

Model-Based Cardiovascular Parameter Estimation in the Intensive Care Unit

Z Samar, T Heldt, GC Verghese, RG Mark

Massachusetts Institute of Technology, Cambridge, MA, USA

Abstract

In this paper, we present a simulation study that aims at estimating parameters of a hemodynamic model using observable data typically available in an Intensive Care Unit (ICU). Tracking model parameters in time reveals disease progression, and hence can be very useful for patient monitoring purposes. However, the observable data is generally not rich enough to allow for reliable estimation of all parameters of the underlying model. This leads to an ‘ill-conditioned’ estimation problem. To overcome this ill-conditioning, we employ subset selection to identify the ‘well-conditioned’ parameters that can be estimated robustly. We attempt to estimate only these parameters while the rest are fixed at prior values. Our results indicate that focusing on the reduced-order estimation problem improves the reliability of the estimates by more than 50%; the scheme is capable of recovering the underlying well-conditioned parameters with reasonable accuracy in both steady-state and transient conditions.

1. Introduction

Modern ICUs are equipped with many patient monitoring devices, each continuously recording signals from the human body. Currently, these signals need to be interpreted by a clinician in order to assess the state of the patient, to formulate physiological hypotheses, and to determine treatment options. In some areas, such as hemodynamic monitoring, enough quantitative knowledge about the underlying physiological system is available to formulate computational models capable of simulating normal and abnormal human physiology. Through tuning, such models could be used to track patient states automatically thus relating properties of the observable data streams directly to the properties of the underlying cardiovascular system. Frequently, however, the observable signals are not rich enough to allow for accurate estimation of *all* the model parameters. This leads to an *ill-conditioned* estimation problem in which the parameter estimates become very sensitive to the noise in the data. To overcome the ill-

conditioning, we employ subset selection, a methodology that improves the conditioning of the estimation problem by reducing its dimensionality. Subset selection identifies those *well-conditioned* parameters that have strongly independent effects on the model output, and thus can be estimated robustly. Attempting to identify only the well-conditioned parameters while the ill-conditioned ones are fixed at prior values greatly improves the reliability of the reduced-order parameter estimates [1].

To evaluate the subset selection-based estimation algorithm in an ICU setting, we conducted a simulation study during both steady-state and transient conditions that used single-cycle waveforms of Arterial Blood Pressure (ABP), Central Venous Pressure (CVP), and Pulmonary Arterial Pressure (PAP) as observable data in a bid to recover parameters of an underlying model.

2. Cardiovascular model

Computational models of the cardiovascular system are conveniently represented in the form of their circuit analogs where pressure $P(t)$ maps to voltage, flow $\dot{q}(t)$ to current, and blood volume $Q(t)$ to charge. The computational model used in our studies is based on Davis’ work [2] and is shown in Figure 1. The model comprises six compartments representing the left and right ventricles (l, r), the systemic arteries and veins (a, v), and the pulmonary arteries and veins (pa, pv). The ventricles are modeled by time-varying compliances connected to inflow and outflow resistances; diodes represent cardiac valves. The rest of the compartments are each modeled by a linear capacitor coupled with a linear resistor. Short-term cardiovascular control was modeled in the form of the arterial baroreflex and implemented as a set-point controller that attempts to maintain mean arterial blood pressure by modifying peripheral resistance, ventricular contractility, venous tone, and heart rate. The implementation of the arterial baroreflex is based on Davis’ extension of deBoer’s work [2]. The appeal in using this model lies in its manageable complexity and its ability to represent short-term cardiovascular dynamics [2, 3].




Figure 1. Circuit analog of the cardiovascular model. P_{th} refers to the intrathoracic pressure.

3. Parameter estimation

3.1. Nonlinear least squares optimization

In the context of parameter estimation, nonlinear least squares optimization arrives at estimates for the parameters of an underlying model by iteratively minimizing the error between the model output and the experimental data. Let

$$\mathbf{r}(\boldsymbol{\theta}) = \hat{\mathbf{y}}(\boldsymbol{\theta}) - \mathbf{y}$$

denote the residual error, where $\hat{\mathbf{y}}(\boldsymbol{\theta}) \in \mathbb{R}^n$ corresponds to the model output, which is a function of the parameter vector $\boldsymbol{\theta} \in \mathbb{R}^m$, and where $\mathbf{y} \in \mathbb{R}^n$ refers to the observation (or data) vector. The cost function we try to minimize is the sum of squares of residual errors

$$\Phi(\boldsymbol{\theta}) = \frac{1}{2}(\mathbf{r}^T \mathbf{r}) \quad (1)$$

Approximating the cost function by its second-order Taylor series expansion around the current best estimate, $\boldsymbol{\theta}_i$, and attempting to minimize this approximation yields the following equation (see [4] for a detailed description and derivation):

