Learning Cyber-Physical Models of Resuscitation

Mayra Melendez

CMU-RI-TR-20-62

November 20, 2020

The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA

Thesis Committee:
Dr. Artur Dubrawski, *chair*
Dr. Jeff Schneider
Nicholas Gisolfi

Submitted in partial fulfillment of the requirements for the degree of Master of Science - Research.

Abstract

Ability to predict outcomes and forecast trajectories of recovery from resuscitated intensive care patients could guide treatment decisions and improve outcomes of care in both clinical and field settings. We develop a machine learning driven cyber-physical model to provide such predictive capabilities by leveraging arterial blood pressure (ABP) waveforms, one of the routinely collected vital signs. A cohort of 51 Yorkshire pigs was subjected to induced slow rate hemorrhage followed by fluid resuscitation. To represent physics of the arterial system and emulate blood pressure dynamics, we combine a two-element Windkessel model with an Unscented Kalman Filter (UKF) to track the instantaneously estimated Windkessel parameters over time. As the arterial pressure waveform exponentially decays during diastole after each pump, we use UKF-tracked Windkessel parameter estimates to identify time windows of ABP waveforms taken from other subjects in the cohort to reconstruct the shapes of the test subject's ABP signal and its moving average. We allow UKF covariance to temporarily increase to account for the effects of treatment such as administering norepinephrine. When evaluated under leave-one-subject-out cross-validation protocol, the model stays within 14+/-5% (mean+/-standard deviation) of mean absolute percentage error when reconstructing the current 250Hz ABP waveforms, and 19+/-6%, 24+/-6%, and 25+/-6% when forecasting at 5, 15 and 30 minute horizons, respectively. Our results demonstrate feasibility of using cyber-physical modeling of hemodynamic waveform data to predict trajectories of resuscitation and therefore timely inform treatment of hemorrhagic patients in both clinical and prolonged field care settings. We also provide a few thoughts for future work, including potential improvements attainable by calibrating neighbor selections and model predictions to the individual subject baselines to improve handling heterogeneity of the subjects, on-line tracking of model performance to estimate confidence in its predictions in real-time, and forecasting eventual outcomes of resuscitation as well as time-to-recovery for those subjects who are likely to recover, to name a few such ideas.

Acknowledgments

I would like to give my greatest and most heartfelt thanks to Dr. Artur Dubrawski, who has shown me incredible kindness, guidance, patience, and of course the fundamentals of data science, and without whom this thesis would not have been possible, and without whom my experience at the Robotics Institute would have been far less enriching. I would also like to thank Dr. Kyle Miller and Anthony Wertz, who served as my mentors throughout my graduate studies, for their suggestions and interesting discussions as the subjects of my thesis changed and evolved. Dr. Michael Pinsky, Dr. Joo Heung Yoon, and Dr. Gilles Clermont have my very deep appreciation and thanks for patiently explaining medical terminology and giving me validation that my research could be interesting and meaningful to the medical community. I also thank everyone in the Auton Lab for their inspiring work and ever-present willingness to help others - too many of you answered quick questions from me to count.

This work was partially supported by the U.S. Army Medical Research and Development Command awards W81XWH-19-C-0083 and W81XWH-19-C-0101, National Institutes of Health award R01HL141916, and Defense Advanced Research Projects Agency award FA8750-17-2-0130.

Contents

When this dissertation is viewed as a PDF, the page header is a link to this Table of Contents.

List of Figures

List of Tables

Chapter 1

Introduction

The Windkessel model is a lumped model of the heart's arterial system. Based on simple differential equations, it uses electrical circuit analogs that have physiological meanings to express in simple, interpretable terms the complex system of blood flow, blood pressure, organ volume, and other components that dictate blood circulation through and around the heart. Simple models such as the arterial Windkessel have proven to be practical in tracking patients' health in clinical settings [11].

The most basic of the Windkessel models, the two-element model, represents the arterial system with a single resistance and a single compliance element, with the heart pumping blood acting as a current source. Its electric circuit analog is shown in Figure 1.1 and simulates the system's blood pressure dynamics. This model, when given a pulsed current, represents the arterial pressure waveform as an exponential decay in diastole, the period of the heartbeat in which the heart relaxes after a pump [12].

The Windkessel model parameters can be used as a means to determine values commonly used in clinical practice such as cardiac output, though their own physiological significance can be telling of the state of the cardiovascular system. Arterial resistance R represents resistance to blood flow in small blood vessels. It captures the relation between the pressure drop across a blood vessel and the flow going through it. Total peripheral resistance captures this feature across the entire vascular bed of blood vessels leaving the heart. Arterial compliance C relates the change in volume of a blood vessel for a given change in pressure. Smaller arteries are less compliant than

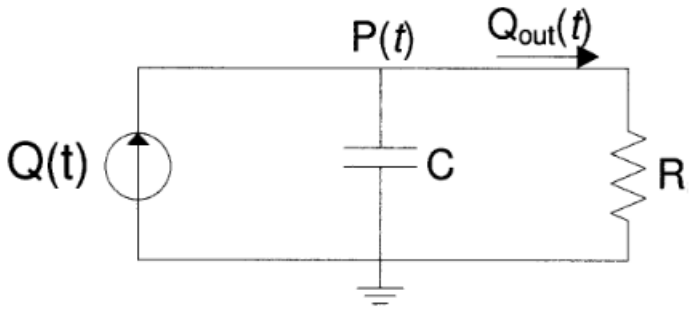

Figure 1.1: The two-element Windkessel model circuit analog. The heart acts as a current source producing cardiac output, with the larger arteries expanding as blood is pumped through them due to their compliance and the smaller vessels providing resistance to blood flow.

larger ones, and the majority of the system's compliance is found in the proximal aorta, the head vessels, and the upper limb vessels [8].

Different numerical methods have been used to solve for the Windkessel parameters. Methods such as subspace model identification [3, 6] and least squares [1] have been used to solve many variations of two-, three-, and four-element Windkessel models with offline data, but have struggled with tracking data in real-time. Kalman filters [9, 14], particularly the unscented Kalman filter, have had success with real-time tracking, both in regards to convergence time and value accuracy [4].

There exist many cardiac predictive modeling techniques using hemodynamic parameters such as the Windkessel parameters. Recently, deep neural networks have been trained on physics-based Windkessel systems and used to predict blood flow and pressure [7], among other parameters [13]. Other more traditional regression techniques, such as support vector and random forest regression, have also been used to calculate hemodynamic characteristics [2].

This thesis research sought to study the efficacy of using the two-element Windkessel model parameters as indicators of health similarity and to use these similarities to reconstruct the health trends of pigs recovering from severe blood loss. This paper will be organized as follows. Chapter 2 will outline the methods used for calculating the Windkessel parameters and using them for arterial pressure regression. Chapter 3 will present and discuss the results of this research. Chapter 4 will summarize the findings and discuss future work.

Chapter 2

Methods

The data used in this research was taken from 51 adult female Yorkshire pigs [10]. Each dataset was taken from a three-phase, several-hour-long experiment intended to obtain vital sign data before, during, and after induced hemorrhagic shock and the ensuing resuscitation. After undergoing surgery to install a pulmonary artery catheter and other sensors, each pig was rested for the first 30 minute phase to monitor its vitals and establish corresponding baseline readings. The second phase consisted of the bleed phase. 15 pigs were bled at 5 mL/min until they reached a mean arterial pressure of 40 mmHg, and the remaining pigs were bled at a faster rate of 20 mL/min until reaching a mean arterial pressure of 30 mmHg. The third and final phase was the resuscitation phase. The pigs were stabilized with resuscitation fluids and monitored as their vital signs returned to the healthy "normal" established in their respective baselines before being euthanized.

The original vital signs data was sampled at 250 Hz and contains the following attributes: time, arterial pressure, continuous cardiac output, central venous pressure, pulmonary artery pressure, arterial oxygen saturation, its corresponding plethysmographic waveform and electrocardiography tracings, mixed venous oxygen saturation, and the outside air pressure.

Accompanying the numerical data were annotations written by the physicians performing the experiments. These annotations contained information such as when different phases of the experiments began and ended, when fluids or medication doses were given, and when the pigs began to crash, among others.

The only vital signs that were studied as part of this research were arterial pressure and continuous cardiac output, as these were the minimum necessary for the Windkessel model described in the next section. Time and its accompanying annotations were used to monitor when different phases of the experiments began and ended and when medical interventions occurred. To reduce the size of the data, as each CSV file pertaining to each pig contains several hundred thousand entries, and remove some of the noise, the vital sign data was downsampled to the median of every 5 entries, resulting in a final sampling rate of 50 Hz. Time was downsampled to every fifth entry.

2.1 Two-Element Windkessel Model

The circuit in Figure 1.1 was solved to determine the characteristic equations for the two-element Windkessel model. Continuity of flow gives

$$\frac{dV(t)}{dt} = Q(t) - Q_{out}(t) \tag{2.1}$$

where $V(t)$ is the volume of blood flowing through the artery, $Q(t)$ is cardiac output and is acting as the current source, and $Q_{out}(t)$ is the flow through the arterioles, represented by the resistor. Considering the flow through the resistor R and its corresponding pressure drop $P(t)$ yields

$$Q_{out}(t) = \frac{P(t)}{R} \tag{2.2}$$

The total arterial compliance C is defined as the ratio of the arterial system's volume change to its pressure change,

$$C = \frac{dV(t)}{dP(t)} \tag{2.3}$$

Combining these above relations yields the governing equation for the current flow, the cardiac output, of the two-element Windkessel model:

$$Q(t) = C\frac{dP(t)}{dt} + \frac{P(t)}{R} \tag{2.4}$$

Equation 2.4 was then rearranged to solve for the rate of change of the arterial

pressure, giving

$$\frac{dP(t)}{dt} = \frac{Q(t) - \frac{P(t)}{R}}{C} \tag{2.5}$$

for each C and R value at time t. Lastly, using Euler's method to account for discretized data, the following update equation was formed:

$$P_{k+1} = P_k + dt\frac{Q_k - \frac{P_k}{R}}{C} \tag{2.6}$$

for each time step k and a constant sampling rate dt. Equation 2.6 served as the state equation for the Kalman filter detailed in the next section, with the arterial pressure and cardiac output being the input values for the characteristic equation of the two-element Windkessel model.

Three major assumptions were made in this research. Firstly, Windkessel models are typically used to model the heart, particularly its pressure, during diastole, the half of the heartbeat when the heart relaxes after having pumped out blood. However, as the purpose of the Kalman filter was to calculate the Windkessel parameters given the pressure and cardiac output over many heartbeats, the above relation was assumed to hold true during the entire heartbeat process to avoid having to isolate the data pertaining to the latter half of every heartbeat. Secondly, the data studied does not supply an instantaneous cardiac output measurement, instead giving an average estimate over a short period of time. Because no other more accurate measurement was available, this average value had to be used. Therefore, the cardiac output in the data was assumed constant across the entire heartbeat. This was also reflected in the state update for the Kalman filter below in Equation 2.6, as only the pressure value is updated using the two-element model. Thirdly, the two-element model requires aortic pressure, which was not available directly. However, it could be approximated by the arterial pressure, the closest measurement available, so this value was used instead.

