

USING NONLINEAR MODEL PREDICTIVE CONTROL TO FIND OPTIMAL THERAPEUTIC STRATEGIES TO MODULATE INFLAMMATION

JUDY DAY

Mathematical Biosciences Institute, The Ohio State University
1735 Neil Ave, 377 Jennings Hall
Columbus, OH 43210, USA

and

Current address: Department of Mathematics and
Department of Electrical Engineering & Computer Science
The University of Tennessee
104 Aconda Court, 1534 Cumberland Avenue
Knoxville, TN 37996-0614, USA

JONATHAN RUBIN

Department of Mathematics, University of Pittsburgh
301 Thackeray Hall
Pittsburgh, PA 15260, USA

GILLES CLERMONT

Department of Critical Care Medicine
University of Pittsburgh Medical Center
3550 Terrace St
Pittsburgh, PA 15261, USA

(Communicated by Mette Olufsen)

ABSTRACT. Modulation of the inflammatory response has become a key focal point in the treatment of critically ill patients. Much of the computational work in this emerging field has been carried out with the goal of unraveling the primary drivers, interconnections, and dynamics of systemic inflammation. To translate these theoretical efforts into clinical approaches, the proper biological targets and specific manipulations must be identified. In this work, we pursue this goal by implementing a nonlinear model predictive control (NMPC) algorithm in the context of a reduced computational model of the acute inflammatory response to severe infection. In our simulations, NMPC successfully identifies patient-specific therapeutic strategies, based on simulated observations of clinically accessible inflammatory mediators, which outperform standardized therapies, even when the latter are derived using a general optimization routine. These results imply that a combination of computational modeling and NMPC may be of practical use in suggesting novel immuno-modulatory strategies for the treatment of intensive care patients.

2000 *Mathematics Subject Classification.* Primary: 92C50, 93C15; Secondary: 93C83.

Key words and phrases. Inflammation, nonlinear model predictive control, immuno-modulation, dosing control.

1. Introduction. Although the inflammatory response is crucial to restoring health following a wide range of biological stresses, uncontrolled systemic inflammation is also the primary cause of organ failure and death in victims of severe trauma, infections, and many other conditions leading to admission to an intensive care unit. Therefore, control of the inflammatory response has become a key focal point in the treatment of critically ill patients. Much of the theoretical work regarding severe inflammation has aimed to elucidate the mechanisms underlying systemic inflammation and to explain how the mediators of inflammation interact with one another across their respective time scales [14, 5, 21]. Significant insight has been acquired from these approaches, including an enhanced understanding of the ways in which manipulations of inflammatory components may combat persistent inflammation. A main result coming from this research confirms that successful therapeutic interventions require appropriate timing relative to the evolution of the inflammatory response [20, 9].

An early clinical approach to controlling inflammation was to target an individual inflammatory mediator [13]. However, there is no one mediator that stands as the source for persistent inflammation [2, 3, 17]. Instead, inflammation involves a cascade of processes, which is initiated by a few key factors but persists as a result of a complicated feedback process involving effectors that are produced later than the instigator. In addition, anti-inflammatory mediators may be present at elevated levels during prolonged inflammation, but their effects on pro-inflammatory mediators may be small or negligible due to the relative amounts of inflammation present in the system [4]. Inversely, immunoparalysis, an overwhelmingly under responsive state resulting from a relative anti-inflammatory excess, may play a large role in predisposing patients to secondary infections that compromise organ recovery [22]. Because the inflammatory response is a complex process involving multiple positive and negative feedback loops, it is extremely difficult to predict the response of the various mediators to perturbations (i.e. to therapy) applied to one or more of the system's components.

As a result of this complexity, there is still much to be done to identify appropriate therapeutic targets to combat excessive and pervasive inflammation and to develop strategies for delivering appropriate interventions in the correct amounts and at the right times [10, 6, 7]. One of the tools that can help optimize complex dose regimens is nonlinear model predictive control (NMPC). NMPC algorithms have mainly been developed and applied for industrial operations involving system processes that can be well described with mathematical models, usually systems of ordinary differential equations. NMPC is advantageous relative to other control algorithms because it combines predictions of the real system state at a future time, based on a mathematical model, with measurements derived from the system to calculate a control move that will help to optimize the desired outcome for a specific process variable. Recently, NMPC has been applied to a variety of biomedical processes including the regulation of glucose supply in diabetic patients and an exploration of optimal dosing of anticancer agents, among others [19, 12, 11].

Motivated by the recent success of NMPC in biomedical settings, we aimed to explore the utility of NMPC for the derivation of optimal therapeutic interventions for the control of inflammation triggered by a pathogenic infection, simulating the ominous clinical problem of severe sepsis [15]. There are two essential components of an NMPC scheme: the process to be controlled and the model predicting the process (sometimes called the predictive model). In the clinical setting of a severe

infection, the process to be controlled would be the time course of the inflammatory response in a patient, while insuring eradication of the infection. To explore the feasibility of applying NMPC, we chose to emulate the patient's inflammatory response with the model developed in [20], such that the current exposition is completely simulation based. This model is referred to as the patient model or virtual patient. Initially, we used a predictive model and a patient model with identical equations, parameter values, and initial conditions, which we refer to as a no mismatch scenario. On each time step, the NMPC algorithm was applied to the predictive model to generate an intervention strategy, and this was directly applied to the patient model. Subsequently, we introduced mismatch between the predictive model and the patient model (i.e. patient-model mismatch) in some parameter values and initial conditions, to achieve a more realistic representation of a clinical scenario.

Although the model developed in [20] is more abstract than those used to simulate insulin or cancer chemotherapeutic agents in past biological applications of NMPC, its biological relevance is supported by previous analysis showing that it reproduces several important observations related to severe systemic inflammation in biological organisms [20, 9]. Our implementation of NMPC in a reduced ODE model for inflammation will pave the road for future applications involving more quantitatively detailed models, such as that presented in [8].

This paper is organized as follows. In Section 2, we give a brief overview of the NMPC framework. In Section 3, we present the equation-based model of inflammation that we use, along with additional methodological details, including the way in which we generate our simulated patient population (i.e. virtual patients) and the specifics of the therapeutic strategies that we implement. The results of our simulations are presented in Section 4, while we conclude with a discussion in Section 5.

2. NMPC overview. The NMPC algorithm must incorporate certain essential elements, which we now briefly describe [18].

I. *The Specification of a Reference Trajectory*

The reference trajectory defines a target path that we would like our process outputs to follow. Specifically, our model includes a variable representing tissue damage/dysfunction that we would like to cause to decrease to zero as quickly as possible. So, our reference trajectory for damage is the constant function, $R_D(t) \equiv 0$.

II. *The Prediction of Process Output*

The reference trajectory will be compared to a prediction of process behavior over a specified duration of time, h , called the prediction horizon. Simulation of the patient model (equations (1)-(4) in Section 3.1) describing the acute inflammatory response to pathogen yields this prediction.

III. *The Definition of an Objective Function*

The objective function quantifies the difference between the reference trajectory and the predicted process output. The objective function we use has the typical weighted sum of squares form (see Section 3.2, equation (5)).

IV. The Computation of a Sequence of Control Moves

Using the predictive model to simulate the system's response to changes in input over time h , we seek a sequence of control moves that will minimize the objective function. The control moves are simply changes in the level of an input into the model, calculated in such a way as to achieve the desired goal. For example, in our model, control moves consist of steps in control terms in equations (3) and (4), which represent pro- and anti-inflammatory therapy, respectively.

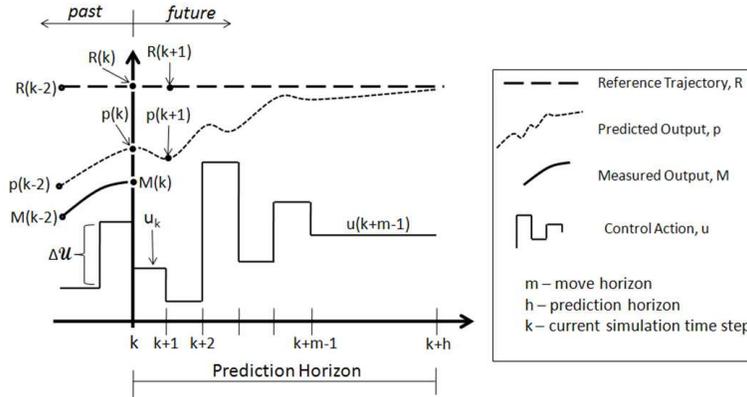


FIGURE 1. MPC Schematic (adapted from Ogunnaike et al. [18], 1994, pg 997).

The number of control moves, m , that will be made during time h can also be specified. If $m < h$, which is typically the case, then the control terms are held constant after the m^{th} control move when the system's response to the input over time h is determined. Only the first control move is actually implemented in the patient and predictive models as the dose for the current time step, after which the algorithm repeats, eventually determining the dose for the next time step. The MPC schematic in Figure 1, adapted from Ogunnaike et al., offers an excellent summary of this process [18].

