A Data Mining Approach for Optimization of Acute Inflammation Therapy

Vladan Radosavljevic*
Temple University
Philadelphia, USA
vladan@temple.edu

Kosta Ristovski*
Temple University
Philadelphia, USA
kosta@temple.edu

Zoran Obradovic
Temple University
Philadelphia, USA
zoran.obradovic@temple.edu

***Abstract*—Acute inflammation is a medical condition which occurs over seconds, minutes or hours and is characterized as a systemic inflammatory response to an infection. Delaying treatment by only one hour decreases patient chance of survival by about 7%. Therefore, there is a critical need for tools that can aid therapy optimization for this potentially fatal condition. Towards this objective we developed a data driven approach for therapy optimization where a predictive model for patients' behavior is learned directly from historical data. As such, the predictive model is incorporated into a model predictive control optimization algorithm to find optimal therapy, which will lead the patient to a healthy state. To save on the cost of clinical trials and potential failure, we evaluated our model on a population of virtual patients capable of emulating the inflammatory response. Patients are treated with two drugs for which dosage and timing are critical for the outcome of the treatment. Our results show significant improvement in percentage of healthy outcomes comparing to previously proposed methods for acute inflammation treatment found in literature and in clinical practice. In particular, application of our method rescued 87% of patients that would otherwise die within 168 hours due to septic or aseptic state. In contrast, the best method from literature rescued only 73% of patients.**

***Keywords*-data mining; therapy optimization; acute inflammation.**

I. Introduction

One of the challenging problems in clinical practice is planning of individualized therapy regimens and durations. The state-of-the-art approach for therapy planning usually follows general rules that are common for all patients. Due to patients' variability, such an approach often leads to sub-optimal treatment that impacts therapy efficacy, toxicity, and patient outcome.

Planning of optimal therapy is especially critical in rapid progression medical conditions like acute inflammation, a systematic inflammatory response syndrome triggered by infection. Acute inflammation is often diagnosed too late and the patient is then treated with broad-spectrum antibiotics and/or intravenous fluids with dosages adjusted manually, even though more specific therapy would be far more effective. Inadequate treatment results in a mortality rate of 30-35%, and for every hour that the administration of appropriate therapy is delayed, the mortality rate increases by about 7% [12]. Therefore, there is a critical need for tools that can aid clinicians in designing optimal strategies for inflammation treatments.

One of the tools that can be used for finding optimal therapy is model predictive control (MPC). MPC algorithms are suitable for medical applications [3], [5], [8], [9] primarily because of their ability to deal with a large number of variables and constraints that are not well addressed with other approaches. The key component of MPC is a predictive model, which is used to predict a patient's future states as a response to the treatment. MPC uses predictions of future states and a set of constraints to calculate an optimal set of future treatments that will guide the patient to a healthy state. The quality of MPC directly depends on the ability of the predictive model to accurately predict the future states. Deterministic predictive models that rely on domain knowledge assumptions often fail when dealing with complex biological processes like sepsis. On the other hand, the hypothesis of our study is that data-driven models can learn patient responses directly from historical data without making any domain-based assumptions.

Data-driven models have been successfully utilized as predictive models in MPC for other medical applications, including the regulation of glucose supply [13], an exploration of optimal dosing of anticancer agents [7], the regulation of mechanical ventilation [11], and defining an optimal anesthesia [14].

Model predictive control for acute inflammation treatment was introduced in [2], where a deterministic predictive model was deployed to find optimal therapy. It was shown that using two drugs in MPC for acute inflammation treatment significantly increased the probability of success of the treatment. On the other hand, treatment success strongly depended on parameter settings in the predictive model.

We developed a data-driven model predictive control approach which uses two control signals (two drugs) for optimization of therapy and overcomes issues of the deterministic model from [2]. As in [2], all evaluations are performed on virtual patients, which is common practice for biomedical control research. Our results show significant improvement in percentage of healthy outcomes comparing to other methods for acute inflammation treatment found in literature and in clinical practice.

