

An Evolutionary Approach to Cancer Chemotherapy Scheduling

Gabriela Ochoa* Minaya Villasana† Edmund K. Burke*

Abstract

In this paper, we investigate the employment of evolutionary algorithms as a search mechanism in a decision support system for designing chemotherapy schedules. Chemotherapy involves using powerful anti-cancer drugs to help eliminate cancerous cells and cure the condition. It is given in cycles of treatment alternating with rest periods to allow the body to recover from toxic side-effects. The number and duration of these cycles would depend on many factors, and the oncologist would schedule a treatment for each patient's condition. The design of a chemotherapy schedule can be formulated as an optimal control problem; using an underlying mathematical model of tumour growth (that considers interactions with the immune system and multiple applications of a cycle-phase-specific drug), the objective is to find effective drug schedules that help eradicate the tumour while maintaining the patient health's above an acceptable level. A detailed study on the effects of different objective functions, in the quality and diversity of the solutions, was performed. A term that keeps at a minimum the tumour levels throughout the course of treatment was found to produce more regular treatments, at the expense of imposing a higher strain on the patient's health, and reducing the diversity of the solutions. Moreover, when the number of cycles was incorporated in the problem encoding, and a parsimony pressure added to the objective function, shorter treatments were obtained than those initially found by trial and error.

keywords: evolutionary algorithms, evolution strategies, objective function, optimal control, cancer chemotherapy, cancer model, cycle-phase-specific drugs.

1 Introduction

Chemotherapy involves using anti-cancer drugs to help control or prevent the growth of cancerous tumours. A cell is considered cancerous when it has lost

*Automated Scheduling, Optimisation and Planning Group, School of Computer Science and IT, University of Nottingham, Nottingham NG8 1BB, UK. gxo@cs.nott.ac.uk, ekb@cs.nott.ac.uk

†Departamento de Cómputo Científico y Estadística, Universidad Simón Bolívar, AP 89000 Caracas 1081-A, Venezuela. mwillasa@usb.ve

its ability to regulate cell growth and division (mitosis). Thus, cancer consists of the rapid uncontrolled growth of malignant cells. The main objective in cancer chemotherapy is to kill the cancerous cells. Chemotherapy creates a damaging range of side-effects, and so it is normally given in cycles of treatment which alternate with rest periods, to allow the body to recover. Several cycles of treatment are needed, as chemotherapy only attacks cells that are actively dividing. At any one time, some cancer cells will be dormant, and may not be killed until a later round of drug treatment. The number and duration of these rounds depends on many factors including the type of cancer, how advanced it is, and the general health of the person being treated. An oncologist schedules the chemotherapy treatment for each person.

Considering the complexity of designing a schedule that achieves certain goals whilst moderating the cancer drug's toxic side-effects, the idea of providing computer-based decision support systems, is appealing. We propose evolutionary algorithms (EAs) as a search tool in a decision support system for designing chemotherapy schedules. Using an underlying mathematical model that captures the essential qualitative features of a cancer tumour, the purpose is to use the chemotherapy to control the system, and drive it into a desirable (minimal) tumour level after which the body could eliminate the remaining cancerous cells. This problem can be formulated as an optimisation problem, specifically an *optimal control problem* which refers to the problem of finding a control scheme for a given dynamical system such that a certain optimality criterion is achieved. The design of chemotherapy schedules has been formulated before from the point of view of optimal control [10, 3, 2], solving the stated optimisation 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 optimise chemotherapy schedules. Petrovski, McCall and colleagues, have extensively and successfully used EAs and other modern heuristics in this domain [11, 13, 12]. Villasana and Ochoa [16], compared the performance of three meta-heuristics (genetic algorithms, evolution strategies, and simulated annealing) in a similar problem. The main difference between the approaches of these two groups of authors, lies in the underlying mathematical model of tumour growth. Whilst Petrovsky et al. considered the Gompertz growth model with linear cell-loss effect [18], without including interactions with the immune system; Villasana et al. employed a more realistic cancer model [17], including the interactions between tumour cells and immune cells; and differentiating between cell phases for subsequent treatment with a cycle-phase-specific drug (see section 2). Another important difference lies in the number of drugs modelled, whilst Petrovsky et al. consider a combination of drugs, Villasana et al. model a single cycle-specific agent. Finally, the solution representation in both groups differs, with the former using binary encoding for values representing drug concentrations at discrete times, and the latter employing real valued vectors representing time lengths of application and resting periods.

The present study extends our previous contribution [16], by comparatively assessing the effects, on both the quality and diversity of the solutions, of dif-

ferent formulations of the objective function. We first considered the relative importance of maintaining low tumour levels versus assuring the patient health (measured by the level of immune cells); and, secondly, the inclusion of a term that keeps at a minimum the tumour levels throughout the course of treatment (in addition to the standard term that measures tumour levels at the end of the treatment). A third study modified both the problem encoding and the objective function by considering treatments with a variable number of cycles, with a parsimony pressure that favours shorter treatments.

The paper is organised as follows. Section 2 gives a brief background on chemotherapy and cycle-phase specific drugs. Section 3 describes the underlying mathematical model of cancer growth; highlighting its advantages and disadvantages as compared to other models in the literature. Thereafter, section 4 describes in detail the mathematical formulation of the problem, including three different studies that extended our previous formulation in [16]. The Methods section (section 5), describes relevant implementation issues such as the problem's encoding, formulation of the different objective functions, type and parameter settings of the evolutionary algorithm used, and the performance measures devised to gauge both the quality and diversity of the solutions. Section 6 presents and discusses our results, and finally section 7 summarises and further discusses our findings.

