

# Acute Inflammation Treatment via Particle Filter State Estimation and MPC

Justin S. Hogg\* Gilles Clermont\*\* Robert S. Parker\*\*\*,<sup>1</sup>

\* Carnegie Mellon – University of Pittsburgh Ph.D. Program in  
Computational Biology (e-mail: jsh32@pitt.edu)

\*\* Critical Care Medicine (e-mail: cler@pitt.edu)

\*\*\* Chemical and Petroleum Engineering (e-mail: rparker@pitt.edu)  
University of Pittsburgh, Pittsburgh, PA 15261 USA

---

**Abstract:** Models of acute inflammatory disease may have the potential to guide treatment decisions in critically ill patients. Model Predictive Control (MPC) leverages the predictive capacity of a model to determine a control strategy that guides a system to a target trajectory. As applied to acute inflammation, MPC might be used to guide a patient from disease to health by monitoring the patient state, computing and applying an optimal intervention strategy, and updating the strategy if the patient state diverges from predictions. A key challenge to the application of MPC is mapping the observable patient state into the complete state space of the model. We propose that a Particle Filter (PF) is a suitable algorithm for state estimation in nonlinear models of acute inflammation. As a proof of concept, we apply MPC and PF to the administration of hemoadsorption (HA) treatment in an 8-state model of endotoxemia in rats. *In silico* tests demonstrate that the PF generates accurate state estimates from limited observations in the presence of noise and parameter uncertainty. Furthermore, we explore the maximal predicted benefits of HA treatment with a standard single column configuration and hypothetical multi-column configurations, where each column has a specificity for a target cytokine. Simulations suggest that two column HA will improve treatment efficacy, but physiological restrictions on HA will limit benefits from higher order configurations.

*Keywords:* Sepsis, Inflammation, Model Predictive Control, Particle Filter, State Estimation.

---

## 1. INTRODUCTION

Inflammation is the body's natural response to infection and trauma (Janeway and Medzhitov, 2002). Macrophage cells play a key role in the initiation and orchestration of the inflammatory response. These immune cells reside in tissues where they monitor the environment for molecules derived from pathogen or damaged cells (Medzhitov and Janeway, 2000). When such molecular patterns are detected, macrophage secrete pro-inflammatory mediators (such as TNF and interleukin(IL)-1) and chemoattract other white blood cells, thus initiating an immune response (Cohen, 2002). TNF and IL-1 induce endothelial cells to express adhesion molecules for neutrophils, a circulating white blood cell. Neutrophils migrate into the tissues, following chemokine gradients, where they scavenge and digest pathogen in a process called phagocytosis (Abraham, 2003). Pro-inflammatory cytokines also trigger an anti-inflammatory wave that suppresses inflammation and returns the system to baseline as the infection or damage is cleared. IL-10 is a powerful anti-inflammatory cytokine that suppresses the expression of pro-inflammatory cytokines and the activity of innate immunity effector cells (Fujiwara and Kobayashi, 2005). Anti-inflammation is an essential regulator of the inflammatory response that thwarts potential deleterious cytotoxic effects of vigorous pro-inflammation. Although the goal of the inflammatory

response is to contain and eliminate the initial biological stressor (*e.g.* infection) and thus restore a healthy state, it also has the potential to push the system into a number of pathological states.

Sepsis, a systemic inflammatory response triggered by infection, is one such pathology (Bone et al., 1992). Sepsis is a common reason for admission to the intensive care unit. Severe sepsis may lead to the failure of multiple organ systems (Singh and Evans, 2006) and death in 40% of patients (Angus et al., 2001). The cascade of organ failures may occur even while the infection is suppressed by antibiotics, demonstrating that uncontrolled inflammation is destructive to organ tissues. Numerous attempts have been made to control sepsis using immunomodulation such as anti-inflammatory treatments (*e.g.* anti-TNF antibodies, IL-10, etc.), but clinical trials have failed to find consistent benefits in randomized populations of septic patients (Dellinger et al. (2008), Vincent et al. (2002)).

