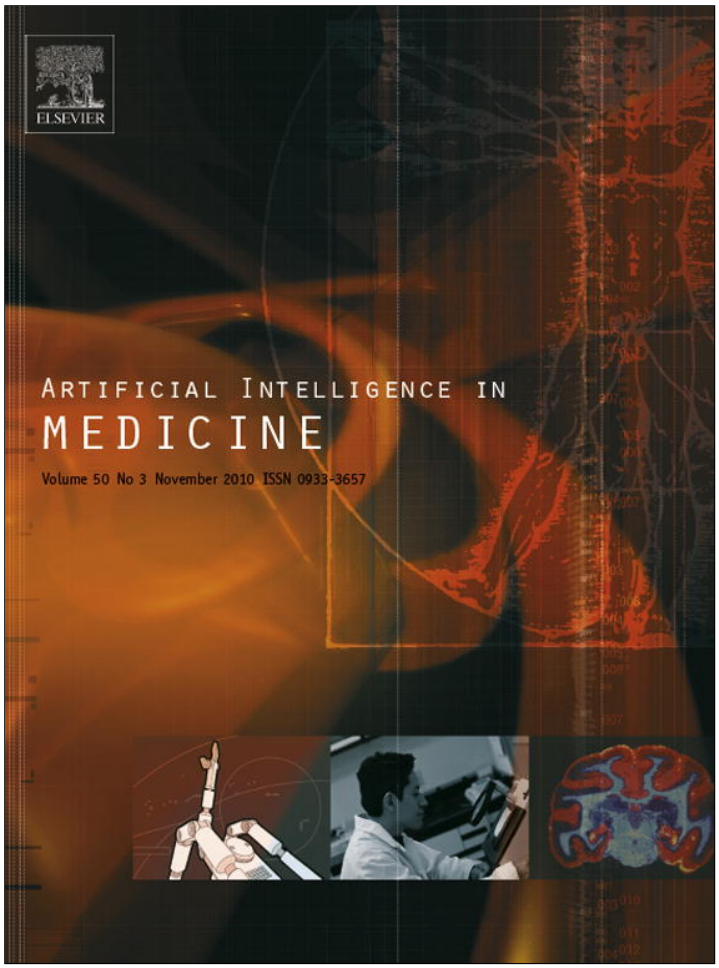

Artificial Intelligence in Medicine 50 (2010) 163–173

Contents lists available at ScienceDirect

Artificial Intelligence in Medicine

journal homepage: www.elsevier.com/locate/aiim

Modeling and optimization of combined cytostatic and cytotoxic cancer chemotherapy

Minaya Villasana [a,*], Gabriela Ochoa [b], Soraya Aguilar [c]

[a] Departamento de Ćomputo Científico y Estadística, Universidad Simon Bolivar, AP 89000 Caracas 1081-A, Venezuela
[b] School of Computer Science, Automated Scheduling, Optimisation, and Planning (ASAP) Research Group, University of Nottingham, Nottingham NG8 1BB, UK
[c] Departamento de Computación, Facultad de Ciencia y Tecnología, Universidad de Carabobo, Valencia, Venezuela

ARTICLE INFO

Article history:
Received 21 July 2009
Received in revised form 26 May 2010
Accepted 30 May 2010

Keywords:
Evolutionary algorithms
Evolution strategies
Singular optimal control
Tumor model
Delay differential equation system
Cytostatic drug
Cytotoxic drug
Combined chemotherapy
Cancer chemotherapy

ABSTRACT

Objectives: This study extends a previous mathematical model of cancer cytotoxic chemotherapy, which considered cycling tumor cells and interactions with the immune system, by incorporating a different type of drug: a cytostatic agent. The effect of a cytostatic drug is to arrest cells in a phase of their cycle. In consequence, once tumor cells are arrested and synchronized they can be targeted with a cytotoxic agent, thus maximizing cell kill fraction and minimizing normal cell killing. The goal is to incorporate the new drug into the chemotherapy protocol and devise optimal delivery schedules.
Methods: The problem of designing efficient combined chemotherapies is formulated as an optimal control problem and tackled using a state-of-the-art evolutionary algorithm for real-valued encoding, namely the *covariance matrix adaptation* evolution strategy. Alternative solution representations and three formulations of the underlying objective function are proposed and compared.
Results: The optimization problem was successfully solved by the proposed approach. The encoding that enforced non-overlapping (simultaneous) application of the two types of drugs produced competitive protocols with significant less amount of toxic drug, thus achieving better immune system health. When compared to treatment protocols that only consider a cytotoxic agent, the incorporation of a cytostatic drug dramatically improved the outcome and performance of the overall treatment, confirming *in silico* that the combination of a cytostatic with a cytotoxic agent improves the efficacy and efficiency of the chemotherapy.
Conclusion: We conclude that the proposed approach can serve as a valuable decision support tool for the medical practitioner facing the complex problem of designing efficient combined chemotherapies.

1. Introduction

The main goal in cancer chemotherapy is to cure the patient, and it is important to do so as efficiently as possible. Several alternatives to enhance chemotherapy treatments have been proposed [1] such as using combinations of toxic drugs, immunotherapy and more recently virotherapy [2]. This article focuses on the use of a cytostatic drug to aid a cytotoxic drug in chemotherapy. There is evidence in the medical literature [3], that this type of combined therapy has increased effectiveness.

The model formulated by Villasana and Radunskaya [4] considers the tumor growth, its interaction with the immune system and the action of a cycle-specific cytotoxic drug. In [5] this model is used in an optimal control formulation of the chemotherapy scheduling problem, which is successfully solved using modern heuristic search methods. A more recent study [6] considered the effect of different terms in the objective function of the optimal control formulation (such as the tumor levels, the immune system level, and the number of treatment cycles) on the overall features and efficacy of the obtained treatments.

Other authors have formulated the design of chemotherapy schedules from the point of view of optimal control [7–9], solving the stated optimization problem either analytically or numerically. However, for increasingly complex and realistic cancer models, analytical or traditional numerical methods are no longer applicable, and some authors have turned to meta-heuristics to optimize chemotherapy schedules. Petrovski, McCall and colleagues, have extensively and successfully used evolutionary algorithms and other modern heuristics in this domain [10–12]. Their work differs from the approach in [5], mainly in the underlying mathematical model of tumor growth. While Petrovski et al. considered the Gompertz growth model with linear cell-loss effect [10], without including interactions with the immune system;

* Corresponding author. Tel.: +58(212)9063386; fax: +58(212)9063364.
E-mail addresses: mvillasa@usb.ve, minayavillasana@hotmail.com (M. Villasana), gxo@cs.nott.ac.uk (G. Ochoa), saguilar@uc.edu.ve (S. Aguilar).

0933-3657/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.artmed.2010.05.009

Villasana et al. employed a more realistic cancer model [4]. Specifically, our model includes the interactions between tumor cells and immune cells; and differentiates between cell phases for subsequent treatment with a cycle-phase-specific drug. In a more recent book chapter, McCall et al. present a survey of approaches employing heuristic search methods to solve the cancer chemotherapy scheduling problem via optimal control. Examples of these approaches include the use of simulated annealing on a model of tumor and host cell interaction, a parallelized genetic algorithm, and multimodal optimization genetic algorithms (see [13] and the references therein for further details). More recently, Liang and colleagues have applied several algorithms to the chemotherapy scheduling problem using optimal control, where the underlying dynamics follows a modification of Martin's original model [14]. In [15], the authors use a genetic algorithm to solve the proposed optimal control problem, while in [16], they combine the genetic algorithm with the forward iterative dynamic programming as the local search in a memetic approach. However, none of these published studies based on heuristics search methods considers a drug that is not cytotoxic nor do the models incorporate the tumor interaction with the body's natural defense system.

Swierniak et al. [17] published a series of models for tumor growth using cycle-phase-specific drugs. The authors also developed analytical relations for the optimal drug scheduling on the simpler of those problems. In their exposition, a model that incorporates a cytostatic drug is included and the numerical solution for the optimal control model using Pontryagin's Maximum Principle is obtained. The optimal solution encountered was bang–bang (i.e. a solution that only takes upon the maximum and minimum values on a bounded range) with non-overlapping applications of the two types of drugs. Some of the models considered were simple enough to be still mathematically tractable, and thus, analytic solutions were readily available. However, the model that took into account the cell arrest, considered a single treatment cycle, while in practice cancer treatments are composed of multiple cycles. For more complex models of tumor growth, and multiple drug applications, mathematical manipulation becomes prohibitive. In consequence, an understanding of the qualitative features of the treatments that would be obtained in those cases is still lacking.