V. Error Prediction Update

An important element of the algorithm lies in the error prediction step. After a control action is implemented in both the predictive and patient models, a measurement $M(k)$ is taken from the patient model and compared to the corresponding quantity $p(k)$ computed from the predictive model, where k is the present time step in the algorithm. In a standard NMPC implementation, $\|M(k) - p(k)\|_2^2$ is minimized as a part of the objective function. In our NMPC scheme, when a mismatch exists, updating is done differently. This is due to the fact that, when the system is only incompletely observed, the variables that can be realistically measured are not necessarily those that we specifically wish to minimize, and hence are not appropriate for inclusion in the objective function. The strategies we implement for handling this situation are discussed in Section 3.2.

3. Designing an NMPC algorithm for the therapeutic control of inflammation. Having described the general components of NMPC schemes, we now present the acute inflammation model we use and the specific details of our customization of an NMPC algorithm. The algorithm we use is a modified version of that developed by Florian et al. [12] (acquired through personal communication). All simulations were conducted using MatLab[®] on a distributed computing platform. The ODE system is numerically integrated using MatLab's Simulink default solver routine, *ode45*. The optimization routine used the *fmincon* built-in MatLab[®] function.

3.1. The model for the acute inflammatory response to pathogenic infection. The model of acute inflammation was previously described in [20] and consists of the ODE system:

$$\frac{dP}{dt} = k_{pg}P \left(1 - \frac{P}{P_\infty}\right) - \frac{k_{pm}s_m P}{\mu_m + k_{mp}P} - k_{pn}f(N^*)P \quad (1)$$

$$\frac{dN^*}{dt} = \frac{s_{nr}R}{\mu_{nr} + R} - \mu_n N^* + PIDOSE(t) \quad (2)$$

$$\frac{dD}{dt} = k_{dn} \frac{f(N^*)^6}{x_{dn}^6 + f(N^*)^6} - \mu_d D \quad (3)$$

$$\frac{dC_A}{dt} = s_c + k_{cn} \frac{f(N^* + k_{cnd}D)}{1 + f(N^* + k_{cnd}D)} - \mu_c C_A + AIDOSE(t), \quad (4)$$

where R and f in (1) - (4) are given by

$$R = f(k_{np}P + k_{nn}N^* + k_{nd}D) \quad \text{and} \quad f(x) = \frac{x}{1 + \left(\frac{C_A}{c_\infty}\right)^2}.$$

In the model, equation (1) represents the evolution of the bacterial pathogen population (P) that instigates the cascade of inflammation. Equation (2) governs the dynamics of the concentration of a collection of early pro-inflammatory mediators such as activated phagocytes and the pro-inflammatory cytokines they produce (N^*). Equation (3) corresponds to a marker of tissue damage/dysfunction (D), which helps to verify response outcomes. Finally, equation (4) describes the evolution of the concentration of a collection of anti-inflammatory mediators (C_A) that inhibit many of the interactions within the system. Table 1 gives the parameter values used in [20], which we refer to as the reference parameter set.

Equations (1) - (4), with AIDOSE=PIDOSE= 0, admit three stable critical points (outcomes) under certain choices of parameter values [20]:

1. Healthy: $(P, N^*, D, C_A) = (0, 0, 0, \hat{C}_A)$, for a small value $\hat{C}_A > 0$.
2. Aseptic: $(P, N^*, D, C_A) = (0, \hat{N}^*, \hat{D}, \hat{C}_A)$ for $\hat{N}^*, \hat{D}, \hat{C}_A > 0$.
3. Septic: $(P, N^*, D, C_A) = (P, N^*, D, C_A)_S$, a point with all components positive.

TABLE 1. Model parameter reference values for the system (1) - (4).

Parameter	Reference Value	Parameter	Reference Value
k_{pm}	0.6/ M units hr	μ_n	0.05/hr
k_{mp}	0.01/ P units hr	k_{nd}	0.02/ D units hr
s_m	0.005 M units/hr	k_{dn}	0.35 units of D /hr
μ_m	0.002/hr	x_{dn}	0.06 N^* units
k_{pg}	Various in range: (0.021-2.44)/hr	μ_d	0.02/hr
p_∞	20x10 ⁶ /cc	c_∞	0.28 C_A units
k_{pn}	1.8/ N^* units hr	s_c	0.0125 C_A units/hr
k_{np}	0.1/ P units hr	k_{cn}	0.04 C_A units/hr
k_{nn}	0.01/ N^* units hr	k_{cnd}	48 N^* units/ D units
s_{nr}	0.08 N_R units/hr	μ_c	0.1/hr
μ_{nr}	0.12/hr		

We label our finite time simulation outcomes based on these three states, with simulations that end with negligible P classified as healthy or aseptic depending on which state (N^* , D , C_A) are approaching, as discussed further in subsection 3.5. Figures 2 (a) and (b) show typical aseptic and septic scenarios, respectively. It is assumed that basic therapy, including the administration of antibiotics, resuscitation with fluids, and so forth, are implicitly modeled in system (1) - (4). This means that the various outcomes mentioned above can occur despite administration of basic treatment.

Input to the NMPC algorithm consists of an anti-inflammatory therapy, present as a source term (+AIDOSE) in equation (4), and a pro-inflammatory therapy, incorporated as a source term (+PIDOSE) in equation (2). Constraints are defined that prevent dosing from going negative, meaning that therapy can be infused into the system but not extracted.

In all of the simulations that we discuss, the total simulation time is 168 hours (1 week). In addition, k is an hourly step, so doses are adjusted on an hourly basis. The goal of the NMPC control algorithm is to identify (virtual) patient-specific therapy dosing profiles that can correct inflammatory responses that, without intervention, would result in either aseptic or septic scenarios.

3.2. The objective function, constraints, and error prediction under mismatch. The objective function J that we use contains terms to minimize damage levels (D), pathogen levels (P), and total therapy AIDOSE and PIDOSE given over the prediction horizon h and takes the form

$$J = \min_{\substack{PIDOSE(t) \\ AIDOSE(t)}} \|\Gamma_D D\|_2^2 + \|\Gamma_P P\|_2^2 + \|\Gamma_{AI} AIDOSE(t)\|_2^2 + \|\Gamma_{PI} PIDOSE(t)\|_2^2 \quad (5)$$

Minimization is done over piecewise constant time courses of AIDOSE and PIDOSE, achieved by a sequence of control moves, as discussed in Section 2. The

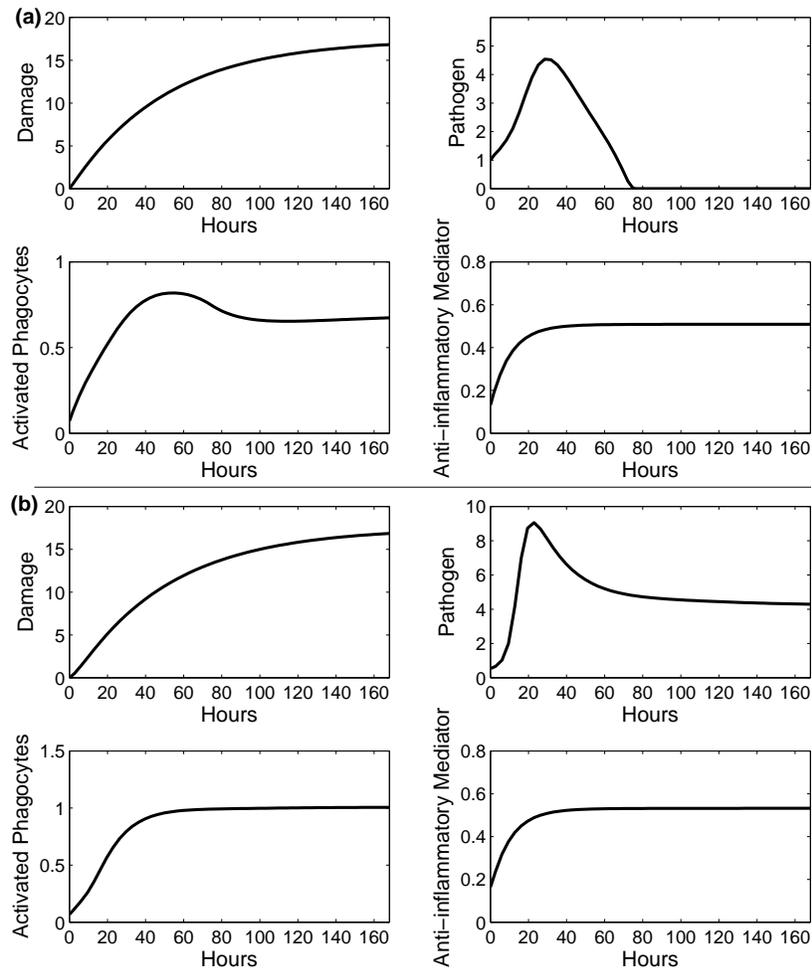


FIGURE 2. Representative placebo simulations of two virtual patients. (a) Simulation of a typical aseptic outcome in the absence of treatment. (b) Simulation of a typical septic outcome in the absence of treatment.

Γ -parameters are the weighting constants and the zero function is used as the reference trajectory for both tissue damage/dysfunction (D) and pathogen (P). From computational experimentation, it became apparent that striving to minimize both damage and pathogen was essential. This, however, introduced the difficult challenge of maintaining a balance between these two objectives. An emphasis on minimizing damage might lead to unrestricted pathogen growth. On the other hand, an emphasis on minimizing pathogen might lead to an overzealous immune response determined to eliminate pathogen as soon as possible whatever the costs, after which

it might be too late to control the excessive inflammation. For all simulations, we chose $\Gamma_D = \Gamma_P = 1$ so as not to emphasize one factor over the other. In addition, we also penalized the actual dose amounts, AIDOSE and PIDOSE, over the prediction horizon. The rationale for penalizing dosing is that drugs have financial costs as well as the potential for harmful side effects. The values for Γ_{AI} and Γ_{PI} are 1 if the corresponding form of therapy is used and 0 otherwise.

