Optimal HIV Treatment by Maximising Immune Response*

REBECCA V. CULSHAW
Department of Mathematics, Clarke College,
1550 Clarke Drive, Dubuque, IA 52001, USA.

SHIGUI RUAN
Department of Mathematics and Statistics,
School of Biomedical Engineering,
Dalhousie University, Halifax, Nova Scotia, B3H 3J5 Canada,
and
Department of Mathematics,
University of Miami, Coral Gables, FL 33124-4250, USA.

RAYMOND J. SPITERI
Department of Mathematics and Statistics,
Faculty of Computer Science,
Dalhousie University,
Halifax, NS B3H 3J5, Canada.

Abstract

We present an optimal control model of drug treatment of the human immunodeficiency virus (HIV). Our model is based upon ordinary differential equations that describe the interaction between HIV and the specific immune response as measured by levels of natural killer cells. We establish stability results for the model. We approach the treatment problem by posing it as an optimal control problem in which we maximise the benefit based on levels of healthy CD4+ T cells and immune response cells, less the systemic cost of chemotherapy. We

*Research supported by the Natural Science and Engineering Research Council (NSERC) and the Mathematics of Information Technology and Complex Systems (MITACS) of Canada.

completely characterise the optimal control and compute a numerical solution of the optimality system via analytic continuation.

Key words: HIV – Immune response – Optimal control – Cytotoxic lymphocytes – Numerical methods.

1 Introduction

The role of the immune response to human immunodeficiency virus (HIV) infection has received much attention in recent years. It is clear that some patients progress to AIDS much more rapidly than others, and the specific immune response to HIV has been shown to be an important determinant of the rate of disease progression (or non-progression).

When HIV invades the body, it targets the CD4+ T cells, often referred to as "helper" T cells. These cells can be considered "messengers", or the command centres of the immune system — they signal other immune cells that an invader is to be fought. The immune response cells, or cytotoxic lymphocytes (CTLs), are the cells that respond to this message and set out to eliminate infection by killing infected cells.

If the immune system is functioning normally, these components work together in an efficient manner and an infection is eliminated in short order, causing only temporary discomfort to the host. However, over time HIV is able to deplete the population of CD4+ T cells. The exact mechanism by which this occurs remains unknown, but several models have been suggested. For a variety of different hypotheses of how this occurs, we refer the reader to papers by Weber and Weiss [27], Nowak and May [21], Haseltine and Wong-Staal [11], and Kirschner, Webb, and Cloyd [17].

The impact of the depletion of CD4+ T cells on the host is that although the natural killer cells may be fit to perform their function of eliminating infection, they are never deployed. This then culminates in a clinical problem wherein the patient becomes vulnerable to infections that a healthy immune system would normally handle.

Although HIV does not target the CTLs directly, it has been noted clinically for some time (see Carr et al. [4], Cocchi et al. [5], Gray et al. [10], Arnaout, Nowak, and Wodarz [1], Musey et al. [20], Ogg et al. [22], Walker et al. [26], Weine et al. [28], Wodarz et al. [29], and references cited therein) that individuals who maintain a high level of CTLs remain healthy longer. The ideal clinical situation would be one in which the patient retains high levels of CD4+ T cells and low viral load. We also wish to maintain a positive population of CTLs so as to ensure that if viral load does rebound, the immune system is able to respond; however the precise level of CTLs required is dependent on the actual viral load. The best drug treatments should establish this result, while keeping adverse effects to a minimum. Virtually all anti-HIV drugs have many common adverse effects (see Lippincott's Nursing Drug Guide [18], for example, for an extensive listing of such adverse effects). Because of this and the fact that these drugs are so costly, it would be best to administer the lowest amount of drug necessary to keep the CD4+ T levels high.

New drug treatments and combinations of drugs are under constant development. The optimal treatment scheme for HIV-positive patients remains the subject of intense debate. The role of the immune response renders the problem particularly challenging. When models do not consider the immune response, they are in essence assuming that there is no significant natural anti-HIV response in the absence of treatment. In other words, the *only* anti-HIV (or pro-CD4+ T cell level) benefit is assumed to come from the drug or drugs administered. In reality, individuals do have a natural anti-HIV response that varies from patient to patient depending on a number of factors (one of which is severity of illness). Presumably, the stronger this natural killer response, the less drug is needed to control infection.

To date, many mathematical models of drug treatment of HIV have been developed. Some models in the mid-1990s focused on modelling AZT treatment (see, for example, Kirschner and Webb [15], or Kirschner and Perelson [13]). Some authors have proposed mathematical models of HIV treatment using control theory. Kirschner, Lenhart, and Serbin [14] established optimal treatment schemes for a scenario in which the drug reduces the rate of viral production. Optimal treatments were, in general, monotonically decreasing over the time interval of treatment. In the case where treatment was initiated 800 days after infection, there was a small peak in the drug level very shortly after initiation, after which the control was monotonically decreasing. In both other scenarios (1000 and 1200 days after infection), the optimal control was monotonically decreasing. Balancing effects to CD4+ T cell counts with drug cost, the earliest treatment was always the best no matter the length of treatment interval.

Fister, Lenhart, and McNally [8] established results for a system similar to that analysed in [14]. Their control was represented by a drug that reduced the infectivity rate. Their results indicated that strength of treatment should balance with duration (that is, the longer the treatment length, the smaller the dose should be). As well, optimal treatment schemes are monotonically decreasing.

In 1998, Wein, Zenios, and Nowak [28] constructed a model which allowed for viral mutation and the ability for the clinician to change treatment at any time during treatment (hence their use of the term "dynamic"). Their model assumed treatment that corresponds to different combinations of reverse transcriptase inhibitors. The authors used what they refer to as a "perturbation technique" to conclude that such dynamic treatment protocols are far preferable to static protocols.

