Biomedical Signal Processing and Control 8 (2013) 90–97

Contents lists available at SciVerse ScienceDirect

Biomedical Signal Processing and Control

journal homepage: www.elsevier.com/locate/bspc

Optimizing thymic recovery in HIV patients through multidrug therapies

V. Costanza[a,*], P.S. Rivadeneira[a], F.L. Biafore[b], C.E. D'Attellis[b,c]

[a] Grupo de Sistemas No Lineales, INTEC(UNL-CONICET), Güemes 3450, 3000 Santa Fe, Argentina
[b] Departamento de Matemática, Facultad de Ingeniería, Universidad Favaloro, Avda. Belgrano 1723, 1093 Buenos Aires, Argentina
[c] Centro de Matemática Aplicada, Escuela de Ciencia y Tecnología, UNSAM, M. de Irigoyen 3100, 1650 San Martín, Pcia. de Buenos Aires, Argentina

ARTICLE INFO

Article history:
Received 14 October 2011
Received in revised form 12 June 2012
Accepted 12 June 2012
Available online 24 July 2012

Keywords:
HIV
Antiretroviral drugs
Optimal control
Nonlinear dynamics
Dynamic Programming

ABSTRACT

An optimal control approach based on an enlarged nonlinear model for the dynamics of HIV infection and thymic function is composed to simulate and evaluate antiretroviral therapies. In addition to the relevant biological agents, an extra state variable is included, associated with the thymus capacity for healthy cells production. The methodology contemplates eventual deleterious effects of drugs over children's thymus recovery. The intake of 'Reverse Transcriptase Inhibitors' and 'Protease Inhibitors' are modeled as two independent control variables, each affecting a different term in the dynamics, so extending the prevailing pure-HAART-therapy analysis. The objective function designed here is also more inclusive than usual, accounting for the costs of the two drug families involved and for the thymus deterioration, in addition to penalizing eventual virus excess and healthy cells deficits. The search for the best combined therapy is treated as an optimal control problem. A hybrid version of Dynamic Programming for continuous and discrete variables is used to treat the problem numerically. Long time-horizons are explored, aiming to avoid typical peaks in drug prescriptions found at the beginning and at the end of the optimization periods. Results indicate that certain combinations of drugs are more convenient than pure protocols when the value of thymus functioning is relevant, specially for children patients.

© 2012 Elsevier Ltd. All rights reserved.

1. Introduction

The role of the thymus in the regeneration of healthy CD4+T cells during HIV infection has been widely studied in the last decade (see [11,13] and the references therein). It is likely that such a regenerative property is more relevant in children [9,10], although in all cases it is accepted that HIV presence deteriorates thymus functioning, and therefore that the thymopoiesis at a nearly constant rate cannot be assured after the appearance of virus, specially in its CXCR4 version [13]. The negative effect of HIV on thymus functioning is partially reversed by the HAART therapy [3,29], and it is also found that the later the therapy is initiated the slower the recuperation of normal behavior will be achieved [27]. But HAART has shown to produce severe side-effects (see [6] and its references) so it is worth to search for milder medications that still reduce the viral load to protocol target values.

Even in the absence of virus, the thymus is one of the first organs to undergo significant age-related degeneration, termed thymic involution. Thymic involution results in a dramatic drop in the production of new T-cells, and is a significant contributing factor in immune senescence. Despite the importance of this subject for human health, the molecular dynamics that operate in the postnatal thymus and mediate thymic homeostasis and involution are largely unknown [17]. Detailed mathematical models have been developed at cellular level [15,29], aiming to reflect all relevant interactions affecting the intervening agents, to the cost of including numerous state variables, most of them difficult to measure (or to observe) continuously. According to applicable findings of these studies, the dynamical models used by the authors in previous publications [5,6] were modified to include variations on healthy cells' regeneration rate and thymus function decay. For instance, it has turned clear that the parameter associated with the thymopoiesis in simple models should become a state variable in itself, reflecting thymic function dynamics. The new variable declines in the presence of HIV infection [25], and this affects the immune response negatively, most in the form of a self-feeding loop. It is known that recovery of thymic function could occur in HIV-1-infected patients on HAART. However, specific mechanisms that contribute to this recovery are still under study [29,18].

Also it is now widely accepted that dynamic models should account for differences between children and adults evolution of the infection and thymic function, and also in their respective responses to medication [9,22,25,29]. In this last respect, this paper emphasizes the importance of evaluating alternative drug families for medication, since their affecting different terms or parameters of the model will provide alternatives to standard protocols. The evaluation of these alternatives will contemplate viral load suppression, thymic function recovery, and medication expenses, giving rise to a combined cost functional with competing individual

* Corresponding author.

1746-8094/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.bspc.2012.06.002

objectives. The relative importance of these individual objectives into the total cost, and the addition or suppression of other objectives, has been discussed in international forums [28].

The model for the dynamics together with the cost functional designed for therapies evaluation pose an optimal control problem. This is a richer viewpoint in comparison with those of previous papers: innovations are included in the model and in the evaluation functional, both containing more variables and individual cost objectives. A special feature in this last respect is the consideration of long optimization time-horizons, in the order of several years, in an attempt to avoid big variations in the drug doses typically prescribed at the beginning and at the end of shorter time-horizons (of around 180 days) [6]. Another improvement in the scope of the optimization problem is the possibility of including two families of drugs into the medication, eventually in variable proportions. The doses admitted here will contain typical members of the Reverse Transcriptase Inhibitors (RTIs) and the Protease Inhibitors (PIs) groups.

The results arising from the numerical treatment of this optimization setup support two main conclusions: (i) a combination of RTI and PI drugs can work better than monodrug protocols when all individual cost objectives are given proper weight, and (ii) less aggressive drugs can partially replace HAART during small infected children medication, especially when avoiding thymus dysfunctionality is assigned a significant value.

From the numerical simulations there arise other observations, whose precise meaning deserve further investigation. For instance, it seems that extending the time-horizon in optimization analyses may conduct to reliable therapies consisting of a constant dose for each intervening drug. This outcome is positive in avoiding big changes in therapy typically appearing at the beginning and at the end of optimization periods of the order of 180 days. But the optimal required doses become higher the longer the horizon is posed, which is not neatly positive in several respects, not the least the increasing money-cost of the medication. A first interpretation would propose that higher doses are required to give more assurance against predictable rebounds of the viral load. However, it is also true that side-effects of higher doses will increase, and that should also be included as a partial cost in the optimization analysis, as done in previous work [6]. Unwanted effects of medication were not considered in this paper due to lack of data concerning both combined therapies and long intervals of application.

The paper is organized as follows: after the Introduction there is Section 2 that explains the model for the dynamics of the control system under study, and in Section 3 the cost functional is designed and discussed. Section 4 gives the details of the calculation techniques used to find optimal therapies, and the numerical results obtained are analyzed in Section 5. Section 6 concludes the paper.

2. Dynamics of infection and thymus function

The behavior of infected patients will be modeled through a continuous-time control system described by the following set of coupled ordinary differential equations:

$$\begin{aligned}\dot{x} &= \lambda - \delta x - \beta(u_1)xz\\ \dot{y} &= \beta(u_1)xz - \mu y\\ \dot{z} &= \kappa(u_2)y - \gamma z\\ \dot{\lambda} &= \eta(u_1,u_2)(\lambda_m - \lambda) - \sigma z\end{aligned} \tag{1}$$

where the states are: healthy CD4+T cells (x), infected CD4 T cells (y), free virus copies or virions (z), and thymopoiesis rate (λ).