When patient-model mismatch is introduced, we need to specify which variables from the patient model will be measured to update the predictive model. We propose that it is fairly realistic to acquire accurate hourly measurements for N^* and C_A and hence these are designated as the variables to be measured. We assume that it is not feasible, however, to measure the variables D and P since damage is impossible to quantify in real time in clinical settings, and it is unlikely that a measurement of the pathogen population could be made at all, much less at every hour. Thus, our NMPC implementation is non-standard, in that the states measured from the process (virtual patient) and the variables appearing in the objective function are disjoint sets. To harness the measurements of N^* and C_A taken from the patient model at the end of each time step, we use these values as our initial conditions for N^* and C_A in the predictive model at the start of the next one-hour time step, with AIDOSE and PIDOSE set to the values prescribed by the NMPC algorithm and with P and D in the predictive model evolving continuously across time steps.

This updating scheme does not address possible discrepancies between the predictive model and patient model with respect to levels of damage and pathogen, both of which are the primary forces capable of pushing the system toward an unhealthy endpoint. Hence, when pathogen levels in the predictive model and in the patient model are vastly different due to differences between virtual patient and predictive model parameter values, an additional updating strategy not based on direct measurements is necessary. We added an update that kicks in when either

- (A) pathogen levels are low in the predictive model but sufficiently high in the patient model, OR
- (B) pathogen levels are high in the predictive model but sufficiently low in the patient model.

Every four hours, pathogen levels in the predictive model and patient model are compared. If (A) occurs, then the predictive model's pathogen level for the next time step is reset to $P = 0.5$, which is a relatively high value. If (B) occurs instead, then the predictive model's pathogen level is reset to zero. These modifications can be looked at as a re-initialization of the pathogen value in the predictive model. Use of the virtual patient's pathogen values in this way is clinically justifiable, reflecting the fact that in a clinical setting, a physician can potentially identify persistent infection based on high fever, positive blood cultures, toxic granulations in neutrophils, or rising procalcitonin levels, while in other cases, a physician may be able to judge, based on the absence/resolution of symptoms, that it is highly unlikely that a patient has a high pathogen load. This rationale, together with the fact that we are not directly setting the predictive model's pathogen state to the exact value present in the virtual patient, implies that this updating strategy is a reasonable and clinically relevant way to alert the algorithm of significant discrepancies between the predictive model and the patient model, corresponding to either case (A) or (B) above.

Constraints on dosing levels must also be specified. This might be construed as an ad hoc regularization procedure, where one wishes to reduce the likelihood of large changes in therapy. The maximum dose amount of anti-inflammatory therapy allowed at a given step is calculated as the difference between the current level of C_A (the anti-inflammatory mediator) and a maximum allowable level of C_A , given by $C_{AMax} := 0.6264$ [20]. When pro-inflammatory therapy is used, the maximum dose amount allowed at a given step is calculated as the difference between the current level of N^* and $N_{Max}^* := 0.5$. This maximum was selected as a value of N^* that is high enough to consistently have an impact on pathogen but not so high that anti-inflammatory feedback cannot rein it in.

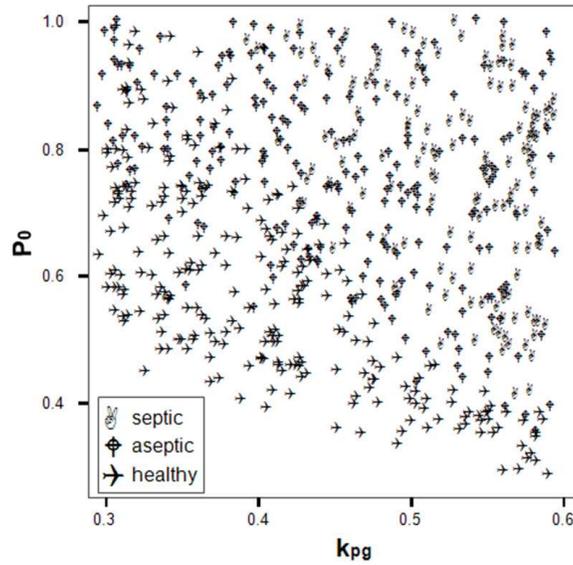


FIGURE 3. Scatterplot showing the distribution of patient outcomes with respect to pathogen growth rate (k_{pg}) and the initial level of pathogen (P_0). For the given ranges of k_{pg} and P_0 the various outcomes (healthy, aseptic, and septic) of the patient profiles in the placebo case are well mixed in the k_{pg} - P_0 plane. Thus, k_{pg} and P_0 are not the primary drivers of patient outcome in the absence of treatment.

3.3. Building a virtual patient population. Several of the reference parameters in equations (1) - (4) and the initial conditions for P and C_A , which we denote P_0 and C_{A0} , respectively, were allowed to vary to reflect virtual patient variability. Virtual patients were generated with individualized parameter profiles as follows. First, we selected a reference parameter set (Table 1) from [20]. Values of five parameters, as well as C_{A0} , were selected independently for each patient, from bounded uniform distributions on ranges defined as the reference values $\pm 25\%$ (Table 2). For the parameter k_{pg} , an experimentally determined range is available: $0.021 - 2.44 \text{ hr}^{-1}$ (Table 1). We restricted the range of k_{pg} values to $0.3 - 0.6 \text{ hr}^{-1}$ and selected P_0 between 0 and 1.0 units (Table 2) such that placebo outcomes in the virtual patient population were not driven primarily by pathogen-related parameters

TABLE 2. Model parameters in which variability was assumed in the patient-model mismatch case. Patient parameters were generated by choosing a random value from a uniform distribution on each given range

Parameter	Patient Parameter Ranges	Description
P_0	0.0 – 1.0	Initial condition of pathogen (P)
C_{A0}	0.0938 – 0.1563	Initial condition of the anti-inflammatory mediator (C_A)
k_{pg}	0.3 – 0.6	Growth rate of pathogen (P)
k_{cn}	0.03 – 0.05	Maximum production of anti-inflammatory mediator (C_A)
k_{nd}	0.015 – 0.025	Activation of phagocytes (N^*) by tissue damage (D)
k_{np}	0.075 – 0.125 (Co-varies w/ k_{nd})	Activation of phagocytes (N^*) by pathogen (P)
k_{cnd}	36.0 – 60.0 (Co-varies w/ k_{cn})	Controls the effectiveness of activated phagocytes (N^*) versus damage (D) in the production of the anti-inflammatory mediator (C_A)
k_{nn}	0.0075 – 0.0125 (Co-varies w/ k_{nd})	Activation of phagocytes (N^*) by already activated phagocytes (N^*) (or the cytokines that they produce)

k_{pg} and P_0 but by the parameter profile as a whole (Figure 3). In some preliminary simulations, in the absence of patient-model mismatch, the virtual patient-specific parameter profile was also used in the underlying predictive model. When patient-model mismatch was introduced, the reference parameter set (Table 1) was used for the predictive model. In addition, some parameters were chosen to co-vary since many molecules are produced by the same biological machinery or actuate their effects through similar pathways. That is, if parameters p_1 and p_2 co-vary and the value chosen for p_1 is $+n\%$ of its reference value, then the generated value for p_2 should also be close to $+n\%$ of its reference value. For example, we specified that the variability in k_{cnd} , the rate of production of the anti-inflammatory mediator (C_A) by damaged tissue (D), was to vary by the same percentage as k_{cn} , the rate of production of C_A by activated phagocytes (N^*), so that the rates at which these sources were producing C_A were balanced.

3.4. Alternative Therapies. A significant advantage of *in silico* simulations is the ability to apply different therapies to the same virtual patient and compare the outcomes. For each virtual patient receiving treatment, we use our NMPC algorithm to generate a dynamic therapeutic profile specific to that virtual patient’s particular evolution as determined by the virtual patient’s unique parameter profile. This is referred to as *targeted therapy*. To gauge how well the targeted therapy achieves our objectives, we compare its performance to the results from the administration of three alternative therapies.

The simplest alternative therapy is *Placebo Therapy*, where no treatment is given. *Static Therapy* is designed to represent the therapy currently given to critically ill patients with severe inflammatory disorders such as, sepsis in the intensive care unit: a consistent dosing regimen of an anti-inflammatory therapy. In practice, we implemented this therapy by creating a dosing profile that gives a small dose (0.005) of the anti-inflammatory therapy (via instantaneous injections) each hour over a period of 72 hours, after which therapy terminates. The control algorithm does not

play a role in static therapy. *Standard Therapy* is generated by applying the control algorithm to the model (1) - (4), with the fixed set of reference parameter values and the initial conditions $(P_0, N_0^*, D_0, C_{A0}) = (0.5, 0, 0, 0.125)$, as specified in [20], to obtain one set of dosing profiles for pro- and anti-inflammatory therapy. This one set of dosing profiles is then administered to all virtual patients. Thus, in contrast to targeted therapy, standard therapy is not patient specific. Since the predictive model's k_{pg} value is a parameter that we are free to tune to get the best possible therapeutic outcomes, we repeated this process for each virtual patient for each of three values for k_{pg} (growth rate of the pathogen) in the underlying predictive model: 0.52, 0.6, and 0.8. Targeted therapy is also determined for each of these k_{pg} values when we introduce mismatch between predictive model and patient model parameter values.