However, none of these papers has included immune response as a specific component. In 1999, Wodarz and Nowak [30] published a four-dimensional ordinary differential equation (ODE) model of the interactions between CD4+ T cells, viral load, and immune response (both precursor or "memory" cells and effector immune response cells). They assumed that treatment negatively affects the population of immune response cells and modelled treatment numerically by running simulations of their ODE model with a parameter that reduced viral infectivity to represent treatment. Their conclusions were that interruption of therapy (to allow the immune response to rebuild after being suppressed by chemotherapy) or antigenic boosts to the immune system would be beneficial to the long-term clinical outcome of the patient.

This paper establishes an optimal control model of HIV treatment, using a single drug that reduces the cellular infection rate and explicitly incorporates the specific anti-HIV immune response as represented by levels of effector and memory CTLs.

In Section 2 we present an untreated model that is based upon that presented by Wodarz and Nowak [30]. We find three equilibria and completely analyse their local stability properties. In Section 3 we present the optimal control problem in which the coefficient of the cellular infection rate is the control. We seek to maximise the performance index, which is the benefit based on CD4+ T cell and CTL levels less the systemic cost of chemotherapy. We characterise the optimal control using Pontryagin's Maximum Principle. In Section 4 we solve the resulting optimality system numerically. In Section 5 we discuss the clinical implications of the results established in this paper.

2 Presentation of an Untreated Model

Here we introduce the ODE modelling of the immune dynamics of an HIV-infected immune system. We note that these equations model an *untreated* individual; treatment will be introduced in the next section via an optimal control.

The system is defined as follows:

$$\frac{dx}{dt} = \lambda - \delta x - \beta xy \tag{2.1}$$

$$\frac{dy}{dt} = \beta' xy - ay - \rho yz \tag{2.2}$$

$$\frac{dz}{dt} = cxyz - hz. \tag{2.3}$$

Variables are defined as follows: $x(t)$ and $y(t)$ are populations of uninfected and infected CD4+ T cells at time t, respectively. We consider viral load as proportional to levels of infected cells, since according to Arnaout et al. [1], "free virus is thought to be short lived relative to infected cells". $z(t)$ is the population of immune response cells at time t. In this model, we consider a single pool of immune response cells. The stability analysis we perform in the following section indicates that our system behaves qualitatively very much the same as that in [30].

Our parameters are interpreted as follows: λ is the source term for healthy CD4+ T cells, δ is their death rate, and β is the rate at which they are infected by virus. In this case, we consider the viral source to be directly from infected cells. The ratio of $\beta' : \beta$ is the proportion of infected cells that survive the incubation period (the time between the new infection of a CD4+ T cell and the time it becomes infectious). Henceforth in this paper, we assume that this ratio is $1 : 1$, or, in other words, that all infected cells survive incubation; but we note that this may not always be the case in reality. a is the death rate of infected cells by means other than elimination by CTLs, and ρ is the rate at which they are killed by CTLs. c is a generation constant for the CTL pool. Since it is an immune response specific to HIV, it is proportional to $y(t)$, the term representing infection level. It is also dependent upon healthy CD4+ T cell help, and levels of CTLs themselves, hence the trilinear term. Finally, h is the death rate for CTLs. Parameter ranges used in simulations, as well as their references, are in Table 1 at the end of this section. Following the analysis in [1] and clinically cited viral load ranges for HIV-positive individuals, we assume that viral load is approximately $(10^6 - 10^9) \times y(t)$.

2.1 Stability Analysis

We find that this system has three equilibria. They are

$$E_0 = \left(\frac{\lambda}{\delta}, 0, 0\right)$$

$$
\begin{aligned}
E_1 &= \left(\frac{a}{\beta'}, \frac{\lambda\beta'}{a\beta} - \frac{\delta}{\beta}, 0\right) \\
\overline{E} &= \left(\frac{\lambda c - \beta h}{c\delta}, \frac{h\delta}{\lambda c - \beta h}, \frac{\beta'(\lambda c - \beta h)}{\rho c\delta} - \frac{a}{\rho}\right).
\end{aligned}
$$

The first is an uninfected equilibrium corresponding to maximal levels of healthy CD4+ T cells and no infected cells or immune response. While at first glance the lack of immune response may seem alarming, we note that the immune response we are modelling here is that which is *specific to HIV*; therefore, in the absence of infection, we should expect no specific immune response.

The second equilibrium E_1 corresponds to positive levels of both healthy and infected cells, but no immune response. Clearly this is not desirable.

The interior equilibrium $\overline{E}$ corresponds to positive levels of all three components—healthy and infected CD4+ T cells, and immune response.

Since the system is three-dimensional, we use the Routh-Hurwitz criteria to establish negativity of the real parts of the roots of the characteristic equation, and therefore stability of the equilibrium under consideration. The general Routh-Hurwitz criteria for an n-dimensional system can be found, for example, in Appendix 2 of Murray [19]. For a three-dimensional system, with a characteristic equation of the form:

$$v^3 + a_1 v^2 + a_2 v + a_3 = 0,$$

the Routh-Hurwitz criteria state that all roots of the characteristic equation have negative real parts (and thus the equilibrium is stable) if and only if

$$a_1 > 0, a_3 > 0, a_1 a_2 > a_3.$$

Analysis of the characteristic equation evaluated at E_0, the healthy equilibrium, reveals that it is stable if $\beta' < \frac{a\delta}{\lambda}$; i.e., the infection rate and/or the fraction of cells surviving incubation is quite low.

