



Stochastic optimal therapy for enhanced immune response

Robert F. Stengel *, Raffaele Ghigliazza

*Department of Mechanical and Aerospace Engineering, P.O. Box CN5263, Princeton University,
School of Engineering and Applied Science, D-202 Engineering Quadrangle, Princeton, NJ 08544, USA*

Received 28 August 2003; received in revised form 22 June 2004; accepted 23 June 2004

Abstract

Therapeutic enhancement of humoral immune response to microbial attack is addressed as the stochastic optimal control of a dynamic system. Without therapy, the modeled immune response depends upon the initial concentration of pathogens in a simulated attack. Immune response can be augmented by agents that kill the pathogen directly, that stimulate the production of plasma cells or antibodies, or that enhance organ health. Using a generic mathematical model of immune response to the infection (i.e., of the *dynamic state* of the system), previous papers demonstrated optimal (open-loop) and neighboring-optimal (closed-loop) control solutions that defeat the pathogen and preserve organ health, given initial conditions that otherwise would be lethal [Optimal Contr. Appl. Methods 23 (2002) 91, Bioinformatics 18 (2002) 1227]. Therapies based on separate and combined application of the agents were derived by minimizing a quadratic cost function that weighted both system response and drug usage, providing implicit control over harmful side effects.

Here, we focus on the effects that corrupted or incomplete measurements of the dynamic state may have on neighboring-optimal feedback control. Imperfect measurements degrade the precision of feedback adjustments to therapy; however, optimal state estimation allows the feedback strategy to be implemented with incomplete measurements and minimizes the expected effects of measurement error. Complete observability of the perturbed state for this four state example is provided by measurement of four of the six possible pairs of two variables, either set of three variables, or all four variables. The inclusion of state estimation extends the applicability of optimal control theory for developing new therapeutic protocols to enhance immune response.

© 2004 Elsevier Inc. All rights reserved.

* Corresponding author. Tel.: +1 609 258 5103; fax: +1 609 258 6109.
E-mail address: stengel@princeton.edu (R.F. Stengel).

Keywords: Stochastic optimal control; Immune system; Drug therapy

1. 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. Initially, the *innate immune system* provides a non-specific tactical response, killing what pathogen it can, inducing inflammation and vasodilation that aids the defense, causing blood coagulation that slows the spread of infection to other parts of the body, and raising the alarm for more complete response. In the process, a *humoral response* is initiated, signaling the presence of extracellular ‘non-self’ organisms and activating B cells to become *plasma cells* that are specific to the intruders’ *antigens*. The plasma cells produce *antibodies* that bind to the antigens, mediating the destruction of pathogens by various modalities [3–5]. The *adaptive immune system* provides a strategic response that is tailored to the primary attack, producing B and T cells, as well as a host of molecules, that defeat specific intracellular pathogens by binding to infected cells and either killing them outright, inducing programmed cell death, or signaling other cells to finish the job. Additional B and T cells with narrowly focused memory also are produced; they can respond rapidly if invading microbes of the same type are encountered again. Innate, humoral, and adaptive immune responses are coupled, even though separate modes of operation can be identified.

Many models of immune response to infection have been postulated [6–9], with recent emphasis on the human-immunodeficiency virus [10–15]. Norbert Wiener and Richard Bellman appreciated and anticipated the application of mathematical analysis to treatment in a broad sense [16,17], and Swan surveyed early optimal control applications to biomedical problems in [18–20]. Optimal control theory was postulated as an organizing principle for natural immune system behavior in [21–24], and it is applied to HIV treatment in [25,26]. Intuitive control approaches are presented in [27–30]. The dynamics of drug response (*pharmacokinetics*) are modeled in [31,32], and control theory is applied to drug delivery in [33–41].

In the remainder, we consider therapy that enhances humoral immune response to a pathogen, such as a toxin or extracellular bacterium. The options available for clinical treatment of the infection are to kill the invading microbes, to neutralize their harmful effects, to enhance the efficacy of immune response, to provide healing care to organs that are damaged by the microbes, or to employ some combination of therapies.

In prior studies, we examined remedial treatments with differing hypotheses about the initial pathogen concentration. If the initial concentration is known precisely [1], the optimizing control history maximizes efficacy of the drug while minimizing its side effects and cost. For the second study [2], a feedback strategy based on a linear perturbation model of response dynamics is derived to account for variations induced by unknown initial infection. The therapy is modified as a function of the difference between the optimal and observed dynamic states over the entire treatment period, assuming that the difference is measured without error. This feedback approach is approximately optimal with zero-mean random disturbance inputs and perfect measurements of the state (e.g., with unmodeled variability of the infectious agent or small errors in the dynamic model itself that are modeled as ‘process noise’) [42]. If, however, the measurements used

for feed-back therapy contain error or are incomplete, then the perturbed state must be estimated to account for the effects of imperfect knowledge.

In this paper, we incorporate a linear-optimal state estimator in the feedback therapy to minimize the effects of measurement error and to account for missing measurements. To provide a consistent pedagogical basis, we use the model of immune response that was employed in [1,2]. The dynamic model and the deterministic optimal solutions are reviewed briefly. We introduce basic concepts of controllability and observability, illustrating how they apply to this problem. The feedback control law is recast as a stochastic neighboring-optimal control problem, whose solution can be partitioned into the control law derived in [2] plus the linear-optimal estimator (or Kalman–Bucy filter [43]) presented here. Examples of therapeutic effects in the presence of measurement error and incompleteness are given.

2. Immune response model

The mathematical model used in [1,2] is an idealization of a generic humoral immune response that subsumes many details into aggregated effects [6]. There are four components in the model's *dynamic state*: concentration of a foreign pathogen that expresses an identifying antigen (x_1), concentration of plasma cells that are specific to the antigen (x_2), concentration of antibodies that bind to the antigen (x_3) and mediate the destruction of the pathogen, and a measure of the health of an organ (x_4) that may be damaged by infection. The model presented in [6] does not account for treatment. We add idealized active and passive immunotherapeutic control agents, u_i , as well as an exogenous input, w_i , to the model: pathogen killer (u_1), plasma cell enhancer (u_2), antibody enhancer (u_3), organ healing factor (u_4), and uncertainties in continuing pathogen input and immune response (w_i , $i = 1-4$). (*Active immunotherapy* strengthens natural immune response, as by enhancing plasma cell and antibody production, while *passive immunotherapy* addresses the effects of infection directly, as in killing the pathogen or healing the infected organ.)

Introduction of the pathogen stimulates the production of plasma cells and antibodies, and it degrades organ health (Fig. 1). Organ health mediates plasma cell production, inferring a relationship between immune response and fitness of the individual. Antibodies bind to the attacking antigens, thereby neutralizing pathogens, triggering an attack by *phagocytic cells*, or activating *complement* proteins that either destroy pathogens directly or identify them for ingestion and digestion by phagocytes [3,4]. Each element of the state is subject to independent control, and new microbes may continue to enter (or be recognized by) the system [1,2], as represented by the following set of ordinary differential equations:

$$\dot{x}_1 = (a_{11} - a_{12}x_3)x_1 + b_1u_1 + w_1, \quad (1)$$

$$\dot{x}_2 = a_{21}(x_4)a_{22}x_1x_3 - a_{23}(x_2 - x_2^*) + b_2u_2 + w_2, \quad (2)$$

$$\dot{x}_3 = a_{31}x_2 - (a_{32} + a_{33}x_1)x_3 + b_3u_3 + w_3, \quad (3)$$

$$\dot{x}_4 = a_{41}x_1 - a_{42}x_4 + b_4u_4 + w_4. \quad (4)$$

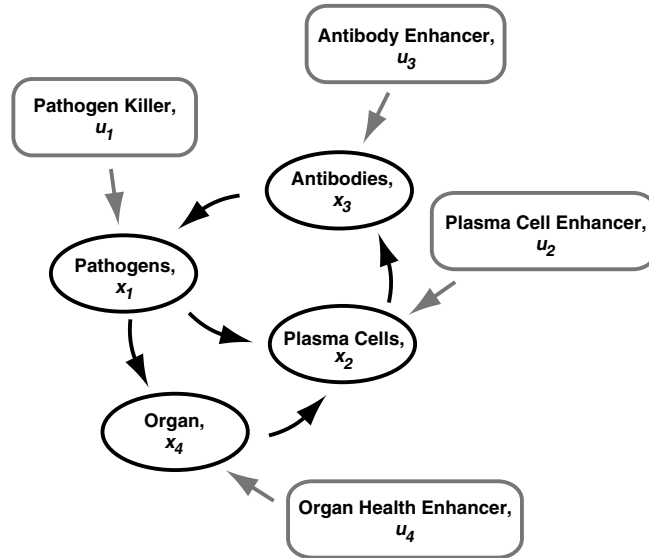


Fig. 1. Diagram of natural and enhanced immune dynamics in response to a pathogenic attack (from [1]).

