

Multi-objective Optimisation of Cancer Chemotherapy Using Evolutionary Algorithms

Andrei Petrovski and John McCall

School of Computer and Mathematical Sciences, The Robert Gordon University,
St. Andrew Street, ABERDEEN, AB25 1HG, United Kingdom
{ap, jm}@scms.rgu.ac.uk

Abstract. The main objectives of cancer treatment in general, and of cancer chemotherapy in particular, are to eradicate the tumour and to prolong the patient survival time. Traditionally, treatments are optimised with only one objective in mind. As a result of this, a particular patient may be treated in the wrong way if the decision about the most appropriate treatment objective was inadequate. To partially alleviate this problem, we show in this paper how the multi-objective approach to chemotherapy optimisation can be used. This approach provides the oncologist with versatile treatment strategies that can be applied in ambiguous cases. However, the conflicting nature of treatment objectives and the non-linearity of some of the constraints imposed on treatment schedules make it difficult to utilise traditional methods of multi-objective optimisation. Evolutionary Algorithms (EA), on the other hand, are often seen as the most suitable method for tackling the problems exhibiting such characteristics. Our present study proves this to be true and shows that EA are capable of finding solutions undetectable by other optimisation techniques.

1 Introduction

Cancer chemotherapy is a highly complex process which controls tumour development by the administration of a cocktail of chemicals in a series of doses over a course of treatment. There is a wide variety of anti-cancer drugs available to oncologists. Due to their high toxicity, these drugs give rise to a variety of side-effects, ranging from cosmetically undesirable through debilitating through to the effects that are themselves life threatening. The oncologist therefore is faced with a complex task of designing a therapy which achieves certain treatment goals whilst limiting the toxic side-effects of the treatment to an acceptable level.

In the treatment of most common cancers multi-drug combinations are usually used. Traditionally, combination treatments are developed through empirical trials of different combinations, dosing, schedules and sequencing. However, since around 35 drugs are in common clinical use nowadays [17], it is evident that an almost infinite number of treatment schedules are conceivable and that the need for the optimisation of chemotherapeutic treatment is indisputable. The number of combinatorial possibilities for multi-drug schedules, coupled to the conflicting nature [13] and non-linearity of the constraints imposed on cancer treatments, make it difficult to solve the

problem of cancer chemotherapy optimisation by means of empirical clinical experimentation or by means of traditional optimisation methods [6]. An alternative approach is to use evolutionary methods of computational optimisation to search for multi-drug treatment schedules that achieve certain treatment objectives and satisfy a number of simultaneous constraints.

A body of work has been established by the authors [7], [8], [10] and [11], where they have applied Genetic Algorithms to find the best (or at least suitable) treatment strategies given a single optimisation objective. In this paper, however, we endeavour to develop this approach further and to address the problem of finding treatment strategies that show a good performance with respect to more than one treatment objective. Thus, the evaluation of different treatment strategies will involve multiple measures (objectives) of performance, which should be optimised simultaneously, even though they may be conflicting in nature. The presence of conflicting objectives gives rise to a set of optimal solutions, known as the Pareto-optimal set. If all objectives are equally important, the conflict between them requires a compromise to be reached. A good solution to such problems involving conflicting objectives and therefore multiple evaluation criteria, should offer suitable, though possibly sub-optimal in the single-objective sense, performance in all objective dimensions [14]. Generally, there exists a multitude of such solutions; hence, the algorithm used to solve a multi-objective optimisation problem should find a wide variety of them, instead of just one.

Evolutionary Algorithms (EA) are a promising choice for solving the multi-objective optimisation problem of cancer chemotherapy for a number of reasons. Firstly, a set of Pareto-optimal solutions can, in principle, be captured in an EA population, thereby approximating the Pareto-optimal set in a single simulation run [2]. Secondly, in general Evolutionary Algorithms are less susceptible to the shape or continuity of the Pareto front than other techniques of multi-objective optimisation [16]. Thirdly, it has been shown by the authors (see [10] and [11]) that the problem of optimising cancer chemotherapy treatment belongs to the class of complex optimisation problems involving such features as discontinuity, multi-modality, non-connected, non-convex feasible regions, and inaccuracy in establishing model parameters. This is precisely the problem area where the methods of evolutionary computation really distinguish themselves from their competitors, thereby reinforcing the potential effectiveness of Evolutionary Algorithms in multi-objective optimisation of chemotherapeutic treatment.

The remaining sections are organised as follows. In section 2 we provide the background information on optimisation of chemotherapeutic treatment, which includes medical aspects of chemotherapy, the formulation of treatment design as a constrained multi-objective optimisation problem, and a description of salient features of Evolutionary Algorithms used in multi-objective optimisation. Section 3 explains implementation details of the evolutionary search for Pareto-optimal treatment schedules. The results of chemotherapy optimisation and their analysis are given in Section 4. Finally, Section 5 summarises the contribution of the present study to cancer chemotherapy and outlines possible directions for its further development.