The interior equilibrium $\overline{E}$ is stable exactly when the equilibrium E_1 is unstable, and vice-versa. In other words, a transcritical bifurcation occurs whereby the stability of the two equilibria switch. Specifically, $\overline{E}$ is stable so long as the following condition holds:

$$\beta' > \frac{ac\delta}{\lambda c - \beta h}.$$

Note that this is also the feasibility condition for the interior equilibrium; i.e., the necessary and sufficient condition for the equilibrium to exist. When this inequality is reversed, $\overline{E}$ is unstable and E_1 is stable. However, under most realistic parameter ranges we find $\overline{E}$ to be the stable equilibrium; see Table 1 at the end of this section for a complete listing of the relevant parameter ranges used as well as their sources. Also, $\overline{E}$ is in fact a spiral point, as the Jacobian evaluated at $\overline{E}$ has one real and two complex conjugate eigenvalues.

We summarise the above in the following proposition.

Proposition 2.1 *The uninfected equilibrium E_0 of the system (2.1)–(2.3) is stable for $\beta' < \frac{a\delta}{\lambda}$. When this inequality is reversed, either E_1 or $\overline{E}$ is stable, depending upon the parameter values. Specifically, for $\beta' < \frac{ac\delta}{\lambda c - \beta h}$, E_1 is stable and $\overline{E}$ is not feasible. When this inequality is reversed, E_1 loses stability and $\overline{E}$ becomes a locally asymptotically stable spiral point.*

We ran numerical simulations using the XPP package [7] for phase-plane analysis of systems of ordinary differential equations, and the results are displayed below. In Figure 1, we observe a plot of the solution $(x(t), y(t), z(t))$ versus time and see that it quickly settles toward a steady-state.

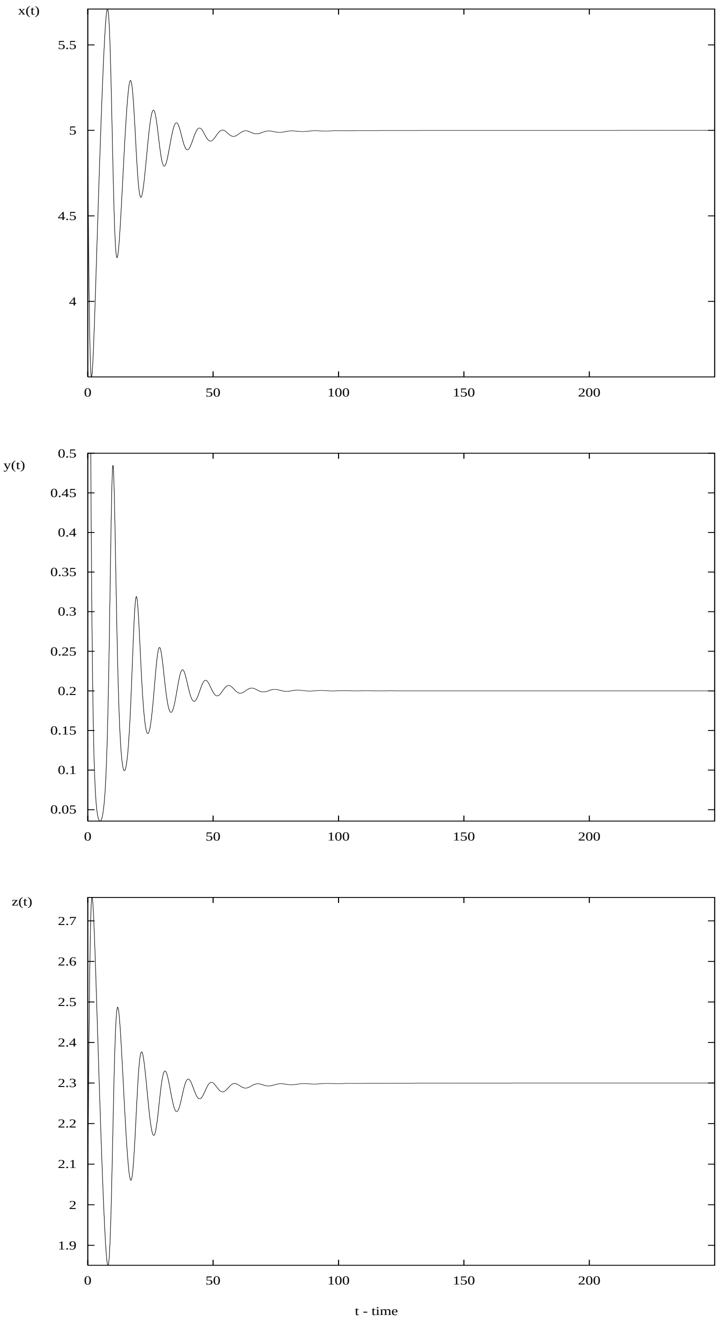

Figure 2.1: $x(t), y(t)$, and $z(t)$ converge to steady-state values.

2.2 Variables and Parameter Ranges Used

Parameter ranges used for numerical simulations of (2.1)–(2.3) are given in Table 1. The parameter ranges for $\lambda, \delta, \beta, a$ were obtained from references [1], [28], [30], [23], [25], [16]. The ranges for the immune response parameters ρ, c, h, were found in [1] and [30]. A wide range of possible parameter values has been suggested for HIV modelling. Part of the reason for this is that it is difficult to assign one set of parameters to individuals showing dramatically different clinical outcomes.

Parameters specific to HIV are generally given in units of cells mm^{-3} day^{-1}. Realistic levels of CD4+ T cells can be found using parameters within the ranges given in Table 1. Note that we also obtain a CD4+ T cell:CTL ratio of about $2:1$. In an individual with a healthy immune system, the ratio of CD4+ T cells to CTLs is usually about $2:1$ as we see here. However, this ratio switches to $1:2$ in advanced HIV infection. Our steady state values do indeed show an increase of CTLs and a decrease of CD4+ T cells from their equilibrium values, although the ratio at equilibrium is not $1:2$. This may be partly due to our implicit assumption that the pool of CTLs represented by $z(t)$ is specific to HIV and does not include any other components of the host immune response. Assuming viral load to be proportional to infected cells and that it is approximately $(10^6 - 10^9) \times y(t)$, we obtain a viral load in the $1,000 - 100,000$ range. This is not unreasonable, considering the nature of the viral load test — "viral load" is an approximation based upon the polymerase chain reaction (as opposed to "viral burden", the actual number of viral particles per millilitre of plasma, a quantity which is much harder to measure).