2.2 Unscented Kalman Filter Design

A Kalman filter was used to calculate the Windkessel parameters at each time step in the baseline and resuscitation phases of the experiments because of its ability to

track and respond to changes in a dynamic system. This research used an Unscented Kalman filter (UKF) following the method proposed by Huang et al [4].

The UKF takes a state vector $\mathbf{x}_k$, observation vector $\mathbf{y}_k$, and input vector $\mathbf{u}_k$ at time k and models a dynamic system as:

$$\begin{aligned} \mathbf{x}_{k+1} &= f(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{w}_k \\ \mathbf{y}_k &= h(\mathbf{x}_k) + \mathbf{v}_k \end{aligned} \tag{2.7}$$

where $f(\cdot)$ and $h(\cdot)$ are the state and observation functions and $\mathbf{w}_k$ and $\mathbf{v}_k$ are the process and measurement noises, respectively. The objective of the filter is to estimate the state $\mathbf{x}_k$ given the observation $\mathbf{y}_k$.

The UKF requires $2n$ sigma points to approximate the distribution of the state vector as a random variable, where n is the total number of states to be estimated [5]. There are many ways to generate sigma points. The way used in this method uses a symmetric set unscented transform and is given by the following expressions:

$$\begin{aligned} s_{k-1}^i &= \hat{\mathbf{x}}_{k-1}^+ + \left(\sqrt{n\mathbf{P}_{k-1}^+}\right)_i^T, i = 1, ..., n \\ s_{k-1}^{n+i} &= \hat{\mathbf{x}}_{k-1}^+ - \left(\sqrt{n\mathbf{P}_{k-1}^+}\right)_i^T, i = 1, ..., n \end{aligned} \tag{2.8}$$

where $(\sqrt{\cdot})_i$ is the ith row of the matrix square root, and $\hat{\mathbf{x}}_{k-1}^+$ and $\mathbf{P}_{k-1}^+$ are the *a posteriori* state estimate and covariance at time $k-1$, respectively.

These sigma points are propagated through the state function, given above in Equation 2.6, averaged, and used to calculate the state covariance, giving,

$$\begin{aligned} s_k^i &= f(s_{k-1}^i, \mathbf{u}_k) \\ \hat{\mathbf{x}}_k^- &= \frac{1}{2n}\sum_{i=1}^{2n} s_k^i \\ \mathbf{P}_k^- &= \frac{1}{2n}\sum_{i=1}^{2n}(s_k^i - \hat{\mathbf{x}}_k^i)(s_k^i - \hat{\mathbf{x}}_k^i)^T + \mathbf{Q}_{k-1} \end{aligned} \tag{2.9}$$

where $\hat{\mathbf{x}}_k^i$ and $\mathbf{P}_k^-$ are the *a prior* state estimate and covariance, respectively, and $\mathbf{Q}_{k-1}$ is the covariance associated with the process noise.

Next, the predicted observation $\hat{\mathbf{y}}_k$, innovation covariance $\mathbf{P_y}$, and cross covariance $\mathbf{P_{xy}}$ are found using:

$$\begin{aligned}
\hat{\mathbf{y}}_k &= \frac{1}{2n}\sum_{i=1}^{2n} h(s_k^i) \\
\mathbf{P_y} &= \frac{1}{2n}\sum_{i=1}^{2n}(h(s_k^i) - \hat{\mathbf{y}}_k^i)(h(s_k^i) - \hat{\mathbf{y}}_k^i)^T + \mathbf{R}_k \\
\mathbf{P_{xy}} &= \frac{1}{2n}\sum_{i=1}^{2n}(s_k^i - \hat{\mathbf{x}}_k^i)(h(s_k^i) - \hat{\mathbf{y}}_k^i)^T
\end{aligned} \tag{2.10}$$

$\mathbf{R}_k$ is the covariance associated with the measurement noise.

Lastly, the Kalman gain and measurement and covariance update equations are given by:

$$\begin{aligned}
\mathbf{K}_k &= \mathbf{P_{xy}}\mathbf{P_y}^{-1} \\
\hat{\mathbf{x}}_k^+ &= \hat{\mathbf{x}}_k^- \mathbf{K}_k(\mathbf{y}_k - \hat{\mathbf{y}}_k) \\
\mathbf{P}_k^+ &= \mathbf{P}_k^- - \mathbf{K}_k\mathbf{P_y}\mathbf{K}_k^T
\end{aligned} \tag{2.11}$$

The new *a posteriori* state estimate $\hat{\mathbf{x}}_k^+$ - which contains the C and R values - are recorded, and it along with the new *a posteriori* state covariance matrix $\mathbf{P}_k^+$ are then used as inputs for the following step of the Kalman filter in the next time step.

The continuous-time governing equation in Equation 2.5 can be discretized - for discrete time data - to give the update rule in Equation 2.6, and thus the UKF must be adjusted accordingly. The state function was instead adjusted to $\dot{\mathbf{x}} = f(\mathbf{x}, \theta, \mathbf{u})$, where $\dot{\mathbf{x}}$ is the time derivative of the continuous-time state variable, and θ is the additional model parameters, C and R, following Equation 2.5. Therefore, the following extended model for the UKF was used:

$$\mathbf{x}'_{k+1} = \begin{pmatrix} \mathbf{x}_{k+1} \\ \theta_{k+1} \end{pmatrix} = \begin{pmatrix} \mathbf{x}_k + f(\mathbf{x}_k, \theta_k, \mathbf{u}_k)dt \\ \theta_k \end{pmatrix} \tag{2.12}$$

The state vector $\mathbf{x}_k$ was $[P_k, Q_k]$, the arterial pressure and cardiac output values at time k; the extended state parameters vector θ was $[C_k, R_k]$, the arterial compliance and peripheral resistance values calculated at that time; and the observation vector $\mathbf{y}_k$ was Q_k, the cardiac output.

Because this UKF method generates an arterial pressure prediction at each time step, along with its calculations for C and R, it was also used as a comparison for the predictive model described in the next section. The root mean square error (RMSE) and mean absolute percentage error (MAPE) between the actual arterial pressure values and those output by the filter were calculated and compared.

This research also studied the effect of adjusting variance around times of drastic external changes to the system. To aid the filter in its ability to capture changes in the system quickly, the process noise estimate was increased at and around times of moderate to severe medical intervention. The process noise covariance was increased for 1000 consecutive data points, the equivalent of 20 seconds, for times listed in the annotations when norepinephrine doses were injected or changed, and when CPR occurred. This 20 second window was chosen somewhat arbitrarily as a time period long enough for noticeable changes to occur and propagate throughout the system, yet short enough to isolate the disturbance, though there is certainly room to experiment with and optimize this time window duration. For the sake of simplicity, this research limited the studied medical interventions to just the cases of the application of norepinephrine and CPR, though a few others were listed in the annotations and could potentially be further investigated. The increased covariance was used for 500 data points prior to the time listed in the annotations, to allow for a buffer for any discrepancies between the listed time and actual time of the intervention, and for 500 data points afterward, to allow time for the change in the system to take effect. We did not have capacity to study effects of the specific settings of these assumptions for this report, but these topics should be investigated further in the future work.

This system, like most other Kalman filters, was extremely sensitive to both initial state values and noise values [in vitro 10], and several design iterations of the filter had to be tested before values that functioned for all the several dozen unique and diverse pigs were settled upon. The initial state estimate $\mathbf{x}_k^-$ used was the means of the arterial pressure and cardiac output values for the entire resuscitation for each pig. The initial state covariance $\mathbf{P}_k^-$ was set as a diagonal matrix with the variance

of the arterial pressure along its diagonal. The measurement noise covariance $\mathbf{R}_k$ was set as the variance of the arterial pressure. While the two initial state estimates and measurement noise covariance varied per pig, being functions of their unique data, the process noise covariance matrix $\mathbf{Q}_k$ was fixed for all pigs and was set to $\begin{bmatrix} 500 & 0 \\ 0 & 500 \end{bmatrix}$ normally, and increased to $\begin{bmatrix} 500 & 0.0001 \\ 0.0001 & 500 \end{bmatrix}$ during medical intervention phases.

2.3 Regression Using Dynamic Time Warping

The utility of using the calculated Windkessel parameters as similarity metrics in predicting the health trends of a recovering pig was also investigated. In order to begin testing this query, firstly, after the Windkessel parameters of each pig's baseline and resuscitation phases were found, the mean and standard deviations of the baseline phases were used to normalize the resuscitation phase time series values. This assured that each resuscitation was personalized by its baseline, which was previously found to substantially reduce the effects of subjectivity between hemodynamic presentation of patients in critical care [10].

The next step was to attempt to reconstruct a test pig's arterial blood pressure waveform during its resuscitation phase. Any of the other vital signs provided could have been reconstructed, but this research focused on arterial pressure because it was one of the two vital signs used and predicted in the UKF implementation, and because it is, in medical practice, easier to obtain than cardiac output, the other vital sign used in the UKF. Rather than attempt to reconstruct an entire resuscitation's pressure time series of tens of thousands of data entries all at once, this research focused on reconstructing much smaller periods, one at a time. Each normalized episode of resuscitation was broken up into non-overlapping windows of 100 entries each, or 2-second-long windows. Each data entry consisted of two values: the C and R values for the pig at that time step. For each window of one test pig, n similar neighbors of windows from other pigs at some point in their respective resuscitations were chosen by minimizing a distance criterion. Various techniques were used to group the normalized C and R values such K-nearest neighbors and Dynamic Time Warping (DTW) with a Euclidean distance metric. The latter was chosen as producing the

best results.

Once the n most similar neighbors for the given window were selected, DTW was used to align the pressure waveforms corresponding to that window of the n neighbors to the pressure waveform of the test pig. The mean of these n waveforms was taken as the predicted pressure profile for the test pig, and the model moved on to the next window and selected a new set of similar pigs, reconstructing the test pig's entire resuscitation phase window by window. Both the predicted pressure waveforms and a smoothed average pressure time series, using a sliding window average, were studied and compared to the original data in a leave-one-subject-out validation scheme. The RMSE and MAPE between the actual and predicted pressure waveforms were also studied and compared to those calculated by the UKF. Different values for the number of other pigs chosen per window were tried before $n = 7$ was settled upon as producing the best results the most consistently.

The same neighbors can be used to construct forecasts of the arterial blood pressure waveforms. In the next section, we include results computed for 5, 15, and 30 minute forecast horizons. These time spans have been selected to fit the typical scenarios of resuscitation, patient monitoring, and clinical decision making in the fast-paced critical care settings.

Chapter 3

Results and Discussion

All test pigs that did not crash had converging and stable arterial compliance C and peripheral resistance R values using the UKF, though the filter took longer for some pigs to reach stability than for others. The UKF caused an initial large spike in one or both Windkessel parameters for most pigs before stabilizing; pigs whose values took longer to settle experienced an additional small number of peaks in the first few minutes of resuscitation data. This is likely due to the tendency of the UKF to have large initial errors in practice [4].