The present study extends our previous work [4–6] with the goal of suggesting more efficient cancer treatments. Specifically, a modification of the model presented in [4] is carried out so that it incorporates a different type of drug, which would act as a cytostatic agent in conjunction with the original cytotoxic agent. The idea behind combining these two agents, is that the cytostatic drug can halt the rapid progression of the cancerous cells through their cell cycle at a certain phase. Thus, when the cells are released, they are mostly arrested in the most vulnerable stage to the action of cytotoxic drugs. The overall strategy is that once cells are arrested and synchronized in the cell cycle, these can be targeted with a cytotoxic agent, thus maximizing cell kill fraction and minimizing normal cell killing. An example of a cycle-phase-specific cytotoxic drug is Taxol (paclitaxel), and an example of a cytostatic drug is Iressa (gefitinib). These are the drugs that were identified and modeled in our approach. Our study proposes and compares several treatment encodings and optimal control formulations of the chemotherapy scheduling problem. We present a detailed analysis of the treatments obtained, and a comparison with previous treatments that do not include the cytostatic agent.

The article is organized as follows. Section 2 presents the mathematical formulation of the problem, including the relevant biomedical background, the mathematical model describing the patient dynamics and the optimal control formulation. Thereafter, Section 3 details the methodology, including the alternative problem encodings and objective functions, and the evolutionary algorithm employed. Section 4 outlines the results, while Section 5 summarizes and discusses the main findings.

2. Problem formulation

2.1. Biomedical background

Most chemotherapy drugs work by attacking cells that are dividing rapidly. Normal cells divide at a self-regulated rate with tight controls in its progression in the cell cycle. In cancer cells, these controls are bypassed giving way to defective cells unable to control their reproduction, thus leading to the formation of a tumor or blood cancer. Chemotherapy drugs interfere with the division of these cells and may cause the cancer to recede completely. The treatment reduces the number of cancerous cells to a minimum level, at which point other mechanisms (e.g. programmed cell death) will remove the remaining tumor cells.

The cell cycle is the process leading to cell division. It encompasses four stages: G_1, S, G_2, and M, where G_1 and G_2 are resting phases (or Gap periods), S is the synthetic period, and M or mitosis is the time during which cells segregate the duplicated DNA material between daughter cells. Cycle-phase-specific drugs are those acting on a specific phase of the cell cycle. These drugs are either *cytotoxic*, or *cytostatic*. Cytotoxic drugs are toxic to the cells, thus killing them, while cytostatic drugs are not aimed at killing cancer cells but rather at stopping them from multiplying and trapping them in the cell cycle progression. When the concentration levels of the cytostatic drug fades, the cells are then released to continue in the cell cycle.

An example of a cytotoxic phase-specific drug is Taxol (paclitaxel) which has shown high efficacy in the treatment of breast, ovarian, head, and neck cancer. The action of this drug is carried through different mechanisms: it inhibits mitosis, induces apoptosis (programmed cell death), and enhances tumor radiosensitivity. Today, paclitaxel is used either as a single agent or accompanied by other drugs. The optimal scheduling and possible drug interactions for paclitaxel are not yet fully understood [18]. An example of a cytostatic drug is Iressa (gefitinib). Gefitinib is the first selective inhibitor of epidermal growth factor receptor's (EGFR) tyrosine kinase domain. Over-expression of EGFR is observed in certain types of carcinomas (for example lung and breast) leading to uncontrolled cell proliferation. Gefitinib inhibits EGFR tyrosine kinase by binding to the adenosine triphosphate (ATP) binding site of the enzyme. Thus the function of the EGFR tyrosine kinase in activating the Ras signal transduction cascade is inhibited, and malignant cells are inhibited. The study presented in [3] confirmed that the combination of these two drugs (paclitaxel and gefitinib) produces higher toxicity for the cancer cells. The Iressa drug acts by inducing a delay in cell cycle progression, with a complete arrest of G_1 cell phase growth after 72 h of treatment (daily dose of a 250 mg tablet). Iressa has been used with Taxol in clinical trials on mice obtaining better results than those treated exclusively with Taxol [19].

2.2. Mathematical model

The patient model used [4] is a competition model of tumor growth that includes the immune system response. The model considers three populations of cells: immune system, tumor during interphase (period comprising G_1 through G_2), and tumor during mitosis. Delay differential equations are used to take into account the phases of the cell cycle.

In the model, $T_1(t)$ and $T_M(t)$ denote the population of tumor cells during interphase and mitosis at time t respectively. $I(t)$ represents the immune system population at time t, that we take

as the cytotoxic T cells (CTL) (see [4] for a full discussion). Let τ be the resident time of cells in interphase. The governing equations for the patient dynamic are:

$$\begin{aligned} T'_I &= 2a_4T_M - (c_1I + d_2)T_I - a_1T_I(t-\tau) \\ T'_M &= a_1T_I(t-\tau) - d_3T_M - a_4T_M - c_3T_MI \\ I' &= k + \frac{\rho I(T_I+T_M)^3}{\alpha + (T_I+T_M)^3} - c_2IT_I - c_4T_MI - d_1I \end{aligned} \tag{1}$$

where $'$ denotes derivatives with respect to time and with initial data given by:

$$\begin{aligned} T_I(t) &= \phi_1(t) \quad \text{for} \quad t \in [-\tau, 0] \\ T_M(t) &= \phi_2(t) \quad \text{for} \quad t \in [-\tau, 0] \\ I(t) &= \phi_3(t) \quad \text{for} \quad t \in [-\tau, 0] \end{aligned} \tag{2}$$

This system (1) can have up to 5 different fixed points depending on the parameter values [4], one of which is always present, namely $(0, 0, k/d_1)$. This fixed point represents the desirable scenario of a tumor-free environment with positive immune population.

2.3. Including the drug actions

The action of a cytotoxic drug can be described as follows: for high concentrations of the drug, it arrests tumor cells in mitosis where they die naturally when they fail to continue in the cycle. This can be modeled by assuming that a tumor cell that encounters the drug is taken out of the cycle and can no longer proliferate. High drug concentration also impairs other cell lines by either destroying cells (as in the case of tumor cells) or diminishing their ability to attack (as in the case for immune system cells). The fraction of cells that are targeted by this drug can be modeled by a term of the form $-k_j(1 - e^{-h_iu})$, where u is the drug concentration and h_i and k_j are parameters that model the effectiveness. The decay rate of paclitaxel is modeled as before ([4,5]) with two separate elimination terms. Thus the decay function is expressed as:

$$decay_u(t) = r_1e^{-\lambda_1 t} + r_2e^{-\lambda_2 t} \tag{3}$$

with r_1 and r_2 are the real non-dimensional constants.

Let u_1 and u_2 be such that the drug concentration at any given time is the linear convex combination $u(t) = r_1u_1(t) + r_2u_2(t)$. Eqs. (4) and (5) of system (7) model this situation with multiple drug applications in time, identified with the function $c_u(t)$, which is the concentration of paclitaxel that goes in the system at time t. With this choice and initial conditions $u_1(0) = 0$ and $u_2(0) = 0$ we get:

$$u(t) = c_u(t) * decay_u(t) \tag{4}$$

where * denotes convolution.

On the other hand, the action of high concentrations of a cytostatic drug is to arrest cells in the interphase compartment. It has been reported [20] that after 3 days of treatment with Iressa there is 100% cell arrest (cytostasis) in phase G_1. Its elimination half life is 48 h.

From half life data we can fit an exponential decay curve of the form $decay_v(t) = A\exp\{-\alpha t\}$. Given that at $t = 0$, $decay_v(0) = 1$, and that at $t = 3$ (48 h), $decay_v(2) = 1/2$, then it is easy to show that the decay function must be $decay_v(t) = \exp(\ln(\sqrt{2})t)$. Thus the ordinary differential equation that governs the dynamics for the concentration of the drug v is

$$\frac{dv}{dt} = -\ln(\sqrt{2})v + c_v(t) \tag{5}$$

where c_v is the amount of drug administered at time t. c_v is regarded in Eq. (5) as a continuous input.

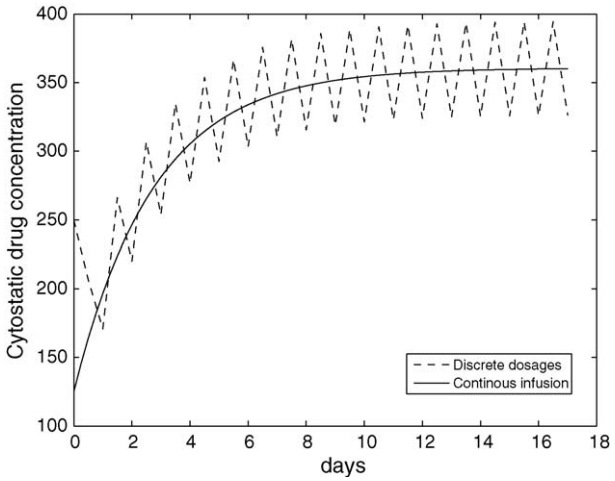

Fig. 1. A simulation of Iressa drug concentration levels using daily doses of 250 mg (dashed lines) and all day infusion of 125 mg (solid lines).