In the simulations reproduced here and in Sections 4 and 5, we use parameters from [30] to display the qualitative behaviour of the systems in parameter space. We note that the stability properties of the models are retained using most values from the ranges given in Table 1.

As a final note, we point out that the main goals of this study are mathematical and qualitative in nature. Much larger and more quantitatively sophisticated models are needed to determine *precise* treatment regimes in terms of days on therapy, exact drug quantities, and so forth. Such models are however a logical extension of the model presented here.

3 The Optimal Control Problem

We would like to maximise levels of healthy CD4+ T cells, as well as levels of CTLs (immune response cells). Also, we want to keep cost—as measured in

Table 1: Variables and Parameters

	Parameters and Variables	*Values*
	Dependent Variables	
$x(t)$	uninfected CD4+ T cell population size	
$y(t)$	infected CD4+ T cell population size/viral load	
$z(t)$	immune response cell population size	
	Parameters and Constants	
λ	source rate of CD4+ T cells	$1 - 10$ cells/day
δ	decay rate of healthy cells	$0.007 - 0.1$ cells/day
β	rate CD4+ T cells become infected	$0.00025 - 0.5$ cells/day
$\beta' : \beta$	proportion of infected cells surviving incubation	≈ 1
a	death rate infected CD4+ T cells, not by CTL killing	$0.2 - 0.3$ cells/day
ρ	rate at which infected cells are killed by CTLs	1/day
c	immune response activation rate	$0.1 - 1$/day
h	death rate of CTLs	$0.1 - 0.15$/day

terms of chemotherapy strength, a combination of duration and intensity—as low as possible. Our control is a function $u(t)$ with values normalised to be between 0 and 1, where $u(t) = 1$ represents totally effective chemotherapy and $u(t) = 0$ represents no treatment. We choose as our control class:

$$U := \{u(t) : u \text{ is Lebesgue-measurable with values between 0 and 1}\}.$$

Mathematically, the optimal control problem is formulated as:

$$\max J[u] = \int_0^T \left(x + z - \frac{Bu^2}{2} \right) dt \tag{3.1}$$

subject to the state system

$$\frac{dx}{dt} = \lambda - \delta x - (1-u)\beta xy \tag{3.2}$$

$$\frac{dy}{dt} = (1-u)\beta' xy - ay - \rho yz \tag{3.3}$$

$$\frac{dz}{dt} = cxyz - hz. \tag{3.4}$$

Since the control reduces the viral replication rate, we multiply our infectivity term βxy by $(1-u)$. In this case, both our cellular infection rate and our viral (infection) production rate are represented by the same term, β, so the drug may represent either a protease or a reverse transcriptase inhibitor drug (see Perelson and Nelson [24] for a detailed description of how to mathematically model different types of anti-HIV drugs).

3.1 Characterisation of an Optimal Control

We invoke Pontryagin's Maximum Principle [9] to determine the precise formulation of our optimal control $u^*(t)$. To do this, we note that our Hamiltonian is given by

$$H = x + z - \frac{Bu^2}{2} + \lambda w_1 - \delta x w_1 - (1-u)\beta xyw_1 + (1-u)\beta' xyw_2$$

$$-ayw_2 - \rho yzw_2 + cxyzw_3 - hzw_3 + v_1 u + v_2(1-u).$$

Here $v_1(t)$ and $v_2(t)$ are penalty multipliers ensuring that $u(t)$ remains bounded between 0 and 1. We also have that $v_1(t)u(t) = 0$ and $v_2(t)(1 - u(t)) = 0$ at the optimal $u^*(t)$.

The $w_j(t),\ j = 1,2,3$, are our adjoint variables; they determine the adjoint system which, together with our state system, determines our optimality system.

We shall consider all possible values for the control, including those on the boundary ($u = 0$ and $u = 1$).

(i) Consider the set $\{t : 0 < u(t) < 1\}$.

Pontryagin's Maximum Principle states that the unconstrained optimal control $u^*(t)$ satisfies

$$\frac{\partial H}{\partial u^*} = 0.$$

So we find $\frac{\partial H}{\partial u}$ and solve for u^* by setting our partial derivative of H equal to zero. Thus,

$$\frac{\partial H}{\partial u^*} = -Bu^* + xy(\beta w_1 - \beta' w_2) + v_1 - v_2 = 0$$

$$\Rightarrow Bu^* = xy(\beta w_1 - \beta' w_2) + v_1 - v_2$$

$$\Rightarrow u^*(t) = \frac{xy(\beta w_1 - \beta' w_2) + v_1 - v_2}{B}.$$

So we find that in this case, where $v_1(t) = v_2(t) = 0$, our optimal control is characterised as:

$$u^*(t) = \frac{xy(\beta w_1 - \beta' w_2)}{B}.$$

To completely characterise $u^*(t)$, we must consider the boundary cases $u^* = 0$ and $u^* = 1$ as well as the non-boundary cases.