The majority of pigs converged to the same range of C and R values once they reached their stable baseline or near-baseline health levels in the latter stages of their resuscitations. The C values converged to 0.5 mL/mmHg on average and R values to 0.5 mmHg·s/mL on average, though distributions of the two shown in Figure 3.2 show that there was a comparatively smaller yet noticeable cluster of secondary C convergence values at 9.7 mL/mmHg. The pigs from the 5 mL/min bleed group had peaks at 0.5 mL/mmHg and 0.5 mmHg·s/mL, and the pigs from the 20 mL/min bleed group had slightly different peaks at 0.6 mL/mmHg and 0.6 mmHg·s/mL. Its ability to reconstruct the arterial pressure waveforms was also investigated by calculating the RMSE and MAPE between each time step's predicted value and the actual value. Across all pigs that did not crash, the average RMSE was 5.49 mmHg, and the average MAPE was 6.19%. An example is shown in Figure 3.3, depicting both several individual heartbeats of and smoothed averages of the actual pressure waveforms and those predicted by the filter, in which the UKF is able to reconstruct each successive

pressure value very well. This is expected due to the high sampling frequency of the data used in the filter, as the filter predicts each state a small amount of time in the future.

Increasing the variance around points of medical interventions either had little to no noticeable effect in revealing sudden value changes, such as in the arterial compliance plot in Figure 3.1, or revealed a sudden change in one or both of the Windkessel parameters, as seen in the accompanying peripheral resistance plot. Plots showing the calculated Windkessel parameters for the same specimen (pig 6 from the 5 mL/min bleed group) without increased variance during its medical interventions are shown in Figure 3.4. There is a slight change in the trend followed by the arterial compliance between the constant and adjusted variance scenarios, but the values largely remain the same. The peripheral resistance without the adjusted variance, however, shows a much more gradual decay to the value it eventually settles at and shows no major immediate changes due to the external influence, the administered norepinephrine dose. Comparing the two scenarios indicates that the increased variance allowed the filter to react more quickly to this change in the pig's system

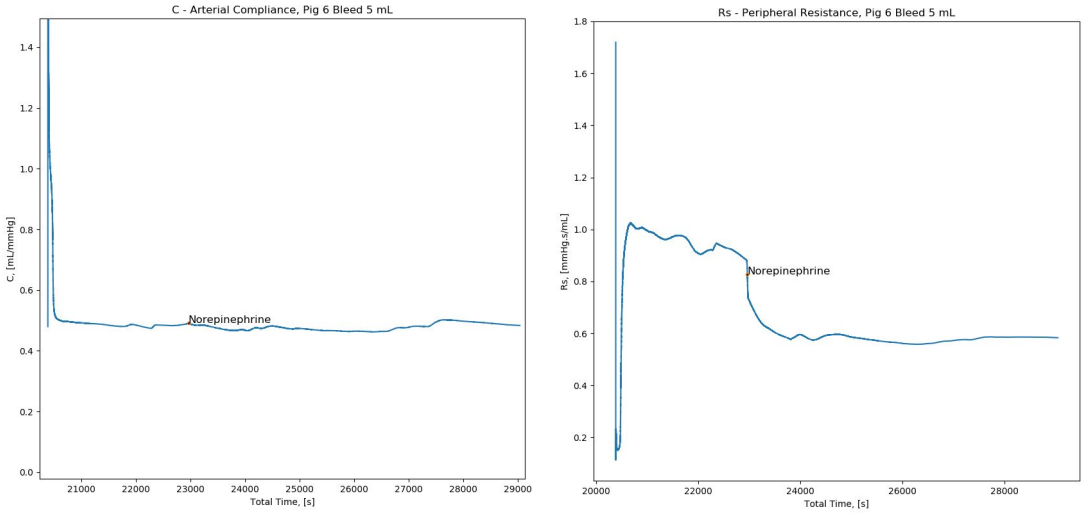

Figure 3.1: Arterial compliance C (left) and peripheral resistance R (right) for pig 6 from the 5 mL/min bleed group. This pig experienced a norepinephrine dose at around 23,000 seconds, causing the filter to increase its variance around that time. Its arterial compliance values did not appear to change much in response to this disturbance, but its peripheral resistance suddenly dipped.

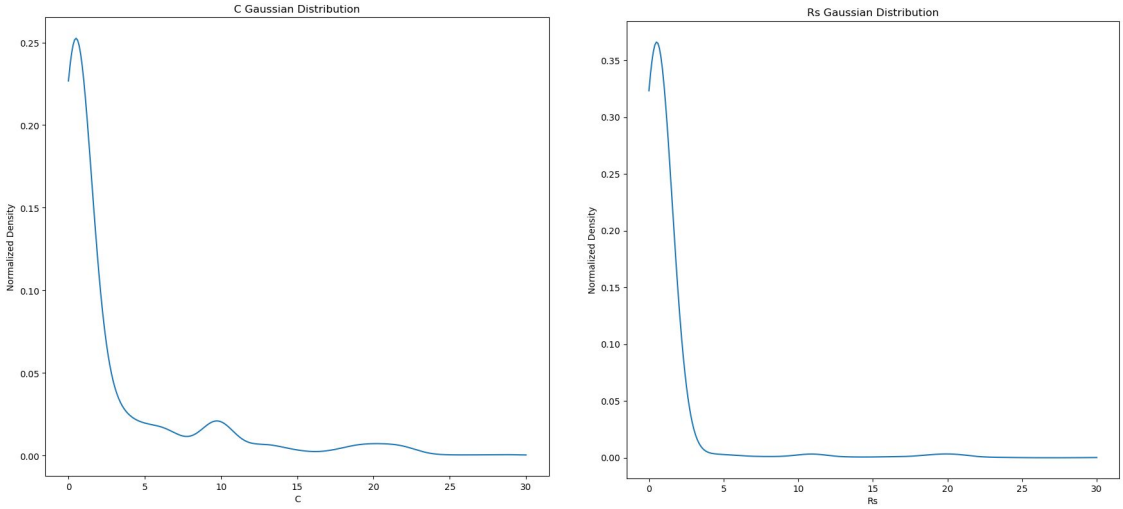

Figure 3.2: The distributions for the calculated C (left) and R (right) values. Arterial compliance has peaks at 0.5, 9.7, and 20.2 mL/mmHg. Peripheral resistance has one major peak at 0.5 mmHg·s/mL and two smaller peaks at 10.9 and 19.9 mmHg·s/mL.

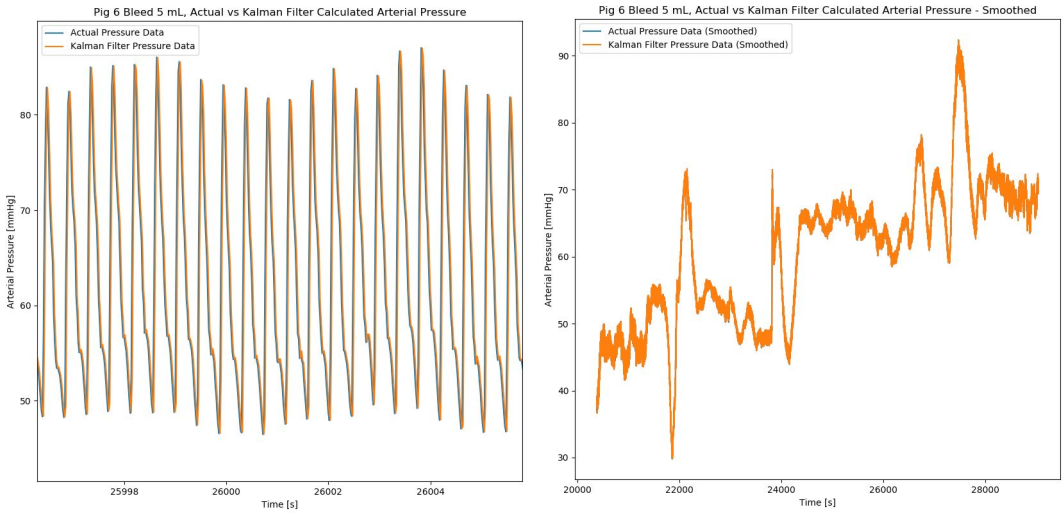

Figure 3.3: Predicted and actual pressure waveforms and smoothed average pressure, calculated from the Unscented Kalman Filter, from the resuscitation of pig 6 from the 5 mL/min bleed group. Because the data sampling rate is high, the UKF is able to calculate the next pressure value in the time series with high accuracy. The RMSE across the entire resuscitation was 3.17 mmHg, and the MAPE was 5.06%.

and adjust the peripheral resistance accordingly. About a third of the pigs showed behavior in one or both of their Windkessel parameters similar to this latter case.

Though the pigs used in these bleed experiments were chosen and raised to be as similar as possible, that there was such disparity in the effects of the increased variance reveals an interesting insight into how different these seemingly homogeneous specimens are from each other when responding to medical trauma.

Because of the large number of test pigs and their apparent heterogeneity, coupled with the observed sensitivity of this filter, it was difficult to find optimal variances for each unique pig, and so the same standard and increased variances were used for all of them. This may have negatively affected the efficacy of the adjusted variance method altogether, as some pigs may have responded better to a higher variance and others better to a lower variance. For example, the arterial compliance for pig 6 from the 5 mL/min bleed group may have also had a sudden change in value that the filter could not capture with the increased variance used. This apparent diversity should be studied further as it has a potential to help characterize response to interventions at the individual patient levels, given that accurately predicting and rapidly characterizing such response is one of the key challenges in clinical practice, especially in fast-paced critical care settings.

For every consecutive 100 data entry, 2-second-long window throughout a test pig's resuscitation, the n most similar pigs to the test pig in that window were selected by minimizing distances between Windkessel parameters using DTW and a Euclidean distance metric. Pigs from both bleed groups (5 mL/min and 20 mL/min) were used when selecting similar pigs. Their arterial pressure waveforms were then aligned to the test pig's waveform via DTW and averaged to produce pressure waveform predictions. A small section of the actual and predicted pressure waveforms are shown in Figure 3.5.

Both the original and predicted pressure waveforms were also averaged over time using a sliding window to create a smoothed average pressure plot to compare how well the predicted average arterial pressure matched that of the original both in value and in trends. A comparison of the average arterial pressures over the course of the resuscitation are shown in Figure 3.6. The root mean square error (RMSE) and mean absolute percentage error (MAPE) between the predicted and actual pressure values were calculated to serve as performance metrics for the model. The average RMSE was 11.56 mmHg, and the average MAPE was 13.93%.