2 Optimisation of Chemotherapeutic Treatment

Amongst the modalities of cancer treatment, chemotherapy is often considered as inherently the most complex [17]. As a consequence of this, it is extremely difficult to find effective chemotherapy treatments without a systematic approach. In order to realise such an approach, we need to take into account the medical aspects of cancer treatment.

2.1 Medical Aspects of Chemotherapy

Drugs used in cancer chemotherapy all have narrow therapeutic indices. This means that the dose levels at which these drugs significantly affect a tumor are close to those levels at which unacceptable toxic side-effects occur. Therefore, more effective treatments result from balancing the beneficial and adverse effects of a combination of different drugs, administered at various dosages over a treatment period.

The beneficial effects of cancer chemotherapy correspond to treatment objectives which oncologists want to achieve by means of administering anti-cancer drugs. A cancer chemotherapy treatment may be either curative or palliative. Curative treatments attempt to eradicate the tumour. It is believed that chemotherapy alone cannot eradicate cancer, but if the overall tumour burden is held below a certain level, other mechanisms (e.g. immune system or programmed cell death) will remove remaining tumour cells. Palliative treatments, on the other hand, are applied only when a tumour is deemed to be incurable. Here the objective is to maintain a reasonable quality of life for as long as possible.

The adverse effects of cancer chemotherapy stem from the systemic nature of this treatment: drugs are delivered via the bloodstream and therefore affect all body tissues. Since most anti-cancer drugs are highly toxic, they inevitably cause damage to sensitive tissues elsewhere in the body. In order to limit this damage, toxicity constraints need to be placed on the amount of drug applied at any time interval, on the cumulative drug dosage over the treatment period, and on the damage caused to various sensitive tissues [17]. In addition to toxicity constraints, the tumour size (i.e. the number of cancerous cells) must be maintained below a lethal level during the whole treatment period for obvious reasons.

The goal of cancer chemotherapy therefore is to achieve the beneficial effects of treatment objectives without violating any of the abovementioned constraints. This problem would not be much different from that of a general class of constrained optimisation problems, was it not for the conflict between treatment objectives. The objectives of curative and palliative treatments conflict with each other in the sense that drug schedules which tend to minimise tumour size are highly toxic and therefore have a negative effect on the quality of patient's life. Moreover, it has been shown that a severe treatment schedule that fails to cure can result in a shorter patient survival time (PST) than a milder palliative treatment [6].

Previously, the conflict between objectives was resolved by addressing each of them separately, that is, treatment strategies were sought which optimised only one of the objectives without considering the other [7]. The choice of the best strategy was left to the decision maker, i.e. the practicing oncologist who treats the patient; the role

of the optimiser was to provide the alternatives to choose from. Although this approach produced some interesting results, it cannot show the whole picture. In particular, the single-objective approach is ineffective in finding versatile treatment schedules that show a reasonably good performance in one objective dimension and, at the same time, can be effectively used with the other objective in mind.

We contend that such versatile treatment schedules will belong to the Pareto-optimal set, which needs to be found by the optimisation algorithm capable of dealing with multi-objective optimisation. We also contend that Evolutionary Algorithms are well-suited for this role. However, before EA can be applied to the multi-objective optimisation problem of cancer chemotherapy, we need to mathematically formulate the objectives of chemotherapeutic treatment and the constraints imposed on it.

2.2 Problem Definition and Related Concepts

In general, a multi-objective optimisation problem (MOP) consists of n decision variables comprising a decision vector $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \Omega \subset \mathfrak{R}^n$, m constraints $g_1(\mathbf{x}), g_2(\mathbf{x}), \dots, g_m(\mathbf{x})$, and k objectives expressed as (non)linear criteria or objective functions $f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_k(\mathbf{x})$. Brought together, the multiple objectives define the evaluation function $F(f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_k(\mathbf{x})): \Omega \rightarrow \Lambda \subset \mathfrak{R}^k$, which, if some of the objectives are in conflict, places a partial, rather than normal, ordering on the search space Ω [14]. In order to mathematically define this partial ordering, a notion of Pareto dominance is introduced in the objective space Λ . (NOTE. In this paper we will be concerned with the problem of maximising the values of the objective functions.)

Definition 1. A decision vector $\mathbf{x} = (x_1, x_2, \dots, x_n)$ is said to **dominate** $\mathbf{x}' = (x'_1, x'_2, \dots, x'_n)$, denoted as $\mathbf{x} \succ \mathbf{x}'$, iff the value of the evaluation function at \mathbf{x} , $F(\mathbf{x})$, is **partially greater** than $F(\mathbf{x}')$, i.e., $\forall i \in \{1, \dots, k\}: f_i(\mathbf{x}) \geq f_i(\mathbf{x}') \wedge \exists j \in \{1, \dots, k\}: f_j(\mathbf{x}) > f_j(\mathbf{x}')$.

