of the system from the malignant into the benign region of the state space. As is seen in Fig. 2, points with small tumor volumes and depressed immunocompetent cell densities lie in the malignant region and it may not suffice to simply minimize the tumor volume. Rather, it is the geometric shape of the stability boundary between the benign and malignant region that matters. Even for a low-dimensional dynamical system like Stepanova’s model, it is generally not possible to give a functional description of the separatrix which is the stable manifold of the unstable equilibrium point of the system. Its tangent space, however, is easily computed, and it thus becomes a good heuristic to choose the coefficients A and B in the penalty term such that the minimization of this term corresponds to moving in the normal direction to this stable manifold from the malignant toward the benign region. Alternatively, one could also choose these coefficients to move into the direction of the unstable eigenvector at the saddle point. For, this is the tangent vector to the path which uncontrolled trajectories will closely follow when in the benign region. In the formulation (17) we have followed the first approach and $\mathbf{v} = (B, A)^T$ is the stable eigenvector of the saddle point oriented so that both A and B are positive numbers (see Fig. 5). It follows from the geometry of the stable manifold that A and B will have the same sign and including in the objective a term of the form $Ap(T) - Br(T)$ gives the correct direction to minimize in the objective. The level sets of this quantity are lines parallel to the tangent space of the stable manifold

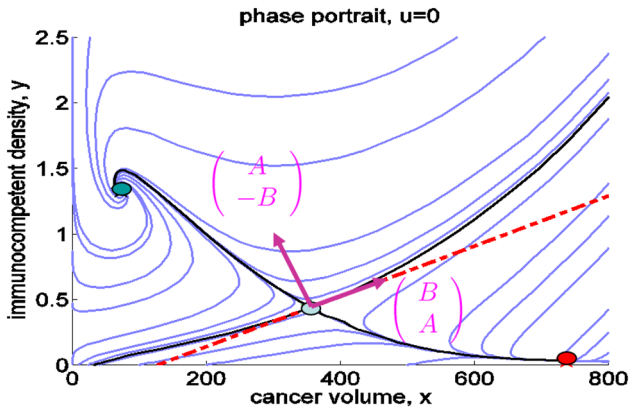


Fig. 5 On the choice of the weights A and B : The figure depicts the phase portrait for the uncontrolled system (5)–(6) with some sample trajectories shown in blue. The dynamics has three equilibria: a locally asymptotically stable focus which is benign (shown as a small green circle), a saddle point (shown as a small grey circle) and a locally asymptotically stable node which is malignant (shown as a small red circle). The stable and unstable manifolds of the saddle point are shown as solid black curves. The tangent vector to the stable manifold of the saddle at the equilibrium point is labeled $v = (B, A)^T$ oriented so that both entries are positive

of the saddle, and minimizing this quantity thus creates an incentive for the system to move into the benign region.

(ii) The integrals of the controls, $\int_0^T u(t)dt$ and $\int_0^T v(t)dt$, measure the total amounts of drugs given and, once again, represent a soft modeling of the side effects of the drugs. Clinical data as to the severity of the drugs should be reflected in the choices for C and D . Naturally, the specific type of tumor and stage of cancer will need to enter into the calibration of these coefficients. In a more advanced stage, higher side effects need to be tolerated and thus smaller values of C would be taken.

(iii) The penalty term ST on the terminal time makes the mathematical problem well-posed. This term, which can be written either under the integral or as a separate penalty term ST , is somewhat unusual. It has to do with the nature of the underlying problem and the emphasis on the regions of attraction. The existence of the asymptotically stable, benign equilibrium point generates controlled trajectories that improve the value $Ap(T) - Br(T)$ of the objective along the trivial controls $u = 0$ and $v = 0$. If no penalty is imposed on the free terminal time, then this creates a “free pass” structure in which the value of the objective can be improved without incurring a cost. As a result, in such a situation an optimal solution may not exist. Intuitively, the controls can switch to $(u, v) = (0, 0)$ immediately as the separatrix is crossed and then take an increasingly longer time as they pass near the saddle point with the infimum arising in the limit $T \rightarrow \infty$. From a practical point of view, this behavior of course is not acceptable. Mathematically, this behavior is easily eliminated by including a penalty term on the final time in the objective. This limits the size of the therapy interval and creates a mathematically well-posed problem.

4 Structure of Optimal Solutions: Bang–Bang (MTD) Versus Singular (BOD) Controls

We discuss general properties of optimal controls for control-affine systems when the L_1 -norm ($k = 1$) is used on the controls and give examples of solutions for some of the problems mentioned earlier.

