BIOINFORMATICS

Vol. 18 no. 9 2002
Pages 1227–1235

Optimal enhancement of immune response

Robert F. Stengel, Raffaele M. Ghigliazza and Nilesh V. Kulkarni*

School of Engineering and Applied Science, Princeton University, Princeton, NJ 08540, USA

Received on September 24, 2001; revised on March 20, 2002; accepted on March 26, 2002

ABSTRACT

Motivation: Therapeutic enhancement of innate immune response to microbial attack is addressed as the optimal control of a dynamic system. Interactions between an invading pathogen and the innate immune system are characterized by four non-linear, ordinary differential equations that describe rates of change of pathogen, plasma cell, and antibody concentrations, and of an indicator of organic health. Without therapy, the dynamic model evidences sub-clinical or clinical decay, chronic stabilization, or unrestrained lethal growth of the pathogen; the response pattern depends on the initial concentration of pathogens in the simulated attack. In the model, immune response can be augmented by therapeutic agents that kill the pathogen directly, that stimulate the production of plasma cells or antibodies, or that enhance organ health. A previous paper demonstrated open-loop optimal control solutions that defeat the pathogen and preserve organ health, given initial conditions that otherwise would be lethal (Stengel *et al.*, 2002). Therapies based on separate and combined application of the agents were derived by minimizing a quadratic cost function that weighted both system response and control usage, providing implicit control over harmful side effects.

Results: We demonstrate the ability of neighboring–optimal feedback control to account for a range of unknown initial conditions and persistent input of pathogens by adjusting the therapy to account for perturbations from the nominal-optimal response history. We examine therapies that combine open-loop control of one agent with closed-loop control of another. We show that optimal control theory points the way toward new protocols for treatment and cure of human diseases.

Contact: stengel@princeton.edu; rghiglia@princeton.edu; nkulkarn@princeton.edu

INTRODUCTION

Infectious microbes trigger a dynamic response of the immune system, in which potentially uncontrolled growth of the invader (or *pathogen*) is countered by various protective mechanisms. The outer perimeter of defense consists of the surface *epithelial layers* of the body, including the *epidermal cells* of the skin and the *mucosal cells* that line the respiratory, gastrointestinal, and genitourinary tract (Lydyard *et al.*, 2000; Janeway, 2001; Thain and Hickman, 2000). The *innate immune system* provides a tactical response, signaling the presence of 'non-self' organisms and activating *B cells* to produce *antibodies* that bind to the intruders' *antigens*. The antibodies identify targets for scavenging cells (e.g. *neutrophils* and *macrophages*) that engulf and consume the microbes, reducing them to non-functioning units. They also stimulate the production of *cytokines*, *complement*, and *acute-phase proteins* that either damage an intruder's plasma membrane directly or that trigger the second phase of immune response. The innate immune system protects against many *extracellular bacteria* or *free viruses* found in blood plasma, lymph, tissue fluid, or interstitial space between cells, but it cannot defeat microbes that burrow into cells, such as *viruses*, *intracellular bacteria*, and *protozoa*.

Strategic response to intracellular microbial assault is provided by the *adaptive immune system*, which produces protective cells that remember specific antigens, that produce antibodies to counter the antigens, and that seek out *epitopes* (or defining regions) of antigens on the surfaces of infected cells. Adaptive immune mechanisms depend on the actions of *B* and *T lymphocyte cells* that become dedicated to a single antibody type through *clonal selection*. *Killer T cells* (or *cytotoxic T lymphocytes*) bind to infected cells and kill them by initiating programmed cell death (*apoptosis*). *Helper T cells* assist naive *B* cells in maturing into *plasma cells* that produce the needed antibody type. Immune cells with narrowly focused memory are generated, ready to respond rapidly if invading microbes with the same antigen epitopes are encountered again. Elements of the innate and adaptive immune systems are shared, and response mechanisms are coupled, even though separate modes of operation can be identified.

Here, we address post-exposure therapy for a clinically diagnosed condition. The options available for clinical treatment of infection once it has been recognized focus on killing the invading microbes, neutralizing their

*To whom correspondence should be addressed.

deleterious effects, enhancing the efficacy of immune response, and providing palliative or healing care to other organs of the body. Few biological or chemical agents have just a single effect; for example, an agent that kills a virus also may damage healthy 'self' cells. A critical function of drug discovery and development is to identify new compounds that have maximum intended efficacy with minimal side effects in the general population. Examples include *antibiotics* as microbe killers; *interferons* as microbe neutralizers; *interleukins, antigens* from *killed* (i.e. *non-toxic*) *pathogens*, and *pre-formed* and *monoclonal antibodies* as immunity enhancers (each of very different nature); and *anti-inflammatory* and *anti-histamine compounds* as palliative drugs.

Many models of immune response to infection have been postulated (Asachenkov *et al.*, 1994; Rundell *et al.*, 1995; Perelson, 1997; Nowak and May, 2000), with recent emphasis on the human-immunodeficiency virus (HIV; Nowak *et al.*, 1995; Perelson *et al.*, 1996; Perelson and Nelson, 1999; Wodarz *et al.*, 2000a,b; Stafford *et al.*, 2000). Norbert Wiener (Wiener, 1948) and Richard Bellman (Bellman, 1983) appreciated and anticipated the application of mathematical analysis to treatment in a broad sense, and Swan (1981) surveys early optimal control applications to biomedical problems. Kirschner *et al.* (1997) offers an optimal control approach to HIV treatment, and intuitive control approaches are presented in Bonhoeffer *et al.* (1997); Wein *et al.* (1998); Wodarz and Nowak (1999) and Wodarz and Nowak (2000). The dynamics of drug response (*pharmacokinetics*) are modeled in van Rossum *et al.* (1986) and Robinson (1986), and control theory is applied to drug delivery in Bell and Katusiime (1980); Carson *et al.* (1985); Chizeck and Katona (1985); Schumitsky (1986); Jelliffe (1986); Polycarpou and Conway (1996); Kwong *et al.* (1996); Gentilini *et al.* (2001) and Parker *et al.* (1996).