(ii) Consider the set $\{t : u(t) = 0\}$. In this case, $v_2 = 0$. Thus, from the definition of the optimal control above, we have

$$0 = \frac{xy(\beta w_1 - \beta' w_2) + v_1}{B}.$$

Since (by definition) $v_1 \geq 0$, we see that the above implies that

$$xy(\beta w_1 - \beta' w_2) \leq 0,$$

so to ensure that u^* is not negative, we use the following notation:

$$s^+ = \max\{s, 0\}.$$

Therefore, on this set,

$$u^*(t) = \frac{xy\left(\beta w_1 - \beta' w_2\right)^+}{B}.$$

(iii) Now consider the set $\{t : u(t) = 1\}$. In this case, $v_1 = 0$. Thus,

$$1 = \frac{xy(\beta w_1 - \beta' w_2) - v_2}{B}.$$

This tells us that $0 \leq v_2(t) = xy(\beta w_1 - \beta' w_2) - B$, or, more precisely,

$$\frac{xy(\beta w_1 - \beta' w_2)}{B} \geq 1 = u^*.$$

So, on this set, we must choose

$$u^*(t) = \min\left\{\frac{xy\left(\beta w_1 - \beta' w_2\right)}{B}, 1\right\}.$$

To conclude, we take all three cases together and we find that we can completely characterise $u^*(t)$ as follows:

$$u^*(t) = \min\left\{1, \frac{xy\,(\beta w_1 - \beta' w_2)^+}{B}\right\}.$$

We summarise the above results in the following proposition.

Proposition 3.1 *The optimal control for the optimal control problem (3.1)-(3.4) is completely characterised by*

$$u^*(t) = \min\left\{1, \frac{xy\,(\beta w_1 - \beta' w_2)^+}{B}\right\}.$$

So we can see that the control is described in terms of levels of circulating healthy and infected cells as well as their related adjoint variables.

3.2 Derivation of the Optimality System

The optimality system is an important part of this problem. It describes mathematically how the system behaves under application of the control. Therefore, we may find how the different populations of cells grow or decay when the individual is treated with optimal therapy as characterised in Section 3.

The optimality system is defined as the state system together with the adjoint system and the optimal control u^*. The adjoint system is given by

$$\begin{aligned}
\frac{dw_1}{dt} &= -\frac{\partial H}{\partial x} \\
\frac{dw_2}{dt} &= -\frac{\partial H}{\partial y} \\
\frac{dw_3}{dt} &= -\frac{\partial H}{\partial z}.
\end{aligned}$$

The final component in the optimality system is the set of *transversality conditions*, which in this case reduce to end conditions on the adjoint variables. They are a consequence of the following result, which can also be found in Fleming and Rishel [9].

Given the maximisation problem $\max J[u] = F(x(T)) + \int_0^T f_0(x,u)\,dt$, subject to the state system $dx/dt = f(t,x,u)$ and such that $x(T)$ belongs to some *target set* $g(x(T))$, we have the following transversality conditions on the adjoint variables

$$w_i(T) = \nabla F(x(T)) + \sum_{i=1}^{k} c_i g_i(x(T)). \tag{3.5}$$

The function F is known as the *terminal cost.*

In our problem, there is no terminal cost, so $F(x(T)) = 0$. We also do not have a target set for our state variables — we have a desired end result, of course, but the final state is in fact free, so the summation term is also zero.

Therefore, the transversality conditions for the adjoint variables are

$$w_i(T) = 0, \ i = 1, 2, 3. \tag{3.6}$$

Therefore, taking the state system together with the adjoint system, the optimal control, and the transversality conditions, we have the following optimality system:

$$\begin{aligned}
\frac{dx}{dt} &= \lambda - \delta x - (1-u)\beta xy && (3.7)\\
\frac{dy}{dt} &= (1-u)\beta' xy - ay - \rho yz && (3.8)\\
\frac{dz}{dt} &= cxyz - hz && (3.9)\\
\frac{dw_1}{dt} &= -1 + \delta w_1 + (1-u)y(\beta w_1 - \beta' w_2) - cyzw_3 && (3.10)\\
\frac{dw_2}{dt} &= (1-u)x(\beta w_1 - \beta' w_2) + aw_2 + \rho z w_2 - cxzw_3 && (3.11)\\
\frac{dw_3}{dt} &= -1 + \rho y w_2 - cxyw_3 + hw_3 && (3.12)\\
u^*(t) &= \min\left\{1, \frac{xy(\beta w_1 - \beta' w_2)}{B}^{+}\right\} && (3.13)\\
w_i(T) &= 0, \ i = 1, 2, 3. && (3.14)
\end{aligned}$$

4 Numerical Results

In this section, we discuss the method for numerically solving the optimality system (3.7)–(3.14) and present the results. We note that this is a two-point boundary-value problem, with separated boundary conditions at times $t = 0$ and $t = T$. It is our aim to solve this problem for the value $T = 100$.

This value was chosen to represent the time (in days) at which treatment is stopped.

It turns out that the problem is quite challenging to solve numerically because our uncontrolled optimality system is (very) unstable in the adjoint variables[1]. We use a finite-difference approach to solve the optimality system (3.7)–(3.14). In particular we employed the software package COLDAE [3], which solves boundary-value differential and differential-algebraic equations (DAEs) by collocation at Gaussian points. These methods are equivalent to sophisticated high-order one-step finite difference schemes. Although ultimately we did not use the DAE capability of the solver, the potential for added flexibility in the problem formulation was useful.

A further complication worth noting is that the optimality system (3.7)–(3.14) is a nonlinear problem; hence we need an initial guess from which to begin a Newton iteration. As is well-known, the convergence of Newton's method depends critically on the initial guess being sufficiently close to the desired solution. In the case of the optimality system (3.7)–(3.14), we were not able to simply divine an initial guess that allowed the Newton iteration to converge. Hence we appealed to the method of analytic continuation [12, 2] to solve the problem. This is a standard and powerful technique used to solve nonlinear boundary-value problems. We now describe our methodology.