2 Biomedical Background

All chemotherapy drugs work by attacking cells that are dividing rapidly. Normal cells divide at a rate that is tightly controlled by the body. However, in cancer cells, the division goes wrong, leading to the uncontrolled production of new cells and the formation of a tumour 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. the immune system and the natural death of cells) will remove the remaining tumour cells.

Cycle-phase-specific drugs are those acting on a specific phase of the cell cycle, which is the process between two cell divisions or mitosis. The cell cycle 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.

An example of a cycle-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 tumour radiosensitivity. Today, Paclitaxel is used either as a single agent or accompanied by other drugs. The optimal scheduling of, and possible drug interactions, for Paclitaxel are not fully understood [6].

In medical practice, there are standard protocols and approved maximum dosages for known commercial drugs (provided by institutions such as the FDA

- *US Food and Drug Administration*). However, it is often the case that the oncologist would have to tailor the treatment according to the patient's characteristics and disease progression in a trial and error procedure.

3 Mathematical Model for Tumour Growth

Cancer is among the most common causes of death in the developed world. It is therefore not surprising that scientists around the world have been trying to accurately model the disease. The overall goal is to gain understanding of the disease, and thereby design better treatments to eradicate it, or at least to improve the patient's quality of life. Different types of models have been proposed, and each contributes in its own way to a better understanding of cancer dynamics.

The patient's model used in this work [17] is a competition model of tumour growth that includes the immune system response, and a cycle-phase-specific drug chemotherapy. The model considers three populations of cells: immune system, tumour during interphase (period comprising G_1 through G_2), and tumour during mitosis. Delay differential equations are used to take into account the phases of the cell cycle. Previously reported models, do not segregate the phases in which cells are vulnerable, instead the devised compartments usually comprise proliferating and non-proliferating cells. Moreover many models do not include the interactions with the immune system. Given that many cancer drugs are cycle-phase-specific, and the immune system plays a vital role in fighting the disease, we argue that a deeper understanding of efficient protocols can be achieved with a model that separates the cell stages and includes interactions with the immune system. The model has limitations at the moment, it considers treatments with a single cancer drug, whereas in medical practice it is common to use drug cocktails (infusions of various drugs during the treatment period). Another aspect not considered by our model, is the phenomenon of drug resistance often occurring in cancer cells. Nevertheless, the model is close enough to many situations, to mean that its study represents significant potential for deepening our understanding of this aggressive and often fatal disease.

In our underlying cancer model, $T_I(t)$ and $T_M(t)$ denote the population of tumour 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 [17] for a full discussion). Let $u(t)$ be the concentration of drug present at time t , and τ be the resident time of cells in interphase. The governing equations for the system with multiple applications of the drug are:

$$\begin{aligned}
T_I' &= 2a_4T_M - (c_1I + d_2)T_I - a_1T_I(t - \tau) \\
T_M' &= a_1T_I(t - \tau) - d_3T_M - a_4T_M - c_3T_MI \\
&\quad - k_1(1 - e^{-k_2u})T_M \\
I' &= k + \frac{\rho I(T_I + T_M)^3}{\alpha + (T_I + T_M)^3} - c_2IT_I - c_4T_MI \\
&\quad - d_1I - k_3(1 - e^{-k_4u})I \\
u_1' &= -\lambda_1u_1 + \mathbf{c}(t) \\
u_2' &= -\lambda_2u_2 + \mathbf{c}(t)
\end{aligned} \tag{1}$$

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

$$\begin{aligned}
T_I(t) &= \phi_1(t) & \text{for } t \in [-\tau, 0] \\
T_M(t) &= \phi_2(t) & \text{for } t \in [-\tau, 0] \\
I(t) &= \phi_3(t) & \text{for } t \in [-\tau, 0] \\
u_1(0) &= 0 \\
u_2(0) &= 0
\end{aligned}$$

The drug free system corresponding to model equations (1) can have up to 5 different fixed points depending on the parameter values (see [17]), one of which is always present, namely $(0, 0, k/d_1)$. This fixed point represents the desirable scenario of a tumour-free environment with a positive immune population.

Paclitaxel has a decay rate that can be modelled with two separate elimination terms: a fast decay rate while the drug is distributed through the blood to the tissues, and a second, slower rate in the peripheral compartment or tissue [14]. Thus the decay function is expressed as:

$$decay(t) = r_1e^{-\lambda_1t} + r_2e^{-\lambda_2t} \tag{2}$$

with r_1 , and r_2 representing real non dimensional constants.

Letting u_1 and u_2 be such that the concentration of drug at any given time is a linear convex combination represented by $u(t) = r_1u_1(t) + r_2u_2(t)$. The last two equations of system (1) model this situation with multiple drug applications in time, identified with the function $\mathbf{c}(t)$, which is the concentration of Paclitaxel that goes in the system at time t . With this choice and initial conditions we get,

$$u(t) = \mathbf{c}(t) * decay(t)$$

where $*$ denotes convolution.

Parameter estimation was performed on the drug free system [17], and the information available for Paclitaxel in [7, 19, 1] was used for estimating the drug terms. The system was then non-dimensionalised and scaled so model quantities are close to unity. Notice that the parameters will vary between tumour types and from patient to patient. The set of parameters used in this study, represents a patient with a rapidly growing tumour and an immune system not able to control the tumour progression, resulting in her/his eventual death if un-treated.