The present analysis examines regimens for applying drugs in a manner that maximizes efficacy while minimizing side effects and cost. That means killing, neutralizing, or limiting growth of the pathogen with small impact on the patient's health and pocketbook. By illustrating optimal enhancement of immune response, control theory can help discriminate between options for attacking the disease directly, stimulating the immune system, suppressing or amplifying ancillary biological processes, prescribing drug dosage and timing, and specifying multiple-drug therapies.

In the remainder, we present a simple model for the response of the innate immune system to infection and to therapy, review the prior method and results of optimization (Stengel *et al.*, 2002), and introduce a significant extension to the optimal control method of enhancing immune response. The earlier results show not only the progression from an initially life-threatening state to a controlled or cured condition but the optimal history of therapeutic agents that produces that condition. Here, the therapeutic method is extended by adding linear–optimal feedback control to the nominal optimal solution. Perturbations from the expected history of immune state may arise from uncertainty about the initial concentration of pathogen or the continuing introduction of new pathogen. Feedback allows the therapy to be adjusted, expanding the range of infectious intensity that can be treated effectively and efficiently.

A MODEL OF NATURAL AND ENHANCED IMMUNE RESPONSE

A simple model of infectious disease is presented in Asachenkov *et al.* (1994) for the principal purpose of 'studying the general picture of the course of a disease and clarifying some observational results.' There are four components to the model's dynamic state:

x_1 = concentration of a pathogen that expresses a specific foreign antigen;
x_2 = concentration of plasma cells that are specific to the foreign antigen;
x_3 = Concentration of antibodies that bind to the foreign antigen;
x_4 = characteristic of a damaged organ [0 = healthy, 1 = dead].

In Asachenkov *et al.* (1994), the first element of the state is loosely referred to as a concentration of 'viruses', by which is meant a concentration of pathogenic antigens; however, too many elements of the adaptive immune system (most notably helper and killer T cells in various states of activation or infection) are missing for the model to represent response to a viral attack. Nevertheless, the Asachenkov model does characterize qualitative behavior of the innate immune system, and we view x_1 as a concentration of extracellular bacteria that are 'virulent.'

We add idealized therapeutic control agents, u_i, as well as an exogenous input, w, to the model of Asachenkov *et al.* (1994), where:

u_1 = pathogen killer;
u_2 = plasma cell enhancer;
u_3 = antibody enhancer;
u_4 = organ healing factor (or health enhancer);
w = rate of continuing introduction of pathogen.

The structural relationship of system variables is illustrated by Figure 1. Introduction of the pathogen stimulates the production of plasma cells and antibodies and degrades organ health. Organ health mediates plasma cell production, inferring a relationship between immune response and fitness of the individual. Antibodies bind to the attacking antigens, thereby killing pathogenic microbes directly, activating complement proteins, or triggering an attack by phagocytic cells, e.g. macrophages and neutrophils. Each element of the state is subject to independent control, and new microbes may continue to enter the system.

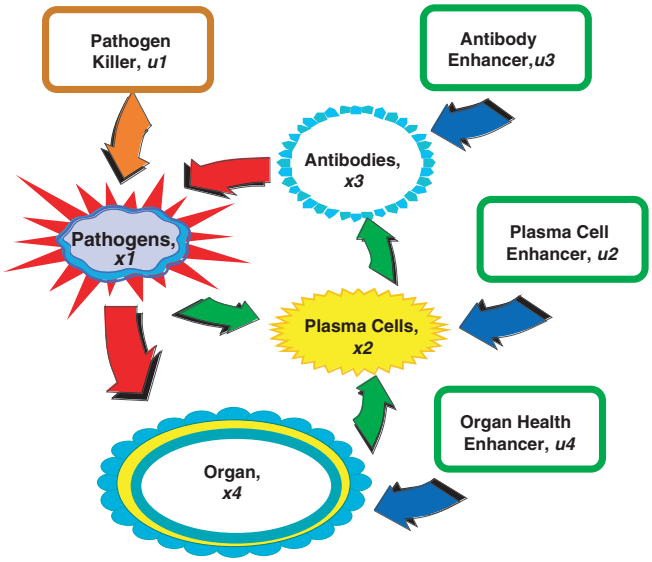

Fig. 1. Innate and enhanced immune response to a pathogenic attack.

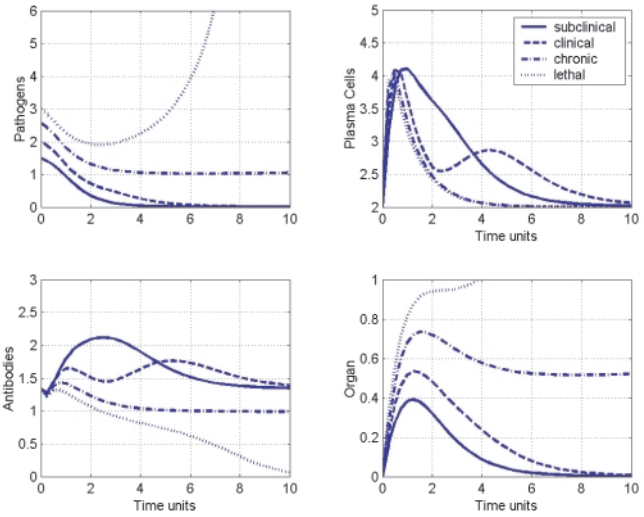

Fig. 2. Natural response to attack by a pathogen (from Stengel *et al.*, 2002).

As in Stengel *et al.* (2002), the four non-linear, ordinary differential equations of the modified dynamic model are:

$$\begin{aligned}
\dot{x}_1 &= (a_{11} - a_{12}x_3)x_1 + b_1u_1 + w \\
\dot{x}_2 &= a_{21}(x_4)a_{22}x_1x_3 \\
&\quad - a_{23}(x_2 - x_2{}^*) + b_2u_2 \\
\dot{x}_3 &= a_{31}x_2 - (a_{32} + a_{33}x_1)x_3 + b_3u_3 \\
\dot{x}_4 &= a_{41}x_1 - a_{42}x_4 + b_4u_4
\end{aligned} \tag{1–4}$$

where the x_i vary with time, except for the steady equilibrium value of $x_2{}^*$, which is 2. Values of the parameters used for this study are: $a_{11} = a_{12} = a_{23} = a_{31} = a_{42} = b_2 = b_3 = 1$, $b_1 = b_4 = -1$, $a_{22} = 3$, $a_{32} = 1.5$, and $a_{33} = a_{41} = 0.5$. $a_{21}(x_4)$ is a non-linear function that describes the mediation of plasma cell generation by the damaged organ:

$$a_{21}(x_4) = \begin{cases} \cos(\pi x_4), & 0 \leq x_4 < \frac{1}{2} \\ 0, & \frac{1}{2} \leqslant x_4 \end{cases} \tag{5}$$

The parameters have been chosen to produce a system that recovers or succumbs to the pathogen (without treatment) as a function of initial conditions during a period of ten time units. Both parameters and time units are abstractions, as no specific disease is addressed. The state and control are always positive because concentrations cannot go below zero, and organ death is indicated when $x_4 = 1$.