*These authors contributed equally

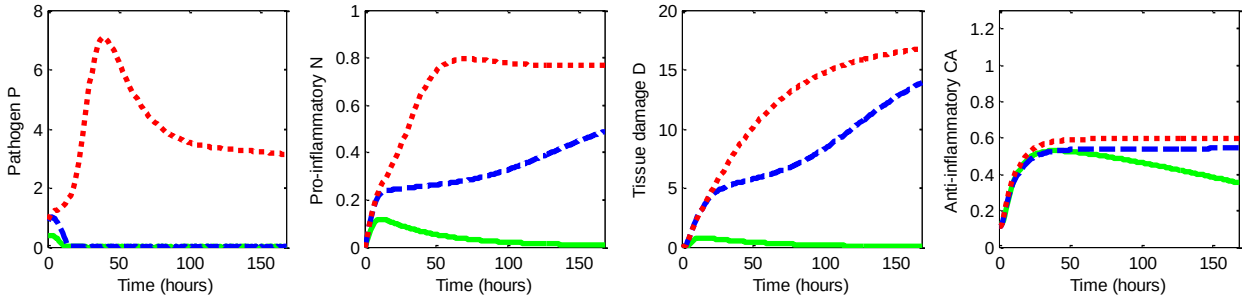

Figure 1. Evolution of pathogen population (P), pro-inflammatory mediators (N), tissue damage (D), and anti-inflammatory mediators (CA) of three virtual patients with healthy (green/solid), aseptic (blue/dashed), and septic (red/dotted) outcomes in the absence of therapy.

II. Virtual Patient Model and Treatment Constraints

To signifcantly reduce the chance of a clinical failure and to save on the costs of clinical trials, biomedical researchers use computer simulations of body processes (often called virtual patients) to perform preliminary tests of hypotheses before they prove them in real patient studies. Virtual patients are generated using a carefully determined mathematical model to simulate the process of interest. A significant advantage of having a virtual patient model for experiments is the possibility of testing different approaches for finding optimal therapies on the same virtual patient and comparing the outcomes. In order to follow a real-life scenario, virtual patient models are accompanied with well-defined constraints in therapy that are in accordance with clinical practice.

A. Patient model

The mathematical model for inflammatory response to severe infection is derived in [10]. This model has not incorporated drug effect on inflammatory response and it was not applicable for acute inflammation treatment. We will use a slightly modified mathematical model recently proposed in [2] that is capable of simulating:

- an evolution of a bacterial pathogen population (P) that initiates the cascade of inflammation,
- dynamics of early pro-inflammatory mediators (N),
- markers of tissue damage/dysfunction (D),
- the evolution of anti-inflammatory mediators (CA),

which are controlled by doses of pro-inflammatory ($PIDOSE$) and anti-inflammatory ($AIDOSE$) therapies. This mathematical model is based on the system of ordinary differential equations (ODE)

$$\begin{aligned}\frac{dP}{dt} &= k_{pg}\left(1-\frac{P}{P_\infty}\right)-\frac{k_{pm}s_mP}{\mu_m+k_{mp}P}- \\ &- k_{pn}f(N)P, \qquad (1)\\ \frac{dN}{dt} &= \frac{s_{nr}R}{\mu_{nr}+R}-\mu N+PIDOSE(t), \qquad (2)\end{aligned}$$

$$\begin{aligned}\frac{dD}{dt} &= \frac{k_{dn}f(N)^6}{x_{dn}^6+f(N)^6}-\mu_dD, \qquad (3)\\ \frac{dCA}{dt} &= s_c+\frac{k_{cn}f(N+k_{cnd}D)}{1+f(N+k_{cnd}D)}-\mu_cCA+ \\ &+ AIDOSE(t), \qquad (4)\end{aligned}$$

where

$$R=f(k_{np}P+k_{nn}N+k_{nd}D),\, f(x)=\frac{x}{1+\left(\frac{CA}{c_\infty}\right)^2}. \qquad (5)$$