The pressure waveforms and averaged pressure plots show the weaknesses and

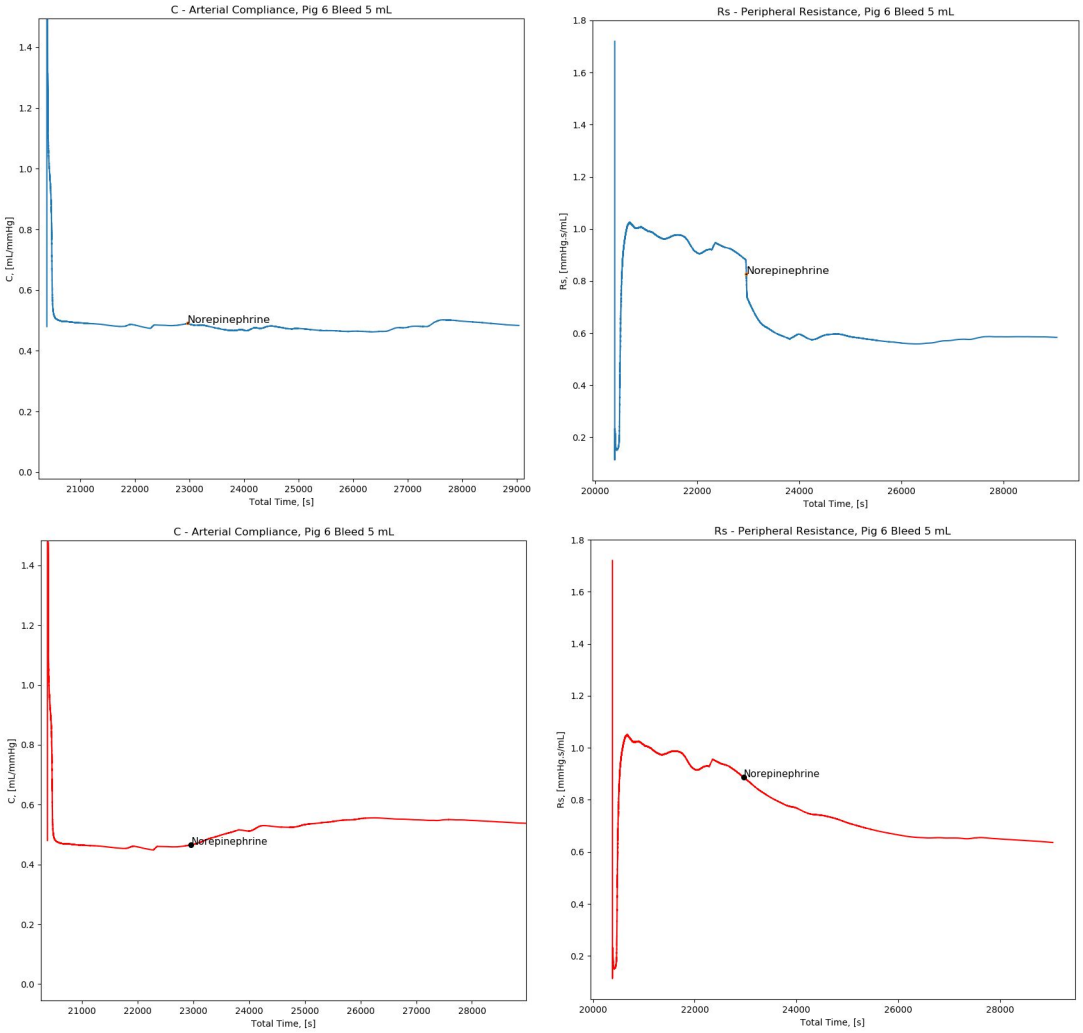

Figure 3.4: Comparisons of the arterial compliances C (left column) and peripheral resistances (right column) for pig 6 from the 5 mL/min bleed group. The top row shows the parameters calculated with increased variance due to the medical intervention (the norepinephrine dose), and the bottom row shows the results without increased variance. There is little change between the two compliance series, with only slight changes in trends, and the two cases remain at and converge to the same values. However, the increased variance clearly produces a drastic change in resistance that is not captured when the intervention is ignored, though both resistances converge to the same values.

strengths of this predictive model. It struggles with consistently recreating the original waveforms with their accurate pressures, notably at the peaks and valleys of each

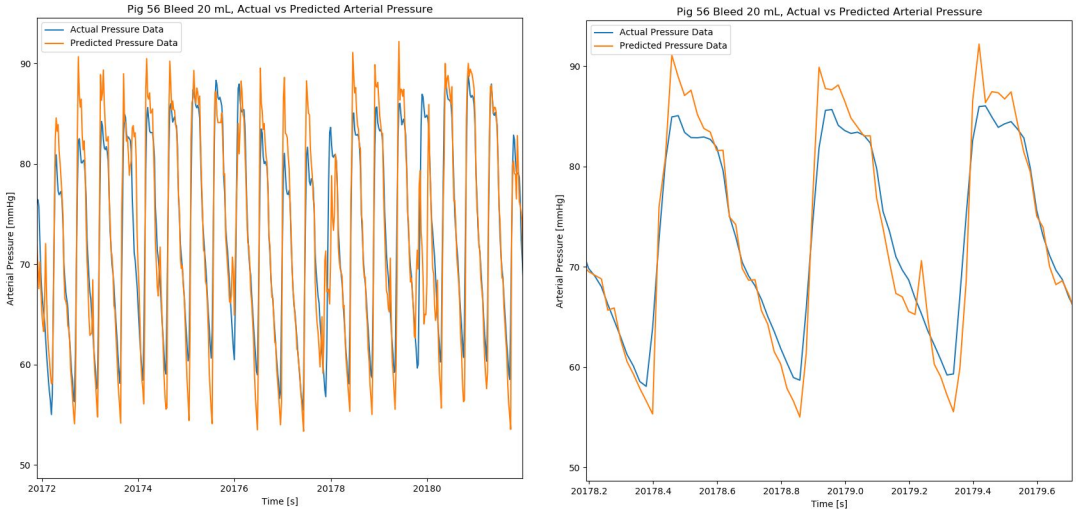

Figure 3.5: Predicted and actual pressure waveforms for a small time period from the resuscitation of pig 56 from the 20 mL/min bleed group. The left image shows general predicted and actual arterial pressure trends over a window of about 10 seconds, and the right image isolates three heartbeats to view closer detail of the predicted waveform's fit. DTW allows predicted and actual heartbeats to be aligned, but the predicted values often significantly overshoot the actual values at the peaks and valleys of each heartbeat.

heartbeat. Figure 3.5 shows that the predicted waveforms can significantly overshoot the actual pressure values. However, the model is able to reproduce the average pressure decently well, particularly the pressure trends. The predicted values in Figure 3.6 align with many of the small peaks and dips in the actual pressure values, especially in the early and later stages of the resuscitation, despite having occasional random spikes and under- or overshooting the actual average values. It is possible that some of these discrepancies could be reduced by independently calibrating model predictions on a set-aside small sample of test pig data, perhaps collected during the stable period before the bleeding experiment start, or at the very beginning of the resuscitation process. Due to time constraints, this intriguing question had to be left for future research.

The majority of the pigs had performances well within a standard deviation of average values outlined in Table 3.1, though the model performed noticeably worse on some pigs than on others. The UKF occasionally output large spikes in Windkessel parameter values that, in a few select cases, aligned with large arterial pressure spikes,

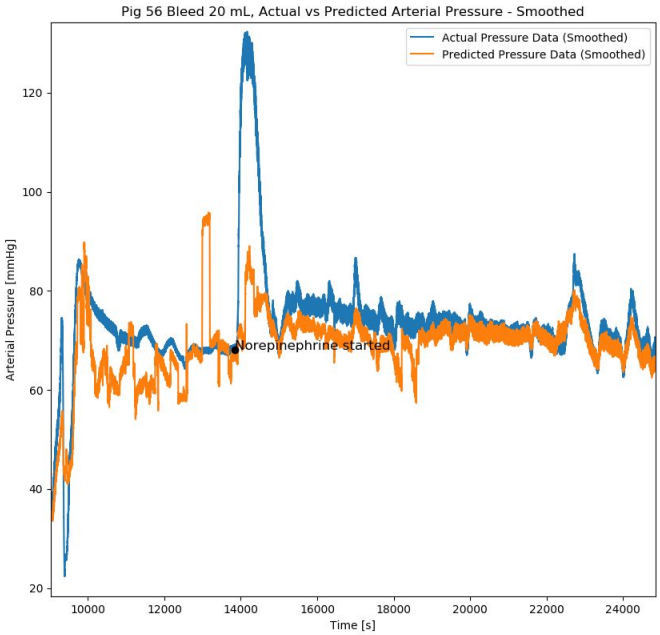

Figure 3.6: Predicted and actual pressure waveform averages over the course of the resuscitation of pig 56 from the 20 mL/min bleed group. The predicted values follow the trends of the actual values, especially in the latter half of the resuscitation, despite a few random peaks and dips. However, the model is unable to capture the large pressure spike early in the resuscitation. The RMSE across the entire resuscitation was 12.01 mmHg, and the MAPE was 9.17%.

as seen in the plots in Figure 3.7. However, the majority of these Kalman filter spikes are placed randomly and do not align with pressure spikes. We have not yet fully determined the sources of these large value spikes generated by the UKF, though the behavior visually indicates a ramp in value caused by poorly tuned variance values in the filter for those few pigs in particular. These drastic ramps in Windkessel parameter values resulted in poorer predictions matching neither in value nor trend to the actual pressure values, as few other pigs had Windkessel parameters that behaved similarly at that point in their resuscitations. These abnormalities in the results further show that the pigs used in the bleed experiments can differ enough to produce unusual results even though they were raised to be as similar as possible. A promising circumstance is that by noticing proximity of neighbors selected for making predictions for the current patient, we could flag predictions with less than favorable support of evidence from close enough peers.

Many of the pigs experienced large spikes in arterial pressure in their resuscitations, such as those seen in Figure 3.6 and 3.7, for unknown reasons. The model is unable to capture and replicate these large spikes. This is likely due to the UKF-generated Windkessel parameters not changing accordingly to reflect those spikes in a meaningful way. Because the cardiac output as measured and presented in the data is an average estimate as opposed to an instantaneous measurement, it rarely has large jumps aligning with those in the pressure data, which may affect the accuracy of the calculated Windkessel parameters at those pressure peaks. Moreover, the physician annotations for the data rarely mention anything at or near these peaks (the norepinephrine dose proceeding the spike in Figure 3.6 is a rare exception), so they cannot be relied upon to adjust the variance manually and aid the filter change quickly to these disturbances in the system. However, since the manual adjustments do improve the regression results, it is entirely possible that adjustments implemented at these spikes - if their onset is known or can be estimated ahead of time - could improve the model.