The specificity of multi-objective optimisation is to find a set of non-dominated decision vectors rather than the global optimum, which might not even exist. For this purpose, the concept of Pareto optimality ought to be used.

Definition 2. The decision vector $\mathbf{x} \in \Omega$ is **Pareto-optimal** iff \mathbf{x} is non-dominated regarding Ω ; formally

$$\neg(\exists \mathbf{x}' \in \Omega: \mathbf{x}' \succ \mathbf{x})$$

Pareto-optimal decision vectors cannot be improved in any objective without causing deterioration of at least one other objective. Such decision vectors comprise the Pareto-optimal set, $P^* \subset \Omega$, in the search space. The mapping of the Pareto-optimal set to the objective function space gives rise to the Pareto front PF^* . The

Pareto front can be non-convex and non-connected; nonetheless, if it is known, or at least approximated reasonably well, the decision maker will be able to select a solution via a choice of acceptable objective performance and, as a result of this, the problem of multi-objective optimisation will be resolved.

Therefore, in order to solve the optimisation problem of cancer chemotherapy, we need to find the set of treatment schedules, which yields the Pareto front in the treatment performance space. This will allow the oncologist to make a decision on which treatment schedule to use, given his/her preferences or certain priorities. In the remainder of this section we will define the decision vectors and the search space for the cancer chemotherapy optimisation problem, specify the constraints, and particularise the optimisation objectives.

Anti-cancer drugs are usually delivered according to a discrete dosage program in which there are n doses given at times t_1, t_2, \dots, t_n [6]. In the case of multi-drug chemotherapy, each dose is a cocktail of d drugs characterised by the concentration levels $C_{ij}, i \in \overline{1, n}, j \in \overline{1, d}$ of anti-cancer drugs in the bloodplasma. Optimisation of chemotherapeutic treatment is achieved by modification of these variables. Therefore, the search space Ω of the chemotherapy optimisation problem is the set of control vectors $\mathbf{c} = (C_{ij})$ representing the drug concentration profiles.

However, not all of these profiles will be feasible as chemotherapy treatment must be constrained in a number of ways. Although the constraint sets of chemotherapeutic treatment vary from drug to drug as well as with cancer type, they have the following general form.

1. Maximum instantaneous dose C_{\max} for each drug acting as a single agent:

$$g_1(\mathbf{c}) = \left\{ C_{\max j} - C_{ij} \geq 0 : \forall i \in \overline{1, n}, \forall j \in \overline{1, d} \right\} \quad (1)$$

2. Maximum cumulative C_{cum} dose for drug acting as a single agent:

$$g_2(\mathbf{c}) = \left\{ C_{\text{cum} j} - \sum_{i=1}^n C_{ij} \geq 0 : \forall j \in \overline{1, d} \right\} \quad (2)$$

3. Maximum permissible size N_{\max} of the tumour:

$$g_3(\mathbf{c}) = \left\{ N_{\max} - N(t_i) \geq 0 : \forall i \in \overline{1, n} \right\} \quad (3)$$

4. Restriction on the toxic side-effects of multi-drug chemotherapy:

$$g_4(\mathbf{c}) = \left\{ C_{s\text{-eff} k} - \sum_{j=1}^d \mathbf{h}_{kj} C_{ij} \geq 0 : \forall i \in \overline{1, n}, \forall k \in \overline{1, m} \right\} \quad (4)$$

The factors \mathbf{h}_{kj} in the last constraint represent the risk of damaging the k^{th} organ or tissue (such as heart, bone marrow, lung etc.) by administering the j^{th} drug.

Estimates of these factors for the drugs most commonly used in treatment of breast cancer, as well as the values of maximum instantaneous and cumulative doses, can be found in [4], [8] or [11].

Regarding the objectives of cancer chemotherapy, we focus our study on the following two. The primary objective is to eradicate the tumour (curative treatment). We define eradication to mean a reduction of the tumour from the initial size to a size below 10^3 cells. Clinical experience shows that other mechanisms (e.g. programmed cell death, a.k.a. apoptosis) are capable of removing remaining tumour cells at this point.