Hemoadsorption (HA) is a blood purification treatment that has been shown to improve short term survival in septic rats (Kellum et al., Peng et al.). The HA device is a column packed with high surface area, bioreactive beads that adsorb proteins in the size range typical of cytokines. A portion of the septic animal's blood is diverted through the HA column, where TNF, IL-6, and IL-10 are removed from the blood (Kellum et al., 2004). The flow rate through the column may be adjusted to achieve maximal benefit,

<sup>1</sup> To whom correspondence should be addressed: +1-412-624-7364

although physiological constraints define a maximum practical rate. The adsorption is not selective and presumably adsorbs other molecules in the size range of cytokines. A dynamic model of cytokine adsorption in the HA column was reported in DiLeo et al. for the case of constant flow.

A number of system models of the inflammatory response have been constructed, ranging from phenomenological low order models that recapitulate asymptotic system behavior to predictive models of varying complexity calibrated to experimental data. Simple phenomenological models of inflammation have been used to explain the multiple outcomes possible with abstract pro-inflammatory and anti-inflammatory feedback loops (Kumar et al., Reynolds et al., Day et al.). More complex models that include measurable biomarkers, such as TNF and IL-10, have been calibrated to experimental data and show potential for the development of predictive models (Chow et al., Daun et al.). Such models serve as a discovery platform for new treatments for acute inflammation (Kumar et al., Clermont et al.). Furthermore, predictive models should have the potential to guide individualized treatment regimens for patients in real time.

Model Predictive Control (MPC) uses the predictive capacity of a model to determine control inputs that guide a dynamic process towards a target trajectory (Muske and Rawlings, 1993). MPC is a receding finite horizon method, where the control decision is based only on the predicted trajectory over a finite time window (*i.e.* horizon). As time advances, the horizon slides and new control decisions are based on the latest system observations. In the case of acute illness, MPC of HA could maintain homeostasis and guide the patient to health based on observations of biomarkers and vital signs.

In current models of acute inflammation, a subset of state variables are either difficult to measure or are abstract variables that do not correspond to a distinct biological observable. Therefore, a method of estimating the model state space given a set of observables is required. This is known as the state estimation problem. Rawlings and Bakshi (2006) describe a variety of methods for nonlinear state estimation in the context of MPC. The unscented Kalman filter (UKF) is commonly implemented for state estimation in non-linear models and assumes normally distributed noise. This is a poor assumption for models of sepsis, where distributions of cytokine measurements are highly positively skewed with variance positively correlated to their mean. Particle Filter (PF) state estimation permits arbitrary noise models and is appealing for inflammation models due to its generality and ease of implementation.

Particle Filtering is a stochastic state estimation method described as “survival of the fittest”. A number of parallel simulations (“particles”) are randomly initialized from a prior distribution of states. Following an observation, simulations are weighted by the posterior probability that the observation was generated by the particle. The particles are then resampled (with replacement) from the weighted distribution. The resampled particle simulations are evolved dynamically until the next observation. The expected state values are given by the weighted average of the particle states. See Bishop (2006) for an overview. The

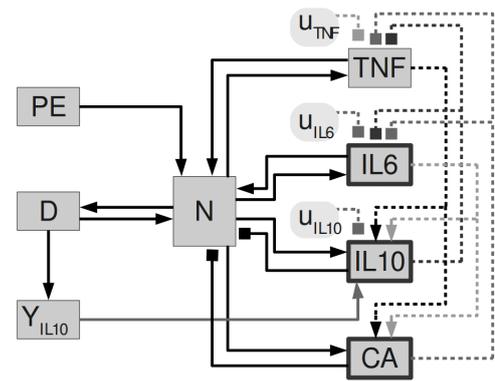


Fig. 1. Block diagram of the Daun et al. (2008) endotoxemia model. Boxes corresponds to a state variables; bold edges indicate observed states. Arrows indicate upregulation, lines ending with squares indicates downregulation. Dotted lines between cytokine blocks represent indirect feedback mechanisms that act through  $N$ . The control variables  $u_{tnf}$ ,  $u_{il6}$  and  $u_{il10}$  correspond to HA filtering.

key limitation of PF is the computational burden as the dimensionality of the system increases.

This manuscript describes an *in silico* demonstration of real-time control of hemoadsorption in acute inflammation using a framework of MPC with PF state estimation.

## 2. METHODS

### 2.1 Endotoxemia model

All simulations in this study are based on an 8-state ODE model of endotoxemia in rats, previously presented by Daun et al. (Fig. 1). The model states are endotoxin concentration,  $PE$ ; number of activated phagocytic cells,  $N$ ; a measure of tissue damage,  $D$ ; an abstract, long-acting anti-inflammatory,  $CA$ ; three key extracellular mediators of acute inflammation: IL-6, TNF, and IL-10; and a first-order filter between  $D$  and IL-10,  $Y_{IL10}$ . Model equations are included in the appendix. The model was calibrated to rat serum concentrations of TNF, IL-10 and IL-6 obtained at 1, 2, 4, 8, 12, and 24 hours after an intravenous injection of 3 or 12 mg/kg endotoxin.

