Optimisation of system dynamics models using genetic algorithms

Patrick E. McSharry [a,b,c,*],

[a] *Department of Engineering Science, University of Oxford, Oxford, OX1 3PJ, UK*

[b] *Mathematical Institute, University of Oxford, Oxford OX1 3LB, UK*

[c] *Commissioned by The World Bank, 1818 H Street, NW, Washington, D.C. 20433*

Abstract

System dynamics models provide a detailed mathematical description of the complex interactions between numerous variables. Software packages, such as *iThink* and *Vensim*, allow policy makers to easily adapt, implement and explore these models. By testing various scenarios through an analysis of the model with different parameter values it is possible to investigate and visualise the long-term effects of changes that could be implemented over the short-term. Genetic algorithms (GA) are shown to provide an efficient and accurate method for identifying optimal scenarios from among the vast number of possible scenarios that are available. The combination of the GA parameter search and human intuition can be utilised to arrive at better strategies for government policy. This approach to optimisation is demonstrated using a model of the malaria-control program in Bolivia.

* Department of Engineering Science, University of Oxford, Oxford, OX1 3PJ, UK
Email address: patrick@mcsharry.net (Patrick E. McSharry).

Preprint submitted to World Bank Report *11 May 2004*

1 Introduction

The availability of user-friendly packages, such as *iThink* [1] and *Vensim* [2], allows a wide range of practitioners to construct and implement system dynamics models. Their easy to understand building-block approach and visual appeal also allows non-experts to easily adapt and employ these models. At present, however, the ability of these packages to provide optimal solutions is limited.

The simplest procedure for identifying optimal solutions, the one currently used by the most advanced system dynamics software package, is to divide the range of each parameter value into discrete points (forming a grid) and to test each combination in turn. This procedure works for systems with a small number of parameters but quickly becomes intractable for larger systems. Another feature of complicated, nonlinear system dynamics models is that the parameter space often exhibits multiple equilibria and many optimisation techniques tend to give sub-optimal solutions.

In this paper, a stochastic optimisation technique based on a Genetic Algorithm (GA) is proposed as a means of addressing both (i) the dimensionality and (ii) multiple equilibria problems. Furthermore, this approach to optimisation does not rely on specific information about the problem (gradients, linearity or continuity) and has been shown to be highly efficient in comparison to competing optimisation techniques. In addition, the GA gives a set of solutions, which may be used to analyse the robustness of the optimal solution.

Providing policy makers with a means of exploring models built in order to facilitate their decision-making gives them the chance to better acquaint themselves with the likely effects of making small changes to the current system or completely switching to an alternative one. The large number of possible changes implies that it would take a long time for one person to investigate all the different scenarios. In contrast, in this paper, a technique for providing a set of optimal scenarios is proposed. Each scenario is specified by a collection of parameter values, to help guide the policy maker towards both relevant and viable outcomes. This approach will still allow for sufficient human intervention without losing the policy makers' interest because of the need for exhaustive searches.

The paper is structured as follows: Section 2 outlines the main challenges underlying optimisation of complicated models; Section 2.1 explains the dimensionality problem, in that many system dynamics models used for policy-making require the simultaneous optimisation of numerous parameters; Section 2.2 discusses the problem of multiple equilibria; Section 2.3 describes the effect of different possible dynamical regimes; Section 3 introduces the GA;

Section 4 introduces the malaria-control model used to illustrate the technique; Section 5 presents the results of applying the GA to the malaria-control model and Section 6 concludes the paper.

2 Optimisation challenges

The purpose of using an optimisation technique is to automatically identify optimal solutions. Usually each parameter value is associated with a range of acceptable values. The optimisation technique searches throughout this parameter space for solutions which optimise (maximise or minimise) some criteria. The main challenges when optimising complex models include: (i) dimensionality, (ii) multiple equilibria and (iii) dynamical regimes. Each of these is discussed in more detail in the following sections.

