Accelerated Optimisation of Chemotherapy Treatment Plans

Robert L Barbour
rlb4@macs.hw.ac.uk

May 26, 2008

I *Robert L Barbour* confirm that this work submitted for assessment is my own and is expressed in my own words. Any uses made within it of the words of other authors in any form e.g., ideas, equations, figures, text, tables, programs etc are properly acknowledged. A list of references employed is included.

Signed : Date :

Abstract

This portfolio represents the preparatory work for an MSc project in Artificial Intelligence. The aim of the project is to investigate the acceleration of Evolutionary Algorithms in the solution of cancer chemotherapy drug scheduling problems. The rationale behind the project and a detailed description are given.

A short discussion on the professional, ethical, legal and social issues is presented, followed by an outline of the dissertation. This shows clearly the main elements of the project including: background; literature survey; design and implementation of the experiments; results and conclusions.

A significant part of the portfolio is the literature survey. This covers the medical background of cancer, chemotherapy and evolutionary algorithms, followed by modelling methods and current solutions for the drug scheduling problem. Methods of improving the efficiency of evolutionary algorithms are considered and provide information for the experimental design.

The conclusion examines which methods can be applied to accelerate solution finding and experimental work for the project is anticipated.

Contents

1 Aims and Objectives **3**
1.1 Aims and Objectives 3

2 Rationale and Description **4**
2.1 Rationale 4
2.2 Projection Description 6

3 Professional Issues **8**
3.1 Professional, Ethical and Social Issues 8

4 Project Outline **9**
4.1 Introduction 9
4.2 Literature Review 10
4.3 Research Methods 11
4.4 Experiments 11
4.5 Results 11
4.6 Conclusions 12

5 Literature Review **13**
5.1 Introduction 13
5.2 Evolutionary Algorithms in Medicine 13
5.3 Modelling Tumours and Chemotherapy 14
5.4 Recent Work on Chemotherapy Scheduling 19
5.5 Accelerating Evolutionary Algorithms 25

6 Conclusions **29**
6.1 Discussion 29
6.2 Conclusion 29

A Ethical Form **35**
A.1 Completed Ethical Issues Form 36

B Risk Form **37**
B.1 Completed Risk Assessment Form 38

GLOSSARY

Adjuvant preventative treatment, often after surgery

Angiogenesis growth of blood vessels

CTL Cytotoxic T Lymphocyte

Cytotoxic toxic to cells

EA Evolutionary Algorithm is a generic term

ES Evolutionary Strategy – an EA which uses integer or real number encoding

GA Genetic Algorithm – an EA which uses binary encoding

GP Genetic Programming – an EA which uses tree structures

LEM Learnable Evolutionary Model

LEMMO Learnable Evolutionary Model for Multi-objective Optimisation

MO Multi-Objective

MRI Magnetic Resonance Imaging

Neo-adjuvant pre-operative treatment

PBIL Population Based Incremental Learning

PD PharmacoDynamic – the effect of a drug on the body

PK PharmacoKinetic – the effect of the body on a drug

PSO Particle Swarm Optimisation.

Regimen Systematic plan (of treatment)

WBC White Blood Cell

Chapter 1

Aims and Objectives

1.1 Aims and Objectives

As the mathematical models for chemotherapy scheduling become increasingly complex, the Evolutionary Algorithms to find drug schedules become more computationally expensive. Accelerating EAs is one possible solution, on which this project will focus.
The objectives are to investigate the use of:

- Fitness Inheritance
- Learnable Evolution Methods (LEM)

to reduce the computation time to find solutions to the drug scheduling problem, without compromising the quality of results.

Chapter 2

Rationale and Description

2.1 Rationale

2.1.1 Cancer

Cancer is a major cause of death in modern society. Anti-cancer drugs have been developed by pharmaceutical research. These drugs attack and kill cancer cells. But they also damage healthy cells. To prevent damage to healthy organs and tissue the drug doses have to be limited. A dose of drug is given, then a period for the body to recovery is followed by another dose of drug and so on in a cyclic fashion. These dose schedules developed from clinical trials and the experience of oncologists in practice. The schedules have become more complex with the increase in the number of drugs and in the use of multi-drug treatments.
Different drugs affect different organs, so limiting the dosage to control side effects also becomes more complicated. With a treatment schedule of ten or more dose/rest periods and 20 or 30 drugs available in different doses, there are a huge number of treatment combinations possible, far too many for clinical trials to evaluate. So it is possible that the drugs available are not being used to best advantage.
Extensive clinical trials are not ideal in the search for better treatment schedules. It would be unethical for oncologists to treat patients with a less than optimum schedule as part of a research programme. Five and ten year survival rates are used to measure progress in cancer treatment. This means clinical trials take a long time to show progress.

2.1.2 Computer Modelling

Computer simulation is a method of investigating problems without physical trials. One method of solving problems by computer is to mimic nature's very successful evolutionary methods using programs called Evolutionary Algorithms (EA). By representing individual possible solutions to a problem as members of a population, allowing mutation and then using the better individuals to

form the next generation, the program gradually evolves good solutions. This usually takes many generations. Given the right conditions EAs can produce better solutions than other methods, but there is a computational expense in the time to simulate the many thousands of generations usually needed. Selection is a key process and is usually based on some measure of an individual's *fitness* i.e. its quality as a solution to the problem. The fitness value is usually the result of an evaluation of some mathematical function which may be a simple formula or a complex mathematical process.

2.1.3 Chemotherapy Scheduling

When EAs are applied to the cancer chemotherapy scheduling problem, each individual in the population often represents one set of drug dose/rest periods i.e. one schedule of treatment (regimen). The fitness is found by applying this treatment schedule to a mathematical model of the tumour growth and the effect of the drugs. The tumour/drug model simulates the interaction of the drugs throughout the treatment period and provides a fitness measure to the EA. This may be a single value, for example tumour size at the end of the simulated treatment, or several values representing the average tumour size and the expected patient survival time, etc.
The EA may use this fitness evaluation function for every member of the population (maybe hundreds) in every generation (tens of thousands). If the fitness is a complex function it can limit the effectiveness of an EA, the time for evaluation creating a bottleneck in the evolutionary simulation.
Developing these simulations of chemotherapy treatment and improving the models and methods is an ongoing process. Ultimately this should lead to a better understanding of cancer treatment and encourage further collaboration between researchers and oncologists. In the long term this may result in clinical trials and treatments schedules derived with the help of computer modelling. This is still some (considerable) distance in the future. Progress will depend on many factors including: better models; faster models; oncologist confidence in the models and simulations, etc.

