

Article

Antifragile Control Systems: The Case of An Anti-Symmetric Network Model of the Tumor–Immune–Drug Interactions

Cristian Axenie¹, Daria Kurz^{2,3}, Matteo Saveriano⁴

¹ Audi Konfuzius-Institut Ingolstadt Laboratory, Technische Hochschule Ingolstadt, Ingolstadt, Germany.; cristian.axenie@audi-konfuzius-institut-ingolstadt.de
² Fakultät für Gesundheit, Universität Witten/Herdecke, Witten, Germany.
³ Interdisziplinäres Brustzentrum, Helios Klinikum München West, Munich, Germany; daria.kurz@uni-wh.de
⁴ Department of Industrial Engineering, University of Trento, Trento, Italy.; matteo.saveriano@unitn.it
* Correspondence: cristian.axenie@audi-konfuzius-institut-ingolstadt.de

Abstract: A therapy’s outcome is determined by a tumor’s response to treatment which, in turn, depends on multiple factors such as the severity of the disease and the strength of patient’s immune response. Gold standard cancer therapies are in most cases fragile when sought to break the ties to either tumor kill ratio or patient toxicity. Lately, research has shown that cancer therapy can be at most robust when handling adaptive drug resistance and immune escape patterns developed by evolving tumors. This is due to the stochastic and volatile nature of the interactions, at the tumor environment level, tissue vasculature, and immune landscape, induced by drugs. Herein, we explore the path towards antifragile therapy control, that generates treatment schemes that are not fragile but go beyond robustness. More precisely, we describe a first instantiation of a control-theoretic method to make therapy schemes cope with the systemic variability in the tumor–immune–drug interactions and gain more tumor kill with less patient toxicity. Considering the anti-symmetric interactions within a model of the tumor–immune–drug network, we introduce the antifragile control framework that demonstrates promising results in simulation. We evaluate our control strategy against state-of-the-art therapy schemes on various experiments and discuss the insights we gained on the potential that antifragile control could have in treatment design in clinical settings.

Keywords: Computational Oncology; Cancer; Antifragility; Control Theory

1. Introduction

Improving therapy outcome in cancer care is still the crux of modern oncology. With cancers which become more robust and resistant to conventional schemes, therapy design undergoes a revolution. Tumors where metronomic therapy paradigms were the gold standard are now treated by sophisticated adaptive control algorithms and machine learning techniques capable to capture their growth and resistance patterns. The goal of such therapies is to induce controlled changes in the tumor environment through the timing and magnitude of the drug administration. Yet, such an intervention determines dramatic perturbations in tumor’s evolution and environment, but also in the structure of the surrounding normal cells vasculature, and, of course, the immune system response [1]. Given the stochastic and complex nature of such processes, therapy schemes tend to fail in the face of the systemic variability and volatility of the effects they have upon the tumor. In other words, they are *fragile*. It is typically desired that such therapies would follow precise and fixed schedules designed by clinicians to control both the administration time and the dose magnitude to obtain a targeted response with minimal toxicity. However, the stochastic and patient-specific nature of the tumour-immune-drug dynamics impedes the development of suitable patient-centered

interventions. There are already efforts combining mechanistic modelling and control algorithms targeted at enhancing the understanding and control of cancer dynamics by exploiting heterogeneous clinical data [2]. Such closed-loop systems can capture [3] and control [4] the underlying dynamics of the physical interactions among the three main ingredients: the tumor, the immune system, and the chemotherapeutic drug and make treatments robust. In the recent seminal work of West et al. [5], *antifragile* therapy was introduced as a new framework to quantify the behavior of tumors in response to treatment variability, and how antifragile therapy can gain from such variability. The fundamental idea, promoted by this excellent work, is to identify and quantify regimes of fragility and antifragility through the curvature of the drug dose–response of the tumor. In other words, analysing the second-order behavior of the tumor dynamics in the presence of the drug, they were able to inform the optimality of a treatment schedule. Such treatment optimality included switching from even schedules in fragile tumor regions to uneven schedules in antifragile tumor regions. Despite the limited but promising results on lung cancer cell lines and 11 drugs, the framework has shed light on how to simultaneously capture the evolution of resistance patterns, the benefits of antifragile dosing, and their temporal evolution. Motivated by the great preliminary results of the antifragile therapy, we propose the development of the adaptive antifragile therapy, which taps into treatments adjustments to control tumor growth and manage toxicity based on tumor response that are clinically feasible and beneficial. While the approach of West et al. [5] allows for a straightforward prediction of optimal dose treatments for a large range of treatment schedules that amount to the same cumulative drug dose, our work will go beyond and enable a control-theoretic framework that can validate the link between the shape of the dose response curve and the treatment scheduling. More precisely, we seek to build a framework for treatment dosing and scheduling that can push the dynamics of the tumor–immune–drug network towards a patient-specific shape of dose response.

1.1. The need for control in cancer therapy

Cancer therapies follow national or international guidelines that are typically outlining the recommendations from the latest and most successful research studies on the cumulative dosing, therapy staging, and therapy combinations. However, the therapy parameters, such as dose frequency and dose magnitude, are typically based on the prescribing physician’s assessment of the specific tumor characteristics (i.e., BRCA1 and BRCA2 gene mutations in breast cancer for instance) and behavior (i.e., multistage carcinogenesis, from tumor genesis to metastasis).

The motivation to use control theory in therapy design stems from a couple of phenomena that describe the intricate dynamics of the tumor–immune–drug network, such as periodicity, resonance, and synchronization. For instance, cell growth is maximal when the periodicity of drug administration is a multiple of the characteristic periodicity of the cell population [6] and more visible in cell-cycle specific drugs [7]. In this case, tumor growth is maximised when the periodicity of environmental perturbations acting upon the tumor is a multiple of the periodicity of the tumor population dynamics [8], and hence beneficial to exploit in drug administration. Nevertheless, such temporal phenomena can become more complex when the dynamics of the drug and the tumor specific to a patient come into play. For instance, the work in [9] introduced a control theoretic approach based on resonance–antiresonance drug pulsing combined with vascularised tumor growth models for drug scheduling optimization and “drug dense paradigm” (in the Norton-Simon sense) that depends on patient’s cytokinetic and angiogenic parameters. Phenomena like the resonance can be used for timing the inter-dose interval in order to maximize the efficacy/toxicity ratio and minimize the toxicity/efficacy ratio (the dual problem).

Finally, when considering population-level interactions, covering both between-population competition (e.g., tumor–immune) and within-population cooperation (i.e.

tumor–tumor), therapy design can exploit various phenomena which span across spatial scales. For instance, as cell population persistence depends on the local level of synchronization between the environmental and population processes (i.e., interaction among extra cellular matrix and tumor) as shown in [10], the therapy could in principle induce resonance through optimal drug dosing interval that maximizes host drug susceptible cells growth and minimizes malignant cells growth.

Optimizing chemotherapy strategies in order to maximize the efficacy/toxicity ratio can be done by exploiting the nonlinear tumor response dynamics and an adaptive control of dosage under patient specific tumor–immune–drug interactions. Yet, the complex dynamics describing patient-level tumor–immune–drug interactions are described by variability, volatility, and randomness. Control theory provides a systematic apparatus to tackle such a landscape, in order to ensure not just tolerating systemic variability but in fact gaining from it. Such a behavior is captured by the recently developed paradigm of antifragile therapy [5]. Instead of focusing on antifragility theory and heuristics to inform treatment scheduling or resistance management plans, we consider a systematic approach to build chemotherapy control systems that are antifragile.

This first computational oncology study is laying the basis for a novel framework to designing control systems that can operate in the presence of the randomness, variability, and volatility of the patient-specific tumour–immune–drug dynamics.

2. Materials and Methods

In this section, we first introduce the tumor–immune–drug interaction model based on the network design of [11], which is further refined with the parametrization suggested in [12] and [13]. This model will be the subject for the control theoretic implementation of therapies based on various state-of-the-art approaches. Afterwards, we formally introduce the antifragile control framework with its mathematical apparatus and controller synthesis details. Finally, we introduce the formalism and synthesis of typically used robust, optimal, and pulsed controllers for cancer therapy.

2.1. Tumor–Immune–Drug Network Model

In order to analyze and evaluate the antifragile control framework, we employ a network model that captures the tumor–immune–drug interactions from [11]. The original model consists of three populations of cells (i.e., tumor population, immune cells population, and normal cells population) which evolve through the mutual connections among themselves and under the presence/absence of the chemotherapeutic drug. The network model used in this study is depicted in Figure 1.

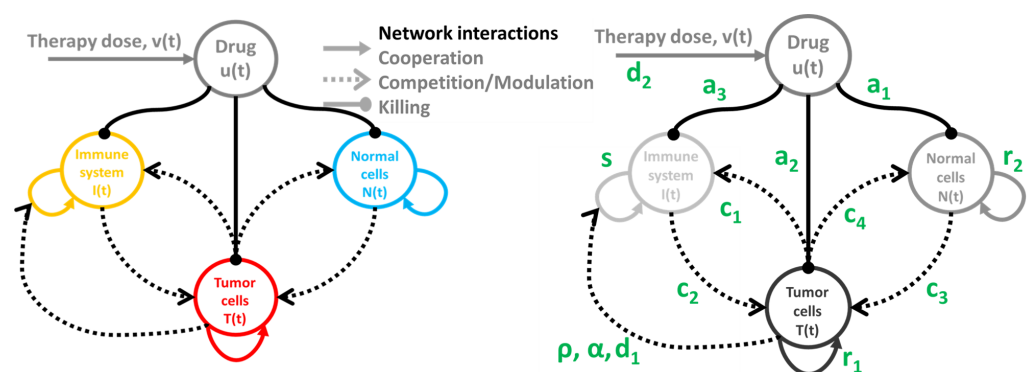


Figure 1. Network model of tumor–immune–drug interactions. Left panel: populations cells interactions; Right panel: differential equations parameters mapped on the network populations interactions.