The control action over the states is exerted by the medication, i.e. by the intake of drugs. Basically, the antiretroviral drugs can be grouped into the following three categories [14]:

(i) *Inhibitors of the reverse transcriptase enzyme (RTIs)*: if RT is inhibited, HIV can enter a cell but will not successfully infect it; a DNA copy of the viral genome will not be made and the cell will not make viral proteins.
(ii) *Protease Inhibitors (PIs)*: if HIV protease is inhibited, cleavage of the viral polyprotein will not occur, and viral particles that lack functional enzymes will be made. The net effect of blocking HIV protease is that noninfectious viral particles are produced.
(iii) *Fusion Inhibitors (FIs)*: these work by inhibiting the binding of HIV to healthy CD4+T-cells (used in patients with multi-drug HIV resistance, not studied in this paper).

The control variables will be restricted in this paper to u_1: dose of Zidovudine AZT (RTI), and u_2: dose of Ritonavir (PI), both in [g/day].

The first three equations (where λ, $\beta(u_1)$ and $\kappa(u_2)$ are kept time-constant, as if they were parameters) constitute a classical simplified model [5,19,20], thus further comments on the meaning of their terms seem unnecessary at this time. However, it is appropriate to explain that the pharmacokinetics (PK) of each drug enters implicitly in the determination of parameters β and κ, since they were experimentally estimated in terms of the doses and not in terms of the efficiencies of the drugs. The efficacies of the drugs are incorporated via dose-dependent parameters that are time-constant. This efficacy is dependent on the PK of the agent and its pharmacodynamics (PD; relation between efficacy and, typically, plasma drug concentration). These are time-dependent processes (as evidenced by the need to re-dose daily). Thus the control parameters in the proposed model are piecewise constant approximations. Since the relations PK/PD of the drugs are not included explicitly, then, when interpreting the resulting optimal strategies, it should be understood that the form of the administration of the drugs must remain the same along the whole time-span under consideration.

The notations $\beta(u_1)$, $\kappa(u_2)$ denote that each one of these parameters is mostly affected by only one of the drug types [1,23,26]. The last equation for λ dynamics is new. It governs the natural thymus homeostasis. No data have been reported describing the ability of HAART agents to cross the blood–thymus barrier and penetrate into the thymus. Therefore the explicit dependence of the thymus variable λ on any of the control variables u_i adopted in the last of Eq. (1) should be designed so as to give a reasonable approximation to the observed qualitative behavior reported up to now. The form adopted here includes

$$\eta(u_1,u_2) \triangleq [1 - \chi u_1(1-u_2)], \tag{2}$$

where χ is a nonnegative parameter. Supposing there is no infection for the moment ($z=0$, and consequently with no medication, i.e. $u_1 = u_2 = 0$, $\eta(u_1,u_2) = 1$), and if the term $(\lambda_m - \lambda)$ is positive (the thymus functioning is below its standard value λ_m), then a self-induced recuperation will be intended by the human organism, i.e. the rate $\dot{\lambda} = d\lambda/dt$ will increase. But the HIV infection, reflected in the present value of the state $z(t) > 0$, will oppose this recovering through the term $-\sigma z$, σ a parameter identified from: (i) the available data for thymus dynamics showing, for instance, that the thymocytes decay to about one half of their original number [15,29] during the first month after infection, together with (ii) the conventionally accepted proportional relation between thymocytes and CD4+T cells. The behavior of $\eta(u_1,u_2)$ in the range of interest of the control doses can be seen in Fig. 1. Setting

$$\chi_{\text{adults}} = 0 \tag{3}$$

implies $\eta_{\text{adults}}(u_1,u_2) \equiv 1$ for any admissible combination of RTI and PI drugs, coincident with the observation that for adults the thymus recovery seems to be independent of drug efficiency inside the organ [15].

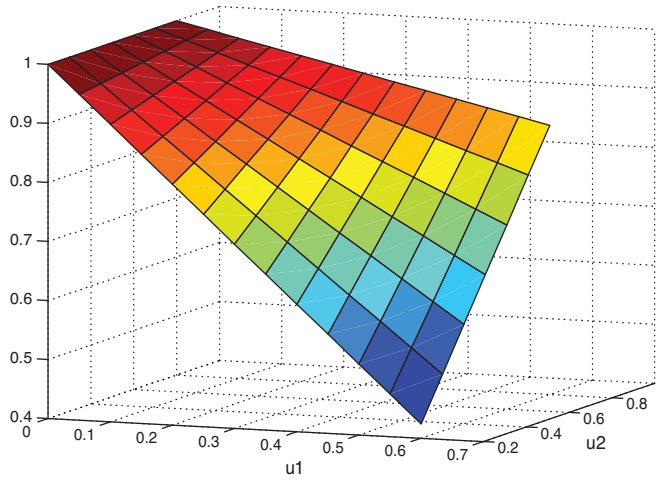

Fig. 1. Function $\eta(u_1, u_2)=[1-\chi u_1(1-u_2)]$ acting as a coefficient for the recuperation term $(\lambda_m-\lambda)$ in the dynamic equation for the thymic function. Here $\chi=\chi_{\text{children}}=1.2$.

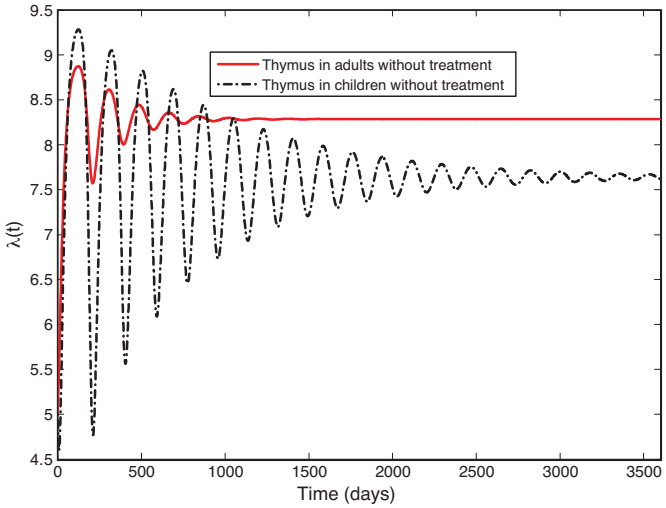

Fig. 2. Asymptotic behavior of thymus function for infected children and adults without medication. Initial conditions as in Eq. (19).