2.1 Dimensionality

Most system dynamics models used for policy-making are complex, often having more than ten parameters that need to be optimised. This means that the simple "brute force" approach of dividing each parameter range into D parts and sampling every possible scenario in turn is not feasible. For example, consider the example of the Bolivian malaria-control model (see Section 4) with $N = 17$ parameters. Suppose that each parameter range has $D = 100$ divisions, then the model must be simulated $D^N = 100^{17}$ times. To get an idea of the amount of time required, suppose that one single run of the model takes one second, then 100^{17} runs will take approximately 3×10^{25} years, which is much longer than the life of the universe ($\approx 10^{10}$ years). This "brute force" procedure is obviously not suited to large models.

2.2 Multiple equilibria

Traditional optimisation techniques rely on specific characteristics of the problem (e.g. gradients, linearity and continuity) to select the next sampling point. Optimisation over a number of parameters can be made difficult by discontinuities and many local maxima (Fig. 1). While it is easy to find the top of the peak once the search is started on that same peak, many techniques fail to discover neighbouring peaks with greater heights. In contrast, stochastic optimisation techniques such as the GA make no assumptions about the problem and the next sampling point is chosen based on stochastic sampling rules rather than deterministic rules.

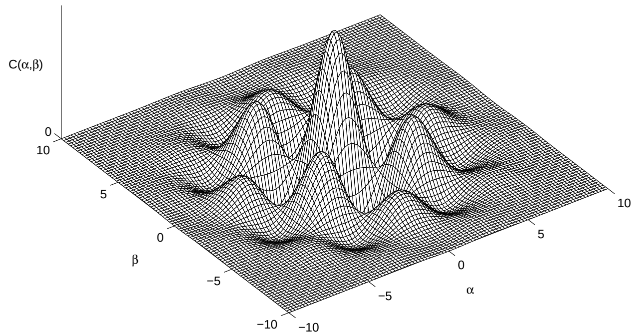

Fig. 1. *Example of multiple local maxima within the two-dimensional parameter surface $C(\alpha, \beta)$ surrounding the global maximum $(\alpha, \beta) = (0, 0)$.*

2.3 Dynamical regimes

Even the simplest model can produce complicated behaviour. This is especially true for nonlinear models which can gives rise to chaos [3]. When attempting to optimise the output of a given system, one must be aware of the rich variety of dynamics that are possible. The range of different behaviour is governed by the number of interacting variables or dimension of the model. Minute changes in the parameter values can cause a dynamical transition or bifurcation from one regime to another. For example, the model may have dynamics which converge to a fixed point, a limit cycle or chaotic dynamics. A small change in one parameter value can often lead to a bifurcation, whereby a fixed point becomes a limit cycle. The optimisation procedure should be complemented by checking each solution so that 'false solutions' due to limit cycle or chaos are avoided.

3 Genetic Algorithm

A Genetic Algorithm (GA) [4,5] provides a stochastic approach to optimising complex models and can efficiently solve each of the three optimisation challenges outlined above (Sections 2.1, 2.2 and 2.3).

The GA provides a faster and more accurate alternative to the currently best available optimisation procedure provided with the *Vensim* or *iThink* software packages ("brute force" optimisation). The GA can be used to search for an optimal set of solutions while attempting to optimise some specific criteria. GAs mimic aspects of biological evolution to find optimal solutions, namely (i) crossover of heritable information, (ii) random mutation and (iii) selection on the basis of fitness between generations. Trial solutions or parameter sets are encoded as data strings ('chromosomes'), which are scored according to how well the corresponding model satisfies the optimisation task. Better scenarios are given higher fitness scores. A population of strings, encoding different parameter sets, is made to evolve over a number of generations by random mutation and crossover between strings at each generation. The algorithm moves towards optimal solutions since the probability of survival to the next generation is dependent on its fitness score. The following basic steps present the general idea:

(1) Generate a population of scenarios; each encoding a string of uniformly distributed random numbers assigned to each parameter value in the model. The random numbers are chosen to be within the relevant range for each parameter.
(2) For each scenario in the population, evaluate its fitness level; a measure of how well the encoded model meets the optimisation criteria. The algorithm is terminated if a predetermined stopping criterion on the fitness is met. The scenario with the best fitness level is returned as representing the optimal scenario.
(3) Selection: by sampling with replacement, select a pool of scenarios with probabilities determined by fitness scores (generally the highest scoring scenario from the previous round is represented at least once).
(4) Crossover and Mutation: from this pool generate new scenarios by mixing parameters from two parents to form new scenarios. A (small) degree of normally distributed, zero mean, random variation is introduced to the parameters.
(5) The resulting scenarios form the next generation. Go to (2.) to calculate fitness.

Due to its random nature, the genetic algorithm can search the space of possible scenarios much faster than the "brute force" technique. For example, starting with a population of 100 scenarios and evolving these over 100 generations would require 100^2 runs of the model in contrast to the 100^{17} runs needed to solve the Malaria by the "brute force" technique.

The success of the algorithm in converging to an optimal solution can greatly depend on the size of the population and the choice of fitness function, stopping criterion, crossover and mutation schedules. The underlying model basis can be very general and can even increase in complexity during a computation, for example by including additional parameters once a certain fitness has been achieved, as a method of exploring very complicated models.

4 Malaria-control model

The GA approach to optimisation will be demonstrated using the malaria-control program in Bolivia. This particular example is relevant for a number of reasons: (i) the underlying link between the proposed policy actions to control malaria and their expected results are relatively well-known, (ii) there has been a dramatic decrease in malaria prevalence in Bolivia over the past five years, suggesting that the policy-makers are aware of which particular actions lead to the desired outcomes and (iii) there is a global interest in reducing malaria, exemplified by the Roll Back Malaria program.

The system dynamics model makes explicit the links between the policy actions pursued and the results obtained, highlighting the fact that the behaviour of malaria prevalence is largely predictable over time [6]. A monitoring and evaluation system could be used to inform the relevant authorities (e.g. the Ministries of Health and Finance) in sufficient time for them to allocate the resources necessary to keep the problem under control permanently. Currently, a lack of resources has given rise to the resurgence of the malaria problem. Utilisation of the model for long-term policy-making could help to prevent this situation and keep the incidences of malaria at a low level.

The model analyses the relations between the health sector policies and underlying epidemiological risks that together determine the level of malaria prevalence in the country. Nonlinearities in the model account for the feedback between infected individuals and the rate of transmission of the disease. The model also describes the Government's 2001-05 strategy and tracks the financing needs contemplated by the plan. The model used for describing the Malaria-control program in Bolivia enables the policy-maker to select from 17 variables when attempting to minimise the Annual Parasitic Index (API). Each of the 17 variables represents a possible financial expenditure (see Table 1 for details).

4.1 Background

In 2001, the Bolivian population amounted to 8,328,700 inhabitants, 3,139,805 (37.7%) of who lived in endemic areas, which were therefore vulnerable to malaria transmission. Approximately 75% of Bolivia is considered endemic area (see map), covering 67 municipalities where the risk of malaria transmission is high (API over 10/1,000 inhabitants) and 52 municipalities where the risk is medium (API of 2 to 9/1,000 inhabitants). In 1999, estimates were that the loss of economic production in Bolivia due to malaria, as a direct consequence of the disease and mortality of the economically active population was equivalent to 3% of GDP. High morbidity indices result in malaria being a public health issue for Bolivia.