The ordinary differential equations (ODE) describing the dynamics of the tumor–immune–drug network are introduced in Equations 1–4. The network interaction model captures competition (i.e., c_1, c_2, c_3, c_4 parameters) and modulation among tumor and

immune and normal cells (i.e., ρ, α, d_1 parameters, as well as the cooperation within each of the three populations (i.e., s, r_1, r_2 parameters), accounting for a basic form of self-excitation. The therapeutic drug acts upon all three cells populations in the network model by killing cells with a certain rate (i.e., a_1, a_2, a_3 parameters), as shown in the mapping in Figure 1-right panel. As visible in the network model, the presence of tumor cells stimulates the immune response, represented by the positive nonlinear growth term $\frac{\rho I(t)T(t)}{\alpha + T(t)}$ within the immune evolution in Equations 1–4.

$$\frac{dN(t)}{dt} = r_2 N(t)(1 - b_2 N(t)) - c_4 T(t)N(t) - a_1(1 - e^{u(t)})N(t) \quad (1)$$

$$\frac{dT(t)}{dt} = r_1 T(t)(1 - b_1 T(t)) - c_2 T(t)I(t) - c_3 T(t)N(t) - a_2(1 - e^{u(t)})T(t) \quad (2)$$

$$\frac{dI(t)}{dt} = s + \frac{\rho I(t)T(t)}{\alpha + T(t)} - c_1 I(t)T(t) - d_1 I(t) - a_3(1 - e^{u(t)})I(t) \quad (3)$$

$$\frac{du(t)}{dt} = v(t) - d_2 u(t) \quad (4)$$

For the readers interested to reproduce the analysis and experiments presented in the manuscript, we provide, in Table 1 in Section 3, the network model parametrization along with explanations on how each parameter impacts the dynamics of the overall model (as well as a link to the programs code). Recall that in our experiments the units of cells were rescaled. Hence, in our analysis one unit represents the carrying capacity of the normal cells in the region of the tumor. Although this depends on the type of tumor, we consider the reasonable scale on the order of 10^{11} cells.

2.1.1. Drug-free evolution

This section presents the drug-free evolution of the model in order to facilitate the analysis and the control design in the upcoming sections. In order to analyze the basic dynamics of the tumor-immune-drug network model, we consider the cumulative drug dose $u(t) = v(t) = 0$ when simulating the ODEs of the system. We start with a relatively large tumor burden $T(0) = 0.25$ and an analysis period of $t = 150$ days (i.e., a typical duration of a chemotherapy schema). The initial tumor size corresponds to a tumor with approximately 0.20×10^{11} cells, or, in other words, a solid tumor of radius between 1.8 and 3.9 cm. To provide the reader some reference, the clinical detection threshold for a solid tumor is typically around 10^7 cells, this sets our chosen initial tumor volume of 0.20×10^{11} above clinical detection level. We also consider a patient with a weak immune system (but above the level of an immune-compromised patient) $I(0) = 0.1$. The immune threshold rate α is chosen inversely related to the immune response curve such that when the number of tumor cells T is equal to α the immune response rate is at half of its maximum. In order to guide our control-theoretic design of therapy, we perform the null-space analysis of the drug-free dynamics. The evolution of each of the three interacting cell populations is depicted in Figure 2.

As we see in Figure 2, once the tumor has a considerable size (i.e., $T(0) = 0.20$), the immune system reacts with a very steep rise time in the first 10 days. The immune response is then modulated by the tumor proliferation and the decay in normal cells production which we see already after 50 days. This behavior is also visible in the configuration of equilibrium points in the dynamic landscape of the system (see Figure 3 for the null-surfaces of each three cell types). Here, from the dynamical system analysis perspective, without any drug input, that is $u = 0$, the system must be either in the basin of a stable tumor-free equilibrium or in the basin of a stable equilibrium at which only a harmlessly small amount of tumor is present.

What we observe in the null-space analysis is that our network model can capture the multifaceted dynamics of the tumor-immune interactions. For instance, the drug-free setup allowed us to identify: the healthy tumor-free equilibrium \mathbf{H} , the unstable




Figure 2. Drug-free evolution of the network model. The immune system (I) reacts by proliferating T-cells when the tumor (T) population (and tumor itself) reaches a threshold triggering immune-surveillance. The normal cells (N) production decreases as the tumor expands. In this case, without any drug administration, the tumor proliferate uncontrollably escaping immune surveillance.

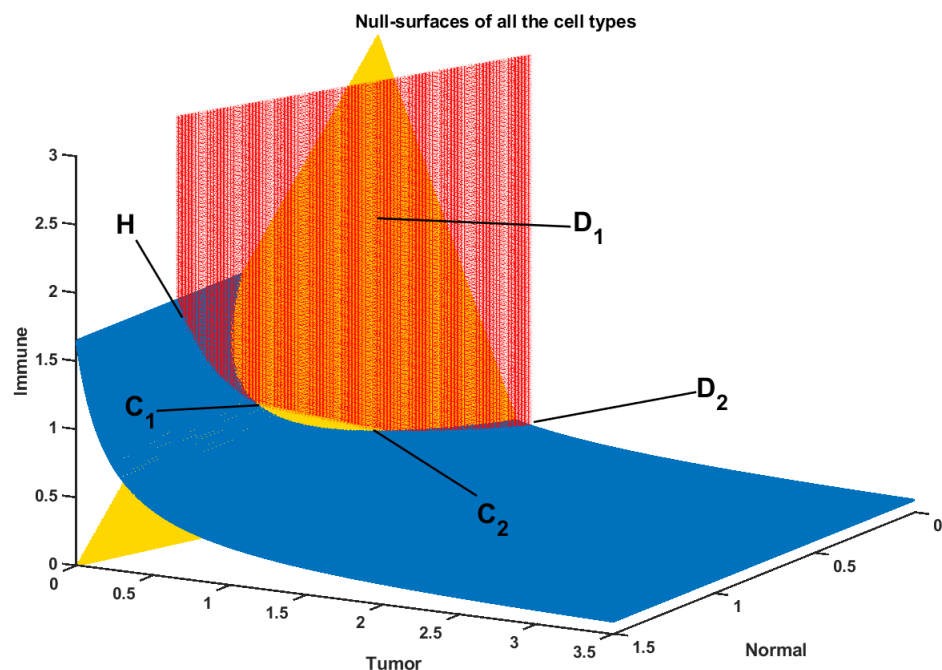


Figure 3. Dynamical null-space analysis of the drug-free model and the prerequisites for therapy control design. We are interested in having two attractors, one for “healthy” state (H) and one for “disease” (C2) so that the control can push the dynamics to reduce tumor size under drug intake.

interior equilibrium C_1 , the “unhealthy” interior stable equilibrium C_2 , and two unstable “dead” equilibria where there are no normal cells D_1 and D_2 , as shown in Figure 3. This description of null-surfaces of the vector fields is useful in the characterization of the long term behavior of all system’s orbits. In our case, in order to consider the patient “cured”, the system must be either in the basin of a stable “healthy” equilibrium or in the basin of a “unhealthy” stable equilibrium at which only a harmlessly small amount of tumor is present.

We further analyzed the model equilibria as a function of the immune cell influx without tumor s and the immune response ρ —see Figure 4. This analysis is very important for the control design as it provides a sufficient characterization of the long term behavior of all system orbits given the competitive, cooperative, and modulation interactions we consider in the model. Additionally, if the “healthy” tumor-free equilibrium is unstable then, according to the dynamics of the model in Equations 1–4, no amount of therapeutic drug will be able to completely eradicate the tumor.

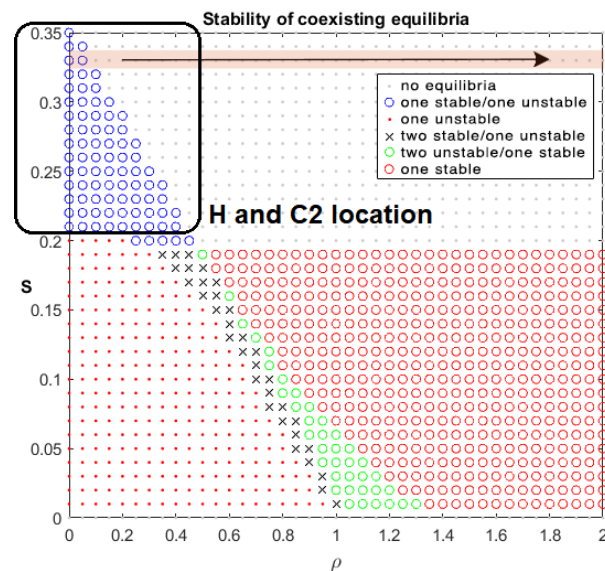


Figure 4. Model equilibria as a function of the immune cell influx without tumor s and the immune response ρ . The control design will revolve around the marked region in the landscape of coexisting equilibria with two stable attractors, one for “healthy” state (H) and one for “disease” (C2).

Our findings are consistent with the experiments and analysis in [12] especially when assessing the reasonable choice of the two stable attracting states, one of which is “healthy” (i.e., the H point in Figure 4) and another which is “diseased” (i.e., the C₂ point in Figure 4). This is an intuitive design as a system with only a healthy attractor state would not need therapy and, conversely, a system with no healthy attractor state will never be cured, so no disease remission would be observed. We will see in the following sections how the basin of attraction of H will be affected by adding the drug in the network model and, later, the control theoretic calculation of the dosing $u(t)$.

We offer the reader the possibility to reproduce the analysis and reuse the models (i.e., tumor-immune-drug and control theoretic models) from the experiments performed in this study by downloading the code available on GitLab¹.

2.2. Antifragile Control

In this section, we describe the mathematical apparatus of antifragile control, going from the theory and principles to control synthesis.