4 Problem Formulation

In the design of an effective chemotherapy, two conflicting objectives are at play:

1. To eradicate the tumour
2. To ensure that the chemotherapy side-effects are maintained at an acceptable level

Therefore, in our formulation, the goal is to design effective treatments with the single agent Paclitaxel on the model described in (model equations 1), so that the conflicting objectives mentioned above, are satisfied. In mathematical terms, the goal is to drive the dynamical system model inside the estimated basin of attraction of the tumour free fixed point, while maintaining the immune system population at an acceptable level.

In our initial formulation [16], the main objective was to minimize the average and final tumour size, and the patient's health was modeled as a restriction on the immune system's level. The problem was, therefore, stated as follows:

$$\begin{aligned} \mathbf{Min} \quad & T_I(t_f) + T_M(t_f) + \frac{1}{t_f} \int_0^{t_f} T_I(t) + T_M(t) dt \\ \mathbf{s.t} \quad & \text{Equations in system (1)} \end{aligned} \quad (3)$$

along with the added restriction:

$$I - \gamma I_{\text{thr}} \geq 0$$

Notice that there is no methodological way to determine the threshold imposed over the immune system. In practice, we want the patient to be as healthy as possible. In our experiments we required that the immune system does not fall below its initial state. The control function $\mathbf{c}(t)$ (that appears in the model equations 1) is the amount of drug introduced into the system as a function of time, determining the scheduling and dosing of the drug.

Pontryagin's Maximum Principle was used to obtain the necessary conditions for an analytical solution to this problem [16]. It turned out that such a solution is prohibitive (as are also numerical solutions) which justifies the use of meta-heuristics 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 [8]. 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).

Below, we describe three studies that extended our initial formulation, mainly by modifying the objective function. Further details and implementation are discussed in section 5.

First Study (OF_1) : The first objective function considers the tumour level deviation, at the end of treatment, from a desired level which is inside the tumor-free basin of attraction. The patient’s health is again modeled as a restriction on the immune system’s level. We assign weights to these contending objectives, and study the effects of different weight combinations (from a discretised grid of weights $w_i \in [0..1]$ and $\sum = 1$) on the quality of the solutions.

Second Study (OF_2) : The second objective function considers not only the tumour levels at the end of treatment, but also the average tumour level throughout the course of treatment, $\frac{1}{T_f} \int_0^{T_f} (T_M + T_I) dt$. This term may be important to prevent spikes for the tumor orbit which can compromise the patients’ health. Such spikes were seen in [3], and the integral term was included in our initial formulation [16] to rule out this undesirable behaviour. We study in detail, here, the effect of including such a term on the quality of the solutions.

Third Study (OF_3) : We modified both the problem encoding, by allowing treatments of a variable number of cycles, and the objective function by favouring shorter treatments. The idea behind this formulation is to automatically find the shortest possible treatment that achieves the desired goals. Shorter treatments would lessen the patient’s burden.

5 Methods

Since the control variable is the amount of drug administered, and solutions are bang-bang, the problem reduces to finding the times where the solution $\mathbf{c}(t)$ switches from “on” to “off”. That is, the times at which we begin and cease administering the drug. Each of these on-off switches constitutes a chemotherapy cycle. In order to admit variable time intervals, these switching times were encoded as real numbers. Two types of control variables are distinguished: administration-time lengths and resting-time lengths (measured in days). These variables are intercalated and concatenated to encode a potential solution to the problem of designing an effective chemotherapy (see figure 1).

The range of values for administration and resting times were set as follows. According to the literature, the maximum tolerated dose for Paclitaxel is 5 days of infusion at 30 mg/m²/day, every three weeks [7] which imposes an upper bound for drug administration times. A lower limit of 3 hour infusions is also a common practice when using Paclitaxel. Thus, the range of values for application-times was set to be $[0.2, 5]^1$. On the other hand, the resting-times were set in the interval $[0, 30]$ days , where 0 means that there is no resting period and the treatment continues, and 30 days (4 weeks) follows the current

¹Notice that 3 hours corresponds to 0.125 as a fraction of a day, however, we decided to round this value to 0.2, as 0.125 was found to be too small to provide visible differences in our simulations.

application time 1	resting time 1	application time 2	resting time 2	...	application time n	resting time n
-----------------------	-------------------	-----------------------	-------------------	-----	-----------------------	-------------------

Figure 1: Schematic view of a candidate solution (control variable). Both the application and resting times are real numbers representing days.

practice in a standard chemotherapy schedule (i.e, infusions taking up to a week and a resting period of at least 3 weeks). An external parameter, NC , indicates the number of treatment cycles.

The course of treatment is simulated starting from a constant initial function outside the tumour-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 tumour cells (in interphase and mitosis) and immune system cells, normalised by a factor of 10^6 . These values are taken as an example, and actually represent a specific patient² with a tumour which cannot be controlled by her/his own immune system. Therefore, the goal is to apply the drug to drive the tumour population inside the tumour-free basin of attraction (which in our simulations is given by $(T_I^*, T_M^*) = (0.3, 0.3)$), while maintaining the immune system level above its initial value ($I_{thr} = 0.9$).