4.1 Necessary Conditions for Optimality

We consider the optimal control problem to minimize the objective (13) over all admissible controls $u \in \mathcal{U}$ subject to the dynamics (1) and terminal constraints. Necessary conditions for optimality of a controlled trajectory (x_*, u_*) are given by the Pontryagin maximum principle [33] (for some recent references, see [34–36]). These conditions are formulated in terms of the (control) *Hamiltonian* H defined as

$$H = \lambda_0 \left(L(x) + \sum_{j=1}^m \theta_j u_j \right) + \left\langle \lambda, F(x) + \sum_{j=1}^m G_j(x) u_j \right\rangle \tag{18}$$

where $\lambda_0 \geq 0$ and $\lambda \in \mathbb{R}^n$ are multipliers which we write as a row-vector. The most important of the conditions of the maximum principle states that the optimal control u_* for almost every time minimizes the control Hamiltonian H pointwise over the control set U along the optimal controlled trajectory x_* and the multipliers λ_0 and λ , i.e., we have that

$$H(\lambda_0, \lambda(t), x_*(t), u_*(t)) = \min_{v \in U} H(\lambda_0, \lambda(t), x_*(t), v). \tag{19}$$

For the problems considered here, the Hamiltonian H is affine in the controls and for each component u_i of u the control set is a compact interval. Hence, the minimum condition splits into m separate scalar minimization problems that are easily solved. Defining the functions $\Phi_j : [0, T] \rightarrow \mathbb{R}$,

$$t \mapsto \Phi_j(t) = \lambda_0 \theta_j + \langle \lambda(t), G_j(x_*(t)) \rangle = \frac{\partial H}{\partial u_j}(\lambda_0, \lambda(t), x_*(t), u_*(t)), \tag{20}$$

it follows that the optimal controls satisfy

$$u_j^*(t) = \begin{cases} 0 & \text{if } \Phi_j(t) > 0, \\ u_j^{\max} & \text{if } \Phi_j(t) < 0. \end{cases} \tag{21}$$

The function Φ_j is called the *switching function* corresponding to the control u_j .

A priori, the control u_j is not determined by the minimum condition at times when $\Phi_j(\tau) = 0$. In such a case, all control values trivially satisfy the minimum condition and thus, in principle, all are candidates for optimality. However, the switching functions are absolutely continuous and if the derivative $\dot{\Phi}_j(\tau)$ does not vanish, then the

control switches at time τ between $u_j = 0$ to $u_j = u_j^{\max}$ if $\dot{\Phi}_j(\tau) \neq 0$. Such a time τ is called a *bang–bang switch*. On the other hand, if $\Phi_j(t)$ were to vanish identically on an open interval I , then, although the minimization property by itself gives no information about the control, in this case also all the derivatives of $\Phi_j(t)$ must vanish and this condition puts strong limitations on the controls. Extremal controls for which the switching function vanishes identically over an open interval I are called *singular*, while the constant controls $u_j = 0$ and $u_j = u_j^{\max}$ are called *bang* controls; controls that only switch between 0 and the maximum control values are called *bang–bang controls*. We mention that higher-order necessary conditions for optimality of singular controls are available, the so-called generalized Legendre–Clebsch conditions (e.g., see [36]), that allow us to distinguish between minimizing and maximizing singular controls.

If the control represents dose rates for the application of some therapeutic agent, then bang–bang controls correspond to treatment strategies that switch between maximum dose therapy sessions and rest-periods, the typical MTD (maximum tolerated dose) type applications of chemotherapy. Singular controls, on the other hand, correspond to time-varying administrations of the agent at intermediate and often significantly lower dose rates. There is growing interest in such structures in the medical community because of mounting evidence that “more is not necessarily better” [42,43] and that a biologically optimal dose (BOD) with the best overall response should be sought. Hence, the question whether optimal controls are bang–bang or singular has a rather direct interpretation and its answer is of practical relevance for the structure of optimal treatment protocols.

Summarizing, solving an optimal control is tantamount to determining the correct sequences and switching times for the controls as concatenations of bang and singular segments. We close this review with two examples of optimal solutions and interpretations of these results.

4.2 Combination Therapy of Chemotherapy with Antiangiogenic Treatment

We consider the model (2)–(3) for angiogenic signaling with an antiangiogenic agent u and a chemotherapeutic agent v . Modeling the effects of these controls using Skipper’s log-linear model and treating the concentrations of the agents as the controls gives the following controlled dynamics:

$$\dot{p} = -\xi p \ln\left(\frac{p}{q}\right) - \varphi p v, \quad (22)$$

$$\dot{q} = bp - dp^{\frac{2}{3}}q - \mu q - \gamma qu - \eta qv. \quad (23)$$

We are taking the point of view that the available amounts of agents have been determined a priori. We add the trivial dynamics $\dot{y} = u$ and $\dot{z} = v$ which keep track of the amounts of agents used and impose the isoperimetric constraints $\int_0^T u(t)dt \leq y_{\max}$ and $\int_0^T v(t)dt \leq z_{\max}$ at the end-point. The terminal time T is free. Once side effects have been modeled in this way, our objective simply becomes to minimize the tumor

volume, i.e., $J(u) = p(T)$. From a practical point of view, the optimal solution then simply answers the following question: *given a priori determined amounts of agents, how can they best be administered in time to maximize the tumor reduction?*

