

# Evolving Novel and Effective Treatment Plans in the context of Infection Dynamics Models: Illustrated with HIV and HAART Therapy

Rebecca Haines<sup>1</sup>, David Corne<sup>2</sup>

<sup>1</sup> SECaM, Harrison Building, University of Exeter, Exeter EX4 4QF, UK  
r.i.haines@ex.ac.uk

<sup>2</sup> MACS, Earl Mountbatten Building, Heriot-Watt University, Edinburgh EH14 8AS, UK  
dwcorne@macs.hw.ac.uk

**Abstract.** Several diseases involve complex interplay between an infection and the body's defences. Concerning AIDS, for example, this corresponds to developments in the immune system's responses and the HIV virus' counter-responses. Treatment for such diseases involves, at specific times, delivery of an agent that inhibits the infection. We hypothesise that: given a credible model of the combined dynamics of infection and response, the timing and quantities involved in treatment can be valuably investigated using that model. In particular, we investigate searching for optimised treatment plans with an evolutionary algorithm (EA). Our test case is a cellular automaton (CA) model of HIV dynamics, extended to incorporate HAART therapy (a favoured HIV treatment). An EA is wrapped around this model, and searches for treatments that maximally delay onset of AIDS, given certain constraints. We find that significant improvements over default HAART strategy are readily discovered in this way.

## 1 Introduction

Several viral and bacterial diseases are characterised in terms of the complex interplay over time between the infection and the body's defences. With AIDS, for example, three distinct phases have been observed, corresponding to developments in the immune system's responses to the attack, and the HIV virus' counter-responses. Any treatment strategy will involve intervention at specific times, with the delivery of an agent that is believed to somehow inhibit the course of the infection. We hypothesise that: given a credible model of the combined dynamics of the infection and the body's response, the timing and quantities involved in any treatment can be valuably investigated in terms of that model. In particular, we investigate the notion of searching for optimised treatment plans with an evolutionary algorithm (EA). Our test case is a cellular automaton (CA) model of the dynamics of HIV infection, which we extend to incorporate the modeling of treatment via HAART therapy (a currently favoured HIV treatment). A simple EA is wrapped around this extended model, and searches for treatment plans that maximise the time taken for AIDS to take hold, in the context of a constraints concerned with harmful effects of prolonged continuous therapy. We

find that, insofar as the models hint at reality, treatment plans that improve on default HAART strategies can be readily discovered by this method.

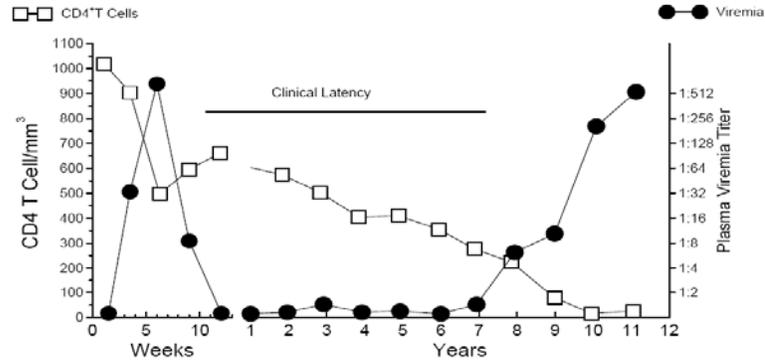
The paper is set out as follows. In section 2 we set out relevant introductory material on HIV/AIDS, mathematical and computational modeling thereof, HAART therapy. Section 3 then details the models and algorithms employed, aimed at providing a replicable account of the HIV and HAART CA model, and the EA that searches using this model within the fitness function. Experiments and their results are described in section 4, and a concluding discussion appears in section 5.

## 2 Background Material

**HIV, AIDS and HAART Therapy:** The Human Immunodeficiency Virus (HIV) has been subject to intense research over the last few years. Although major progress has been achieved in understanding different aspects of the virus-host interaction, the mechanisms by which HIV causes AIDS still remain unexplained. HIV can lie dormant for years, invisible to the body's surrounding immune defenses.