5.1 Proposed Objective Functions

Three objective functions were considered:

Objective Function 1 (OF_1) : The equation $|T_M - T_M^*| + |T_I - T_I^*|$ was used for measuring the distance between the final tumour level to the desired target value. This term penalizes excursions, both upper and lower, from the desired level ($(T_I^*, T_M^*) = (0.3, 0.3)$). The objective function can, therefore, be formulated as the combination of the two goals: $J_1 = r_1(|T_M - T_M^*| + |T_I - T_I^*|) + r_2(\text{immune restriction})$, with r_1 , and r_2 real positive constants, such that $r_1 + r_2 = 1$. We wish to assess the relative importance of each term, thus we systematically explored the range of factors between $1/4$ and $3/4$ with a step of $1/4$, that is, pairs (r_1, r_2) , taken from the set $\{(1/4, 3/4), (1/2, 1/2), (3/4, 1/4)\}$. The set of weights $(0,1)$ and $(1,0)$ were not tested as they do not consider both goals simultaneously. We express the immune restriction as the violation of the threshold imposed on the immune system, which is written mathematically as:

$$\text{Immune Restriction} = \begin{cases} 0 & \text{if } I(t) > I_{thr} \\ I_{thr} - I(t) & \text{if } I(t) \leq I_{thr} \end{cases} \quad (4)$$

where $I_{thr} = 0.9$.

²According to the oncologists consulted during Dr. Villasana's doctorate degree (personal communication).

This study assumed a constant number of cycles, set to 12, which gives 24 switching times. Therefore, candidate solutions are vectors of 24 real numbers. We found, empirically, [16] that 12 treatment cycles were enough to drive the tumour towards the target value.

Objective Function 2 (OF_2) : The objective function was stated as $J_2 = r_1(|T_M - T_M^*| + |T_I - T_I^*|) + r_2(\text{immune restriction}) + \frac{1}{T_f} \int_0^{T_f} (T_M + T_I) dt$. The problem encoding and number of treatment cycles was set as in the first study.

Objective Function 3 (OF_3) : We incorporated, within the problem encoding, an additional integer variable representing the number of cycles. This parameter was allowed to vary in the range of 6 to 12. These values were set according to our simulations, since we noticed that under 6 cycles the system could not enter the basin of attraction, and less than 12 cycles were in general able to reach this basin. Therefore, restricting this range would reduce the search space and thus produce a faster search. An additional term was added to the objective function to penalise long treatments. Thus, the objective function used is the following:

$$J_3 = (|T_M - T_M^*| + |T_I - T_I^*|) + (\text{immune restriction}) + NC/k$$

where NC is the number of treatment cycles, and $k = 120$ is a constant selected to properly scale this term to the same order of magnitude of the remaining terms in the objective function. The appropriate order of magnitude is roughly $10 * NC$, since 12 cycles is the maximum. We use this scale as a reference.

5.2 An Evolutionary Algorithm

To solve the minimization problems described above, we selected the de-randomised Evolution Strategy (ES) with covariance matrix adaptation (CMA-ES)[5]. This is a state of the art evolutionary algorithm, that was found to have convergence velocity improvements over other evolutionary strategies on a large function optimization test suite [4]. Notice also that this was, by far, the best performing algorithm in our previous study [16], where we compared it with a Genetic Algorithm with real-number encoding, and a Simulated Annealing algorithm with various neighborhood operators. Moreover, the authors[5] provide a freeware, modular, and well documented Matlab implementation, with useful default settings for its strategy parameters. Specifically, the number of offspring λ , has a value of $\lambda = 4 + \lceil 3 \ln N \rceil$ (where N is the problem size, in our case 24); the number of parents, μ , is set to $\mu = \lfloor \lambda/2 \rfloor$; and the weights (w_i, \dots, w_μ) for weighted recombination, are given by:

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

for $i = 1 \dots, \mu$. We selected these default values. As the stopping criteria, a fixed number of iterations was set for each objective function individually after observing very little decrease in its evaluation through successive iterations. Notice the plateaus seen in figures 2, and 3, which also reveal the very fast speed of convergence of this algorithm.

5.3 Performance Measures

In order to assess the relative merits of the proposed objective functions, two sets of measures were devised. The first set gauges the quality of the solutions, whereas the second set measures the diversity of the best solutions across several runs. The interest of gauging the diversity of the obtained solutions, lies in the consideration that providing different solutions with similar quality, could be an advantage for a decision maker using our proposed automated system. In such a case the medical practitioner would have several treatment suggestions that he or she could assess according to other external factors not considered in the model.

We ran 10 experiments for each objective function. We are aware that 10 replicas is a small number for statistical analysis purposes. Thus 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 DDE (Delay Differential Equations) system for large periods of time. Thus, a single run of the evolutionary algorithm took in the order of three days to complete on a up to date PC (Pentium 4, 3.4 GHz). Performing any extensive statistical analysis of the results was, therefore, not feasible on our current implementation.

5.3.1 Quality Measures

The best individual at the end of each of the 10 runs was taken in order to report some simple statistics (mean, maximum, minimum, and standard deviation), of the measures described below:

Area under the Solution Curve (*AUC*) : given a solution vector, *AUC* is the integral under the control variable that represents the total amount of drug given during the course of treatment. This quantity is important because a treatment schedule that minimises the total amount of drug may be generally preferred.

Tumour Deviation from the desired level (*TD*) : is the quantity calculated as $|T_M - T_M^*| + |T_I - T_I^*|$, that is the amount of the tumour level deviation from its desired target at the end of treatment.

Immune System Health (*ISH*) : is the average immune system's level, calculated with the difference

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

this measure accounts for the average immune level through the course of the entire treatment minus the established threshold (I_{thr}). Notice that ISH gives the average deviation above the threshold, but it does not give information about possible immune population drops below I_{thr} .

In general, we should favour treatments with low AUC and TD values, and high ISH values.

5.3.2 Diversity Measures

In order to measure the diversity of the treatments obtained across several runs, we considered the best solution at the end of each of the 10 runs for each objective function, and calculated the following metrics:

Deviation from Average Best Solution (DAB) : is a vector containing the standard deviations of each of the treatment cycles (application and resting times), from the average best treatment values.