The idea behind analytic continuation is to imbed the given problem in a family of related problems that depend on a parameter. We illustrate how we did this for the optimality system (3.7)–(3.14). In this case, we use the parameter T. For $T = 100$, the parameterized problem of course reduces to the original problem. However, we note that the problem is easily solvable for another parameter value; namely, $T = 1$. The solution to the problem for $T = 1$ can then be used as an initial guess to the solution to a nearby problem; e.g., for $T = 1 + \Delta T$ for ΔT sufficiently small. With the proper choice of neighbouring problems, this process can be continued until the desired problem is solved. The successive values of the parameter chosen are referred to as a *homotopy path.*

We solved the optimality system (3.7)–(3.14) by making the change of variable $\tau = t/T$ and hence transforming it to the interval $[0, 1]$, leading to the system:

$$\frac{dx}{d\tau} = T(\lambda - \delta x - (1-u)\beta x y)$$

[1]Most notably the variations of the adjoint variable $w_2(t)$ *even in the controlled case* are very large, indicating a high degree of sensitivity of the performance index $J[u]$ to the changes in $y(t)$; see Figure 4.2 below.

$$
\begin{aligned}
\frac{dy}{d\tau} &= T((1-u)\beta' xy - ay - \rho yz) \\
\frac{dz}{d\tau} &= T(cxyz - hz) \\
\frac{dw_1}{d\tau} &= T(-1 + \delta w_1 + (1-u)y(\beta w_1 - \beta' w_2) - cyzw_3) \\
\frac{dw_2}{d\tau} &= T((1-u)x(\beta w_1 - \beta' w_2) + aw_2 + \rho z w_2 - cxzw_3) \\
\frac{dw_3}{d\tau} &= T(-1 + \rho y w_2 - cxyw_3 + hw_3) \\
u^*(\tau) &= \min\left\{1, \frac{xy(\beta w_1 - \beta' w_2)}{B}^{+}\right\} \\
w_i(1) &= 0,\ i = 1, 2, 3.
\end{aligned}
$$

Using T as our continuation parameter, we were able to converge to the solution of the problem for $T = 1$ with a zero initial guess. We then chose $\Delta T = 0.1$ and incremented T in this way until we reached our desired final value of $T = 100$. The solutions are displayed below in Figures 4.1–4.3. We note that this is not necessarily the only or most efficient homotopy path possible; however, it was sufficient to produce an accurate answer in an acceptable amount of time.

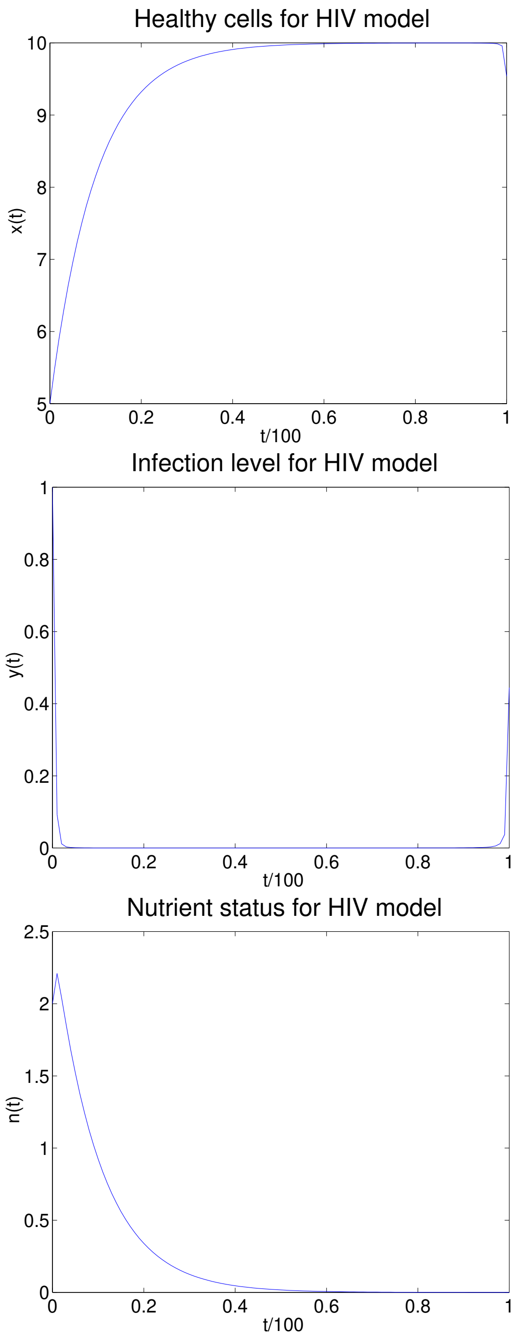

Figure 4.1: The states $x(\tau)$, $y(\tau)$, and $z(\tau)$ of the optimality system.