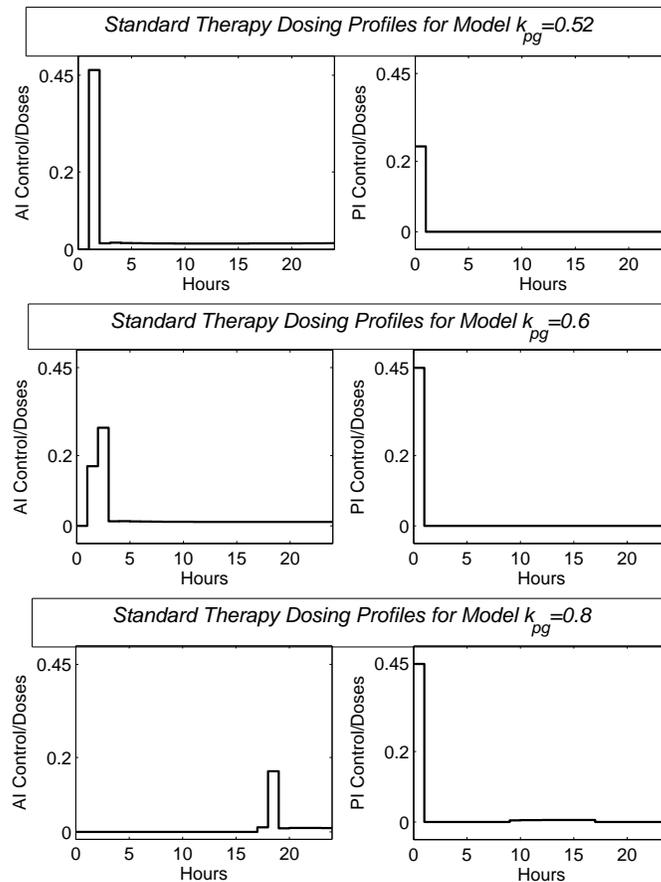


FIGURE 4. Three different standard therapy dosing profiles, with pro- and anti-inflammatory dosing schedules. Standard therapy was calculated using initial conditions $(P_0, N_0^*, D_0, C_{A0}) = (0.5, 0, 0, 0.125)$ and a model k_{pg} value of 0.52 (top panel), 0.6 (middle panel), or 0.8 (bottom panel). The same pair of doses, both pro- and anti-inflammatory, was administered to all patients receiving treatment.

We initially chose $k_{pg} = 0.52$ in the predictive model, since this value is above the bifurcation point where healthy, aseptic, and septic states co-exist ($k_{pg} = 0.514$) and lies within the range of virtual patient k_{pg} values. A later choice of $k_{pg} = 0.8$ in the predictive model, although above the upper bound of the range for k_{pg} patient values we used, allowed for a more aggressive therapeutic approach. Although results with this value were favorable for targeted therapy, the outcomes for standard therapy suffered. Thus, we subsequently explored the k_{pg} value of 0.6 in the predictive model, which is the upper bound of the patient k_{pg} range. When we display our results for standard therapy and targeted therapy (in the mismatch case), we show outcomes for all three k_{pg} values. Figure 4 displays the three sets of dosing profiles that were calculated for the standard therapies explored, using initial conditions $(P_0, N_0^*, D_0, C_{A0}) = (0.5, 0, 0, 0.125)$.

3.5. Additional practical issues. In our simulations, we based the intervention time, or time of onset of therapy, on the level of N^* , which denotes the early pro-inflammatory signals. During the course of an infection, if a virtual patient's N^* level rises above a certain threshold, the virtual patient is considered to show clinical manifestations warranting treatment. This implies a biomarker driven approach to initiating therapeutic intervention. This method, however, is not without its complications. For instance, how should the N^* threshold be chosen? In the current exploration, an N^* threshold of 0.05 was selected, based on the finding that N^* levels generally do not exceed this threshold in simulations yielding healthy outcomes. Of the 1000 virtual patients we generated and simulated, 620 exhibited values of N^* that exceeded this threshold and thus received treatment.

As we worked through the process of customizing the algorithm, one of the main goals was to make the setup and assumptions more realistic. For instance, we noticed that the prescribed amount of anti-inflammatory therapy given to virtual patients would sometimes cause the levels of the anti-inflammatory mediator to stay elevated for very long periods of time. This problem usually happened in scenarios when inflammation was very high after the eradication of pathogen. Consequently, to bring inflammation down, the anti-inflammatory therapy would be given continually, in an attempt to essentially saturate the system with as much C_A as possible for as long as possible. This situation would be avoided clinically for fear of secondary infections, and hence we put a mechanism into place so that if the level of C_A remained consistently elevated for more than 48 hours, the maximum allowable amount of C_A was reduced by half. This mechanism compares the C_A measurement of the current step with that of the previous step and if the difference is close (within 0.001), then this implies that the level of C_A has remained elevated from one time step to the next. A counter is then incremented to keep track of how many consecutive times this occurs.

It was also necessary to choose an algorithm for determining the outcome (i.e. healthy, septic, or aseptic) of an individual simulation once an entire therapy dosing profile was administered to a virtual patient. Outcome was determined from the values of the variables at the end of the simulation time of 168 hours; however, when the outcome could not be determined at 168 hours because the system had not yet reached a steady state, we took the values of the variables at 168 hours and integrated the system with $AIDOSE = PIDOSE = 0$ for an additional 300 (simulated) hours, by which time the solution invariably settled to a steady state. From such simulations, we systematically determined virtual patient outcomes and tallied the results. A virtual patient outcome was labeled septic if pathogen levels were above

TABLE 3. Summary for NPMC Simulation Setup

Therapies used	No Mismatch (Case 1): Anti-inflammatory only No Mismatch (Case 2) & Mismatch Case: Anti- and Pro-inflammatory
Total number of virtual patients generated	1000
Treatment Intervention Time	After N^* levels in the patient are greater than or equal to 0.05
Number of virtual patients whose N^* levels surpassed the threshold for receiving treatment	620
Control parameter values	$m=2$ and $h=12$
Γ -weighting constants in objective function given by Equation (5)	All weights equal to 1
Maximum allowable levels of C_A : C_{AMax}	0.6264
Maximum allowable levels of N^* : N^*_{Max}	0.5
Maximum duration for elevated C_A levels	48 hours
Patient initial conditions	$(P_0, N^*, D, C_A) = (P_0\text{-random}, 0, 0, C_{A0}\text{-random})$ (See Table 2 for ranges for $P_0\text{-random}$ and $C_{A0}\text{-random}$)
<u>The following are applicable in the <i>mismatch</i> simulations only:</u>	
Model initial conditions	$(P_0, N^*, D, C_A) = (0.5, 0, 0, 0.125)$
Pathogen update check	Every 4 hours
k_{pg} value in underlying model	Various explored: 0.52, 0.6, or 0.8

a threshold of 1.0 and damage and activated phagocytes were also above their designated thresholds of 1.0 and 0.05, respectively. If pathogen levels were not above threshold, yet damage and activated phagocyte levels were, then the virtual patient outcome was labeled aseptic, in accordance with the definitions of these physiologic states mentioned earlier. Otherwise, a virtual patient was labeled healthy.

Finally, after experimenting with a number of values for the prediction horizon, h , and the move horizon, m , we chose $h = 12$ hours, since this was long enough to capture essential model dynamics, and $m = 2$, which provided moderately aggressive dosing. Table 3 summarizes the setup details for the virtual patient simulations with and without the presence of patient-model mismatch.

3.6. Predictors of the response to a control intervention. It is of interest to identify initial conditions and model parameters predictive of the response to a targeted strategy in virtual patients that would otherwise evolve to either a septic or aseptic outcome in the absence of therapy. We therefore constructed a variety of standard statistical classifiers, to identify such predictors. We considered as potential predictors all initial conditions and model parameters that varied among virtual

TABLE 4. No Mismatch Case 1 and Case 2 Simulation Results and Comparison to Alternative Therapies. Standard therapy results are given over three different k_{pg} values for the underlying model: 0.52, 0.6, and 0.8. The number of patients out of 620 is given in parenthesis (unless otherwise specified). In the first three rows, this number is followed by a number that represents either the quantity of patients coming from an unhealthy placebo category into the healthy category due to the effects of treatment or the number coming from the healthy placebo category into either of the unhealthy categories.