A system of stochastic differential equations was induced by the deterministic model with a multiplicative noise process. The stochastic equation for each state  $w_i$  is:

$$dw_i(t) = f_i(\mathbf{w}(t))dt + \sigma w_i(t)d\mathcal{W}_t \quad (1)$$

where  $\mathbf{w}$  is the state space vector and  $\mathcal{W}_t$  is a Wiener process, and  $f_i(\mathbf{w}(t))$  is the deterministic derivative function. This stochastic system can be integrated with the Euler-Maruyama method: the deterministic portion is integrated and a noise term is sampled from the distribution  $w_i(t)\sigma\sqrt{\Delta t}\mathcal{N}(0,1)$ . For small  $\Delta t$ ,  $\sigma$  the solution is approximated by integrating the deterministic portion and multiplying by a sample from a  $\log\text{-}\mathcal{N}(0, \sigma^2\Delta t)$  distribution. This latter approach was used for simulation as it avoids the possibility of generating a negative sample. The parameter  $\sigma = 0.05/\text{hr}$  was selected to balance the stability of the system with a non-trivial variation in the output trajectory. The deterministic ODE was integrated

using the CVODE library, while the stochastic noise was applied at fixed intervals of  $\Delta t = 0.1\text{hr}$ . State noise was applied independently to each state. Observation noise was applied to the trajectories of the measured cytokines TNF, IL-10, and IL-6. Observation noise was multiplicative and  $\log\mathcal{N}(\mu = 0, \sigma^2 = 0.15^2)$  distributed.

## 2.2 Hemoadsorption

Hemoadsorption was applied to the model system as a first-order elimination of cytokines with rate  $\frac{\text{flow rate}}{\text{blood volume}}$ . A more detailed model of HA was presented by DiLeo et al. (2009). However, this model assumes a fixed flow rate for the duration of HA, which is not suitable for our purposes. Three configurations of HA are simulated (Table 1). The single-column configuration is the only option currently available in the lab. The present work examines *in silico* the potential utility of having multiple HA columns with selective cytokine adsorption capabilities.

## 2.3 Particle Filter Implementation and Initialization

A particle filter was implemented in MATLAB according to a standard algorithm (Bishop, 2006). While the endotoxemia model was calibrated to mean data in a rat population, individual rats respond differently to endotoxin challenge due to genetic and environmental variations. We assume that individual parameters are drawn from a prior distribution centered around the best-fit population parameters. If individual variation is neglected, PF state estimates would be biased towards the population mean. To eliminate this bias, each particle  $p$  was initialized with randomized parameter  $\theta_i^{(p)}$  from the prior distribution  $\theta_i \cdot \log\mathcal{N}(0, 0.15^2)$ . Future work may improve estimates of the prior distribution. We also assume the endotoxin dose and administration time is unknown. This reflects the uncertainty in a clinical environment where a patient is admitted at some time following an infection. Each particle was initialized with a random endotoxin dose from the uniform distribution over  $[0, 15]$  mg/kg, and a random time delay before hemoadsorption from the uniform distribution over  $[0, 6]$  hrs. Each particle is simulated for the length of its delay before measurements are obtained and the HA protocol is administered. Hourly measurements of IL-6, TNF, and IL-10 were passed to the PF to obtain state estimates for the hidden variables.

## 2.4 Model Predictive Control

A schematic of the in silico experiment combining MPC, PF, HA, and simulated endotoxic rat is shown in Fig. 2. MPC is a popular control methodology for biomedical systems (Parker and Doyle III, 2001), based on its ability to robustly manage subject-model mismatch in a variety of disease case studies. A nonlinear MPC algorithm was

Table 1. Hemoadsorption configurations

HA configuration	cytokine specificity		
	column 1	col. 2	col. 3
1	TNF, IL-6, IL-10	–	–
2	TNF, IL-6	IL-10	–
3	TNF	IL-10	IL-6