Moment of Inertia ($Inertia$) : proposed by Morrison and De Jong [9], is inspired by concepts from mechanical engineering, specifically on the moment of inertia which measures the mass distribution of an object. The centroid of a set of p points in a k -dimensional space has coordinates given by $c_i = \frac{\sum_{j=1}^p x_{i,j}}{p}$, for $i = 1, 2, \dots, k$, where $x_{i,j}$ is the i^{th} coordinate in the j^{th} point. The moment of inertia of the set of p points, is given by:

$$Inertia = \sum_{i=1}^k \sum_{j=1}^p (x_{i,j} - c_i)^2$$

The higher the values in the vector DAB , the higher the diversity of the set of best solutions. Similarly, the higher the value of I the higher the diversity of this set.

6 Results

Table 1 summarises the quality measures for the three objective functions explored. We found no observable differences between the measures of the three weight combinations ($r1, r2$) considered for OF_1 , which suggest that these weights do not greatly affect the portion of the search space being explored by this function. In consequence, and for the sake of simplicity, we report results for an equal distribution of weights (that is, $r1 = r2$). As for OF_3 , where the number of cycles of a treatment was also subject to evolution, solutions with 10, 11 and 12 cycles were obtained and reported.

The efficient convergence behaviour of the CMA-ES can be appreciated in Figures 2 and 3 for OF_1 (with different weight combinations) and OF_2 , respectively. Notice that the convergence for OF_2 is somewhat slower as compared

Table 1: Quality measures for OF_1 and OF_2 with 12 cycles, and for OF_3 with 10, 11, and 12 cycles. AUC represents the total amount of drug given, TD the tumour level deviation from the desired target, and ISH the immune system health. Treatments with low AUC and TD values, and high ISH values, are preferred.

		OF_1	OF_2	$OF_3(10)$	$OF_3(11)$	$OF_3(12)$
AUC	Mean	40.526	40.519	38.3579	39.3832	40.9365
	Max	41.202	40.601	40.5871	40.3464	43.3180
	Min	40.008	40.427	33.1596	35.4110	39.9468
	Std	0.370	0.056	2.6189	1.4673	1.0934
TD	Mean	0.009	0.010	0.0484	0.0209	0.0184
	Max	0.010	0.011	0.1744	0.1060	0.0739
	Min	0.008	0.010	0.0085	0.0087	0.0085
	Std	0.001	0.000	0.0592	0.0302	0.0245
ISH	Mean	9.2944	11.5093	8.4338	8.5121	9.0263
	Max	11.0767	12.5324	9.5778	10.8347	11.2354
	Min	7.6793	11.1424	7.3929	6.0307	7.1073
	Std	1.3467	0.4442	0.8147	1.7208	1.3229

to that of OF_1 , requiring more iterations to reach a plateau. Notice that these curves show the objective function values. This measure is not adequate to compare the solution's quality, since the weights have an effect on the mean value of the objective function, and this does not properly reflect the tumor levels at the end of treatment. Thus, although in Figure 2, it appears that the weight combination producing the best performance is $(r_1, r_2) = (1/4, 3/4)$, a closer look revealed that all the weight combinations produced similar values for the three quality measures considered (i.e. tumour levels at the end of treatment (TD), total amount of drug given (AUC), and immune system health (ISH)).

We see from Table 1, that there is no observable quantitative difference in AUC and TD between the chemotherapy schedules obtained with OF_1 and OF_2 (that incorporates the integral term). The main differences are seen in the qualitative features, and diversity measures of the schedules. Table 2 shows the diversity measures of the solutions obtained with OF_1 and OF_2 . The first row shows the moment of inertia ($Inertia$), and the remaining rows the standard deviations of each of the treatment cycles (application and resting times), from the average best treatment values. Notice that, the $Inertia$ value is lower for OF_2 . Moreover, across all the cycles the diversity in the schedules produced by OF_2 is much lower, reaching zero for most application periods. In both OF_1 and OF_2 , there is greater variability in the resting periods as compared to the application periods. These qualitative differences can be further appreciated in Figures 4 and 5, which illustrates the best obtained schedules with OF_1 and OF_2 , respectively (where best means lower tumour value at the end of treatment).

The treatment obtained with OF_2 shows a more regular pattern, with short resting periods of about 10 days in the first half of the cycle, and longer resting periods towards the end of the treatment. OF_2 stresses the importance of

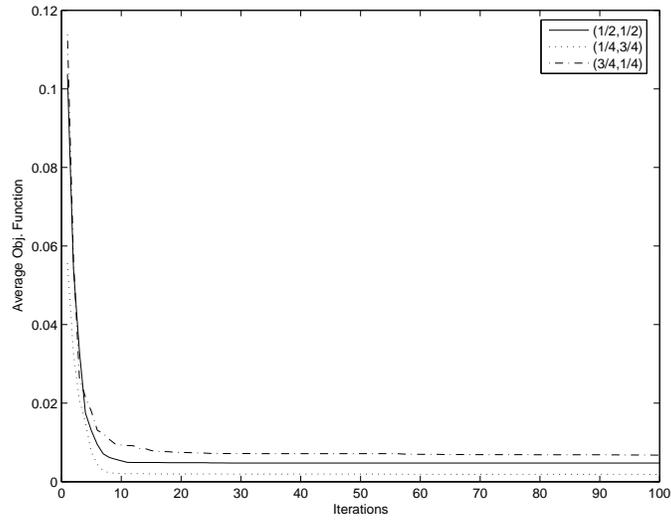


Figure 2: Average best performance of the algorithm (CMA-ES) using OF_1 for the three different combinations of weights considered.