Children have a better production of naive T-cells and thymocytes, which suggests, in terms of parameters,

$$\lambda_{m_{\text{children}}} > \lambda_{m_{\text{adults}}}. \tag{4}$$

But also the efficiency of drugs inside a child's thymus is different than in adults. In previous work we have only considered the action of RTIs. But if only RTI drugs are being dispensed to children patients, it is found that their efficiency inside the thymus is low [29], i.e. such a medication is not enough to produce the same thymus recovery effect than in adults (this is the explanation for $\eta_{\text{children}}(u_1, u_2)<1$ for $u_1>0$, Eq. (2) and Fig. 1). The same authors indicate that the combination with PI drugs enhances the performance of RTI medication, and that is the reason for the term $(1-u_2)$ multiplying u_1 in Eq. (2). The combination $\chi u_1(1-u_2)$ works as a typical interactive term between the two drugs: $(1-u_2)<1$ reduces the detrimental action due to u_1 alone. In order to keep $0<\eta_{\text{children}}(u_1, u_2)\leq 1$ in the range of interest, the fitted value for the relevant parameter was

$$\chi_{\text{children}} = 1.2. \tag{5}$$

The asymptotic value for the thymus function in an infected non medicated patient $(u_1=u_2\equiv 0)$ will be

$$\overline{\lambda} \triangleq \lambda_m - \frac{\sigma}{\overline{z}}, \tag{6}$$

where the symbol $\overline{z}$ represents the 'ill' equilibrium value of the viral load. The point $(\overline{x}, \overline{y}, \overline{z}) = ((\mu\gamma/\beta\kappa), (\lambda-\delta\overline{x}/\mu), (\kappa\overline{y}/\gamma)) \in \mathbb{R}^3$ describes the values approached by the first three state variables when the patient is critically infected and non medicated. At such values the dynamics would become asymptotically stable, tending to a terminal situation with $\dot{x}=\dot{y}=\dot{z}=0$. By assuming that $\overline{z}$ is approximately the same for children and for adults, the fact that the thymus function of infected children is more susceptible of becoming dysfunctional [4,10] will be interpreted as the need for

$$\overline{\lambda}_{\text{children}} < \overline{\lambda}_{\text{adults}},$$

which in view of Eqs. ((4)–(6)) would only be possible provided that also

$$\sigma_{\text{children}} > \sigma_{\text{adults}}.$$

The remaining nominal values for the parameters in Eq. (1) are then recovered from previous versions of the model and fitted to meet the differences between children and adults discussed in the text and the real-life data reported in the citations above:

$$\begin{aligned} &\lambda_{m_{\text{adults}}} = 9\,\text{cells mm}^{-3}\,\text{day}^{-1},\\ &\lambda_{m_{\text{children}}} = 9.5\,\text{cells mm}^{-3}\,\text{day}^{-1}\\ &\sigma_{\text{adults}} = 2\times 10^{-5}\,\text{cells copies}^{-1}\,\text{mm}^{-3}\,\text{day}^{-2},\\ &\sigma_{\text{children}} = 6\times 10^{-5}\,\text{cells copies}^{-1}\,\text{mm}^{-3}\,\text{day}^{-2}\\ &\delta = 0.009\,\text{day}^{-1},\quad \mu = 0.3\,\text{day}^{-1},\quad \gamma = 0.6\,\text{day}^{-1}. \end{aligned} \tag{7}$$

The proposed expression for $\beta(u_1)$

$$\beta(u_1) \approx \beta_0 - \alpha_1 u_1 - \alpha_2 u_1^2 \tag{8}$$

was substantiated from qualitative and quantitative data in [1,23,26]. The parameters were estimated through standard least-squares regression techniques, by assuming that for a short period (of the order of one week, see [1,23]) after therapy has begun, x remains approximately constant (say $x(t)\approx\tilde{x}\triangleq x(0)$). Then, the equations for $\dot{y}$ and $\dot{z}$ in the model (1) become linear and their solution imply:

$$z(t) \approx (k_1 e^{a_1 t} + k_2 e^{a_2 t}) z(0), \tag{9}$$

with coefficients depending on $x(0)$, $y(0)$, β_0, α_1, α_2. Data in the quoted literature, actually corresponding to treatments of recently discovered infections, produce the following estimates

$$\begin{aligned} &\beta_0 = 4\times 10^{-6}\,\text{ml copies}^{-1}\,\text{day}^{-1},\quad \alpha_1 = 0.88\times 10^{-6},\\ &\alpha_2 = 0.3\times 10^{-6}. \end{aligned} \tag{10}$$

Analogously, PI drugs are known to affect preferentially the pseudo parameter κ in a form which has also been adapted to empirical data [2,21,24] to obtain:

$$\kappa(u_2) \approx \kappa_0 - \alpha_3 u_2 - \alpha_4 u_2^2, \tag{11}$$

$$\begin{aligned} &\kappa_0 = 8\times 10^{-2}\,\text{copies cells}^{-1}\,\text{day}^{-1},\quad \alpha_3 = 40.943,\\ &\alpha_4 = 1.589. \end{aligned} \tag{12}$$

Some aspects of the behavior of thymus function variable λ for children and adults, in concordance with the comments above, can be visualized in Figs. 2–4. Figs. 2 and 3 also show that from a lower value $\lambda(0)$ (revealing some deterioration of the thymus

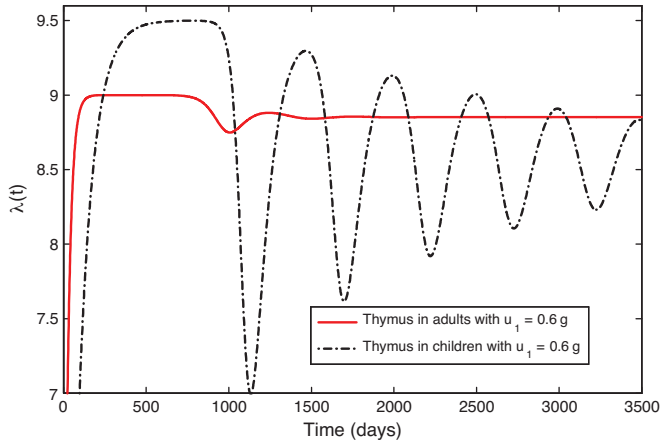

Fig. 3. Asymptotic behavior of thymus function for infected children and adults with several constant PI therapies in addition to a constant RTI doses $u_1 \equiv 0.6$. Initial conditions as in Eq. (19).

function due to infection), its recuperation in children is obtained after stronger oscillations than in adults. This is most probably due to the children's greater thymus volume, which allows for a quick response, giving rise to more abundant thymocytes than adults do. Children thymus generates comparatively more CD4+T cells, able to be infected and to reproduce more virus copies, which in turn deteriorate thymus more significantly, and so starting the recuperation cycle again. Perhaps these strong oscillations explain early thymus dysfunctionality in children when not properly medicated. Fig. 4 shows that combined RTI plus PI therapies in children could restore thymus function to a higher value and more smoothly than with pure RTI drugs. The question is which the best proportion for such combinations is when taking a multi-objective viewpoint. This will be attempted in the rest of the paper.