The state components, x_i , vary with time, and the resting value of plasma cell concentration, x_2^* , is 2. The organ is healthy when $x_4 = 0$ and dead when $x_4 \geq 1$. Values of the parameters used for

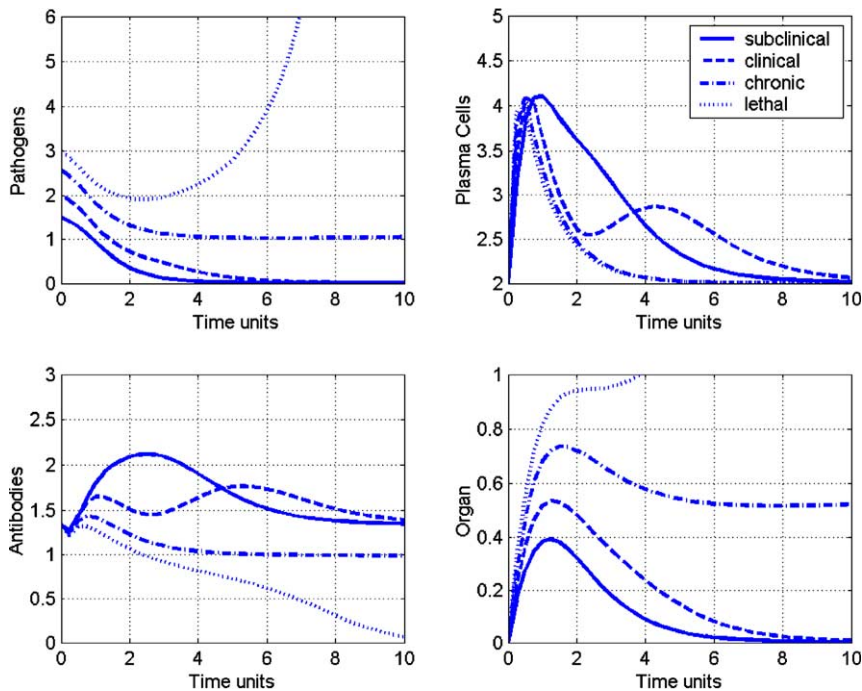


Fig. 2. Natural response to attack by a pathogen (from [1]).

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 $a_{33} = a_{41} = 0.5$. In this model, immune response is mediated by the damaged organ through $a_{21}(x_4)$, which regulates the production of plasma cells:

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

The parameters produce a system that recovers or succumbs to the pathogen without treatment, depending on initial conditions, during a period of ten time units. Parameters and time units are abstractions in this generic model. The state and control are always positive because concentrations cannot go below zero.

Fig. 2 shows typical uncontrolled response to increasing levels of pathogen concentration at the start of the time period. We may assume some initial period of microbial infection and growth prior to beginning the simulated immune response at zero time. The subclinical response would not require medical examination, while the clinical case could warrant medical consultation but is self-healing. The chronic case presents an unstable equilibrium with degraded organ health, while the lethal case results in death of the organ.

3. Deterministic optimal therapeutic protocol

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. This multi-objective, positive-definite scalar cost function of many variables allows tradeoffs between important factors to be adjusted through the relative weighting of individual components. Concordant responses tend to reinforce each other while conflicting responses compete in the development of an optimal regimen. The cost function is evaluated over the fixed time interval $[t_o, t_f]$ and at the end of the treatment interval,

$$J = \frac{1}{2}(s_{11}x_{1_f}^2 + s_{44}x_{4_f}^2) + \frac{1}{2} \int_{t_o}^{t_f} (q_{11}x_1^2 + q_{44}x_4^2 + r_{11}u_1^2 + r_{22}u_2^2 + r_{33}u_3^2 + r_{44}u_4^2) dt, \tag{6}$$

where the subscript f denotes quantities at the final time, t_f . The cost-function elements 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 (s_{ii} , q_{ii} , or r_{ii}) that establishes the relative importance of the factor in the treatment cost. These coefficients could reflect financial cost of treatment, or they could represent physiological ‘cost’ such as virulence, toxicity, or discomfort. The resulting treatment protocol balances speed, efficacy, and cost of treatment against implicit side effects. The optimization is ‘deterministic’ because all elements of the model and cost function are assumed to be known without error.

The immune response model (Eqs. (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 $(n \times 1)$ state, \mathbf{u} is the $(m \times 1)$ control, and \mathbf{w} is an $(s \times 1)$ process uncertainty vector. The scalar cost function (Eq. (6)) takes the general form,

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

$L[\mathbf{x}(t), \mathbf{u}(t), t]$ is called the *Lagrangian*, and the s_{ii} , q_{ii} , and r_{ii} of Eq. (6) are diagonal elements of the matrices \mathbf{S}_f , \mathbf{Q} , and \mathbf{R} . We adjoin the dynamic constraint to the Lagrangian in the *Hamiltonian* of the system through the $(n \times 1)$ *adjoint vector*, $\boldsymbol{\lambda}(t)$:

$$H[\mathbf{x}(t), \mathbf{u}(t), \mathbf{w}(t), \boldsymbol{\lambda}(t), t] = L[\mathbf{x}(t), \mathbf{u}(t), t] + \boldsymbol{\lambda}^T(t) \mathbf{f}[\mathbf{x}(t), \mathbf{u}(t), \mathbf{w}(t)]. \tag{9}$$

The necessary conditions for optimizing the cost function with respect to control are expressed by the three Euler–Lagrange equations [42]:

$$\dot{\boldsymbol{\lambda}}(t) = - \left\{ \frac{\partial H[\mathbf{x}(t), \mathbf{u}(t), \mathbf{w}(t), \boldsymbol{\lambda}(t), t]}{\partial \mathbf{x}} \right\}^T = - \left[\frac{\partial L(t)}{\partial \mathbf{x}} \right]^T - \mathbf{F}^T(t) \boldsymbol{\lambda}(t), \tag{10}$$

$$\boldsymbol{\lambda}(t_f) = \left\{ \frac{\partial \phi[\mathbf{x}(t_f)]}{\partial \mathbf{x}} \right\}^T = \mathbf{P}_f \mathbf{x}(t_f), \tag{11}$$

$$0 = \left[\frac{\partial H[\mathbf{x}(t), \mathbf{u}(t), \mathbf{w}(t), \boldsymbol{\lambda}(t), t]}{\partial \mathbf{u}} \right]^T = \left[\frac{\partial L(t)}{\partial \mathbf{u}} \right]^T + \mathbf{G}^T(t) \boldsymbol{\lambda}(t). \tag{12}$$

$\mathbf{F}(t)$ and $\mathbf{G}(t)$ are the time-varying *Jacobian matrices*, $\partial \mathbf{f} / \partial \mathbf{x}$ and $\partial \mathbf{f} / \partial \mathbf{u}$, evaluated along the deterministic optimal history. Eqs. (7) and (10)–(12) must be satisfied concurrently, specifying a *two-point boundary-value problem* that is solved numerically. A *steepest-descent method* is used to generate successive approximations that converge to the optimal control history $\mathbf{u}^*(t)$ [1].