Figure 2 (from Stengel *et al.*, 2002) shows typical uncontrolled response to increasing levels of initial pathogen concentration. Conceptually, the sub-clinical response would not require medical examination, while the clinical case warrants medical consultation but is self-healing without intervention. Pathogen concentration stabilizes at non-zero value in the chronic case, which is characterized by permanently degraded organ health, and it diverges in the lethal case, killing the organ. The 'lethal' simulation of Figure 2 is allowed to continue past the point at which x_4 exceeds one for illustrative purposes only. The mathematical model is seen to have a stable equilibrium when $\mathbf{x} = 0$ and a neutrally stable equilibrium in the neighborhood of the chronic solution at the end of the period.

TREATMENT COST FUNCTION AND THE NOMINAL–OPTIMAL CONTROL POLICY

The model of Equations (1–4) allows us to simulate innate immune response to pathogenic attack and to therapy, but it does not tell us what the therapeutic protocol should be. The optimal therapeutic protocol is derived by minimizing a treatment cost function, J, that penalizes large values of pathogen concentration, poor organ health, and excessive application of therapeutic agents over the fixed time interval $[t_0, t_f]$ and at the end of the treatment interval:

$$\begin{aligned}
J = \frac{1}{2}\left(p_{11}x_{1_f}^2 + p_{44}x_{4_f}^2\right) + \frac{1}{2}\int_{t_0}^{t_f} (q_{11}x_1^2 + q_{44}x_4^2 \\
+ r_{11}u_1^2 + r_{22}u_2^2 + r_{33}u_3^2 + r_{44}u_4^2)\mathrm{d}t
\end{aligned} \tag{6}$$

The cost function variables are squared to amplify the effects of large variations and to de-emphasize contributions of small variations. Each squared element is multiplied by

a coefficient (p_{ii}, q_{ii}, or r_{ii}) that establishes the relative importance of the factor in the treatment cost. These coefficients could reflect financial cost, or they could represent physiological 'cost,' such as virulence, toxicity, or discomfort. The weighting coefficients provide a mechanism for trading one variation against the others in defining the treatment protocol, balancing speed and efficacy of treatment against implicit side effects.

The disease dynamic model (Equations 1–4) can be expressed in vector form,

$$\dot{\mathbf{x}}(t) = \mathbf{f}[\mathbf{x}(t), \mathbf{u}(t), \mathbf{w}(t)] \tag{7}$$

where $\mathbf{x}$ is the state, $\mathbf{u}$ is the control, and $\mathbf{w}$ is an exogenous disturbance. The scalar cost function (Equation 6) takes the general form,

$$\begin{aligned} J &= \phi\left[\mathbf{x}(t_f)\right] + \int_{t_0}^{t_f} L\left[\mathbf{x}(t), \mathbf{u}(t)\right] \mathrm{d}t \\ &= \frac{1}{2}\left\{\mathbf{x}^{\mathrm{T}}(t_f)\mathbf{P}_f\mathbf{x}(t_f)\right. \\ &\quad \left. + \int_{t_0}^{t_f} [\mathbf{x}^{\mathrm{T}}(t)\mathbf{Q}\mathbf{x}(t) + \mathbf{u}^{\mathrm{T}}(t)\mathbf{R}\mathbf{u}(t)]\mathrm{d}t\right\} \end{aligned} \tag{8}$$

$L[\mathbf{x}(t), \mathbf{u}(t), t]$ is called the *Lagrangian*, and the p_{ii}, q_{ii}, and r_{ii} are diagonal elements of the matrices $\mathbf{P}$, $\mathbf{Q}$, and $\mathbf{R}$. Defining the *Hamiltonian* of the system,

$$\begin{aligned} &H[\mathbf{x}(t), \mathbf{u}(t), \mathbf{w}(t), \boldsymbol{\lambda}(t), t] \\ &\quad = L[\mathbf{x}(t), \mathbf{u}(t), t] + \boldsymbol{\lambda}^T(\mathbf{t})\mathbf{f}[\mathbf{x}(t), \mathbf{u}(t), \mathbf{w}(t)] \\ &= \frac{1}{2}(q_{11}x_1^2 + q_{44}x_4^2 + r_{11}u_1^2 + r_{22}u_2^2 + r_{33}u_3^2 + r_{44}u_4^2) \\ &+\lambda_1[(a_{11} - a_{12}x_3)x_1 + b_1u_1 + w] \\ &+\lambda_2[a_{21}(x_4)a_{22}x_1x_3 - a_{23}(x_2 - x_2{}^*) + b_2u_2] \\ &+\lambda_3[a_{31}x_2 - (a_{32} + a_{33}x_1)x_3 + b_3u_3] \\ &+\lambda_4[a_{41}x_1 - a_{42}x_4 + b_4u_4] \end{aligned} \tag{9}$$

the necessary conditions for optimizing the cost function with respect to control are expressed by the three *Euler–Lagrange equations* (Stengel, 1994):

$$\begin{aligned} \dot{\boldsymbol{\lambda}}(t) &= -\left\{\frac{\partial H[\mathbf{x}(t), \mathbf{u}(t), \mathbf{w}(t), \boldsymbol{\lambda}(t), t]}{\partial \mathbf{x}}\right\}^T \\ \boldsymbol{\lambda}(t_f) &= \left\{\frac{\partial \phi[x(t_f)]}{\partial \mathbf{x}}\right\}^T \\ 0 &= \frac{\partial H[\mathbf{x}(t), \mathbf{u}(t), \mathbf{w}(t), \boldsymbol{\lambda}(t), t]}{\partial \mathbf{u}}. \end{aligned} \tag{10a–12a}$$