All variables used in the mathematical model except output signals P, N, D, CA are parameters with valid ranges specified in [2]. Virtual patients were generated by random initialization of three parameters in ODE (k_{pg}, k_{cn}, and k_{nd}) and by random initialization of the initial conditions for P and CA from uniform distribution on valid ranges. Variability in the population of virtual patients is obtained by changing parameter values and initial conditions. All other parameters were fixed to referent values as in [2]. In all of the simulations, t is an hourly step that starts from 0. At $t = 0$, outputs and parameters are set to initial conditions. Then, all four outputs evolve according to ODE through the simulation time of 168 hours (one week). According to [2] there are three possible outcomes at the end of simulation time, which are shown in Figure 1. A patient is in healthy state if $P = 0$, $N = 0$, $D = 0$, and $CA > 0$ at the end of simulation. The aseptic death state of the patient is defined as $P = 0$, $N > 0$, $D > 0$, and $CA > 0$. The third possible outcome is septic death, where all outputs are non-zero. Although conceptual, ODE is capable of modeling the complex effect of pathogen P on the patient. From ODE, a large P leads to the development of a pro-inflammatory response N. A large N indicates faster elimination of pathogen P. However, a large N damages tissue D and therefore mobilizes a negative feedback, or anti-inflammatory response CA, which lowers N [2]. Also, CA inhibits damage to tissue D that may be caused by N. Leading the patient to a healthy state is a difficult challenge of maintaining a balance between objectives $P = 0$ and $D = 0$ (if these conditions are satisfied N will eventually

be 0 so there is no need to have $N = 0$ as an objective). An emphasis on minimizing D by increasing CA (with $AIDOSE$) might lead to unrestricted pathogen P growth. On the other hand, an emphasis on minimizing pathogen P by increasing N (with $PIDOSE$) might lead to a pro-inflammatory response aimed to eliminate pathogen P as soon as possible, after which it might be too late to control the tissue damage D. Both therapy timing and therapy dosage are critical for finding optimal treatment.

B. Treatment constraints

Well-designed constraints on drug dosage levels would make model predictive control applied to inflammation therapy to be close to the clinical practice. We know that giving large amounts of a drug at once can cause patients death due to overdose. It is also known from clinical practice that keeping anti-inflammatory doses at a high level for a longer period of time may play a large role in predisposing patients to secondary infections that might lead to death. We followed [2] to formulate constraints:

- $0 \leq PIDOSE \leq PIDOSE_k^{MAX}$ - the difference between the current level of $N = N_k$ and $N_{max} = 0.5$,
- $0 \leq AIDOSE \leq AIDOSE_k^{MAX}$ - the difference between the current level of $CA = CA_k$ and a maximum allowable level of CA, initialized to $CA_{max} = 0.6264$,
- *saturation of anti-inflammatory mediator*: the situation with CA saturated for very long time would be avoided clinically for fear of secondary infections that compromise organ recovery. Therefore, if the level of CA remained consistently elevated for more than 48 hours, the CA_{max} was reduced by half.

III. Model Predictive Control

The objective of model predictive control in acute inflammation treatment is to compute optimal values of drug doses that will probably lead the patient to a healthy state over time. In order to do that, we need to define reference trajectory (healthy state) such that minimizing the difference between estimated future patient state (output of predictive model) and reference trajectory might lead to success in treatment. We will use the same reference trajectory as in [2] which is defined as $D = 0$ and $P = 0$. Suppose that we want to find optimal therapy at time point k. We need to define following terms:

- *Prediction horizon* is a time window of length p which dictates how far in the future we want to predict.
- *Predicted patient states* are the future patient states $\{(\hat{P}, \hat{N}, \hat{D}, \hat{CA})_{k+j}, j = 1, \ldots, p\}$ predicted using a predictive model which simulates patient behavior.
- *Control horizon* is a time window of length $c < p$ which dictates the number of future control signals to be found.
- *Future control signals* $\{(\hat{PIDOSE}, \hat{AIDOSE})_{k+j}, j = 0, \ldots, c-1\}$ are determined by the optimization algorithm. We consider that future control signals from $k + c$ to $k + p - 1$ stay constant having the value $(\hat{PIDOSE}, \hat{AIDOSE})_{k+c-1}$.