Hence, the Department of Epidemiology of the Ministry of Health and Social Prevision has drafted the National Malaria Program, defining actions to improve malaria indicators. The Program contemplates the following aspects: (1) early diagnostic and treatment of sick people, (2) prevention and vector control measures, (3) research capacity, (4) implementation of an information and malaria surveillance subsystem, (5) consolidation of the program at the prefectural and municipal level, and (6) an information, communication and education strategy. Implementation of the Program during 1999-2001 has resulted in a significant decrease of the API, from 24.8 in 1998 to 4.9 in 2001. In 2001, 15,765 cases of malaria were detected, i.e. a significant decrease as compared to the 74,350 cases registered in 1998. Basically, these achievements are the result of the application of incisive combined strategies, such as the early diagnostic and treatment with follow-up, selective vector control and preventive communal actions, especially in the Amazon. These strategies are also applied in many other countries of the region, within the framework of the anti-Malaria initiative in the Americas that was launched by the PHO/WHO in 1998. The early diagnostic and immediate treatment approach have given rise to a reduction of malaria deaths from 29 in 1994 to 4 in 2000. However, it is recognized that there is a high degree of sub-registration of mortality due to malaria.

4.2 Current situation

In Bolivia, the following two parasitic species prevail in malaria infections: plasmodium vivax and plasmodium falciparum, the latter of which can give rise to serious malaria if it is not adequately treated. Besides the significant decrease of the API, from 24.8 in 1998 to 4.9 to 2001, the decreased prevalence of malaria caused by plasmodium falciparum with respect to the total positive cases reveals a change in the epidemiological distribution of malaria in Bolivia and a relative reduction in the seriousness of the situation (the percentage of people with malaria due to falciparum dropped from 15.4% in 1998 to 5.1% in 2001). If the API and Annual Vivax Index (AVI) are compared to the Annual Blood Test Index (ABTI), it becomes clear that the number of notified plasmodium vivax infections is decreasing with the number of examined blood samples, which reflects complete coverage of plasmodium vivax reservoirs.

Characterization of the factors that perpetuate transmission makes it possible to determine present and potential control measures. In order to ensure sustainability of these measures, selection of adequate methods still has to be improved and resources must be reoriented and reallocated. In spite of the

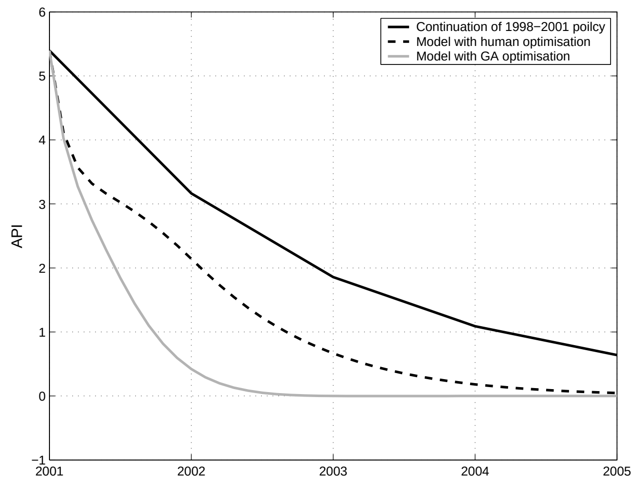

Fig. 2. *Comparison of predicted API (2001-2005) for three different scenarios: (i) continuation of the same policy used between 1998 and 2001 (black line), (ii) human optimised solution of the model (dashed line) and (iii) GA optimised solution of the model (grey line).*

irrefutable achievements in malaria control, official reports do not include any information on a-symptomatic cases of malaria, the prevalence of which is estimated to be even higher than notified cases of malaria.

The amount of resources assigned to the malaria control program has been increased in recent years, with an increasingly important presence of international funding. Depending on costs, concentrating efforts on control of the disease might be a more cost-effective investment than the traditional approach in the anti-vector fight.

5 Results

This section explores the performance and reliability of the GA solution, showing that the GA approach can be used to improve policy-making. The benefits of using the GA to optimise the model are demonstrated by comparing the predicted API values with other alternatives for obtaining a solution in Section 5.1. The structure of the optimal GA solution given by expenditure allocated to each of the 17 variables is described in Section 5.2. Finally, in Section 5.3, the robustness of the GA approach is demonstrated by exploring a set of solutions that give similar low values of API.