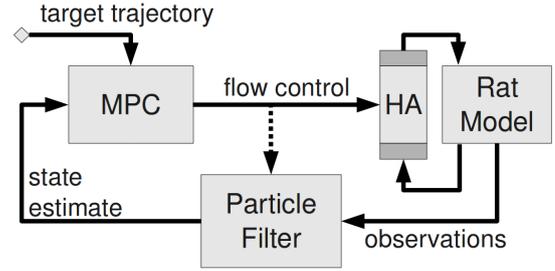


Fig. 2. Schematic of the in-silico experiment with a simulated endotoxemic rat, Particle Filter state estimates, and Model Predictive Control of hemoadsorption.

implemented using HA column flow rate(s) as the manipulated variable(s), and damage ( $D$ ) as the controlled variable (target of  $D = 0$ ). As  $D$  is not measurable, the PF state estimated was used at each observation time point. The objective function for the algorithm was:

$$\begin{aligned} \min_{\mathbf{u}(k|k)} \quad & \|\Gamma_y(\mathbf{Y}(k+1|k) - \mathbf{R}(k+1))\|_2^2 + \\ & \|\Gamma_u \mathbf{U}(k|k)\|_2^2 + \|\Gamma_{\Delta u} \Delta \mathbf{U}(k|k)\|_2^2 \quad (2) \\ \text{s.t.} \quad & \mathbf{U}(k|k) \leq 72\text{ml/hr} \end{aligned}$$

These terms penalize predicted error in the controlled variable ( $\mathbf{Y}$ ) from the reference ( $\mathbf{R}$ ), the use of large HA flow rates ( $\mathbf{U}$ ), and changes in the flow rates ( $\Delta \mathbf{U}$ ). Standard statistical notation is employed throughout (prediction at time  $k+1$  given information up to time  $k$ ), and the  $\Gamma$  weights are selected such that  $\Gamma_y = \mathbf{I} > \Gamma_u = \frac{0.02}{c} \mathbf{I} > \Gamma_{\Delta u} = \frac{0.01}{c} \mathbf{I}$ . The number of HA channels used in the intervention is  $c$ . Minimization was performed using *fmincon* in MATLAB (©2009, The MathWorks, Natick, MA). The MPC time step was  $\Delta t = 1\text{ hr}$ , the prediction horizon was  $p = 6$  steps, and the move horizon was  $m = 2$ . The size of  $\mathbf{U}(k|k)$  was  $m * c$ , with the HA configuration establishing the number of adjustable flow rates. For computational efficiency, model predictions were based only on the deterministic portion of the endotoxemia model. The weighted averages of particle parameters were used as parameters for the model prediction.

## 3. RESULTS

### 3.1 Particle Filter State Estimation

The performance of a Particle Filter depends on the number of particles. To determine a sufficient number of particles for subsequent experiments, we evaluated the PF performance for various numbers of particles. Performance was based on the average relative error ( $E_{\text{absrel}}$ ) over all states and all observed time points:

$$E_{\text{absrel}} = \frac{1}{8T} \sum_{s=1}^8 \sum_{k=0}^T \frac{|y_s^{\text{est}}(t_k) - y_s^{\text{true}}(t_k)|}{y_s^{\text{true}}(t_k)} \quad (3)$$

Here  $T$  is the number of observed time points, and  $y_s$  is the value of state  $s$ . Fig. 3 shows PF performance versus the number of particles over 21 simulations with randomized endotoxin dose and time interval before first observation. As expected, the performance improved as the number

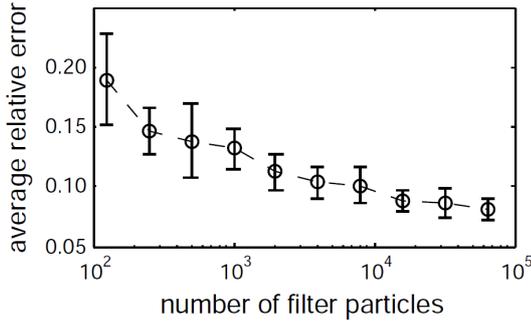


Fig. 3. Particle Filter performance is dependent on the number of particles. Performance was measured as the average relative error over all model states at all observed time points. Error bars show the 95% confidence interval of the mean ( $t$ -distr.,  $N = 21$ ).