Iressa's commercial presentation consists of 250 mg pills. In order to model this presentation as a continuous variable, we adjusted the maximum drug concentration to follow the trend illustrated in Fig. 1. This figure represents a simulation of daily doses of Iressa. The infusion rate can be adjusted so that continuous infusion time intervals can be directly associated with a discrete dosage. It should be noted that the overall features of the graph in Fig. 1 coincide with other reports [21], in that the saturation levels are reached within 7–10 days. From [21] we know that after 3 days of administration a 100% cell arrest is observed, at concentration level of roughly 330 mg.

In the mathematical model, we need to express the effect of the drug with respect to concentration levels. A simple approach would be to consider the effect of the drug on the tumor level to be linear with respect to concentration levels. Given that the concentration level after 3 days is approx. 330 mg, we obtain that the effect of the drug with respect to drug concentration (v) is:

$$E(v) = \begin{cases} -\frac{1}{330}v + 1 & 0 \le v \le 330 \\ 0 & v > 330 \end{cases} \tag{6}$$

This linear relation is an approximation to concentration level curves encountered in published data [21] from human trials. Other functions were tested (logistic and exponential) to capture the main features of the data but we found the linear function to be a sufficiently good approximation.

Therefore the model equations, with respect to the original model (Section 2.2), are modified as follows to account for the incorporation of the drugs:

$$\begin{aligned} T'_I &= 2a_4T_M - (c_1I + d_2)T_I - \max\left(0, -\frac{1}{330}v + 1\right)a_1T_I(t-\tau) \\ T'_M &= \max\left(0, -\frac{1}{330}v + 1\right)a_1T_I(t-\tau) - d_3T_M - a_4T_M \\ &\quad - c_3T_MI - k_1(1 - e^{-k_2u})T_M \\ I' &= k + \frac{\rho I(T_I+T_M)^3}{\alpha + (T_I+T_M)^3} - c_2IT_I - c_4T_MI - d_1I - k_3(1 - e^{-k_4u})I \\ u'_1 &= -\lambda_1u_1 + c_u(t) \\ u'_2 &= -\lambda_2u_2 + c_u(t) \\ v' &= -\ln(\sqrt{2})v + c_v(t) \end{aligned} \tag{7}$$

The first and second equation of system (7) have a modified cycling term that is inspired from the model in [17]. This term reflects that

the proliferating fraction of cells depends on the concentration level of the cytostatic drug. The parameter estimation was performed on the drug free system [4]. The information available for paclitaxel in [18,22,23] was used for estimating the paclitaxel drug terms. The system was re-scaled using the same non-dimensionalization as in [4]. The maximum drug concentration after 3 days of treatment (330 mg) was used as the scaling factor for the drug concentration of v. The remaining model parameter values used are the same as in [5].

2.4. Optimal control problem

The problem of designing a chemotherapy protocol can be viewed as an optimal control problem. The goal is to eradicate the tumor while maintaining acceptable levels for the immune system. Mathematically this means using the drugs introduced into the system as control functions to drive the tumor into the tumor-free fixed point's basin of attraction, while keeping the immune system above a certain threshold.

The general optimal control problem can be stated as follows:

$$\begin{array}{ll} \text{Min} & J(T_I(t), T_M(t), I(t), u(t), v(t)) \\ \text{s.t} & \text{Equations in system (7)} \end{array} \quad (8)$$

Pontryagin's Maximum Principle was used to obtain the necessary conditions for an analytical solution to this problem [5]. It turned out that such a solution is prohibitive (as are also numerical solutions) which justifies the use of heuristic search algorithms (such as evolutionary algorithms) in our approach. The analysis also revealed that the problem is singular (the Hamiltonian's gradient does not provide information about the control when it is zero). This occurs when the controls appear linearly in the state equations [24]. In consequence, formulations of the objective function that do not have the control variable explicitly, will not change the singular property of the problem. Since the amount of drug (the control variable) in this formulation is bounded below and above, the candidate solutions are bang–bang, which means that the optimal control switches from one extreme to the other at certain times (i.e. is never strictly in between the bounds).

In [17], an optimal control problem was solved numerically for a three compartment system that modeled two drugs: one cytostatic and one cytotoxic. The solution considered a single treatment cycle and the optimal solution consisted of non-overlapping applications of first the cytostatic drug followed by the cytotoxic drug. This is intuitive, first the cytostatic drug is used to synchronize the cells and then they are targeted with a cytotoxic drug. One of the objectives of the present study is to verify that this non-overlapping strategy is conserved for more realistic systems and multiple cycle therapies.

3. Methods

3.1. The evolutionary algorithm

Evolutionary computing is the contemporary term given to a well established area within computer science and artificial intelligence [25]. As the name suggests, it draws inspiration from the process of natural evolution. The power of evolution is evident in the diverse species that make up our world, and their wonderful adaptations. The fundamental metaphor of evolutionary computing relates the power of natural evolution to a trial-and-error style of problem solving. The algorithms involved within evolutionary computing are termed *evolutionary algorithms*, of which several variants can be found in the literature. The *covariance matrix adaptation* evolution strategy (CMA-ES) is a state-of-the-art evolutionary algorithm for difficult non-linear non-convex optimization problems in continuous domain. The algorithm is typically applied to unconstrained or bounded constraint optimization problems, and search space dimensions between three and a hundred. This algorithm is, therefore, perfectly suited to tackle our optimal control problem in which the search space is composed of a set of real numbers representing drug application times (as described in the next subsection).

CMA-ES has proved to be a particularly reliable and and highly competitive evolutionary algorithm for both local optimization [26] and global optimization [27–29]. Moreover, it was the best performing approach in our previous studies of chemotherapy scheduling [5,6]. Evolution strategies are primarily based on mutation and selection as search operators [30]. As usual with evolutionary algorithms, the operators are applied in a loop for a certain number of iterations, until a termination condition is met. For real-valued search spaces, mutation is generally performed by adding a small normally distributed random value to each vector component. The step size or mutation strength (i.e. the standard deviation of the normal distribution) is often governed in evolution strategies by *self-adaptation*, which refers to a specific way of varying evolutionary parameters during a run, namely, the parameters are included in the chromosomes and evolve with the solutions. Individual step sizes for each coordinate, or correlations between coordinates, are either governed by self-adaptation or by covariance matrix adaptation [26]. Covariance matrix adaptation attempts to "derandomize" the process of mutation rate adjusting. Unlike previous methods, it is deterministic rather than based on mutation and selection. This scheme uses the path that has been followed by evolution so far to (i) adapt the step size, a scaling parameter that tunes the granularity of the search, by comparing the actual path length to that of a random walk; (ii) modify the covariance matrix of the multivariate Gaussian distribution in order to increase the likelihood of beneficial moves. A single Gaussian distribution is maintained, centered at a linear combination of the parameters. CMA-ES does not require tedious parameter tuning as it provides robust default parameters: the population size is set is set to $\lambda = 4 + \lfloor 3 \ln N \rfloor$ (where N is the problem size); and the initial step size to a third of the parameters range. The version used here (Matlab implementation [31]) incorporates weighted recombination [32], with weights $(w_i, \ldots, w_\mu)$ given by:

$$w_i = \ln \frac{\lambda + 1}{2} - \ln i$$

for $i = 1, \ldots, \mu$.

3.2. Problem encoding

In our optimal control formulation, the control functions are the amount of drug concentration introduced into the system as a function of time. These concentrations determine the scheduling and dosing of both drugs. Since solutions are *bang–bang*, the problem reduces to finding optimal switching times for both drugs. Following the encoding used in [5], where real numbers are used to encode the switching times (from "on" to "off" and viceversa), an extension of the coding is considered to incorporate the new drug in a similar way. Four types of control variables are, therefore, distinguished: administration-time lengths and resting-time lengths (measured in days) for both drugs. For each drug, these variables (application and resting times) are intercalated and concatenated. Two additional variables are incorporated to account for the number of cycles of each drug, *PU* and *PV*, which are integer numbers in the range from 6 to 12 cycles in accordance to previous results [6]. So this representation deals with a variable length encoding. Fig. 2 illustrates this encoding, which we termed *overlapping-permitting*, as it allows overlapping (simultaneous) applications of the cytostatic and cytotoxic drugs.