For the monotherapy problem when $z_{\max} = 0$, i.e., no chemotherapy is given, we have given a complete global solution to this problem in form of a regular synthesis of optimal controls [13,15]. This solution specifies the optimal control as a feedback function of the state (p, q, y) for arbitrary initial conditions. Figure 6 gives a two-dimensional representation of this solution into the (p, q) -space with p shown as the vertical and q as the horizontal variable. For a medically realistic initial condition (with the carrying capacity larger than the tumor volume), optimal controls start with a full-dose segment (shown as solid green curves in the figure) until an optimal singular arc \mathcal{S} is reached. This segment is short in time. Using the conditions of the maximum principle, the curve \mathcal{S} can be described explicitly as [13,15]

$$\mu + dp^{\frac{2}{3}} = bx(1 - \ln x), \quad x = \frac{p}{q}. \tag{24}$$

Then, and for a rather long time, the optimal controlled trajectory follows this singular arc \mathcal{S} (shown as the solid blue curve) until all antiangiogenic agents are used up. Because of after effects there still exists a small interval on which the tumor volume decreases further shown as a dashed green curve. We note that, and different from the impression given in the phase portrait, the times along the individual segments are as described above. The system has a strong differential-algebraic character and as a result the dynamics along the full and no dose segments occurring in the optimal synthesis is very fast generating the almost horizontal lines seen in Fig. 6. Optimal controls are strongly characterized by the singular control which follows a specific time-varying concentration of the agents which maintains an optimal relation between the tumor volume p and its carrying capacity q . It is the functional relation defined by the singular arc which provides ideal conditions for tumor shrinkage.

The optimal solution for the monotherapy problem provides the mathematical foundation on which optimal controls for combination therapies with antiangiogenic treatments can be computed. This holds both for combinations with radiotherapy [44] and chemotherapy [32,45]. Here, we still discuss the latter case. Below we state the optimal concatenation structure for a typical scenario which includes an initial condition when the carrying capacity significantly exceeds the tumor volume, when there is ample supply of antiangiogenic inhibitors (whose side effects generally are low) and a more limited total dose of chemotherapy (which generally has much higher side effects) to be given. In such a situation, optimal controls (u_*, v_*) consist of concatenations of four segments of the following form:

$$(u_*(t), v_*(t)) = \begin{cases} (u_{\max}, 0) & \text{for } 0 \leq t < t_1, \\ (u_{\text{sing}}(p(t), q(t)), 0) & \text{for } t_1 \leq t < \tau, \\ (u_{\text{sing}}(p(t), q(t)), v_{\max}) & \text{for } \tau \leq t < t_3, \\ (0, v_{\max}) & \text{for } t_3 \leq t < T, \end{cases} \tag{25}$$

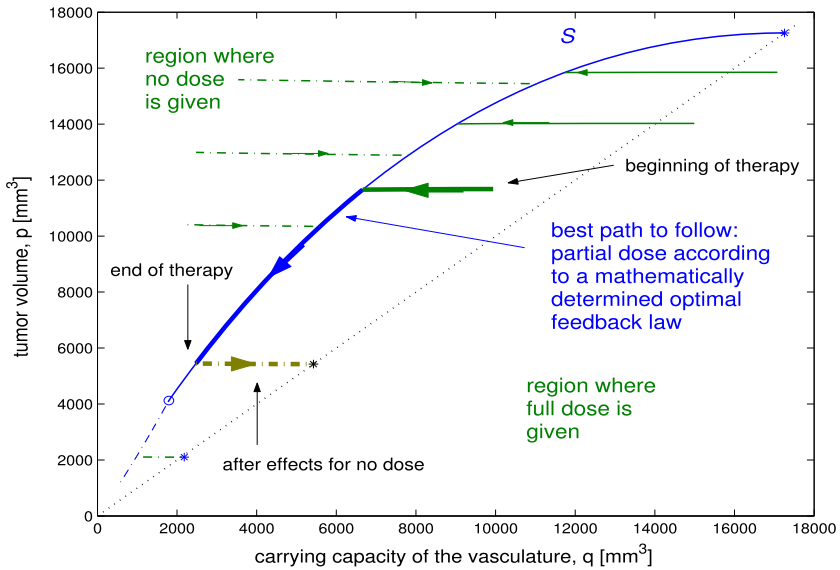


Fig. 6 Sketch of the synthesis of optimal controlled trajectories for the monotherapy problem: the figure shows projections of optimal controlled trajectories into the (p, q) -subspace of the state space with the tumor volume p plotted vertically and the carrying capacity q plotted horizontally. The control is given by full dose, $u = u_{\max}$, along the solid green curves, by no dose, $u = 0$, along the dash-dotted green curves, and by the singular control $u = u_{\text{sing}}$ along the singular arc S shown in blue. The small blue circle along S indicates the point where the singular control saturates, i.e., reaches its maximum value. The portion of the curve S shown dash-dotted is where the singular control is inadmissible

Initially, no chemotherapy is given, while antiangiogenic drugs are given at full dose until the optimal singular arc S for the monotherapy problem is reached. At this point, administration switches to the optimal singular control for the monotherapy problem until a specific time τ is reached. At this time chemotherapy commences in a single full dose (MTD) session. In the medical literature this has been called the ‘optimal therapeutic window’ [30]. Also, because of the onset of chemotherapy, the formula for the optimal singular control adjusts and this is seen as a small kink in Fig. 7 which gives a sample of the time-evolution for optimal controls u_* and v_* . Then, both controls are given until all available drugs are exhausted. In Eq. (25) we have formulated the controls for the case that the antiangiogenic inhibitors run out first, but this depends on the total amounts available and could also be reversed. Figure 8 illustrates the shape of the optimal controlled trajectory in (p, q) -space.