2.1.4 Aims and Objectives

In the light of this, EA acceleration is one aspect worthy of consideration, on which this project will focus. The objectives are to investigate the use of:

- Fitness Inheritance
- Learnable Evolution Methods (LEM)

to reduce the computation time to find solutions to the drug scheduling problem, without compromising the quality of results.

2.2 Projection Description

Both Fitness Inheritance and Learnable Evolution Models (LEM) as described later (section 5.5) can be used to reduce the number of fitness evaluations computed to find a solution. EAs will be used with a kinetic tumour/drug model in this investigation.
The kinetic model used is written in Java and is provided by a researcher in the field. It uses the Gompertz growth model to compute fitness functions for a binary encoded multi-drug scenario. Some re-engineering of the code will be necessary. Integer encoding of the drug schedule will replace binary encoding and these changes will be tested and validated before use.
The model will be the basis for three sets of experiments: one to solve a drug scheduling problem using an EA, (the control); one to solve the same problem using fitness estimation for a proportion of the time; and one to to find solutions using LEM.
There are various input parameters which are required by the model. Some will be provided by the individual problem, initial tumour size and growth rate for example. Others may remain constant in the model: fitness function penalties for tumour size; drug dose and toxicity; risk factors for major organs etc.
Initial EA control experiments will establish the operating envelope for these parameters.

Control Set

Having set the model parameters which will remain fixed, the EA will be used to solve five problems covering a range of tumour sizes, growth rates and drug selection. These same problems will be used by the other algorithms.
The EA is a generational, elitist, cross-over plus mutation algorithm with tournament selection. Termination criteria will be set after initial experiments to establish the quality of solutions.

Fitness Estimation

This will be an EA which uses an estimate of fitness, the weighted average of the parents' fitness, for a proportion of the time, in place of actual fitness evaluations. The algorithm used will be an implementation of that detailed in Section 5.5.1 of the literature survey. The proportion of fitness estimation will be varied to investigate its influence on results.

Learnable Evolutionary Models (LEM)

The LEM algorithm will be that described in Section 5.5.2 of the literature review, alternating between machine learning and evolutionary methods until a solution is reached. Finding appropriate termination criteria for switching between the two methods and also a final criterion to halt the LEM will require some investigation.

Results

Analysis of results will compare the quality of the results from the three experiments over the range of problems examined. Comparison of the efficiency of the methods will be made on the basis of number of fitness evaluations and also overall run times. The quality of solutions will also be examined to ensure that efficiency comparisons are made on an equitable basis.

Chapter 3

Professional Issues

3.1 Professional, Ethical and Social Issues

The use of the Java code for the tumour model is acknowledged and will be attributed to the author.
Otherwise there are no professional, ethical or social issues in this project.

Chapter 4

Project Outline

4.1 Introduction

4.1.1 Aims anddddddddddd Objectives

Aim and Objectives

To explore the feasibility of accelerating Evolutionary Algorithms (EA) in chemotherapy drug scheduling. Objectives are to reduce the computational effort to find good solutions and maintain diversity of treatment schedules.

- Estimating fitness function values
- Learning from previous generations

4.1.2 Cancer and Chemotherapy

Cancer

.

Chemotherapy

.

4.1.3 Drug Scheduling

.

Modelling

.

Implementing Models with EAs

.

4.1.4 The need for EA Efficiency

.

4.2 Literature Review

.

4.2.1 EAs in Medicine

.

Background showing wide range of EA applications

.

4.2.2 Modelling Tumours and Chemotherapy

.

Description of model types

.

4.2.3 Improving EAs

.

4.2.4 Review of Speed-up Strategies

.

Fitness Inheritance

.

Learning

.

4.2.5 Solving the scheduling problem: Current Results

.

4.3 Research Methods

.

4.3.1 Description of model chosen

.

4.3.2 Source code for model and alterations

.

4.3.3 Design of Experiments

.

4.4 Experiments

.

4.4.1 Control Experiment

.

4.4.2 Experiment 1 – Fitness Inheritance

.

4.4.3 Experiment 2 – LEM

.
.

4.5 Results

.

4.5.1 Present results of experiments

.

4.5.2 Discussion and Interpretation

.

4.5.3 Compare and Contrast

.

4.6 Conclusions

.

4.6.1 Summary of Achievements

.

4.6.2 Aims and Objectives Achieved?

.

4.6.3 Contribution by this work

.

4.6.4 Future work

.

Chapter 5

Literature Review

5.1 Introduction

Cancerous tumours are characterised by rapidly dividing cells. Chemotherapy is the application of drugs in an attempt to reduce or eliminate the tumour, without compromising the immune system or causing unacceptable side effects to the patient.
Clare et al.[4] report that breast cancer mortality rates in the USA have only improved slightly since 1970, despite the extensive use of chemotherapy. The use of mathematical modelling in cancer treatment is reviewed. The Gompertz model for solid tumours is commonly used. It does not model the natural history of breast cancer, which often shows dormant phases and growth spurts. Unfortunately there is limited scope to improve natural history modelling as there is little clinical data available – once a tumour is diagnosed there may be one previous mammogram, but little else to establish the dormancy/growth spurt history. Cited by Clare et al., Norton and Simon[21, 22], were among the first to model chemotherapy using the Gompertzian model for tumour growth, regression and re-growth. Clonal resistance is also discussed, citing Goldie et al.[8, 9] as putting tumour drug resistance on a mathematical basis and also suggesting alternating drug chemotherapy.
Smith et al.[31] investigate the use of neoadjuvant (pre-operative) chemotherapy to increase breast conservation and improve survival. It's use has been increasing for 30 years and this comparative study of two different drugs shows that one is significantly better than the other. Five percent of the women actually got worse during the trial before being withdrawn. In human terms experiments *in vivo* are so very expensive compared to *in silico*.

5.2 Evolutionary Algorithms in Medicine

Evolutionary Algorithms (EAs) are computer programs which mimic nature by making a series of small random steps to gradually find the solution to a problem

and are used in a wide range of applications in medicine.
For example, the reliable diagnosis of Parkinson's disease is notoriously difficult. An EA classification technique by Smith et al.[33] can identify a Parkinson's disease symptom, bradykinesia, by the analysis of the figure copying ability of patients.
A useful improvement to MRI scanners would be to reduce the length of the two metre long tunnel which causes a claustrophobic reaction in some patients. MRI Magnet design is a very complex problem addressed by Yuan et al.[41]. An EA design shows the feasibility of reducing the length to around one metre.
Medical practitioners of *keyhole surgery* benefit from a simulator developed by Bosman[3] to help acquire the difficult practical skills necessary to safely insert a catheter and guide wire into the artery of a patient. The simulator uses an EA model which outperforms a problem-specific analytical algorithm.
The three dimensional problem of beam direction design for radiotherapy treatment is complex. Using a multi-objective EA, Goldbarg et al.[7] develop a model which generates options for an oncologist to select a treatment regimen.