In scalar form, the equations are

$$\begin{aligned} \dot{\lambda}_1 &= -[q_{11}x_1 + \lambda_1(a_{11} - a_{12}x_3) \\ &\quad + \lambda_2 a_{21}(x_4)a_{22}x_3 - \lambda_3 a_{33}x_3 + \lambda_4 a_{41}] \\ \dot{\lambda}_2 &= \lambda_2 a_{23} - \lambda_3 a_{31} \\ \dot{\lambda}_3 &= \lambda_1 a_{12}x_1 - \lambda_2 a_{21}(x_4)a_{22}x_1 + \lambda_3 a_{33}x_1 \\ \dot{\lambda}_4 &= -[q_{44}x_4 + \lambda_2 \frac{\partial a_{21}}{\partial x_4} a_{22}x_1x_3 - \lambda_4 a_{42}] \end{aligned} \tag{10b}$$

$$\begin{aligned} \lambda_1(t_f) &= p_{11}x_1(t_f) \\ \lambda_2(t_f) &= 0 \\ \lambda_3(t_f) &= 0 \\ \lambda_4(t_f) &= p_{44}x_4(t_f) \end{aligned} \tag{11b}$$

$$\begin{aligned} r_{11}u_1 + \lambda_1 b_1 &= 0 \\ r_{22}u_2 + \lambda_2 b_2 &= 0 \\ r_{33}u_3 + \lambda_3 b_3 &= 0 \\ r_{44}u_4 + \lambda_4 b_4 &= 0. \end{aligned} \tag{12b}$$

The Euler–Lagrange equations include a linear, ordinary-differential equation whose integral is the *adjoint vector*, $\boldsymbol{\lambda}(t)$ (Equation 10), a terminal boundary condition that specifies $\boldsymbol{\lambda}(t_f)$ at the end of the interval (Equation 11), and a *stationarity condition* on the control throughout the interval (Equation 12). Here, the disturbance, $\mathbf{w}(t)$, is treated as a known parameter. Equations (7) and (10–12) must be satisfied concurrently, specifying a *two-point boundary-value problem* that is solved numerically. The solution is necessarily iterative because the system model is non-linear, initial conditions are given for $\mathbf{x}$, and terminal conditions are given for $\boldsymbol{\lambda}$. The solution is initiated by solving Equation (7) with a starting guess for the control history, $\mathbf{u}_o(t)$ in $[t_0, t_f]$. Equation (10) is solved, integrating back from the end conditions specified by Equation (11). In general, the remaining necessary condition for optimality, Equation (12), is not satisfied, so a *steepest-descent method* is used to generate successive approximations of the optimal control history $\mathbf{u}^*(t)$ from

$$\mathbf{u}_k(t) = \mathbf{u}_{k-1}(t) - c\frac{\partial H(\mathbf{x}, \mathbf{u}, \mathbf{w}, \boldsymbol{\lambda}, t)}{\partial \mathbf{u}} \tag{13}$$

where c is a small positive constant and k is the iteration index (Stengel *et al.*, 2002). For k sufficiently large, $\partial H/\partial \mathbf{u}$ tends to zero, and $\mathbf{u}(t)$ converges to the optimal control history that is applied in Equations (1–4) or (7).

Nominal–optimal solutions computed for otherwise-lethal initial conditions and unit cost function weights are presented in Figure 3 (from Stengel *et al.*, 2002). Finding the control history that minimizes the cost function typically requires 10–20 steepest-descent iterations. The

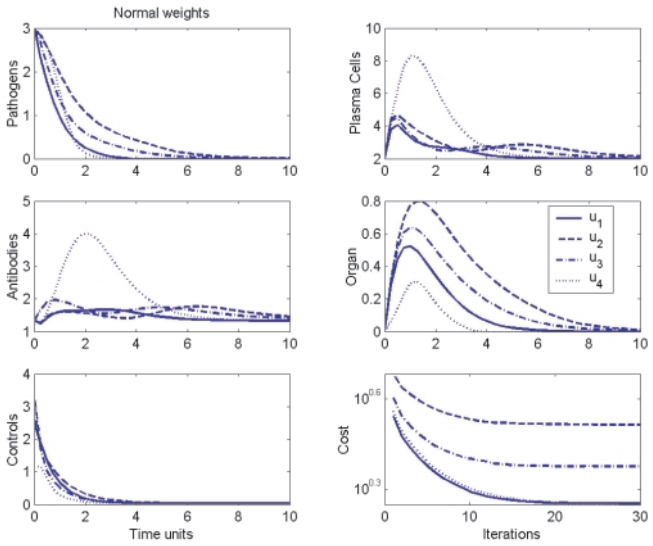

Fig. 3. Optimal therapies with unit cost-function weights and scalar controls (from Stengel *et al.*, 2002).

example shows that each of the therapeutic agents used separately can defeat the pathogen and maintain organ health with varying (but important) participation of the innate immune system. All of the therapeutic protocols specify an initially strong dose of the agent followed by exponential decay. These results infer that combination therapies can be even more effective in defeating the pathogen and maintaining organ health than treatment by a single agent, and the present neighboring-optimal approach confirms the inference.

While the drug concentration decays over time in Figure 3, this is not a pharmacokinetic effect, as the model contains no dynamics of drug uptake. The reduction in therapeutic level is prescribed by the optimization procedure alone. If the drug is consumed or eliminated at a rate greater than that shown in the figure, then additional dosage is required to maintain the level prescribed by the optimization.

If the initial concentration of pathogen is changed from its nominal value, the nominal–optimal therapy is no longer optimal, and a new regimen must be defined to retain optimality. For small increase in initial pathogen concentration, the combination of the innate immune system and the nominal–optimal control policy prevails, and the pathogen is defeated, though the response history is no longer optimal. For a strong enough assault, the combination of immune response and therapy is insufficient, and the pathogen grows without bound, killing the organ. The therapeutic protocol must be adjusted to accommodate the change, either through continued re-evaluation of the nominal–optimal policy or through a simpler mechanism for modifying the policy in proportion to deviations from the expected response history. In the remainder of the paper, we show that linear–optimal feedback control provides a good mechanism for adjusting the therapy, and it provides a simple means for introducing additional therapeutic agents that are not included in the nominal–optimal protocol.