Deterministic optimal solutions computed for otherwise-lethal initial conditions and unit cost function weights are presented in Fig. 3. The cost convergence of the numerical solution also is shown in the figure. The example shows that each of the therapeutic agents used separately can defeat the pathogen and maintain organ health with varying participation of the immune system. The rate of response, the final response error, the maximum value of degraded organ health, and the amount of drug used can be controlled by increasing or decreasing the values of \mathbf{S}_f , \mathbf{Q} , and \mathbf{R} accordingly.

If the initial concentration of pathogen is changed from its nominal value, the computed therapy is no longer optimal, and a new regimen must be defined to restore optimality. For a strong enough assault, the combination of immune response and nominal 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 reevaluation of the deterministic optimal policy or through a simpler mechanism for modifying the policy in proportion to deviations from the expected response history. In [2], 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 were not included in the nominally optimal protocol.

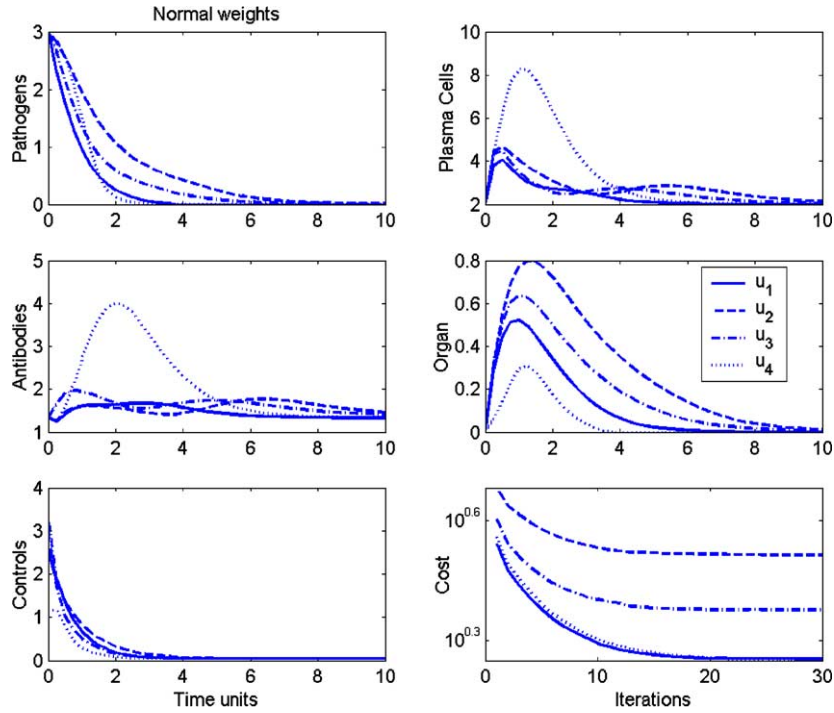


Fig. 3. Four optimal therapies against lethal attack with unit cost-function weights and each therapeutic agent, u_{1-4} applied separately (from [1,2]).

4. Deterministic neighboring-optimal therapeutic protocol

Adjustments to the optimal therapy that account for different levels of infection can be based on the solution of a neighboring-optimal control problem [42]. State and control histories are 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 $\Delta\mathbf{u}(t)$:

$$\mathbf{x}(t) = \mathbf{x}^*(t) + \Delta\mathbf{x}(t), \tag{13}$$

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

Eq. (7) can be expanded as,

$$\begin{aligned} \dot{\mathbf{x}}(t) &= \dot{\mathbf{x}}^*(t) + \Delta\dot{\mathbf{x}}(t) = \mathbf{f}\{\mathbf{x}^*(t) + \Delta\mathbf{x}(t), [\mathbf{u}^*(t) + \Delta\mathbf{u}(t)], [\mathbf{w}^*(t) + \Delta\mathbf{w}(t)]\} \\ &\approx \mathbf{f}[\mathbf{x}^*(t), \mathbf{u}^*(t)] + \mathbf{F}(t)\Delta\mathbf{x}(t) + \mathbf{G}(t)\Delta\mathbf{u}(t) + \mathbf{L}(t)\Delta\mathbf{w}(t) \end{aligned} \tag{15}$$

recognizing that the disturbance, $\mathbf{w}(t)$, also may have nominal and perturbation components. Because the deterministic optimal solution satisfies Eq. (7), the dynamics of small perturbations are closely approximated by the linear, time-varying equation

$$\Delta\dot{\mathbf{x}}(t) = \mathbf{F}(t)\Delta\mathbf{x}(t) + \mathbf{G}(t)\Delta\mathbf{u}(t) + \mathbf{L}(t)\Delta\mathbf{w}(t). \tag{16a}$$

For our simple model, each control element perturbs a single state element, and the control-effect matrix, $\mathbf{G}(t)$, is diagonal. With the process uncertainty acting directly on each element, $\mathbf{L}(t) = \mathbf{I}_4$, and this equation can be written explicitly as

$$\begin{aligned} \begin{bmatrix} \Delta \dot{x}_1 \\ \Delta \dot{x}_2 \\ \Delta \dot{x}_3 \\ \Delta \dot{x}_4 \end{bmatrix} &= \begin{bmatrix} (a_{11} - a_{12}x_3) & 0 & -a_{12}x_1 & 0 \\ a_{21}(x_4)a_{22}x_3 & a_{23} & a_{21}(x_4)a_{22}x_1 & \frac{\partial a_{21}}{\partial x_4} a_{22}x_1x_3 \\ -a_{33}x_3 & a_{31} & a_{31}x_1 & 0 \\ a_{41} & 0 & 0 & -a_{42} \end{bmatrix} \begin{bmatrix} \Delta x_1 \\ \Delta x_2 \\ \Delta x_3 \\ \Delta x_4 \end{bmatrix} \\ &+ \begin{bmatrix} b_1 & 0 & 0 & 0 \\ 0 & b_2 & 0 & 0 \\ 0 & 0 & b_3 & 0 \\ 0 & 0 & 0 & b_4 \end{bmatrix} \begin{bmatrix} \Delta u_1 \\ \Delta u_2 \\ \Delta u_3 \\ \Delta u_4 \end{bmatrix} + \begin{bmatrix} \Delta w_1 \\ \Delta w_2 \\ \Delta w_3 \\ \Delta w_4 \end{bmatrix}. \end{aligned} \quad (16b)$$

Optimal histories for this model of perturbed response are derived from the same conditions as before (Eqs. (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 (Eq. (8)),

$$\begin{aligned} \Delta^2 J &= \phi[\Delta \mathbf{x}(t_f)] + \int_{t_0}^{t_f} L[\Delta \mathbf{x}(t), \Delta \mathbf{u}(t)] dt \\ &= \frac{1}{2} \left\{ \Delta \mathbf{x}^T(t_f) \mathbf{S}_f \Delta \mathbf{x}(t_f) + \int_{t_0}^{t_f} [\Delta \mathbf{x}^T(t) \mathbf{Q} \Delta \mathbf{x}(t) + \Delta \mathbf{u}^T(t) \mathbf{R} \Delta \mathbf{u}(t)] dt \right\} \end{aligned} \quad (17)$$

and the Hamiltonian (Eq. (9)) is expressed as,

$$H[\Delta \mathbf{x}(t), \Delta \mathbf{u}(t), \Delta \boldsymbol{\lambda}(t)] = L[\Delta \mathbf{x}(t), \Delta \mathbf{u}(t)] + \Delta \boldsymbol{\lambda}^T(t) [\mathbf{F}(t) \Delta \mathbf{x}(t) + \mathbf{G}(t) \Delta \mathbf{u}(t)], \quad (18)$$

where $\Delta \boldsymbol{\lambda}(t)$ is the adjoint vector for the linearized system and the disturbance effect is neglected. The Euler–Lagrange equations (Eqs. (10)–(12)) can be applied to the variational system described by Eqs. (16)–(18). From Eq. (12), the control perturbation is

$$\Delta \mathbf{u}^*(t) = -\mathbf{R}^{-1} \mathbf{G}^T(t) \Delta \boldsymbol{\lambda}(t). \quad (19)$$