2.2.1. Preliminaries

In this subsection, we offer a short primer on the **principles, theory, and design of antifragile control systems**, with a focus on the cancer therapy problem. As coined in the book of Taleb [14], **antifragility** is a property of a system to **gain from uncertainty, randomness, and volatility**, opposite to what fragility would incur. An **antifragile system’s response** to external perturbations is **beyond robust**, such that small stres-

¹ Code available at: <https://gitlab.com/akii-microlab/antifragile-therapy-ctrl>

sors can strengthen the future response of the system by adding a strong **anticipation component**.

When looking at control systems, inducing such a behavior to a feedback control loop accounts for a **novel design and synthesis** approach where: 1) **redundant over-compensation** can bring the system in an overshooting mode that builds extra-capacity and strength in anticipation; 2) **structure-variability** can induce stressors and carry inherent information that emerges only under volatility and randomness of the system dynamics altered by the application of a high-frequency switching control signal; and 3) **high-frequency control activity** opposite to a “tight control” which only eliminates the beneficial effect of noise and volatility of the closed-loop dynamics. We hereby guide the reader on how the above concepts relate to the therapy control problem and their practical instantiation.

Antifragile control is designed to capture and exploit the relations between the nonlinearity of drug dose–response and the properties of the outcomes (i.e., tumor kill ratio, toxicity ratio and their average and variations) under the complex tumor–immune–drug dynamics. But how can one quantify antifragility? In his work, Taleb [15] provided a clear mathematical and, rather geometrical, description: “**antifragility**” is described as the mathematical property of a **local convex response** and its generalization, whereas the designation “**fragility**” is its opposite, the **locally concave response**.

This perspective was further refined by West et al. in [5] where the drug dose–response shape was used to denote a fragile or antifragile therapy scheme, respectively. More precisely, considering a typical drug dose–response, namely the sigmoidal-shaped Hill function [16], indicating the percent of cells which survive a given dose, the tumor’s response to treatment is antifragile if the **curvature is negative**, otherwise is fragile. In short, one can consider that **antifragility is a second order effect**, whereas, for instance the average is the first order effect.

Provided this intuition of how antifragile control can be designed, we elaborate, in the following subsections, the formalism and the controller synthesis for therapy control.

2.2.2. Formalism

As shown in the previous section, the shape of the drug dose–response is a useful heuristic to detect antifragility and, implicitly, inform tumor evolution and optimal dosing plans within a therapy schedule. Yet, assessing the drug dose–response in complex tumor–immune–drug models, such as the one we consider in Equations 1– 4, is highly dependent on the cell populations dynamics and interactions. For instance, let’s consider the tumor cells. In solid cancers, they form heterogeneous populations with varying drug sensitivities that depend on multiple factors such as the cell cycle stage [17], the presence of geno- or phenotype features [18], and the environmentally mediated drug resistance [19].

Under the antifragile control of drug dose, the tumor–immune–drug model is a dynamical system whose state space is not a vector space but rather a curved state manifold (i.e., the drug dose–response surface), more precisely a place and a dynamics to push the system towards [20]. To get an intuition, revisit the attractor states analysis in Figure 3, where the intersecting planes of the tumor, immune, and normal cell populations evolution describe the geometrical properties of the model dynamics. We aim at designing a controller that forces the dynamics to evolve on a prescribed state manifold with certain geometrical properties. Control systems on state manifolds is a relatively new research branch in non-linear control theory, with a very well formalized mathematical apparatus cast in the Riemannian geometry [21,22].

For our therapy control problem, we want to push the tumor–immune–drug system to the desired dynamics of the drug dose–response surface, turning this in a tracking control problem. Recall that we want to have two attractor states lying on manifolds whose curvature (i.e., second derivative) is informative on the joint populations’ dynamics. More precisely, the curvature of the imposed dynamics surface can be assessed on

Riemannian manifolds by using, for instance, the Penot generalized directional derivative and the Clarke generalized gradient [23]. This enables use to formulate the **main pillars of antifragile control: redundant overcompensation, structure variability, and bounded high-frequency control activity**, in terms of Riemannian geometric control [24]. For instance, having the advantage that translation in Euclidian space is a canonical transport with acceleration and velocity in Riemannian space, we can apply the exponential map to perform updates along the shortest path in the relevant direction in unit time while remaining on the desired drug dose–response manifold region [25].

Dynamical systems on manifolds

In the following, we recall notions from manifold calculus and illustrate how these tools prove useful in describing the system-theoretic design of an antifragile controller. We describe **redundant overcompensation, structure variability, and bounded high-frequency control activity** using **manifold calculus** which: 1) carries the advantage that is coordinate-free; 2) relies on the embedding of a manifold into a larger dynamical space; 3) supports simpler definitions of control of non-linear systems whose states belong to curved manifolds [26].

Let Q denote a Riemannian manifold. At each point $x \in Q$ is possible to define the tangent space $T_x Q$, i.e., a real vector space containing all possible directions tangentially passing through x . The disjoint union of the tangent spaces of Q forms the tangent bundle, formally defined as $TQ := \{(x, v) | x \in Q, v \in T_x Q\}$. The Riemannian manifold Q is equipped with a bilinear, positive-definite inner product operator $\langle \cdot, \cdot \rangle_x : T_x Q \times T_x Q \rightarrow \mathbb{R}$. This also defines a local norm $\|x\|_x := \sqrt{\langle v, v \rangle_x}, v \in T_x Q$. The Riemannian gradient of a function $f : \mathbb{R} \rightarrow Q$ at a point $x \in Q$ is $\frac{df}{dx}$. The exponential map exp of a manifold Q is a function $exp : TQ \rightarrow Q$ which takes a point-velocity pair $(x, v) \in TQ$ and maps it to a point on the manifold Q , namely $exp_x(v) \in Q$. The inverse, the manifold logarithm log is defined only locally and returns a vector $v = log_x(y) \in T_x Q$ such that $exp_x(v) = y$, given $x, y \in Q$. In order to compute a distance $d(x, y)$ on the manifold Q , we define a geodesic arc as the manifold-theoretic counterpart of a straight line on a curved manifold. On a Riemannian manifold, the distance between two points may be calculated as $d(x, y) = |log_x(y)|_x$ and implicitly the gradient of the squared distance as $\frac{d^2(x, y)}{dx} = -2log_x(y)$, given, of course, that the logarithm is defined. Finally, we assume the manifold Q to be equipped with a metric connection (which makes it easier to compute the covariant derivative involving simply an inner product). The covariant derivation operator $\nabla_v w$ of a vector field $w_x \in T_x Q, x \in Q$ in the direction of a vector $v \in T_x Q$ is closely related to the transport map $T_{x, y}$ as one can express the derivation in terms of parallel transport as

$$\nabla_v w \approx \frac{T_{(\gamma(h) \rightarrow x)}(w_{\gamma(h)}) - w_x}{h} \quad (5)$$

$$\gamma(0) = x \in Q, \dot{\gamma}(0) = v \in T_x Q \quad (6)$$

$$h \ll 1 \quad (7)$$

In this study, the purpose of the antifragile controller is to drive the state $q(t) = [I(t), T(t), N(t), u(t)]^T$ of the tumor–immune–drug system toward a reference point on the desired/reference configuration $r(t)$ —capturing the dynamics of the system in the antifragile region of the drug dose–response surface (see West et al. [5]). The first step in the controller design is the definition of a control error function, quantifying how far is the system state from the reference.

Control error function

The control error function, $\varphi : Q \rightarrow \mathbb{R}$, describes the distance between the reference state (i.e., the tumor–immune–drug system dynamics in antifragile region on the desired drug dose–response shape) denoted by r and the actual configuration/state q on a state

manifold Q . In our case, the reference dynamics is determined by the position of the tumor-immune-drug system trajectory on the drug dose-response curve that describes the survival rate S vs. drug dose u . Here, a lower survival is better as it accounts to more tumor kill and reaching the antifragile control region. In our case, we consider φ a uniformly quadratic function with constant L . This implies that, for all $\epsilon > 0$, there exist $b_1 \geq b_2 > 0$ such that

$$b_2 \left\| \frac{d\varphi(q, r)}{dq} \right\|^2 \leq \varphi(q, r) \leq b_1 \left\| \frac{d\varphi(q, r)}{dq} \right\|^2 \quad (8)$$

278 In our case, we consider the configuration $q(t) = [I(t), T(t), N(t), u(t)]^T$ and the reference $r(t) = [I_S(t), T_S(t), N_S(t), u_S(t)]^T$. The control error φ is the most important
279 component in ensuring the **redundant overcompensation** through its temporal evolution $\varphi(t)$ and its velocity $\dot{\varphi}(t)$.
280
281

282 **Parallel transport map**

Given the configuration $q(t)$ and the reference $r(t)$, as defined above, a linear map $P_{(q,r)} : T_r Q \rightarrow T_q Q$ is a parallel transport map if it is compatible with the error function φ , that is

$$\frac{d\varphi(q, r)}{dr} = -P_{(q,r)} \frac{d\varphi(q, r)}{dq} \quad (9)$$

Intuitively, the parallel transport map describes the actual dynamics transformation that pushes the system towards a desired manifold along a geodesic. Now, given the transport map P , the velocities (i.e., first derivatives) of both the configuration q , marked as \dot{q} and reference r , marked as \dot{r} , can be compared to build a velocity error, i.e.,

$$\dot{e} = \dot{q} - P_{(q,r)} \dot{r} \in T_q Q \quad (10)$$

Given the previous notations, we can now compute the time derivative of the error function φ such as

$$\frac{d}{dt} \varphi(q(t), r(t)) = \frac{d\varphi(q(t), r(t))}{dq(t)} \dot{e}(t) \quad (11)$$

283 These geometrical objects serve as tools to introduce the system-theoretic synthesis of
284 the antifragile controller in the following section.