Other attempts were made using only one of the Windkessel parameters, either arterial compliance or the peripheral resistance, to select similar pigs. This was done to investigate if one parameter were more indicative of the health of the pig than the other or than both used together. To compare each similarity selection method, the RMSE and MAPE computed with the test pig's arterial pressure and the respective predicted pressures for the entire resuscitations were compared and averaged across all pigs that did not crash. These are summarized in Table 3.1 and Table 3.2. With regards to RMSE, on average across all pigs, using only the peripheral resistance provided the best prediction results of the three methods with an average RMSE of 11.52 mmHg, though the differences in errors across all of the pigs are slight, with using both compliance and resistance giving 11.56 mmHg and only compliance giving 11.63 mmHg. Using both Windkessel parameters gives the best average MAPE at 13.93%. For comparison, the UKF had errors about half these values, with an average RMSE of 5.49 mmHg and MAPE of 6.19%. Performances for pigs from just the 5 mL/min group and just the 20 mL/min group were similar across both groups, which may be because the model does not use data from the bleed portion of the experiment, though it may also indicate that the rate of blood loss does not affect the Windkessel parameters enough to have an effect on this model. However, the model on average

performed worse for pigs from the 5 mL/min bleed group in both RMSE and MAPE.

Average RMSE of all pigs			
Bleed Group	Both C and R [mmHg]	C [mmHg]	R [mmHg]
All pigs	11.56	11.63	11.52
5 mL/min	12.40	12.10	12.45
20 mL/min	11.19	11.42	11.11
Std. Dev.	3.39	3.81	4.05
Var.	11.47	10.39	14.53

Table 3.1: The average RMSE when predicting arterial pressure for the entire resuscitation and selecting similar pigs based on combinations of the Windkessel parameters. Using only the peripheral resistance provides the best results for all pigs by a slight margin, though all three methods provide comparable results. The averages are also shown for pigs only from the 5 mL/min and 20 mL/min bleed groups, as are the standard deviations and variances for all pigs from both combined groups.

Average MAPE of all pigs			
Bleed Group	Both C and R [%]	C [%]	R [%]
All pigs	13.93	14.41	14.11
5 mL/min	15.36	15.25	15.78
20 mL/min	13.31	14.04	13.34
Std. Dev.	5.25	5.52	6.31
Var.	0.28	0.30	0.40

Table 3.2: The average MAPE when predicting arterial pressure for the entire resuscitation and selecting similar pigs based on combinations of the Windkessel parameters. Using both the arterial compliance and peripheral resistance provides the best results for all pigs by a slight margin, though all three methods provide comparable results. The averages are also shown for pigs only from the 5 mL/min and 20 mL/min bleed groups, as are the standard deviations and variances for all pigs from both combined groups.

Pigs that crashed were excluded from the group calculations of model performance because the model was completely unable to capture the steeply decreasing arterial pressure trends, as seen in Figure 3.8. 6 such pigs were removed from the average RMSE and MAPE calculations as all had over double their respective average values. It may be possible to train separate Kalman filters and DTW regressors, or to use

discrepancies between our model predictions and actually measured blood pressure trends on these subjects to be able to predict which pigs and when will crash, but no such attempts were made due to time constraints and the small number of examples.

To test the validity of increasing the variance of the UKF when calculating the Windkessel parameters, additional experiments were done in which no changes to the filter were made during medical interventions, and these newly calculated Windkessel parameter time series were used for the same arterial pressure regression method detailed above. The RMSE across all resuscitations of using both Windkessel parameters when selecting similar pigs and with no changes due to interventions was 12.86 mmHg, and the corresponding MAPE was 14.69%, larger errors than all other attempts in which medical interventions were taken into account. Though the error increase is slight, this does suggest that altering the dynamic system in the UKF due to extreme external influences better reflects medical changes in the pigs and provides better similarity criteria, and that further refinement of acknowledging these external influences in the filter, and perhaps also in the regressor, could provide better regression results. The results for all test pigs used, excluding pigs that crashed, are shown in Table A.1.

In addition to reconstructing arterial pressure waveforms, the model's ability to forecast a test pig's arterial pressure was investigated. For every 2-second-long window, the n most similar neighbors were selected using both Windkessel parameters as already described, including the increased variance due to medical interventions, and those pigs' waveforms from a certain amount of time in the future were combined to forecast how the test pig's arterial pressure would change. Each window and its similar pigs were used to forecast the test pig's arterial pressure 5, 15, and 30 minutes in the future. An example of a smoothed average 5-minute forecast is shown in Figure 3.9. The RMSE and MAPE are 18.70 mmHg and 26.59% for the 5-minute forecast, 19.46 mmHg and 27.74% for the 15-minute forecast, and 19.55 mmHg and 28.20% for the 30-minute forecast.

Each method, like the original reconstruction model, is able to forecast the general pressure trends decently well. Because of time constraints, further refinement of the forecasting method could not be investigated. However, that the model is able to forecast the pressure trends significant amounts of time into the future is a promising result and shows that given more time and research, models such as this can be

utilized in care settings to forecast a patient's vital sign trends and responses to treatment.

Despite the inability to track extremely large and sudden spikes in pressure and the occasional random peaks and dips in its predictions, that the model is able to track and forecast the average arterial pressure trends well shows that the health trends of pigs with similar Windkessel parameters can be indicative of each other. Moreover, it shows that predictive accuracy can be increased by manually accounting for periods of moderate to severe medical intervention. This model uses only a simple distance metric when selecting similar pigs based on their Windkessel parameters. Had other selection criteria been implemented, such taking into account the presence or absence of medical interventions or weighting the pressure data of the selected pigs by some similarity factor to the test pig, the predicted pressure values may have shown more refinement and closer similarities to the original.

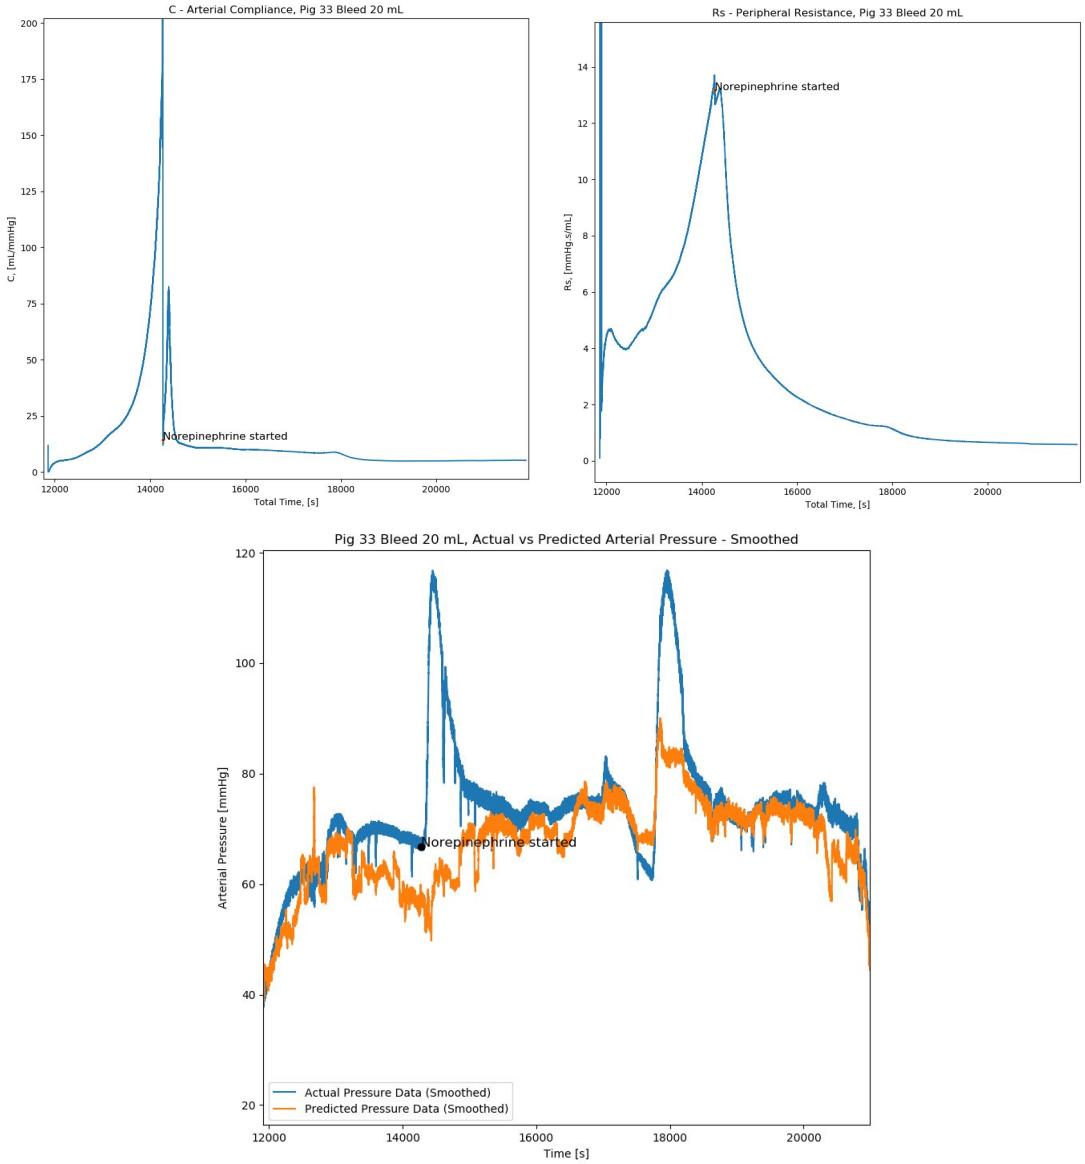

Figure 3.7: Arterial compliance C (upper left) and peripheral resistance R (upper right) for pig 33 from the 20 mL/min bleed group, and its corresponding predicted and actual pressure waveform averages (bottom). The spikes in compliance and resistance values seem to match the first large spike in arterial pressure, though most other pigs with similar pressure spikes did not have corresponding Windkessel parameter spikes. These abnormally large compliance and resistance values around the 14,000 second mark caused the regression model to perform poorly at that time as there were few other pigs that behaved similarly at that stage in their resuscitations. Interestingly, the Windkessel parameter values show small increases in value at the 18,000 second mark, seemingly corresponding to the second large arterial pressure spike. The RMSE across the entire resuscitation was 16.61 mmHg, and the MAPE was 10.79%.

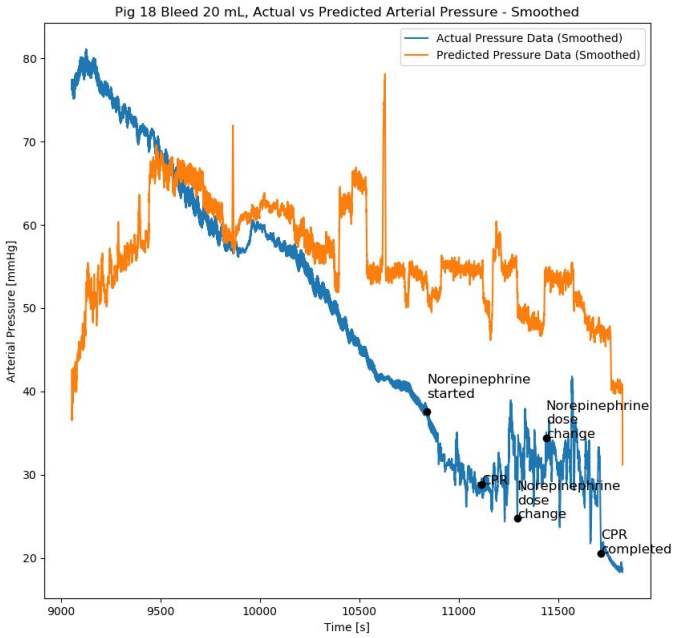

Figure 3.8: Predicted and actual pressure waveform averages for pig 18 from the 20 mL/min bleed group. This pig was unable to be resuscitated as indicated by the numerous CPR interventions near the end of the data available. The model is unable to predict the arterial pressure values for pigs that crash like this. The RMSE across the entire failed resuscitation was 28.86 mmHg. Note that these substantial mismatches are included in the results reported, diminishing overall performance metrics, even though these cases are relatively easy to identify and exclude. Note also that we could track discrepancy between predicted and actual values of arterial blood pressure and use these residuals to build detectors that could issue early warnings about patients trending towards significant health crises.