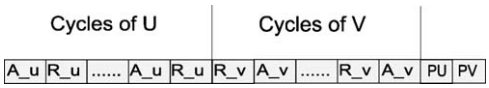

Fig. 2. Schematic view of a candidate solution (control variable). This first encoding termed *overlapping-permitting*, has two sections, the first section corresponds to the coding of the cytotoxic drug while the second to the cytostatic drug. The cytotoxic drug has cycles beginning with an application, while the cytostatic drug's cycle begins with a resting period. Application and resting times for both drugs are real numbers representing days. Two additional variables are incorporated to account for the number of cycles of each drug, *PU* and *PV*.

It is of interest, however, to find solutions that will have the same qualitative features as those found in [17]. For this purpose, a second encoding, that we termed *non-overlapping* representation was also considered, where the solutions always deliver non-overlapping treatment cycles, that is, the two drugs are never simultaneously applied. In this encoding, a treatment cycle consists of the application of the cytostatic drug, followed by the application of the cytotoxic drug, and a resting period. An additional variable, *P*, accounts for the number of cycles (allowed to vary from 6 to 12 cycles). Fig. 3 illustrates the non-overlapping representation. Notice that this encoding induces a smaller search space, since a single resting time for each cycle is encoded.

The resting times were set in the interval [0, 30], where 0 means that there is no resting period and the treatment continues; and 30 follows the current practice of standard chemotherapy schedule (i.e. infusions taking up to a week and a resting period of at least 3 weeks). The maximum tolerated dose for paclitaxel is 5 days of infusion at 30 mg/m^2/day, every 3 weeks [18] which imposes an upper bound for drug administration times. A lower limit of 3 h infusions is also a common practice when using paclitaxel. Thus, the range of values for application times was set as [0.2, 5]. In the case of Iressa, the bounds for the application times were set as: [1–5]. From [21] we know that after 3 days of administration a 100% cell arrest is observed. We used 5 days as an upper bound to provide greater scheduling variety, and to be able to achieve steady state concentration levels.

The course of treatment is simulated starting from a constant initial function outside the tumor-free basin of attraction. Specifically, the initial conditions were set as $T_I(0)$, $T_M(0)$, $I(0) = (1.3, 1.2, 0.9)$, where these values represent the populations of tumor cells (in interphase and mitosis) and immune system cells, respectively, all normalized by a factor of 10^6. This represents a patient with a tumor which cannot be controlled by her/his own immune system. Therefore, the goal is to apply the drugs to drive the tumor population inside the tumor-free basin of attraction (which in our simulations is represented by the point $(T_I^*, T_M^*) = (0.3, 0.3)$), while maintaining the immune system level above its initial value ($I_{thr} = 0.9$). See [5] for more details on the computation of the basin of attraction for the model.

3.3. Objective functions

The formulation of the objective function is crucial for setting up the optimal control problem. Our previous study of cytotoxic only chemotherapy [6], suggests relevant terms to be considered in the objective function. Some additional terms (such as those described fifth in the next list) were also explored for designing synchronized cytotoxic/cytostatic combination therapies:

1. $|T_I(t_f) - 0.3|$ and $|T_M(t_f) - 0.3|$, this term is present in all the objective functions considered since it states that we are minimizing the distance of the tumor from a desired state inside the basin of attraction. Here, $T_M(t_f)$ and $T_I(t_f)$ are the tumor level in mitosis and interphase at the end of the treatment.
2. The integral term, $(1/T_f)\int_0^{t_f}(T_M + T_I)\,dt$ accounts for the average tumor level throughout the course of treatment. This term may be important to prevent spikes for the tumor orbit which can compromise the patients' health. Such spikes were seen in [8], and the integral term was included in the initial formulation [5] to rule out this undesirable behavior.
3. The immune restriction expressed as the violation of the threshold imposed to the immune system, written mathematically as:

$$F(I) = \begin{cases} 0 & \text{if } I(t) > I_{thr} \\ I_{thr} - I(t) & \text{if } I(t) \le I_{thr} \end{cases} \tag{9}$$

where $I_{thr} = 0.9$. This term was included in all the formulations studied, as the immune system level accounts for the patient's health.4. The number of treatment cycles of the cytotoxic drug, *PU*, and the number of cycles of the cytostatic drug, *PV*, were also included in the minimization process. The idea behind these terms is to automatically find treatments with the lowest number of treatment cycles, which still achieve the desired goals. Shorter treatments would lessen the patient's burden.5. The sum of the treatment resting times for both the cytotoxic and the cytostatic drugs were also considered, *SRU* and *SRV*, respectively. The rationale behind these terms is the following: maximizing the resting time of the cytotoxic drug (*SRU*) will extend the time given to the patient to recover, while maximizing the resting time of the cytostatic drug (*SRV*) will extend the time given to the system to liberate the trapped cells.

An exploration of several combinations of these factors was carried out. However, to maintain clarity in the exposition we will focus the analysis on those objective functions that produced better chemotherapy treatments. All the components in the sums representing the objective functions described below, were considered equally, that is they were normalized to have magnitudes in similar ranges. The following three objective functions were considered:

Objective Function 1 (OF_1) is expressed mathematically as:

$$|T_I(t_f) - 0.3| + |T_M(t_f) - 0.3| + \frac{1}{T_f}\int_0^{T_f}(T_M + T_I)\,dt + F(I) + PU \tag{10}$$

Objective Function 2 (OF_2) is similar to OF_1, with the important difference that it does not incorporate the integral term. It is written mathematically as:

$$|T_I(t_f) - 0.3| + |T_M(t_f) - 0.3| + F(I) + PU \tag{11}$$

Objective Function 3 (OF_3) is similar to OF_1, but incorporates the term ($-SRU$). As mentioned earlier, *SRU* accounts for the total resting time of the cytotoxic drug during treatment. Maximizing this value will help in reducing the simultaneous application of the two drugs, when using the overlapping-permitting representation:

$$|T_I(t_f) - 0.3| + |T_M(t_f) - 0.3| + \frac{1}{T_f}\int_0^{T_f}(T_M + T_I)\,dt + F(I) - SRU \tag{12}$$

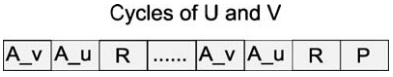

Fig. 3. Schematic view of a candidate solution (control variable). This second encoding, termed *non-overlapping* has cycles consisting of the application of the cytostatic drug, followed by the application of the cytotoxic drug, and a resting period. Application and resting times for both drugs are real numbers representing days. An additional variable, *P*, accounts for the number of treatment cycles.

3.4. Experimental setup

We conducted experiments for the two proposed encodings (overlapping-permitting and non-overlapping-permitting) and the three objective functions described above. The size of the search space depends on the encoding used (see Section 3.2). The first (overlapping-permitting) encoding consists of 50 numbers, of which the first 24 are application and resting times for the cytotoxic drug, the following 24 application and resting times for the cytostatic drug, and the last two correspond to the number of effective cycles of each drug. The second encoding consists of 37 numbers, of which the first 36 correspond to the application times and the resting times for the drugs, and the last, is the number of effective cycles of both drugs. Recall that the application and resting times are measured in days in the ranges discussed in Section 3.2, whereas the effective number of cycles is integers in the range 6–12.

The stopping criteria was set as a fixed number of iterations, specifically, 100 iterations. This value was selected after observing very little progression in the performance curves during a set of preliminary experiments in which 10 executions for each problem encoding were allowed to run for 300 iterations. The evolutionary algorithm parameters were set according to the suggestions and default values given by the CMA-ES Toolbox [31] (see also Section 3.1). Specifically, the number of offspring was set to $\lambda = 4 + \lfloor 3 \ln N \rfloor$ (where N is the problem size). In our case the problem sizes are $N = 50$ for the first encoding and $N = 37$ for the second, therefore, $\lambda = 15$ and $\lambda = 14$, were respectively used.

We ran 10 replicas for each combination objective function—problem encoding. Notice that the objective function 3 (OF_3) was combined with the overlapping-permitting representation only, as the intention was to reduce the amount of overlap in that case. We are aware that 10 replicas is a small number for statistical analysis purposes. In consequence, we rely on a study of typical runs, and stress the qualitative value of the results. It is worth mentioning that each function evaluation required the integration of a delay differential equation system for large periods of simulated time. Thus, a single replica of the evolutionary algorithms took in the order of 5 h to complete on an Intel Xeon CPU 3.00 GHz (2 processors). Performing any extensive statistical analysis of the results was, therefore, not practical on our current implementation. In order to asses how the results would vary for a larger number of runs, we conducted an additional set of 10 runs for one of the studied scenarios (the overlapping-permitting representation with the objective function 2). The conclusions and results (overall behavior) were maintained in this larger set.