5.1 *Predictions of API*

The benefits of using a Genetic Algorithm (GA) to optimise the system dynamics model describing the malaria-control program is best illustrated by comparing three different policy-making scenarios: (i) continuation of the same policy used between 1998 and 2001, (ii) human optimised solution of the model and (iii) GA optimised solution of the model (Fig. 2). Predictions of API between 2001 and 2005 for each of these three scenarios demonstrates that:

- continuation of the 1998-2001 policy is sub-optimal
- the system dynamics model can be used to guide policy such as to further minimise API
- human optimisation of the model can be used to outperform the 1998-2001 policy
- GA optimisation of the model gives the best results

The initial 1998-2001 policy and the GA solution may be viewed as upper and lower bounds on the solutions obtained by human optimisation. But while the human effort could take many days of an experts time, the GA is sure to quickly (taking approximately an hour) provide the optimal solution. Furthermore, the GA optimisation requires only basic information about the model, such as the range of each parameter and the existence of any budget constraints.

5.2 *The GA solution*

Many investment strategies for allocating one million US$ were found to be capable of reducing the API to less than or equal 0.75 within one year. The optimal solution, giving an API of 0.42, is shown in Fig. 3. Note that this solution suggests spending most on Uniforms and Protective Equipment (index 13) and least on IEC Posters (index 4).

5.3 *Robustness*

The GA approach has been shown to provide a single good solution (Fig. 3), thereby guiding policy for reducing API. Given a set of solutions, all possible of reducing API, it is interesting to ask how similar these solutions are in terms of their suggested financial allocations. If all these solutions suggest a similar strategy, then one can be more confident that this policy is robust to small changes in the amounts allocated. Alternatively, if these solutions correspond

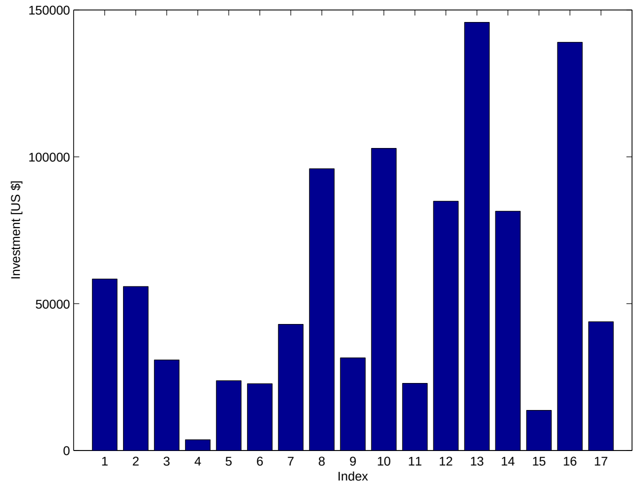

Fig. 3. *Investment solution for the malaria model showing the optimal way to divide the budget of one million US$ on the 17 investment targets. See Table 1 for a description of the investment target indices.*

to extremely different financial allocations, then this implies that the solution is extremely sensitive to the amounts allocated.

In order to establish the robustness of the GA optimal solution, all solutions which are capable of providing an API ≤ 0.75 within one year were investigated. An analysis of the mean and standard deviation of these 99 solutions with API ≤ 0.75 (Fig. 4 and Table 1) shows that all these solutions are in agreement about investing most of the finances, US$ 132396 or 13.2% of the entire budget, on Uniforms and Protective Equipment. The next highest expenditure is US$ 91631 or 9.2% of the budget on Laboratories. The smallest expenditure is US$ 26576 or 2.7% on IEC Posters. The means and standard deviations for the other investment targets suggest that these are approximately of equal importance, having been allocated approximately US$ 50000 or 5% of the budget.