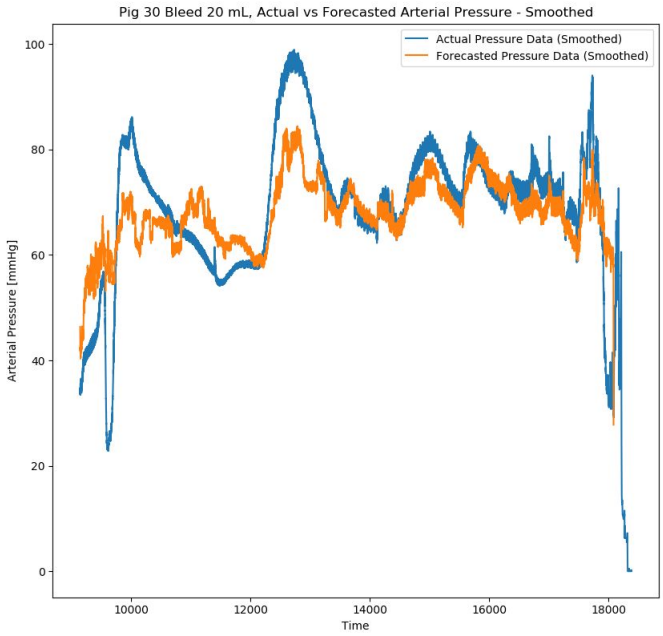

Figure 3.9: Forecasted and actual pressure waveform averages over the course of the resuscitation of pig 30 from the 20 mL/min bleed group when the model attempts to forecast 5 minutes into the future. The forecasted values follow the trends of the actual values, though they seem to be shifted forward slightly as seen in the several larger pressure peaks in orange appearing later than their respective peaks in blue. The RMSE across the entire resuscitation was 11.93 mmHg, and the MAPE was 12.96%.

Chapter 4

Conclusions and Further Work

This research sought to calculate the two-element Windkessel parameters and use them to predict arterial pressure over the course of the resuscitation of a pig recovering from severe blood loss. An Unscented Kalman Filter was used to calculate the Windkessel parameters and reconstruct the arterial pressure waveforms of each pig with high accuracy. It was also able to track the dynamic system of the pigs' health using manually adjusted variances with moderate success, being unable to track large spikes in arterial pressure without corresponding trends in cardiac output or mentions of causes by the physicians performing the experiments. Dynamic Time Warping was used to select pigs similar to a test pig using a Euclidean distance metric and combine their pressure waveforms into a prediction of the test pig's arterial pressure over the course of its resuscitation. This predictive model was able to follow the overall trends of the test pig's average arterial pressure, though it was unable to predict large spikes. It was also shown that adjusting the variance in the Kalman filter during medical interventions produced better overall predictive accuracy than when these external influences to the system were ignored. Moreover, in addition to reconstructing the original pressure waveforms successfully, this model was also able to forecast the pressure trends of the test pigs several minutes into the future. This model has set the framework for predictive health modeling using the Windkessel model parameters as similarity features.

This work set the basis for an analytic algorithm for predicting a patient's recovery from medical trauma and their response to treatment using a simple physical heart

model that many physicians are familiar with. Though it focused only on predicting the arterial blood pressure, which in of itself is a highly useful vital sign for physicians to monitor, other implementations of similar models could be used to predict other vital signs as well. A further developed model would be incredibly valuable in both hospital and field critical care settings, particularly models that could be trained and tuned on data with a wide variety of common medical interventions, and as such could be tailored to specific medical applications. Predictive models such as this could alert physicians to evolution of patients' health seen through the lens of high frequency vital signs, and what their expected outcomes are and could therefore greatly aid in the division of care and resources in bustling and fast-paced medical settings.

There are several directions this research can be taken next. Further studies could be done on the utilization of a Kalman filter to track changes in the Windkessel parameters more accurately, particularly during annotated and/or apparent severe changes to the dynamics of the system. If given knowledge of sudden spikes in arterial pressure, the Kalman filter may be able to respond more quickly and produce Windkessel values more reflective of the change, though there are challenges in this approach in that there is no gold standard for calculating Windkessel parameters and what their values should be in critical care patients. More research could be done into the forecasting abilities of the Kalman filter; it is able to reconstruct pressure waveforms well, and its dynamic state tracking abilities could prove useful in forecasting changes to this state, either by itself or combined with regressors like the DTW method implemented in this research. Separate filters could also be tuned only on subjects that crashed and were unable to be resuscitated to begin building the framework for predicting fatal trends as opposed to just stable recoveries. This approach can also be tried with higher complexity Windkessel models, such as the three- and four-element models, both of which have been shown to model the cardiac system more accurately than the two-element in many cases.

Other techniques can also be studied for cyber-physical modeling of hemodynamic data for regression. When selecting similar pigs, predicting, and forecasting blood pressure values, weighting by similarity in Windkessel parameters can be implemented. Weight criteria could include the presence or absence of specific types of medical interventions, along with numerical parameters directly related to the data such as

its variance or standard deviation. Other values can also be used to determine the most similar pigs, such as additional vital sign waveforms, medical history data, and the rate of change in the Windkessel parameters, instead of just their raw values at that time. Closer attention could be paid to the similar neighbors whose waveforms are used for regression; if some or all of the most similar neighbors fall past some threshold considered too dissimilar, or if the considered neighbors vary too much in their Windkessel parameters, then the corresponding predicted pressure windows could be flagged as less confident or even rejected. There is also much room to study what metrics are best for selecting similar subjects other than just simple Euclidean distance, and learning such metrics from data appears very appealing in this context.

Appendix A

Appendix

Average RMSE of each pig using four methods					
Bleed Group [mL/min]	ID	Both C and R [mmHg]	C [mmHg]	R [mmHg]	No Interventions [mmHg]
5	2	17.61	17.08	17.35	18.43
5	3	12.46	11.64	12.30	14.00
5	4	24.55	23.01	25.17	24.90
5	5	15.84	12.10	15.06	15.29
5	6	9.15	9.96	9.25	8.08
5	7	10.79	10.38	8.89	10.47
5	8	7.70	7.85	8.58	8.07
5	9	9.46	9.35	10.05	8.22
5	10	9.22	9.57	8.34	8.74
5	11	6.91	7.11	7.07	9.51
5	12	10.48	10.67	10.47	11.3
5	13	7.37	6.91	7.32	7.90
5	14	12.25	13.02	12.73	12.47
5	15	12.71	13.00	13.37	13.52
5	16	19.46	19.53	20.83	18.82
20	6	12.44	11.81	12.39	12.64
20	7	12.48	12.88	11.89	12.46

20	8	9.27	10.36	10.35	10.81
20	9	8.94	8.89	10.21	10.32
20	10	13.34	13.89	12.66	20.50
20	11	7.62	8.02	6.62	11.3
20	14	8.11	7.89	8.84	7.87
20	21	9.75	10.81	9.67	8.53
20	23	16.58	16.74	18.72	16.53
20	24	9.62	9.15	7.56	10.62
20	26	10.58	10.66	10.50	11.94
20	27	10.62	10.89	11.40	12.93
20	29	18.00	18.55	19.02	19.48
20	30	10.09	10.43	10.38	13.06
20	32	11.52	13.66	11.99	14.46
20	33	13.28	12.35	12.58	15.86
20	36	12.29	11.88	12.50	15.33
20	37	13.07	14.36	12.75	14.50
20	38	10.63	12.04	8.86	13.20
20	41	9.06	8.95	9.13	8.35
20	42	12.27	13.79	11.15	12.58
20	43	10.32	10.13	11.31	17.68
20	44	9.36	9.43	7.61	13.45
20	45	9.94	10.15	10.36	11.00
20	47	12.32	13.33	9.90	17.71
20	48	9.02	8.92	8.82	11.36
20	49	11.57	14.06	10.87	10.92
20	53	12.53	11.54	13.52	13.68
20	56	12.52	11.64	17.48	12.19

20	58	14.11	9.67	8.79	15.35
20	60	13.32	13.65	12.07	15.35
20	61	9.10	8.27	7.82	9.14
20	64	7.59	9.84	6.44	7.07
20	65	9.23	9.70	14.52	10.45

Table A.1: The average RMSE for each pig. The The third, fourth, and fifth columns indicate different combinations of Windkessel parameters used to select similar pigs for each window, and the Windkessel parameters were calculated with increased variance in the UKF around times of moderate to severe medical interventions. The last column used both the arterial compliance C and peripheral resistance R Windkessel values when selecting similar pigs, and the parameters were calculated without increased variance during medical interventions.

Average MAPE of each pig using four methods					
Bleed Group [mL/min]	ID	Both C and R [%]	C [%]	R [%]	No Interventions [%]
5	2	21.69	21.27	22.01	21.59
5	3	14.71	13.48	14.77	15.34
5	4	19.63	18.86	22.05	19.67
5	5	18.44	16.90	17.37	19.20
5	6	12.68	14.43	12.46	12.04
5	7	12.34	9.29	8.00	9.71
5	8	9.02	9.58	10.03	9.59
5	9	13.84	14.41	15.02	11.85
5	10	12.64	13.53	12.16	11.48
5	11	7.38	7.88	7.51	8.47
5	12	10.92	11.27	11.93	10.95
5	13	9.26	8.45	9.22	9.36
5	14	19.01	20.15	19.46	18.50
5	15	18.60	19.45	20.48	22.75
5	16	30.01	29.75	34.26	29.35

20	6	14.55	15.44	16.40	16.86
20	7	11.37	11.90	11.13	10.74
20	8	14.15	16.23	16.21	16.64
20	9	11.59	13.83	15.07	15.36
20	10	12.21	11.46	12.16	11.48
20	11	8.76	9.30	7.51	8.47
20	14	10.09	9.89	19.46	18.50
20	21	14.99	16.93	14.16	10.94
20	23	28.56	29.07	32.73	29.70
20	24	13.76	13.10	10.48	11.43
20	26	16.10	15.88	16.37	13.99
20	27	9.68	9.51	9.91	9.88
20	29	30.75	33.51	35.00	40.35
20	30	11.39	11.61	11.95	11.97
20	32	17.25	21.33	18.60	21.36
20	33	10.79	10.21	9.90	10.20
20	36	11.85	12.19	11.93	12.41
20	37	16.04	18.69	15.60	17.13
20	38	16.23	18.08	13.05	12.72
20	41	11.72	12.09	11.50	10.31
20	42	12.97	16.48	11.04	13.06
20	43	9.17	9.25	9.78	12.41
20	44	13.02	13.71	10.43	13.19
20	45	10.93	11.79	10.25	12.69
20	47	12.74	15.29	9.67	18.25
20	48	10.54	10.74	9.91	10.82
20	49	10.51	14.26	10.08	9.29
20	53	11.77	10.42	12.66	11.20
20	56	9.17	8.13	17.45	10.30
20	58	17.69	12.29	11.43	13.36
20	60	11.09	12.22	10.17	12.22
20	61	12.08	11.09	9.29	10.26

20	64	9.07	12.77	8.01	17.11
20	65	7.80	8.86	14.64	7.99

Table A.2: The average MAPE for each pig. The The third, fourth, and fifth columns indicate different combinations of Windkessel parameters used to select similar pigs for each window, and the Windkessel parameters were calculated with increased variance in the UKF around times of moderate to severe medical interventions. The last column used both the arterial compliance C and peripheral resistance R Windkessel values when selecting similar pigs, and the parameters were calculated without increased variance during medical interventions.