$$\mathbf{H}(\boldsymbol{\theta})|_{\boldsymbol{\theta}_i} \cdot (\boldsymbol{\theta}_{i+1} - \boldsymbol{\theta}_i) = -\mathbf{J}(\boldsymbol{\theta})^T \mathbf{r}(\boldsymbol{\theta})|_{\boldsymbol{\theta}_i} \quad (2)$$

where $\mathbf{J} \in \mathbb{R}^{n \times m}$ is the Jacobian (or gradient) matrix of the error vector with respect to the parameter vector:

$$J_{ij} = \frac{\partial r_i(\boldsymbol{\theta})}{\partial \theta_j}$$

and $\mathbf{H} \in \mathbb{R}^{m \times m}$ is the Hessian matrix of second-order partial derivatives of the cost function with respect to the parameter vector:

$$H_{ij} = (\mathbf{J}^T \mathbf{J})_{ij} + \sum_{l=1}^n r_l \cdot \frac{\partial^2 r_l}{\partial \theta_i \partial \theta_j}$$

In the presence of small residuals, the Gauss-Newton approximation to the Hessian states that the terms containing the residuals can be ignored, thus giving the following approximation:

$$H_{ij} = (\mathbf{J}^T \mathbf{J})_{ij}$$

Substituting the Hessian matrix with its approximation in Equation 2 yields the following equation:

$$\mathbf{J}^T \mathbf{J} \cdot (\boldsymbol{\theta}_{i+1} - \boldsymbol{\theta}_i) = -\mathbf{J}^T \mathbf{r} \quad (3)$$

Equipped with an initial estimate $\boldsymbol{\theta}_0$, Equation 3 can be solved iteratively until a pre-defined exit criterion is obtained.

To illustrate the problem of ill-conditioning, we follow the reasoning presented by Burth and co-workers [1]. If the matrix \mathbf{J} is rank-deficient, it has at least one eigenvalue at zero; its column space does not span the entire \mathbb{R}^m space. Consequently, the parameter update vector can be arbitrarily varied in the direction of any $\boldsymbol{\vartheta}$ that belongs to the null-space of \mathbf{J} , without affecting the error criterion to first order:

$$\begin{aligned} \mathbf{J}^T \mathbf{J} \cdot (\boldsymbol{\theta}_{i+1} - \boldsymbol{\theta}_i + \boldsymbol{\vartheta}) &= \mathbf{J}^T \mathbf{J} \cdot (\boldsymbol{\theta}_{i+1} - \boldsymbol{\theta}_i) + \mathbf{J}^T \mathbf{J} \cdot \boldsymbol{\vartheta} \\ &= \mathbf{J}^T \mathbf{J} \cdot (\boldsymbol{\theta}_{i+1} - \boldsymbol{\theta}_i) \\ &= -\mathbf{J}^T \mathbf{r} \end{aligned}$$

Thus, if \mathbf{J} is singular, then the model parameters are not uniquely determinable from the available observation data; such an estimation problem is said to be over-parameterized.

Typically though, \mathbf{J} is not exactly singular, but nearly so, with its largest singular value much greater than its smallest. Nearness to singularity is measured by the condition number, $\kappa(\mathbf{J})$, which, for real and symmetric matrices, is given by the ratio of the largest to the smallest eigenvalues. This nearness to singularity gives an ill-conditioned problem, in which small numerical errors or noise in the underlying data can radically modify the solution. To overcome the problem of ill-conditioning, we turn to subset selection to determine which parameters should be discarded from the estimation formulation in order to improve the conditioning of the problem.

3.2. Subset selection

Conceptually, subset selection is a two-step process. The first step determines ρ , the number of parameters that

can be estimated robustly from the available data; the second step determines which ρ parameters can be identified. The number of well-conditioned parameters is determined from the structure of the Hessian eigenspectrum: if the Hessian eigenspectrum contains ρ large eigenvalues and $m - \rho$ small ones, then this indicates that the Hessian has a numerical rank of ρ , and that only ρ parameters should be included in the estimation formulation. A QR-decomposition with column pivoting or a singular value decomposition of the Hessian matrix can be used to determine which ρ parameters should be included in the reduced-order estimation problem [5]. The subsequent estimation problem involves the use of reduced-dimension Jacobian and Hessian matrices, indicated by \mathbf{J}_ρ and \mathbf{H}_ρ , respectively. \mathbf{J}_ρ contains ρ columns of the original Jacobian matrix that are strongly independent. This selection results in a small condition number for \mathbf{J}_ρ and \mathbf{H}_ρ .

When applied to the problem of estimating cardiovascular parameters using single-cycle waveforms of ABP, CVP, and PAP, the subset selection algorithm returned the following well-conditioned or “active” parameters:

1. Distending blood volume (DBV)
2. Peripheral resistance (R_a)
3. Right end-diastolic compliance (C_r^{ed})

4. Description of the estimation problem

As we are interested in judging the performance of the estimation algorithm, we must know the true values of the underlying parameters. We therefore used our computational model to produce synthetic target data, which were then treated as *measurements* to which we applied the estimation algorithm.