Having confirmed the robustness of the solution, the policy-maker can be more confident that the predicted values of the API will be realised in the future. This ability to examine the sensitivity of the solution to small changes is extremely important in accounting for the uncertainty in the predictions.

6 Conclusion

The varied challenges surrounding optimisation of system dynamics models have been investigated. A range of different possible dynamical regimes, fixed points, limit cycles and chaos, complicate the process of optimising the output of the model. The number of parameters over which the optimisation

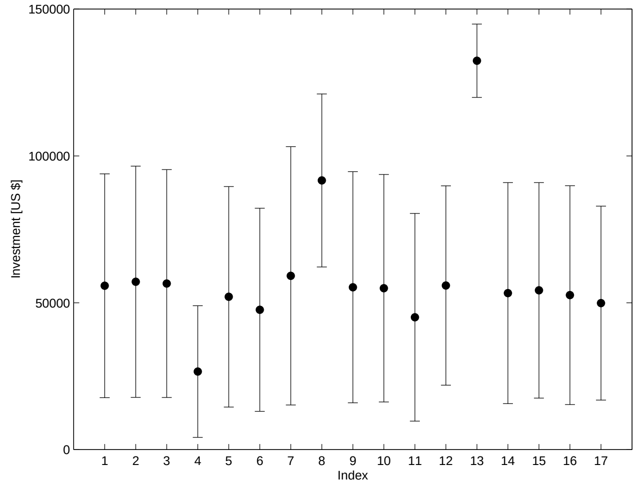

Fig. 4. *Robustness of investment solutions giving an API $\leq$ 0.75 within one year, showing one standard deviation above and below (error-bars) the mean (●). See Table 1 for a description of the investment target indices.*

Table 1
Analysis of solutions giving an API $\leq$ 0.75 within one year.

Index	Investment Target	Mean/US$	Standard Deviation/US$
1	Clorhidrate Mefloquine	55783	38096
2	Cloroquine	57145	39355
3	Fumigation Equipment	56546	38799
4	IEC Posters	26576	22409
5	IEC Radio	52031	37543
6	IEC TV	47592	34574
7	KO Tab Insecticide	59162	43960
8	Laboratories	91631	29448
9	Mosquito Nets	55274	39344
10	Piretroide Insecticide	54933	38720
11	Prevention Trained Volunteers	45065	35357
12	Primaquine	55870	33947
13	Uniforms and Protective Equipment	132396	12493
14	Quinine Biclorhidrate	53275	37636
15	Quinine Sulfate	54231	36679
16	Fumigation Technicians	52605	37251
17	Laboratory Technicians	49885	33011

takes place in many system dynamics models suggests that "brute force" approaches are not adequate. In addition, it may be extremely difficult, if not impossible to obtain information about the gradient of the model. For these reasons, a stochastic optimisation technique based on a Genetic Algorithm, was applied in order to determine a collection of solutions. A system dynamics model for malaria-control in Bolivia was used to demonstrate this approach. A comparison of the 1998-2001 policy, the human-optimised model and the GA-optimised model demonstrated that the GA approach can outperform the others in terms of providing more accurate results and being more efficient. An analysis of a collection of solutions in the form of investment strategies was used to determine the robustness of the optimal solution. This showed that the GA optimal solution is robust to small changes in the investment allocations.

The combination of user friendly system dynamics packages such as *iThink* and *Vensim* with a Genetic Algorithm for determining optimal solutions has important implications for policy making. Having easy access to relevant solutions makes it easier for the policy makers to explore and experiment with the model while incorporating their own intuition before deciding on a final plan of action.

Appendix A: Operations manual

The task of constructing and optimising a system dynamics model can be separated into four phases: (a) data collection, (b) *iThink* model construction, (c) *iThink* model conversion to *Matlab* and (d) GA optimisation. All four phases are explained in the following sections.