The following is summarised from [1—5]. HIV is sphere shaped with a diameter of 1/10,000 of a millimetre. Infection begins when an HIV particle enters the body and enters a CD4+ T lymphocyte cell. Since these cells are key to the immune response, this is a central reason for HIV's devastating impact. A small number of cells harbour HIV in a stable, inactive form, so that HIV can lie dormant for years, invisible to the body's immune defenses. Once in the cytoplasm of a cell, HIV reverse transcriptase converts viral RNA into DNA. This stage is important to drug therapy development and modeling as seven of the antiviral drugs approved in the US for HIV treatment (AZT, ddC, etc.) function by interfering with this stage of the viral life cycle. The HIV DNA travels to the cell's nucleus and is spliced into the host's DNA. Once incorporated into the cell's genes, HIV DNA is now a 'provirus'. The cell is no coerced into manufacturing more virus particles; viral RNA and associated proteins are packaged and released from the lymphocyte surface, taking with them a swatch of lymphocyte membrane containing viral surface proteins. These proteins will then bind to the receptors on other immune cells facilitating continued infection.

Broadly, the infection is characterized by three stages: In stage 1, lasting only a few weeks, there is an initial increase in the viral load followed by a sharp decrease. A large amount of HIV is in the peripheral blood and the immune system produces antibodies and cytotoxic lymphocytes. This stage is present in the left hand side of figure 1 which is plotted from clinical data. In stage 2, which can last 10 or more years, HIV level in peripheral blood drops very low but people remain infectious. HIV is not dormant during this stage, but is actually very active in the lymph nodes, wherein many T helper cells are infected and die and a large quantity of virus is produced. This stage is the mid-area of figure 1. Finally, in stage 3, the immune system starts to fail (see right end of figure 1). HIV becomes more prevalent and varied, leading to destruction of more T helper cells. The body fails to continually replace these. As the immune system continues to fail, mild symptoms develop. As the immune system deteriorates the symptoms worsen, resulting in the onset of AIDS.



**Fig. 1.** The graph of patient cell and viral measurements from [18]. Notice the change in horizontal scale following the 10-weeks point. The first 10 weeks show the initial phase of infection following which the density of healthy cells (squares) recovers, but the HIV virus ‘incubates’ for several years until an acceleration in viral count that heralds the onset of AIDS.

Currently there is no defined cure for HIV. However there are several drug therapies. A vaccine would be ideal, however, future prospects for this are poor due to HIV’s rapid mutation rate. For now, many scientists are instead concentrating on improving therapy. At present, there are fifteen drugs licensed for HIV treatment [6]. But, individual therapies each have their drawbacks, as, for example, HIV develops specific resistance to it. Currently, theory and clinical trials indicate that maximal viral suppression is through highly active anti-retroviral therapy (HAART), which consists of triple therapy with two nucleoside analogues and a protease inhibitor [6].

This study will concentrate on HAART as the drug therapy to model since at this time it is the most effective drug at reducing the time for onset of AIDS. However, due to the development of drug resistance, evolution of viral strains, or poor patient compliance, this combination of drugs fail to effectively contain the virus long-term in the majority of patients. Recent research has investigated Structured Treatment Interruptions [19, 20] in which treatment is interrupted by periods of several weeks. This will be further discussed later, and forms the basis of our experiments.

**Modelling HIV, AIDS and Potential Therapies:** Several mathematical models of HIV infection have been proposed (for example [8-10]), but all fail to describe some important aspects of the infection’s dynamics. Typically, models proposed so far have problems in maintaining biological plausibility at the same time as producing a qualitative match to the known dynamics of the concentrations of healthy and dead cells, the source of the main difficulty being that characteristic dynamics occur at two quite distinct timescales present in clinical observations (see figure 1, showing timescales over the first few weeks, and over several years, respectively). A recent example is [10], who integrate a CA approach with a graph percolation model, focussing on HIV’s long-term dynamics and the distribution of incubation periods, but this model does not account for the characteristic dynamics of the first few weeks of infection.

However, the recent model reported by Dos Santos & Coutinho [11], based on cellular automata, clearly shows the different time scales of the infection and has a

broad qualitative agreement to the density of healthy and infected cells observed *in vivo*. This ability to demonstrate the qualitative agreement over two time scales sets it apart as a promising candidate for further research, although [12] notes that this qualitative agreement is quite parameter sensitive. Nevertheless, Dos Santos and Coutinho's approach is clearly of interest (as is also pointed out in [12]), and unique in capturing the multi-timescale dynamics with apparently biologically plausible rules; we believe this model is sufficient as the basis of this study in terms of proving feasibility of our approach, while meanwhile we are investigating (with some success) improvements and alternatives to this model [13]).