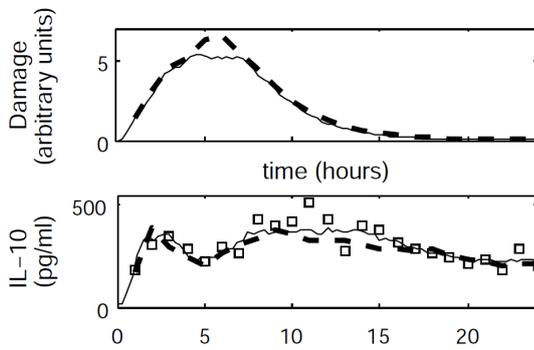


Fig. 4. An example of Particle Filter state estimation on simulated data. Square points are observed data fed to the PF; dashed lines show the state estimate generated by PF; solid lines show the true state.

of particles increased. Average relative error dropped to about 10% with 4000 particles. Increasing the particle numbers to 8000 and 16,000 further reduced the error, but at the expense of computational time. We chose 4000 particles for subsequent experiments as a balance between reasonable accuracy and computational time. 24 hours of simulated time required about 10 CPU minutes. Note that average relative error is below noise levels of the observed states. Fig. 4 shows an example of PF state estimation (4000 particles) on simulated data. Estimates are shown for a hidden state,  $D$ , and an observed state, IL-10.

### 3.2 Hemoadsorption Control

MPC of HA was simulated with 1, 2 or 3 columns. See Table 1 for descriptions of each configuration. MPC performance was evaluated by the average relative reduction in damage area-under-the-curve (damage AUC) for the 24 hour interval following endotoxin administration:

$$\Delta^{rel} = \frac{\int_0^{24} (D_{HA}(t) - D_{sham}(t)) dt}{\int_0^{24} D_{sham}(t) dt}, \quad (4)$$

where  $D_{HA}$  is the damage trajectory with HA and  $D_{sham}$  is the damage trajectory without HA. Endotoxin dose was 8 mg/kg and the first observation occurred at 30

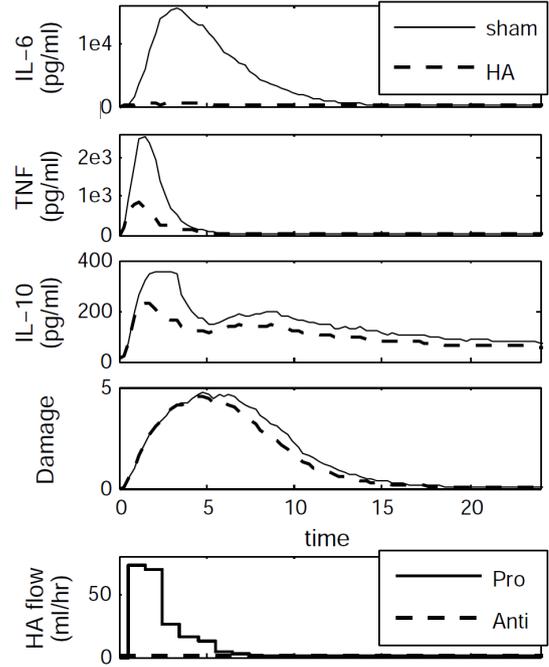


Fig. 5. Sample simulation results from paired sham and HA control experiments. HA was applied using configuration 2.

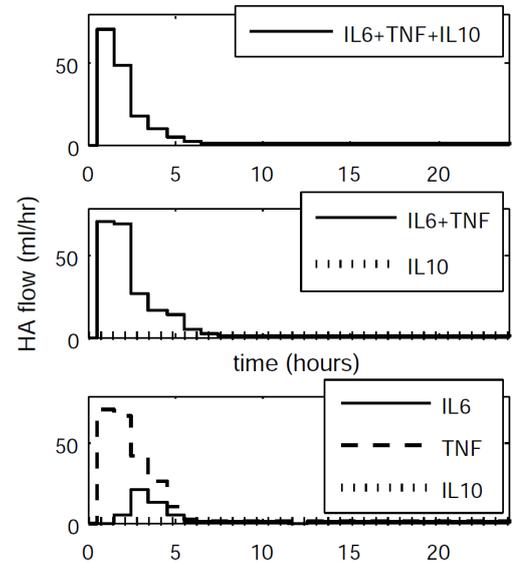


Fig. 6. Typical HA column flows for 1, 2 and 3 column configurations. In all cases, MPC applies aggressive flow at the start of treatment and then drops to zero as further intervention has negligible effect.

minutes (information withheld from PF). Identical noise was applied to paired HA and sham simulations to limit stochastic variation. Each configuration was simulated  $N = 14$  times. Fig. 5 shows typical trajectories for paired HA and sham simulations using configuration 2.