NEIGHBORING–OPTIMAL CONTROL POLICY

Enhancement of the optimal therapy can be based on the solution of a neighboring–optimal control problem, as presented in Stengel (1994) and elsewhere. Actual state and control histories can always be represented as sums of the optimal histories derived from the iterative procedure, $\mathbf{x}^*(t)$ and $\mathbf{u}^*(t)$, and deviations from those histories, $\Delta\mathbf{x}(t)$ and $\Delta\mathbf{u}(t)$,

$$\mathbf{x}(t) = \mathbf{x}^*(t) + \Delta\mathbf{x}(t) \tag{14}$$

$$\mathbf{u}(t) = \mathbf{u}^*(t) + \Delta\mathbf{u}(t) \tag{15}$$

Neglecting the disturbance effect, Equation (7) can be expanded as,

$$\begin{aligned}\dot{\mathbf{x}}(t) &= \dot{\mathbf{x}}^*(t) + \Delta\dot{\mathbf{x}}(t)\\ &= \mathbf{f}([\mathbf{x}^*(t) + \Delta\mathbf{x}(t)], [\mathbf{u}^*(t) + \Delta\mathbf{u}(t)])\\ &= \mathbf{f}[\mathbf{x}^*(t), \mathbf{u}^*(t)] + \frac{\partial\mathbf{f}}{\partial\mathbf{x}}\Delta\mathbf{x}(t) + \frac{\partial\mathbf{f}}{\partial\mathbf{u}}\Delta\mathbf{u}(t) + \cdots\\ &\simeq \mathbf{f}[\mathbf{x}^*(t), \mathbf{u}^*(t)] + \mathbf{F}(t)\Delta\mathbf{x}(t) + \mathbf{G}(t)\Delta\mathbf{u}(t)\end{aligned} \tag{16}$$

where $\mathbf{F}(t)$ and $\mathbf{G}(t)$ are the time-varying *Jacobian matrices* evaluated along the nominal–optimal history. Because the nominal–optimal solution satisfies Equation (7), the dynamics of the perturbed state are closely approximated by the linear, time-varying equation

$$\Delta\dot{\mathbf{x}}(t) = \mathbf{F}(t)\Delta\mathbf{x}(t) + \mathbf{G}(t)\Delta\mathbf{u}(t) \tag{17a}$$

when perturbations are small deviations from the nominal solution. This equation can be written explicitly as

$$\begin{bmatrix}\Delta\dot{x}_1\\ \Delta\dot{x}_2\\ \Delta\dot{x}_3\\ \Delta\dot{x}_4\end{bmatrix} = \begin{bmatrix}(a_{11} - a_{12}x_3) & 0 & -a_{12}x_1 & 0\\ a_{21}(x_4)a_{22}x_3 & -a_{23} & a_{21}(x_4)a_{22}x_1 & \frac{\partial a_{21}}{\partial x_4}a_{22}x_1x_3\\ -a_{33}x_3 & a_{31} & -(a_{32} + a_{33}x_1) & 0\\ a_{41} & 0 & 0 & -a_{42}\end{bmatrix} \times \begin{bmatrix}\Delta x_1\\ \Delta x_2\\ \Delta x_3\\ \Delta x_4\end{bmatrix} + \begin{bmatrix}b_1 & 0 & 0 & 0\\ 0 & b_2 & 0 & 0\\ 0 & 0 & b_3 & 0\\ 0 & 0 & 0 & b_4\end{bmatrix}\begin{bmatrix}\Delta u_1\\ \Delta u_2\\ \Delta u_3\\ \Delta u_4\end{bmatrix} \tag{17b}$$

Optimal histories for this model of perturbed response are derived from the same conditions as before (Equations 10–12), but we redefine the cost function and Hamiltonian as functions of the perturbation variables. The new cost function is the second variation of the original cost function,

$$\Delta^2 J = \phi[\Delta\mathbf{x}(t_f)] + \int_{t_0}^{t_f} L[\Delta\mathbf{x}(t), \Delta\mathbf{u}(t)]\mathrm{d}t$$
$$= \frac{1}{2}\left\{\Delta\mathbf{x}^{\mathrm{T}}(t_f)\mathbf{P}_f\Delta\mathbf{x}(t_f) + \int_{t_0}^{t_f}[\Delta\mathbf{x}^{\mathrm{T}}(t)\mathbf{Q}\Delta\mathbf{x}(t) + \Delta\mathbf{u}^{\mathrm{T}}(t)\mathbf{R}\Delta\mathbf{u}(t)]\mathrm{d}t\right\} \quad (18a)$$

or

$$\Delta^2 J = \frac{1}{2}(p_{11}\Delta x_{1_f}^2 + p_{44}\Delta x_{4_f}^2) + \frac{1}{2}\int_{t_0}^{t_f}(q_{11}\Delta x_1^2 + q_{44}\Delta x_4^2 + r_{11}\Delta u_1^2 + r_{22}\Delta u_2^2 + r_{33}\Delta u_3^2 + r_{44}\Delta u_4^2)\mathrm{d}t \quad (18b)$$

and the Hamiltonian is expressed as,

$$H[\Delta\mathbf{x}(t), \Delta\mathbf{u}(t), \Delta\boldsymbol{\lambda}(t)] = L[\Delta\mathbf{x}(t), \Delta\mathbf{u}(t)] + \Delta\boldsymbol{\lambda}^{\mathrm{T}}(t)[\mathbf{F}(t)\Delta\mathbf{x}(t) + \mathbf{G}(t)\Delta\mathbf{u}(t)] \quad (19)$$

where $\Delta\boldsymbol{\lambda}(t)$ is the adjoint vector for the linearized system. The Euler–Lagrange equations (Equations 10–12) can be applied to the variational system described by Equations (17–19). From Equation (12), the control perturbation can be expressed as

$$\Delta\mathbf{u}^*(t) = -\mathbf{R}^{-1}\mathbf{G}^{\mathrm{T}}(t)\Delta\boldsymbol{\lambda}(t). \quad (20)$$

Furthermore, the terminal condition for the adjoint vector (Equation 11) is of the form,

$$\Delta\boldsymbol{\lambda}(t_f) = \mathbf{P}(t_f)\Delta\mathbf{x}(t_f) \quad (21)$$

and as $\Delta\boldsymbol{\lambda}$ and $\Delta\mathbf{x}$ are adjoint, eq. 21 applies over the entire interval.