The terminal condition for the adjoint vector (Eq. (11)) is a function of the terminal state error and cost sensitivity,

$$\Delta \boldsymbol{\lambda}(t_f) = \mathbf{S}(t_f) \Delta \mathbf{x}(t_f). \quad (20)$$

As $\Delta \boldsymbol{\lambda}$ and $\Delta \mathbf{x}$ are adjoint, Eq. (20) applies over the entire interval; hence, the optimal adjustment of therapy is a *linear feedback control law* [42],

$$\Delta \mathbf{u}^*(t) = -\mathbf{R}^{-1} \mathbf{G}^T(t) \mathbf{S}(t) \Delta \mathbf{x}(t) = -\mathbf{C}(t) \Delta \mathbf{x}(t), \quad (21)$$

where $\mathbf{C}(t)$ is the time-varying *optimal gain matrix*. Simultaneous solution of the linear, ordinary differential equations for $\Delta \mathbf{x}(t)$ and $\Delta \boldsymbol{\lambda}(t)$ leads to a non-linear *matrix Riccati equation* for $\mathbf{S}(t)$:

$$\dot{\mathbf{S}}(t) = -\mathbf{F}^T(t) \mathbf{S}(t) - \mathbf{S}(t) \mathbf{F}(t) + \mathbf{S}(t) \mathbf{G}(t) \mathbf{R}^{-1} \mathbf{G}^T(t) \mathbf{S}(t) - \mathbf{Q}, \quad \mathbf{S}(t_f) = \mathbf{S}_f. \quad (22)$$

The dimensions of all matrices follow from the original problem specification. The feedback gain matrix, $\mathbf{C}(t)$, is calculated from Eqs. (21) and (22) just once using the values of $\mathbf{F}(t)$ and $\mathbf{G}(t)$ calculated along the deterministic optimal therapeutic history. From Eqs. (13) and (14), the optimal control policy accounts for previously unknown initial condition perturbations in the form

$$\mathbf{u}(t) = \mathbf{u}^*(t) - \mathbf{C}(t)\Delta\mathbf{x}(t) = \mathbf{u}^*(t) - \mathbf{C}(t)[\mathbf{x}(t) - \mathbf{x}^*(t)], \tag{23}$$

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

The effect of neighboring-optimal therapy is illustrated by a scalar control example from [2]. The cost-function weighting matrices, \mathbf{Q} and \mathbf{R} , are the same for both the nominal and neighboring optimizations. \mathbf{Q} is an identity matrix, the diagonal term of \mathbf{R} corresponding to the single control is one, and all other elements of \mathbf{R} are zero. The nominal therapy, $u_1^*(t)$, controls the pathogen and preserves organ health with the assumed initial pathogen concentration but not with a 50% higher microbial assault (Fig. 4). 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 to the nominally optimal control

$$u_1(t) = u_1^*(t) - c_{11}(t)\Delta x_1(t) - c_{12}(t)\Delta x_2(t) - c_{13}(t)\Delta x_3(t) - c_{14}(t)\Delta x_4(t) \tag{24}$$

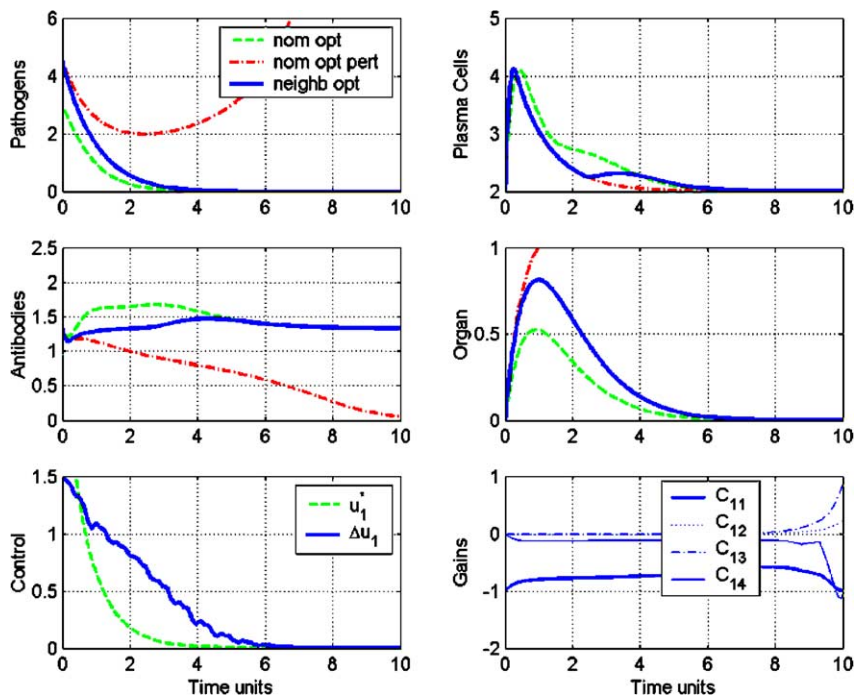


Fig. 4. Nominal and neighboring-optimal control with pathogen killer, u_1 and increased initial concentration of pathogen (from [2]).

produces a response that parallels the original one. The additional infusion of pathogen killer provides a stronger response to the pathogen and preserves antibody concentration. Maximum degradation of organ health is larger than for the nominally optimal case, but health eventually is restored as $x_4(t)$ goes to zero with increasing time. The feedback gains that provide this beneficial effect also are shown in the figure. The predominant effect is feedback of pathogen concentration to pathogen killer through c_{11} , and there is a small continuing feedback of the organ health indicator through c_{14} .

5. Controllability and observability of the immune response model

A principal goal is to reduce and eventually eliminate the pathogen concentration, preferably by the end of the treatment period. There are two means of doing this for the present model. The first is to deliver the state to a *basin of attraction* surrounding this model's stable equilibrium point. Once in the proper neighborhood, the state can decay to its desired value through natural immune response without further therapy, the speed of convergence being dependent on the strength of natural immune response. However, not all immune response models possess satisfactory stable equilibria. The second approach is to force the state to zero by therapeutic action with or without the existence of a stable equilibrium. If there is a control history that transfers a system from an arbitrary initial condition to $\mathbf{x} = \mathbf{0}$ in finite time regardless of its stability, the system is said to be *completely controllable* [42]. In the present case, we may be satisfied to know that the pathogen concentration alone could be driven to zero in finite time, satisfying a *scalar terminal constraint* on its value and letting the other elements remain free. A terminal constraint on x_1 can be satisfied if

$$\int_{t_0}^{t_f} \boldsymbol{\mu}^T(t) \mathbf{G}(t) \mathbf{G}^T(t) \boldsymbol{\mu}(t) dt \neq 0, \quad (25)$$

where $\boldsymbol{\mu}(t)$ is an adjoint vector that expresses sensitivity to the terminal constraint. The sensitivity is propagated back from the terminal condition by,

$$\dot{\boldsymbol{\mu}}(t) = -\mathbf{F}^T(t) \boldsymbol{\mu}(t), \quad \boldsymbol{\mu}(t_f) = [m \ 0 \ 0 \ 0]^T \quad (26)$$

and m is a non-zero constant [42]. The condition is satisfied for both the original model (Eq. (7)) and its local linearization (Eq. (16)).

To explore the effects of incomplete or imperfect measurements on the effectiveness of feedback therapy, we introduce the $(r \times 1)$ measurement, $\mathbf{z}(t)$, and its perturbation from some constant value, $\mathbf{z}^*(t)$, here assumed to be zero. Thus, $\mathbf{z}(t)$ equals $\Delta \mathbf{z}(t)$; the perturbation measurement is defined to be a linear combination of the $(n \times 1)$ perturbation state vector, $\Delta \mathbf{x}(t)$, plus the $(r \times 1)$ *perturbation measurement error vector*, $\Delta \mathbf{n}(t)$:

$$\Delta \mathbf{z}(t) = \mathbf{H} \Delta \mathbf{x}(t) + \Delta \mathbf{n}(t) = \mathbf{H}[\mathbf{x}(t) - \mathbf{x}^*(t)] + \Delta \mathbf{n}(t). \quad (27)$$

For this example, the $(r \times n)$ output matrix, \mathbf{H} , contains 1 where an element of the state is measured directly and 0 elsewhere. With a scalar measurement of $\Delta x_1 (r = 1)$, $\mathbf{H} = [1 \ 0 \ 0 \ 0]$, while for measurement of all state elements ($r = 4$), \mathbf{H} is the (4×4) identity matrix, \mathbf{I}_4 .