5.3 Modelling Tumours and Chemotherapy

5.3.1 Oncology and Mathematics

Gatenby and Maini[6] bemoan the lack of biologists' and oncologists' enthusiasm for mathematical modelling. Despite a huge volume of oncological data and molecular biology theory there is no overall model for considering this data. According to Gatenby, *"Articles in cancer journals rarely feature equations. Clinical oncologists and those who are interested in the mathematical modelling of cancer seldom share the same conference platforms"*. Referring to genetic and molecular biology Gatenby observes *"Critical parameters that emerge from mathematical modelling focus attention on issues that require further theoretical and experimental work"*. This is also relevant in cancer chemotherapy.

5.3.2 Biological Aspects

Cancer is not a single disease; the biology of cancers is described in[14, 34]. Consider first the free growth of a tumour. This does not proceed at a fixed rate. Once a tumour is established, initial growth rate is high, whilst the tumour receives nutrients from local blood vessels. Tumour cells can produce proteins which stimulate angiogenesis, the production of blood vessels in the tumour which supply food, oxygen and remove waste. As the tumour grows outwards the centre may consist of dead cells. If no angiogenesis occurs, the tumour size is limited by the nutrition available. In any case, the growth rate reaches a maximum and then reduces until the tumour attains a maximum size when it is in equilibrium with the nutrient supply.
But the body's immune system reacts to cancer. White blood cells include Cytotoxic T Lymphocytes (CTL) which can stimulate programmed cell death

(apoptosis) of cancerous and other faulty cells.This is distinct from necrosis (uncontrolled cell death) which leads to inflammation and health problems. Apoptosis is triggered by CTL detecting a faulty cell, binding with it and expressing a surface enzyme which *kills* the target cell – by triggering a mechanism which causes the cell to self-destruct in a controlled fashion, allowing debris to be removed by other processes.
Cells have a life cycle consisting of periods of replication, quiescence and death. The replication period or *cell cycle* is characterised by four phases:

- G_1 – Growth and preparation for replication;
- S phase – Synthesis;
- G_2 – Preparation for mitosis ;
- M phase – Mitosis.

Quiescent cells are in stage G_0.
M phase is the shortest part of the cycle and mitosis actually takes place in four (sub-)phases, progressing from a single parent cell in Prophase, through Metaphase and Anaphase to Telophase, at the end of which two daughter cells exist.
If the DNA in a cell is damaged, the levels of protein p53 increase in the cell, allowing time for repairing the DNA by keeping the cell in phase G_1. If a cell is severely damaged, apoptosis in induced. Because of p53's ability to initiate apoptosis in faulty cells, it has a key role in the control of cancer. However many cancers have p53 mutations, which cannot perform this function.
Tumour cells and normal cells proliferate by the same cell cycle/mitosis mechanism, but different cell types replicate at different rates. Cancer is characterised by uncontrolled cell production. High rates are seen in the digestive tract, hair follicles and in tumours. The cytotoxic drugs used in chemotherapy usually target the rate of cell replication by affecting cells in the M-phase of the cell cycle. Unable to distinguish tumour cells from other fast replicating cells, drug-induced apoptosis also affects the normal cell population. This explains the well-known chemotherapeutic side effects of hair-loss and nausea.

5.3.3 Mathematical modelling

Modelling of chemotherapy treatment includes:

- tumour growth
- tumour cells killed by drug(s)
- constraint(s) on drug application, dose rates and amounts etc.
- meeting the objective(s) of treatment

Tumour Growth

Whilst there is a proliferation of models for tumour development, many are at the molecular or genetic level and have not been used for tumour chemotherapy modelling. Several general models of growth can be applied to tumours – Petrovski and McCall[26] cite Marusic et al.[16] as listing four common models of growth:

- exponential: $f(N) = \lambda N$
- von Bertalanffy: $f(N) = \lambda \sqrt[3]{N^2} - \mu N$
- Verhulst: $f(N) = \lambda N - \mu N^2$
- Gompertz: $f(N) = \lambda N - \mu N ln N$

where N is the number of tumour cells and λ and μ are growth parameters.
Marusic et al found the von Bertalanffy equation inadequate and the Gompertz model the most appropriate for spheroid growth.
The Gompertz equation is currently the basis for much work in this field.
Useful models are constructed by adding features to a basic growth model to account for the effect of drugs and the immune system and so on.

Cell Kill (Cell Loss)

The rate at which a cytotoxic drug kills cells will depend on its efficacy and its concentration at the kill site.

Constraints

Due to the toxicity of chemotherapy drugs, there are limits on drug dose, both on instantaneous dose rate and total dose over a period of treatment. These limits are to protect healthy cells and control side effects of the drug. Constraints may also be applied to meet objective requirements, for example tumour size must never get any larger than a certain size.

Objectives

Tumour reduction is the primary aim but there may be details of how to measure the success of that. For example the primary objective may be for a small tumour at the end of treatment (curative), or minimising the average tumour size throughout treatment (palliative).

5.3.4 A Basic Model

To a growth equation, add a 'cell kill' or 'cell loss' term to model the effect of the drug on the tumour cell population.
If κ is the effectiveness and c the concentration of a cytotoxic drug, adding the 'kill' term $-\kappa c N$ to the growth equation gives [26] :

$$N = f(N) - \kappa cN \tag{5.1}$$

where $f(N)$ is a tumour growth function.
Using the Gompertz growth model and differentiating equation (5.1) McCall et al.[19] show that this becomes :

$$\frac{dN}{dt} = N(t)\lambda ln\left(\frac{\Theta}{N(t)}\right) - \kappa c(t) \tag{5.2}$$

where $N(t)$ is the number of tumour cells at time t, $c(t)$ is the drug concentration at time t and λ, Θ are tumour parameters.
This equation can be extended to cover multiple drugs by summing the kill terms for each drug.