The best solution at the end of each of the 10 runs (for each combination problem encoding—objective function) was stored, and boxplots were produced to illustrate the magnitude and distribution of the best objective values. Additionally, the following treatment quality measurements were considered:

Area under the tumor curve: given a solution vector, this measure calculates the sum of integrals under the state variables that represents the total amount of tumor cells during the course of treatment. A treatment schedule that minimizes the total amount of tumor cells is clearly preferable, therefore this magnitude is a direct measure of the efficiency of the treatment.

Immune system health: is the the average immune system's level, calculated with the difference:

$$ISH = \int_0^{T_f} I(t)\,dt - I_{thr} * T_f$$

It corresponds to the average deviation of the immune level, through the course of the entire treatment, with respect to the established threshold (I_{thr}). This measure accounts for the overall patients' health.

Both measures described above are embedded into the objective function formulation. However, they are reported separately to give more information about the treatments, and to allow a direct comparison among the alternative objective functions. Other quantities of interest that can help to describe the quality and structure of the solutions are:

Total application time for the cytotoxic drug: This magnitude corresponds to the area under the curve for the control variable, and measures the total amount of toxic drug in the system. In general it is desirable to eradicate the tumor giving the least possible amount of drug, as this would reduce the toxicity on healthy tissue.

Duration of treatment: This measure corresponds to the total length of the treatment in days. A reduction of this magnitude would represent both a more efficient treatment and a better quality of life for the patient.

In general, one should favor treatments with low values for the area under the tumor curve, total application time for the cytotoxic drug, and duration of treatment, while having high values for the immune system health.

4. Results

4.1. Comparing alternative encodings and objective functions

Since the studied objective functions have different terms, and thus produce different ranges of values, they cannot be directly compared. In order to illustrate the convergence behavior, and the spread of the solutions obtained with each objective function and problem encoding, Figs. 4–6 show, respectively, for each of function: (a) the convergence behaviors (trace) of a typical (median) and the best run and (b) the magnitude and distribution (boxplot) of the objective function values. Notice that for objective function 3 (Fig. 6), results are shown only for the overlapping-permitting representation. Recall that this function was especially designed to implicitly reduce the simultaneous application of the two drugs in this case. This is, of course, not necessary for the other representation, in which the non-overlapping application condition is imposed by the encoding. Figs. 4–6, suggest that 100 iterations are enough to obtain reasonably good convergence of the method. It can be noticed that the non-overlapping representation has a faster convergence (Figs. 4 and 5), which can be explained by the smaller search space induced by this encoding. These figures also show that the overlapping-permitting encoding can produce solutions with lower (better) objective fitness values, but with a much wider spread of values; while the non-overlapping encoding produces a more stable set of solutions (more noticeable for the objective function 2, Fig. 5).

The spread of solutions is also small for objective function 3, as can be seen when considering the small range of values in the vertical axis of Fig. 6(b). This small variation in the objective values suggests that the non-simultaneous application of the two drugs, enforced by OF_3 in the first representation and given by definition in the second, imposes a strong constraint on the design space of the drug schedules. This observation is indirectly supported by the distribution of the number of treatment cycles obtained by each function and encoding (Table 1). Recall that the number of cycles was allowed to vary between 6 and 12, and 10 runs were conducted for each combination of objective function and solution encoding. As indicated in Table 1, both OF_3 and the non-overlapping encoding, produced solutions with 12 cycles only; while OF_1 and OF_2 with the overlapping-permitting representation were able to produce shorter treatments with 10 and 11 cycles. This is done, however, at the expense of higher toxic drug doses

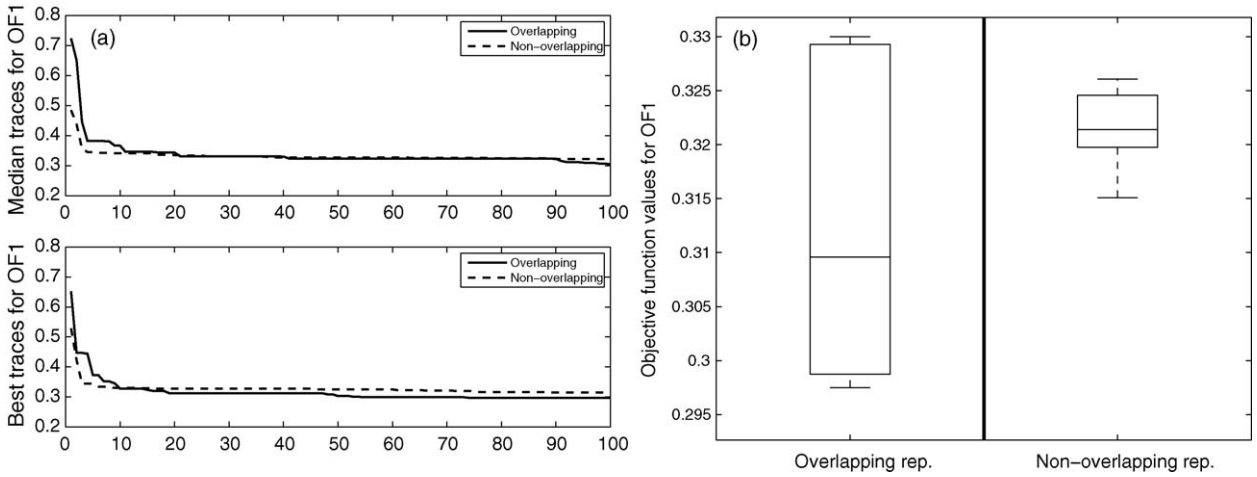

Fig. 4. Objective function 1 (OF_1): (a) typical (top) and best (bottom) trace and (b) distribution of objective function values at the end of the run. The two problem representations, *overlapping-permitting* and non-overlapping-permitting, are displayed.

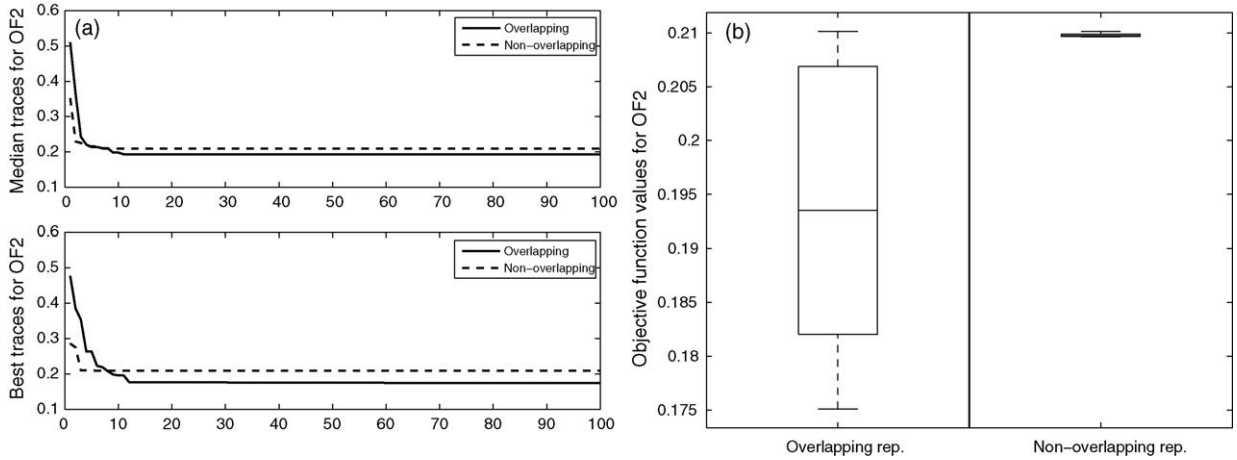

Fig. 5. Objective function 2 (OF_2): (a) typical (top) and best (bottom) trace and (b) distribution of objective function values at the end of the run. The two problem representations, *overlapping-permitting* and non-overlapping-permitting, are displayed.