In order to simulate the response of a tumour to chemotherapy, a number of mathematical models can be used [10]. The most popular is the Gompertz growth model with a linear cell-loss effect [17], which has been extensively validated in clinical trials:

$$\frac{dN}{dt} = N(t) \cdot \left[I \ln\left(\frac{\Theta}{N(t)}\right) - \sum_{j=1}^d \mathbf{k}_j \sum_{i=1}^n C_{ij} \{H(t-t_i) - H(t-t_{i+1})\} \right] \quad (5)$$

where $N(t)$ represents the number of tumour cells at time t ; I, Θ are the parameters of tumour growth, $H(t)$ is the Heaviside step function; \mathbf{k}_j are the quantities representing the efficacy of anti-cancer drugs, and C_{ij} denote the concentration levels of these drugs. One advantage of the Gompertz model from the computational optimisation point of view is that the equation (5) yields an analytical solution after the substitution $u(t) = \ln(\Theta/N(t))$ [5]. Since $u(t)$ increases when $N(t)$ decreases, the primary optimisation objective of tumour eradication can be formulated as follows [9]:

$$\underset{\mathbf{c}}{\text{maximise}} \quad f_1(\mathbf{c}) = \int_{t_1}^{t_n} \ln\left(\frac{\Theta}{N(t)}\right) dt \quad (6)$$

subject to the state equation (5) and the constraints (1)-(4).

The second objective of cancer chemotherapy is to prolong the patient survival time (PST) maintaining a reasonable quality of life during the palliation period. If we denote the PST as T , then the second objective becomes:

$$\underset{\mathbf{c}}{\text{maximise}} \quad f_2(\mathbf{c}) = \int_{t_1}^T dt = T \quad (7)$$

again subject to (1)-(5).

Therefore, the evaluation function of the multi-objective optimisation problem of cancer chemotherapy takes the form of a two-dimensional vector function $F(\mathbf{c}) = [f_1(\mathbf{c}), f_2(\mathbf{c})]^T$, which maps the decision vectors $\mathbf{c} \in \Omega$ to the objective function space $\Lambda \subset \Re^2$ using the objectives (6) and (7).

As we mentioned in the previous section, these objectives are conflicting in nature. The conflict between objectives manifests itself in the fact that small tumours are more likely to be successfully eliminated, whereas it is much easier to palliate a large tumour [6]. Thus, in order to pursue the first treatment objective the maximum tolerable amount of drugs has to be administered at the start of treatment. The best palliative strategy, on the other hand, is to allow the tumour to grow up to the maximum size and then to maintain it at that level using only a necessary amount of drugs.

Taking this into account and considering the number of constraints imposed on chemotherapeutic treatment, it is not difficult to see that the traditional approaches to multi-objective optimisation of cancer chemotherapy (such, for example, as the weighting or constraint methods) are likely to fail. Our previous experiments with traditional optimisation methods (the complex and Hooke & Jeeves techniques) showed a lack of robustness in finding feasible solutions even in the case of single-objective optimisation [11]. Moreover, all traditional methods require several optimisation runs to obtain an approximation of the Pareto-optimal set. As the runs are performed independently from each other, synergies between them cannot be easily exploited, which may lead to substantial computational overhead [16].

Therefore, the necessity of a specialised optimisation technique to deal with the cancer chemotherapy MOP is evident. Recently, Evolutionary Algorithms (EA) have become established as an alternative to traditional methods. The major advantages of EA are: 1) the ability to effectively search through large solution spaces; 2) the ability to overcome the difficulties faced by the traditional methods mentioned above; and 3) the ability to approximate the Pareto-optimal set in a single run. In the following section we briefly discuss the salient features of Evolutionary Algorithms.

2.3 Evolutionary Multi-objective Optimisation

Evolutionary Algorithms entail a class of stochastic optimisation methods that simulate the process of natural selection. Although the underlying principles are quite simple, these algorithms have proven to be in general robust and powerful [1]. A large number of applications of EA to hard, real-world MOPs, the survey of which is given in [2], suggest that multi-objective optimisation of cancer chemotherapy is the problem set where Evolutionary Algorithms might excel.

As with any MOP, the problem of cancer chemotherapy optimisation involves two independent processes. The first process is the search through the solution space for the Pareto-optimal set. The search space of cancer chemotherapy MOP is very large [7], which makes the multi-directional and synergetic features of EA extremely helpful. The second process is decision-making, i.e. the selection of a suitable compromise solution from the Pareto-optimal set.

Depending on the order of performing these processes, the preferences of the decision maker (the oncologists in our case) can be made known either before, during or after the search process [14]. In the case of a priori preference articulation, the objectives of the given MOP are aggregated into a single objective that implicitly includes preference information (in the form of objective weights for example). This

approach requires profound domain knowledge, which is not available for the optimisation problem of cancer chemotherapy [3].

If the search process is performed without any preference information given by the oncologist, then we are applying a posteriori preference articulation. Here, the search results in a set of candidate treatment schedules (ideally the Pareto-optimal set of treatments), from which the final choice is made by the oncologist. The main drawback of the latter approach is that it entirely excludes the domain knowledge, which in some cases might substantially reduce the size of the search space or/and its complexity. However, in a general case of cancer chemotherapy optimisation such a reduction is not advisable [8], which supports the suitability of the a posteriori approach.