5.3.5 Swan's Review 1990

Swan[35] conducts a wide ranging review of optimal control modelling in cancer chemotherapy scheduling, characterising the main types of models as:

1. Cell cycle models
2. Cell population models
3. Kinetic tumour growth models

Cell Cycle models

In examining cell cycle modelling, Swan gives the main assumptions:

- Cell cycle mechanisms apply to cancerous cells
- Cells not in the cycle are in stage G_0
- One model unusually assumes cells move from G_0 directly to the S phase

Stating that many people believe chemotherapy can be effective in some phase of the cell cycle, Swan conducts critical reviews of the work of about a dozen researchers in the area of optimal control of drug scheduling. Built on a sound basis in biology, it is clear Swan favours cell cycle model development over other models and believes that development of the modelling of drug resistance in tumours will be of key importance in future work.

Cell population models

Several models containing both tumour and other cell populations are discussed. Swan reviews four pieces of work which model either two populations – tumour and normal cells – or a single population of tumour cells.
To some extent this work is similar to the kinetic models discussed below, but attention is drawn to a more sophisticated modelling of drug concentration and toxicity.

Kinetic tumour growth models

Classed as kinetic tumour growth models, the Gompertz, Verhulst and Cox equations are examined in detail. Key issues discussed are the nature of the cell loss term and drug concentration and delivery.
The cell loss term is related to drug concentration. Using a constant rate of cell loss proportional to the drug concentration assumes that experimental results carry over into human tumours – an assumption that Swan states is not accepted by many clinicians. An alternative term is proposed which limits the cell loss rate when the drug concentration reaches a (saturation) level.
The drug delivery and it's concentration at the tumour site is based on assumptions – instantaneous mixing with blood plasma and immediate delivery to the tumour. Incorporation of some delay into models is suggested.
Discussing the optimal control theory, Swan describes earlier work which examined cyclic chemotherapy for myeloma, a bone marrow cancer affecting the cells in blood plasma. Due to the toxicity of the cyclic treatment a case is then made for low level continuous drug infusion as a better solution, in contrast to dose/rest treatment.
Swan critically reviews works on optimal control as a tool for chemotherapy scheduling and highlights common assumptions including:

- Tumours are homogeneous
- No biological justification for Gompertz growth – just consistency with data
- Growth rate parameters are assumed – no human data available
- Drug kills a constant fraction of tumour cells, irrespective of tumour size
- Drug concentration is proportional to drug delivered to patient
- All tumour cells are susceptible to drug action (i.e. no drug resistance)
- Instantaneous mixing of drug in blood plasma
- Instantaneous delivery of drug to cancer site

Swan urges the exploration of other equations and theories. The lack of verification of growth models with clinical data is a recurring theme in the review.

5.4 Recent Work on Chemotherapy Scheduling

The following sections describe some of the recent work on the three model types discussed by Swan : Cell Cycle, Cell Populations and Growth Models.

5.4.1 Cell Cycle models

Estimation of cell-cycle parameters is considered by Swierniak et al.[36] and Panetta and Fister[24] investigate the effects of drugs on the immune system, aiming to maintain the White Blood Cell (WBC) count in the spirit of clinical oncologists who use WBC count to determine the next step in chemotherapy treatment.
Ochoa et al.[23] study chemotherapy optimisation using previous cell cycle specific drug modelling, [40] and improving the implementation of earlier work[39]. Three objective functions are modelled: Final tumour size; Average tumour size during treatment and Preference for shorter treatment period.
An Evolutionary Strategy Algorithm implements the model and the quality of solutions found is measured by:

- Area under the curve of drug dose i.e. total drug administered
- Tumour size deviation from the target level
- Immune system level above the threshold
- Diversity of solutions (deviation from average)

High computational cost precluded a comprehensive statistical analysis. The main conclusion is two-fold: results are dependent on the modelling system chosen and the desired goals in the the form of the objective functions.

It seems that the expansion of work in cell cycle modelling anticipated by Swan has been slow to get under way. Studies are largely confined to the use of a single drug. Ochoa et al. require no less than 22 parameter values to derive results, which together with the high computational cost for these mathematically and numerically intensive algorithms may explain the current level of research activity in this area.

5.4.2 Cell Population models

PharmacoKinetic (PK) and PharmacoDynamic (PD) Modelling by Iliadias and Barbolosi[11] considers a single drug and attempts to optimise drug administration in the first cycle of chemotherapy. Iliadias and Barbolosi echo Swan's concern over the difficulty of modelling drug toxicity. The approach is to use Pharmacokinetics (effect of body on drug) to model drug concentration and elimination. Two differential equations are used. One considers the mass balance in the blood plasma. The second equation models the drug elimination.

Pharmacodynamic (effect of drug on body) equations model the tumour growth using the Gompertz equation, the cell-loss term being directly proportional to drug concentration.
The toxicity equation models normal White Blood Cell (WBC) turnover with a cell-loss term for the toxic effect of drugs. This term is proportional to the WBC concentration and the drug concentration, but the drug concentration is time-shifted by a fixed amount to model the lag for reduced WBC production from bone marrow (5 days in the simulation).
Constraints are applied to limit drug concentration and total drug exposure.
Numerical solutions show an improvement over the clinical protocols used.
With a model requiring 18 estimated parameters and considering only a single drug dose, it may be the comparison with a clinical protocol is somewhat premature.

Matveev and Savkin[17] develop a theoretical model which includes the effect of the tumour on normal cells. Citing Afenya[1] and Zeitz[42] the effects of cytotoxic drugs on tumour and normal cells are considered explicity and also the effect of cancerous cells on normal cells. The model consists of differential equations for :

1. Tumour growth and regression by drugs
2. Normal cell growth and effect of drugs
3. Normal cell growth and effect of tumour

The tumour growth equation is Gompertzian and the sum of the effect of the drugs on the tumour is added.
There are several populations of normal cells and the second differential equation uses exponential growth and adds terms for cell loss due to drugs and cell loss due to the tumour, for each population.
The third equation covers the remaining normal cell populations where the effect of the tumour is taken into account, but the drug has no significant effect.

The constraints are :

- Drug concentrations at the tumour site kept below maximum levels
- Normal cell populations maintained above a required minimum level

Assumptions are mostly those enumerated by Swan, with the addition of two covering individual drugs damaging only specific cell populations of normal cells and that drugs have little effect on the cells attacked severely by the cancer cells. No justification is given for these assumptions.
The objective is to minimise the tumour at the end of the treatment period.
The analysis finds that the optimal treatment is cyclical for a single drug. Maintaining the maximum dose until the normal cell population falls to it's minimum

level, then allowing recovery followed by another period of high dosage, and repeating this cycle until the end of treatment. No problem is given to illustrate the theory.
Matveev and Savkin propose to investigate the multi-drug problem in future work and suggest the inclusion of pharmacokinetic equations to model the effects of infusion rates on drug concentration.