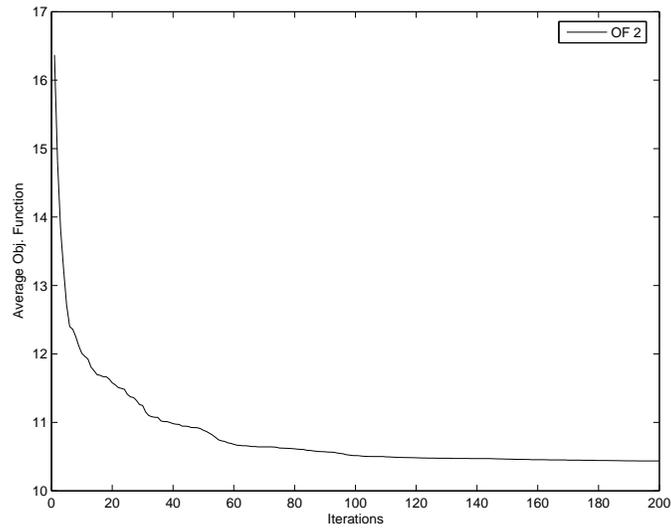


Figure 3: Average best performance of the algorithm (CMA-ES) using OF_2 (which includes the integral term).

Table 2: Diversity measures for OF_1 and OF_2 with 12 cycles. *Inertia* refers to the Moment of Inertia. *A* stands for Application periods, and *R* for resting periods.

	Period	OF_1	OF_2
<i>Inertia</i>		9498.3	437.4
Cycle 1	A	0.1929	0.1610
	R	5.2004	0.2387
Cycle 2	A	1.1581	0
	R	8.5673	0.2427
Cycle 3	A	1.3401	1.0443
	R	8.1281	0.2752
Cycle 4	A	1.0475	0
	R	8.7046	0.2599
Cycle 5	A	1.7431	0
	R	8.8508	0.4638
Cycle 6	A	0.6741	0.0014
	R	9.0996	6.3292
Cycle 7	A	1.1146	0
	R	9.6003	0.0284
Cycle 8	A	1.6767	0.0049
	R	7.8766	0
Cycle 9	A	0.5452	1.0535
	R	12.2226	0.0152
Cycle 10	A	1.8878	1.7673
	R	10.1883	0
Cycle 11	A	1.0541	1.9505
	R	10.7173	0
Cycle 12	A	1.4727	0
	R	10.5228	0

minimizing the tumour levels from the beginning of the treatment, whilst OF_1 does not reinforce this behaviour. We hypothesize that this relaxation on initial minimization explains the greater variability in the solutions found when using OF_1 . Thus, while OF_2 produces more regular treatments; OF_1 has the potential of producing a variety of schedules which increases the options available to the oncologist when designing a chemotherapy schedule. Notice that having regular patterns of treatments (with fixed resting and application times) could also be considered to be advantageous from the patient's point of view, given logistic and personal circumstances.

In our experiments, we did not observe the spikes of the tumour orbit reported in [3], that can compromise the patients' health. Thus, the necessity of the integral term can be questioned. Furthermore, when we analysed the patient outcome under the best treatments, the one obtained with OF_1 is, on

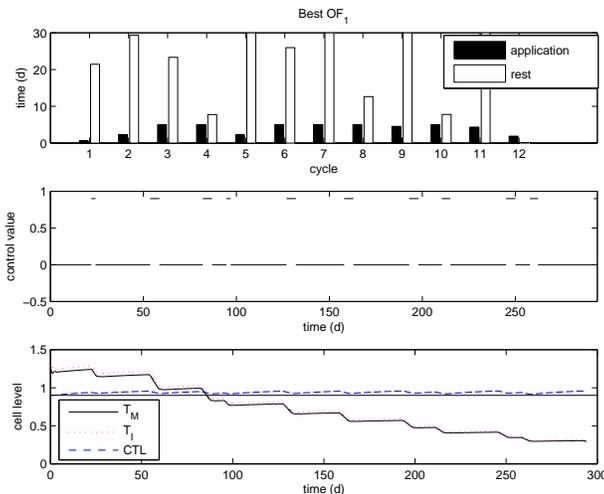


Figure 4: The best chemotherapy schedule obtained with OF_1 . (a) Application and resting times for each cycle. (b) Application times (horizontal lines at level 1), and resting times (horizontal lines at level 0). (c) Behavior of the dynamical system across the treatment time: T_M = tumour level in mitosis, T_I = tumour level in interphase, and CTL = immune system level.

average, more considerate of the immune system, in the sense that the immune level fluctuates around a higher level than the initial state. Meanwhile the best treatment obtained with OF_2 is more severe with the immune system at the initial stages of the treatment, even minimally violating the immune restriction (see Figure 6). However, the values for ISH in table 1 are higher than those for OF_1 , but this measure reflects the average immune system dynamics, hiding its particular features and possible violations of the imposed restriction. We conjecture that the need for the integral term for chemotherapy scheduling is dependent on the underlying model dynamics. For the mathematical model in [3], Jeff’s phenomenon was explained through the appearance of these spikes in the solutions for their proposed optimal control formulation (which was also singular). The mathematical model used here explains such a phenomenon as instability with respect to the delay parameter [15, 17].