Therapy Type:	Placebo	Static	Standard Therapies:			No Mismatch Targeted Anti-inflammatory	No Mismatch Targeted Anti- & Pro-inflammatory
			$k_{pg}=0.52$	$k_{pg}=0.6$	$k_{pg}=0.8$		
Percentage Healthy:	40% (251)	46% (282; 36)	60% (370; 119)	67% (417; 191)	52% (322; 159)	55% (341; 99)	95% (588; 337)
Percentage Aseptic:	37% (228)	32% (195; 1)	19% (118; 0)	22% (138; 25)	48% (298; 88)	21% (133; 4)	5% (32; 0)
Percentage Septic:	23% (141)	23% (143; 4)	21% (132; 0)	11% (65; 0)	0% (0; 0)	24% (146; 5)	0% (0; 0)
Percentage Harmed (out of 251):	n/a	2% (5/251)	0% (0/251)	10% (25/251)	35% (88/251)	4% (9/251)	0% (0/251)
Percentage Rescued (out of 369):	n/a	10% (36/369)	32% (119/369)	52% (191/369)	43% (159/369)	27% (99/369)	91% (337/369)

patients (see Section 3.3 and Table 2). The classification variable was restoration of health by targeted therapy (Yes/No). A logistic regression (LR) classifier performed consistently better than other classifiers in its ability to identify responders to targeted therapy (as judged by receiver operating characteristic (ROC) performance), and we therefore report the LR results.

4. **Results.** Outcomes of the 620 patients qualifying for therapy are presented in Tables 4 and 5. In the placebo case, 40% (251) fell into the healthy category and 60% (369) fell into the unhealthy categories, with 37% (228) and 23% (141) having aseptic and septic outcomes, respectively. We also provide the percentages of patients harmed and rescued under the various treatments. Harmed patients are those that were in the healthy state under placebo but were instead redirected to an unhealthy state under a non-placebo treatment. On the other hand, rescued patients are those that were in an unhealthy state under placebo but were redirected to the healthy state under a non-placebo treatment. Since there were 251 healthy

patients and 369 unhealthy patients under placebo, we use these totals to determine percentages of patients harmed and rescued.

4.1. Static and standard therapies (Table 4). The overall outcomes of static therapy were marginally better than placebo, with 46% of patients achieving healthy outcomes. This weak improvement resembles results attained in current clinical practice, unfortunately [16]. Even though static therapy rescued 10% (36/369) of unhealthy patients, it also harmed 2% (5/251) of otherwise healthy patients.

The administration of standard therapy yielded better overall outcomes than the placebo case or static therapy for each k_{pg} considered, even though it was not individually tailored (Table 4 and Section 3.4). Standard therapy rescued 32% (119/369), 52% (191/369), and 43% (159/369) of unhealthy patients, for k_{pg} values of 0.52, 0.6, and 0.8, respectively. However, within standard therapies, only a k_{pg} value of 0.52 did not harm any patients, whereas values of 0.6 and 0.8 respectively harmed 10% (25/251) and 35% (88/251) of otherwise healthy patients, resulting in aseptic outcomes instead.

4.2. No mismatch (Table 4). We next present results from the use of targeted therapy, first without patient-model mismatch and second in the more realistic case when mismatch is present. In the absence of mismatch, the predictive model will forecast patient dynamics perfectly, which is useful for algorithmic development and for determining the optimal performance of our control algorithm. Since the patient state and the predictive model state are the same at every step in this case, no updating of predictive model states is needed. In the first set of simulations under the no mismatch setup, we experimented with anti-inflammatory therapy alone (Table 4, Case 1). The results obtained demonstrate the need for an additional therapeutic option (Table 4, Case 2) to correctly modulate the immune response, even in the absence of patient-model mismatch.

4.2.1. No Mismatch Case 1: Anti-inflammatory therapy only. In this initial no mismatch case, motivated by the clinical practice of targeting a single inflammatory mediator [16, 1], we set the source term PIDOSE to 0 for each time step and allowed only AIDOSE to take on nonzero values according to the output of the NMPC algorithm. This strategy achieved an overall healthy resolution for 55% (341) of patients compared to 40% (251) under placebo (Table 4). Even though 27% (99/369) of unhealthy patients were rescued, approximately 4% (9/251) of healthy patients were harmed under this targeted treatment regimen, with four becoming aseptic and five septic. Note that the outcomes of the standard treatment (with $k_{pg} = 0.52$) not only yielded better overall results than this strategy but also did not harm any patients.

In septic scenarios, the control algorithm deems the anti-inflammatory therapy irrelevant, since curbing inflammation will not help eliminate pathogen; thus, only aseptic patients could be helped under this treatment. Figure 5 shows the outcome of targeted therapy, without mismatch, applied to the aseptic patient #26, where anti-inflammatory therapy was successfully able to redirect the response to a healthy outcome. On the other hand, Figure 6 shows results for another aseptic patient (#91) who was not helped with therapy, most likely due to the increased infection severity compared to patient #26. Figure 7 shows an unfortunate example of a therapy-driven septic outcome demonstrated with virtual patient #473. The five therapy-driven septic outcomes and four therapy-driven aseptic outcomes were due to the suppression of inflammation by the anti-inflammatory treatment early

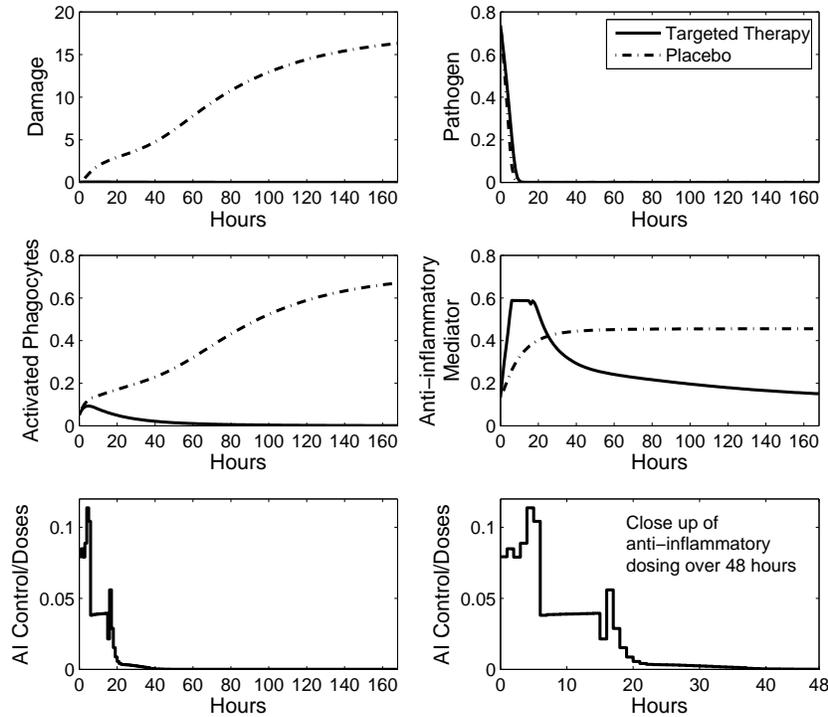


FIGURE 5. The NMPC algorithm is applied to an aseptically infected virtual patient (#26), using anti-inflammatory therapy alone and no patient-model mismatch. Solid curves: Targeted therapy; Dashed curves: Placebo. An anti-inflammatory therapy regimen generated by the algorithm successfully redirects an otherwise aseptically infected response to the healthy state.

in the response. While therapy helped to minimize tissue damage, the pathogen gained the upper hand during the resulting immuno-suppressed state. Once the pathogen levels became sufficiently elevated, the algorithm discontinued the anti-inflammatory therapy and the inflammatory mediators were free to respond fully to the infection. However, by this time, either the pathogen could not be reined in (septic outcomes) or the amount of inflammation needed to successfully eradicate the pathogen was excessive and could not be controlled even by continuing the anti-inflammatory treatment (aseptic outcomes).

4.2.2. *No Mismatch Case 2: Anti- and pro-inflammatory therapy.* Since anti-inflammatory therapy alone could not prevent septic outcomes and was also unsuccessful for many aseptically infected patients, we added pro-inflammatory therapy, which could boost the immune response. Specifically, we allowed PIDOSE to take on nonzero values in equation (2) and reran the algorithm on each patient. The use of combined therapies did not harm any of the patients and 91% (337/369) of unhealthy patients

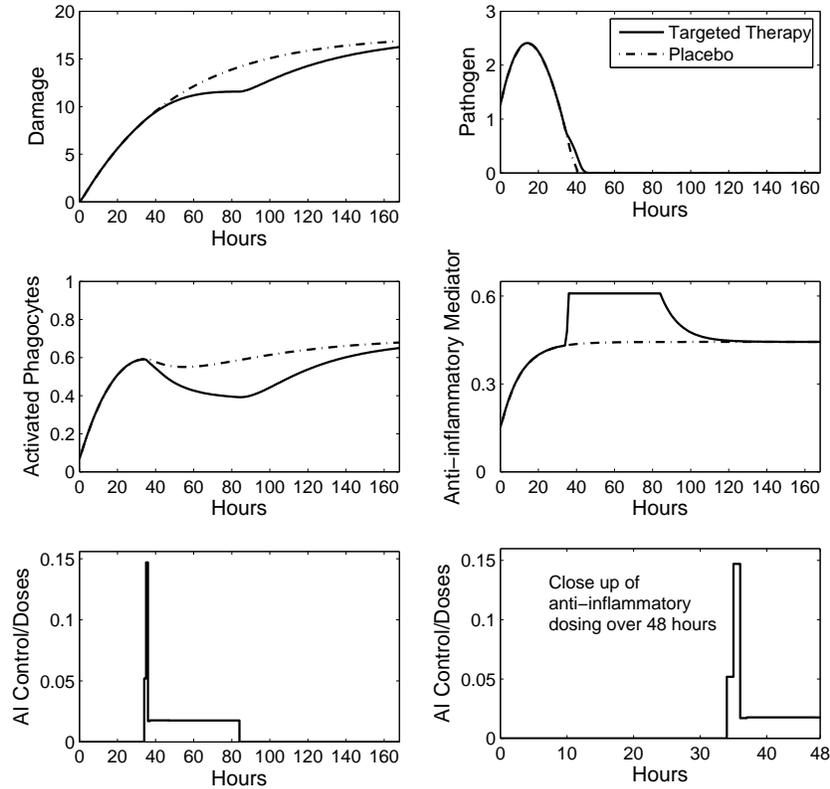


FIGURE 6. The NMPC algorithm is applied to another aseptically infected virtual patient (#91), using anti-inflammatory therapy alone and no patient-model mismatch. Solid curves: Targeted therapy; Dashed curves: Placebo. Unlike the results shown in Figure 5, anti-inflammatory therapy fails to redirect this aseptically infected response to the healthy state, due to the higher levels of pathogen growth seen in this virtual patient.