Also, the process of decision-making may overlap with that of search. This means that after each optimisation step, a number of alternative treatment schedules (temporary Pareto-optimal set) are presented, on the basis of which the oncologist specifies further preference information, thereby guiding the search process. Such an approach is known as progressive preference articulation [14] and is a promising way to combine the advantages of the previous two. One example of how it can be used in the context of cancer chemotherapy is to optimise the modification of an existing treatment schedule rather than a schedule itself [11]. However, in this paper we concentrate our efforts on the optimisation of treatment schedules themselves as this is a more general problem. In solving this problem we do not wish to restrict the search process in any way, since a priori information on whereabouts of the Pareto-optimal set in the search space is unavailable. Therefore, hereafter we need to resort to a posteriori preference articulation approach to multi-objective optimisation.

Having established the strategic aspects of the method that is to be utilised for solving the cancer chemotherapy MOP, we now need to specify the implementation details. A general Evolutionary Algorithm can be presented as follows.

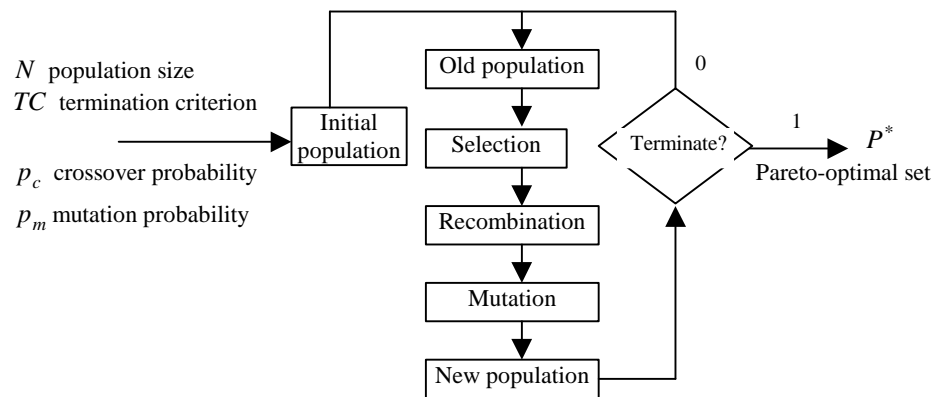


Fig. 1. Input, Output, and Internal Structure of a Generic Evolutionary Algorithm

This general structure holds for most EA implementations. The distinctive feature of Evolutionary Algorithms applied to multi-objective optimisation, however, is that they require addressing the following specific issues [15]. The first issue is how to

accomplish fitness assignment, and consequently selection, given a vector-valued evaluation function $F : \Omega \rightarrow \Lambda$. In contrast to single-objective optimisation, where the fitness function takes into account only one optimisation objective, the fitness function of a multi-objective EA needs to map a k -dimensional objective function space to scalar numbers in such a way as to guide the search process to the Pareto-optimal set. Secondly, the diversity of an EA population has to be maintained more than ever in order to achieve a well distributed and well spread set of non-dominated solutions, in addition to preventing premature convergence.

A body of work has been established setting up various fitness assignment methods, selection techniques, and population diversifying schemes [2], [5], [14], [16]. As a consequence of this, many implementations of multi-objective EA are now available. In spite of this variety, however, there is no clear guideline on which EA implementation is suited to which sort of problem in the sense of ensuring that the derived solutions are the best available [15]. Thus, the choice is subjective and is often based on the developer's attempt to integrate the domains of the optimisation problem and that of the implementation algorithm [14].

Among the different implementation algorithms that have been proposed in the literature and have been used by EA practitioners, we have chosen and will base our study on the Strength Pareto Evolutionary Algorithm (SPEA) thoroughly described in [16]. This algorithm combines promising aspects of various multi-objective EA and has shown a superior performance on a number of test problems [15]. In the next section we describe how it can be applied to the multi-objective optimisation problem of cancer chemotherapy.

3 Evolutionary Search for Optimal Treatment Schedules

The search process aiming at finding non-dominated (with respect to the treatment objectives specified in Section 2.1) chemotherapy schedules is the main part of computational optimisation of chemotherapeutic treatment. The decision-making process is, of course, based on the results of this search, but it is left to oncologists and therefore lies outside the scope of the present paper.