285 2.2.3. Control synthesis

286 In the current section, we formally introduce the three key components for antifragile
287 control synthesis, namely redundant overcompensation, structure variability, and
288 bounded high-frequency control activity.

289 **Redundant overcompensation**

We start the synthesis of the redundant overcompensation component which will induce the anticipation capabilities of the antifragile controller. From a control theoretic point of view, this accounts for a combined Proportional Derivative (PD) control action. Let's start with the more general form of our tumor-immune-drug model as a dynamical system in the Riemannian manifolds framework such as

$$\dot{q}(t) = v(t), t \geq 0 \quad (12)$$

$$\nabla_{\dot{q}} \dot{q} = \mathbb{S}(t, q(t), v(t)) + F(q(t)) \quad (13)$$

290 where q is the state, \dot{q} is the velocity of the state, \mathbb{S} is the time varying-state-transition
291 map $\mathbb{S} : \mathbb{R} \times TQ \rightarrow TQ$, ∇ is the Riemannian covariant derivative, and F the control law.

As in “traditional” PD control design, we need the error function between the state q (i.e. configuration) and the reference r , as well as the anticipation component. We consider the previously introduced error function φ on the manifold Q . The **antifragile proportional gain** K_φ is a smooth self-adjoint positive tensor filed on the manifold Q and $K_q : T_q Q \rightarrow T_q^* Q$ is the **antifragile derivative gain**. Putting all together, the antifragile PD control law for therapy design is computed as

$$F(q) = K_\varphi \varphi(q) + K_q \dot{\varphi}(q) \quad (14)$$

where $q(t) = [I(t), T(t), N(t), u(t)]^T$. This law is consistent with similar approaches for simple PD control on manifolds [27]. As a trademark of antifragile control, the first term in the Equation 14, $K_\varphi \dot{\varphi}(q)$, **seeks to anticipate** (not to predict) a higher level of error than the previous maximum through a **redundant overcompensation** that builds extra-capacity through the choice of the transport map that determines the term $K_q \dot{\varphi}(q)$. This component is responsible to drive the system closer the the desired dynamics, so we now evaluate the (Lyapunov) stability of the closed loop given the control law $F(q)$. We start by selecting a Lyapunov candidate function

$$V(\varphi, q) = \varphi + \frac{1}{2} \|\dot{q}\|^2 \quad (15)$$

that is positive definite in q and \dot{q} . Then, the derivative of V is calculated as

$$\frac{dV(\varphi, q)}{dt} = \frac{d}{dt} \left(\varphi + \frac{1}{2} \|\dot{q}\|^2 \right) \quad (16)$$

$$= \nabla_{\dot{q}} \varphi + \frac{1}{2} \nabla_{\dot{q}} \|\dot{q}\|^2 \quad (17)$$

$$= \langle d\varphi, \dot{q} \rangle + \langle \nabla_{\dot{q}} \dot{q}, \dot{q} \rangle \quad (18)$$

$$= \langle d\varphi, \dot{q} \rangle + \langle -K_\varphi \varphi(q) - K_q \dot{\varphi}(q), \dot{q} \rangle \quad (19)$$

$$= -\langle K_q \dot{\varphi}(q), \dot{q} \rangle \quad (20)$$

292 Then, given that $V(\varphi, q)$ is positive definite and $\frac{dV(\varphi, q)}{dt}$ is negative semi-definite, the
 293 closed-loop system is stable in the Lyapunov sense. For the asymptotic stability one can
 294 use the LaSalle invariance principle.

295 Stucture variability

296 While there are many advanced approaches, such as adaptation based on response
 297 identification and state observation or absolute stability techniques, the most straight-
 298 forward way to deal with uncertainty is to keep certain limitations by “brute force”.
 299 However, any carefully maintained equality eliminates one “uncertainty dimension”.
 300 The theory of **variable structure control** (VSC) developed around this principle [28]
 301 and opened up a wide new area of development known as **sliding mode control** (SMC),
 302 that is characterized by a discontinuous control action which changes structure upon
 303 reaching a set of predetermined switching surfaces [29,30]. The main advantages of
 304 using SMC in our therapy control design are listed below:

- 305 • the motion equation of the sliding mode [31] can be designed linear and homo-
 306 geneous, despite that the tumor-immune-drug model is governed by nonlinear
 307 equations,
- 308 • the sliding surface does not depend on the process dynamics, but it is determined
 309 by parameters selected by the designer [32], i.e., desired trajectory of the system in
 310 the antifragile region of the drug dose-response curve (e.g., Hill function),
- 311 • once the sliding motion occurs (i.e., the system dynamics are on the surface), the
 312 system has invariant properties which make the motion independent of certain
 313 system parameter variations, uncertainty, and disturbances [33]. Hence, the system
 314 performance can be completely determined by the dynamics of the sliding manifold.

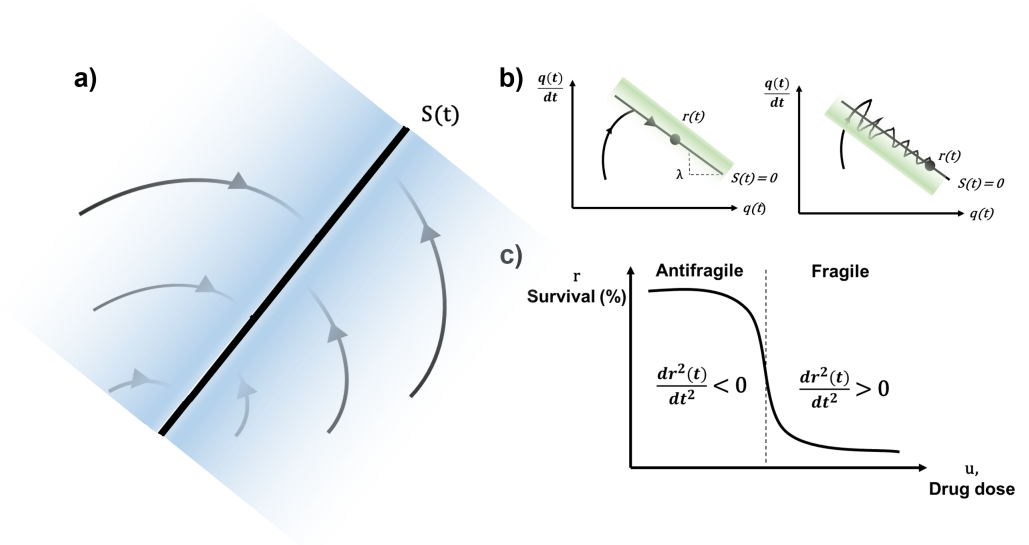


Figure 5. Variable structure control using SMC. a) the sliding condition, from initial states, the system is pushed towards the surface S ; b) the dynamics on the sliding surface; c) drug dose–response curve with fragile and antifragile regimes depending on curvature (see West et al. [5]).

The first design element in the variable structure SMC is the choice of the sliding surface. Starting from the control error function φ , we define S as a time varying surface in the state space / configuration space q , where

$$S(q, t) = \left(\frac{d}{dt} + \lambda \right)^{n-1} \varphi, \quad \lambda > 0. \quad (21)$$

For instance, if $n = 2$ then $S(q, t) = \dot{\varphi} + \lambda\varphi$, if $n = 3$ then $S(q, t) = \ddot{\varphi} + 2\lambda\dot{\varphi} + \lambda^2\varphi$, and so on. As we see in Figure 5 b), the sliding surface is a curve in (q, \dot{q}) space of slope λ and containing the time-varying reference configuration $r(t)$. Recall that, for the “redundant overcompensation”, the system follows the time varying state transition \mathbb{S} . In SMC, the controller needs to “force” the system trajectories to “move” while still pointing towards the surface as depicted in Figure 5 a,b) and theoretically proven in [31]. In other words, the curvature needs to decrease along the system trajectories such that

$$\frac{1}{2} \frac{d^2 S}{dt^2} \leq -\eta |S|, \quad \eta > 0. \quad (22)$$

315 The sliding condition in Equation 22 makes the surface S an invariant set, which is both
 316 a place in the (q, \dot{q}) space as well as a dynamics that the system will follow once reaching
 317 the surface, i.e., when $\mathbb{S} = S$ (see Equation 13).

Particularized to antifragile therapy control, the desired configuration can be given by a vector $r(t)$ which will be “pushed”, given the control law F to the convex (i.e.. antifragile) region of the drug dose–response depicted in Figure 5 c) and mathematically described as

$$r = r_{min} + \frac{r_{max} - r_{min}}{1 + \left(\frac{u}{\mu}\right)^{-n}}. \quad (23)$$

318 For the definition of the survival in Equation 23, we consider the Hill function with r_{min}
 319 minimal survival and r_{max} maximal survival, μ is the inflexion point of r beyond which
 320 increases of drug have less impact on survival, and n the Hill exponent. Note that the
 321 Hill function is a commonly employed mathematical model used to parameterize dose
 322 response assays [34], although other functions might be used. In our experiments, we
 323 consider the parametrization in [5], where: $r_{min} = 20$, $r_{max} = 100$, $n = 10$, $\mu = 10$.

324 The reference state $r(t)$ (i.e., the configuration) is computed such that the dynamics
 325 of the tumor-immune-drug dynamics network model is driven to the antifrangible region
 326 of the survival curve in Figure 6 a) using an antifrangible control law synthesized through
 327 the combined effect of the anticipation $K_q \dot{\phi}(q)$ and the variable structure component
 328 $\beta \operatorname{sgn}(S)$ of the controller, which ensure the curvature decrease along the system state
 329 trajectory (see green region in Figure 6 b)).




Figure 6. Antifrangible principle. a) The drug dose–response basis for the imposed dynamics of the network. The survival function assumes that the system needs to provide a drug dose such that the value of r decreases; b) The derivative of the survival function. This demonstrates that the inflexion point describes the point where the dosage (control law) can drive the tumor-immune-drug system in the antifrangible region (see West et al. [5]).