3. Total cost associated to a therapeutic strategy

A typical objective functional, representing the "total" cost to minimize among all acceptable therapies, may be designed as follows

$$J(u) = Q(t_0, T, x_0, y_0, z_0, \lambda_0, u) + K(x, y, z, \lambda)_T = \int_{t_0}^{T} [a_1 z(t) + a_2 f(u_1(t), u_2(t)) + a_3(\lambda(t) - \lambda_m)^2]dt + a_4 z^2(T). \quad (13)$$

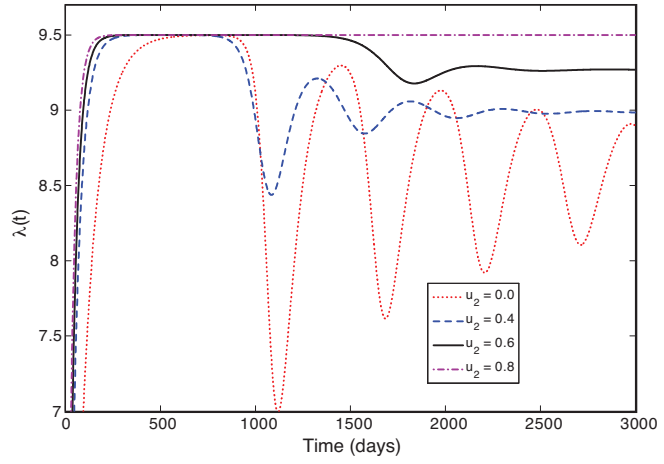

Fig. 4. Asymptotic behavior of thymus function for infected children and adults with several constant PI therapies in addition to a constant RTI doses $u_1 \equiv 0.6$. Initial conditions as in Eq. (19).

This is a typical form for the objective function, consisting of a "trajectory cost" $Q(t_0, T, x_0, y_0, z_0, \lambda_0, u)$ (usually expressed as the integral of the Lagrangian function $L(x, y, z, \lambda, u)$, which models the cost differentials occurring during treatment), and a "final penalty" $K(x, y, z, \lambda)_T$ associated with the departure of the final states from the desired (target) after therapy, which has been maintained quadratic (as in previous work of the authors, for eventual comparison of results), i.e.

$$K(x, y, z, \lambda)_T = K(x(T), y(T), z(T), \lambda(T)) = a_4 z^2(T). \quad (14)$$

The term $a_2 f(u_1, u_2)$ associated with the cost of the drugs used for therapy should be zero if medication is absent ($u_1 = u_1 = 0$). If there were no other (positive) terms in the trajectory cost Q (see for instance [16] and some of the references therein), it may happen that the optimal policy results 'no medication'. Adding state-trajectory costs (a_1 and a_3 not simultaneously null) will prevent from total drug interruption at any intermediate time within optimal therapies. Such interruptions are recognized to be inappropriate [7]. The choice made here takes the form

$$L(x(t), y(t), z(t), \lambda(t), u(t)) = a_1 z(t) + a_2 f(u_1(t), u_2(t)) + a_3(\lambda(t) - \lambda_m)^2 \quad (15)$$

The Lagrangian in this case includes the term $a_1 z(t)$ penalizing the existing of a nontrivial virus load. The values of a_1 and a_4, although adjustable as all evaluation parameters, have been kept high enough as to guarantee that the viral load descend below the 50 copies/ml in the first 180 days for every admissible optimal medication strategy.

The term $f(u_1, u_2)$ represents the effective financial cost of the doses prescribed by the chosen therapy. The simplest form for the function f is just the sum of the individual costs of each drug, i.e.

$$f(u_1, u_2) = m_1 u_1 + m_2 u_2, \quad (16)$$

with m_i denoting the current price of drug u_i in dollars per gram (in Argentina, March 2011). It should be made clear that the value of $f(u_1, u_2)$ represents just the expense in money for buying the drugs. Other indirect costs associated with the medication, like those commonly considered as 'side-effects' (see [6]), or the propensity to encourage drug resistance, should reasonably be taken into account in a more complete formulation of the dynamics and cost objective $J(u)$. The setup here is just an approximation to the problem.

The coefficient a_3 weights the importance given to thymus deterioration in the whole. The square in $(\lambda - \lambda_m)^2$ reflects the current presumption [25] that a too active thymus is also detrimental in curing the infection, above all in children.

The initial time t_0 is taken as zero since both the dynamics and the cost functional are autonomous. Then the final time T becomes the time-horizon for the problem. A value of $T \gg 180$ days was used in all calculations, in an attempt to find stable therapies and to avoid typical high doses at the beginning and end of short successive optimization periods (in the order of 180 days).

The values of the coefficients adopted after previous considerations were:

$$a_1 = 0.01, \quad a_2 = 1, \quad a_{3_{\text{children}}} = 100, \quad a_{3_{\text{adults}}} = 50, \quad a_4 = 0.001, \quad m_1 = 15.4\,\text{dollars}\,g^{-1}, \quad m_2 = 21.6\,\text{dollars}\,g^{-1}, \quad T = 420 \to 1200\,\text{days}. \quad (17)$$

4. Numerical treatment

4.1. Discretizing the state and control spaces

The structure of a continuous-time optimal control problem is not completely functional to the handling of HIV medication in

real life. Notwithstanding the patient's health undergoes a continuous deterioration, possibly following a model like Eq. (1), the assessment of the situation can only be made through periodic (discontinuous) blood analyses and doctor inspections, and the present administration of drugs can only be made through discrete amounts and changed every some time-period according to prescriptions. This means that, even when the system evolves in continuous time and the cost objective can theoretically be posed in the same context, the discrete nature of:

(i) measurements' availability or observability of the state variables (x, y, z, λ),
(ii) the discrete nature of admissible control values,
(iii) the existence of restrictions in the admissible control values, which hinders the smoothness of control trajectories, and
(iv) the usual delay between physician interventions and control decisions, it is reasonable to consider a mixed continuous/discrete approach to the numerical treatment.

In what follows, the values of the states and control variables will be then discretized according to the following scheme

$$\begin{aligned}
&\mathcal{X} \triangleq \{x_L, x_L + \Delta x, x_L + 2\Delta x, \cdots, x_U\}\\
&\mathcal{Y} \triangleq \{y_L, y_L + \Delta y, y_L + 2\Delta y, \cdots, y_U\}\\
&\mathcal{Z} \triangleq \{z_L, z_L + \Delta z, z_L + 2\Delta z, \cdots, z_U\}\\
&\mathcal{L} \triangleq \left\{\lambda_L, \lambda_L + \Delta\lambda, \lambda_L + 2\Delta\lambda, \cdots, \lambda_U\right\}\\
&\mathcal{U}_1 \triangleq \{u_{1_L}, u_{1_L} + \Delta u_1, u_{1_L} + 2\Delta u_1, \cdots, u_{1_U}\}\\
&\mathcal{U}_2 \triangleq \{u_{2_L}, u_{2_L} + \Delta u_2, u_{2_L} + 2\Delta u_2, \cdots, u_{2_U}\}\\
&\mathcal{U} \triangleq \mathcal{U}_1 \times \mathcal{U}_2
\end{aligned} \tag{18}$$

where the values of the lower (L) and upper (U) bounds and the grid size (Δ) for each variable should depend on real constrains on the appreciation of measurement devices, the possibilities of dose subdivision, calculation capabilities, and expectations. However it is important to note that the adopted values for y_L, z_L and p_L should be strictly greater than zero to avoid an unrealistic stagnation of the variables y, z (see Eq. (1)). This will be discussed further, together with the numerical trials and results.