Mathematical optimization of the combined therapy strongly supports the following two qualitative conclusions: (i) there is an optimal relation between the tumor volume p and its carrying capacity q along which the tumor decrease is maximized. This relation is given by the optimal singular arc S of the monotherapy problem. (ii) Along this path there is an optimal point so that chemotherapy is given in one full dose MTD administration as this point is reached. These properties reflect the ideas of ‘pruning’ the vasculature and of a ‘therapeutic window’ [30,31] proposed in the medical literature.




Fig. 7 Illustration of the qualitative shape of optimal controls u_* (left) and v_* (right) for combination treatments with chemotherapy and antiangiogenic inhibitors for medically realistic initial values: after a brief initial period of full dose therapy, administration of the antiangiogenic agent follows the singular control until, at a specific time τ , chemotherapy commences. At that time, the dose of the antiangiogenic agent adjusts to the presence of the chemotherapy. As the dose of the antiangiogenic agent intensifies along the singular arc for the monotherapy problem, in the presence of chemotherapy which also has antiangiogenic effects, the dose increases in a small jump at time τ . The control follows the adjusted singular control until all inhibitors are exhausted. Chemotherapy is given in one full dose segment commencing at time τ

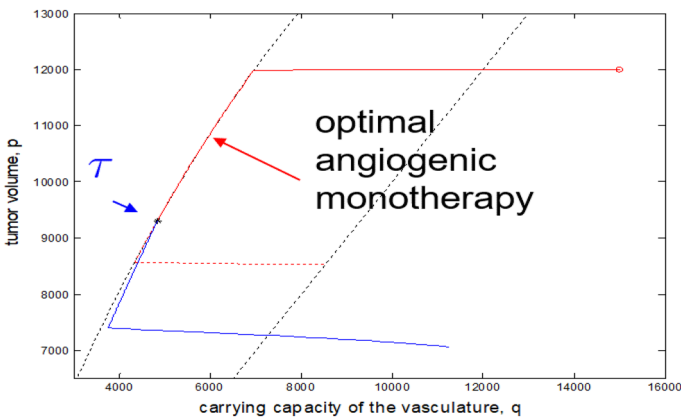


Fig. 8 “Optimal therapeutic window”: the red curve depicts the optimal controlled trajectory for the antiangiogenic monotherapy problem (the control is given by full dose therapy along the vertical portion and then follows the singular arc where the tumor volume decreases). Chemotherapy then commences at an optimal time τ and the optimal trajectory is shown as solid blue curve. (The remaining optimal monotherapy trajectory is shown as dotted red curve.) When all antiangiogenic agents are exhausted, chemotherapy continues until all drugs have been administered. The junction point corresponds to the kink in the blue trajectory. Chemotherapy still lowers the tumor volume, but as there are no more antiangiogenic agents given, the vasculature increases again

We remark that the optimal controls in this model (and this is also the case in many other models where the concentration is taken as the control) are discontinuous. Obviously, this a quirk of the modeling as concentrations are continuous functions. This discrepancy, however, is easily removed by adding a model for the pharmacokinetics of the drug which smoothes out the transitions between the segments while following

the same structure. The simplified modeling with the controls representing concentrations has, from a mathematical point of view, the advantage of a lower dimensional state-space and, more importantly, the analysis of singular controls is simpler (c.f., [46,47]). This makes the overall structure of solutions more tractable. Our paper [48] gives an in depth discussion of these aspects related to the pharmacometrics of the drugs.

4.3 Combination Therapy of Chemotherapy with an Immuno-Stimulatory Agent

We conclude with an example of a numerically computed optimal solution for Stepanova's model with strongly targeted chemotherapy u [15,40,41,49]. Again we treat the concentration of the agent as the control and use Skipper's log-linear model to describe the effect of treatment. This simply adds a term $-\kappa pu$ to the dynamics (5) of the tumor volume. Choosing the objective functional in the form (17) with $N = 0$ (as we do not consider an immune boost), optimal chemotherapy protocols follow the concatenation structure **1s01** with **1** representing a full dose segment, **s** denoting administration following a singular control and **0** standing for a rest-period of the treatment.

Figure 9 gives a characteristic example of a solution: the initial state $z_0 = (p_0, r_0)$ is characterized by a high tumor volume and depressed immuno-competent density close to the values of the malignant equilibrium point. In the figure, the unstable equilibrium point is marked with a green dot and the black curves crossing in it are its stable and unstable manifolds, respectively. Initially, chemotherapy is given at full dose until the state of system has crossed over into the benign region. Then it follows a singular arc where the concentration of the drug is dropped significantly through lower dose administrations. The full singular curve is indicated by a dotted red curve in the figure. Finally, as the state of the system is in the benign region, chemotherapy is stopped and the uncontrolled dynamics (i.e., the actions of the immune system) lead to further reductions in the tumor volume. Interestingly, toward the end, optimal chemotherapy protocols still apply a short chemotherapy segment. This indeed seems to be mirrored by what has proven a particularly effective administration protocol also in medical practice.