With respect to our third study, OF_3 , the intention was to let the evolutionary algorithm design not only the course of treatment for a pre-fixed number of cycles, but also to provide the minimum number of cycles needed to achieve the desired goals, which would lower the patient’s burden by minimizing the treatment length. In this study, the number of cycles was encoded as an additional “gene” at the end of the chromosomal representation, and this parameter was allowed to vary in the range of 6 to 12. We conducted enough simulations, to accumulate at least 10 experimental results for each of the number of cycles predicted by the algorithms (10, 11, and 12). On average, we observed that 30%

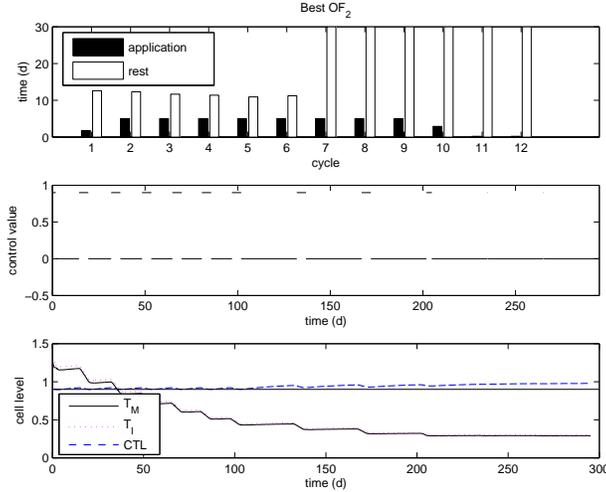


Figure 5: The best chemotherapy schedule obtained with OF_2 (including the integral term). (a) Application and resting times for each cycle. (b) Application times (horizontal lines at level 1), and resting times (horizontal lines at level 0). (c) Behaviour of the dynamical system across the treatment time: T_M = tumour level in mitosis, T_I = tumour level in interphase, and CTL = immune system level.

of the experiments resulted in treatments consisting of 10 cycles, 45% of treatments had 11 cycles, and 25% were 12 cycles long. Therefore, approximately 75% of all outcomes predicted shorter treatment schedules than previously used. In this experiment it makes sense to monitor the health of the immune system, because shorter treatments can take a toll on the body's ability to sustain them.

The quality measures obtained with OF_3 are outlined in table 1. Notice that, as expected, the lower the number of cycles the lower the amount of drug is used by the schedule (AUC). All the obtained solutions reached tumour levels inside the desired basin of attraction. Moreover, the immune system restriction was not violated throughout the course of the simulated treatments with 11 cycles, while small movements below the $I_{thr} = 0.9$ level were observed in one experiment with 12 cycles, and another with 10 cycles. Comparing these results with those of the OF_1 , we observe that there is no observable quantitative differences. This shows that it is possible to attain the desired goals with less treatment cycles than those previously used. This would imply a treatment reduction of approximately 6 months. The ISH values obtained for the experiments with 12 cycles are similar to those obtained with OF_1 and OF_2 , because the number of cycles are the same in each case. Meanwhile, this measure is lower for experiments with 10 and 11 cycles. This is due to a reduction in treatment times, because the amount of drug needed to drive the tumour to the desired level seems to be roughly around 40 (see AUC values in table 1). So, there is less time for the immune system to recover with these shorter treatments. The

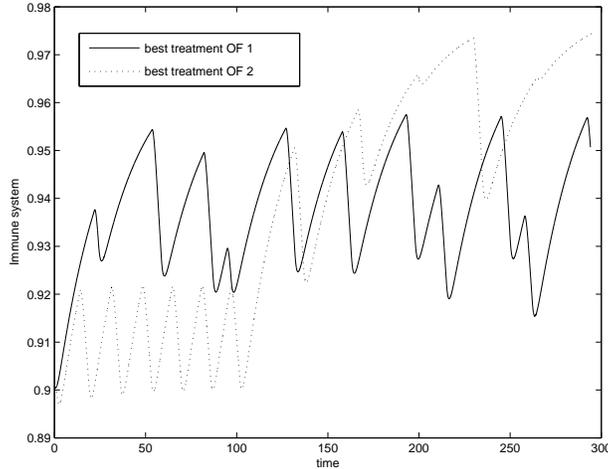


Figure 6: Immune system performance when minimizing with OF_1 and OF_2 .

diversity of the obtained solutions, was similar to that observed for OF_1 . Thus, for the sake of brevity we report the *Inertia* measures only (Table 3). Notice that it has the same order of magnitude as OF_1 , and greater diversity for OF_3 with 11 cycles.

Table 3: Diversity measure (*Inertia*) for OF_3 and solutions with 10, 11, and 11 cycles

	10 Cycles	11 Cycles	12 Cycles
<i>Inertia</i>	7526.1	7722.2	5481.8

7 Discussion and Conclusions

We have studied the effects, on the quality and diversity of the solutions, of different objective functions in an optimal control formulation of cancer chemotherapy. A highly competent evolutionary algorithm (CMA-ES) was used to address the formulated search problem, and the patient’s dynamic was simulated through a mathematical model of tumour growth that includes interactions with the immune system and multiple applications of a single cycle-phase specific drug. The goal of the chemotherapy is to eradicate the tumour, while maintaining the drug side-effects above an acceptable level. These conflicting objectives were captured in a single-objective function (to be minimised) with several terms: (a) the deviation of the tumour level at the end of treatment, from a desired (low) level, (b) the tumour level throughout the course of treatment, and (c) the amount of violation to a threshold imposed on the immune