The search for non-dominated treatment schedules is accomplished using the SPEA approach. Multi-drug chemotherapy schedules, represented by decision vectors $\mathbf{c} = (C_{ij}), i \in \overline{1, n}, j \in \overline{1, d}$, are encoded as binary strings. Using the EA terminology, the individual space \mathbf{I} (a discretized version of Ω) can then be expressed as a Cartesian product

$$\mathbf{I} = A_1^1 \times A_1^2 \times \dots \times A_1^d \times A_2^1 \times A_2^2 \times \dots \times A_2^d \times \dots \times A_n^1 \times A_n^2 \times \dots \times A_n^d \quad (8)$$

of allele sets A_i^j . Each allele set uses a 4-bit representation scheme

$$A_i^j = \{a_1 a_2 a_3 a_4 : a_k \in \{0,1\} \forall k \in \overline{1,4}\} \quad (9)$$

so that each concentration level C_{ij} takes an integer value in the range of 0 to 15 concentration units. In general, with n treatment intervals and up to 2^p concentration levels for d drugs, there are up to 2^{npd} individual elements. Henceforth we assume that $n=10$ and that the number of available drugs is restricted to three, one of which is strong but highly toxic, another is medium, and the last one is less toxic at the expense of reduced efficacy. In our study we experiment with the following drugs: Taxotere (strong), Adriamycin (medium), and Cisplatinum (weak), which are commonly used in multi-drug treatment of breast cancer. The values $n=10$ and $d=3$ result in the individual (search) space of power $|\mathbf{I}|=2^{120}$ individuals, referred to as chromosomes.

Thus, a chromosome $x \in \mathbf{I}$ can be expressed as

$$x = \{a_1 a_2 a_3 \dots a_{4nd} : a_k \in \{0,1\} \forall k \in \overline{1,4nd}\} \quad (10)$$

and the mapping function $m: \mathbf{I} \rightarrow \mathbf{C}$ between the individual \mathbf{I} and the decision vector \mathbf{C} spaces can be defined as

$$C_{ij} = \Delta C_j \sum_{k=1}^4 2^{4-k} a_{4d(i-1)+4(j-1)+k}, \quad \forall i \in \overline{1,n}, j \in \overline{1,d} \quad (11)$$

where ΔC_j represents the concentration unit for drug j . This function symbolizes the decoding algorithm to derive the decision vector $\mathbf{c} = m(x)$ from a chromosome x . If this vector violates any of the constraints (1)-(4), a penalty is applied to the values of the objective functions. The evaluation function F and the penalties yield the following augmented objective vector:

$$[f_1(\mathbf{c}) - \sum_{j=1}^m P_j \cdot \max^2\{-g_m(\mathbf{c}), 0\}, f_2(\mathbf{c}) - \sum_{j=1}^m P_j \cdot \max^2\{-g_m(\mathbf{c}), 0\}]^T \quad (12)$$

on the basis of which the fitness value is assigned to x .

The fitness assignment procedure is a two-stage process that uses two interacting populations - the external set \overline{P} , which stores the individuals representing a non-dominated front among all solutions considered so far, and the EA population of chromosomes P . The first stage is to rank the elements of \overline{P} , and the second is to evaluate the chromosomes in P . The full description of the fitness assignment procedure and of other auxiliary SPEA elements is given in [16]. All that remains to be specified here is the parameter settings of the SPEA algorithm: namely, the population size N , the maximum number \overline{N} of elements in the external set, the probabilities of crossover (p_c) and mutation (p_m), and the maximum number of generations TC , which serves as the stopping criterion.

In the choice of these parameters we will adhere to the values used in the previous work [7]. This will allow us to make an unbiased comparison between the single- and multi-objective approaches to chemotherapy optimisation. Moreover, in [12] the authors have shown that with the following values of crossover and mutation probabilities the efficiency of evolutionary search significantly improves. Taking all these into account, the SPEA parameters will be set to the values: $TC = 10000$, $p_c = 0.6$ and $p_m = 0.1$. Regarding the N/\bar{N} ratio, we have chosen it to be 50/5 as the population size $N = 50$ has proved to be efficient in the previous studies (see [7], [10] and [11]), and the size $\bar{N} = 5$ is deemed to be sufficient to provide the required density of solutions on the Pareto front without drastically reducing selection pressure of SPEA.

4 Results

In order to illustrate the results of cancer chemotherapy optimisation, a number of typical treatment scenarios are usually considered. For instance, in [7] the authors apply single-optimisation to three scenarios that very often occur in practice: 1) cure possible, eradication treatment is applied; 2) cure impossible, eradication treatment is applied; 3) cure impossible, palliative treatment is applied. However, when the single-objective approach is used, it is often necessary to assign priorities to each objective. Generally, the primary objective of cancer treatment is to eradicate the tumour. In cases when the eradication is possible (the first treatment scenarios in [7]), different treatment schedules are merited on the basis of how quickly they can achieve this goal. If, on the other hand, cure is impossible (the second treatment scenario) but the treatment objective remains unchanged, then a single-objective optimisation algorithm is likely to yield a solution far from optimum. In the latter case the palliative treatment gives much better results [7].