If the effect of an immune boost is incorporated into the system, this structure is generally preserved. Figure 10 gives an example of an optimal control for the combination therapy with the chemotherapeutic agent shown in red and the immune boost in green. Essentially, the immune boost is only given toward the end to steer the system into a better terminal state. This figure also clearly shows the lowering of the dose rate for the cytotoxic drug along the singular segment.

5 Brief Comments on the Numerical Methods

It has not been our aim to discuss numerical methods in optimal control in this paper. For completeness sake, however, we include some brief comments on the methods that were used by us and our co-workers to arrive at the solutions shown in this paper.

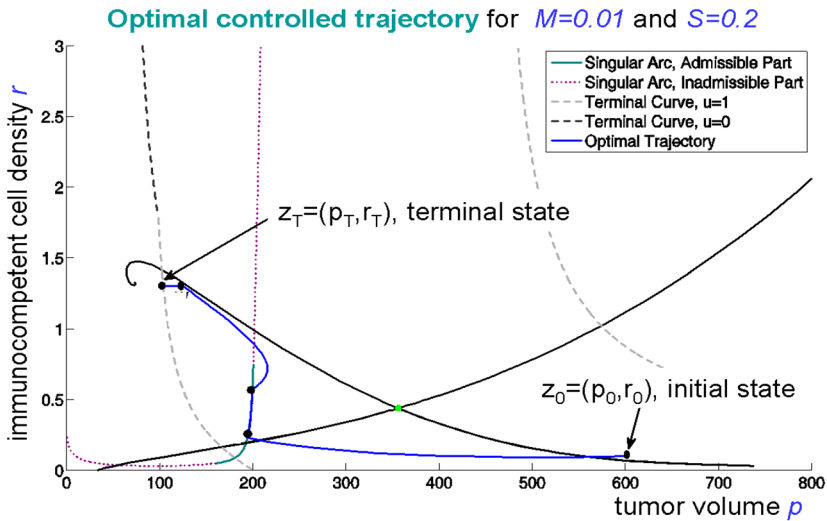


Fig. 9 Example of a numerically computed optimal controlled trajectory (shown as the solid blue curve) for $(M, S) = (0.01, 0.2)$: The initial point, all junctions (where the control switches), and the terminal point of the optimal trajectory are marked by black dots. Starting at $z_0 = (p_0, r_0)$, initially full dose chemotherapy is used. Along that trajectory, the system enters from the malignant region (below the stable manifold of the saddle point) into the benign region (above the stable manifold of the saddle point). The stable and unstable manifolds of the saddle point (marked by a green dot) are the black curves which intersect transversally in the saddle point. The singular arc is the U-shaped curve shown as solid green segments where the singular control is admissible and as dotted red segments where the singular control is inadmissible. Once the optimal controlled trajectory meets the singular arc, the control switches to become singular and the trajectory follows the singular arc (at much reduced dose rates, c.f., Fig. 10). Then, at a certain time, chemotherapy stops and the trajectory follows the dynamics of the uncontrolled system approaching, and ultimately tracing, the unstable manifold of the saddle point. As the state lies in the benign region, this leads to an overall improvement of the objective value. The optimal control then ends with another brief full dose chemotherapy session to reach the optimal terminal point

The problem is simpler if an L_2 -type objective is employed in the control as this makes the Hamiltonian H of the optimal control problem strictly convex and optimal controls are continuous. In this case, standard shooting methods can be applied to the combined system of dynamics and adjoint equations to compute numerical extremals. Their local optimality is easily verified or excluded by computing conjugate points using a second-order approach which tests the existence of a solution of the associated matrix Riccati differential equation as described in [37]. Other approaches are based on discretizations or pseudo-spectral methods (e.g., see [50–52]). For the solutions shown in Sect. 4.3 we used GPOPS (General Pseudo-spectral OPTimal control Software), an open-source MATLAB optimal control software that implements the Gauss hp-adaptive pseudospectral methods (<http://www.gpops.org/>, [53]). These methods approximate the state using a basis of Lagrange polynomials and collocate the dynamics at the Legendre–Gauss nodes [50]. The continuous-time optimal control problem is then transformed into a finite-dimensional nonlinear programming problem that is being solved using well-known and standard algorithms.

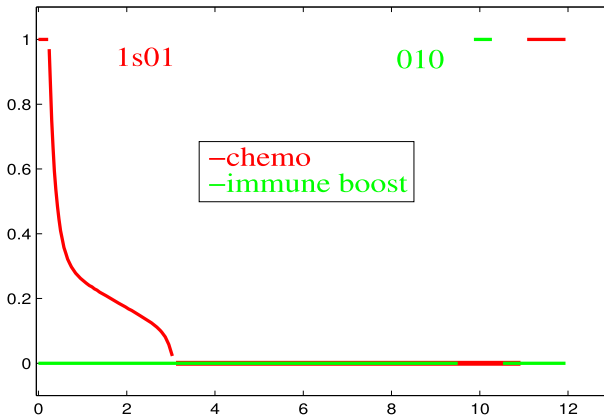


Fig. 10 Example of optimal controls for combination of chemotherapy with an immune boost: This figure gives the time (horizontal axis) evolution of the chemotherapy protocol described in Fig. 9. The maximum chemotherapy dose rate (vertical axis) has been normalized to 1. As can be seen, the time-intervals when chemotherapy is at full dose are brief and the optimal dose rates drop significantly along the much longer singular segment. The rest-period is the most dominant time-period (almost 8 units out of $T = 12$)