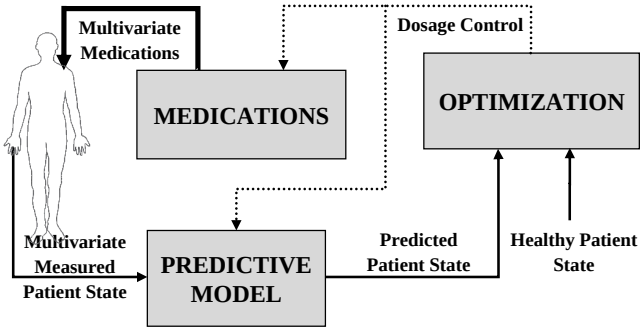

Figure 2. Model predictive control scheme for optimal therapy

The objective of therapy is to find a sequence of control signals from k to $k+c-1$ such that it satisfies all constraints and also minimizes the difference between predicted patient's states and reference trajectory over prediction horizon

$$\min_{AIDOSE,PIDOSE} \sum_{j=1}^{p} \left(w_D \hat{D}_{k+j}^2 + w_P \hat{P}_{k+j}^2 \right) + \\ + w_c \sum_{j=1}^{p} \left(\hat{PIDOSE}_{k+j-1}^2 + \hat{AIDOSE}_{k+j-1}^2 \right), \quad (6)$$

subject to

$$0 \leq \hat{AIDOSE}_{k+j-1} \leq AIDOSE_{k+j-1}^{MAX},$$
$$0 \leq \hat{PIDOSE}_{k+j-1} \leq PIDOSE_{k+j-1}^{MAX},$$

where w_D, w_P, and w_c are weighting constants that have to be chosen in advance.

We apply to the patient just the first term $(\hat{AIDOSE}, \hat{PIDOSE})_k$ from the control sequence obtained by minimization, observe the new state at time point $k+1$, and repeat the optimization procedure to obtain new control sequences. The new state at time point $k+1$ is calculated as the output of the patient model since the state at time point k is observed and control signals are found by optimization. MPC scheme is illustrated in Figure 2.

A. Predictive Model

As we emphasized in the introduction, the quality of an MPC directly depends on the ability of the predictive model to accurately predict patients' future states. The predictive model has to be carefully designed in order to be able to handle high variability in patients' behaviors, which is the greatest challenge in this application. When using data-driven predictive models there are typically two steps involved in predictive model design: 1) defining model structure and 2) inferring model parameters from the set of training data. Each of these steps we will consider separately. The resulting model will be named *LearnedMPC*.

1) Structure of LearnedMPC predictive model: Virtual patient state is represented by the four outputs P, N, D, and CA. The goal of the predictive model is to accurately predict future states using past states along with past control signals as inputs. To predict four-dimensional state with a single model is not feasible in this application because of the complex interactions among P, N, D, and CA. Instead, we propose splitting the predictive model into four sub-models, each of which is responsible for prediction of one of the outputs P, N, D, and CA, keeping the same set of inputs for each of the sub-models. If we denote $\mathbf{y}_k$ and $\mathbf{u}_k$ to represent patient's state and control signal at time point k respectively, it can be written

$$\mathbf{y}_k = (P_k, N_k, D_k, CA_k)^T, \quad (7)$$

$$\mathbf{u}_k = (AIDOSE_k, PIDOSE_k)^T. \quad (8)$$

Having observed patients' states and control signals up to time point k, the sub-model responsible for predicting $\hat{P}_{k+j}$ over prediction horizon can be represented by