There have been several attempts to model drug therapy and continued progress has led to the development of various types of therapy. *Multiple* drug therapy is considered to increase long term survival moreso than single-drug therapy. The appearance of virus resistance against HAART is apparently more delayed than with combined drug therapy [6], but long-term survival data with HAART is unknown. Mathematical models have difficulties with single treatment models and also have difficulties with unifying multiple therapies into one model [6]. This (as mentioned above) is again a difficulty in capturing the dynamics at different time scales. However, Sloot et al's system [6] echoes the nature of dos Santos and Coutinho's model and extends to model a variety of therapies; they show good qualitative agreement to clinical treatment data. We base our work on such a combined HIV/HAART model.

### 3 Models and Algorithms

**The Cellular Automaton / HAART Therapy Model:** First we describe our implementation of the HIV Infection model [11]; this is how our model operates during periods when treatment is *off*; further rules below, following [6], define how the model operates when simulated therapy is in operation.

As [11], we use a 2D synchronous CA with a square lattice, and the Moore neighbourhood (a cell's neighbourhood are its N, NE, E, SE, S, SW, W and NW neighbours). Cells can be in one of four states: healthy (H), infected-1 (I1), infected-2 (I2), or dead (D). Initially, cells are mainly H, but a small random proportion  $p_{HIV}$  are infected. We use  $p_{HIV} = .025$ , in common with [11]. 'HAART off' operation then continues according to the following four rules; one application of these rules to every grid cell results in simulating 1 week's progress. Biological justification of these rules is in [11]; we provide brief notes on this point.

**Rule 1** - If an H cell has at least one I1 neighbour, or if has at least  $k$  I2 neighbours, then it becomes I1. Otherwise, it stays healthy.

**Rule 2** – An I1 cell becomes I2 after  $\tau$  time steps (simulated weeks). (to operate this the CA maintains a counter associated with each I1 cell).

**Rule 3** - An I2 cell becomes D.

**Rule 4** – A D cell becomes: H, with probability  $p_{repl}(1 - p_{Infec})$ ; I1, with probability  $p_{repl}p_{Infec}$ ; otherwise, it remains D

Rule 1 models spread of infection by contact prior to development of immune response. As [11], we use  $k = 4$ . An I2 cell represents one against which an immune response has finally been developed (note that each infected cell will harbour a different HIV strain), which happens after roughly  $\tau$  weeks; as [11], we use  $\tau = 4$ . Rule 4 simulates the replenishment of new healthy cells in the context of maintained infection. In common with [11] we use  $p_{repl} = 0.99$  and  $p_{Infec} = 0.00001$ .

Now we turn to the integration of simulated HAART. The following rules, gleaned from [6], are designed to model the effects of HAART's anti viral-replication activity. These rules become operative in the model when therapy is switched on, replacing the corresponding rules above. Note that in the case of HAART-Rule 1, it is convenient to write and implement it as one which updates its *neighbours* on the basis of its own state, rather than updating a cell on the basis of its neighbours' states:

**HAART-Rule 1:** Given an I1 cell, the following is done to its neighbours at the next time step: with prob.  $p_{response}(t)$  (where  $t$  is current week of infection), a randomly chosen  $N$  ( $7 \geq N \geq 0$ ) of neighbouring H cells become I1; otherwise, all 8 become I1.

**HAART-Rule 3** - An I2 cell becomes dead after 2 time steps.

Meanwhile, rules 2 and 4 remain unchanged. HAART-Rule 1 models the principle that the drug inhibits viral replication, resulting in less efficient infection, with  $N$  representing the effectiveness of the HAART drugs. Without therapy, Rule 1 provides (implicitly) that all H neighbours of an I1 cell will become infected at the next time step. On therapy, with probability  $p_{response}(t)$ , this falls to  $N$ . The smaller  $N$ , the more efficient the therapy. Meanwhile, the larger  $p_{response}(t)$ , the more effective is the therapy. This models the response function for the HAART drug therapy, and incorporates the fact that the therapy will not immediately influence all of the I1 cells, but rather it will affect some of them at each time step [3]. From [3] we obtain a suitable  $p_{response}(t)$  function which is closely related to clinical data,  $p_{response}(t) = 0.95(1 - (t - 300)/300)$ . Also, given the increasingly high quality of the HAART drug, and observations that show relatively prolonged long-term survival for the highly inhibitive drugs [3], we can use  $N = 0$ . The model maintains the possibility for other values of  $N$ , which may become useful in prolonged treatments involving different flavours of HAART.