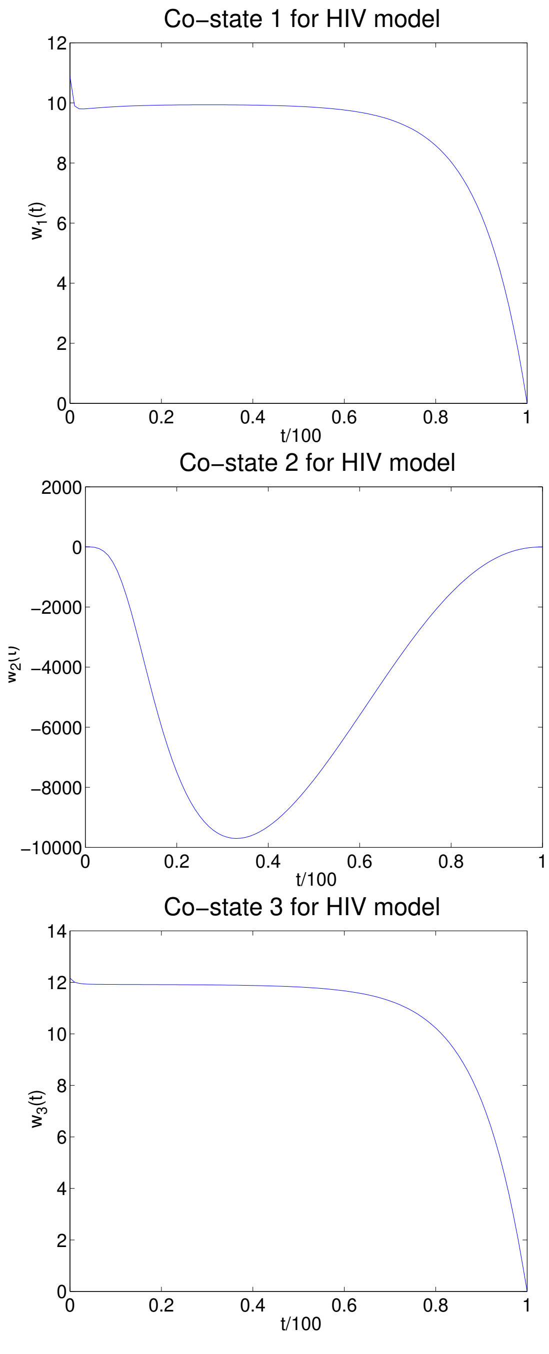

Figure 4.2: The adjoint variables $w_1(\tau)$, $w_2(\tau)$, and $w_3(\tau)$ of the optimality system.

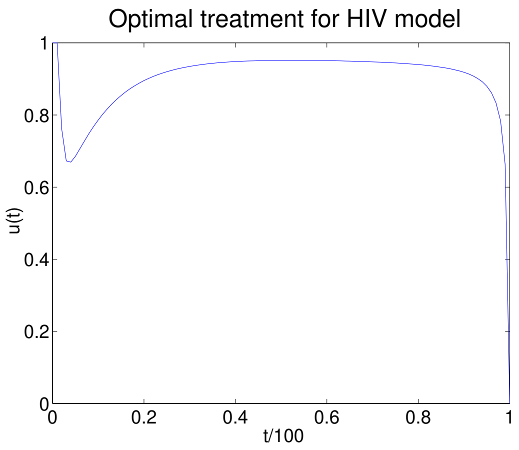

Figure 4.3: The optimal control $u^*(\tau)$.

5 Discussion

We sought to determine optimal treatment strategies that would maximise not only healthy cells but immune response cells as well. We considered first an untreated, three-dimensional ODE model of the interaction between healthy and infected CD4+ T cells and the natural anti-HIV immune response. We found that, depending on the parameter space in which we work, the system either tends to eliminate infection, tends toward an equilibrium with high infection and no specific immune response, or tends toward a balance between infection, healthy cells, and the immune response.

We formulated the optimal control problem to keep both healthy CD4+ T cell and CTL levels high whilst minimising drug cost. The control itself is described in terms of both healthy and infected CD4+ T cells and their corresponding adjoint variables. We note that the existence and continuity of the optimal control are proven in [6].

We then solved the optimality system numerically via analytic continuation in the parameter T, the length of the treatment interval. Referring again to [6], we note that existence of an optimal control is proven using the concept of a *maximising sequence* — that is, a sequence of solutions $(x^{(r)}, u^{(r)})$ to the control problem that converge to the optimal solution (x^*, u^*) as $t^{(r)} \to T$ (see [9] for a much more detailed discussion of existence properties of optimal controls). Referring back to Section 4, we note that continuation in T exactly corresponds to this theoretical concept.

We found that the control starts out at its boundary value of $u = 1$, corresponding to treatment at full strength. After an initial decrease, the optimal treatment grows high again and drops sharply to zero at the final time. The resolution of this sharp drop contributes significantly to the difficulty in obtaining a numerical solution of (3.7)–(3.14).

We can see that the optimal drug treatment protocol has a very desirable effect upon the population of healthy cells. They increase to near their maximal level for almost the entire length of treatment. The sharp drop off at the end is presumably because the cessation of drug enables the infection to rebound and destroy CD4+ T cells.

Infection level decreases to very low levels, but is never eradicated. However, at the end of the treatment schedule, when the drug is no longer given, the infection level begins to rise again. When the infection is low, so too is the specific immune response.

We note that the specific immune response is always maintained at a positive level — it is never eliminated. We also note that an increase in infection is followed by a corresponding increase in the immune response, which then serves to suppress infection (by killing off infected cells). Once the infection is low, the immune response is not needed at such high levels and this is why it too drops off. We note the initial decrease in the control with interest. This occurs at roughly the same time as the immune response is high, indicating that during periods of effective immune responsiveness, less medication is needed to control infection. We suggest that this may indicate that high/low or on/off drug treatment schemes may work well to keep infection under control, provided a sufficient immune response can be maintained. We note that on/off treatment was first modelled mathematically in [15]. As well, implementing treatments that enhance a patient's natural immune response may be beneficial as an alternative to quite such high levels of drug therapy. Boosting host immune response via immunotherapy was investigated in [16]; this is the only work done on this topic to date of which we are aware. We believe that this area deserves further exploration.

Comparing our results with those established in [14] and [8], we find that our control does behave somewhat differently from drugs used to control systems not explicitly modelling immune response. In [14], the control was either monotone decreasing from its maximum value, or peaked slightly just after near the initiation of treatment and then dropped off. However, we observe that our control actually *decreases* soon after initiation of treatment, only to rise again, remain close to constant, and drop rapidly near the end. We believe that this initial drop is directly dependent upon the action of the immune response, which occurs shortly after treatment initiation in response to the high infection level. That is, our optimal treatment is actually reduced for a period of time while the immune response of the host takes over. This indicates that stimulation of the immune response by some means other than continual administration of anti-HIV drugs should be considered seriously in a clinical setting. Treatment strategies such as interruption of drug therapy to allow the immune response to rebuild should also be considered. This can be tested clinically via drug trials, but also mathematically using a periodic control.

References

[1] Arnaout, R. A., Nowak, M. A., and Wodarz, D., HIV-1 dynamics revisited: biphasic decay by cytotoxic T lymphocyte killing? *Proc. Roy. Soc. Lond. B* **265** (2000), 1347–1354.