4.2. Hybrid Dynamic Programming

From now on the optimal treatment of a patient with a recently discovered infection, for instance with initial conditions near

$$\begin{aligned}
&x_0 = 850\,\text{cells mm}^{-3}, \quad y_0 = 41\,\text{cells mm}^{-3},\\
&z_0 = 3760\,\text{copies ml}^{-1}, \quad \lambda_0 = 5\,\text{cells mm}^{-3}\,\text{day}^{-1},
\end{aligned} \tag{19}$$

is illustrated. Dynamic Programming was implemented for a range of states around these initial conditions, covering expected behaviors of the patient under different medication strategies. The adopted discretization or spacing (Δ) of the variables and their corresponding lower (L) and upper (U) bounds are listed bellow:

$$\begin{aligned}
&\Delta x = 50, \quad \Delta y = 10, \quad \Delta z = 500, \quad \Delta\lambda = 0.1,\\
&\Delta u_1 = \Delta u_2 = 0.2,
\end{aligned} \tag{20}$$

$$x_L = 500, \quad y_L = 1, \quad z_L = 10, \quad \lambda_L = 4, \quad u_{1_L} = u_{2_L} = 0, \tag{21}$$

$$\begin{aligned}
&x_U = 1000, \quad y_U = 51, \quad z_U = 5010, \quad \lambda_U = 10,\\
&u_{1_U} = u_{2_U} = 0.8
\end{aligned} \tag{22}$$

It must be remarked that the lower thresholds y_L and z_L are given strictly positive values. This is to avoid that, after an eventual rounding of their values, the discretized trajectories reach any point with $y = z = 0$, since in that case the optimal strategy would continue with $u \equiv 0$ until the end, which is certainly erroneous (and dangerous). Indeed, the real system never reaches $y = z = \lambda = 0$ from an initial condition different from the unstable equilibrium, which means that any remaining infection ($y > 0, z > 0$) will grow if $u = (0, 0)$ and this growth should eventually be controlled with some $u_i > 0$. The Δ spacings used here to make each variable discrete might result too demanding given the non-uniformity of the data acquisition met in real-life scenarios. In applications, the physician will surely resort to interpolations, and to the help of some professional in informatics to make compatible the accuracy of results with the samples analyses and computational capabilities. The upper limits on the doses of each drug were adopted on safety grounds. Their sum, namely 1.6 g, was the maximum value for each individual drug admitted by any of the essays quoted before, whose data were used to estimate the parameters of the dynamics.

The cost takes now a slightly different form

$$\begin{aligned}
\mathcal{J}(u) \triangleq \sum_{k=0}^{T/h} \int_{t_k}^{t_{k+1}} &[a_1 z(t) + m_1 u_1(t) + m_2 u_2(t)\\
&+ a_3(\lambda(t) - \lambda_m)^2]dt + a_4\left(z\left(\frac{T}{h}\right) - \hat{z}\right)^2,\\
t_k \triangleq t_0 + hk,&
\end{aligned} \tag{23}$$

where $x(t)$, $y(t)$, $z(t)$, $\lambda(t)$ must be understood, in each interval $[t_k t_{k+1})$, as the rounded result of the state-transition function $\phi(t, t_0, x, y, z, \lambda, u(\cdot))$ associated with the continuous-time model (1), namely

$$(x(t), y(t), z(t), \lambda(t))' = \phi(t, t_k, x_k, y_k, z_k, \lambda_k, \tilde{u}_k), \tag{24}$$

$$(x_{k+1}, y_{k+1}, z_{k+1}, \lambda_{k+1})' = \text{round}(\phi(t_{k+1}, t_k, x_k, y_k, z_k, \lambda_k, \tilde{u}_k)), \tag{25}$$

the symbol $\tilde{u}_k$ being interpreted as the piecewise-constant control trajectory

$$u(t) = (u_1(t), u_1(t)) \triangleq \tilde{u}_k(t) \equiv u_k = (u_{1_k}, u_{2_k}) \in \mathcal{U} \ \forall t \in [t_k, t_{k+1}), \tag{26}$$

and where 'round' acts over the values $(x(t_{k+1}), y(t_{k+1}), z(t_{k+1}), \lambda(t_{k+1}))' = \phi(t_{k+1}, t_k, x_k, y_k, z_k, \lambda_k, \tilde{u}_k)$ in a 'safe' way, precisely

$$\begin{aligned}
x_{k+1} &\triangleq \text{closest smaller value next to } x(t_{k+1}) \text{ in } \mathcal{X}\\
y_{k+1} &\triangleq \text{closest bigger value next to } y(t_{k+1}) \text{ in } \mathcal{Y}\\
z_{k+1} &\triangleq \text{closest bigger value next to } z(t_{k+1}) \text{ in } \mathcal{Z}\\
\lambda_{k+1} &\triangleq \text{closest value next to } \lambda(t_{k+1}) \text{ in } \mathcal{L}
\end{aligned} \tag{27}$$

This means that, even when the time, state, and control spaces are discretized, the values of the states at a time t_{k+1} are calculated by integrating the continuous-time dynamics of the control system (i.e. through the state-transition function ϕ of the system) by using the known states at time t_k as initial conditions, and then rounding the states resulting from integration. That explains the adjective 'Hybrid' assigned to the Dynamic Programming technique used as numerical approach to the problem in this paper. The adopted value of $h = 30$ days takes into consideration the observed 'peak time' (of approximately 20 days) occurring in the state variables and sensitivities (see [5]). The possibility of a hidden acute infection period is discarded this way, since at least the results of a blood analysis reflecting the situation will come at some intermediate point.

The numerical scheme proceeds 'backwards in time'. Calculations start by assuming that the cost of an instantaneous process at

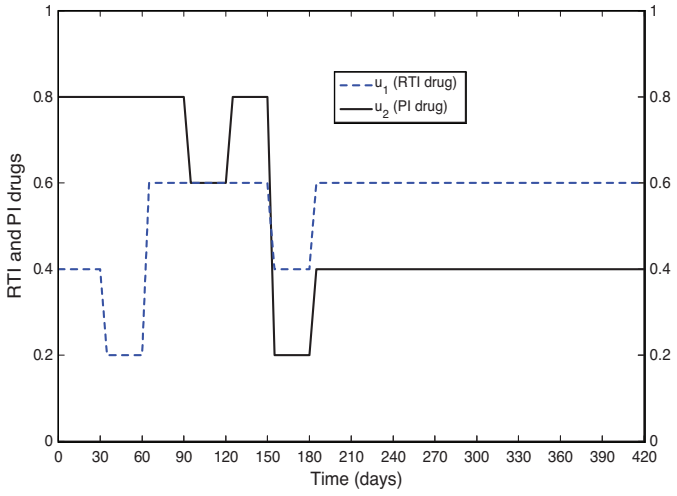

Fig. 5. Optimal strategy (therapy) for adults, among all admissible combinations of RTI and PI drugs. Optimization horizon $T = 420$ days.