The change to rule 3 reflects the fact that, during HAART therapy, an I2 cell will take longer to die while viral propagation is inhibited. Finally, we note various parameters that have not yet been specified. We use a lattice size of  $700 \times 700$  - as noted later, this is a minimal size at which the basic HIV infection model yields quite repeatable results. Our CA boundary conditions are *fixed* - e.g. a cell at the right hand edge of the grid has no right hand neighbours. This is unlike the [11] case, but we feel it better models the physical conditions.

**Structured Treatment Interruptions:** Rather than prolonged, continual HAART treatment, many believe that structured treatment interruptions (STIs - scheduled periods without treatment) may *improve* the immune system by increasing HIV's sensitivity to antiretroviral drugs. Interruptions draw viral mutant selection pressure away from drug resistance, thus breeding a more drug-susceptible virus [5].

Implementing STIs is straightforward in our HIV/HAART model, and is done by switching between the normal ruleset while off treatment and the HAART ruleset during treatment. STI based therapy is not the same as Intermittent Therapy, which

has shorter cycles of switching on and off the antiviral drugs, such as ‘five days on two days off’ and has in some research proved ineffective and lead quickly to drug resistance. Meanwhile STI (with on/off periods measured in several weeks) has proven beneficial in trials. By optimising STIs we can investigate vastly more possible treatments that can be done with clinical trials.

**Optimisation:** We use a very simple evolutionary algorithm [14—17] (EA). With a population of 10 randomly generated solutions, it continues for  $\nu$  generations as follows: a parent is selected by binary tournament selection, and then subject to mutation. If the mutant is fitter than (or equally as fit as) the current worst in the population, then it replaces this worst. This represents a first-choice method, in the context of having very limited opportunity to experiment with several different algorithms owing to the very time-complex nature of the evaluation function in this case. On the basis of the results we feel both that this algorithm design suffices for now, and that alternatives will be worth investigating in future when we come to understand the landscapes involved in this and related problems.

## 4 Experiments and Results

A number of experiments were done to investigate optimizing STI plans using a simple evolutionary algorithm. Each candidate solution (*chromosome*) specified a schedule of treatment interruptions to start from the 300-week stage; this is the mean point at which the viral load is at a threshold that indicates HAART therapy should begin. The fitness function consisted of running the CA for 69 simulated weeks, using the chromosome to indicate when to simulate the therapy. Fitness itself is the week number at which the density of infected and healthy cells becomes equal (called the ‘set point’) and the healthy density does not recover past this point in later weeks. Such a point essentially heralds the onset of AIDS, while earlier weeks at which the ‘set point’ may be reached do not develop into AIDS owing to the introduction of treatment. There are various ‘repeating-pattern’ STI schemes that have been investigated in clinical trials, and limited simulations have found these tend to yield fitnesses between 400 and 500 in our model.

It is worth noting the time complexity of the experiments (which severely limited what we have been able to do so far). Given an  $N$  by  $N$  CA grid, a single iteration involves  $N^2$  cell updates, each of which involves several operations. A run of a black-box search algorithm for  $\nu$  evaluations would therefore have a complexity of  $690\nu N^2$  cell updates. Ideally, to compensate for the fact that multiple runs of the CA do not have identical behaviour, it would make sense for a single evaluation to return the mean result from  $s > 1$  simulations. But, initial tests and observations [11, 13] show that variation is also markedly reduced with grid size. Here we chose to raise  $N$  to a level where variation became acceptably low ( $N = 700$ ), allowing us to set  $s=1$ . Each evaluation still involved  $> 3 \times 10^8$  cell updates, and certain pragmatics and resource issues<sup>1</sup> mean that we can only here report on (at most) 50-evaluation runs. The main reason for high  $N$ , rather than  $s>1$ , is that lower  $N$  can lead to simulations that

---

<sup>1</sup> Now solved! If fortunate enough for this to be accepted, we will have runs of up to 1000 evals.

do not always follow the stages of infection, quickly settling down to an 'all-healthy' state. This would of course corrupt the evaluation of STI plans and the optimisation.