Problems with an L_1 -type objective which contain optimal singular arcs are more involved numerically as the singular control is only optimal on an embedded submanifold of positive codimension. Such a structure is inherently difficult to locate numerically. Various algorithms, including commercial programs, we tried were simply unable to deal with this situation. For the antiangiogenic monotherapy problem (see Fig. 6) the complete optimal synthesis was determined by us theoretically by analytical means [13]. The computation of optimal controlled trajectories then amounts to the ‘evaluation of a function’ and is a matter of seconds (c.f., also [54]). For the combination therapies, based on the monotherapy results, the computations were made by Maurer. States, controls and adjoint variables were computed using discretization methods which transcribed the optimal control problem into a large-scale nonlinear programming problem which was implemented in the framework of the Applied Mathematical Modeling Programming Language AMPL and is linked to various optimization codes. Either the interior-point method IPOPT based on [55] or the optimization code NUDOCCS developed in [56] in conjunction with the arc-parametrization method [57] were used. On the other hand, if optimal controls are bang–bang (such as for the cell-cycle specific methods for chemotherapy), the numerical problem is simple and basic algorithms which minimize over the switching times (e.g., gradient procedures [58]) can be used. A verification of second-order conditions for optimality can easily be built into this algorithm [15].

6 Conclusions

In his paper, we formulated an optimal control problem for a general dynamical system which is nonlinear in the state and affine in the controls. The presented framework encompasses a number of models for treatment of biomedical problems, especially

cancer therapies, as well as the more common models in epidemiology. The controls account for either therapeutic or epidemiological interventions. The terms in the cost functional represent various aspects and constraints that need to be taken into consideration when the dynamics is considered as an optimal control problem.

We discussed, and using examples from our own work illustrated how the solutions to these problems share certain common characteristics, what differences can be expected to appear and, most importantly, how the choice of the objective affects the form of the solutions both quantitatively and qualitatively. Specifically, it has been emphasized that omitting some important terms in the objective, like, for example, penalties on the tumor size during therapy—not an uncommon procedure in the literature—may lead to dubious solutions which indicate to wait and administer cancer treatment only toward the end of the therapy period. Clearly, often mathematical features of the solutions (e.g., continuous versus discontinuous protocols for the dosage of a drug) are preprogrammed through the choice of the objective functional. This is acceptable, provided one is aware that this is an outcome of the mathematical formulation and not an indication of properties of the underlying dynamical system or biology.

The purpose of this article was to review and somewhat organize and illustrate work done by us on this topic as well as to point out some common pitfalls and errors which often are made when biological phenomena are put into the mathematical framework of optimal control theory.

Acknowledgements We are grateful to four anonymous referees of this paper for their useful comments which helped us improve our presentation.

Data Availability Statement All data generated or analyzed during this study are included in this published article.

References

1. Eisen, M.: *Mathematical Models in Cell Biology and Cancer Chemotherapy*. Lecture Notes in Biomathematics, vol. 30. Springer, Berlin (1979)
2. Kimmel, M., Swierniak, A.: An optimal control problem related to leukemia chemotherapy. *Sci. Bul. Sil. Tech. Univ.* **65**, 120–130 (1983)
3. Swan, G.W.: *Applications of Optimal Control Theory in Medicine*. Marcel Dekker, New York (1984)
4. Swan, G.W.: General applications of optimal control theory in cancer chemotherapy. *IMA J. Math. Appl. Med. Biol.* **5**, 303–316 (1988)
5. Swan, G.W.: Role of optimal control in cancer chemotherapy. *Math. Biosci.* **101**, 237–284 (1990)
6. Swierniak, A.: Optimal treatment protocols in leukemia—modelling the proliferation cycle. In: *Proceedings of the 12th IMACS World Congress, Paris, vol. 4*, pp. 170–172 (1988)
7. Martin, R., Teo, K.L.: *Optimal Control of Drug Administration in Cancer Chemotherapy*. World Scientific Press, Singapore (1994)
8. Swierniak, A.: Cell cycle as an object of control. *J. Biol. Syst.* **3**, 41–54 (1995)
9. de Pillis, L.G., Radunskaya, A.: A mathematical tumor model with immune resistance and drug therapy: an optimal control approach. *J. Theor. Med.* **3**, 79–100 (2001)
10. Ledzewicz, U., Schättler, H.: Optimal bang-bang controls for a 2-compartment model in cancer chemotherapy. *J. Optim. Theory Appl.* **114**, 609–637 (2002)
11. Ledzewicz, U., Schättler, H.: Analysis of a cell-cycle specific model for cancer chemotherapy. *J. Biol. Syst.* **10**, 183–206 (2002)