Thus, one strong drawback of single-objective chemotherapy optimisation is that the choice of the desired treatment outcome needs to be made when the treatment starts and the cost of a mistake may be unacceptably high. In order to overcome this difficulty, we now consider the main objectives of chemotherapeutic treatment - tumour eradication and prolongation of PST - simultaneously for each potential chemotherapy schedule. Figures 2, 3 and 4 show three multi-drug treatment schedules from the set of Pareto-optimal decision vectors found by SPEA.

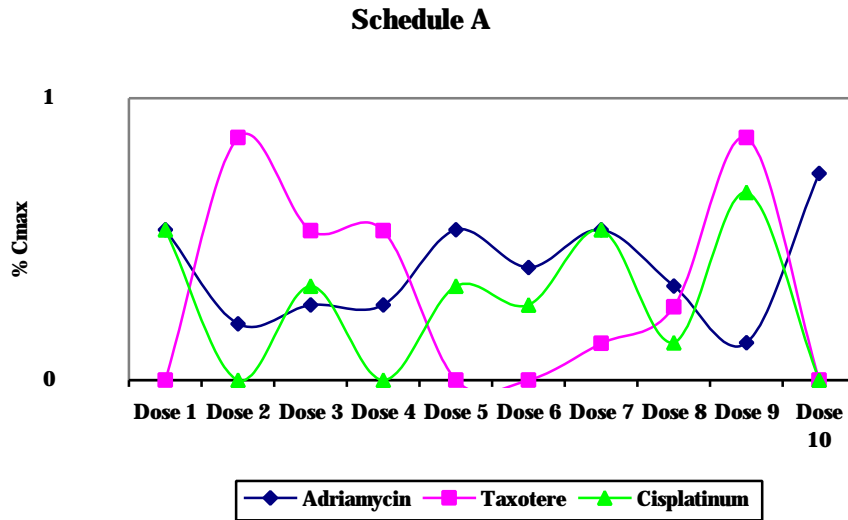


Fig. 2. Schedule achieving an agreeable balance between treatment objectives




Fig. 3. Schedule excelling in minimising the size of tumour burden




Fig. 4. Schedule excelling in prolongation of the patient survival time

Table 1 gives the values that quantitatively characterise the treatment schedules presented above.

Table 1. Comparison between different Pareto-optimal treatment schedules

Optimal treatment	Constraints satisfaction	$f_1(\mathbf{c})$	$f_2(\mathbf{c})$	Average tumour size in terms of $N(t_0)$	PST
Schedule A	All constraints are satisfied	4.3230	4.3370	0.7986	36 weeks
Schedule B	All constraints are satisfied	5.5864	3.3835	0.7380	35 weeks
Schedule C	All constraints are satisfied	3.1151	7.0297	0.8998	38 weeks

As can be seen from this table, the schedules B and C yield a good value of one optimisation objective ($f_1(\mathbf{c})$ and $f_2(\mathbf{c})$ respectively) at the expense of relatively poor performance in the other objective's dimension. This is reminiscent to the single-objective optimisation [7], which would favour the schedule B as a candidate for the tumour eradication treatment strategy and would reserve the schedule C for palliative treatment.

From the treatment profile corresponding to the schedule B (see Figure 3) we can observe that the good result in tumour eradication is achieved by administering high doses of Taxotere and Cisplatin towards the end of the treatment period. There is a danger in doing this however. If the tumour eradication strategy fails to achieve the desired outcome, studies show that the tumour can re-grow again and reach the lethal size in shorter time than it would have done had a milder palliative strategy been used

[6]. The treatment schedule C, on the other hand, prolongs the PST to a greater extent by keeping the number of tumour cells at an acceptable level (Figure 4 shows that this is done by administering relatively small dosages of anti-cancer drugs), but might miss a chance to completely eradicate it.

The major result of using the multi-objective approach to cancer chemotherapy optimisation is in finding the schedule A. As can be seen from its values of the optimisation objectives in Table 1, it is an agreeable compromise between the other two schedules. Although the schedule A is outperformed by the schedules B and C in the single-objective sense, we believe that it represents a more versatile treatment strategy. Our interpretation of its mode of action is that it makes attempts to eradicate the tumour using high doses of Taxotere - the most efficacious drug available (see Figure 2). Failing these attempts, the schedule A switches to the palliative regime similar to that of the schedule C.

The final remark we would like to make in this section is that the information on the effects and on the modes of delivery of the specified drugs has been given to us by our collaborating oncologists or taken from [4]. In our experiments with single-objective optimisation of cancer chemotherapy in [7], [8] and [11], we managed to successfully emulate the outcome of actual clinical trials using the mathematical model (5) and the constraints (1)-(4). This gives us a reason to believe that the results presented in this paper are also viable and that our approach to multi-objective optimisation of cancer treatment can be used in real life situations. In order to prove this we intend to ask clinicians to try our schedules developed for more complex treatment scenarios with the help of the Oncology Workbench [7], [8].