Single column HA reduced damage AUC by 7% on average. The controller applied maximum flow to the column initially and then tapered the flow to zero as the benefit of additional HA became negligible. Dual column HA

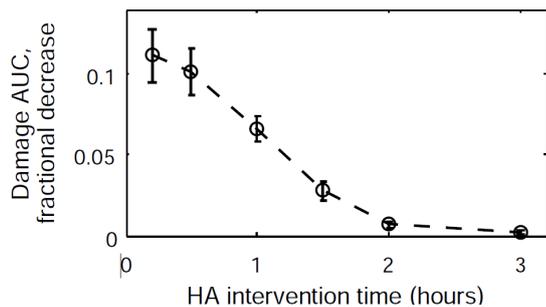


Fig. 7. Treatment delay decreases HA efficacy. A simulated 8 mg/kg endotoxin dose was followed by HA (configuration 2) after a variable time delay. Prompt intervention reduced Damage AUC by 11%, but efficacy dropped rapidly past one hour. Error bars show 95% confidence of means ( $t$ -distr.,  $N = 21$ ).

reduced damage AUC by 11% on average, significantly more than the single column configuration. No additional benefit was observed with the triple column configuration. A comparison of the control schemes for the various configurations is shown in Fig. 6. Maximal flow is applied to the TNF/IL-6 column in the two column configuration, while flow to the IL-10 column is zero. The improved damage reduction in this configuration demonstrates a benefit in filtering pro-inflammatory TNF and IL-6 while retaining the anti-inflammatory IL-10. In the triple column configuration each cytokine is filtered by a separate column. It might be expected that the performance of three columns should meet or exceed the performance of the two column arrangement, but the constraint that the sum of column flows not exceed 72 ml/hr prevents the controller from filtering TNF and IL-6 with the same intensity as in the two column scheme. The controller first aggressively filters TNF, to the exclusion of IL-6, but then applies flow to the IL-6 column as the flow to the TNF column decreases.

In order to determine the importance of timely HA intervention, MPC control of HA was initiated at various times following endotoxin administration. In all cases, the dose and timing of endotoxin was hidden from the PF. Fig. 7 shows the relationship between HA intervention timing and HA performance. Prompt intervention at 12 minutes post endotoxin dosing resulted in 12% damage AUC reduction on average. The performance of HA diminished rapidly after one hour and was near zero by hour two. This simulation result suggests that HA intervention requires timely application during the onset of acute inflammation.

#### 4. DISCUSSION AND SUMMARY

Our results demonstrate that MPC with PF state estimation is suitable for online control of hemoabsorption in acute inflammation. Furthermore, simulations suggest that two-column HA devices may reduce tissue damage more than conventional HA, while three or more columns will have diminishing impact due to physiologic limitations on the sum of flows through the columns. Since the currently available HA device adsorbs a broad spectrum of cytokines, these results provide motivation for the development of cytokine-specific columns.

PF state estimation and MPC computation was fast enough for online implementation. On a typical PC, state estimation with 4000 particles and MPC required less than 1 minute of computation time per hour of simulated time. The bulk of the computation is consumed by particle simulation. Computational cost may become prohibitive in models with larger state spaces if larger particle sets are required. Future work is needed to determine scalability to higher dimensional models of inflammation.

This work assumed that real-time measurements of the observables were available at point-of-care. In current practice, IL-6 is the only cytokine measurement that is available bedside. Boyle et al. (2006) were able to produce semi-quantitative measurements of cytokines in under 45 minutes. Efforts to develop point-of-care inflammatory profiles are underway, but the time requirement is currently a barrier to online control implementation.

The model used to simulate the endotoxemic rat was also used by the MPC controller to predict trajectories. While parameter uncertainty was introduced to mimic the genetic and environmental variability in the population, we had the benefit of a predictive model with correct mathematical structure. Future work is needed to determine whether performance is robust to subject-model mismatch. At present, alternative models calibrated to our dataset are not available for mismatch analysis. A more detailed model of acute inflammation developed by Chow et al. (2005) is a candidate for future comparison.