were rescued. Only 5% (32) of patients overall remained aseptically infected and none remained septic, with 100% of the septic patient population rescued (Table 4). Figure 8 shows a successful outcome for patient #2 under the dual therapy regime. In such successful interventions, initial pro-inflammatory dosing knocks out the pathogen, while subsequent anti-inflammatory dosing reins in the inflammation. If a clinician had full information about each patient, this would be the obvious therapy of choice, tailored specifically for each patient. Overall, 95% (588) of the 620 patients receiving dual therapy resolved to healthy outcomes, making this approach much more favorable than using anti-inflammatory therapy alone as in the previous results. In addition, this strategy was superior to the standard and static therapies. However,

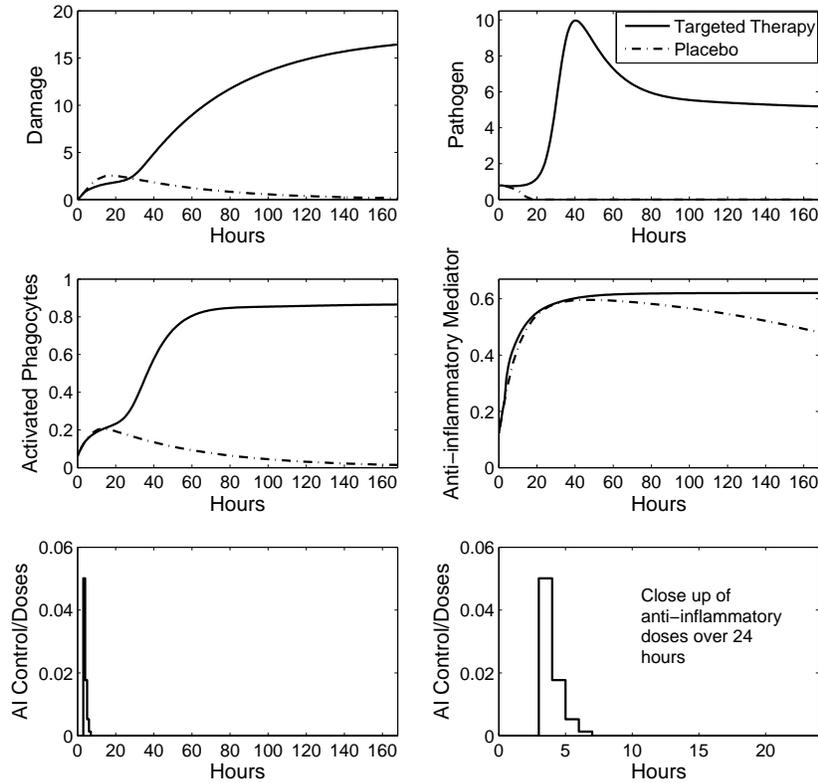


FIGURE 7. The NMPC algorithm is applied to a virtual patient (#473) who would have survived without treatment. Anti-inflammatory treatment alone is used and there is no patient-model mismatch. Solid curves: Targeted therapy; Dashed curves: Placebo. The use of anti-inflammatory therapy suppresses the immune system in an attempt to minimize damage; however, this prevents a successful response, causing the pathogen to grow such that the immune system is unable to control it. By the time the immune system begins to respond, it is too late, showing that the therapy can have adverse effects.

the fact that no patient-model mismatch was included makes these results less impressive, since the presence of mismatch is inevitable and potentially severe in a real world setting.

4.3. Patient-Model mismatch. We next considered the more clinically relevant scenario of patient-model mismatch, reflecting the reality that only highly incomplete observations of a patient's inflammatory response are available. Mismatch was introduced by using the reference parameters listed in Table 1 in the underlying predictive model, instead of patient-specific parameters. To correct for the mismatch

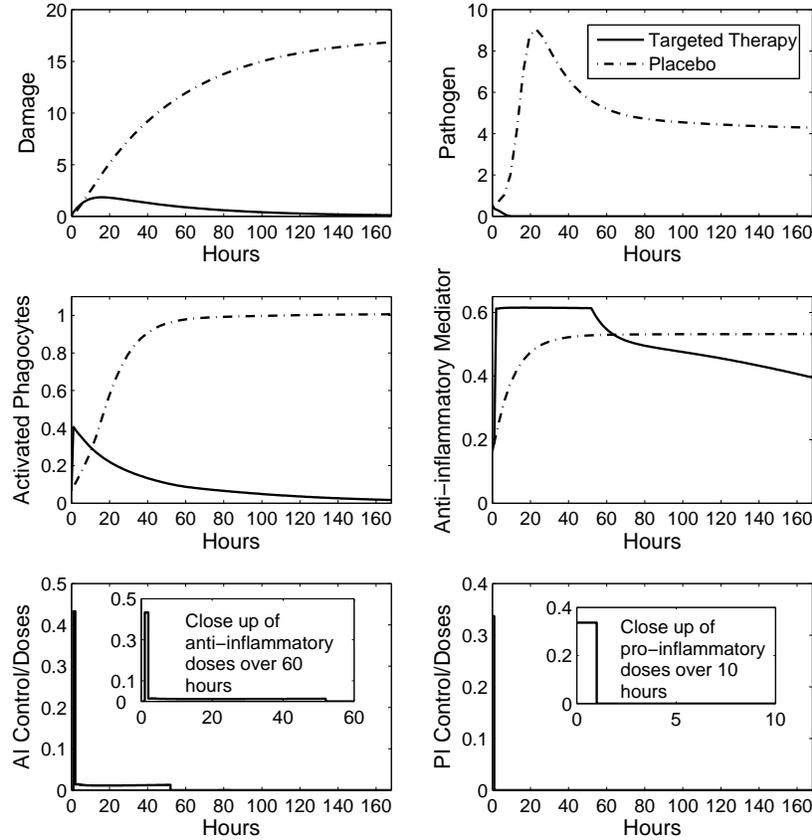


FIGURE 8. The NMPC algorithm is applied to patient #2, who was septic under placebo, with both pro- and anti-inflammatory components and with no patient-model mismatch. Solid curves: Targeted therapy; Dashed curves: Placebo. The therapy found by the algorithm successfully changes the outcome from septic to healthy. This figure is representative of the No Mismatch Case 2 results for all virtual patients who were septic under placebo.

introduced between patient and model, the algorithm had to rely on updates from hourly virtual patient measurements of N^* and C_A and indirect pathogen updates every four hours, as described in Section 3. Only dual therapy was considered under the mismatch construct, permitting both AIDOSE and PIDOSE to take on positive values.

The results show that under targeted therapy with the higher k_{pg} values of 0.6 and 0.8, more than twice as many patients (82% and 83% vs. 40%) achieved healthy outcomes as in the placebo case (Table 5). In addition, targeted therapy with these higher k_{pg} values helped more septic patients. Overall, only 10% (61) of patients

TABLE 5. Patient-Model Mismatch Simulation Results and Comparison to Alternative Therapies. Both targeted and standard therapy results are given over three different k_{pg} values for the underlying model: 0.52, 0.6, and 0.8. The number of patients out of 620 is given in parenthesis (unless otherwise specified). In the first three rows, this number is followed by a number that represents either the quantity of patients coming from an unhealthy placebo category into the healthy category due to the effects of treatment or the number coming from the healthy placebo category into either of the unhealthy categories.