12. Swierniak, A., Ledzewicz, U., Schättler, H.: Optimal control for a class of compartmental models in cancer chemotherapy. *Int. J. Appl. Math. Comput. Sci.* **13**, 357–368 (2003)
13. Ledzewicz, U., Schättler, H.: Antiangiogenic therapy in cancer treatment as an optimal control problem. *SIAM J. Control Optim.* **46**(3), 1052–1079 (2007)
14. Fernández, L.A., Pola, C.: Optimal control problems for the Gompertz model under the Norton–Simon hypothesis in chemotherapy. *Discr. Cont. Dyn. Syst. B* **24**(6), 2577–2612 (2019). <https://doi.org/10.3934/dcdsb.2018266>
15. Schättler, H., Ledzewicz, U.: *Optimal Control for Mathematical Models of Cancer Therapies*. Interdisciplinary Applied Mathematics, vol. 42. Springer, New York (2015)
16. Norton, L.: A Gompertzian model of human breast cancer growth. *Cancer Res.* **48**, 7067–7071 (1988)
17. Wheldon, T.E.: *Mathematical Models in Cancer Research*. Hilger Publishing, Boston (1988)
18. d’Onofrio, A.: Fractal growth of tumors and other cellular populations: linking the mechanistic to the phenomenological modeling and vice versa. *Chaos Solitons Fractals* **41**, 875–880 (2009)
19. Hahnfeldt, P., Panigrahy, D., Folkman, J., Hlatky, L.: Tumor development under angiogenic signaling: a dynamical theory of tumor growth, treatment response, and postvascular dormancy. *Cancer Res.* **59**, 4770–4775 (1999)
20. Stepanova, N.V.: Course of the immune reaction during the development of a malignant tumour. *Biophysics* **24**, 917–923 (1980)
21. Moore, H., Ledzewicz, U., Strauss, L.: Optimization of combination therapy for chronic myeloid leukemia with dosing constraints. *J. Math. Biol.* **77**(5), 1533–1561 (2018). <https://doi.org/10.1007/s00285-018-1262-6>
22. Kuznetsov, V.A., Makalkin, I.A., Taylor, M.A., Perelson, A.S.: Nonlinear dynamics of immunogenic tumors: parameter estimation and global bifurcation analysis. *Bull. Math. Biol.* **56**, 295–321 (1994)
23. de Vladar, H.P., González, J.A.: Dynamic response of cancer under the influence of immunological activity and therapy. *J. Theor. Biol.* **227**, 335–348 (2004)
24. d’Onofrio, A., Gandolfi, A.: Tumour eradication by antiangiogenic therapy: analysis and extensions of the model by Hahnfeldt et al. (1999). *Math. Biosci.* **191**, 159–184 (2004)
25. Dunn, G.P., Old, L.J., Schreiber, R.D.: The three E’s of cancer immunoediting. *Annu. Rev. Immunol.* **22**, 322–360 (2004)
26. Chadwick, K.H., Leenhouts, H.P.: *The Molecular Theory of Radiation Biology*. Springer, Berlin (1981)
27. Wein, L.M., Cohen, J.E., Wu, J.T.: Dynamic optimization of a linear-quadratic model with incomplete repair and volume-dependent sensitivity and repopulation. *Int. J. Radiat. Oncol.* **47**(4), 1073–1083 (2000)
28. Skipper, H.E.: On mathematical modeling of critical variables in cancer treatment (goals: better understanding of the past and better planning in the future). *Bul. Math. Biol.* **48**, 253–278 (1986)
29. Shimoda, S., Nishida, K., Sakakida, M., Konno, Y., Ichinose, K., Uehara, M., Nowak, T., Shichiri, M.: Closed-loop subcutaneous insulin infusion algorithm with a short-acting insulin analog for long-term clinical application of a wearable artificial endocrine pancreas. *Front. Med. Biol. Eng.* **8**(3), 197–211 (1997)
30. Jain, R.K.: Normalizing tumor vasculature with antiangiogenic therapy: a new paradigm for combination therapy. *Nat. Med.* **7**, 987–989 (2001)
31. Jain, R.K., Munn, L.L.: Vascular normalization as a rationale for combining chemotherapy with antiangiogenic agents. *Princ. Pract. Oncol.* **21**, 1–7 (2007)
32. d’Onofrio, A., Ledzewicz, U., Maurer, H., Schättler, H.: On optimal delivery of combination therapy for tumors. *Math. Biosci.* **222**, 13–26 (2009). <https://doi.org/10.1016/j.mbs.2009.08.004>
33. Pontryagin, L.S., Boltyanskii, V.G., Gamkrelidze, R.V., Mishchenko, E.F.: *The Mathematical Theory of Optimal Processes*. Macmillan, New York (1964)
34. Bonnard, B., Chyba, M.: *Singular Trajectories and Their Role in Control Theory*. Mathématiques and Applications, vol. 40. Springer, Paris (2003)
35. Bressan, A., Piccoli, B.: *Introduction to the Mathematical Theory of Control*. American Institute of Mathematical Sciences, San Francisco (2007)
36. Schättler, H., Ledzewicz, U.: *Geometric Optimal Control*. Interdisciplinary Applied Mathematics, vol. 38. Springer, New York (2012)
37. Schättler, H., Ledzewicz, U., Maurer, H.: Sufficient conditions for strong local optimality in optimal control problems with L_2 -type objectives and control constraints. *Discr. Cont. Dyn. Syst. B* **19**(8), 2657–2679 (2014). <https://doi.org/10.3934/dcdsb.2014.19.2657>