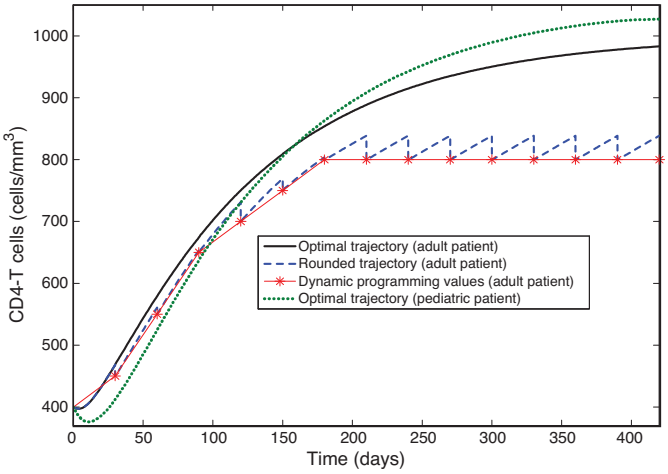

Fig. 6. x-Trajectories for children and adults resulting from their corresponding optimal medications. x-Values used in Dynamic Programming calculations are also shown for adults.

time T coincides with the final penalization K_T of Eq. (13), for each admissible final state $(x, y, z, \lambda)_T$. Then, the procedure determines the optimal strategies for a process of duration Δt starting at $T - \Delta t$ (i.e. for the last stage); and this is done for each 'initial' state $(x, y, z, \lambda)_{T-\Delta t}$. The resulting optimal controls are stored, identified with each $(x, y, z, \lambda)_{T-\Delta t}$. The procedure is repeated for each 'previous' stage until arriving to the initial time, and the optimal controls are stored for every state at every stage. In this way, at any time and for each admissible state of the patient, the optimal control strategy from there to the end can be recovered from the stored data, which makes the methodology robust with respect to inaccuracies between model and patient. At each stage the control can be 'corrected' in terms of the real state of the patient. This is the usual meaning of 'feedback control'. For further details see for instance [12].

5. Numerical results

The first calculations made were for adult patients, from initial conditions as in Eq. (19), resulted in an optimal control strategy as shown in Fig. 5. The optimal doses of RTI stabilized at 0.6 g after the critical period of 180 days, while PI did the same at 0.4 g. This was a good result in some aspects: first, both drugs were present in a nontrivial combination at all stages; and second, both of them stayed at a constant level from more than the last half of the period under consideration, which suggests that the medication $u^*_{\text{adults}}(t) \equiv (0.6, 0.4)$ could become safe for approximately $t > 180$ days, without need to repeat calculations for successive optimization periods. But the preference for RTIs over PIs was somewhat expected since u_1 acts by diminishing directly both the viral load and the thymus function dysfunctionality (and consequently lowering the costs weighted by a_1, a_3, and a_4), while increasing the drug cost weighted by a_2 at a lower rate than u_2 does (PIs are more expensive). This result is reinforced when the cost of thymus illness is not given preponderance in the composition of the cost, as was decided to be the case for adults.

Things change for children, as can be seen in Fig. 9, due to the different set of parameters used in calculations, as indicated in Eqs. ((7) and (17)). The optimal medication for children stabilizes at $u^*_{\text{children}}(t) \equiv (0.4, 0.6)$ for $t > 180$ days, reversing the roles of RTI and PI in the optimal therapy for adults. This shows that an alternative combination of drugs, less aggressive to the thymus functioning in children, can result optimal when the weights of the partial costs are tuned as to reflect a comprehensive evaluation of the situation.

The state trajectories for $x(t)$, $z(t)$ and $\lambda(t)$, corresponding to the optimal control strategies, both for adults and children, are depicted in Figs. 6–8. They show a very acceptable performance, judging from common-sense expectations for a good medication, i.e. it should quickly abate viral load while at the same time recuperates healthy cells and thymus function. For illustration purposes, and only in the adults case, curves for the values of the states x, z, λ that are generated in the intermediate calculations of the optimal control strategy by the Dynamic Programming scheme are added in Figs. 6–8. These values, obtained after applying the rounding conventions announced in Eq. (27), are shown to be conservative, since the continuous-line curves (real behavior of the patient due to optimal medication) tend to the desired values more quickly and accurately than the remaining ones (Dynamic Programming rounded approximations).

Partial costs due to optimal medications in children and adult patients are shown in Table 1, for $a_3 = 30$ and $a_3 = 100$. In both cases the sum of the asymptotic optimal doses is the same for children and adults, although for $a_3 = 100$ such a sum is higher (1 g instead of 0.8 g for $a_3 = 30$), indicating that more concern for thymus function will require more drug, all other objectives being equally weighted.

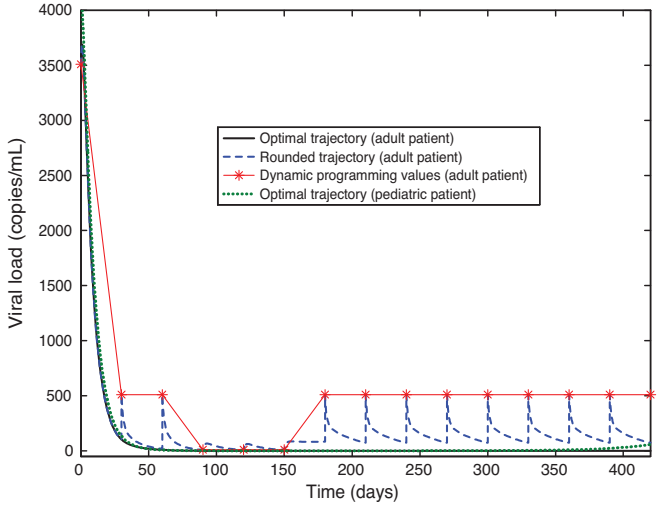

Fig. 7. z-Trajectory for children and adults resulting from their corresponding optimal medications. z-Values used in Dynamic Programming calculations are also shown for adults.

Table 1
Cost values for children and adult treatments for $a_3 = 30, 100$; $T = 420$ days. The asymptotic optimal value of the drug u_i appears as a superscript of the partial cost associated with such u_i.

	a_3	$\int a_1 z(t)dt$	$\int a_3(\lambda - \overline{\lambda}_m)^2 dt$	$\int m_1 u_1 dt$	$\int m_2 u_2 dt$	Cost of medication
Children	30	476.30	16622	$2495^{(0.4)}$	$3889^{(0.4)}$	6383
Adult		416.81	8253	$3142^{(0.6)}$	$1814^{(0.2)}$	4956
Children	100	393.72	53546	$2495^{(0.4)}$	$5184^{(0.6)}$	7679
Adult		309.90	25493	$3788^{(0.6)}$	$3888^{(0.4)}$	7676

The cost of medication for children when $a_3 = 30$ is 1.28 times the cost for adults, but this is not too expensive to pay for as much as 16, 622/8253 = 2.0141 times the all along deviation of the thymus function. For $a_3 = 100$ this argument is even stronger but the cost of medication rises both for children and adults to reach around 7680 dollars in 420 days. As $a_{3_{\text{children}}}$ increases the recommended PI doses increases and surpasses the one for RTIs. This could be rephrased by saying that emphasizing thymus concern in children patients would amount in increasing proportion of PI drugs replacing RTIs. This policy would be reinforced if the hypothesis of PI retarding the switching time from R5 to X4 virus [8] is confirmed by further empirical data.