5 Discussion

In this paper we addressed the problem of multi-objective optimisation of cancer chemotherapy. A number of different objectives of chemotherapeutic treatment can be defined. Although some of these objectives need to be considered simultaneously in order to develop an effective treatment, in the past they were optimised independently from each other. Our present study attempts to cover this gap and utilises the evolutionary methods of computational optimisation to find a solution to a complex optimisation problem with two conflicting objectives. The solution is sought in the form of a Pareto-optimal set, which is approximated by an optimisation run of the Strength Pareto Evolutionary Algorithm. The resultant set found during our experiments includes not only the treatment schedules discovered in the previous studies with the help of the single-objective approach, but a number of new schedules as well that have not been detected before. Thus, the multi-objective approach to chemotherapy optimisation reveals additional treatment strategies that can be more suitable in certain cases, thereby assisting in the decision-making process. The number of such strategies increases when more doses are administered during treatment or more anti-cancer drugs are used.

Therefore, one possible direction for future work is to explore more complex treatment scenarios. This might necessitate the introduction of additional objectives in the evaluation function and the enlargement of the EA population and of the

external set. It will be interesting for us to experiment with different N/\sqrt{N} ratios in order to see the effect of this parameter on the effectiveness of the EA search. Another direction is to develop a software tool on the basis of the Oncology Workbench described in [7] and [8] that will allow the oncologist to refine or change treatment preferences during the optimisation run. This would integrate the search and the decision-making processes, resulting in a more efficient and reliable choice of chemotherapeutic treatment.

References

1. Baeck, T., Fogel, D., Michalewicz, Z. (eds.): Handbook of Evolutionary Computation. Oxford University Press, (1997)
2. Coello, C.: An Updated Survey of Evolutionary Multiobjective Optimization Techniques: State of the Art and Future Trends. Proceedings of the 1999 Congress on Evolutionary Computation. IEEE Service Center, Washington, D.C., (1999) 3-13
3. Cassidy, J., McLeod, H.: Is it possible to design a logical development plan for an anti-cancer drug. *Pharmaceutical Medicine*, (1995), **9** 95-103
4. Dearnaley, D., et al.: Handbook of adult cancer chemotherapy schedules. The Medicine Group (Education) Ltd., Oxfordshire, (1995)
5. Fonseca, C., Fleming, P.: An overview of evolutionary algorithms in multi-objective optimization. *Evolutionary Computation*, (1995) **3**(1) 1-16
6. Martin, R., Teo, K.: Optimal Control of Drug Administration in Cancer Chemotherapy. World Scientific, Singapore New Jersey London Hong Kong (1994)
7. McCall, J., Petrovski, A.: A Decision Support System for Cancer Chemotherapy Using Genetic Algorithms. Proceedings of the International Conference on Computational Intelligence for Modelling, Control and Automation, Vol. 1. IOS Press (1999) 65-70
8. McCall, J., Petrovski, A.: OWCH - A Decision Support System for Designing Novel Cancer Chemotherapies. Proceedings of the First International Symposium on Soft Computing in Biomedicine. ICSC Academic Press (1999) 504-510
9. Murray, J.: Some optimal control problems in cancer chemotherapy with a toxicity limit. *Mathematical Biosciences*, (1990) **100** 49-67
10. Petrovski, A., McCall, J.: Computational Optimisation of Cancer Chemotherapies Using Genetic Algorithms. In: John, R., Birkhead, R. (eds.): *Soft Computing Techniques and Applications*. Series on Advances in Soft Computing. Physica-Verlag (1999) 117-122
11. Petrovski, A.: An Application of Genetic Algorithms to Chemotherapy Treatment. PhD thesis, The Robert Gordon University, Aberdeen, U.K., (1999)
12. Petrovski, A., Wilson, A., McCall, J.: Statistical identification and optimisation of significant GA factors. Proceedings of the 5th Joint Conference on Information Sciences. Atlantic City, USA, (2000) Vol. 1, 1027-1030
13. Swan, G.: Role of Optimal Control theory in Cancer Chemotherapy. *Mathematical BioScience*, (1990) **101** 237-284
14. Veldhuizen, D., Lamont, G.: Multiobjective Evolutionary Algorithms: Analyzing the State-of-the-Art. *Evolutionary Computation*, (2000) **8**(2) 125-147
15. Zitzler, E., Deb, K., Thiele, L.: Comparison of Multi-objective Evolutionary Algorithms: Empirical Results. *Evolutionary Computation*, (2000) **8**(2) 173-195
16. Zitzler, E.: Evolutionary Algorithms for Multi-objective Optimization: Methods and Applications. PhD thesis, Swiss Federal Institute of Technology (ETH), Zurich, (1999)
17. Wheldon, T.: Mathematical models in cancer research. Adam Hilger, Bristol Philadelphia (1988)