$$\hat{P}_{k+j} = F_P(\hat{\mathbf{y}}_{k+j-1}, \ldots, \hat{\mathbf{y}}_{k+1}, \mathbf{y}_k, \ldots, \mathbf{y}_{k+j-n_{py}}, \hat{\mathbf{u}}_{k+j-1}, \ldots, \hat{\mathbf{u}}_k, \mathbf{u}_{k-1}, \ldots, \mathbf{u}_{k+j-n_{pu}}, \boldsymbol{\beta}_P), \quad (9)$$

where $j = 1, \ldots, p$. F_P is a function with unknown parameters $\boldsymbol{\beta}_P$ that models the input-output relation; n_{py} and n_{pu} are time lags for state and control signals respectively. Sub-models responsible for prediction of $\hat{N}_{k+j}$, $\hat{D}_{k+j}$, and $\hat{CA}_{k+j}$ have similar functional form but each of them has its own function F_N, F_D, F_{CA} and parameters $\boldsymbol{\beta}_N$, $\boldsymbol{\beta}_D$, $\boldsymbol{\beta}_{CA}$ respectively.

The next step in designing the predictive model is definition of input-output functional relations F_P, F_N, F_D, or F_{CA} in terms of model parameters. Predictive models based on machine learning tools such as neural networks [6] or Gaussian processes [4] have been successfully used in industrial applications. However, developing a non-linear model is a very difficult task [1]. Due to the lack of a superposition principle for non-linear models, the amount of data required to train a non-linear model is much larger than that for a linear model. If the process is multivariate, the difference in the amount of data required is even larger. Having in mind a realistic scenario of limited availability and high cost of clinical data, we used linear functional forms for F_P, F_N, F_D, or F_{CA}.

2) Training of LearnedMPC predictive model: The critical aspect of the predictive model design is the availability of representative training data to learn unknown model parameters $\boldsymbol{\beta}_N$, $\boldsymbol{\beta}_N$, $\boldsymbol{\beta}_D$, and $\boldsymbol{\beta}_{CA}$. Our objective was to address a real-life scenario in which data available for training of the predictive model come from clinical trials done on a small group of diverse patients observed in time. Accordingly, a small set of $N_{training}$ virtual patients with hourly observations for one week (168 hours) was generated from ODE equations. To generate a sequence of observations for a virtual patient we need to know model parameters, initial conditions and a control sequence. Initial conditions and parameters are randomly generated following allowable ranges while a control sequence was carefully chosen.

In industrial applications, control sequences are usually generated randomly so that they span the whole operational range and adequately characterize the response of the system. Random generation of treatments is not a clinically relevant scenario. Instead we propose the following approach. For each of the $N_{training}$ virtual patients we used its own mathematical model as a predictive model in MPC. Such predictive models give perfect prediction of the patients' future states, as their predictions are identical to future observations at every time point (ideal predictive models). Ideal control sequences for each virtual patient would be obtained by minimizing objective function (6). Ideal control sequences are not realistic in clinical practice and are also unsuitable for learning data-driven models because they do not contain enough dynamics to sufficiently represent the system response. Therefore, we used a more clinically realistic scenario, such that for an observed patient's state at time point k, control signals at k may not be ideal but they should be reasonably close to ideal. This is modelled such that at each time point k random Gaussian noise is added to $AIDOSE$ and $PIDOSE$ values found by the MPC with the ideal predictive control. Then, instead of treating patients with the ideal control sequences we treated patients with non-ideal ones, which gave a wider range of system response.

IV. Evaluation

A. Model order of LearnedMPC

We need to determine the number of past signals (model order) used as inputs to the predictive model. The predictive model, which uses time lag set to 1, has not provided satisfactory results in terms of predictive accuracies. It was not able to distinguish between the beginning and the end of the state sequence, which also resulted in simulations that ended in unhealthy states. On the other hand, the predictive model built with time lag set to 2 had much better predictive power so we chose the model order of 2 to be applied to all signals involved.