Therapy Type:	Placebo	Static	Standard Therapies:			Mismatch Targeted Therapies:		
			$k_{pg}=0.52$	$k_{pg}=0.6$	$k_{pg}=0.8$	$k_{pg}=0.52$	$k_{pg}=0.6$	$k_{pg}=0.8$
Percentage Healthy:	40% (251)	46% (282; 36)	60% (370; 119)	67% (417; 191)	52% (322; 159)	60% (369; 118)	82% (510; 261)	83% (513; 278)
Percentage Aseptic:	37% (228)	32% (195; 1)	19% (118; 0)	22% (138; 25)	48% (298; 88)	19% (120; 0)	8% (49; 2)	17% (107; 16)
Percentage Septic:	23% (141)	23% (143; 4)	21% (132; 0)	11% (65; 0)	0% (0; 0)	21% (131; 0)	10% (61; 0)	0% (0; 0)
Percentage Harmed (out of 251):	n/a	2% (5/251)	0% (0/251)	10% (25/251)	35% (88/251)	0% (0/251)	1% (2/251)	6% (16/251)
Percentage Rescued (out of 369):	n/a	10% (36/369)	32% (119/369)	52% (191/369)	43% (159/369)	32% (118/369)	71% (261/369)	75% (278/369)

remained in the septic outcome with $k_{pg} = 0.6$ and there were no septic outcomes with $k_{pg} = 0.8$. Further analysis shows that under $k_{pg} = 0.6$, 55% (77/141) of septic patients (under placebo) had a healthy outcome under targeted therapy and only 2% (3/141) of these had an aseptic outcome instead. Correspondingly, under $k_{pg} = 0.8$, 96% (135/141) of septic patients (under placebo) had a healthy outcome under targeted therapy and only 4% (6/141) of these had an aseptic outcome instead. With a model k_{pg} value of 0.52, however, NMPC conferred little gain, as targeted results were nearly identical to standard therapy results with the same k_{pg} value (third column, Table 5). While the targeted therapy results with $k_{pg} = 0.52$ helped relatively few patients, no patients were harmed by this treatment. Under model k_{pg} values of 0.6 and 0.8, however, although the overall results were better, 1% of patients (2/251) and 6% of patients (16/251), respectively, were harmed, resolving to the aseptic condition instead.

Since the presence of noise is inevitable in measurement data, we next incorporated 5% Gaussian noise to each hourly measurement of N^* and C_A to observe the effects of measurement noise on the targeted therapy results discussed above. For each of the three k_{pg} values, six sets of simulations were run on the virtual patient

cohort and the percentages of healthy, aseptic, and septic patients were averaged over the six trials. The inclusion of noise necessitated a very minor modification to the NMPC algorithm, to avoid missing cases where C_A levels were consistently elevated but the measurement of C_A was lower due to noise. For all three k_{pg} values, the outcomes with this addition of noise were quite similar to simulation results without noise.

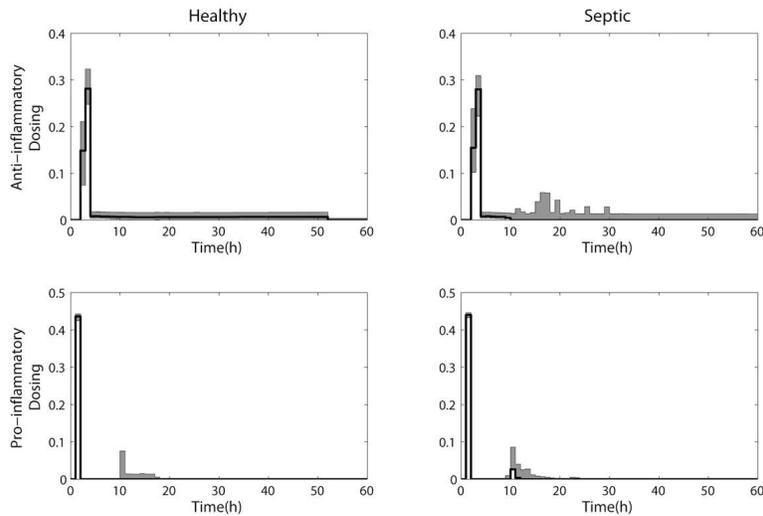


FIGURE 9. Pro-inflammatory and anti-inflammatory dosing profiles that either redirected septic placebo outcomes to healthy (successful) or sustained a septic outcome (unsuccessful). The doses for successful (left panels) and unsuccessful (right panels) control strategies are similar, with an early pro-inflammatory peak that helps control pathogen (bottom row) and a later anti-inflammatory peak that reins in inflammation (top row). Median (solid line) and 2.5%-97.5% dosing range (grey band) are depicted. Successful (left panels) and unsuccessful (right panels) control strategies both feature late surges in pro-inflammatory components, while unsuccessful control strategies feature late surges in anti-inflammatory components as well. Successful control of aseptic patients do not show such a late surge in the pro-inflammatory component (data not shown).

4.4. Comparing intervention strategies and predicting controllability. Although the targeted therapies generated by the NMPC algorithm were individualized, the resulting dosing profiles nonetheless presented the same general features (Figure 9). Not surprisingly, the control strategy first targets pathogen destruction by enhancing pro-inflammation, then modulates anti-inflammation to mitigate excessive inflammation and restore health. This pattern of two slightly offset large dosage peaks is preserved even when the control strategy is unsuccessful. In unsuccessful cases, one also notices a later surge of both pro- and anti-inflammatory dosing, a feature absent in successful control strategies. Moreover, in aseptic cases,

the anti-inflammatory therapy saturates at our preset maximum level. Indeed, this anti-inflammatory cut-off can explain the lack of controllability of several aseptic cases, where it is plausible that allowing a more incisive anti-inflammatory therapy could thwart excessive pro-inflammatory activity. At the intermediate k_{pg} value, $k_{pg} = 0.6$, the cumulative dose of pro-inflammatory therapy varied by 47% across outcomes, with higher mean values for cases yielding septic outcomes, while the cumulative anti-inflammatory dose varied by 100%, with cases in which aseptic death occurred in spite of treatment receiving the highest doses. The cumulative dose of the pro-inflammatory intervention doubled (0.23 to 0.47) with increasing k_{pg} (0.52 to 0.8). The opposite trend held for the anti-inflammatory intervention (1.20 to 0.90), likely reflecting the fact that elevated pro-inflammatories intrinsically promote additional anti-inflammatory activation, as evident in equation (4), resulting in less need for externally applied anti-inflammatory dosing. Together, these variations yielded a three-fold variation in the relative cumulative dosing across k_{pg} .

To identify initial conditions and model parameters predictive of the response to control in cases that would otherwise evolve to septic or aseptic death at $k_{pg} = 0.6$, we designed several classifiers, including logistic regression (LR), naive Bayes, a support vector machine (linear kernel), a neural network (NN), and a variety of decision trees (DT). All of these models were run using 10-fold cross validation. Very consistently, P_0 and k_{pg} were the strongest predictors of non-controllability for both septic and aseptic outcomes. For lower P_0 and k_{pg} values, k_{dn} and k_{cn} were also important predictors. That is, whereas P_0 and k_{pg} did not determine a patient's outcome in the absence of therapy (Figure 3), they nonetheless were dominant factors in determining responsiveness to therapy within each outcome class. We found that LR, NN and DT performed well as classifiers, with an ROC of 0.922 to 0.949, respectively, and misclassification rates of 7.9% to 8.4%. When we omitted P_0 and k_{pg} from the best classifiers, C_{A0} and the parameters k_{nm} and k_{cn} were also discriminatory of outcome, but the ROC dropped to 0.676 for the best LR classifier, with 20% misclassification.

5. Discussion. We have shown that incorporation of an NMPC algorithm into a model of the acute inflammatory response to pathogenic infection allows for the derivation of therapeutic interventions that can produce healthy resolutions for virtual patients that would have otherwise faced septic or aseptic outcomes. If we supply the model with complete information about each virtual patient, with parameters selected from anywhere within a physiological range, then the NMPC routine almost always produces a healthy outcome, as long as both pro- and anti-inflammatory dosing are included. If we consider a more realistic scenario in which only clinically available virtual patient information is supplied to the predictive model, then, although the results are less perfect than in the no-mismatch case, the NMPC algorithm again significantly enhances the likelihood of healthy outcomes across the virtual patient population, an advantage maintained in the presence of noisy patient measurements. The success rate of the algorithm depends on the pathogen growth rate, k_{pg} , used in the predictive model, with use of a high growth rate leading to a therapy with a strong pro-inflammatory component that prevents septic outcomes, and use of a more moderate growth rate minimizing the rate of aseptic outcomes.

Therapeutic dosing profiles can be quite similar on trials that lead to different outcomes. Differences in dynamic dosing are subtle indeed, while outcomes

themselves are strikingly influenced by treatment. This finding argues strongly for individually titrated therapy, especially given the modest effects observed from constant, untitrated anti-inflammatory therapy and the diverse outcomes obtained from even a standardized dynamic dosing profile.

Although clinicians practice titrated patient care in modern intensive care units to support organ systems and combat infections, existing guidelines as to how to treat patients with potentially lethal infections are generic, with very limited room for individualized titration [10]. In particular, targeted immunomodulation remains an elusive goal [7, 6]. Even when appropriate drugs are available, there is no unifying concept as to how these should be used and combined to improve patient outcome. Model-based immunomodulation of complex inflammatory diseases thus represents an extraordinary opportunity for very significant advances in the care of the critically ill. Indeed, the idea of using a control-based algorithm to generate appropriate therapeutic regimens makes particular sense in light of the fact that the acute inflammatory response features multiple interconnected nonlinear feedback loops that would vastly complicate the design of successful therapies under clinical conditions.

In this work, we implement control algorithms in a four-equation reduced ordinary differential equation model for the acute inflammatory response [20]. Reynolds et al. briefly explored the possibility of modulating the outcome of a simulated infection, in the context of this model, by altering the anti-inflammatory mediator levels at a particular time point. It was found that most of the alterations perturbed the system from an otherwise healthy resolution to an unfavorable state. The NMPC algorithm that we implemented yields a much more refined approach to the design of therapeutic strategies. In particular, our results illustrate the need for multimodal therapeutic strategies to successfully modulate a complex inflammatory response.