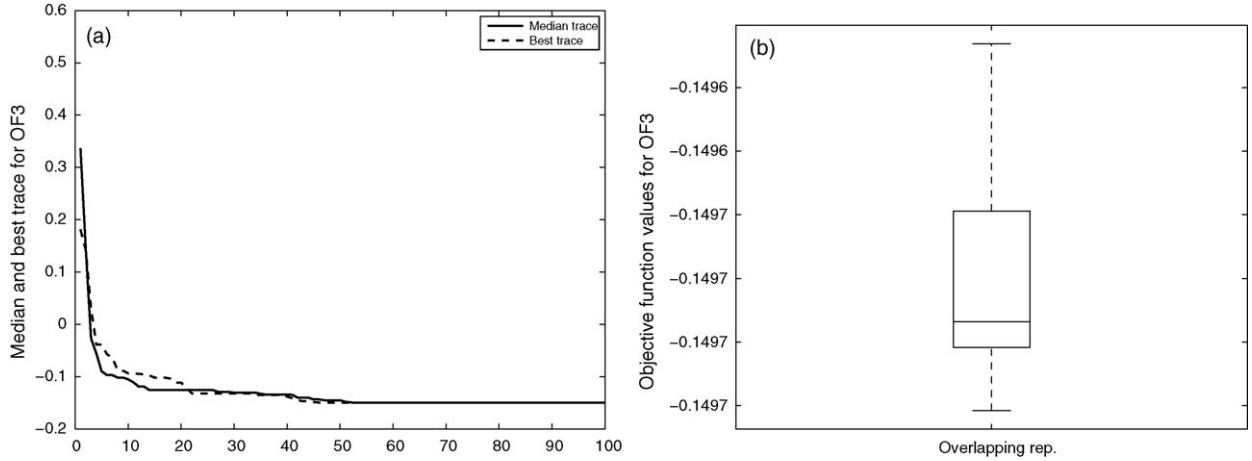

Fig. 6. Objective function 3 (OF_3): (a) typical (solid line) and best (dotted line) trace and (b) distribution of objective function values at the end of the run. Only the *overlapping-permitting* representation is used for this objective function.

Table 1
Distribution of the number of treatments cycles of the cytotoxic drug obtained with the different objective functions, and problem encodings.

Number of cycles	Overlapping			Non-overlapping	
	OF_1	OF_2	OF_3	OF_1	OF_2
10	5	4	0	0	0
11	2	4	0	0	0
12	3	2	10	10	10

Table 2
Overlapping times of the applications of the cytotoxic and cytostatic drugs for the three objective functions studied.

Overlapping	OF_1	OF_2	OF_3
Min	0.31	0	0
Max	12.1	13.96	10.29
Mean	4.99	4.74	3.16
Std	4.22	4.54	2.97

during treatment, as will be discussed below. It is worth noticing, also, that the number of cycles is correlated with the total duration of a treatment, but the exact duration in days also depends on the length of each cycle.

When inspecting the average overlapping times of the two drugs, produced when using the overlapping-permitting encoding in combination with the three objective functions (Table 2), we found that it is smaller when using the objective functions 2 and 3. Indeed, some of the solutions obtained with these functions, were non-overlapping. As expected, OF_3 produced the least overlapping.

Fig. 7 allows us to directly compare the quantitative features of the treatments obtained with each combination of objective function and solution encoding. With respect to the area under the tumor curve reduction (Fig. 7(a)), the two encodings produced comparative solutions when used in combination with objective functions 1 and 2. The objective function 1 (OF_1), shows the lowest tumor area at the end of treatment, which can be explained by the presence of the integral term in this function (see Section 3.3). This term enforces a minimization of the tumor level throughout the whole treatment.

The objective function 3, appears the least efficient from the three objective functions. A closer examination of the tumor dynamic across the treatment, reveals that the solutions obtained from this function (OF_3) decrease much slower than its consorts. This is illustrated in Fig. 8, where a comparison of the tumor dynamics for typical solutions obtained with OF_1 and OF_3, is presented. This slower decline towards zero leads to a greater value of the area under the tumor curve, regardless of the fact that the solution for OF_3 provides a smaller tumor burden at the end of the treatment. However, OF_3 achieves higher immune system health (part (b) in Fig. 7), at the expense of much longer treatment durations (part (d) in Fig. 7).

The most striking observation from Fig. 7, is that the solutions with non-overlapping drug applications achieve a similar efficacy

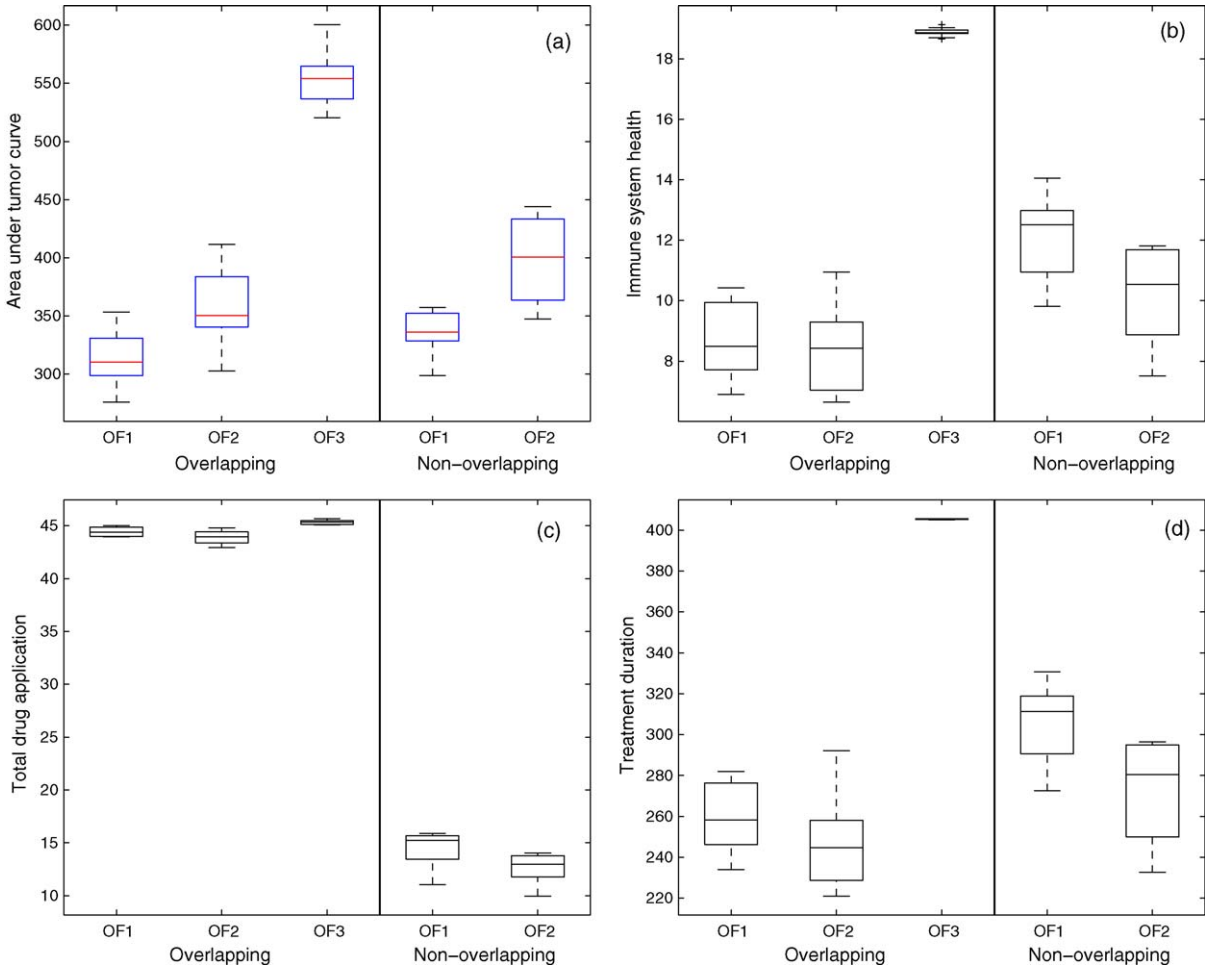

Fig. 7. Comparison of the three objective functions (OF), using both representations, with respect to: (a) area under the tumor, (b) immune system health, (c) total drug application and (d) total treatment duration.

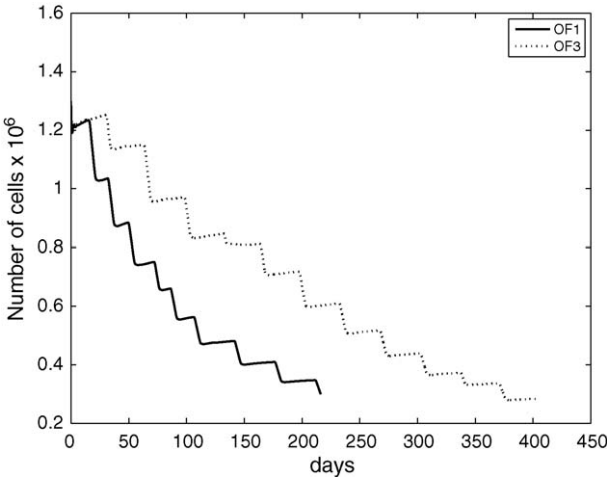

Fig. 8. Comparison of the system dynamics for tumor in interphase using the overlapping-permitting representation with OF_1 and OF_3.

and improved immune system health, when compared to their overlapping-permitting consorts, but with far less amount of toxic drug applied to the system. We suggest that this increased efficiency is due to the clear-cut synchronization between the two drugs. This is consistent with the strategy used in practice of first arrest and synchronize the cells, and then target them a cytotoxic agent, thus maximizing cell kill fraction and minimizing normal cell killing. We observed, however, that the non-overlapping-permitting representation produced treatments with a larger number of cycles, and longer duration in days (part (b) in Fig. 7). We suggest that this increase in treatment duration is also due to the strict synchronization (i.e. non-simultaneous application) of the two drugs.