We require that the perturbation state be *completely observable* in the measurements, so that it can be used for effective feedback control. Paralleling the terminal-state-constraint condition (Eq. (25)), the requirement is satisfied if the *observability Grammian matrix*, \mathbf{O} ,

$$\mathbf{O} = \int_{t_o}^{t_f} \Phi^T(t, t_o) \mathbf{H}^T \mathbf{H} \Phi(t, t_o) dt \quad (28)$$

is non-singular [44]. This condition does not specify *how* the state estimate should be made; it merely tells us whether an estimate *can* be made. The quality of the estimate depends upon the measurement error and the design of the estimator.

If the system matrix, \mathbf{F} , were constant, the state transition matrix, $\Phi(t, t_o)$, that appears in \mathbf{O} would be computed as

$$\Phi(t_2, t_1) = e^{\mathbf{F}(t_2 - t_1)} = e^{\mathbf{F}\Delta t}. \quad (29)$$

However, \mathbf{F} is time-varying and the time increment varies during the integration, so the exact solution for $\Phi(t, t_o)$ is more complex [42]. Taking small time steps, δt , and using the state transition property,

$$\Phi(t_2, t_1) = \Phi(t_2, t_1 + \delta t) \Phi(t_1 + \delta t, t_1) \quad (30)$$

the time-varying state transition matrix required in Eq. (28) can be generated incrementally using Eq. (29) with slowly changing, stepwise-constant \mathbf{F} .

The observability Grammian has been calculated along the deterministic optimal history shown in Fig. 3 for various combinations of measurements. If all elements or any three elements of the state are measured, \mathbf{O} is non-singular, and the system is completely observable. Conversely, no single scalar measurement is satisfactory. For this example, the system is completely observable when four of the six possible pairs are measured: (x_1, x_2) , (x_1, x_3) , (x_1, x_4) , (x_3, x_4) .

6. Stochastic neighboring-optimal therapeutic protocol

With unknown initial conditions or process uncertainty and measurements that contain error, the optimization problem is no longer deterministic, and a cost function that assumes perfect characterization of the system cannot be optimized with certainty. Rather than minimizing the deterministic cost function, J , subject to a deterministic dynamic constraint, we minimize the *expected value* of the cost, $E(J)$, subject to a stochastic constraint. While the expected value has a rigorous definition [42], it suffices to view $E(J)$ as the most likely value of the original cost function given probabilistic descriptions of the uncertainties in $\mathbf{x}(t_o)$, $\mathbf{w}(t)$, and $\Delta \mathbf{n}(t)$.

The small-perturbation assumption allows us to separate the deterministic and stochastic solutions; hence the expected value of the cost can be expressed as

$$E(J) \approx J^* + E(\Delta^2 J). \quad (31)$$

The deterministic optimal solution establishes a baseline that is unaffected by uncertainty, and it is sufficient for us to examine stochastic optimality of the perturbation system alone. With the further assumption that all random processes are satisfactorily described by their means and covariances, the goal devolves to optimizing a *linear-quadratic-Gaussian (LQG) problem*: the system is

linear, the control cost function is quadratic, and the uncertainty is Gaussian. The optimal LQG solution minimizes $E\Delta(J^2)$.

The optimizing solution for the LQG problem consists of a feedback control law that generates the control perturbation, $\Delta\mathbf{u}(t)$, based on an estimate of the state perturbation, $\Delta\hat{\mathbf{x}}(t)$, rather than the state perturbation itself [42]:

$$\Delta\mathbf{u}(t) = -\mathbf{C}(t)\Delta\hat{\mathbf{x}}(t). \quad (32)$$

Furthermore, the optimizing control law and estimator can be designed separately: the control gain matrix, $\mathbf{C}(t)$, is unchanged (Eqs. (21) and (22)), while the optimal estimator is a *Kalman–Bucy filter*.

As derived in [43] and described in [42], the optimal filter takes the form

$$\Delta\dot{\hat{\mathbf{x}}}(t) = \mathbf{F}(t)\Delta\hat{\mathbf{x}}(t) + \mathbf{G}(t)\Delta\mathbf{u}(t) + \mathbf{K}(t)[\Delta\mathbf{z}(t) - \mathbf{H}\Delta\hat{\mathbf{x}}(t)], \Delta\hat{\mathbf{x}}(t_o) = \Delta\bar{\mathbf{x}}(t_o), \quad (33)$$

where $\Delta\bar{\mathbf{x}}(t_o)$ is the expected value of the initial perturbation. The $(n \times r)$ *linear-optimal filter gain matrix*, $\mathbf{K}(t)$, is

$$\mathbf{K}(t) = \mathbf{P}(t)\mathbf{H}^T\mathbf{N}^{-1}(t). \quad (34)$$

$\mathbf{P}(t)$ is the solution to the $(n \times n)$ *matrix Riccati equation for estimation* (the dual of Eq. (22)):

$$\dot{\mathbf{P}}(t) = \mathbf{F}(t)\mathbf{P}(t) + \mathbf{P}(t)\mathbf{F}^T(t) + \mathbf{L}(t)\mathbf{W}(t)\mathbf{L}^T(t) - \mathbf{P}(t)\mathbf{H}^T\mathbf{N}^{-1}\mathbf{H}\mathbf{P}(t), \mathbf{P}(t_o) = \mathbf{P}_o. \quad (35)$$

\mathbf{P}_o represents the initial uncertainty in the state estimate, $\mathbf{W}(t)$ is the $(s \times s)$ spectral density matrix of process uncertainty, and $\mathbf{N}(t)$ is the $(r \times r)$ spectral density matrix of the measurement uncertainty. We note that $\Delta\hat{\mathbf{x}}(t)$ is an estimate of the mean value of $\Delta\mathbf{x}(t)$, while $\mathbf{P}(t)$ is the covariance of the estimate error. Together, they define the estimated probability density function of $\Delta\mathbf{x}(t)$. Because the model used to define the optimal estimator is linear, the estimator gain matrix, $\mathbf{K}(t)$, depends on the relative magnitudes of \mathbf{W} and \mathbf{N} , not their absolute values [42]. The estimator is applied to the output of a non-linear system; hence, the state estimate is approximately optimal for small disturbances and measurement errors but suboptimal for larger values.

The stochastic neighboring-optimal control policy (Eq. (32)) consists of the deterministic neighboring-optimal control policy (Eq. (21)) operating on the output of the Kalman–Bucy filter (Eqs. (33)–(35)). The result is added to the deterministic nominal control history, $\mathbf{u}^*(t)$, to derive the stochastic neighboring-optimal therapeutic protocol:

$$\mathbf{u}(t) = \mathbf{u}^*(t) - \mathbf{C}(t)\Delta\hat{\mathbf{x}}(t). \quad (36)$$

7. Results of stochastic neighboring-optimal therapy

There are many combinations of measurements, control variables, cost function weights, and estimation statistics that could be considered. For illustration, we present five examples that display effects of measurement dimension and measurement error. In all cases, we assume that the control variable is the pathogen killer, u_1 ; hence, the result that assumes perfect measurement of the entire state (Fig. 4) presents a baseline against which the effects of estimation can be compared. While the estimators are designed under the assumption that \mathbf{W} and \mathbf{N} take their indicated




Fig. 5. Therapeutic effect with full-state measurement and linear-optimal state estimation based upon low signal-to-noise ratio ($\mathbf{W} = \mathbf{I}_4$, and $\mathbf{N} = 20\mathbf{I}_4$).

values, the corresponding disturbances and measurement error are not simulated in the first three figures, allowing us to inspect the dynamic effect of state estimation. Furthermore, we assume that $\mathbf{P}_o = \mathbf{W}$ and $\Delta\mathbf{x}_o = \mathbf{0}$ in calculating perturbation-state estimates.