38. Ferreira, M.M., Ledzewicz, U., do Rosario de Pinho, M., Schättler, H.: A model for cancer chemotherapy with state space constraints. *Nonlinear Anal.* **63**(5), 2591–2602 (2005)
39. Schättler, H.: A local field of extremals for optimal control problems with state constraints of relative degree 1. *J. Dyn. Control Syst.* **12**, 563–599 (2006)
40. Ledzewicz, U., Naghnaeian, M., Schättler, H.: Optimal response to chemotherapy for a mathematical model of tumor-immune dynamics. *J. Math. Biol.* **64**, 557–577 (2012). <https://doi.org/10.1007/s00285-011-0424-6>
41. Ledzewicz, U., Faraji Mosalman, M.S., Schättler, H.: Optimal controls for a mathematical model of tumor-immune interactions under targeted chemotherapy with immune boost. *Discr. Contin. Dyn. Syst. B* **18**, 1031–1051 (2013). <https://doi.org/10.3934/dcdsb.2013.18.1031>
42. Hanahan, D., Bergers, G., Bergsland, E.: Less is more, regularly: metronomic dosing of cytotoxic drugs can target tumor angiogenesis in mice. *J. Clin. Invest.* **105**(8), 1045–1047 (2000)
43. Pasquier, E., Ledzewicz, U.: Perspective on “more is not necessarily better”: metronomic chemotherapy. *Newsl. Soc. Math. Biol.* **26**(2), 9–10 (2013)
44. Ledzewicz, U., Maurer, H., Schättler, H.: Optimal combined radio- and antiangiogenic cancer therapy. *J. Optim. Theory Appl.* **180**, 321–340 (2019). <https://doi.org/10.1007/s10957-018-1426-y>
45. Ledzewicz, U., Maurer, H., Schättler, H.: Optimal and suboptimal protocols for a mathematical model for tumor anti-angiogenesis in combination with chemotherapy. *Math. Biosci. Eng.* **8**(2), 307–323 (2011). <https://doi.org/10.3934/mbe.2011.8.307>
46. Ledzewicz, U., Maurer, H., Schättler, H.: Minimizing tumor volume for a mathematical model of anti-angiogenesis with linear pharmacokinetics. In: Diehl, M., Glineur, F., Jarlebring, E., Michiels, W. (eds.) *Recent Advances in Optimization and Its Applications in Engineering*, pp. 267–276. Springer, Berlin (2010)
47. Ledzewicz, U., Schättler, H.: Singular controls and chattering arcs in optimal control problems arising in biomedicine. *Contr. Cybern.* **38**, 1501–1523 (2009)
48. Ledzewicz, U., Schättler, H.: On the role of pharmacometrics in mathematical models for cancer treatments. *Discr. Contin. Dyn. Syst. B* **1**, 2 (2017). <https://doi.org/10.3934/dcdsb.2020213>. (published online)
49. Ledzewicz, U., Olumoye, O., Schättler, H.: On optimal chemotherapy with a strongly targeted agent for a model of tumor-immune system interactions with generalized logistic growth. *Math. Biosci. Eng.* **10**(3), 787–802 (2012). <https://doi.org/10.3934/mbe.2013.10.787>
50. Benson, D.A., Huntington, G.T., Thorvaldsen, T.P., Rao, A.V.: Direct trajectory optimization and costate estimation via an orthogonal collocation method. *J. Guid. Control Dyn.* **29**(6), 1435–1440 (2006)
51. Zhang, W., Ma, H.: The Chebyshev–Legendre collocation method for a class of optimal control problems. *Int. J. Comput. Math.* **85**(2), 225–240 (2008)
52. Bedrossian, N.S., Bhatt, S., Kang, W., Ross, I.M.: Zero-propellant maneuver guidance. *IEEE Control Syst. Mag.* **29**, 53–73 (2009)
53. Rao, A.V., Benson, D.A., Huntington, G.T., Francolin, C., Darby, C.L., Patterson, M.A.: *User’s Manual for GPOPS: A MATLAB Package for Dynamic Optimization Using the Gauss Pseudospectral Method*. University of Florida report (2008)
54. Ledzewicz, U., Schättler, H.: Optimal and suboptimal protocols for a class of mathematical models of tumor anti-angiogenesis. *J. Theor. Biol.* **252**, 295–312 (2008)
55. Wächter, A., Biegler, L.T.: On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math. Program.* **106**, 25–57 (2006)
56. Büskens, C., Maurer, H.: SQP-methods for solving optimal control problems with control and state constraints: adjoint variables, sensitivity analysis and real-time control. *J. Comput. Appl. Math.* **120**, 85–108 (2000)
57. Maurer, H., Büskens, C., Kim, J.H.R., Kaya, Y.: Optimization methods for the verification of second-order sufficient conditions for bang-bang controls. *Opt. Control Appl. Methods* **26**, 129–156 (2005)
58. Duda, Z.: A gradient method for application of chemotherapy models. *J. Biol. Syst.* **3**, 3–11 (1995)