[2] Ascher, U. M., Mattheij, R. M. M., and Russell R. D., *Numerical Solution of Boundary Value Problems for Ordinary Differential Equations*, Prentice-Hall, Englewood Cliffs, NJ (1988).

[3] Ascher, U. M. and Spiteri, R. J., Collocation software for boundary value differential-algebraic equations. *SIAM J. Sci. Comput.* **15**(4) (1994), 938–952.

[4] Carr, A., Emery, S., Kelleher, A., Law, M., and Cooper, D.A., CD8+ lymphocyte responses to antiviral therapy of HIV infection, *J. AIDS Hum. Retrovir.*, **13** (1996), 320–326.

[5] Cocchi, F., et al., Identification of RANTES, MIP-1 alpha, MIP-1 beta as the major HIV-suppressive factors produced by CD8+ T-cells, *Science*, **270** (1995), 1811–1815.

[6] Culshaw, R. V., *Immune Response Models of HIV Infection and Treatment*, Ph.D. thesis, Dalhousie University, 2002.

[7] Ermentrout, B., *Simulating, Analyzing, and Animating Dynamical Systems: A Guide to XPPAUT for Researchers and Students*, SIAM, Philadelphia (2002).

[8] Fister, K. R., Lenhart, S., and McNally, J. S., Optimizing chemotherapy in an HIV model, *Elect. J. of Differential Equations* **32** (1998), 1–12

[9] Fleming, W. and Rishel, R. *Deterministic and Stochastic Optimal Control*, Springer-Verlag, New York (1975).

[10] Gray, C. M., Lawrence, J., Schapiro, J. M., Altman, J. D., Winters, M. A., Crompton, M., Loi, M., Kundu, S. K., Davis, M. M., and Merigan, T. C., Frequency of class I HLA-restricted anti-HIV CD8+ T cells in individuals receiving highly-active antiretroviral therapy, *J. Immunol.* **162** (1999), 1780–1788.

[11] Haseltine, W. A. and Wong-Staal, F., The molecular biology of the AIDS virus, *Scientific American* (Oct. 1988), 52–62

[12] Keller, H. B., *Numerical Solution of Two-Point Boundary Value Problems*, Regional Conference Series in Applied Mathematics, No. 24, SIAM, Philadelphia (1976).

[13] Kirschner, D. E. and Perelson, A. S., A model for the immune system response to HIV: AZT treatment studies, in *"Mathematical Population Dynamics: Analysis of Heterogeneity, Vol. 1: Theory of Epidemics"*, O. Arino, D. Axelrod, M. Kimmel, and M. Langlais, Eds., Wuerz Pub. Ltd., Winnipeg, Canada, 1995, pp. 295–310.

[14] Kirschner, D. E., Lenhart, S., and Serbin, S., Optimal control of the chemotherapy of HIV, *J. Math. Biol.* **35** (1997), 775–792.

[15] Kirschner, D. E. and Webb, G., A model for treatment strategy in the chemotherapy of AIDS, *Bull. Math. Biol.*, **58** (1996), 167–190.

[16] Kirschner, D. E. and Webb, G., Immunotherapy of HIV-1 Infection, *J. Biol. Systems*, **6** (1998), 71–83.

[17] Kirschner, D. E., Webb, G., and Cloyd, M., Model of HIV-1 disease progression based on virus-induced lymph node homing and homing-induced apoptosis of CD4+ lymphocytes, *J.AIDS*, **24** (2000), 352–362.

[18] Klach, A. M., *Lippincott's Nursing Drug Guide 2001*, Lippincott, Williams and Wilkins, 2001.

[19] Murray, J. D., *Mathematical Biology*, Springer-Verlag, Berlin-Heidelberg (1989).

[20] Musey, L., et al., Cytotoxic T cell responses, viral load and disease progression in early HIV-type 1 infection, *N. Engl. J. Med.*, **337** (1997), 1267–1274.

[21] Nowak, M. A. and May, R. M., Mathematical biology of HIV infections: Antigenic variation and diversity threshold, *Math. Biosci.*, **106** (1991), 1–21.

[22] Ogg, G. S., et al., Decay kinetics of human immunodeficiency virus-specific effector cytotoxic T lymphocytes after combination antiretroviral therapy, *J. Virol.* **73** (1999), 797–800.

[23] Perelson, A. S., Kirschner, D. E., and DeBoer, R., Dynamics of HIV infection of CD4+ T cells, *Math. Biosci.*, **114** (1993), 81–125.

[24] Perelson, A. S. and Nelson, P. W., Mathematical analysis of HIV-1 dynamics in vivo, *SIAM Rev.* **41** (1999), 3–44.

[25] Spouge, J. L., Shrager, R. I., and Dimitrov, D. S., HIV-1 infection kinetics in tissue cultures, *Math. Biosci.*, **138** (1996), 1–22.

[26] Walker, C.M., Moody, D.T., Stites, D.P., and Levy, J.A., CD8+ lymphocytes can control HIV infection in vitro by suppressing virus replication, *Science*, **234** (1986), 1563–1566.

[27] Weber, J. N. and Weiss, R. A., HIV infection: The cellular picture, *Scientific American* (Oct. 1988), 101–109

[28] Wein, L. M., Zenios, S. A., and Nowak, M. A., Dynamic multidrug therapies for HIV: A control theoretic approach, *J. Theor. Biol.* **185** (1997), 15–29.

[29] Wodarz, D., Klenerman, P., and Nowak, M. A., Dynamics of cytotoxic T-lymphocyte exhaustion, *Proc. Roy. Soc. Lond. B*, **265** (1998), 191–203.

[30] Wodarz, D. and Nowak, M., Specific therapies could lead to long-term immunological control of HIV, *Proc. Natl. Acad. Sci.*, **96** (1999), 464–469