4.2. Comparing system dynamics

Fig. 9 illustrates the dynamics of the system following good schedules obtained with the overlapping-permitting (a) and the non-overlapping representation (b). The *x*-axis indicates treatment duration in days. The horizontal thick lines in the upper graphs of the plots, represent the application times of the cytotoxic and the cytostatic drugs. The resting times are represented by the empty spaces in between drug applications. The left plot ((a) in Fig. 9) shows overlapping between the applications of the static and toxic drugs in some of the cycles (upper plot), whereas this is not the case in the upper portion of the right plot ((b) in Fig. 9), as expected. The lower portion of each graph describes the resulting dynamics for each of the state variables (T_I, T_M, and I) when the protocol shown on the upper portion is employed. Notice that the immune system (I) is always above its threshold value (its initial level) with small oscillations that correspond to the toxic drug application. The tumor levels in mitosis and interphase (T_I, T_M) are closely correlated and decrease steadily during the treatment progression. The figures suggest that both representations produce a similar tumor dynamics. The overlapping-permitting representation (plot (a) in Fig. 9) produces a shorter treatment of 220 days, while the alternative representation (plot (b) produces a longer treatment (300 days)). However, although not visually clear, the total amount of cytotoxic drug administered in (a) is 44.35, while in (b) is much smaller: 15.43.

4.3. Comparing cytotoxic versus combined cytotoxic/cytostatic treatments

One crucial aspect in this study is to asses the impact on the treatment efficiency of administering the two drugs in conjunction (cytotoxic and cytostatic) versus administering the cytotoxic drug only. Figs. 10 and 11 compare the dynamics of the tumor (specifically the tumor levels in interphase) using examples of good treatments obtained with the model incorporating only the cytotoxic drug [5,6], and the current model including also the cytostatic drug. In Fig. 10 the time frame was fixed to 120 days of treatment, and the dynamics is plotted according to the drug schedule up to that time. The curves clearly indicate the increased efficiency of the combined treatment, evidenced by the lower tumor levels.

In this case the amount of cytotoxic drug introduced with the combined therapy is 22.8 versus 18.5 administered during the same time frame using a single cytotoxic agent. Even though the schedule administers 18% less cytotoxic drug with the combined therapy, the difference in tumor levels is roughly 30% less tumor burden when the combined therapy is employed.

The treatments are different with and without the cytostatic agent. Therefore, up to any given time, the amount of toxic drug introduced into the system in each case is not necessarily the same. In Fig. 11, the system of equations was integrated until a maximum quantity of drug administered was reached (15 units), giving a qualitative idea of the increased efficacy with respect to drug dosing, corroborating the results above.

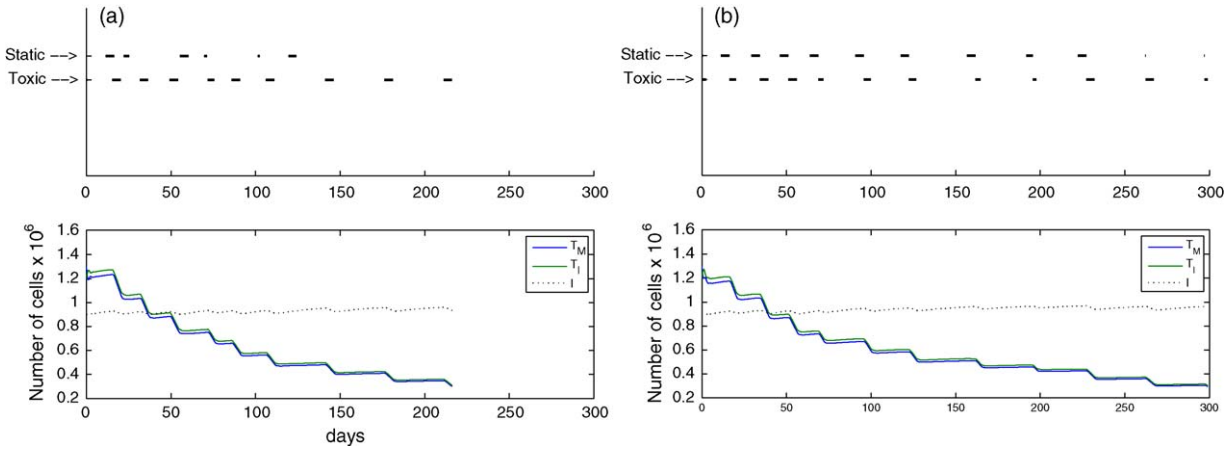

Fig. 9. Comparison of the system dynamics when minimizing the objective function 1. (a) Using the overlapping-permitting representation and (b) using the non-overlapping representation. The total amount of cytotoxic drug administered in (a) is 44.35, while in (b) is much smaller: 15.43.

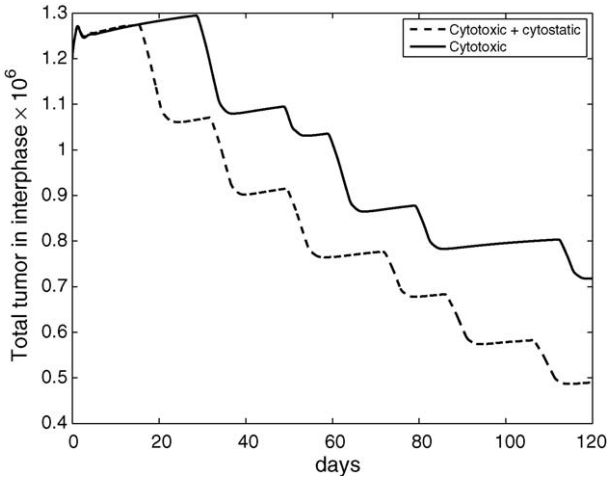

Fig. 10. Efficiency with respect to time. Comparing the efficiency of a combination treatment (cytostatic and cytotoxic drugs), versus a treatment with the cytostatic drug exclusively. The curves show the tumor level dynamics for a fixed time frame of 120 days.

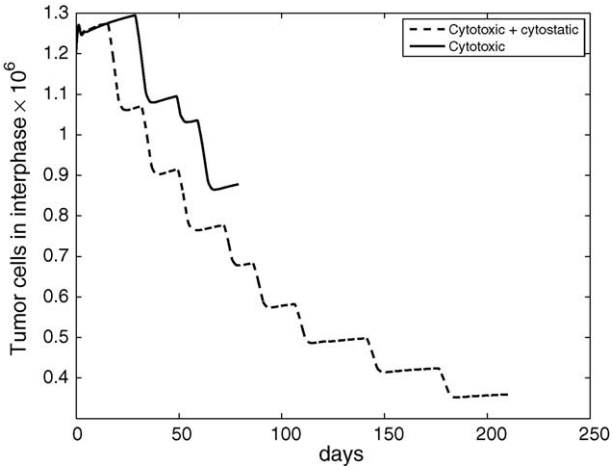

Fig. 11. Efficiency with respect to a fixed amount of drug. Comparing the efficiency of a combination treatment (cytostatic and cytotoxic drugs), versus a treatment with the cytostatic drug exclusively. The curves show the tumor level dynamics for a fixed amount of drug.

5. Discussion

An extension of a previous mathematical model of cancer cytotoxic chemotherapy is introduced by including a different type of drug, namely a cytostatic agent. The design of combined chemotherapies incorporating these two types of drugs was formulated as an optimal control problem, which was successfully solved by a state-of-the-art evolutionary algorithm for real-valued representation. Alternative objective functions and solution encodings were proposed and tested. Our main findings can be summarized as follows:

- The three objective functions studied for the overlapping-permitting representation provide competitive solutions, with different qualitative and quantitative features. The first two functions OF_1 and OF_2 are similar in the quality of the solutions, rendering some with non-simultaneous drug applications (or very small overlapping).
- The encoding that enforced non-overlapping (simultaneous) application of the two types of drugs, produced not only competitive protocols as compared to the overlapping-permitting encoding in terms of the tumor levels during treatment, but was able to do so with significant less amount of toxic drug, thus achieving better immune system health. Our results, therefore, confirm what was shown for more simplistic cancer models and a single cycle treatment [17].
- Finally, and most importantly, when compared to treatment protocols that only consider a cytotoxic agent, the incorporation of a cytostatic drug dramatically improved the outcome and performance of the overall treatment. Our model confirmed *in silico* what medical researchers have shown *in vivo* [19,3], that the combination of a cytostatic with a cytotoxic agent improves the efficacy and efficiency of the chemotherapy.