Data collection

The first step to obtaining a realistic model is to collect all the relevant data. This involves an analysis of the different dynamical variables (stocks in *iThink*) and the influences between them. For this reason there should be some iteration between this step and the model construction. An initial value will be required for each dynamical variable. This phase usually involves contact with experts in the field to establish the important stocks. Furthermore many of the parameters may vary slightly over time, e.g. prices. This information should also be provided to the model builder.

Model construction

The *iThink* package provides an intuitive, dynamic and graphic modelling approach. It facilitates the integration of a range of sub-models and sub-processes. The simplicity of the graphical interface and the ease of integration means that a number of different people can be involved in developing different sub-models at the same time. A thorough description of the *iThink* model for the malaria-control program is given in [6].

Model conversion from iThink to Matlab

The *Matlab* program `ithink2matlab.m` was developed in order to automate the conversion of a given *iThink* model into *Matlab*. This may be achieved by carrying out the following steps:

- construct an *iThink* model without using characters $, &, % or accents
- view the underlying equations of the *iThink* model
- choose **Select All** from the **Edit** menu
- copy and paste into the Microsoft package *Wordpad*
- select **No Wrap** from **View** > **Options** > **Word Wrap**
- save as an ascii file, e.g. `filename.txt`

Within the *Matlab* environment, the execution of the command:

```
ithink2matlab(filename, optparam)
```

performs the task of reading in the *ithink* model from `filename.txt` and uses the parameters listed in `filename.par` to generate a *Matlab* model `filename.m` and matlab derivative file `derivsfilename.m`. In addition, it generates a matlab script `optfilename.m`, which is a function of the parameters and gives the final value of the specific parameter chosen for optimisation.

GA optimisation

In the case of the malaria-control program, the file `malaria.m` may be executed using

```
F = malaria(a)
```

where `a` is a vector containing the values of the 17 financial investments. The program runs the model, evolving all the stocks over time. Inside the file `malaria.m`, a call to `optmalaria.m` calculates the values of the variable selected for optimisation, API. The output `F` gives the value of the API at the final time step of the model simulation.

In summary, the file `malaria.m` calculates the final value of the API for any given investment scenario that is encoded in the vector `a`. The aim is now to minimise `F` with respect to `a` while maintaining the constraint that the sum over all the values in `a` must equate with the financial budget.

The genetic algorithm (GA) is then used to optimise the results of running the script `malaria.m` with respect to different input parameters. This is achieved by running the command

```
[X, P] = ga(parspace,'malaria')
```

where `X` is the optimal GA solution and `P` is the population of solutions in the final generation of the GA. The matrix `parspace` gives the range of each of the 17 parameters used in the optimisation.

Appendix B: Future work

Many optimisation problems within the field of policy making require the solution to satisfy a financial restriction given by a budget constraint. The speed of the GA optimisation could be increased by constructing a genetic algorithm (GA) which directly incorporates the constraint imposed by such a fixed budget. Currently, the optimisation spends a lot of its time ensuring that this constraint is met. This could be achieved by only allowing the population of solutions to evolve within the space set by the budget constraint.

References

[1] www.hps-inc.com.

[2] www.vensim.com.

[3] E. Ott, T. Sauer, and J. A. Yorke. *Coping with Chaos.* John Wiley & Sons, Inc., New York, 1994.

[4] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning.* Addison-Wesley, Reading, Mass., 1989.

[5] L. Davis. *The Handbook of Genetic Algorithms.* Van Nostrand Reingold, New York, 1991.

[6] J. Newman, Velasco M. A., L. Martin, and A.-M Fantini. A system dynamics approach to monitoring and evaluation at the country level: an application to the evaluation of malaria-control programs in bolivia. In *Proceedings of the Fifth Biennial World Bank Conference on Evaluation and Development, "Evaluating Development Effectiveness: Challenges and the Way Forward*, Washington, D.C., 2003. Operations Evaluation Department.