**Structured Treatment Experiments:** We report here on three experiments, differing in terms of the space of possible structured treatment interrupts (STIs) that were explored. Clinical studies show that there is a serious risk of constant HAART therapy leading to drug resistant HIV strains emerging. These then dominate the viral strand population resulting in *acceleration* of the onset of AIDS. It is thought that STI, in which constant prolonged therapy is avoided, can *harness* this to the patient's advantage. When the virus population is no longer subject to selection pressure from the drugs (i.e during a therapy-free period), the population reverts to non-resistant forms.

To broadly illustrate the issues, and modelling this for illustration via a binary string, where '1' means 8 weeks of treatment and '0' means 8 weeks without treatment, the following shows a possible 48-week treatment plan: '110000'; We will assume that treatment starts at week 300 following initial infection. In this simple plan, which is actually the standard HAART therapy plan (repeated each 48 weeks following consultation and review), there is 16 weeks' treatment followed by 32 weeks off treatment. The alternative 48-week plan '110110' potentially provides a dangerous level of treatment from the viewpoint of developing resistance, while a plan such as '000101' perhaps starts dangerously late. Of particular interest is the effect of such patterns over a much longer period, and the potential use of treatments without a repeating pattern, but structured to interact with the dynamics of the infection in such a way as to indefinitely delay the onset of AIDS.

We do not explicitly model the effect whereby resistant strains increase during treatment (and reduce in treatment-free periods), but we incorporate a bias against overly prolonged stretches of treatment by enforcing a limit on 'treatment-weeks' in each experiment. The problem we address is therefore to find optimised structured treatment interrupt (STI) plans under a maximal treatment-week constraint.

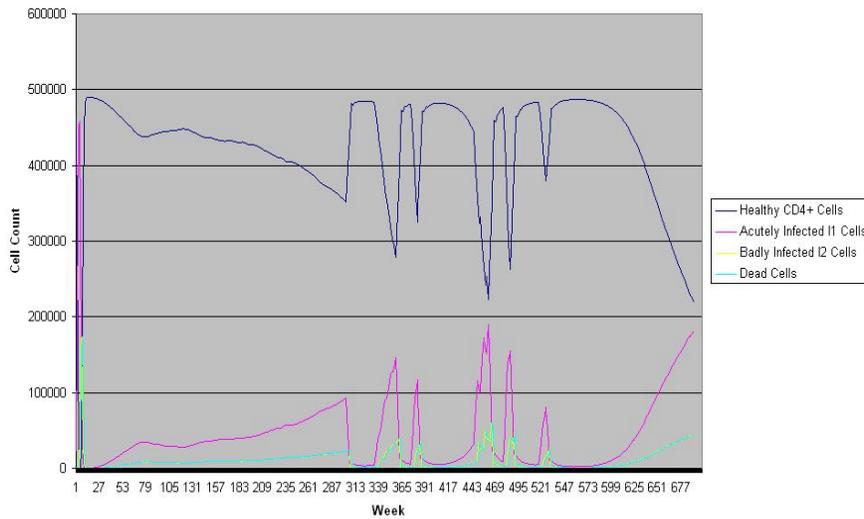
**Encodings and Operators:** In expt 1, an initial exploratory test, we evolved 10-bit binary chromosomes for 20 iterations (evaluations), where a '0' stood for 8 weeks treatment, and a '1' stood for 16 weeks off treatment. These were subject to a maximum of 6 stretches of treatment. Mutation was binary flip of a randomly chosen gene.

In expt 2, we used 20-bit binary chromosomes, where '0' stood for 16 weeks of treatment and '1' stood for 16 weeks without treatment, constrained to have precisely 6 1s. (hence searching a space of STI plans each spanning 6 years and involving precisely 96 weeks' treatment, which represents a clearly safe 'density' of HAART therapy). The 16 weeks on and 16 weeks off stretches correspond with practice in some clinical trials [19] In this case, mutation was a single swap of a randomly chosen '1' with a randomly chosen '0'. This experiment also ran for only 20 iterations.