Putting all together, and hence combining the **redundant overcompensation** of the PD component in the antifrangible design together with the **variable structure** of the SMC, we have the following expression of the control law,

$$F(q) = K_\phi \phi(q) + K_q \dot{\phi}(q) + \beta \operatorname{sgn}(S), \quad (24)$$

330 where β is the SMC control gain and $\operatorname{sgn}(S)$ is the sign function applied to the chosen
 331 sliding surface S , parametrized such that the condition in Equation 22 is fulfilled.

332 Bounded high-frequency activity

The final step in our controller design is the computation of the control law F that verifies the sliding condition in Equation 22 and is discontinuous, in order to account for uncertainty, volatility, and disturbances. Although the choice of the control law is done with a trade-off between control bandwidth and tracking precision in mind (see the last term in Equation 24), we need to ensure the presence of the low-amplitude high-frequency “stressors” that will “shake” the system path towards the antifrangible region (see Figure 5 b, c). The standard SMC control law will cause the controlled system chattering due to the switching action of the control law [23], which is typically a

function of the form $\beta \operatorname{sgn}(S)$, where the SMC control gain $\beta = \beta(q, \dot{q})$ chosen such that Equation 22 is satisfied and

$$\operatorname{sgn}(S) = \begin{cases} 1 & \text{if } S > 0 \\ 0 & \text{if } S \leq 0 \end{cases}. \quad (25)$$

To reduce a high chattering value, which might excite high-frequency structural modes, neglected time-delays, or other modelling uncertainties, a saturation function can be chosen to replace the sgn function such as

$$\operatorname{sat}(S) = \begin{cases} \operatorname{sgn}(S) & \text{if } |S| \geq \Phi \\ \frac{S}{\Phi} & \text{if } |S| < \Phi \end{cases}. \quad (26)$$

This ensures that the control error φ is maintained within a guaranteed precision ϵ calculated as $\epsilon = \frac{\Phi}{\lambda}$. We build the sliding surface S as a linear combination of the control error φ and its velocity $\dot{\varphi}$ such as $S = \dot{\varphi} + \lambda\varphi$. Then, the control law in Equation 24 becomes

$$F(q) = K_{\varphi}\varphi(q) + K_{\dot{\varphi}}\dot{\varphi}(q) + \beta \operatorname{sat}(\dot{\varphi}(q) + \lambda\varphi(q)). \quad (27)$$

333 Finally, given the antifragile controller we synthesized, we will show, throughout
 334 our experiments in next section, that depending on the curvature of the drug-response
 335 reference dynamics on the manifold we impose, we can identify regions of “fragile”
 336 response in which the synthesized control law has low fluctuations, and regions of
 337 “antifragile” response in which the control law has large dose fluctuations, consistent
 338 with the work of West et al. [5]. But in order to achieve this behavior the choice of φ
 339 and T is very important [35]. For instance, φ determines how fast $q(t)$ reaches $r(t)$ in the
 340 topology induced by φ , whereas T determines how “simple” is the control law [23,36].
 341 This will subsequently impact the computation of the sliding surface S and the overall
 342 characteristics of the control law F , which in our model accounts for the drug dose v .

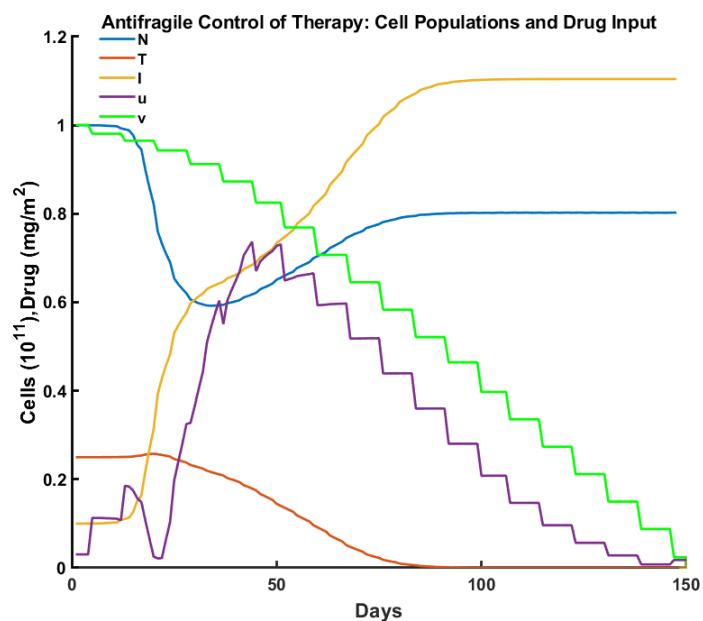


Figure 7. Antifragile control therapy. Individual cell populations evolution under the computation of the antifragile control law of drug dose v . The antifragile drug dose timing and duration is computed based on the control law in Equation 27.

Note that the survival dose curve parametrizes the reference configuration $r(t)$ in the control problem (see Figure 6 a)) such that the antifragile region in Figure 6 b) is attained by minimizing the error function φ given the actual state of the tumor-immune-drug network model in Equations 1–4. In the antifragile control design the anticipation control part (i.e., PD controller) mainly contributes to the normal stabilization of the tumor-immune-drug system, forcing the tracking errors in the boundary layer of the sliding surface through a redundant overcompensation and high control activity (see Figure 7). After entering the boundary layer of the surface (see Figure 5 b)), the tracking performance is dominantly controlled by the SMC control part, described by $\beta \text{sat}(\dot{\varphi}(q) + \lambda \varphi(q))$, which pushes the trajectory on the sliding surface (see Figure 5 a, b)).

2.3. State-of-the-art Control Algorithms in Cancer Therapy

Control-theoretic therapy design is used to describe treatment protocols which have the potential to be more efficient (i.e., maximize tumor kill, minimize patient toxicity) than standard static periodic protocols now in use [37]. As discussed in Section 1.1, static therapy plans cannot cope with the highly nonlinear phenomena emerging in tumor-immune-drug network models, such as periodic oscillations in tumor and normal cells generation [6,8], the resonance–antiresonance in drug response and toxicity [2,9,38], or the synchronization in tumor–environment cells interactions under drug pulses [10,39]. In order to handle such phenomena intrinsic in the tumor-immune-drug network, many control theoretic approaches were developed since the 80's, with notable work in [12,40] and up to recent work in [41,42].

Although employing different control laws, the core objective of all the approaches is to move the system into the basin of attraction of a “healthy” stable equilibrium state under the impact of drug dosing and timing [43] which is also suitable for surgical intervention [44]. In a more general view, this accounts for a state motion towards a reference manifold under the control law action, which then turns the manifold in an invariant set (i.e., the system is described by the manifold's equation).

2.3.1. Optimal control

In the optimal control synthesis the goal is to determine the therapy dose function $v(t)$ in Equation 4, representing the chemotherapy administration schedule [12,13,45]. This is computed such that that kill rate off the tumor cell population is as high as possible, with the constraint that the killing rate of normal cells is minimized. Although for our experiments we consider the simplified approach in [12], the general optimal control problem can be stated as follows.

The goal is to find the control variable, in our case the therapy dose $v(t)$ in Equation 4, and the (possibly free) final time t_f , corresponding to the end of therapy (e.g., after 150 days), that solves the following optimization problem

$$\begin{aligned} & \underset{v, t_f}{\text{minimize}} && J(v, t_f) = \varphi(x(t_f), t_f) \\ & \text{subject to} && \dot{x}(t) = f(x(t), v(t), t), \quad t_0 \leq t \leq t_f, \\ & && g(x(t)) \geq 0. \end{aligned} \quad (28)$$

In Problem 28, objective function $J = \varphi(x(t_f), t_f)$ is the tumor burden of the patient given the changes in the interacting cell populations N, T, I described by $x(t)$ and their evolution $f(x(t), v(t), t)$ under the drug administration. The state constraint $g(x(t))$ is used to maximize tumor kill and keep normal cells above a threshold. This formulation keeps the tumor cell population T lower at t_f at the price of large oscillations.

We update the problem such that we can weight each component of the objective function to avoid oscillations in the network model of Equations 1–4. We then rewrite the objective function J as a weighted combination of tumor burden at therapy end $T(t_f)$, the summed tumor burden over the treatment period $\int_0^{t_f} T(t) dt$, the maximum

tumor burden over the treatment course $\max_{t \in t_0, t_f} T(t)$, and, of course, the drug dose concentration $u(t)$ during the therapy duration. The updated optimal control objective function J is given as

$$J(v, t_f) = w_1 T(t_f) + w_2 \int_0^{t_f} T(t) dt + w_3 \max_{t \in t_0, t_f} T(t) + w_4 u(t), \quad (29)$$

where w_i are weighting constants, chosen as $w_1 = 1500, w_2 = 150, w_3 = 1000, w_4 = 40$. To evaluate the effect of the optimal control law has upon the tumor-immune-drug model we plot evolution of the system in Figure 8. The comparative evaluation and discussion will be done in Sections 3 and 4.




Figure 8. Optimal control therapy. Individual cell populations evolution under optimal computation of the drug dose v . The drug dosing amplitude, timing, and duration is computed as a constrained optimization problem that slows down and, eventually, disrupts tumor growth. When the tumor (T) decays under the detectable size (≈ 80 days), the normal cells population (N) restarts the normal proliferation pattern, and immune reaction (I) is enhanced to support the complete tumor extinction.

2.3.2. Robust control

Another line of research, which goes beyond the optimal control landscape, is the variable structure control, such as bang-bang control [46]. This approach allows for seamless compensation of un-modelled dynamics in the tumor-immune-drug model and parametric uncertainties alike [47,48]. Such models capture the dynamics of the phase field model of tumor growth [28,29] and the regularities of the cell populations states to control the phase of the drug concentration. In our study, we considered the model in [13] with parametrization considerations from [46].