The proposed approach can be applied to drug combination chemotherapies, provided that the model is suitably identified for a specific type of cancer, patient characteristic and drugs' effect. Each patient has his or her own set of biological parameters. In order to tailor this model to a patient, one must first identify the specific tumor and drug kinetics and the patient's response to chemotherapy. These values are embedded in the model parameters as well as the initial conditions. Good parameter estimation is always a delicate issue given both ethical implications and lack of reliable methods. The tumor model parameters used in our numerical simulations can be considered as a case study. This example served, however, to illustrate the suitability of the proposed approach for designing chemotherapy schedules based on the two types of drug modeled. Although our approach has validated previous medical results with regards to combined cytotoxic and cytostatic drug therapies, clearly, further validation in clinical trials of the insights suggested by our approach would need to be carried out.

As future work, we are interested in exploring the performance of alternative evolutionary algorithm and heuristic search methods for solving the formulated optimal control problem. An additional line of research is to incorporate, into the tumor model, the effect of drug resistance observed in medical chemotherapy. The design of cancer chemotherapies may be considered as multi-objective problem, since there are at least two opposing goals: to eliminate the tumor cells as effectively as possible, while maintaining the patient health at an acceptable level. We have preliminary explored multi-objective formulations of our current model, but have obtained unsatisfactory results in terms of the quality of the resulting treatments. This line of research, although worth exploring, will probably require some significant reformulation of our current model, and the identification of efficient state-of-the-art algorithms for continuous multi-objective optimization.

Finally, we argue that mathematical and computational approaches such as the model proposed here, can become valuable decision support tools for the medical practitioner facing the complex and critical problem of designing and testing efficient combined chemotherapies, where an in-silica exploration can provide elements for understanding the characteristics of effective scheduling.

References

[1] Baxevanis CN, Perez SA, Papamichail M. Combinatorial treatments including vaccines, chemotherapy and monoclonal antibodies for cancer therapy. Cancer Immunology Immunotherapy 2009;58(3):317–24.
[2] Kumar S, Gao L, Yeagy B, Reid T. Virus combinations and chemotherapy for the treatment of human cancers. Current Opinion in Molecular Therapeutics 2008;10(4):371–9.
[3] Park JK, Lee S, Kang J, Nishio K, Saijo N, Kuh H. Synergistic interaction between gefitinib (iressa zd1839) and paclitaxel against human gastric carcinoma cells. Anticancer Drugs 2004;15(8):809–19.
[4] Villasana M, Radunskaya A. A delay differential equation model for tumor growth. Journal of Mathematical Biology 2003;47:270–94.

[5] Villasana M, Ochoa G. Heuristic design of cancer chemotherapies. IEEE Transactions on Evolutionary Computation 2004;8(6):513–21.
[6] Ochoa G, Villasana M, Burke E. An evolutionary approach to cancer chemotherapy. Genetic Programming and Evolvable Machines 2007;8(4):301–18.
[7] Panetta J, Adam J. A mathematical model of cycle-specific chemotherapy. Mathematical and Computer Modelling 1995;22(2):67–82.
[8] de Pillis L, Radunskaya A. A mathematical tumor model with immune resistance and drug therapy: an optimal control approach. Journal of Theoretical Medicine 2001;3:79–100.
[9] de Pillis L. The dynamics of an optimally controlled tumnor model: a case study. Mathematical and Computer Modeling 2003;37(11):1221–44.
[10] Petrovski A, McCall J. Multi-objective optimisation of cancer chemotherapy using evolutionary algorithms. In: Zitzler E, Deb K, Thiele L, Coello CAC, Corne D, editors. Evolutionary multi-criterion optimization, First international conference, EMO 2001, Proceedings, vol. 1993 of LNCS. Berlin: Springer; 2001. p. 531–45.
[11] Petrovski A, Sudha B, McCall J. Optimising cancer chemotherapy using particle swarm optimisation and genetic algorithms. In: Parallel problem solving from nature—PPSN VIII, vol. 3242 of LNCS. Berlin: Springer; 2004. p. 633–41.
[12] Petrovski A, Shakya S, McCall J. Optimising cancer chemotherapy using an estimation of distribution algorithm and genetic algorithms. In: Cattolico M, editor. Genetic and evolutionary computation conference GECCO 2006, vol. 1. New York: ACM Press; 2006. p. 413–8.
[13] McCall J, Petrovski A, Shakya S. Computational intelligence in bioinformatics. Hoboken, NJ: Wiley; 2007. p. 263–96, Chapter 15.
[14] Martin RB. Optimal control drug scheduling of cancer chemotherapy. Automatica 1992;28(6):1113–23.
[15] Tse S-M, Liang Y, Leung K-S, Lee K-H, Mok T-S-K. A memetic algorithm for multiple drugs cancer chemotherapy schedule optimization. IEEE Transcations on Systems Man and Cybernetics—Part B Cybernetics 2007;37(1):84–91.
[16] Liang Y, Leung K-S, Mok T-S-K. A novel evolutionary drug scheduling model in cancer chemotherapy. IEEE Transactions on Information Technology in Biomedicine 2006;10(2):237–45.
[17] Swierniak A, Polanski A, Kimmel M. Optimal control problem arising in cell-cycle-specific cancer chemotherapy. Cell Proliferation 1996;29:117–39.
[18] Hardman J, Limbird L. Goodman and Gilman's the pharmacological basis of therapeutics, 9th ed., New York: McGraw Hill, Health Profession Division; 1996.
[19] Ciardiello F, Caputo R, Bianco R, Damiano V, Pomatico G, De Placido S. Antitumor effect and potentiation of cytotoxic drugs activity in human cancer cells by zd-1839 (iressa), an epidermal growth factor receptor-selective tyrosine kinase inhibitor. Clinical Cancer Research 2000;6(5):2053–63.
[20] Piechocki MP, Yoo GH, Dibbley SK, Amjad EH, Lonardo F. Iressa induces cytostasis and augments fas-mediated apoptosis in acinic cell adenocarcinoma overexpressing her2/neu. International Journal of Cancer 2006;119:441–54.
[21] Shintani S, Li C, Mihara M, Yano J, Terakado N, Nakashiro K, et al. Gefitinib ('iressa', zd1839), an epidermal growth factor receptor tyrosine kinase inhibitor, up-regulates $p27^{KIP1}$ and induces g1 arrest in oral squamous cell carcinoma cell lines. Oral Oncology 2004;40:43–51.
[22] Zoli W, Flamigni A, Frassineti G, Bajorko P, De Paola F, Milandri CEA. In vitro activitity of taxol and toxotere in comparison with doxorubicin and cisplatin on primary cell cultures of human breast cancers. Breast Cancer Research and Treatment 1995;34:63–9.
[23] Chuang L, Lotzova E, Cook K, Cristoforoni P, Morris M, Wharton T. Effect of new investigational drug taxol on oncolytic activity and stimulation of human lymphocytes. Gynecologic Oncology 1993;49:291–8.
[24] Luus R. Iterative dynamic programming vol. 110 of monographs and surveys in pure and applied mathematics. Boca Raton: Chapman and Hall; 2000.
[25] Eiben A, Smith J. Introduction to evolutionary computing. Berlin: Springer-Verlag; 2003.
[26] Hansen N, Ostermeier A. Completely derandomized self-adaptation in evolution strategies. Evolutionary Computation 2001;9(2):159–95.
[27] Hansen N, Kern S. Evaluating the CMA evolution strategy on multimodal test functions. In: Yao X, et al., editors. Parallel problem solving from nature PPSN VIII vol 3242 of LNCS, Berlin: Springer; 2004. p. 282–91.
[28] Auger A, Hansen N. A restart CMA evolution strategy with increasing population size. In: McKay B, et al., editors. The 2005 IEEE international congress on evolutionary computation (CEC'05), vol. 2. Piscataway, NJ: IEEE Press; 2005. p. 1777–84.
[29] Auger A, Hansen N. Performance evaluation of an advanced local search evolutionary algorithm. In: McKay B, editor. The 2005 IEEE international congress on evolutionary computation (CEC'05), vol. 2. Piscataway, NJ: IEEE Press; 2005. p. 1769–76.
[30] Rechenberg I. Evolution strategy: optimization of technical systems by means of biological evolution. Stuttgart: Fromman-Holzboog; 1973.
[31] Hansen N. The CMA evolution strategy; 1996, Website, http://www.lri.fr/hansen/cmaesintro.html.
[32] Hansen N. The CMA evolution strategy: a comparing review. Vol. 192 of towards a new evolutionary computation. Advances on estimation of distribution algorithms. Berlin: Springer; 2006. p. 75–102, Chapter 4.