The third experiment was designed after prolonged consideration of data from recent STI clinical trials. In short, the essential notion behind STI is that the off-treatment periods 'jump start' the immune system to take active control over the virus. However this revived defence can only last for short periods of time, hence a recent leaning in trials towards 8 weeks 'off' periods [20]. Meanwhile, this and other summaries and reviews indicate that a 16-week or longer period on treatment may be relatively safe; after 2 years on continuous HAART treatment, 10% of patients de-

velop resistance (in which the HAART drugs no longer have any effect), and this increases to 20% after 4 years and 30% after 6 years of continuous treatment. For the third experiment, we therefore evolved chromosomes of length 20 where a '1' represented 16 weeks of treatment, and a 0 represented 8 weeks off treatment. The number of 1s was maintained at 12, using the swap operator, and contiguous periods on treatment were limited to 80 weeks (5 1s). In this case we ran for 50 iterations.

In all cases, the chromosome's treatment plan is implemented from week 301 onwards, matching clinical practice in terms of the level of viral load at that point.



**Fig. 2:** Effects of an optimal STI treatment from expt 3. Infected cell counts do not start to approach the set point until > week 690 (though came close around week 460).

**Results:** In the first experiment, we decided to ignore results in which the set point was *never* reached. There were two such cases in experiment 1. Although representing an ideal solution, we felt it more likely to arise from random effects than a reliable effect of the evolved plan. The best solution had a fitness of 554, and encoded the following STI plan (for ease of comparing results of expts 1, 2 and 3 in these illustrations a 0 or 1 always refers to 8 weeks off or on respectively):

1100000010010000001.

This begins with a standard treatment but then continues without therapy for almost a year until two stretches of therapy separated by 16 weeks. This leads to a period of complex dynamics, which seems to result in a recovery of healthy cells and a year of stability until unrecovered-onset appears at around two years following the last stretch. We have to omit this figure for space considerations. In experiment 2, the best solution had a fitness of 578, and encoded the following STI plan:

000000001111000000000001111001111000000

We omit the plot for space reasons, but report similar dynamics in which the treatment seems to exploit the dynamics that occur when infected and healthy-cell counts are close, resulting in a series of healthy periods interspersed with near-onset periods.

Finally, in experiment 3, all solutions in the final population had fitness 690 (indicating no onset of the set point throughout the simulation). We could not therefore ‘ignore’ such cases as with experiment 1. Given that each of these solutions was from a random initial configuration, in a model of robust size, and involving slightly different chromosomes (treatment plans); meanwhile, the initial population in this case tended to have fitnesses around 550, owing to the enforced high density of treatment, so overall we consider this result to be robust. An example optimal solution is as follows, with the treatment shown in figure 3.

111111000111111111110001101111011

## 5 Concluding Discussion

In this preliminary investigation we have found support for the possibility of evolving STI plans that exploit the dynamics of infection and treatment as simulated in a model. The work reported so far has many limitations, but each has been or is being addressed in ongoing work. We have not yet evolved treatment plans that we are sure are robust (as would be evidenced, for example, by several simulations per evaluation); although 700 by 700 grids tend to produce repeatable results in the HIV model, the parameter sensitivity of this particular model is known [12, 13], and we do not really know the sensitivity yet of the *combined* HIV/HAART model as a function of grid size. Longer runs would also be desirable, of individual simulations, but in particular of the EA, to provide greater opportunity to find interesting and high-quality results.

Perhaps the most material limitation is the validity of the HIV/HAART model itself. The Dos Santos and Coutinho HIV infection model [11] is unique in providing a good qualitative agreement with real-data on cell counts, but has been criticised for its robustness [12, 13]. The qualitative agreement, arising from simple and apparently biologically plausible rules, is a considerable achievement, and the lack of robustness to parameters is not in itself a denial of the model’s validity (real disease dynamics is not necessarily ‘robust’ in this sense), but this is associated with questions about the way that the biological observation has been translated into rule probabilities and other attributes. For example, in the model, an infected cell always dies after 5 weeks, rather than different cells varying in their survival time around a mean.

However, we do not claim any great degree of validity for this model (and are developing others which seem more robust in related work [13]), but stress the promise of the overall approach, which consists of wrapping a complex disease/treatment dynamics model within an optimisation framework to yield high-quality suggested treatment plans for further consideration, and thus potentially narrowing down on novel strategies that may otherwise never have gone on to the stage of clinical trials. In ongoing work we are repeating and extending these experiments for longer runs, and including multiple simulations per run, and also plan working with encodings finer-grained stretches of treatment. On the modelling side, we are also considering

the use of more robust HIV/treatment models that may arise from an approach investigated in [13]. We also plan to investigate this general approach for other diseases and associated treatments, especially where a robust model exists that reflects all the relevant dynamics with good qualitative accuracy, *and* (unlike the situation with HIV) allows us to computationally quickly simulate or calculate a treatment plan.