system. It is worth noticing that, since this threshold on the immune system reflects the patient's state of health, treatments with different severities can be obtained by modulating this value. The effect of these terms on the quality and characteristics of the treatments produced, was carefully analysed. We found that the relative weights of terms (a) reaching tumour level close to the desired target, and (c) securing immune system level above the pre-established threshold; did not produce an observable effect on the quality and features of the obtained solutions. On the other hand, including a term that considers the tumour level throughout the course of treatment (b), produced treatments with similar quality measures, but with different features (for example, a more regular treatment pattern). This term was also found to double the computational time, and drastically decrease the variability of the solutions obtained. Thus, the requirement of this term in a formulation depends on the underlying model dynamics, and treatment goals. Finally, when the number of treatments cycles was incorporated in the problem encoding, and a parsimony pressure was included to the objective functions, the proposed approach obtained shorter treatments (with a lower number of treatment cycles) than those initially found by a trial and error procedure.

This study testifies that the outcome from a computational tool supporting the design of cancer chemotherapy schedules, greatly depends on the formulation of the desired treatment goals, and the modelling of patient dynamics. This confirms that these systems are in no way a substitute for the practitioner but rather a decision support tool at their disposal. However, the potential versatility of such decision support systems, serving as test-beds for newly discovered drugs and able to be tailored to each patient needs, should encourage their improvement. Motivated by this line of thinking we are currently exploring a multi-objective formulation that would produce, not a single outcome, but a set of alternative treatments, leaving the final decision to the practitioner.

References

- [1] L. Chuang, E. Lotzova, K. Cook, P. Cristoforoni, M. Morris, and T. Wharton, *Effect of new investigational drug taxol on oncolytic activity and stimulation of human lymphocytes*, *Gynecologic Oncology* **49** (1993), 291–298.
- [2] L. de Pillis and Radunskaya A., *The dynamics of an optimally controlled tumor model: A case study*, *Mathematical and Computer Modeling* **37** (2003), no. 11, 1221–1244.
- [3] L. de Pillis and A. Radunskaya, *A mathematical tumor model with immune resistance and drug therapy: An optimal control approach*, *Journal of Theoretical Medicine* **3** (2001), 79–100.
- [4] N. Hansen and A. Ostermeier, *Completely derandomized self-adaptation in evolution strategies*, *Evolutionary Computation* **9** (2001), no. 2, 159–195.
- [5] Nikolaus Hansen and Andreas Ostermeier, *Adapting arbitrary normal mutation distributions in evolution strategies: the covariance matrix adaptation*, *Proc. of the 1996 IEEE Int. Conf. on Evolutionary Computation* (Piscataway, NJ), IEEE Service Center, 1996, pp. 312–317.

- [6] J. Hardman and L. Linbird, *Goodman and Gilman's the pharmacological basis of therapeutics*, 9th ed., McGraw-Hill, New York, 1996.
- [7] J. Hardman and L. Linbird, *Goodman and Gilman's the pharmacological basis of therapeutics*, 9 ed., Mc Graw Hill, Health Profesion Division, 1996.
- [8] R. Luus, *Iterative dynamic programming*, Monographs and Surveys in Pure and Applied Mathematics, vol. 110, Chapman and Hall, 2000.
- [9] R.W. Morrison and K.A. De Jong, *Measurement of population diversity.*, Artificial Evolution, 2001, pp. 31–41.
- [10] J.C. Panetta and J. Adam, *A mathematical model of cycle-specific chemotherapy*, Mathl. Compt. Modeling **22** (1995), no. 2, 67–82.
- [11] A. Petrovski and J. McCall, *Multi-objective Optimisation of Cancer Chemotherapy Using Evolutionary Algorithms*, First International Conference on Evolutionary Multi-Criterion Optimization, Springer-Verlag. Lecture Notes in Computer Science No. 1993, 2001, pp. 531–545.
- [12] A. Petrovski, S. Shakya, and John McCall, *Optimising cancer chemotherapy using an estimation of distribution algorithm and genetic algorithms*, GECCO 2006: Proceedings of the 8th annual conference on Genetic and evolutionary computation, vol. 1, ACM Press, 2006, pp. 413–418.
- [13] A. Petrovski, B. Sudha, and J. McCall, *Optimising cancer chemotherapy using particle swarm optimisation and genetic algorithms*, Parallel Problem Solving from Nature - PPSN VIII (Birmingham, UK), LNCS, vol. 3242, Springer-Verlag, 2004, pp. 633–641.
- [14] J. Sear, *Clinics in anaesthesiology*, vol. 2, Intravenous Anaesthesiology, no. 1, ch. 12, p. 223, W. B. Saunders Company, 1984.
- [15] M. Villasana, *A delay differential equation model for tumor-immune system interactions*, Ph.D. thesis, Claremont Graduate University, 2001.
- [16] M. Villasana and G. Ochoa, *Heuristic design of cancer chemotherapy*, IEEE trans. Evolutionary Computation **8** (2004), no. 6, 513–521.
- [17] M. Villasana and A. Radunskaya, *A delay differential equation model for tumor growth*, Journal of Mathematical Biology **47** (2003), 270–294.
- [18] T. Wheldon, *Mathematical models in cancer research*, Adam Hilger, Bristol Philadelphia, 1998.
- [19] W. Zoli, A. Flamigni, G. Frassinetti, P. Bajorko, F. De Paola, C. Milandri, D. Amadori, and A. Gasperi-Campani, *In vitro activity of taxol and taxotere in comparison with doxorubicin and cisplatin on primary cell cultures of human breast cancers*, Breast Cancer Research and Treatment **34** (1995), 63–69.