In a clinical setting, an endotoxin dose is not limited to a single bolus at the onset of infection. In sepsis, for example, a constant source of endotoxin can be delivered from an ongoing infection in the abdomen. In order to determine the effect of unmeasured endotoxin step disturbances, simulations were performed with a constant source of endotoxin (results not shown). The resulting state estimates for the measured variables were accurate, but biased for the unmeasured variables. The endotoxin variable (PE) settled at a positive steady state, while the state estimate tended to zero. As a preliminary attempt to correct the bias, a modified PF was tested. The modified PF included a random variable to model an endotoxin source. 30% of particles were initialized with an endotoxin source with a rate drawn from an exponential distribution. The modified PF was able to estimate a non-zero endotoxin steady state, but bias was not eliminated in all cases. Further study is required to develop an unbiased observer.

#### ACKNOWLEDGEMENTS

Experimental data for model calibration was generated in the laboratory of our collaborator, Dr. John Kellum, Critical Care Medicine, University of Pittsburgh School of Medicine. Cytokine adsorption characteristics of the HA device were obtained from the laboratory of Dr. Bill Federspiel, Dept. of Bioengineering, University of Pittsburgh. Funding for this work was provided by NIH grant HL-76157; JSH also received funding through NIH T32 EB009403, part of the HHMI-NIBIB Interfaces Initiative.

#### REFERENCES

Abraham, E. (2003). Neutrophils and acute lung injury. *Crit Care Med*, 31(4 Suppl), S195-S199.

- Angus, D.C., Linde-Zwirble, W.T., Lidicker, J., Clermont, G., Carcillo, J., and Pinsky, M.R. (2001). Epidemiology of severe sepsis in the united states: analysis of incidence, outcome, and associated costs of care. *Critical Care Medicine*, 29(7), 1303–1310.
- Bishop, C.M. (2006). *Pattern Recognition and Machine Learning*. Springer, New York, NY.
- Bone, R.C., Balk, R.A., Cerra, F.B., et al. (1992). Definitions for sepsis and organ failure and guidelines for the use of innovative therapies in sepsis. *Chest*, 101(6), 1644–1655.
- Boyle, M.D.P., Hess, J.L., Nuara, A.A., and Buller, R.M. (2006). Application of immunoproteomics to rapid cytokine detection. *Methods (San Diego, Calif.)*, 38(4), 342–350.
- Chow, C.C., Clermont, G., Kumar, R., et al. (2005). The acute inflammatory response in diverse shock states. *Shock (Augusta, Ga.)*, 24(1), 74–84.
- Clermont, G., Bartels, J., Kumar, R., et al. (2004). In silico design of clinical trials: a method coming of age. *Crit Care Med.*, 32, 2061–2070.
- Cohen, J. (2002). The immunopathogenesis of sepsis. *Nature*, 420(6917), 885–891.
- Daun, S., Rubin, J., Vodovotz, Y., Roy, A., Parker, R., and Clermont, G. (2008). An ensemble of models of the acute inflammatory response to bacterial lipopolysaccharide in rats: Results from parameter space reduction. *Journal of Theoretical Biology*, 253(4), 843–853.
- Day, J., Rubin, J., Vodovotz, Y., Chow, C.C., Reynolds, A., and Clermont, G. (2006). A reduced mathematical model of the acute inflammatory response II. *Journal of Theoretical Biology*, 242(1), 237–256.
- Dellinger, R., Levy, M., Carlet, J., et al. (2008). Surviving sepsis campaign: International guidelines for management of severe sepsis and septic shock: 2008. *Intensive Care Medicine*, 34(1), 17–60.
- DiLeo, M., Kellum, J., and Federspiel, W. (2009). A simple mathematical model of cytokine capture using a hemoabsorption device. *Annals of Biomedical Engineering*, 37(1), 222–229.
- Fujiwara, N., and Kobayashi, K. (2005). Macrophages in inflammation. *Current Drug Targets - Inflammation & Allergy*, 4, 281–286.
- Janeway, C.A., J. and Medzhitov, R. (2002). Innate immune recognition. *Annu.Rev.Immunol.*, 20, 197–216.
- Kellum, J.A., Song, M., and Venkataraman, R. (2004). Hemoabsorption removes tumor necrosis factor, interleukin-6, and interleukin-10, reduces nuclear factor-kappaB DNA binding, and improves short-term survival in lethal endotoxemia. *Critical Care Medicine*, 32(3), 801–805.
- Kumar, R., Chow, C., Bartels, J., Clermont, G., and Vodovotz, Y. (2008). A mathematical simulation of the inflammatory response to anthrax infection. *Shock*, 29(1), 104–111.
- Kumar, R., Clermont, G., Vodovotz, Y., and Chow, C.C. (2004). The dynamics of acute inflammation. *Journal of Theoretical Biology*, 230(2), 145–155.
- Medzhitov, R. and Janeway, C., J. (2000). Innate immunity. *N.Engl.J.Med.*, 343(5), 338–344.
- Muske, K.R. and Rawlings, J.B. (1993). Model predictive control with linear models. *AIChE Journal*, 39(2), 262–287.
- Parker, R.S. and Doyle III, F.J. (2001). Control-relevant modeling in drug delivery. *Adv. Drug Deliv. Rev.*, 48, 211–228.
- Peng, Z., Carter, M.J., and Kellum, J.A. (2008). Effects of hemoabsorption on cytokine removal and short-term survival in septic rats. *Critical Care Medicine*, 36(5), 1573–1577.
- Rawlings, J.B. and Bakshi, B.R. (2006). Particle filtering and moving horizon estimation. *Computers & Chemical Engineering*, 30(10-12), 1529–1541.
- Reynolds, A., Rubin, J., Clermont, G., et al. (2006). A reduced mathematical model of the acute inflammatory response: I. *Journal of Theoretical Biology*, 242(1), 220–236.
- Singh, S. and Evans, T. (2006). Organ dysfunction during sepsis. *Intensive Care Medicine*, 32(3), 349–360.
- Vincent, J., Sun, Q., and Dubois, M. (2002). Clinical trials of immunomodulatory therapies in severe sepsis and septic shock. *Clinical Infectious Diseases*, 34(8), 1084–1093.