Average RMSE of each pig when forecasting different horizons				
Bleed Group [mL/min]	ID	5 min [mmHg]	15 min [mmHg]	30 min [mmHg]
5	2	19.25	21.51	20.91
5	3	12.04	15.00	16.82
5	4	24.13	25.53	23.62
5	5	19.93	17.85	18.34
5	6	13.17	17.82	17.76
5	7	13.91	15.55	17.11
5	8	17.95	17.74	19.51
5	9	12.72	15.2	16.33
5	10	13.65	16.70	15.32
5	11	12.24	16.82	18.15
5	12	12.11	16.03	18.06
5	13	11.21	16.15	15.46
5	14	15.81	18.57	18.97
5	15	13.81	18.51	18.99
5	16	23.66	25.23	23.84
20	6	16.21	20.17	21.63
20	7	14.75	17.92	18.03

20	8	13.69	17.74	17.99
20	9	14.86	18.39	19.45
20	10	18.87	22.29	24.33
20	11	11.25	15.07	16.77
20	14	12.28	16.79	16.93
20	21	12.99	16.12	18.03
20	23	21.76	26.14	26.78
20	24	11.44	14.55	15.93
20	26	14.15	15.00	17.05
20	27	15.15	20.37	20.68
20	29	22.62	24.20	26.27
20	30	16.10	22.23	22.55
20	32	14.94	19.49	18.79
20	33	17.76	22.42	23.84
20	36	16.22	21.7	22.36
20	37	16.20	21.74	22.36
20	38	14.77	18.11	20.93
20	41	13.11	17.11	19.35
20	42	15.77	17.89	18.35
20	43	15.78	20.25	22.03
20	44	12.93	16.75	17.32
20	45	15.47	20.24	21.93
20	47	14.05	18.01	19.38
20	48	14.55	16.13	17.96
20	49	14.33	19.00	19.85
20	53	16.06	21.43	22.97
20	56	16.32	19.2	20.54

20	58	13.97	16.7	17.94
20	60	16.03	19.22	20.85
20	61	14.96	19.77	21.36
20	64	12.84	15.65	16.42
20	65	14.70	17.35	18.91

Table A.3: The average RMSE for each pig when forecasting using both Windkessel parameters and increased variance in the UKF around times of moderate to severe medical interventions. The third, fourth, and fifth columns represent forecasting 5, 15, and 30 minutes in the horizon.

Average MAPE of each pig when forecasting different horizons				
Bleed [mL/min]	Group ID	5 min [%]	15 min [%]	30 min [%]
5	2	27.20	25.12	31.55
5	3	14.36	20.14	21.87
5	4	20.48	25.92	22.46
5	5	25.01	27.20	29.32
5	6	19.98	25.20	25.90
5	7	14.36	16.61	18.88
5	8	22.29	22.76	22.98
5	9	18.64	22.90	24.48
5	10	18.37	22.76	23.28
5	11	15.46	20.18	23.07
5	12	15.71	20.38	22.43
5	13	16.42	21.70	20.66
5	14	25.8	25.62	26.09
5	15	21.11	25.35	32.67
5	16	35.41	34.75	28.33
20	6	21.60	23.59	24.10
20	7	17.34	19.44	19.56

20	8	13.74	18.55	19.12
20	9	20.86	24.69	24.20
20	10	19.93	23.89	24.15
20	11	16.75	20.37	24.49
20	14	15.04	22.86	24.11
20	21	18.72	23.01	25.73
20	23	35.90	44.48	48.09
20	24	18.02	19.43	20.02
20	26	21.68	18.21	19.93
20	27	17.99	21.61	21.67
20	29	40.60	46.53	46.99
20	30	18.72	24.91	23.21
20	32	23.20	33.07	32.22
20	33	16.23	23.18	21.81
20	36	15.10	23.36	25.41
20	37	20.81	29.25	29.81
20	38	23.57	26.43	28.55
20	41	17.32	24.14	25.58
20	42	17.78	22.94	23.12
20	43	14.29	20.86	21.09
20	44	18.05	24.69	26.49
20	45	16.09	24.43	25.85
20	47	15.15	20.00	21.03
20	48	14.77	22.29	23.09
20	49	14.20	19.84	20.07
20	53	15.93	21.37	23.48
20	56	16.60	18.65	19.73

20	58	15.37	20.77	21.35
20	60	17.01	19.22	19.88
20	61	19.53	23.47	23.94
20	64	14.16	18.87	19.05
20	65	14.00	18.32	19.16

Table A.4: The average MAPE for each pig when forecasting using both Windkessel parameters and increased variance in the UKF around times of moderate to severe medical interventions. The third, fourth, and fifth columns represent forecasting 5, 15, and 30 minutes in the horizon.

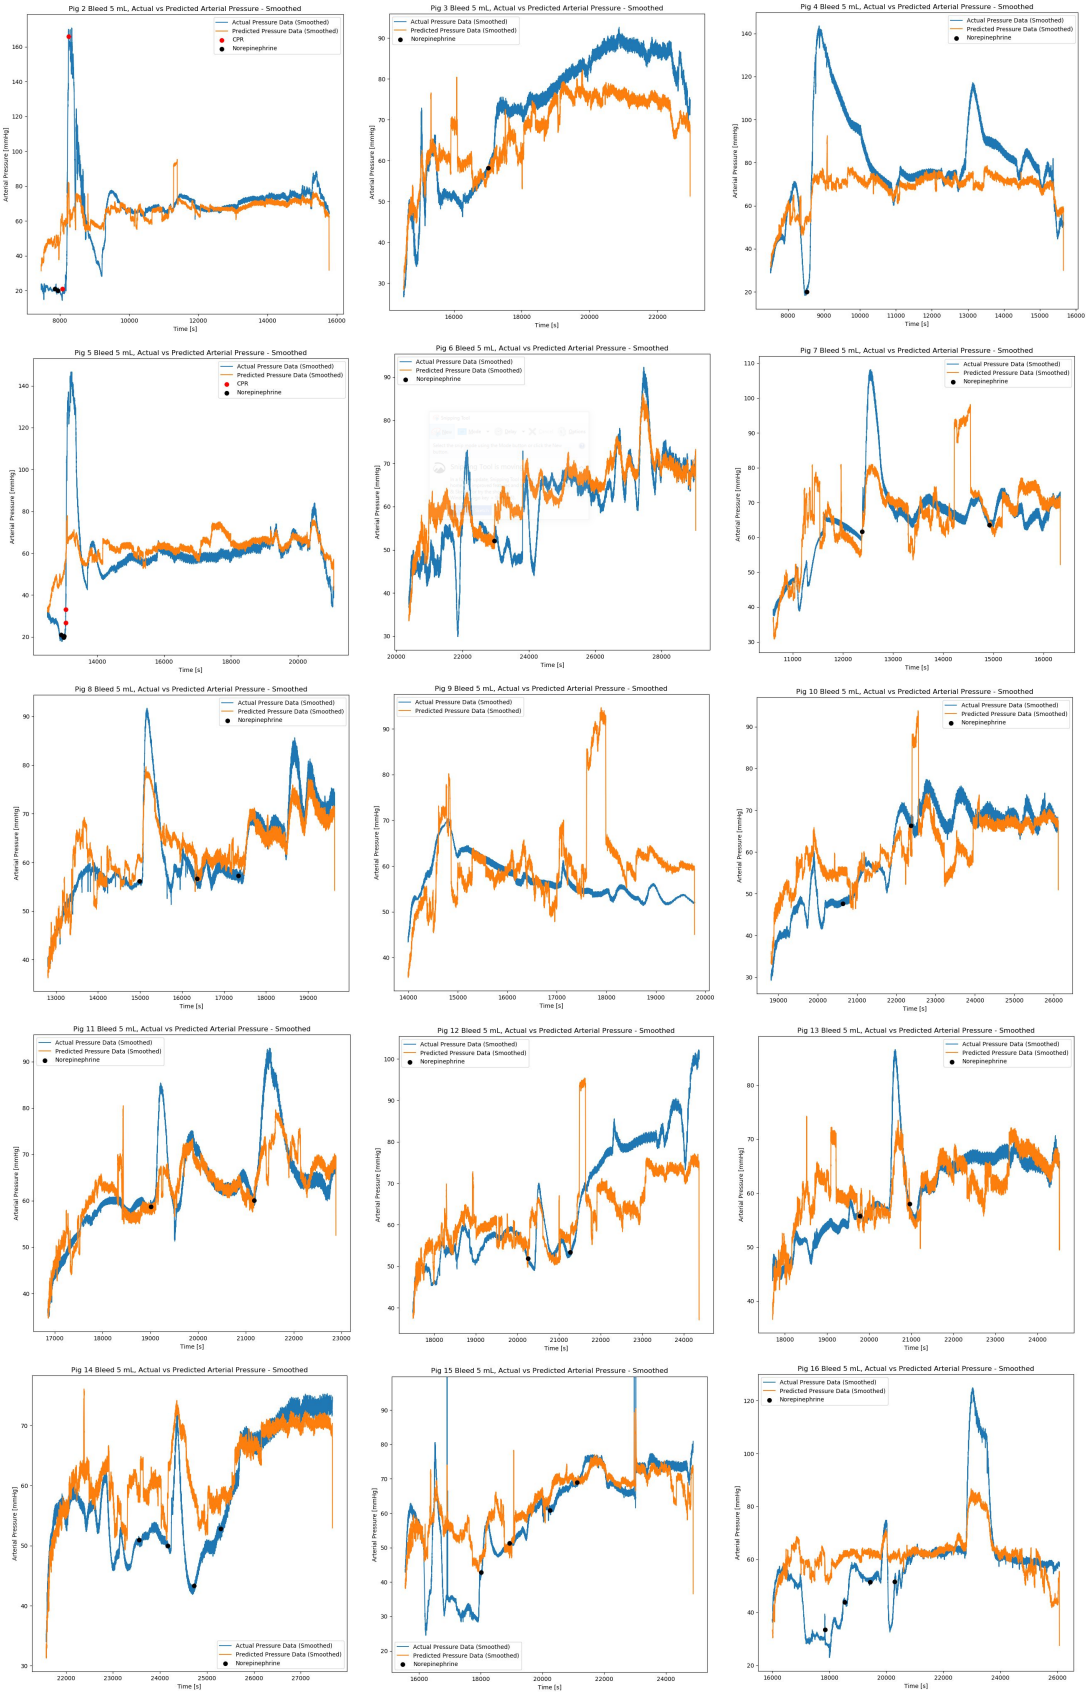

Figure A.1: All smoothed predicted versus actual arterial pressure plots from the 5 mL/min bleed group.