The importance of good quality measurements is portrayed by Figs. 5 and 6, which illustrate the effects of large ratios of $|\mathbf{W}|$ and $|\mathbf{N}|$. In Fig. 5, $\mathbf{W} = \mathbf{I}_4$, and $\mathbf{N} = 20\mathbf{I}_4$, signifying a low signal-to-noise ratio for all 4 measurements. As a consequence, the filter gains are low, and the state estimation is sluggish. The feed-back control is too small and too late to prevent the pathogen concentration from diverging. If the measurement error is reduced to $\mathbf{N} = \mathbf{I}_4$ (not shown), the filter gains are higher, and the feedback control is able to arrest pathogen growth; however, organ health is degraded to an unacceptable level, and the therapy is not successful. With $\mathbf{W} = \mathbf{I}_4$, and $\mathbf{N} = \mathbf{I}_4/20$, the feedback effect is strong enough to overcome the pathogen while preventing death of the organ (Fig. 6), illustrating successful therapy. The critical difference is that the estimates of Δx_1 and Δx_4 are fast enough to command a higher level of drug delivery early in the response, attacking the pathogen directly and preserving the generation of antibodies.

Successful therapy also can be achieved when just two measurements are available. Fig. 7 presents a case in which Δx_1 and Δx_3 are measured, with $\mathbf{W} = \mathbf{I}_4$, and $\mathbf{N} = \mathbf{I}_2/20$. Thus, there are no direct measurements of plasma cell concentration or organ health, yet their values are estimated by the filter. The filter gains are higher than in the previous case. Therapeutic results are comparable to those achieved with full-state measurement (Fig. 6), as the feedback control history is very nearly the same. The example confirms the complete observability provided by the measurement pair and the ability to provide effective therapy with a reduced measurement set.

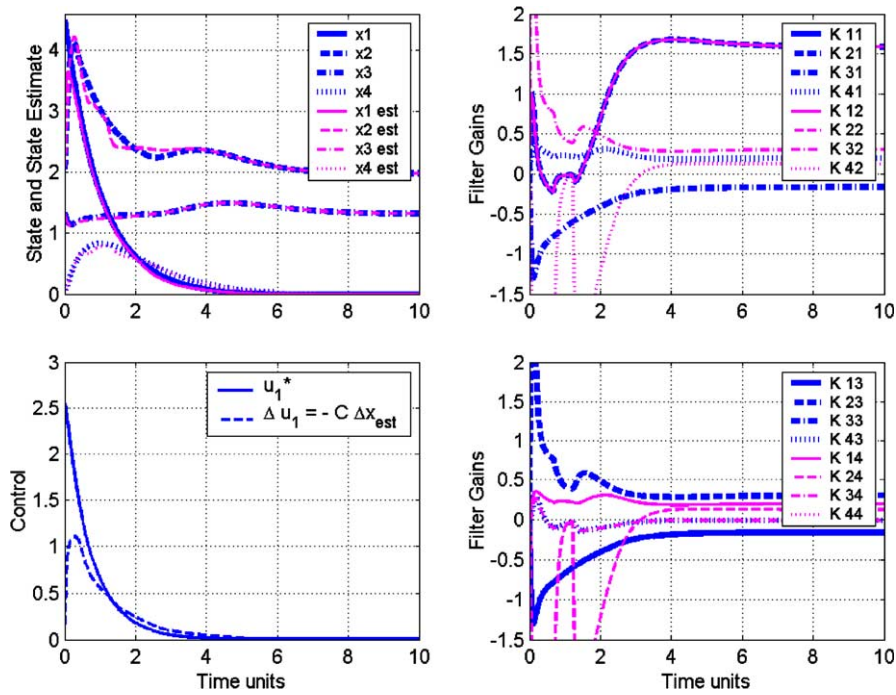


Fig. 6. Therapeutic effect with full-state measurement and linear-optimal state estimation based upon high signal-to-noise ratio ($\mathbf{W} = \mathbf{I}_4$, and $\mathbf{N} = \mathbf{I}_4/20$).

Simulating white-noise random disturbance and measurement error illustrates the effects of stochastic-neighboring-optimal therapy on immune response in uncertain settings. Once again, the initial pathogen concentration is 50% higher than the nominal value, so feedback therapy is essential to non-lethal response. In both cases, there is a continuing, random influx of pathogen (Δw_1), as well as uncertainty in the production rates of plasma cells, antibody, and net effects on organ health ($\Delta w_2, \Delta w_3, \Delta w_4$); measurements also are subject to error. Ten simulations are conducted; the means and standard deviations as well as the individual traces are shown.

With noisy full-state measurement ($\mathbf{W} = \mathbf{I}_4$, and $\mathbf{N} = \mathbf{I}_4/20$), there is considerable variation in response (Fig. 8). However, the mean response is very similar to that obtained with no simulation of variability (Fig. 6), and the pathogen is readily defeated. One cause for concern is that the peak degradation of organ health may be greater than it is in the average case, though death of the organ does not occur in any of the 10 cases. Of course, there is equal likelihood that actual response will be better, not worse.

Measuring just Δx_1 and Δx_3 ($\mathbf{W} = \mathbf{I}_4$, and $\mathbf{N} = \mathbf{I}_2/20$), results are similar, but state-estimate variability is greater, and the organ dies in one of the 10 cases (Fig. 9). As shown in [1], the peak decline in organ health could be reduced by increasing the weight on organ health variations in the nominal control cost function (Eq. (6)). Variations from the nominal response could be reduced by increasing the weight on organ health perturbations in the perturbation control cost function (Eq. (17)). Either of these approaches would increase prescribed drug delivery, eliminat-

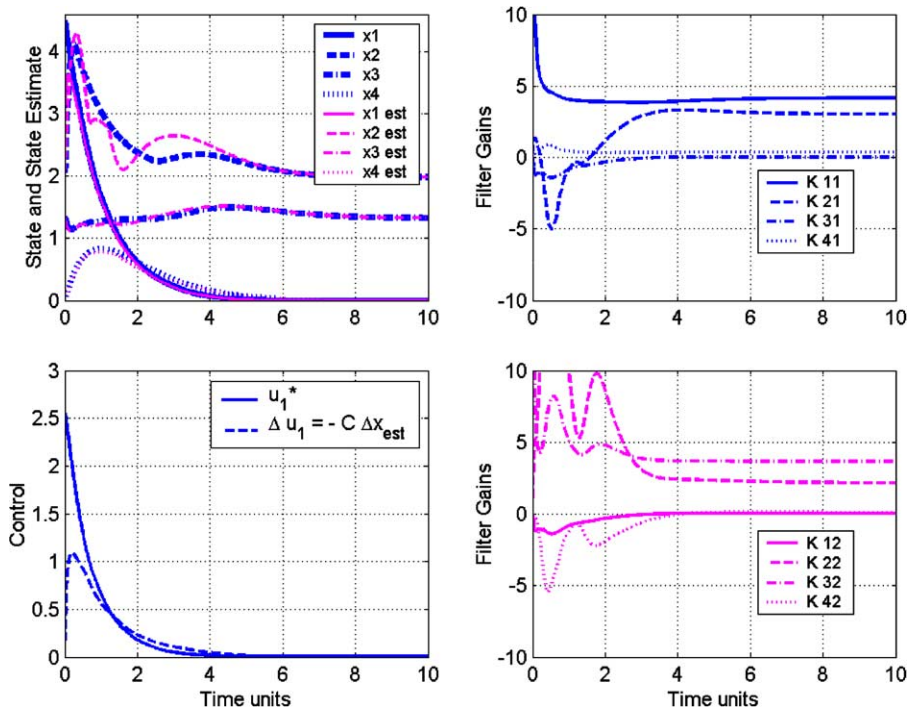


Fig. 7. Therapeutic effect with two measurements (Δx_1 and Δx_3) and linear-optimal state estimation based upon high signal-to-noise ratio ($\mathbf{W} = \mathbf{I}_4$, and $\mathbf{N} = \mathbf{I}_2/20$).

ing the unsatisfactory excursion shown in Fig. 9. Thus, there are effective design mechanisms for reducing the likelihood of unsatisfactory response due to measurement uncertainty or incompleteness.