Models which involve populations of tumour cells and healthy cells interacting with each other and with chemotherapy drugs seem to have a better basis in biology than kinetic growth models. However analysis of single drug, single dose treatments with theoretical models is perhaps best described as research in the early stages.

5.4.3 Kinetic Growth models

McCall and Petrovski[18] present a decision support system (OWCH) for oncologists seeking novel chemotherapy regimens. The main theme is the constrained optimisation of multi-drug chemotherapy using a kinetic model with the Gompertz equation. The treatment schedule is encoded in the GA by a control vector of binary values. This is simply the dose for each drug for each time interval, coded as four bits representing a dose of between 0 and 15 units. Constraints are set to limit drug concentration and total exposure, tumour size during treatment. The toxic side effects on other organs are also constrained using a potency parameter for each drug. Two GAs are used. An objective function minimising tumour size determines if a cure is possible, defined as reducing the tumour to under 1,000 cells in three consecutive intervals. If this is unsuccessful a second GA determines a palliative regimen and estimates Patient Survival Time (PST). Both GAs apply penalties to their objective functions based on violations of constraints.
OWCH allows the oncologist to explore a variety of chemotherapy regimens using these GAs.

In[26] Petrovski and McCall further explore the problem of optimisation posed by multi-drug chemotherapy drug with GAs using exponential, von Bertalanffy, Verhulst and Gompertz growth equations.
The encoding used is essentially the same as in[18]. An alternative *delta* or *sign coding* is also described which is similar but allows a value of -1 to be utilised when searching for modifications to a treatment vector, although no details of its use are given.
A fitness function is required to measure effectiveness of treatment. This takes account of tumour size, constraint violation and whether the goal of the treatment is a curative (tumour eradication)or palliative (prolonging survival time) one.

Experiments are described which use all four growth models and both treatment goals. These are compared to a common chemotherapy schedule (CAF) and show better performance. (Although how the CAF values were obtained is not clear). The authors claim the various growth models do not affect the GA's ability to derive treatment regimens. There is no discussion on the relative quality of solutions found.

A Particle Swarm Optimisation (PSO) algorithm is compared with a GA by Petrovski et al.[28] in experiments simulating and optimising chemotherapy treatment schedules. The multi-drug kinetic model and binary encoded GA are essentially those used by the authors in [18]. The GA uses elitist, roulette wheel selection with two point crossover followed by mutation.
The PSO algorithm uses a population of 50 particles and iterates through a sequence of adjusting each particle's velocity on the basis of it's own best, the local best and the global best fitness, to produce a new generation. Parameters are set using guidance given by Trelea[38], equal weighting being assigned to the local and global terms.
Another version of PSO is also considered – a 'local best' algorithm which concentrates on a neighbourhood of 20% of the population for each particle.
The three algorithms are run 30 times, using the same initial populations as a basis to ensure a fair comparison. It is shown that the PSO algorithms find feasible solutions in considerably fewer generations than the GA and are of better quality in terms of fitness function values, the global PSO performing better than the local algorithm. Analysis shows the improvements in performance and quality found by the PSO algorithms are statistically significant.
The authors suggest that the disparate nature of the areas of feasible solutions within the problem search space may be responsible for the PSO algorithms' success. If good solutions are near the boundary of these regions, it may be that the GA recombination operations can push the search out of the feasible areas. Whereas the PSO tends to keep within the feasible solution regions utilising each particle's *memory* of its performance. Little comment is made on the 'local best' algorithm and it is not clear whether it offers any advantage over the global algorithm. Perhaps further investigations into its ability to find good local solutions would be worth pursuing.
Petrovski et al. conclude that PSO is a realistic tool for chemotherapy scheduling problems, although no further work on this seems to have emerged as yet.

5.4.4 GA performance tuning

Using a problem from their previous work[18] seven GA performance factors are identified by Petrovski et al.[29]. They are: penalty coefficients 1 and 2, probability of mutation, probability of crossover, selection pressure parameters 1 and

2 and a number of breaking points. Analysis of these factors for this problem of multi-drug chemotherapy optimisation using a kinetic model with gompertz growth shows that only the mutation and crossover probabilities are significant in GA performance.

Extending their previous work on GA performance[29] Petrovski et al.[25] state that in order to tune the factors affecting a GA's performance, an efficiency measure is needed. The number of generations Ψ to reach a solution is found to be a suitable parameter. Since Ψ is a random variable and the distribution of $\log(\Psi)$ is approximately Gaussian, this enables analysis of variance to be used to study GA factor tuning. Although the problem is said to be the same as that examined in[29], it is unclear if the GA is the same, as eight factors affecting efficiency are now listed: probability of crossover, probability of mutation, selection method, crossover method, mutation method, creep mutation step, fitness normalisation slope and population size.
The tuning process starts with screening to reduce the eight factors to only those which will affect the results with statistical significance. This reduced set of factors is then used in a regression model to find optimal values.
The re-coding of the problem from binary to integer is introduced as an advanced representation of the solution space. Strictly this converts the GA into an ES.
This chemotherapy optimisation problem ES is tuned using the above procedure. It is found that the mutation probability, creep mutation step and population size are the significant factors to tune and optimal values for these are derived. The results show the tuned ES performs better than the previously tuned GA, both in Ψ and fitness values achieved.
The authors conclude that whilst the ES is unaffected by the crossover rate, it is sensitive to mutation probability and step size and tuning these is worthwhile to improve performance.

Tan et al.[37] find that intermediate tumour size rather than final size can be a significant consideration in chemotherapy. A single drug kinetic model is described in terms of differential equations cited by Tan as due to Martin[15] using a Gompertz growth equation and constraints imposed by drug toxicity. It is aimed to reduce the tumour size by 50% every three weeks. This adds three constraints to the model and is described as a method to combat tumour cells becoming drug resistant, by keeping the tumour burden as low as possible, especially early in the treatment.
A multi-computer distributed EA for chemotherapy drug scheduling is implemented. Encoding is with pairs of variables representing dose level and starting day and the model is computed using numerical differentiation. An interesting feature of the distributed EA is the sharing of sub-populations between computers in the system with a migration rate of 0.5% between populations to simulate natural evolution niches and (infrequent) elite individuals. The two best solu-

tions found involve no drug application for the first 18 and 41 days respectively, but the authors make no comment on the clinical acceptability of doing nothing for this period of 'treatment'. Tan concludes that distributed EAs can find good solutions and that the unconstrained solution is better than the three constraint one and therefore reducing intermediate tumour size is not beneficial.