In our therapy design problem, the control signal v is restricted to be between a lower and an upper bound of dosage $0 \leq v(t) \leq v_{max}$, as suggested by oncology guidelines. Additionally, in a special case of the optimal control in Problem 28, v switches from one extreme to the other (i.e., is strictly never in between the bounds). This is referred to as a bang-bang therapy solution, explored in both [13,46]. In order to

formulate the bang-bang control synthesis, we derive the control Hamiltonian² of the optimal control problem (see Equations 30) in order to identify the switching function of the system (i.e., ensuring that the drug entering the patient at time t is bounded).

$$H = \lambda_1 \left(\frac{dI}{dt} \right) + \lambda_2 \left(\frac{dT}{dt} \right) + \lambda_3 \left(\frac{dN}{dt} \right) + \lambda_4 \left(\frac{du}{dt} \right) + \eta, \quad (30)$$

where the functions λ_i satisfy the co-state variables equations

$$\begin{aligned} \frac{d\lambda_1}{dt} &= -\lambda_1 \left(\frac{\rho T}{\alpha + T} - c_1 T - d_1 - a_1(1 - e^{-u}) \right) + \lambda_2 c_2 T, \\ \frac{d\lambda_2}{dt} &= -\lambda_1 \left(\frac{\rho \alpha I}{(\alpha + T)^2} - c_1 I \right) - \lambda_2 (r_1 - 2r_1 b_1 T - c_2 I - c_3 N - a_2(1 - e^{-u})) + \lambda_3 c_4 N, \\ \frac{d\lambda_3}{dt} &= \lambda_2 c_3 N - \lambda_3 (r_2 - 2r_2 N^2 - c_4 T - a_3(1 - e^{-u})) - \eta(t), \\ \frac{d\lambda_4}{dt} &= -e^{-u} (a_1 \lambda_1 I + a_2 \lambda_2 T + a_3 \lambda_3 N). \end{aligned} \quad (31)$$

The η is chosen to ensure that the normal cells N are above 75% of the tumor-free normalized carrying capacity for this experiment such that

$$\eta(t) = \begin{cases} 1 & \text{if } N \leq 0.75 \\ 0 & \text{otherwise} \end{cases}. \quad (32)$$

We can now rewrite the robust control equation for the drug dosing as

$$\frac{\partial H}{\partial v} = \lambda_4, \quad (33)$$

which is independent of the control variable v . If we assume that the amount of drug entering the patient at time t is bounded above and satisfies $0 \leq v(t) \leq v_{max}$, then the output of the bang-bang control is given by

$$v(t) = \begin{cases} 0 & \text{if } \lambda_4 > 0 \\ v_{max} & \text{if } \lambda_4 < 0 \\ \text{singular} & \text{if } \lambda_4 = 0 \end{cases}. \quad (34)$$

As we can see in Equation 34, the co-state variable λ_4 is the switching function for the tumor-immune-drug network model we consider. In this case, the drug should be injected at the maximum rate when λ_4 is negative and should be ceased when λ_4 is positive. To evaluate the effect the robust control law has upon the tumor-immune-drug model, we plot the evolution of the system in Figure 9. The comparative evaluation with the other state of the art methods and the antifragile control will be done in Sections 3 and 4.

2.3.3. Pulsed control

Pulsed control therapy assumes that the administration of the drug into the patient body is modeled by a train of Dirac impulses. This approach is the traditional approach to chemotherapy and was approached by multiple studies [4,39,49]. In order to parametrize the control system one needs to specify: 1) the total number of treatment sessions N (basically the total number of impulses); 2) the time interval between treatment sessions Tn (inter-impulse time interval); and 3) the drug dose administered in each session An (accounting for the amplitudes of the impulses).

² Note that, for the therapy control design we use the control Hamiltonian that describes the conditions for optimizing some scalar function, basically the Lagrangian with respect to a control variable, and not the dynamics of the system itself.




Figure 9. Robust control therapy. Individual cell populations evolution under the computation of the robust control law of drug dose v based on Equations 31. The bang-bang drug dose timing and duration is computed such that the tumor population (T) decreases rapidly after treatment cessation (≈ 50 days). This decrease in tumor size is also determined by the exponentially increasing immune response (I) and marks also the normal cell (N) proliferation/death (i.e., through toxicity) pattern.

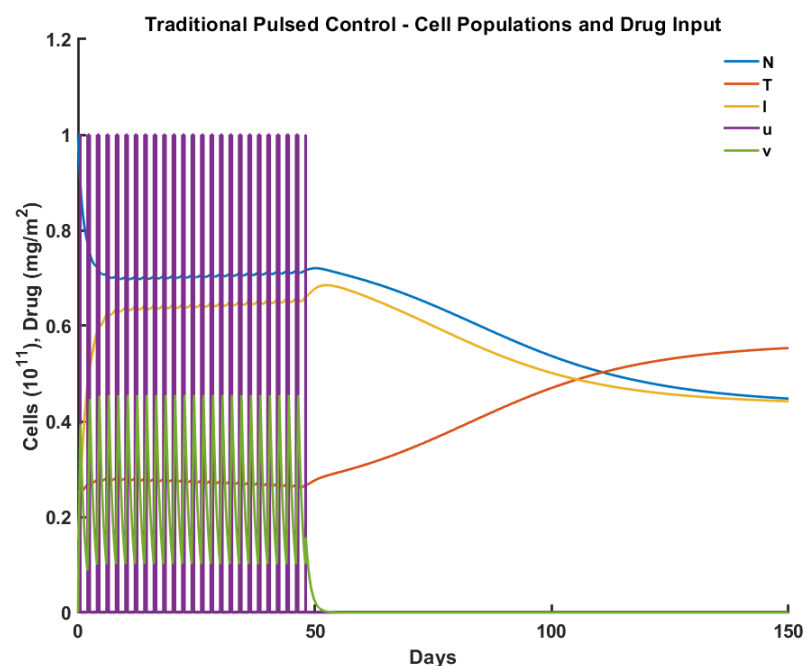


Figure 10. Pulsed control therapy. Individual cell populations evolution under the computation of the robust control law of drug dose v . The pulsed control computes densely timed, equally sized, and dosed drug administrations which only trigger short decays in tumor (T) proliferation, which at the end of the therapy resumes its uncontrolled growth which suppresses the normal cells (N) production and a weakening immune response (I).

409 Despite its simplicity, this control method is not applied in isolation [4]. The control
410 variable v , represented by a train of Dirac impulses, is typically computed in the optimal

control framework [39]. Additionally, the control design problem may be reduced to a finite-dimensional optimization problem that can be handled by a suitable solution, since the control variable influences the model in a linear fashion and we follow an a-priori known number of impulses. To evaluate the effect the pulsed control law has upon the tumor-immune-drug model, we plot evolution of the system in Figure 10. The comparative evaluation with the other state-of-the-art methods and the antifragile control will be done in Section 4. Finally, the fundamental optimal impulsive control sequence is utilized in a receding horizon strategy to produce a feedback control law. This strategy entails applying to the patient only the first control action of the sequence and repeating the entire dynamic optimization process beginning with the patient's state at the subsequent discrete time instant as shown in [4].

Our study departs from the traditional control problem formulation and goes beyond the analysis of the tumor-immune-drug model in order to have a more complete understanding of the implications of the existence of basins of attraction in the treatment process and design a therapy control scheme which goes beyond robustness and reaches antifragility, as we described in the previous section.

3. Experiments

In order to evaluate the antifragile control synthesis, we simulated our tumor-immune-drug network model without and under the impact of drug administration. For the simulation, we used the Matlab environment on a standard PC. In order to implement the robust and optimal control approaches, we used the library developed by [50]. The model parametrization of the tumor-immune-drug model is done as given in [12,13]. The full list of parameters of the tumor-immune-drug network model is given below in Table 1.

In order to evaluate the therapy controllers, we considered metrics of interest for the clinical use practice over a typical 150 days of chemotherapy. The overall evaluation is given in Table 2. The time to tumor elimination describes the number of days until the tumor decreased under a threshold of 0.05×10^{11} cells. For the drug impact assessment, we defined the maximum drug concentration within the patient body over the therapy duration and the total drug administered over the entire therapy duration. Finally, we also considered the impact the therapy has upon the tumor and the normal cells. This is characterized by the maximum number of tumor cells and the minimum number of normal cells, which describe the maximum size of the tumor during the therapy and the highest kill of normal cells during the therapy, respectively.

The pulsed control is a standard approach for chemotherapy that uses a train of impulses to deliver the drug. The parametrization of this type of control assumes the choice of the number of impulses, the inter-impulse time interval, and the drug dose administered through each impulse. For our experiments, we considered 25 impulses of amplitude 1 administered over the first 50 days of therapy (see Figure 10).

The robust control implementation was based on the computation of the control Hamiltonian of the optimal control problem and parametrized according to the work in [46]. The drug dosage was restricted to a maximum dose of 1 mg/L but with a varying administration time intertwined with no drug administration for an overall 18 effective administration days of the 150 total therapy days, as shown in Figure 9.

The optimal control problem was formulated based on an objective function that weighted the tumor size evolution, the total tumor size, the maximum tumor size along the therapy, and the drug concentration—see Equation 29. The optimal dosing was irregular in both amplitude and duration, determining a rather fast decrease of the overall concentration in the first days of therapy, as depicted in Figure 8.