B. Number of patients to train LearnedMPC

We used a population of virtual patients with hourly measurements of their states and control signals observed for one week (168 hours). It is important to emphasize that for learning our data driven predictive model we do not use any knowledge about the patients models used to generate data. Our training data set contains only virtual patients' states and applied control signals in each state.

We noticed that if we applied non ideal control sequence (see Section III-A2) for an entire week, many patients would end up in a non-healthy state. Instead, we applied the noisy

Table I
AVERAGE PERFORMANCE OF *LearnedMPC* ON VALIDATION SET OF 50 PATIENTS FOR WELL BALANCED TRAINING SEST OF DIFFERENT SIZE.

$N_{training}$	Healthy (%)	Aseptic (%)	Septic (%)
3	48.2 ± 19.4	25.4 ± 18.0	26.4 ± 14.8
6	53.2 ± 24.1	20.0 ± 11.6	26.8 ± 18.4
9	75.6 ± 22.3	14.8 ± 9.4	9.6 ± 16.2
12	83.4 ± 7.0	12.0 ± 4.1	4.6 ± 9.3
15	83.4 ± 6.4	11.6 ± 4.6	5.0 ± 10.1
18	88.3 ± 2.3	11.0 ± 1.4	0.63 ± 2.0
21	88.6 ± 2.5	10.6 ± 1.0	0.8 ± 2.5
24	87.2 ± 1.68	11.8 ± 1.8	1.0 ± 2.5

control strategy just in the most critical first 10 hours of treatment. After 10 hours we continued with the ideal control strategy.

From the clinical aspect, where data are limited and expensive, it is important to analyze the minimum number of patients required in the training set that allows the learned predictive model to achieve satisfactory performance. To investigate how the performance of the learned model changes with respect to training data size we performed the following analysis. We created a pool of 65 virtual patients along with the corresponding non-ideal control sequences for each of them. Among 65 patients there were 8 with septic, 14 with aseptic, and 43 with healthy outcomes. We selected $N_{training}$ patients from this pool. In addition, we assumed that the selected set of patients was well balanced, so there were equal numbers of septic, aseptic, and healthy patients. Performances of MPCs with predictive models trained on $N_{training}$ patients were evaluated for different values of $N_{training}$ on a validation set of 50 patients created independently from the training data. We performed analysis for $N_{training} = 3, 6, 9, \ldots, 24$. For each $N_{training}$ we repeated experiments 10 times and reported results as the percentage of septic, aseptic, and healthy outcomes on average along with standard deviation. Results are reported in Table I. Models learned on fewer patients have less success in treatment. Stable performance with a high percentage of healthy outcomes is achievable starting from 18 patients in the training set.

C. *Parameters of* LearnedMPC

Based on results reported at Table I we chose a predictive model trained on 18 patients as a predictive model in MPC. The length of the prediction horizon is chosen as a balance between two opposite requirements. The prediction horizon should be long enough to capture the total effect of therapy. On the other hand predictive model is not reliable over too long prediction horizon because of accumulated prediction error. We found optimal prediction horizon to be 5. The control horizon was set to 2, which resulted in less aggressive therapies. Larger values of control horizon give faster response but the system is less robust to predictive model uncertainty [1]. To compensate for different scaling objective functions, weights w_P, w_D, and w_c were set to 3, 1, and 1 respectively.

D. *Dataset*

To evaluate the proposed *LearnedMPC* we generated a population of $N_{test} = 500$ virtual patients by randomly choosing parameters of the mathematical model. Patient state was classified based on the values of the outputs at the end of the simulation time of 168 hours. The resulting population consisted of 298 virtual patients classified as healthy, 117 classified as aseptic, and 85 classified as septic without any treatment applied. Among them, 321 patients were selected to receive therapy according to the criterion that their value of N exceeded 0.05 at some point [2].