Because Equations (10) and (17) are linear, ordinary differential equations, the neighboring optimization is subject to a linear dynamic constraint (Equation 17), and the optimal control policy is a *linear feedback control law* (Stengel, 1994):

$$\Delta\mathbf{u}^*(t) = -\mathbf{C}(t)\ \Delta\mathbf{x}(t) \quad (22)$$

$\mathbf{C}(t)$ is the time-varying *optimal gain matrix* given by,

$$\mathbf{C}(t) = \mathbf{R}^{-1}\mathbf{G}^{\mathrm{T}}(t)\mathbf{P}(t) \quad (23)$$

and $\mathbf{P}(t)$ is the solution to a *matrix Riccati equation* (Stengel, 1994):

$$\dot{\mathbf{P}}(t) = -\mathbf{F}^{\mathrm{T}}(t)\mathbf{P}(t) - \mathbf{P}(t)\mathbf{F}(t) + \mathbf{P}(t)\mathbf{G}(t)\mathbf{R}^{-1} \times\mathbf{G}^{\mathrm{T}}(t)\mathbf{P}(t) - \mathbf{Q}, \quad \mathbf{P}(t_f) = \mathbf{P}_f \quad (24)$$

The dimensions of all matrices follow from the original problem specification. Thus, the feedback gain matrix, $\mathbf{C}(t)$, is calculated just once for the nominal–optimal therapeutic history. From Equations (14) and (15), the optimal control policy accounts for previously unknown initial condition perturbations in the form

$$\mathbf{u}(t) = \mathbf{u}^*(t) - \mathbf{C}(t)\Delta\mathbf{x}(t) = \mathbf{u}^*(t) - \mathbf{C}(t)[\mathbf{x}(t) - \mathbf{x}^*(t)] \quad (25a)$$

or

$$\begin{bmatrix} u_1(t) \\ u_2(t) \\ u_3(t) \\ u_4(t) \end{bmatrix} = \begin{bmatrix} u_1{}^*(t) \\ u_2{}^*(t) \\ u_3{}^*(t) \\ u_4{}^*(t) \end{bmatrix} - \begin{bmatrix} c_{11}(t) & c_{12}(t) & c_{13}(t) & c_{14}(t) \\ c_{21}(t) & c_{22}(t) & c_{23}(t) & c_{24}(t) \\ c_{31}(t) & c_{32}(t) & c_{33}(t) & c_{34}(t) \\ c_{41}(t) & c_{42}(t) & c_{43}(t) & c_{44}(t) \end{bmatrix} \times \left[\begin{bmatrix} x_1(t) \\ x_2(t) \\ x_3(t) \\ x_4(t) \end{bmatrix} - \begin{bmatrix} x_1{}^*(t) \\ x_2{}^*(t) \\ x_3{}^*(t) \\ x_4{}^*(t) \end{bmatrix}\right] \quad (25b)$$

where $\mathbf{x}(t)$ is the state measured at the time therapy is applied, $\mathbf{x}^*(t)$ is its nominal–optimal value, and $\mathbf{u}(t)$ is the total therapy applied at time t. The optimal treatment protocol is prescribed by the nominal–optimal control history, the time-dependent gain matrix, and the difference between the observed response and the nominal–optimal response.

APPLICATION TO ENHANCED IMMUNE RESPONSE

To illustrate the effect of nominal–plus neighboring–optimal therapy, we first compute nominal–optimal treatment for a given initial concentration of pathogen, then increase the infectious load to a point where the no-longer-optimal therapy fails. The effects of neighboring–optimal control are then shown. Holding the initial pathogen concentration at its original value, we also consider a case with a continuing influx of the infectious agent. The effects of single- and multi-agent feedback therapies are demonstrated for the more stressing cases.

Single-agent therapies

Our previous study (Stengel *et al.*, 2002) revealed that the pathogen killer, u_1, and antibody enhancer, u_3, were the most effective individual controls, so we focus on those two here. The cost–function weighting matrices, $\mathbf{Q}$ and $\mathbf{R}$, are the same for both the nominal and neighboring optimizations. Except as noted, $\mathbf{Q}$ is an identity matrix,

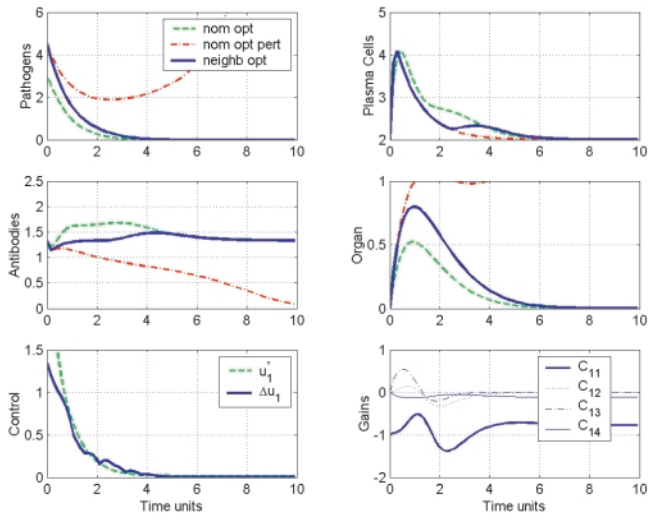

Fig. 4. Comparision of control responses with nominal and increased initial concentrations of pathogen. Nominal initial pathogen with nominal–optimal control (dash), increased initial pathogen with nominal–optimal control (dot–dash), and increased initial pathogen with nominal–optimal plus neighbouring–optimal control (solid) of pathogen killer (u_1).

the diagonal term of **R** corresponding to the single control is one, and all other elements of **R** are zero.