Finally, for the antifragile control the drug dosing was computed such that the survival curve (i.e., tumor size vs. drug dose) points towards the antifragile region, that is the convex region of the curvature of the survival curve (i.e., second derivative of Hill function). The antifragile control law, composed of the PD-like anticipation component

Model parameter	Parameter (initial) value	Impact on dynamics
Number of normal host cells N	1.0	Interacts with I and T w/wo u
Number of tumor cells T	0.25	Interacts with N and I w/wo u
Number of host immune cells I	0.1	Interacts with N and T w/wo u
Drug concentration at the tumor location u	0.01	induces cell kill by toxicity (normal+tumor)
Drug administration v	0.0	modulates the drug concentration at the tumor
Fraction normal cell kill a_1	0.2	dose-related weight of normal cell kill (toxicity)
Fraction tumor cell kill a_2	0.3	dose-related weight of tumor cell kill
Fraction immune cell kill a_3	0.1	dose-related weight of immune cell kill
Tumor growth rate r_1	1.5	typically $c_2 + c_3$
Normal cells growth rate r_2	1.0	per capita growth
Carrying capacity of tumor cells b_1	1.0	weights tumor cells self-excitation
Carrying capacity of normal cells b_2	1.0	weights normal cells self-excitation
Tumor-immune competition factor c_1	1.0	competition
Immune-tumor competition factor c_2	0.5	modulation
Tumor-normal competition factor c_3	1.0	competition
Normal-tumor competition factor c_4	1.0	competition
Immune cells death rate d_1	0.2	regulation through $\frac{s}{d_1}$
Drug influx modulation d_2	1.0	rate of change/decay of u
Immune threshold rate α	0.3	related to the immune response curve
Immune response rate ρ	0.01 - 1.0	immune-compromised to healthy

Table 1: Tumor–immune–drug network model parameters.

Criteria	Control	Pulsed	Robust	Optimal	Antifragile
Time to tumor elimination (days)		150+	68	100	65
Max drug concentration (mg/L)		0.455	0.987	0.659	0.736
Total drug administered (mg/L)		15.00	12.00	15.00	14.31
Max tumor cells ($\times 10^{11}$ cells)		0.553	0.258	0.257	0.251
Min normal cells ($\times 10^{11}$ cells)		0.447	0.706	0.750	0.592

Table 2: Therapy control algorithms evaluation: evaluation criteria and candidate approaches.

464 and the SMC-like variable structure component, pushes the dynamics of the tumor–
465 immune–drug model in a state where the drug concentration ensures a fast elimination
466 of the tumor—see Figure 7. For the detailed parametrization of the simulations please
467 check the source code available on GitLab ³.

468 **4. Discussion**

469 In clinical practice, it is typical to follow a set treatment regimen in which a steady
470 dosage is given on a regular basis. There are, however, a large number of noteworthy
471 cases where the tumor responds better to a “volatile” treatment plan than continuous
472 therapy. Such a volatile (or synonymously an “uneven”) dose control can actually push
473 the tumor growth dynamics towards instability and, under the joint effect of drug and
474 immune system, towards elimination. We explored these cases in our experiments using
475 an asymmetric networks of the tumor–immune–drug interactions.

476 In our study, we were interested to explore how antifragile control can offer a
477 systematic approach to control drug dosage in “volatile” treatment plans where the
478 drug “unevenness” can determine a fast tumor elimination with a good trade-off in
479 drug quantity administration and collateral normal cell damage (i.e., toxicity induced
480 cell death). We summarize our findings in Table 2. Here we notice how each control
481 approach determines gains in one of the multiple dimensions of the therapy regimen
482 efficiency. When considering time to tumor elimination, the antifragile controller is able
483 to reduce the tumor under the threshold of 0.05×10^{11} cells in under 65 days whereas
484 the robust control reaches the same tumor burden in 68 days, as depicted in Figure
485 9 - red trace. This demonstrates that the antifragile control can compensate for the
486 anti-symmetric interactions in the tumor–immune–drug model and push the system
487 towards a faster tumor reduction (see therapy day 65 in Figure 7).

488 Considering the drug administration, the traditional pulsed approach stands out,
489 through the lowest maximum drug concentration among all methods. This comes also
490 with the price of having just a limited impact on the tumor growth dynamics, which
491 starts to diverge once the regimen reaches the maximum delivery program, around day
492 52 in Figure 10. All the other methods determine a higher maximum drug concentration,
493 with up to 80% more, but with the better outcome of reducing the tumor burden in
494 between 65 up to 80 days—see Figures 7, 8, and 9. Interestingly, the most drug quantity
495 efficient approach is the robust control where only 12 mg/L were administered during
496 the therapy, but due to its delivery pattern (i.e., frequency and timing) reached the
497 highest maximum drug concentration of 0.987 mg/L. The antifragile control, utilized a
498 moderate total drug dose of 14.31 mg/L distributed over the entire 150 days of therapy
499 for a maximum dose of 0.736 mg/L, as seen in Figure 7.

500 Considering the individual evolution of the cell populations in the tumor–immune–
501 drug model, we can see that the pulsed therapy is the one which determined the largest
502 tumor size of 0.553×10^{11} cells at day 150 of the therapy regimen, due to its limited
503 delivery schedule and the conservative dosage. This also determines its inefficiency
504 in tumor elimination and, of course, a high loss in the preservation of the normal
505 cells, of which many were killed due to a high toxicity level. This is visible through
506 the large loss of normal cells of 0.447×10^{11} cells at the end of therapy. On the other
507 side of the spectrum, the highest tumor kill strength was produced by the antifragile
508 control, with close results by the robust and the optimal control approaches. This is,
509 of course, motivated by the weighted objective functions in the case of optimal and
510 robust controllers, and the advantages of the anticipation and the high-frequency control
511 activity of the antifragile control. With a minimal tumor burden of 0.251×10^{11} , the
512 antifragile controller is the one which elicited a sustained, high frequency control activity
513 for the entire duration of the therapy, but with moderate drug administration and
514 concentration—see Figure 7.

³ Code available at: <https://gitlab.com/akii-microlab/antifragile-therapy-ctrl>

Finally, when considering the therapy damage to normal cells, quantified by the minimum number of normal cells, the optimal control stands out with a maximum of 0.750×10^{11} cells, followed by the robust and antifragile approaches. Overall, the antifragile control provides a fast reduction of the tumor burden, given a moderate drug administration and drug concentration, with the highest tumor kill and moderate collateral damage of the normal cells through drug induced toxicity.

Our study demonstrates that through the antifragile control framework we can find a trade-off between the time to reduce the tumor burden and the drug administration. This happens while maximising the tumor kill and conserving the normal cells given the complex anti-symmetric interactions among the tumor, the immune system, and the chemotherapeutic agent. Although this is an initial instantiation of antifragile control, we believe that such control framework can gain from the uncertainty and volatility describing the system's dynamics. This can be achieved by building a strong anticipation component and a versatile structure change driving the closed-loop system towards the reference behavior.

5. Conclusion

Cancer therapy control is a complex intervention which revolves around understanding the complex and intricate interactions among the tumor, the immune system, and the chemotherapeutic drug and breaking the tie towards minimal tumor burden. Aside the traditional pulsed therapy schemes, multiple control-theoretic approaches, such as robust and optimal control, were developed with the aim to capture the complex tumor-immune-drug dynamics and compute an optimal and robust drug dosing that minimizes the tumor and preserves as much as possible normal cells. However, there is always a trade-off between tumor reduction and normal cells kill. Antifragile control emerges as a control-theoretic framework capable of gaining from: 1) the uncertainty describing tumor-immune-drug interactions, 2) the volatility of drug response curves in patient populations, and 3) the variability in tumor response, immune capabilities, and drug resistance patterns of a patient-tumor pair. This is achieved through the formal implementation in a control-theoretic framework of the three main pillars: 1) redundant overcompensation, 2) structure-variability, and 3) high-frequency control activity. These three ingredients are concretely implemented in a control structure cast in the Riemannian geometry where a PD control-like structure is combined with a SMC control-like component to elicit strong anticipation and convergent dynamics. As our experiments emphasize, this novel control-theoretic approach has the potential to go beyond the robust and optimal control with the advantage that the closed-loop system doesn't need to model uncertainty, adapt to changes, or seek optimality. Antifragile control anticipates future variability in system dynamics and exploits drug dosing volatility for a fast tumor burden reduction with a trade-off in drug concentration and collateral normal cell kill, through the induced drug toxicity. Our next steps will focus on a broader formalization of the antifragile control theory and its instantiation in a data-driven control loop, where the tumor-immune-drug dynamics are learnt from data and not given by a parametric model. This assumes simultaneously learning tumor growth curves, pharmacokinetics, and immune responses. We are confident that this approach will make antifragile control more versatile and our initial results in the present study and the one in [3] are encouraging. Finally, we believe that the antifragile control of non-parametric system models can be highly relevant, as it could additionally capture aspects such as drug resistance, specific immune cell types dynamics, and the incorporation of the spatial heterogeneity of tumors. We believe that such an approach will strengthen oncologists' therapy tool-set by informing therapy design and patient tailoring.

Author Contributions: CA designed the research, designed the framework, designed the experiments, implemented simulation software, analyzed and interpreted the data and results, DK designed the experiments, analyzed and interpreted the data and results, MS designed the

framework, implemented simulation software, analyzed and interpreted the data and results. CA, DK, and MS wrote and revised the manuscript.

Conflicts of Interest: The authors declare no conflict of interest.