For longer optimization horizons things become more complex. Figs. 10 and 11 depict respectively the z- and λ-trajectories corresponding to the optimal medications calculated for children and a very long T (=1200 days), using different values for a_3. The asymptotic recommended drug doses are higher than for shorter optimization periods, explained by the need to avoid significant rebounds in the viral load. This is an issue that deserves further study, above all considering that model parameters are no longer trusted to remain constant after such long periods of treatment, and the same can be expected of the partial cost weights. Perhaps

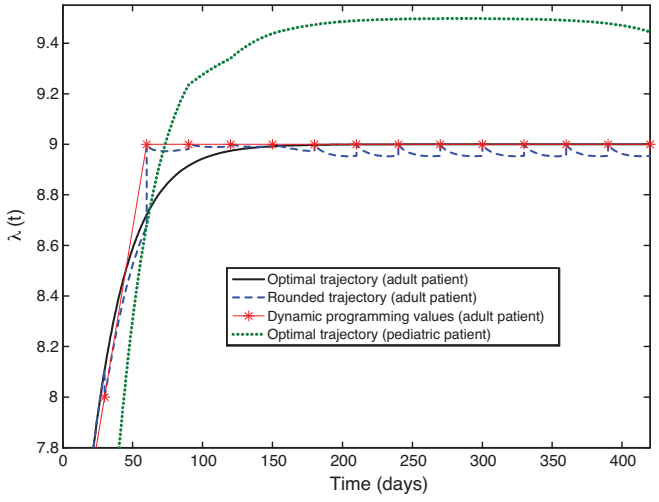

Fig. 8. λ-Trajectories for children and adults resulting from their corresponding optimal medications. λ-Values used in Dynamic Programming calculations are also shown for adults.

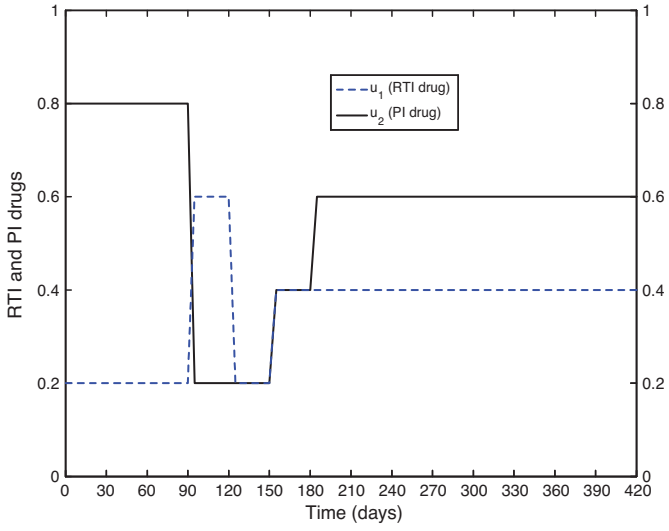

Fig. 9. Optimal strategy (therapy) for children, among all admissible combinations of RTI and PI drugs. Optimization horizon $T = 420$ days.

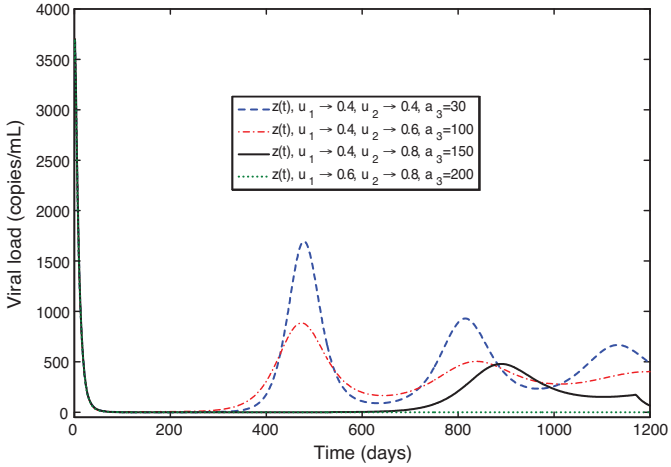

Fig. 10. z-Trajectory for children resulting from optimal combined RTI plus PI medications corresponding to different weights a_3 assigned to thymus' deficiency cost.

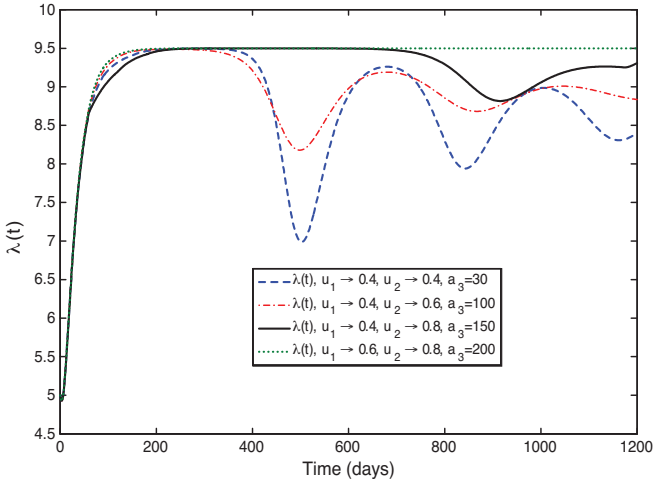

Fig. 11. λ-Trajectory for children resulting from optimal combined RTI plus PI medications corresponding to different weights a_3 assigned to thymus' deficiency cost.

long optimization horizons would require then nonautonomous formulations for both the dynamics and the cost function.

6. Conclusions

An optimal control methodology, suitable for assessing drug combinations in HIV treatment of children and adult patients, has been presented. The problem was posed in the context of deterministic nonlinear systems, although there are at present uncertainties concerning the dynamics of human thymus under viral infection and highly active medication, that would probably require an stochastic setup in the future. It has been found that, providing the model be accurate enough, different drug combinations could result better than monotherapies, and that the best combinations are not completely obvious but discovered through mathematical analysis and numerical calculation. The choice for the doses of two different drugs was to be determined every 30 days, and were to be taken constantly during the next period. The criteria for deciding what was best were expressed as a number of conflicting objectives, and modeled as partial terms of a cost functional, which entails evaluations during the whole optimization time-horizon and at the end of it. Classical Dynamic Programming proved to be useful in finding optimal medication strategies in the presence of many degrees of freedom, after proper adjustment to the mixed continuous-discrete characteristics of this situation. A valuable feature of this methodology is its implicit 'closed-loop', or 'feedback control', character. This means that whenever the state of the patient, reflected by an upcoming blood analysis, is not what was expected, nonetheless the optimal medication corresponding to the 'corrected' state is stored as a part of the numerical results already obtained. No further calculations are needed, since Dynamic Programming provides the optimal control strategies for any initial or intermediate state (discretization allowed).

The problem had essentially two different setups: one for adult patients and another one for just-born children. Parameters for the dynamics and for the evaluation functional of each setup were tuned to reflect existing knowledge, which is still little more than incipient. Results, though, are promising: an inverse tendency in the proportion between the RTI and PI families of drugs seem to be recommendable for children in comparison with that for adults. The appropriate decision, when the concern for thymus function grows in appreciation, would then be to prescribe more PIs than RTIs for children, and the opposite ratio for adults. Another positive outcome is the stabilization of the composition of the optimal doses after some 180 days towards a mix with significant PI share. This is regarded as 'milder' acceptable therapies, especially well received by children patients.