According to Clare et al.[4] the Goldie-Coldman[8] model of clonal resistance showed that by the time it is detected a tumour will have some drug resistant mutant cells present. This clearly means that there is no benefit in delaying drug application.

Liang et al.[13] follow Tan et al.[37] using a single drug kinetic model after Martin[15] but using an exponential tumour growth equation with constraints imposed by drug toxicity and three constraints to the model. The authors modify Martin's third equation for cumulative toxicity to allow for the reduction in total accumulated toxicity by metabolic drug elimination during treatment.
An EA with adaptive elitist crossover and mutation is described. The encoding is unusual and each chromosome consists of two parts: an initial dose then the repetition of a pattern e.g. a dose of followed by a short rest period repeated for the rest of the treatment period. This effectively builds in an additional constraint on the system, but giving the first drug dose on day one eliminates the 'do nothing' period which Tan et al. found.
Liang et al. justify the use of the kinetic model by claiming that *"some experts"* suggest that tumour characteristics such as the distribution of tumour cells among the phases of the cell cycle, the metabolic activity in the tumour, etc. can be accommodated by the adjusting values of parameters λ and κ described in section 5.3.4. The experts remain anonymous. It is also stated that *"anti-cancer drugs have been shown to kill cells by first-order kinetics"* to justify the use of the cell kill term in equation 5.1 in section 5.3.4. No evidence is offered for this view, which is diametrically opposite to that of Swan given in section 5.3.5.
Results show improvements over those previously reported by Tan and others by a substantial margin, reducing the tumour size to a few hundred cells.

5.5 Accelerating Evolutionary Algorithms

5.5.1 Fitness Inheritance

Smith et al.[32] suggest that if fitness evaluation is computationally expensive there is the possibility of avoiding this calculation for every member of the population. In contrast with Grefenestette and Fitzpatrick's 1985[10] proposal cited by[32] to only partially evaluate fitness of all the population, Smith suggests that the full evaluation of only part of the population may be beneficial. The simple idea presented is that a child is given the average fitness of the parents, described as *average inheritance*, to distinguish it from *Proportional inheritance*, a weighted average based on the contributions of genes from each parent. Tests conducted on a simple problem called *OneMax* where the fitness of a 64 bit binary string is the count of set bits in the string showed that efficient convergence is achieved even with only 10% evaluations. A more realistic problem is presented using an EA analysis of aircraft routing through a threat field. The fitness function is the length of the route with penalties added for threats detected. This is inverted to give a fitness maximisation goal. Experiments use proportional inheritance in an EA using roulette wheel selection and evaluation of all the initial population. With a population of 250, the best results are obtained by evaluating (surprisingly) only one individual per generation. It is suggested that the use of inheritance allows high selection pressure and still achieves convergence.
Sastry et al.[30] investigate convergence times, population sizing, optimal inheritance proportions and speed increases from a theoretical standpoint and verify and extend Smith's work[32]. However only the *OneMax* problem with average fitness inheritance is covered and so this work does not represent any substantial practical advance.
Ducheyne[5] cites a well-known suite of Multi-objective Optimisation (MO) problems by Zitzler (1999)[44] for testing GAs. Use of this suite establishes that the Pareto-front must be convex and continuous for fitness inheritance to be effective. It is suggested that running a GA for a few generations can determine if the Pareto-front is convex and continuous. If so, use of the weighted average fitness inheritance is recommended, as in a few MO problems it is better than average, although usually they are both the same.

The simplicity of fitness inheritance makes it an attractive option for improving EA efficiency.

5.5.2 Machine Learning Assistance

Baluja [2] proposes to combat the tendency of EA populations to gravitate to local rather than global solutions by the use of sub-populations instead of a single large population. On the basis of a binary encoding for a population of chromosomes, Population Based Incremental Learning (PBIL) creates a Probability Vector (PV) to represent the population. This PV has the same number

of features as the population chromosome has genes. Each PV feature value is the frequency of that bit in the population i.e. the probability of that bit being set. The PBIL algorithm conducts operations on the PV in the way GA operations are applied to individuals of the population.
The PBIL algorithm is generational:

- Initialise a PV with values of 0.5

  Repeat

- Generate a population of individuals from the PV, evaluate all fitnesses

- Use the fittest individual and the Learning Rate to update the PV with some probability, Mutation_Rate (typically 0.02) by a small fixed amount, Mutation_Shift (typically 0.05).

  Until a termination condition is met.

Mutation is important, particularly in the later stages of search, whilst crossover is more significant in the early stages.
Of the PBIL parameters of population, learning rate, mutation rate and mutation shift, only learning rate is discussed, since the other parameters correspond with those in GAs and Baluja expects tuning methods will carry over. The learning rate affects the speed with which the PV approaches a good individual. It also influences the areas of the search space to be investigated. Two extensions to the basic PBIL algorithm are suggested – using several individuals to update the PV instead of one, and using the poorest individual to update the PV by adjusting away from it. This second option may not be successful at the end of a search when the best and worst individuals are close together.
An extensive analysis is conducted of six PBIL algorithm variants and two GAs on 12 problems including the classic Travelling Sales Person and Bin packing problems.
Baluja concludes that the PBIL Algorithm benefits from simplicity of implementation and a reduction in the number of fitness evaluations and can outperform GAs in the range of problems examined.

In[27] Petrovski et al. compare the performance of a Genetic Algorithm (GA) and a PBIL algorithm (after Baluja [2])which uses a binary chromosome representation. The GA is that described in an earlier work [25] using fitness proportionate selection. Using the same population size (100) as the GA and a learning rate of 0.3, the PBIL algorithm achieves considerably better efficiency; finding feasible solutions in under 6,700 fitness function evaluations compared to the GA's 32,900. The value of fitness found in this experiment is not given, the comparative results being presented as a graph of run length distribution. A second experiment to compare the quality of solutions uses the tuned GA and an adjusted learning rate of 0.05 to find solutions after 200,000 fitness evaluations. The improvement in fitness values found by the PBIL algorithm is shown

to be statistically significant. Petrovski et al. state that *"this result can have practical implications."* But there is no mention of the extent of the trade-off which may exist between getting fast results in 7,000 evals and good results after 200,000 evals. So the practical implications are nebulous, due to lack of some absolute values.