570

- 571 1. Nia, H.T.; Munn, L.L.; Jain, R.K. Physical traits of cancer. *Science* **2020**, 370, eaaz0868.
- 572 2. Schättler, H.; Ledzewicz, U. Optimal control for mathematical models of cancer therapies. *An application of geometric methods* **2015**.
- 573 3. Kurz, D.; Sánchez, C.S.; Axenie, C. Data-driven Discovery of Mathematical and Physical Relations in Oncology Data using Human-understandable Machine Learning. *Frontiers in Artificial Intelligence* **2021**, 4.
- 574 4. Belfo, J.P.; Lemos, J.M. *Optimal Impulsive Control for Cancer Therapy*; Springer, 2020.
- 575 5. West, J.; Strobl, M.; Armagost, C.; Miles, R.; Marusyk, A.; Anderson, A.R. Antifragile therapy. *bioRxiv* **2020**.
- 576 6. Kim, M.; Gillies, R.J.; Rejniak, K.A. Current advances in mathematical modeling of anti-cancer drug penetration into tumor tissues. *Frontiers in oncology* **2013**, p. 278.
- 577 7. McDonald, E.; El-Deiry, W.S. Cell cycle control as a basis for cancer drug development. *International journal of oncology* **2000**, 16, 871–957.
- 578 8. Hu, X.; Jang, S.R.J. Dynamics of tumor-CD4+–cytokine–host cells interactions with treatments. *Applied Mathematics and Computation* **2018**, 321, 700–720.
- 579 9. Agur, Z.; Kheifetz, Y. Resonance and anti-resonance: from mathematical theory to clinical cancer treatment design. *Handbook of Cancer Models with Applications to Cancer Screening, Cancer Treatment and Risk Assessment* **2005**.
- 580 10. Agur, Z.; Kheifetz, Y. Optimizing Cancer Chemotherapy: From Mathematical Theories to Clinical Treatment. In *New Challenges for Cancer Systems Biomedicine*; Springer, 2012; pp. 285–299.
- 581 11. Pillis, L.d.; Radunskaya, A. Modeling Immune-Mediated Tumor Growth and Treatment. In *Mathematical Oncology 2013*; Springer, 2014; pp. 199–235.
- 582 12. De Pillis, L.G.; Radunskaya, A. The dynamics of an optimally controlled tumor model: A case study. *Mathematical and computer modelling* **2003**, 37, 1221–1244.
- 583 13. De Pillis, L.G.; Radunskaya, A. A mathematical tumor model with immune resistance and drug therapy: an optimal control approach. *Computational and Mathematical Methods in Medicine* **2001**, 3, 79–100.
- 584 14. Taleb, N.N. *Antifragile: Things that gain from disorder*; Vol. 3, Random House, 2012.
- 585 15. Taleb, N.N. (anti) fragility and convex responses in medicine. International Conference on Complex Systems. Springer, 2018, pp. 299–325.
- 586 16. Goutelle, S.; Maurin, M.; Rougier, F.; Barbaut, X.; Bourguignon, L.; Ducher, M.; Maire, P. The Hill equation: a review of its capabilities in pharmacological modelling. *Fundamental & clinical pharmacology* **2008**, 22, 633–648.
- 587 17. Gaffney, E.A. The application of mathematical modelling to aspects of adjuvant chemotherapy scheduling. *Journal of mathematical biology* **2004**, 48, 375–422.
- 588 18. Fedorinov, D.S.; Lyadov, V.K.; Sychev, D.A. Genotype-based chemotherapy for patients with gastrointestinal tumors: focus on oxaliplatin, irinotecan, and fluoropyrimidines. *Drug Metabolism and Personalized Therapy* **2021**.
- 589 19. Paraiso, K.H.; Smalley, K.S. Fibroblast-mediated drug resistance in cancer. *Biochemical pharmacology* **2013**, 85, 1033–1041.
- 590 20. Bejenaru, A.; Udriste, C. Riemannian optimal control. *arXiv preprint arXiv:1203.3655* **2012**.
- 591 21. Lee, J.M. *Riemannian manifolds: an introduction to curvature*; Vol. 176, Springer Science & Business Media, 2006.
- 592 22. Bloch, A.M. An introduction to aspects of geometric control theory. In *Nonholonomic mechanics and control*; Springer, 2015; pp. 199–233.
- 593 23. Zou, L.; Wen, X.; Karimi, H.R.; Shi, Y. The identification of convex function on Riemannian manifold. *Mathematical Problems in Engineering* **2014**, 2014.
- 594 24. Bullo, F.; Murray, R.M. Proportional derivative (PD) control on the Euclidean group. *Caltech Reports* **1995**.
- 595 25. Bécigneul, G.; Ganea, O.E. Riemannian adaptive optimization methods. *arXiv preprint arXiv:1810.00760* **2018**.
- 596 26. Fiori, S. Manifold Calculus in System Theory and Control—Fundamentals and First-Order Systems. *Symmetry* **2021**, 13, 2092.

- 625 27. Fiori, S.; Cervigni, I.; Ippoliti, M.; Menotta, C. Synchronization of dynamical systems on
626 Riemannian manifolds by an extended PID-type control theory: Numerical evaluation.
627 *Discrete and Continuous Dynamical Systems-B* **2022**.
- 628 28. Guo, Y.; Xu, B.; Zhang, R. Terminal sliding mode control of mems gyroscopes with finite-time
629 learning. *IEEE Transactions on Neural Networks and Learning Systems* **2020**, *32*, 4490–4498.
- 630 29. Colli, P.; Gilardi, G.; Marinoschi, G.; Rocca, E. Sliding mode control for a phase field system
631 related to tumor growth. *Applied Mathematics & Optimization* **2019**, *79*, 647–670.
- 632 30. Ouyang, P.; Acob, J.; Pano, V. PD with sliding mode control for trajectory tracking. *Robotics*
633 *and Computer-Integrated Manufacturing* **2014**, *30*, 189–200.
- 634 31. Slotine, J.J.E.; Li, W.; others. *Applied nonlinear control*; Vol. 199, Prentice hall Englewood Cliffs,
635 NJ, 1991.
- 636 32. DeCarlo, R.A.; Zak, S.H.; Matthews, G.P. Variable structure control of nonlinear multivariable
637 systems: a tutorial. *Proceedings of the IEEE* **1988**, *76*, 212–232.
- 638 33. Utkin, V. Variable structure systems with sliding modes. *IEEE Transactions on Automatic*
639 *control* **1977**, *22*, 212–222.
- 640 34. Meyer, C.T.; Wooten, D.J.; Paudel, B.B.; Bauer, J.; Hardeman, K.N.; Westover, D.; Lovly, C.M.;
641 Harris, L.A.; Tyson, D.R.; Quaranta, V. Quantifying drug combination synergy along potency
642 and efficacy axes. *Cell systems* **2019**, *8*, 97–108.
- 643 35. Maithripala, D.S.; Berg, J.M. An intrinsic PID controller for mechanical systems on Lie
644 groups. *Automatica* **2015**, *54*, 189–200.
- 645 36. Zhang, Z.; Sarlette, A.; Ling, Z. Integral control on Lie groups. *Systems & Control Letters* **2015**,
646 *80*, 9–15.
- 647 37. Lecca, P. Control Theory and Cancer Chemotherapy: How They Interact. *Frontiers in*
648 *Bioengineering and Biotechnology* **2021**, *8*, 621269.
- 649 38. Li, D.; Xu, W.; Guo, Y.; Xu, Y. Fluctuations induced extinction and stochastic resonance effect
650 in a model of tumor growth with periodic treatment. *Physics Letters A* **2011**, *375*, 886–890.
- 651 39. Ren, H.P.; Yang, Y.; Baptista, M.S.; Grebogi, C. Tumour chemotherapy strategy based on
652 impulse control theory. *Philosophical Transactions of the Royal Society A: Mathematical, Physical*
653 *and Engineering Sciences* **2017**, *375*, 20160221.
- 654 40. Swan, G.W. Role of optimal control theory in cancer chemotherapy. *Mathematical biosciences*
655 **1990**, *101*, 237–284.
- 656 41. Carrere, C. Optimization of an in vitro chemotherapy to avoid resistant tumours. *Journal of*
657 *Theoretical Biology* **2017**, *413*, 24–33.
- 658 42. Irurzun-Arana, I.; Janda, A.; Ardanza-Trevijano, S.; Trocóniz, I.F. Optimal dynamic control
659 approach in a multi-objective therapeutic scenario: Application to drug delivery in the
660 treatment of prostate cancer. *PLOS Computational Biology* **2018**, *14*, 1–16. doi:10.1371/jour-
661 nal.pcbi.1006087.
- 662 43. Uthamacumaran, A. A review of dynamical systems approaches for the detection of chaotic
663 attractors in cancer networks. *Patterns* **2021**, *2*, 100226.
- 664 44. Axenie, C.; Kurz, D. Chimera: Combining mechanistic models and machine learning
665 for personalized chemotherapy and surgery sequencing in breast cancer. *International*
666 *Symposium on Mathematical and Computational Oncology*. Springer, 2020, pp. 13–24.
- 667 45. Wang, S. Optimal control for cancer chemotherapy under tumor heterogeneity. 2019 IEEE
668 58th Conference on Decision and Control (CDC). IEEE, 2019, pp. 5936–5941.
- 669 46. Ledzewicz, U.; Maurer, H.; Schättler, H. Bang-bang and singular controls in a mathematical
670 model for combined anti-angiogenic and chemotherapy treatments. *Proceedings of the 48h*
671 *IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control*
672 *Conference*. IEEE, 2009, pp. 2280–2285.
- 673 47. Ledzewicz, U.; Schättler, H. Optimal bang-bang controls for a two-compartment model in
674 cancer chemotherapy. *Journal of optimization theory and applications* **2002**, *114*, 609–637.
- 675 48. Ledzewicz, U.; Schättler, H.; Gahrooi, M.R.; Dehkordi, S.M. On the MTD paradigm and
676 optimal control for multi-drug cancer chemotherapy. *Mathematical Biosciences & Engineering*
677 **2013**, *10*, 803.
- 678 49. Panetta, J.C. A mathematical model of periodically pulsed chemotherapy: tumor recurrence
679 and metastasis in a competitive environment. *Bulletin of mathematical Biology* **1996**, *58*, 425–
680 447.
- 681 50. Kelly, M. An Introduction to Trajectory Optimization: How to do your own Direct Collocation.
682 *SIAM Review* **2017**, *59*, 849–904.