We attempted to estimate the well-conditioned parameters in both steady-state and transient conditions. For the estimation problem using steady-state data, we generated target data using randomized parameters: each parameter, θ_i , was perturbed using a Gaussian distribution $N(\theta_i^0, 10\%\theta_i^0)$, where θ_i^0 is the nominal parameter value, to simulate a target data. In an attempt to investigate the benefit of using subset selection, this target data was used in two different estimation schemes. In the first scheme, only the active parameters were estimated while the rest were fixed at their nominal values. In the second one, all the parameters were estimated, i.e. subset selection was not applied.

Next, we generated transient data by simulating several cases of hemorrhage each lasting 30 minutes. The data was generated using the same randomization scheme as outlined above. In addition, each waveform was perturbed with additive noise from a Gaussian distribution $\sim N(0, 1.7\%m_i)$, where m_i refers to the cycle-average of the waveform, in order to assess the robustness of the resultant parameter estimates.

5. Results and discussion

Table 1 summarizes the relative errors incurred in estimating the active parameters from steady-state waveform data under a full-order estimation scheme and a reduced-order subset selection-based estimation scheme. For the full-order scheme, 50% of the runs failed to converge, whereas all the trials converged for the subset selection-based scheme. Moreover, we observe that reducing the dimensionality of the estimation problem improves the mean reliability of the active parameter estimates from an average relative error of 13.8% to 6.9%. The average relative error incurred in estimating the non-active parameters under the full-order estimation scheme was 14.9%.

Table 1. Relative estimation error for the active parameters using steady-state waveform data.

Parameter	Estimation problem	
	full-order (mean \pm st. dev.)	reduced-order (mean \pm st. dev.)
DBV	(13.3 \pm 11.4)%	(6.4 \pm 4.8)%
R_a	(16.1 \pm 10.2)%	(7.4 \pm 5.8)%
C_r^{ed}	(15.3 \pm 10.7)%	(7.1 \pm 7.1)%

Table 2 summarizes the relative errors incurred in the estimation process using the hemorrhage data. The results show that the algorithm performs well in recovering the underlying active parameters from transient data with a mean estimation error of less than 8%. In Figure 2, we display

Table 2. Relative estimation error for the active parameters using transient data.

Parameter	(mean \pm st. dev.)
DBV	(4.1 \pm 3.0)%
R_a	(1.6 \pm 1.4)%
C_r^{ed}	(7.7 \pm 8.1)%

the individual estimation results for the hemorrhage data, where we have used a first-point calibration at the beginning of each hemorrhage simulation to obtain absolute values for the peripheral resistance R_a [3]. The correlation coefficients for the three active parameters are 0.98, 0.97, and 0.54 for ΔDBV , R_a , and C_r^{ed} , respectively.

The estimation errors that occur when recovering the reduced set of parameters are mainly due to the bias introduced by fixing the values of the ill-conditioned parameters. In an attempt to reduce the error between model output and observable data, the estimation algorithm distorts the active parameter estimates in order to compensate for the fixing of the ill-conditioned parameters.




Figure 2. Estimated vs. actual parameter values.

6. Conclusions

Our simulation results indicate that the use of subset selection to identify and estimate only the well-conditioned parameters significantly improves the reliability of the estimates. The algorithm was able to recover the active parameters reasonably well in both steady-state and transient conditions and therefore holds promise for tracking hemodynamically important parameters in an ICU setting. Future work will focus on the use of real patient data to evaluate the estimation algorithm.

Acknowledgements

This work was supported in part through the National Aeronautics and Space Administration (NASA) through the NASA Cooperative Agreement NCC 9-58 with the National Space Biomedical Research Institute and in part by Grant RO1 EB001659 from the National Institute of Biomedical Imaging and Bioengineering.

References

- [1] Burth M, Verghese G, Vélez-Reyes M. Subset selection for improved parameter estimation in on-line identification of a synchronous generator. *IEEE Transactions on Power Systems* 1999;14(1):218–225.
- [2] Davis T, Mark R. Teaching physiology through simulation of hemodynamics. *Computers in Cardiology* 1991;17:649–652.
- [3] Samar Z. Cardiovascular Parameter Estimation using a Computational Model. Master's thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, June 2005.
- [4] Bard Y. *Nonlinear Parameter Estimation*. Academic Press, 1974.
- [5] Golub G, Klema V, Stewart G. Rank degeneracy and least squares problems. Technical Report TR-456, Department of Computer Science, University of Maryland, College Park, MD, 1976.

Address for correspondence:

Roger G. Mark, MD, PhD
 77 Massachusetts Avenue, Room E25-505
 Cambridge, MA 02139, USA
 rgmark@mit.edu