It is not unusual for drugs to have adverse side effects, limiting the aggressiveness with which an infection can be treated. As a preliminary step toward understanding what impact side effects may have on optimal therapy, the model is modified so that the pathogen killer also degrades organ health. Eq. (4) becomes

$$\dot{x}_4 = a_{41}x_1 - a_{42}x_4 + b_4u_4 + b_5u_1 + w_4, \tag{37}$$

where b_5 is 0.1, 10% of the pathogen killer’s efficacy. With no change in the cost function and measuring just Δx_1 and Δx_3 for feedback control, the peak of the mean value of organ health exceeds 1, indicating organ death. Increasing the cost-function weighting on organ health (q_{44}) by a factor of 10 provides a satisfactory solution, as shown in Fig. 10. The organ health peak is equivalent to that shown in Fig. 9, indicating a period of severe illness but not death before the pathogen is defeated. The cost function increase causes the nominal control history to be modified. The initial drug dose is increased, and its rate of decay is also higher; this pattern allows the antibodies to grow more quickly and to higher value – thus helping to defeat the pathogen – while not degrading organ health.




Fig. 8. Therapeutic effect with simulated disturbance and measurement error, full-state measurement, and linear-optimal state estimation based upon high signal-to-noise ratio ($\mathbf{W} = \mathbf{I}_4$, and $\mathbf{N} = \mathbf{I}_4/20$).

8. Conclusion

Optimal control theory has an important role to play in the development of new therapeutic protocols for treating infection. Prior studies demonstrated that numerical optimization of non-linear models combined with neighboring-optimal feedback of perfect information can define therapies for enhancing natural response of the immune system. Here, we show that similar results can be achieved when only imperfect or incomplete measurements are available. A principal requirement is that the perturbation state be completely observable from the measurements. For the system studied, observability is provided by measuring any three components of the 4-dimensional state, as well as by four of six possible pairs of measurements. Thus, the analysis reveals the minimum number of measurements required for feedback therapy. A linear-optimal state estimator then provides the information required for feedback control. Nevertheless, the quality of measurements is important. It is shown that the signal-to-noise ratio must be high enough to provide timely state estimates, else the pathogen concentration may diverge and cause organ health to fail. Variable response caused by uncertainty in immune response and measurement error can cause transient excursions that may be harmful to the patient. Fortunately, stoch-

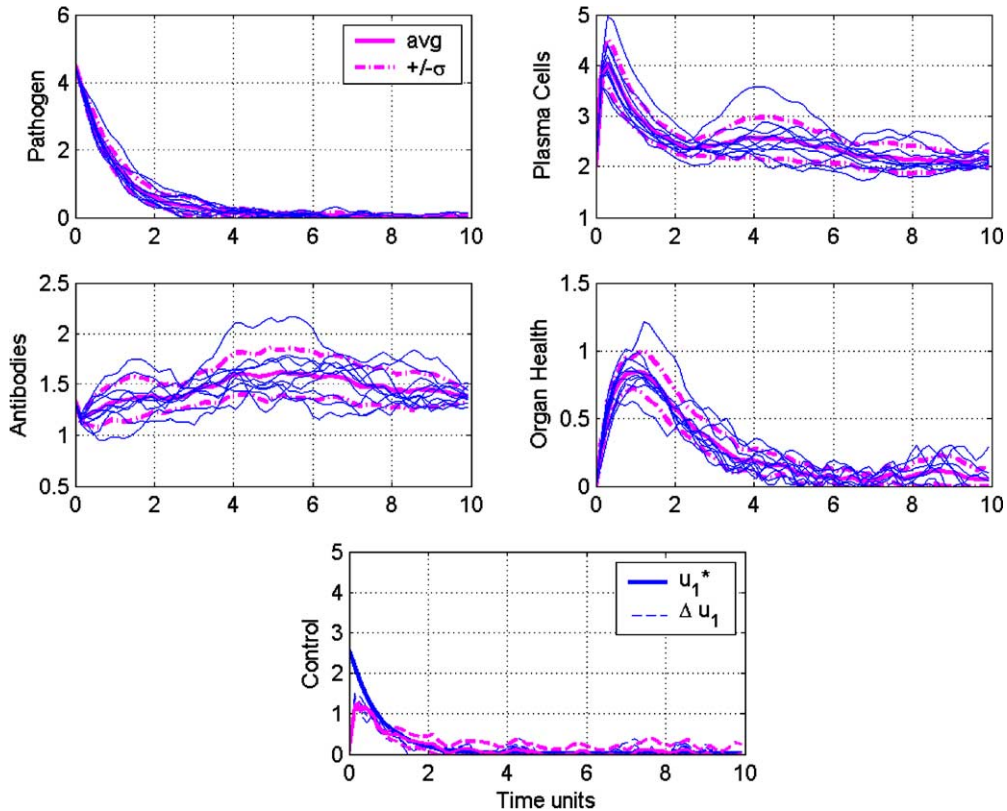


Fig. 9. Therapeutic effect with simulated disturbance and measurement error, two measurements (Δx_1 and Δx_3), and linear-optimal state estimation based upon high signal-to-noise ratio ($\mathbf{W} = \mathbf{I}_4$, and $\mathbf{N} = \mathbf{I}_2/20$).

astic neighboring-optimal control theory provides strong mechanisms for designing therapies to overcome these complications, enhancing immune response to infection while tolerating uncertainty in infectious loading, measured response, and immune system dynamic models.

The biomedical implications of these results are heavily weighted toward developing a sound foundation for mathematically inspired clinical treatment rather than the creation of new models for simulating natural processes. The model of disease dynamics was taken from the literature, and we are not aware that the model has been compared quantitatively to empirical evidence. In implementation, accurate biodynamic models are an imperative, but model development is not the topic of this paper. Nevertheless, our results illustrate a considerable degree of robustness to uncertainties in initial condition, varying infection rate, and measurement error. To the extent that the disturbance input can be viewed as a surrogate for ‘process noise’ (i.e., error in the dynamic model), our results demonstrate the ability to respond to considerable variation in dynamic state histories, even when only two of the four state elements are measured and those measurements are noisy.

Assessing organ health is a difficult but important issue, as the optimal therapy takes organ health into account. Given a model of dynamic coupling between organ health, pathogenic assault, and immune system response, the Kalman–Bucy filter provides a way of estimating the




Fig. 10. Therapeutic effect with simulated disturbance and measurement error, two measurements (Δx_1 and Δx_3), and linear-optimal state estimation based upon high signal-to-noise ratio ($\mathbf{W} = \mathbf{I}_4$, and $\mathbf{N} = \mathbf{I}_2/20$). Pathogen killer has 10% adverse side effect on organ health. Cost function weighting on organ health (q_{44}) increased to 10.

organ health metric without direct measurements. Three of the nine possible measurement sets that we considered do not use an organ health measurement, but all provide an estimate of organ health. Quantifying organ health itself should be the subject of future research, as it is clearly implicated in the development of the optimal policy. Even if not directly observable, it may be possible to deduce organ health by indirect measurements, such as urinalysis, blood test, body temperature, skin discoloration, patient reports of pain and feeling, and a host of developing non-invasive techniques. This additional information could readily be added to the optimal treatment protocol creating a useful synergy: mathematical modeling, state estimation, and optimal control can form a strong basis for creating a quantitative foundation for better understanding organ health.