According to Michalski[20] EAs can be inefficient for complex problems. This is partly due to the random nature of evolution. In contrast with Baluja's PBIL, where each generation stands alone, Michalski alternates Machine Learning (ML) to guide the production of each new generation from previous ones, with conventional EA methods to get the best of both worlds.
The main steps in the process are:

1. Generate an initial population by some method, or randomly.
2. Using a fitness function evaluation as criterion, identify two groups – one of High Performance (HP) and one of Low Performance (LP) in the population. Apply an ML method to distinguish between the groups. Generate new individuals and use ML to select HP individuals. Combine these with the previous generation HP group and select to form a new generation. Repeat until ML mode termination condition is met.
3. If LEM termination condition is met Stop.
4. Use EA methods of mutation, crossover and selection to generate a new population. Repeat until EA termination condition is met.
5. Repeat steps 2 and 4 alternately until LEM termination condition is met.

Methods in ML mode for generation and selection of individuals can be similar to EA methods for these operations.
Michalski concludes that ML can significantly improve EA performance and can also sequentially reduce the search subspaces which may contain a solution.

Jourdan et al.[12] cite Michalski[20] as part inspiration to address the need to accelerate EA solutions in multi-objective problems like large water distributions networks, which have complex and slow simulations to evaluate fitness. Finding a suitable quality measure is not trivial and the S metric by Zitzler[43] is used as a reference point for comparison. The largest pipe diameters (i.e. the maximum cost) and the smallest pipe diameters (i.e. the biggest head loss) are used to compute the metric. Adapting the LEM of Michalski[20] and a decision tree algorithm for machine learning necessitated addressing the problems of defining a training set for learning and how to use the rules generated by the decision tree to generate new population members for the EA. Focusing on one objective at a time and alternating between the two objectives simplified the

training set selection. Applying the (positive) learned rules produces new individuals for the next evolutionary phase. Experiments running a learning phase every ten generations show improved results in fewer evaluations. An interesting counter-intuitive conclusion is that the reduction in required evaluations seems to increase with the problem size.

It seems in the field of chemotherapy drug schedule optimisation, there are possibilities for improving efficiency by tuning GAs, fitness inheritance and machine learning methods.

Chapter 6

Conclusions

6.1 Discussion

Several methods of optimising chemotherapy regimens have been investigated. They are often computationally expensive. This project aims to find ways to reduce the time to find chemotherapy drug schedules, by conducting experiments as described earlier. These will be derived from the study of existing methods found in the literature, namely fitness estimation and machine learning techniques. Success will be measured in terms of reduction in fitness evaluations required in comparison with the control EA.

6.2 Conclusion

Experiments will determine whether these methods can reduce the number of fitness function evaluations and thus reduce EA computation time in the drug scheduling problem, while maintaining the quality of solutions.

Bibliography

[1] E. Afenya. Acute leukemia and chemotherapy: a modelling viewpoint. *Mathematical Biosciences*, 138:79–100, 1996.

[2] Shummet Baluja. Population-based incremental learning: A method for integrating genetic search based function optimization and competitive learning. Technical report, Carnegie Mellon University, Pittsburgh, PA, USA, 1994.

[3] Peter A. N. Bosman and Tanja Alderliesten. Evolutionary algorithms for medical simulations: a case study in minimally-invasive vascular interventions. In *GECCO '05: Proceedings of the 2005 workshops on Genetic and evolutionary computation*, pages 125–132, New York, NY, USA, 2005. ACM.

[4] Susan E. Clare, Faina Nakhlis, and John Carl Panetta. Molecular biology of breast metastasis: The use of mathematical models to determine relapse and to predict response to chemotherapy in breast cancer. *Breast Cancer Research*, 2:430–435, 2000.

[5] Els Ducheyne, Bernard De Baets, and Robert De Wulf. Is fitness inheritance useful for real-world applications? In C. M. Fonseca, editor, *Proceedings of the Second Conference on Evolutionary Multi-Criterion Optimisation, EMO 2003*, pages 31–42, Faro, Portugal, 2003. Springer-Verlag.

[6] Robert A. Gatenby and Philip K. Maini. Cancer summed up. *Nature*, 421:321, January 2003.

[7] Marco C. Goldbarg, Elizabeth F. G. Goldbarg, and Carlos R. A. Mendes. Selecting beam directions in radiotherapy with an evolutionary algorithm. In *SAC '08: Proceedings of the 2008 ACM symposium on Applied computing*, pages 1420–1421, Fortaleza, Ceara, Brazil, March 2008. ACM.

[8] J. H. Goldie and A. J. Coldman. A mathematical model for relating the drug sensitivity of tumours to their spontaneous mutation rate. *Cancer Treat Rep*, 63:1727–1733, 1979.

[9] J. H. Goldie, A. J. Coldman, and G. A. Gudauskas. Rationale for the use of alternating non-cross-resistant chemotherapy. *Cancer Treat Rep*, 66:439–449, 1982.

[10] J. J. Grefenstette and J. M. Fitzpatrick. Genetic search with approximate function evaluations. In *Proceeding of an International Conference on Genetic Algorithms and Their Application*, pages 112–120, 1985.

[11] Athanassios Iliadis and Dominique Barbolosi. Optimising drug regimens in cancer chemotherapy by an efficacy-toxicity mathematical model. *Computers and Biomedical Research*, 33:211–226, 2000.

[12] Laetitia Jourdan, David Corne, Dragan Savic, and Godfrey Walters. Hybridising rule induction and multi-objective evolutionary search for optimising water distribution systems. In *Fourth International Conference on Hybrid Intelligent Systems*, pages 434–439. HIS '04, December 2004.

[13] Yong Liang, Kwong-Sak Leung, and Tony Shu Kam Mok. A novel evolutionary drug scheduling model in cancer chemotherapy. *IEEE Transactions on Information Technology in Biomedicine*, 10(2):237–245, April 2006.

[14] F. MacDonald and C. H. J. Ford. *Molecular biology of Cancer*, chapter 1. BIOS Scientific Publishers Ltd., 1997.

[15] R. B. Martin. Optimal control drug scheduling of cancer chemotherapy. *Automatica*, 28:1113–1123, 1992.

[16] Miljenko Marusić, J. Bajzer, Zeljokoand Freyer, and Stanimir Vuk-Pavlović. Analysis of growth of multicellular tumour spheroids by mathematical models. *Cell Proliferation*, 27:73–94, 1994.

[17] Alexey S. Matveev and Andrey V. Savkin. Influence of tumours on normal cells and optimal chemotherapy regimens: The case of several drugs and toxicity constraints. *Mathematical Medicine and Biology*, 22:143–162, 2005.