The effects of treatment with the pathogen killer alone are shown in Figure 4. The nominal therapy, $u_1{}^*(t)$, controls the pathogen and preserves organ health with the assumed initial pathogen concentration but not with the increased microbial assault. The principal reasons for failure are that the therapy decays exponentially with time and the damage to the organ allows the antibody concentration to drop off as well. Adding the neighboring–optimal control

$$\begin{aligned}\Delta u_1(t) = &-c_{11}(t)\Delta x_1(t) - c_{12}(t)\Delta x_2(t)\\ &-c_{13}(t)\Delta x_3(t) - c_{14}(t)\Delta x_4(t)\end{aligned} \qquad (26)$$

to the nominal–optimal control, $u_1{}^*(t)$, produces a response that parallels the original one. The additional infusion of u_1 provides a stronger response to the pathogen and preserves antibody concentration. Maximum degradation of organ health is larger, but health eventually is restored. The feedback gains that provide this beneficial effect are also shown in the figure. The predominant (and lasting) effect is possible feedback of pathogen concentration to pathogen killer through c_{11}, though there is also a small continuing feedback of organ health through c_{14}. All of the gains possess a two-time-unit transient near the beginning of the history.

Basing the therapy on antibody enhancer alone, $u_3{}^*(t)$, produces similar nominal and off-nominal results (Figure 5; see Internet supplement). The response with neighboring–optimal control, given by $u_3(t) = u_3{}^*(t) + \Delta u_3(t)$, where

$$\begin{aligned}\Delta u_3(t) = &-c_{31}(t)\Delta x_1(t) - c_{32}(t)\Delta x_2(t)\\ &-c_{33}(t)\Delta x_3(t) - c_{34}(t)\Delta x_4(t)\end{aligned} \qquad (27)$$

is successful in defeating the pathogen, but only if we ignore the fact that the organ health indicator briefly reaches one, signifying organ death. Had the initial condition been slightly smaller or the cost-function weight on organ health been higher, the treatment would have succeeded. The response is generally slower than that of the previous case, and build-up of both plasma cells and antibodies is greater than before. Plasma cell and antibody concentrations evidence a secondary response that begins after five time units. The intensity of feedback effect can be increased by decreasing the cost function weight on control, r_{33}, which increases control gains, quickening the response and preserving long-term organ health (not shown). Nevertheless, the feedback gains all approach zero over time, precluding lasting protection without redefining the control solution.

Combined therapies

It is straightforward to compute the nominal–optimal history using one control variable and the neighboring–optimal control law using another. In such a case, the cost function control weighting matrices, **R**, for the two solutions are not the same. We examine some combined therapies based upon the pathogen killer, antibody enhancer, and organ health enhancer.

The response and control gains are like those of the previous case when nominal–optimal pathogen killer is combined with neighboring–optimal antibody enhancer or when the nominal control is antibody enhancer, and the feedback control is pathogen killer (Figure 6 and 7; see Internet supplement). The net responses of these two cases are quite similar, though the control profiles and gains differ. While the organ is not killed in either case, the cost function weight on organ health (q_{44}) could be increased to widen the margin between ill health and organ death.

Feedback enhancement of organ health is unsuccessfully combined with nominal–optimal enhancement of antibodies (Figure 8; see Internet supplement). The effect of u_4 on the pathogen is indirect, stimulating organ health, which enhances plasma cell production, which enhances antibody production, which kills the pathogen. Antibody concentration does not increase enough to prevent runaway of the infection, and the organ is quickly killed. Examining the feedback gains, we see that the control is insensitive to the pathogen perturbation (via c_{41}) until it is too late to make a difference.

A much better result is obtained when the penalty on organ health enhancer use is reduced ($r_{44} = 0.1$) and

that feedback control is combined with nominal–optimal pathogen killer (Figure 9; see Internet supplement). The feedback gains now recognize the importance of the pathogen petrubation and mount an early defense. Although plasma cell and antibody petrubations have a negligible effect on the feedback control, both of their concentrations remain strong throughout the episode. Furthermore, the direct effect of u_4 on the rate of change of x_4 prevents its excursion from being much worse than in the nominal–optimal case.

Response to continuing infection

We may anticipate circumstances in which infection is characterized not only by a large initial concentration at the beginning of a treatment period but by the continuing re-infection of microbes sequestered in regions not directly modeled by Equations (1–4). As a preliminary look at this problem, we add a constant disturbance of $w = 0.5$ to the simulation and return to single-agent control with the pathogen killer. The nominal-optimal control is computed for $w = 0$, but the feedback gains are computed for a second optimal history that assumes $w = 0.5$. Thus, the gains are not the same as those shown in Figure 4. Applying the nominal–optimal control to the disturbed system results in divergence, even though the initial infection is unchanged (Figure 10; see Internet supplement). The characteristic loss of antibodies and organ degradation are apparent. Neighboring–optimal control prevents the divergence, and system response is nearly nominal. The feedback control senses the continuing intrusion and injects a bias control that effectively cancels it. As a consequence, both plasma cells and antibodies are kept at higher levels, adding to the defense.

CONCLUSION

There are compelling reasons to apply control theory to the treatment of disease, as mathematical models of the disease processes alone do not reveal possibly counter-intuitive approaches to therapy. Optimal control methods are particularly well suited to the problem because they show the best that can be done within assumptions and provide a framework within which alternatives can be evaluated. The critical challenge is not to develop new theory but to apply what we know in a reasoned manner. Choosing elements of cost functions, specifying available control variables, and most importantly, using credible, reliable models of pathogenic attack, direct effects of therapy, and immune system response must be the focus. Here, we demonstrate how numerical optimization of nonlinear models can be combined with neighboring–optimal feedback control of linear models to suggest single- and multi-agent therapies that enhance the natural response of the innate immune system. We also show that theory cannot be applied uncritically, that numbers and values make important differences. Ultimately, a combination of mathematics and empiricism can solve real problems and improve the quality of life.

ACKNOWLEDGEMENTS

This research has been supported by a grant from the Alfred P. Sloan Foundation. Mr Ghigliazza has received partial support from the Burroughs-Wellcome Fund for Biological Dynamics.

SUPPLEMENTARY MATERIAL ON THE INTERNET

All figures for this paper can be found at http://www.princeton.edu/∼stengel/bioinfor.pdf.

REFERENCES

Asachenkov,A., Marchuk,G., Mohler,R. and Zuev,S. (1994) *Disease Dynamics*. Birkhauser, Boston.

Bell,D.J. and Katusiime,F. (1980) A time-optimal drug displacement problem. *Optimal Control Applications and Methods*, **1**, 217–225.

Bellman,R.E. (1983) *Mathematical Methods in Medicine*. World Scientific, Singapore.