Acknowledgment

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

References

- [1] R.F. Stengel, R. Ghigliazza, N. Kulkarni, O. Laplace, Optimal control of innate immune response, *Optimal Contr. Appl. Methods* 23 (2002) 91.
- [2] R.F. Stengel, R. Ghigliazza, N. Kulkarni, Optimal enhancement of immune response, *Bioinformatics* 18 (2002) 1227.
- [3] C.A. Janeway, P. Travers, M. Walport, M. Shlomchik, *Immunobiology*, Garland, London, 2001.
- [4] P.M. Lydyard, A. Whelan, M.W. Fanger, *Instant Notes in Immunology*, Springer, New York, 2000.
- [5] M. Thain, M. Hickman, *The Penguin Dictionary of Biology*, Penguin Books, London, 2000.
- [6] A. Asachenkov, G. Marchuk, R. Mohler, S. Zuev, *Disease Dynamics*, Birkhauser, Boston, 1994.
- [7] A. Rundell, H. HogenEsch, R. DeCarlo, Enhanced modeling of the immune system to incorporate natural killer cells and memory, in: *Proc. Amer. Control Conf.*, Seattle, 1995, p. 255.
- [8] A.S. Perelson, *Immunology for physicists*, *Rev. Modern Phys.* 69 (1997) 1219.
- [9] M.A. Nowak, R.M. May, *Virus Dynamics: Mathematical Principles of Immunology and Virology*, Oxford University, Oxford, 2000.
- [10] M.A. Nowak, R.M. May, R.E. Phillips, S. Rowland-Jones, D.G. Lalloo, S. McAdam, P. Klenerman, B. Köppe, K. Sigmund, C.R.M. Bangham, A.J. McMichael, Antigenic oscillations and shifting immunodominance in HIV-1 infections, *Nature* 375 (1995) 606.
- [11] A.S. Perelson, A.V. Neumann, M. Markowitz, J.M. Leonard, D.D. Ho, HIV-1 dynamics in vivo: virion clearance rate, infected cell lifespan, and viral generation time, *Science* 271 (1995) 1582.
- [12] A.S. Perelson, P.W. Nelson, Mathematical analysis of HIV-1 dynamics in vivo, *SIAM Rev.* 41 (1999) 3.
- [13] D. Wodarz, R.M. May, M.A. Nowak, The role of antigen-independent persistence of memory cytotoxic T lymphocytes, *Int. Immun.* 12 (2000) 467.
- [14] M.A. Stafford, Y. Cao, D.D. Ho, L. Corey, A.S. Perelson, Modeling plasma virus concentration and CD4+ T cell kinetics during primary HIV infection, *J. Theor. Biol.* 203 (2000) 285.
- [15] D. Wodarz, M.A. Nowak, CD8 memory, immunodominance, and antigenic escape, *Eur. J. Immun.* 30 (2000) 2704.
- [16] N. Wiener, *Cybernetics: or Control and Communication in the Animal and the Machine*, Technology, Cambridge, 1948.
- [17] R.E. Bellman, *Mathematical Methods in Medicine*, World Scientific Press, Singapore, 1983.
- [18] G.W. Swan, Optimal control applications in biomedical engineering – a survey, *Opt. Contr. Appl. Methods* 2 (1981) 311.
- [19] G.W. Swan, *Applications of Optimal Control Theory in Medicine*, Marcel Dekker, New York, 1984.
- [20] G.W. Swan, Role of optimal control theory in cancer therapy, *Math. Biosci.* 101 (1990) 237.
- [21] A.S. Perelson, Applications of optimal control theory to immunology, in: R.R. Mohler, A. Ruberti (Eds.), *Recent Developments in Variable Structure Systems Economics and Biology*, Springer, Berlin, 1978, p. 272.
- [22] A.S. Perelson, M. Mirmirani, G.F. Oster, Optimal strategies in immunology, I: B-cell differentiation and proliferation, *J. Math. Biol.* 3 (1976) 325.
- [23] A.S. Perelson, M. Mirmirani, G.F. Oster, Optimal strategies in immunology, II: B memory cell production, *J. Math. Biol.* 5 (1978) 213.
- [24] A.S. Perelson, B. Goldstein, S. Rocklin, Optimal strategies in immunology, III: the IgM-IgG switch, *J. Math. Biol.* 10 (1980) 209.
- [25] D. Kirschner, S. Lenhart, S. Serbin, Optimal control of the chemotherapy of HIV, *J. Math. Biol.* 35 (1997) 775.
- [26] L.M. Wein, S.A. Zenios, M.A. Nowak, Dynamic multidrug therapies for HIV: a control theoretic approach, *J. Theor. Biol.* 185 (1997) 15.
- [27] S. Bonhoeffer, R.M. May, G.M. Shaw, M.A. Nowak, Virus dynamics and drug therapy, *Proc. Natl. Acad. Sci. USA* 94 (1997) 6971.
- [28] L.M. Wein, R.M. D'Amato, A.S. Perelson, Mathematical analysis of antiretroviral therapy aimed at HIV-1 eradication or maintenance of low viral loads, *J. Theor. Biol.* 192 (1998) 81.
- [29] D. Wodarz, M.A. Nowak, Specific therapy regimes could lead to long-term immunological control of HIV, *Proc. Natl. Acad. Sci. USA* 96 (1999) 14464.

- [30] D. Wodarz, K.M. Page, R.A. Arnaout, A.R. Thomsen, J.D. Lifson, M.A. Nowak, A new theory of cytotoxic T-lymphocyte memory: implications for HIV treatment, *Philos. Trans. Roy. Soc. B* 355 (2000) 329.
- [31] J.M. van Rossum, O. Steyger, T. van Uem, G.J. Binkhorst, R.A.A. Maes, Pharmacokinetics by using mathematical systems dynamics, in: J. Eisenfeld, M. Witten (Eds.), *Modelling of Biomedical Systems*, Elsevier Science Publishers, 1986, p. 121.
- [32] D.C. Robinson, Principles of pharmacokinetics, in: R. Maronde (Ed.), *Topics in Clinical Pharmacology and Therapeutics*, Springer, New York, 1986, p. 1.
- [33] D.J. Bell, F. Katusiime, A time-optimal drug displacement problem, *Optimal Contr. Appl. Methods* 1 (1980) 217.
- [34] E.R. Carson, D.G. Cramp, L. Finkelstein, D. Ingram, *Control System Concepts and Approaches in Clinical Medicine, Computers and Control in Clinical Medicine*, Plenum, New York, 1985, p. 1.
- [35] H.J. Chizeck, P.G. Katona, Closed-loop control, in: E.R. Carson, D.G. Cramp (Eds.), *Computers and Control in Clinical Medicine*, Plenum Press, New York, 1985, p. 95.
- [36] A. Schumitsky, Stochastic control of pharmacokinetic systems, in: R.F. Maronde (Ed.), *Topics in Clinical Pharmacology and Therapeutics*, Springer, New York, 1986, p. 13.
- [37] R.W. Jelliffe, Clinical applications of pharmacokinetics and control theory: planning, monitoring, and adjusting regimens of aminoglycosides lidocaine digitoxin and digoxin, in: R.F. Maronde (Ed.), *Topics in Clinical Pharmacology and Therapeutics*, Springer, New York, 1986, p. 26.
- [38] M.M. Polycarpou, J.Y. Conway, Modeling and control of drug delivery systems using adaptive neural control methods, in: *Proc. Amer. Control Conf.*, Seattle, 1996, p. 781.
- [39] G.K. Kwong, K.E. Kwok, B.A. Finegan, S.L. Shah, Clinical evaluation of long range adaptive control for mean arterial blood pressure regulation, in: *Proc. Amer. Control Conf.*, Seattle, 1996, p. 786.
- [40] A. Gentilini, M. Morari, C. Bieniok, R. Wymann, T. Schnider, Closed-loop control of analgesia in humans, in: *Proc. IEEE Conf. Decision and Control*, Orlando, 2001, p. 861.
- [41] R.S. Parker, F.J. Doyle, III, J.E. Harting, N.A. Peppas, Model predictive control for infusion pump insulin delivery, in: *Proc. 18th Ann. Conf. IEEE Engineering in Medicine and Biology Soc.*, Amsterdam, 1996, p. 1822.
- [42] R.F. Stengel, *Optimal Control and Estimation*, Dover Publications, New York, 1994.
- [43] R.E. Kalman, R.S. Bucy, New results in linear filtering and prediction, *ASME Trans. J. Basic Eng.* 82D (1961) 95.
- [44] T. Kailath, *Linear Systems*, Prentice-Hall, Englewood Cliffs, 1980.