The variations of optimal medication strategies were also assessed when evaluating long periods, the increasing of the total amount of drugs being the expected consequence of minimizing rebound risks for viral load. Still, it has to be realized that: (i) the longer the horizon under study, the greater the uncertainties are in the number and values of parameters and weight coefficients needed for a definitive formulation, and (ii) long-term side-effects have not explicitly been contemplated in the present methodology.

References

[1] P. Cahn, I. Cassetti, R. Wood, P. Phanuphak, L. Shiveley, R.C. Bethell, J. Sawyer, Efficacy and tolerability of 10-day monotherapy with apricitabine in antiretroviral-naive, HIV-1 infected patients, AIDS 20 (9) (2006) 1261–1268.

[2] A. Cato, J. Qian, A. Hsu, B. Levy, J. Leonard, R. Granneman, Multidose pharmacokinetics of ritonavir and zidovudine in human immunodeficiency virus-infected patients, Antimicrobial Agents and Chemotherapy 42 (1998) 1788–1793.

[3] R. Cheynier, R. Bordi, M.L. Dion, R. Woods, J. Montaner, F. Harris, M. Lederman, R.P. Sekaly, Immunological improvement following HAART coincides with restored thymic function, in: IAS Conference on HIV Pathogenesis and Treatment, Unite des Virus Lents, Institut Pasteur, Paris, France, Antiviral Therapy 8 (Suppl. 1) (2003) (abstract no. 15).

[4] R. Consolini, A. Legitimo, M. Milani, Primary thymic endocrine failure in HIV-1-infected children, Pathobiology 68 (6) (2000) 251–257.

[5] V. Costanza, P.S. Rivadeneira, F.L. Biafore, C.E. D'Attellis, A closed-loop approach to antiretroviral therapies for HIV infection, Biomedical Signal Processing and Control 4 (2009) 139–148.

[6] V. Costanza, P.S. Rivadeneira, F.L. Biafore, C.E. D'Attellis, Taking side effects into account for HIV medication, IEEE Transactions on Biomedical Engineering 57 (9) (2010) 2079–2089.

[7] R.T. Davey, HIV-1 and T cell dynamics after interruption of highly active antiretroviral therapy (HAART) in patients with a history of sustained viral suppression, Proceedings of the National Academy of Sciences of the United States of America 98 (26) (1999) 15109–15114.

[8] P. Delobel, K. Sandres-Saune, M. Cazabat, C. Pasquier, B. Marchou, P. Massip, J. Izopet, R5 to X4 switch of the predominant HIV-1 population in cellular reservoirs during effective highly active antiretroviral therapy, Journal of Acquired Immune Deficiency Syndromes 38 (4) (2005) 382–392.

[9] D.C. Douek, R.A. Koup, R.D. McFarland, J.L. Sullivan, K. Luzuriaga, Effect of HIV on thymic function before and after antiretroviral therapy in children, in: 7th Conf. Retrovir. Oppor. Infect., vol. 7, San Francisco, CA, January 30–February 2, 2000, p. 99.

[10] D.C. Douek, R.D. McFarland, P.H. Keiser, E.A. Gage, J.M. Massey, B.F. Haynes, M.A. Polis, A.T. Haase, M.B. Feinberg, J.L. Sullivan, B.D. Jamieson, J.A. Zack, L.J. Picker, R.A. Koup, Changes in thymic function with age and during the treatment of HIV infection, Nature 396 (6712) (1998) 690–695.

[11] R.H. Fang, A.D. Colantonio, C.H. Uittenbogaart, The role of the thymus in HIV infection: a 10-year perspective: the thymus in discordant immunological and virological response, AIDS 22 (2) (2008) 171–184.

[12] P. Faurre, M. Depeyrot, Elements of System Theory, North Holland, Amsterdam, 1977.

[13] R. Hazra, C. Mackall, Thymic function in HIV infection, Current HIV/AIDS Reports 2 (1) (2005) 24–28.

[14] C. Hoffmann, B. Kamps, HIV Medicine, Flying Publisher, 2007, www.HIVmedicine.com.

[15] D.E. Kirschner, R. Mehr, A.S. Perelson, Role of the thymus in pediatric HIV-1 infection, Journal of Acquired Immune Deficiency Syndromes 18 (1998) 95–109.

[16] J.H. Ko, W.H. Kim, C.C. Chung, Optimized structured treatment interruption for HIV therapy and its performance analysis on controllability, IEEE Transactions on Biomedical Engineering 53 (3) (2006) 380–386.

[17] N.R. Manley, Mechanisms controlling thymic homeostasis, involution, and rebound, Project Number: 1P01AI076514-01A1, 2009. http://projectreporter.nih.gov/project_info_details.cfm?aid=7569071.

[18] M.J. Míguez-Burbano, M. Nair, J.E. Lewis, J. Fishman, The role of alcohol on platelets thymus and cognitive performance among HIV-infected subjects: are they related? Platelets 20 (4) (2009) 260–267.

[19] M.A.M.A. Nowak, R.M. May, Pathogenesis-mathematical models of HIV and SIV infections, AIDS 7 (1993) 3–18.

[20] M.A.M.A. Nowak, R.M. May, Virus Dynamics: Mathematical Principles of Immunology and Virology, Oxford University Press, New York, 2000.

[21] P. Palumbo, H. Wu, E. Chadwick, P. Ruan, K. Luzuriaga, J. Rodman, R. Yogev, Virologic response to potent antiretroviral therapy and modeling of HIV dynamics in early pediatric infection, The Journal of Infectious Diseases 196 (1) (2007) 23–29.

[22] A.S. Perelson, D.E. Kirschner, R. de Boer, Dynamics of HIV infection of CD4+T cells, Mathematical Biosciences 114 (1993) 81–125.

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

[24] A.S. Perelson, Modelling viral and immune system dynamics, Nature. Reviews/Immunology 2 (2002) 28–36.

[25] S. Resino, M. Leal, R. Correa, M. Genebat, V. Campos, M.A. Muñoz-Fernández, Immune reconstitution in HIV-infected children on antiretroviral therapy, Immunology 24 (1) (2005) 44–54.

[26] P.J. Ruane, G.J. Richmond, E. DeJesus, C.E. Hill-Zabala, S.C. Danehower, Q. Liao, J. Johnson, M.S. Shaefer, Pharmacodynamic effects of zidovudine 600 mg once/day versus 300 mg twice/day in therapy-naïve patients infected with human immunodeficiency virus, Pharmacotherapy 24 (3) (2004) 307–312.

[27] A. Saitoh, K.K. Singh, S. Sandall, C.A. Powell, T. Fenton, C.V. Fletcher, K. Hsia, S.A. Spector, Association of CD4+T-lymphocyte counts and new thymic emigrants in HIV-infected children during successful highly active antiretroviral therapy, Journal of Allergy and Clinical Immunology 117 (4) (2006) 748–752.

[28] M.F. Thomas, J.M. Vreeland, A socio-economic model for outcome assessment in health care management, in: Proceedings of the Academy of Accounting and Financial Studies, vol. 3(2), Las Vegas, 1998.

[29] P. Ye, A.P. Kourtis, D.E. Kirschner, Reconstitution of thymic function in HIV-1 patients treated with highly active antiretroviral therapy, Clinical Immunology 106 (2003) 95–105.