## Acknowledgements

We thank Jonathan Owen for his preliminary research work involving the Dos Santos and Coutinho model, we thank Martin Oates for useful discussions about the model, and we thank Ron Yang for partial supervision of the first author.

## References

- [1] Laurie Garrett (1995) New View of How HIV Works , *Newsday* - January 12, 1995, <http://ww2.aegis.org/news/newsday/1995/>
- [2] The national institute of infectious diseases fact sheet <http://www.albany.edu/AIDS/howhiv.html>
- [3] Kevin Robert Frost, Amfar AIDS research- <http://www.amfar.org/cgi-bin/iowa/bridge.html>
- [4] Herbert W. Hethcote, Modeling HIV Transmission and AIDS in the United States, Springer-Verlag, 1992
- [5] Medic Eric S. Daar, M.D and Jay W. Marks, M.D. Human Immunodeficiency Virus (HIV Management), 2004 [http://www.medicinenet.com/human\\_immunodeficiency\\_virus\\_hiv\\_aids/article.htm](http://www.medicinenet.com/human_immunodeficiency_virus_hiv_aids/article.htm)
- [6] Peter Sloot, Fan Chen, and Charles Boucher (2002) Cellular Automata Model of Drug Therapy for HIV Infection, 2002, in *Proc. 5<sup>th</sup> Int'l Conf. On Cellular Automata for Research and Industry*, Springer LNCS 2493, pp. 282—293.
- [7] Kirschner, D. (1996) Using Mathematics to Understand HIV Immune Dynamics, [http://www.cnd.mcgill.ca/courses\\_mackey/pdf%20files/kirschner.pdf](http://www.cnd.mcgill.ca/courses_mackey/pdf%20files/kirschner.pdf)
- [8] Perelson, A.S. and Nelson, P.W. (1999) Mathematical analysis of HIV-1 dynamics in vivo, *SIAM Review*, **41**(1): 3—44.
- [9] Wodarz, D. and Nowak, M.A. (2002) Mathematical models of HIV pathogenesis and treatment, *Bioessays*, **24**(12): 1178—1187.
- [10] Kamp, C. and Bornholdt, S. (2002) From HIV infection to AIDS: a dynamically induced percolation transition?, *Proc. Royal Society B: Biological Sciences*, **269**(1504): 235—2040.
- [11] dos Santos, R., Coutinho, S. (2001) Dynamics of HIV Infection: A Cellular Automata Approach, *Phys. Rev. Lett.*, **87**(16): 168102-1, 4pp.
- [12] Strain, M.C. and Levine, H. (2002) Comment on “Dynamics of HIV infection: a cellular automata approach”, *Phys. Rev. Lett.*, **89**(21): 219805.
- [13] Corne, D., Frisco, P. (2006) Dynamics of HIV Infection studied with cellular automata and conformon-P systems, *Biosystems*, submitted.
- [14] Fogel, L.J., A.J. Owens, and M.J. Walsh (1966) *Artificial Intelligence Through Simulated Evolution*, John Wiley, New York.
- [15] Holland, J. H. (1975) *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor, MI.
- [16] Schwefel, H.-P. (1981) *Numerical Optimization of Computer Models*, John Wiley, Chichester, U.K.
- [17] Goldberg. D.E. (1989) *Genetic algorithms in search, optimization and machine learning*, Addison-Wesley.
- [18] Pantaleo, G., Graziosi, G. Fauci, A.S. (1993) *New England J. Medicine*, **328**(327).
- [19] Coffey, S. (2003) Structured Treatment Interruption (STI) HIV InSite's Coverage of the 10th Conf. on Retroviruses and Opportunistic Infections, *HIV InSite*, <http://hivinsite.ucsf.edu/InSite?page=cf10croi-04>
- [20] Franco Lori<sup>2</sup>, Andrea Foli and Julianna Lisziewicz (2002) Structured treatment interruptions as a potential alternative therapeutic regimen for HIV-infected patients: a review of recent clinical data and future prospects, *Journal of Antimicrobial Chemotherapy* **50**, 155-162.