E. *Baseline Methods*

We compared our *LearnedMPC* to models in [2]:

- no treatment model (*Placebo*),
- the therapy that resembles the one currently used in the intensive care units: a consistent dosing of an anti-inflammatory therapy; we implemented this therapy by giving a small $AIDOSE = 0.005$ each hour over a period of 72 hours, after which therapy terminates (*Static*),
- MPC with a predictive model set to the ODE with parameters from a single patient [2] (*Mismatch*).

F. *Results*

Results on patients who received therapy are reported in Table II. Along with percentage of healthy, aseptic and septic outcomes we differentiate two possible effects of the therapy on the patient: 1) harmed - outcome without therapy was healthy whereas outcome after treatment was aseptic/septic; 2) rescued - outcome without therapy was either aseptic or septic whereas outcome after treatment was healthy.

LearnedMPC achieved better accuracy than all alternatives in terms of all accuracy measures while at the same time keeping the number of harmed patients to 0. A therapy profile for the rescued virtual patient is illustrated in Figure 3. Most of the successful therapies obtained by *LearnedMPC* exhibit a similar pattern: high level pro-inflammatory dose was applied at the early stage in order to reduce the level of pathogens, at which point *LearnedMPC* modulated anti-inflammation dose to alleviate inflammation and restore health.

V. CONCLUSION AND FUTURE WORK

We have developed a completely new approach for acute inflammation treatment based on the use of a machine learning technique. Results obtained from experiments conducted on virtual patients have undoubtedly shown that our method outperformed all existing alternatives over all accuracy measures. In addition, the treatment based on our method has not been harmful to healthy patients, which is very important

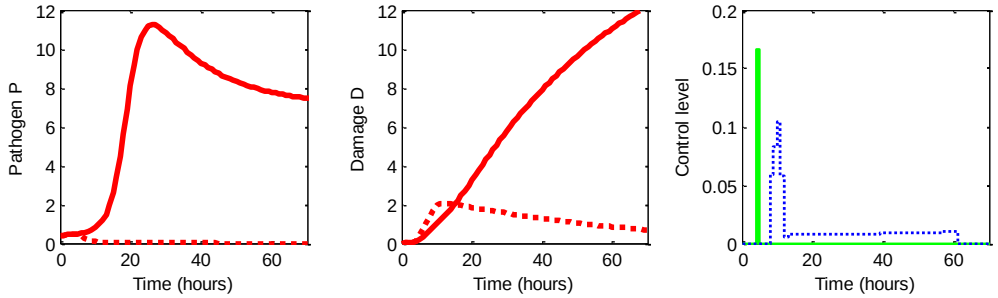

Figure 3. $LearnedMPC$ (dotted) vs. $Placebo$ (solid) on rescued patient: evolution of pathogen P and tissue damage D; control sequence of $PIDOSE$ (solid) and $AIDOSE$ (dotted) found by $LearnedMPC$.

Table II
DISTRIBUTIONS OF THE OUTCOMES OF PATIENTS WHO RECEIVED THERAPY FOR: NO TREATMENT MODEL (*Placebo*), CONSTANT AIDOSE (*Static*), MPC WITH PREDICTIVE MODEL SET TO ODE WITH PARAMETERS FROM A SINGLE PATIENT (*Mismatch*), AND OUR MPC APPROACH (*LearnedMPC*).

	Healthy	Aseptic	Septic	Harmed (out of 119)	Rescued (out of 202)
Placebo	119/321 (37.07%)	117/321 (36.45%)	85/321 (26.48%)	N/A	N/A
Static	140/321 (43.61%)	96/321 (29.91%)	85/321 (26.48%)	3/119 (2.52%)	24/202 (11.88%)
Mismatch	267/321 (83.18%)	50/321 (15.58%)	4/321 (1.25%)	0/119 (0%)	148/202 (73.27%)
LearnedMPC	294/321 (91.59%)	27/321 (8.41%)	0/321 (0%)	0/119 (0%)	175/202 (86.63%)

for clinical practice. We have shown that using a pool of 18 diverse patients was enough to make a general and efficient data-driven predictive model, which is consistent with clinical practice. The obtained results provide evidence that potential solutions for acute inflammation treatment can be based on the joint work of domain scientists and the machine learning community.