Our approach has several limitations. Our representation of the inflammatory response is extremely simplified and our lumped model variables and parameters cannot be mapped directly to quantities accessible in clinical or laboratory settings. It would be advantageous to incorporate the algorithm into a more detailed model that can give quantitative predictions about specific mediators of the acute inflammatory response (e.g. [8]). In addition, the pharmacokinetics of existing or potential therapies could be incorporated to generate specific predictions with respect to particular dosing profiles. The customizations of our NMPC algorithm and features of the controller that we used were based on pragmatic clinical considerations, yet some decisions were admittedly arbitrary. For example, although we limited the intensity of anti-inflammatory therapy to cap the risk of secondary infections associated with states of elevated anti-inflammation, it appears clear that allowing more intensive anti-inflammatory therapy would rescue more cases otherwise destined for aseptic death. Ideally, such limitations on therapy would be implemented much more precisely, based on clinical or experimental data. Another issue for future consideration is the construction of an optimal objective function, which is nontrivial, even in this very simplified context. We experimented with a wide combination of weights and penalty terms, arriving at our objective function empirically. We believe that, for clinical applications, biological input will be essential in constructing suitable objective functions, and that there could be a good argument made for the inclusion of time-varying weights in the objective functions used.

The translation of biomarker-targeted methods, such as we have implemented in this work, to clinical practice will require empirically validated predictions. In a realistic experimental protocol, a mathematical model that forecasts several immune mediators fairly well from a mouse/rat model of endotoxin challenge would be used as the underlying model, predicting the immune response of each animal. A small set of initial observations of individual animals subjected to an infectious challenge would be taken to generate an ensemble of candidate models for each subject. The NMPC algorithm, applied to the predictive ensemble of models, would suggest a therapeutic intervention that would be implemented in each animal, after which measurements of various analytes would be taken to update the ensemble (adaptive control). After repeated iterations of this procedure, comparisons of outcomes in animals receiving such targeted therapy would be made to those observed in animals either receiving a standardized therapy or no therapy at all.

In conclusion, based on the results presented in this paper and other biomedical applications of feedback control, we maintain that the utilization of feedback control methods to assist in treating the critically ill is a strategy worth further exploration. In practice, clinicians currently implement a sort of dosing strategy algorithm, in which they iteratively determine the next treatment step based on all of the information available to them. To utilize their knowledge and expertise most effectively, it is logical to propose the incorporation of model-based control algorithms as a tool to guide them in this process of selecting patient-specific treatment strategies.

Acknowledgments. The authors thank Jeff Florian and Robert Parker of the University of Pittsburgh's Chemical Engineering Department for the suggestion to apply NMPC, specifically, to the problem at hand, for making an algorithm available, and for their very helpful assistance in introducing the authors to the NMPC methodology and the specific algorithm. In addition, the authors thank Senthil Natarajan for making it possible for many of the initial NMPC simulations to be run on PittGrid, a computing network resource of the University of Pittsburgh. This work was supported by NIH grant R01-GM67240 (JR and GC), HL80926(GC) and NSF Agreement Number 0112050 (JD).

REFERENCES

- [1] D. Annane, V. Sebille, C. Charpentier, P. E. Bollaert, B. Francois, J. M. Korach, G. Capellier, Y. Cohen, E. Azoulay, G. Troche, P. Chaumet-Riffaut and E. Bellissant, *Effect of treatment with low doses of hydrocortisone and fludrocortisone on mortality in patients with septic shock*, JAMA, **288** (2002), 862–871.
- [2] R. C. Bone, *The search for a magic bullet to fight sepsis*, JAMA, **269** (1993), 2266–2267.
- [3] R. C. Bone, *Why sepsis trials fail*, JAMA, **276** (1996), 565–566.
- [4] R. C. Bone, *Immunologic dissonance: A continuing evolution in our understanding of the systemic inflammatory response syndrome (SIRS) and the multiple organ dysfunction syndrome (MODS)*, Ann. Intern. Med., **125** (1996), 680–687.
- [5] C. C. Chow, G. Clermont, R. Kumar, C. Lagoa, Z. Tawadrous, D. Gallo, B. Betten, J. Bartels, G. Constantine, M. P. Fink, T. R. Billiar and Y. Vodovotz, *The acute inflammatory response in diverse shock states*, Shock, **24** (2005), 74–84.
- [6] A. S. Cross, S. M. Opal, K. Bhattacharjee, S. T. Donta, P. N. Peduzzi, E. Furer, J. U. Que and S. J. Cryz, *Immunotherapy of sepsis: Flawed concept or faulty implementation?*, Vaccine, **17** (Suppl. 3) (1999), S13–S21.
- [7] A. S. Cross and S. M. Opal, *A new paradigm for the treatment of sepsis: Is it time to consider combination therapy*, Ann. Intern. Med., **138** (2003), 502–505.
- [8] S. Daun, J. Rubin, Y. V. Vodovotz, A. Roy, R. S. Parker and G. Clermont, *An ensemble of models of the acute inflammatory response to bacterial lipopolysaccharide in rats: Results from parameter space reduction*, J. Theor. Biol., **253** (2008), 843–853.

- [9] J. Day, J. Rubin, Y. Vodovotz, C. C. Chow, A. Reynolds and G. Clermont, *A reduced mathematical model of the acute inflammatory response II. Capturing scenarios of repeated endotoxin administration*, *J. Theor. Biol.*, **242** (2006), 237–256.
- [10] R. P. Dellinger, M. M. Levy, J. M. Carlet, J. Bion, M. M. Parker, R. Jaeschke, K. Reinhart, D. C. Angus, C. Brun-Buisson, R. Beale, T. Calandra, J. F. Dhainaut, H. Gerlach, M. Harvey, J. J. Marini, J. Marshall, M. Ranieri, G. Ramsay, J. Sevransky, B. T. Thompson, S. Townsend, J. S. Vender, J. L. Zimmerman and J. L. Vincent, *Surviving Sepsis Campaign: International guidelines for management of severe sepsis and septic shock: 2008*, *Crit. Care Med.*, **36** (2008), 296–327.
- [11] F. Doyle, L. Jovanovic, D. Seborg, R. S. Parker, B. W. Bequette, A. M. Jeffrey, X. Xia, I. K. Craig and T. McAvoy, *A tutorial on biomedical process control*, *J. Proc. Con.*, **17** (2007), 571–594.
- [12] J. A. Florian, Jr., J. L. Eiseman and R. S. Parker, *Nonlinear model predictive control for dosing daily anticancer agents using a novel saturating-rate cell-cycle model*, *Comput. Biol. Med.*, **38** (2008), 339–347.
- [13] J. Klastersky and R. Capel, *Adreno-corticosteroids in the treatment of bacterial sepsis: A double-blind study with pharmacological doses*, *Antimicrobial Agents Chemother.*, **10** (1970), 175–180.
- [14] R. Kumar, G. Clermont, Y. Vodovotz and C. C. Chow, *The dynamics of acute inflammation*, *J. Theor. Biol.*, **230** (2004), 145–155.
- [15] M. M. Levy, M. P. Fink, J. C. Marshall, E. Abraham, D. Angus, D. Cook, J. Cohen, S. M. Opal, J. L. Vincent and G. Ramsay, *2001 SCCM/ESICM/ACCP/ATS/SIS International Sepsis Definitions Conference*, *Crit. Care Med.*, **31** (2003), 250–256.
- [16] D. Lipiner-Friedman, C. L. Sprung, P. F. Laterre, Y. Weiss, S. V. Goodman, M. Vogeser, J. Briegel, D. Keh, M. Singer, R. Moreno, E. Bellissant and D. Annane, *Adrenal function in sepsis: The retrospective Corticus cohort study*, *Crit. Care Med.*, **35** (2007), 1012–1018.
- [17] J. C. Marshall, E. Deitch, L. L. Moldawer, S. Opal, H. Redl and P. T. van Der, *Preclinical models of shock and sepsis: What can they tell us?*, *Shock*, **24** (Suppl. 1) (2005), 1–6.
- [18] B. Ogunnaike and W. Ray, “Process Dynamics, Modeling, and Control,” Oxford University Press, New York, 1994.
- [19] R. S. Parker, F. J. Doyle, III and N. A. Peppas, *The intravenous route to blood glucose control*, *IEEE Eng. Med. Biol. Mag.*, **20** (2001), 65–73.
- [20] A. M. Reynolds, J. Rubin, G. Clermont, J. Day, Y. Vodovotz and B. Ermentrout, *A reduced mathematical model of the acute inflammatory response: I Derivation of the model and analysis of anti-inflammation*, *J. Theor. Biol.*, **242** (2006), 220–236.
- [21] Y. Vodovotz, C. C. Chow, J. Bartels, C. Lagoa, J. M. Prince, R. M. Levy, R. Kumar, J. Day, J. Rubin, G. Constantine, T. R. Billiar, M. P. Fink and G. Clermont, *In silico models of acute inflammation in animals*, *Shock*, **26** (2006), 235–244.
- [22] H. D. Volk, P. Reinke and W. D. Docke, *Clinical aspects: From systemic inflammation to ‘immunoparalysis’*, *Chem. Immunol.*, **74** (2000), 162–177.

Received March 30, 2010; Accepted May 21, 2010.

E-mail address: judyday@gmail.com

E-mail address: jonrubin@pitt.edu

E-mail address: clermontg@ccm.upmc.edu