## Appendix A. ENDOTOXEMIA MODEL EQUATIONS

The deterministic endotoxemia model is defined by the equations below. The control variables are  $u_{il6}$ ,  $u_{tnf}$ , and  $u_{il10}$ , which depend on the HA column configuration.

$$\begin{aligned} \dot{P}E &= -d_{PE} \cdot PE \\ \dot{N} &= \frac{k_N \cdot R(PE, D)}{x_N + R(PE, D)} - d_{N}N \\ \dot{D} &= \frac{k_D \cdot N^6}{x_{D,N}^6 + N^6} - d_D D \\ \dot{C}_A &= s_{C_A} + k_{C_A}N - d_{C_A}C_A \\ \dot{I}L6 &= \frac{k_{il6} \cdot N^4}{x_{il6,N}^4 + N^4} \cdot f_{up}^{il6}(IL6, TNF) \cdot f_{down}^{il6}(IL10, C_A) \\ &\quad - (d_{il6} + u_{il6}) \cdot IL6 \\ \dot{T}NF &= k_{tnf}N^{\frac{3}{2}} \cdot f_{up}^{tnf}(TNF) \cdot f_{down}^{tnf}(IL6, IL10, C_A) \\ &\quad - (d_{tnf} + u_{tnf}) \cdot TNF \\ \dot{I}L10 &= s_{il10} + \frac{k_{il10} \cdot N^3}{x_{il10,N}^3 + N^3} \cdot f_{up}^{il10}(IL6, TNF) \\ &\quad + Y_{il10} - (d_{il10} \cdot f_{down}^{il10}(IL10) + u_{il10}) \cdot IL10 \\ \dot{Y}_{il10} &= \frac{k_{Y_{il10}} \cdot D^4}{x_{Y_{il10},D}^4 + D^4} - d_{Y_{il10}}Y_{il10} \end{aligned}$$

where:

$$R(PE, D) = (k_{N,pe}PE + k_{N,D}D) \cdot f_{up}^N(TNF, IL6) \cdot f_{down}^N(CA, IL10)$$

and regulatory functions  $f_{up}$  and  $f_{down}$  are defined:

$$\begin{aligned} f_{up}^w(z_1, \dots, z_n) &= 1 + \sum_{i=1}^n \frac{k_{w,z_i} Z_i^{h_{w,z_i}}}{(x_{w,z_i})^{h_{w,z_i}} + z_i^{h_{w,z_i}}} \\ f_{down}^w(z_1, \dots, z_n) &= \prod_{i=1}^n \frac{(x_{w,z_i})^{h_{w,z_i}}}{(x_{w,z_i})^{h_{w,z_i}} + z_i^{h_{w,z_i}}} \end{aligned}$$