VI. ACKNOWLEDGEMENT

This work was funded, in part, by DARPA grant [DARPA-N66001-11-1-4183] negotiated by SSC Pacific grant.

REFERENCES

[1] Eduardo F. Camacho and Carlos Bordons, *Model predictive control*, Springer-Verlag, London, 2004.

[2] Gilles Clermont, Jonathan Rubin, and Judy Day, 'Using nonlinear model predictive control to find optimal therapeutic strategies to modulate inflammation', *Mathematical Biosciences and Engineering*, **7**(4), 739–763, (2010).

[3] Frank Doyle, Lois Jovanovič, and Dale Seborg, 'I. Glucose control strategies for treating type 1 diabetes mellitus', *Journal of Process Control*, **17**, 572–576, (2007).

[4] J. Kocijan, R. Murray-Smith, C.E. Rasmussen, and A. Girard, 'Gaussian process model based predictive control', in *American Control Conference, 2004. Proceedings of the 2004*, pp. 2214–2219. IEEE, (2004).

[5] YuanYuan Li, Scott C. Lenaghan, and Mingjun Zhang, 'A data-driven predictive approach for drug delivery using machine learning techniques', *PLoS ONE*, **7**(2), e31724, (2012).

[6] Hansen L.K. Nørgaard M., Ravn O., Poulsen N.K., *Neural networks for modeling and control of dynamic systems*, Springer, London, 2000.

[7] Sarah L Noble, Eric Sherer, Robert E Hannemann, Doraiswami Ramkrishna, Terry Vik, and Ann E Rundell, 'Using adaptive model predictive control to customize maintenance therapy chemotherapeutic dosing for childhood acute lymphoblastic leukemia.', *Journal of theoretical biology*, **264**(3), 990–1002, (2010).

[8] Robert S. Parker, 'II. Modeling for anti-cancer chemotherapy design', *Journal of Process Control*, **17**(7), 576–582, (2007).

[9] Joao Pinheiro and Joao Lemos, 'Multi-drug therapy design for HIV-1 infection using Nonlinear Model Predictive Control', in *Mediterraanean Conference on Control and Automation*, pp. 485–490, (2011).

[10] Angela Reynolds, Jonathan Rubin, Gilles Clermont, Judy Day, Yoram Vodovotz, and G Bard Ermentrout, 'A reduced mathematical model of the acute inflammatory response: I. Derivation of model and analysis of anti-inflammation.', *Journal of theoretical biology*, **242**(1), 220–36, (2006).

[11] F.T. Tehrani and J.H. Roum, 'Intelligent decision support systems for mechanical ventilation', *Artificial Intelligence in Medicine*, **44**(3), 171–182, (2008).

[12] Steven W Thiel, Jamie M Rosini, William Shannon, Joshua A Doherty, Scott T Micek, and Marin H Kollef, 'Early prediction of septic shock in hospitalized patients.', *Journal of hospital medicine : an official publication of the Society of Hospital Medicine*, **5**(1), 19–25, (2010).

[13] Youqing Wang, Howard Zisser, Eyal Dassau, Lois Jovanovic, and Francis J. Doyle, 'Model predictive control with learning-type set-point: Application to artificial pancreatic β-cell', *AIChE Journal*, **56**(6), 1510–1518, (2009).

[14] S. Yelneedi, S. Lakshminarayanan, and GP Rangaiah, 'A comparative study of three advanced controllers for the regulation of hypnosis', *Journal of Process Control*, **19**(9), 1458–1469, (2009).