[18] John McCall and Andrei Petrovski. A decision support system for cancer chemotherapy using genetic algorithms. In M. Mouhammadian, editor, *Proceeding of the International Conference on Computational Intelligence for Modelling, Control and Automation*, volume 1, pages 65–70, Vienna, Austria, 1999. IOS Press : ISBN 90-5199-474-5.

[19] John McCall, Andrei Petrovski, and Siddhartha Shakya. *Computational Intelligence in Bioinformatics*, chapter 12 Evolutionary Algorithms for Cancer Chemotherapy Optimisation. IEEE Press, 2007.

[20] Ryszard S. Michalski. Learnable evolution model: Evolutionary processes guided by machine learning. *Machine Learning*, 38(1-2):9–40, January 2000.

[21] L. Norton and R. Simon. Growth curve of an experimental solid tumour following radiotherapy. *J Natl Cancer Inst*, 58:1735–1741, 1977.

[22] L. Norton and R. Simon. Tumor size, sensitivity to therapy, and design of treatment schedules. *Cancer Treat Rep*, 61:1307–1317, 1977.

[23] Gabriela Ochoa, Minaya Villasana, and Edmund K. Burke. An evolutionary approach to cancer chemotherapy scheduling. *Genetic Programming and Evolvable Machines*, 8(4):301–308, December 2007.

[24] John Carl Panetta and K. Renee Fister. Optimal control applied to cell-cycle-specific cancer chemotherapy. *Siam J. Appl. Math.*, 60(3):1059–1072, 2000.

[25] Andrei Petrovski, Alexander Brownlee, and John McCall. Statistical optimisation and tuning of ga factors. In *Proceedings of the IEEE Congress on Evolutionary Computation*, volume 1, pages 758–764, Edinburgh, Scotland, 2005. IEEE Press.

[26] Andrei Petrovski and John McCall. *Soft Computing Techniques and Applications*, chapter Computational Optimisation of Cancer Chemotherapies Using Genetic Algorithms, pages 117–122. Physica Verlag, 2000.

[27] Andrei Petrovski, Siddhartha Shakya, and John McCall. Optimising cancer chemotherapy using an estimation of distribution algorithm and genetic algorithms. In *GECCO '06: Proceedings of the 8th annual conference on Genetic and evolutionary computation*, volume 1, pages 413–418, Seattle, Washington, USA., July 8-12 2006. ACM Press.

[28] Andrei Petrovski, Bhavani Sudha, and John McCall. Optimising cancer chemotherapy using particle swarm optimisation and genetic algorithms. In Xin Yao, editor, *Parallel Problem Solving from Nature VIII, 8th International Conference Proceedings, Lecture Notes in Computer Science*, volume 3242, pages 633–641. Springer–Verlag ISBN: 3-540-23092-0 ISSN: 0302-9743, 2004.

[29] Andrei Petrovski, Alex Wilson, and John McCall. Statistical identification and optimisation of significant ga factors. In *Proceedings of the Fifth Joint Conference on Information Sciences, JCIS 2000*, volume 1, pages 1027 – 1030, Atlantic City, USA, 2000.

[30] Kumara Sastry, David E. Goldberg, and Martin Pelikan. Don't evaluate, inherit. In Lee Spector, Erik D. Goodman, Annie Wu, W. B. Langdon, Hans-Michael Voigt, Mitsuo Gen, Sandip Sen, Marco Dorigo, Shahram Pezeshk, Max H. Garzon, and Edmund Burke, editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 551–558, San Francisco, California, USA, 7-11 2001. Morgan Kaufmann.

[31] Ian C. Smith, Steven D. Heys, and et al. Neoadjuvant chemotherapy in breast cancer: Significantly enhanced response with docetaxel. *Clinical Oncology*, 20(6):1456–1466, march 2002.

[32] Robert E. Smith, B. A. Dike, and S. A. Stegmann. Fitness inheritance in genetic algorithms. In *Proceedings of the 1995 ACM symposium on Applied computing ISBN:0-89791-658-1*, pages 340–345, Nashville, Tennessee, United States, 1995. ACM.

[33] Stephen L. Smith, Patrick Gaughan, David M. Halliday, Quan Ju, Mabil M. Aly, and Jeremy R. Playfer. Diagnosis of parkinson's disease using evolutionary algorithms. *Genetic Programming and Evolvable Machines*, 8(4):433–477, December 2007.

[34] Robert Souhami and Jeffrey Tobias. *Cancer and its Management*, chapter 3. Blackwell Science Ltd., 1998.

[35] George W. Swan. Role of optimal control theory in cancer chemotherapy. *Mathematical Biosciences*, 101:237–284, 1990.

[36] A. Swierniak, A. Polanski, and M. Kimmel. Optimal control problems arising in cell-cycle-specific cancer chemotherapy. *Cell Proliferation*, 29(3):117–139, 1996.

[37] T. C. Tan, T. H. Lee, J. Cai, and Y. H. Chew. Automating the drug scheduling of cancer chemotherapy via evolutionary computation. In *Proceedings of the 2002 Congress on Evolutionary Computing*, volume 1, pages 908–913, May 2002.

[38] Ioan Cristian Trelea. The particle swarm optimization algorithm: convergence analysis and parameter selection. *Inf. Process. Lett.*, 85(6):317–325, 2003.

[39] Minaya Villasana and Gabriela Ochoa. Heuristic design of cancer chemotherapies. *IEEE Transactions on Evolutionary Computation*, 8(6):513–521, December 2004.

[40] Minaya Villasana and Ami Radunskaya. A delay differential equation model for tumour growth. *Journal of Mathematical Biology*, 47(3):270–294, August 2003.

[41] Bo Yuan, Marcus Gallagher, and Stuart Crozier. Mri magnet design: search space analysis, edas and a real-world problem with significant dependencies. In *GECCO '05: Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 2141–2148, New York, NY, USA, 2005. ACM.

[42] S. Zeitz and C. Nicolini. Mathematical approaches to optimisation of cancer chemotherapy. *Bull. Math. Biol.*, 41:305–324, 1979.

[43] Eckart Zitzler. Evolutionary algorithms for multi-objective optimisation: Methods and applications. Master's thesis, Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, November 1999.

[44] Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary Computing*, 8(2):173–195, 2000.

Appendix A

Ethical Form

A.1 Completed Ethical Issues Form

Appendix B

Risk Form

B.1 Completed Risk Assessment Form