Bonhoeffer,S., May,R.M., Shaw,G.M. and Nowak,M.A. (1997) Virus dynamics and drug therapy. *Proc. Natl Acad. Sci. USA*, **94**, 6971–6976.

Carson,E.R., Cramp,D.G., Finkelstein,L. and Ingram,D. (1985) Control system concepts and approaches in clinical medicine. In Carson,E.R. and Cramp,D.G. (eds), *Computers and Control in Clinical Medicine*. Plenum Press, New York, pp. 1–26.

Chizeck,H.J. and Katona,P.G. (1985) Closed-loop control. In Carson,E.R. and Cramp,D.G. (eds), *Computers and Control in Clinical Medicine*. Plenum Press, New York, pp. 95–151.

Gentilini,A., Morari,M., Bieniok,C., Wymann,R. and Schnider,T. (2001) Closed-loop control of analgesia in humans. *Proc. IEEE Conf. Decision and Control*. Orlando, pp. 861–867.

Janeway,C.A. (2001) In Travers,P., Walport,M. and Shlomchik,M. (eds), *Immunobiology*. Garland, London.

Jelliffe,R.W. (1986) Clinical applications of pharmacokinetics and control theory: planning, monitoring, and adjusting regimens of aminoglycosides, lidocaine, digitoxin, and digoxin. In Maronde,R.F. (ed.), *Topics in Clinical Pharmacology and Therapeutics*. Springer, New York, pp. 26–82.

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

Kwong,G.K., Kwok,K.E., Finegan,B.A. and Shah,S.L. (1996) Clinical evaluation of long range adaptive control for mean arterial blood pressure regulation. *Proc. Am. Control Conf.*. Seattle, pp. 786–790.

Lydyard,P.M., Whelan,A. and Fanger,M.W. (2000) *Instant Notes in Immunology*. Springer, New York.

Nowak,M.A. and May,R.M. (2000) *Virus Dynamics: Mathematical Principles of Immunology and Virology*. Oxford University Press, Oxford.

Nowak,M.A., May,R.M., Phillips,R.E., Rowland-Jones,S., Lalloo,D.G., McAdam,S., Klenerman,P., Köppe,B., Sigmund,K.,

Bangham,C.R.M. and McMichael,A.J. (1995) Antigenic oscillations and shifting immunodominance in HIV-1 infections. *Nature*, **375**, 606–611.

Parker,R.S., Doyle,III,F.J., Harting,J.E. and Peppas,N.A. (1996) Model predictive control for infusion pump insulin delivery. *Proceedings of the 18th Annual Conference of IEEE Engineering in Medicine and Biology Soc*. Amsterdam, pp. 1822–1828.

Perelson,A.S., Neumann,A.V., Markowitz,M., Leonard,J.M. and Ho,D.D. (1996) HIV-1 dynamics in vivo: virion clearance rate, infected cell lifespan, and viral generation time. *Science*, **271**, 1582–1586.

Perelson,A.S. (1997) Immunology for physicists. *Reviews of Modern Physics*, **69**, 1219–1267.

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

Polycarpou,M.M. and Conway,J.Y. (1996) Modeling and control of drug delivery systems using adaptive neural control methods. *Proc. Am. Control Conf*. Seattle, pp. 781–785.

Robinson,D.C. (1986) Topics in Clinical Pharmacology and Therapeutics. In Maronde,R.F. (ed.), *Principles of Pharmacokinetics*. Springer, New York, pp. 1–12.

Rundell,A., HogenEsch,H. and DeCarlo,R. (1995) Enhanced modeling of the immune system to incorporate natural killer cells and memory. *Proc. Am. Control Conf.*. Seattle, pp. 255–259.

Schumitsky,A. (1986) Stochastic control of pharmacokinetic systems. In Maronde,R.F. (ed.), *Topics in Clinical Pharmacology and Therapeutics*. Springer, New York, pp. 13–25.

Stafford,M.A., Cao,Y., Ho,D.D., Corey,L. and Perelson,A.S. (2000) Modeling plasma virus concentration and CD4+ T cell kinetics during primary HIV infection. *J. Theo. Biol.*, **203**, 285–301.

Stengel,R.F. (1994) *Optimal Control and Estimation*. Dover, New York, (originally published as *Stochastic Optimal Control: Theory and Application*, Wiley, New York, 1986).

Stengel,R.F., Ghigliazza,R., Kulkarni,N. and Laplace,O. (2002) Optimal control of innate immune response. *Opt. Control Appl. & Meth.*, **23**, 91–104.

Swan,G.W. (1981) Optimal control applications in biomedical engineering—a survey. *Opt. Control Appl. & Meth.*, **2**, 311–314.

Thain,M. and Hickman,M. (2000) *The Penguin Dictionary of Biology*. Penguin, London.

van Rossum,J.M., Steyger,O., van Uem,T., Binkhorst,G.J. and Maes,R.A.A. (1986) Pharmacokinetics by using mathematical systems dynamics. In Eisenfeld,J. and Witten,M. (eds), *Modelling of Biomedical Systems*. Elsevier, Amsterdam, pp. 121–126.

Wein,L.M., D'Amato,R.M. and Perelson,A.S. (1998) Mathematical analysis of antiretroviral therapy aimed at HIV-1 eradication or maintenance of low viral loads. *J. Theo. Biol.*, **192**, 81–98.

Wiener,N. (1948) *Cybernetics: or Control and Communication in the Animal and the Machine*. Technology Press, Cambridge.

Wodarz,D. and Nowak,M.A. (1999) Specific therapy regimes could lead to long-term immunological control of HIV. *Proc. Natl Acad. Sci. USA*, **96**, 14464–14469.

Wodarz,D. and Nowak,M.A. (2000) CD8 memory, immunodominance, and antigenic escape. *Eur. J. Immun.*, **30**, 2704–2712.

Wodarz,D., May,R.M. and Nowak,M.A. (2000a) The role of antigen-independent persistence of memory cytotoxic T lymphocytes. *Int. Immun.*, **12**, 467–477.

Wodarz,D., Page,K.M., Arnaout,R.A., Thomsen,A.R., Lifson,J.D. and Nowak,M.A. (2000b) A new theory of cytotoxic T-lymphocyte memory: implications for HIV treatment. *Phil. Trans. Roy. Soc.* B, **355**, 329–343.