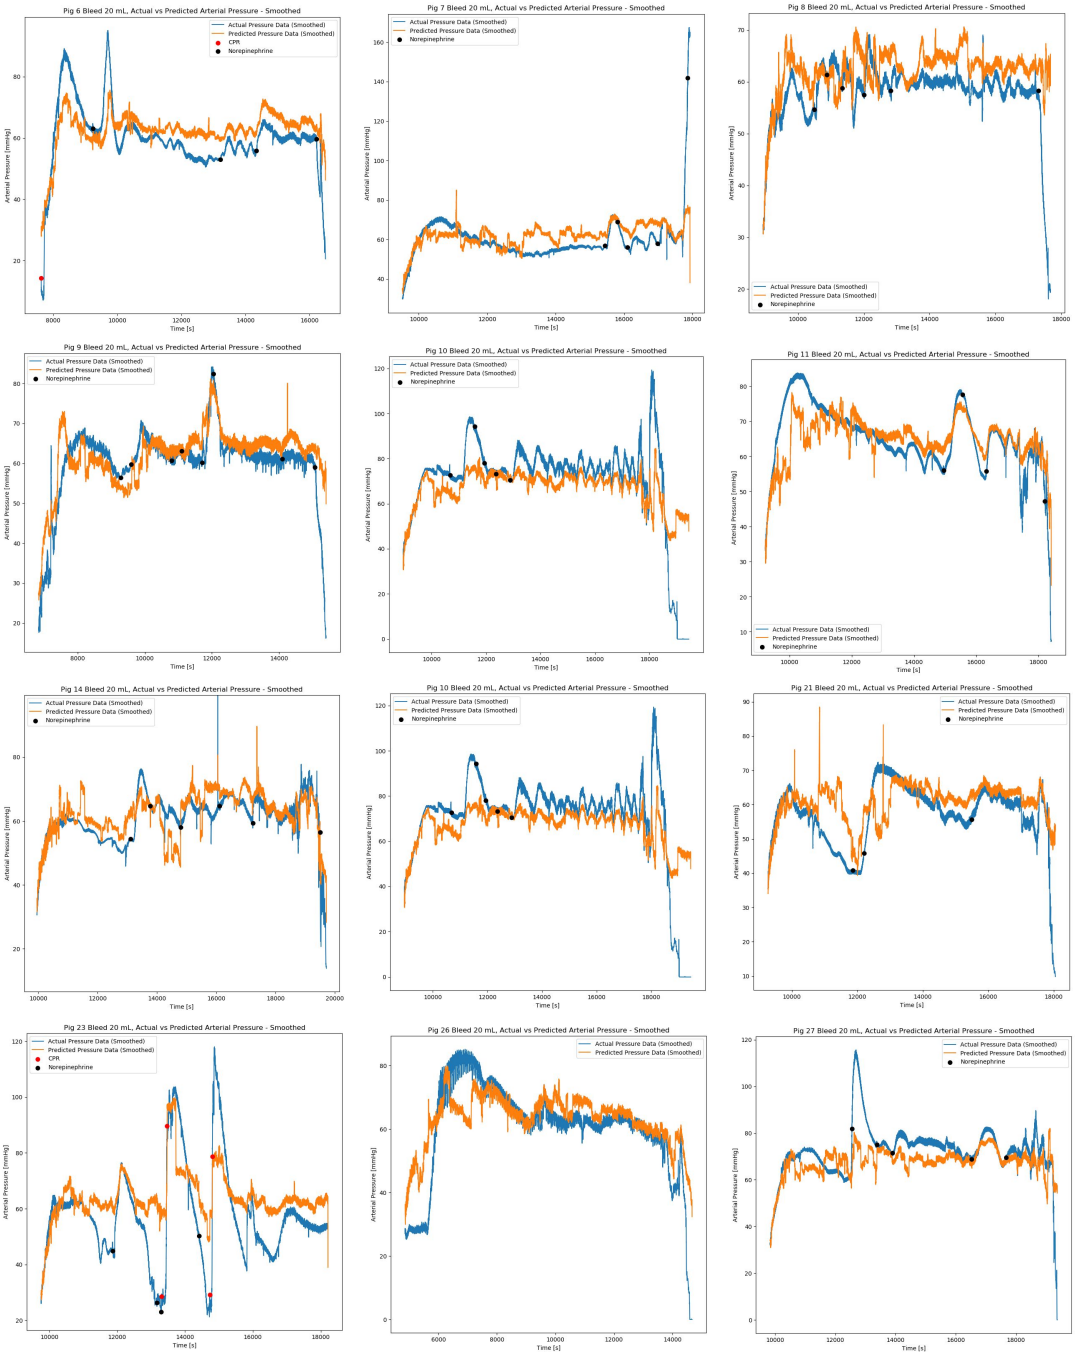
Pig 6 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 7 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 8 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 9 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 10 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 11 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 14 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 10 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 21 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 23 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 26 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 27 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Actual Pressure Data (Smoothed)
Predicted Pressure Data (Smoothed)
CPR
Norepinephrine
Arterial Pressure [mmHg]
Time [s]

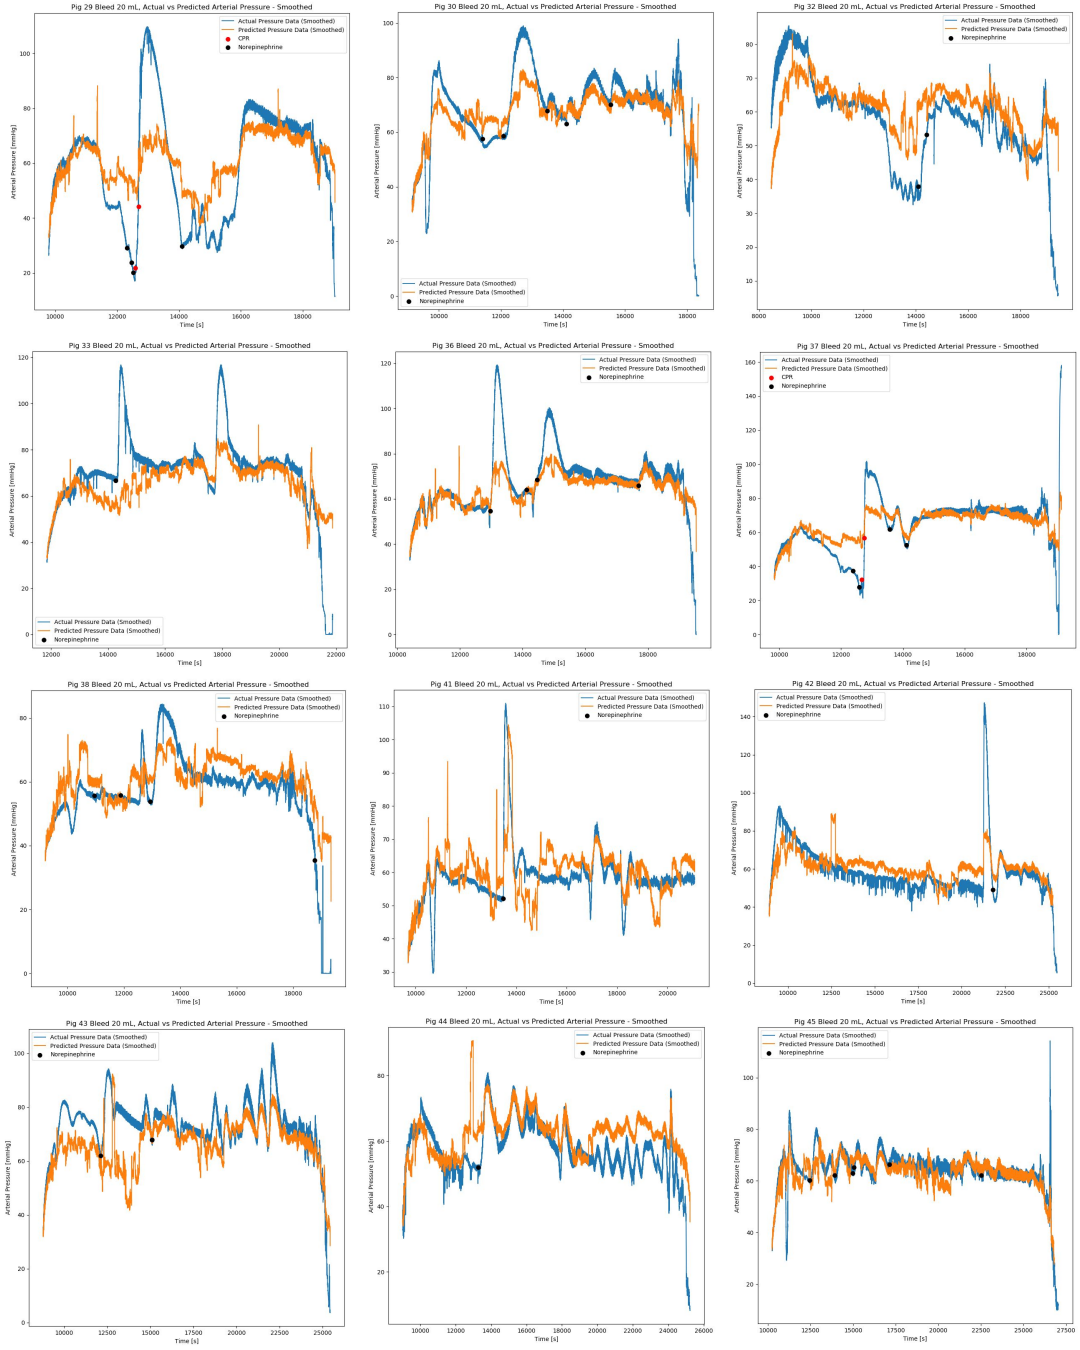
Pig 29 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 30 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 32 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 33 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 36 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 37 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 38 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 41 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 42 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 43 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 44 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Pig 45 Bleed 20 mL, Actual vs Predicted Arterial Pressure - Smoothed
Actual Pressure Data (Smoothed)
Predicted Pressure Data (Smoothed)
CPR
Norepinephrine
Arterial Pressure [mmHg]
Time [s]

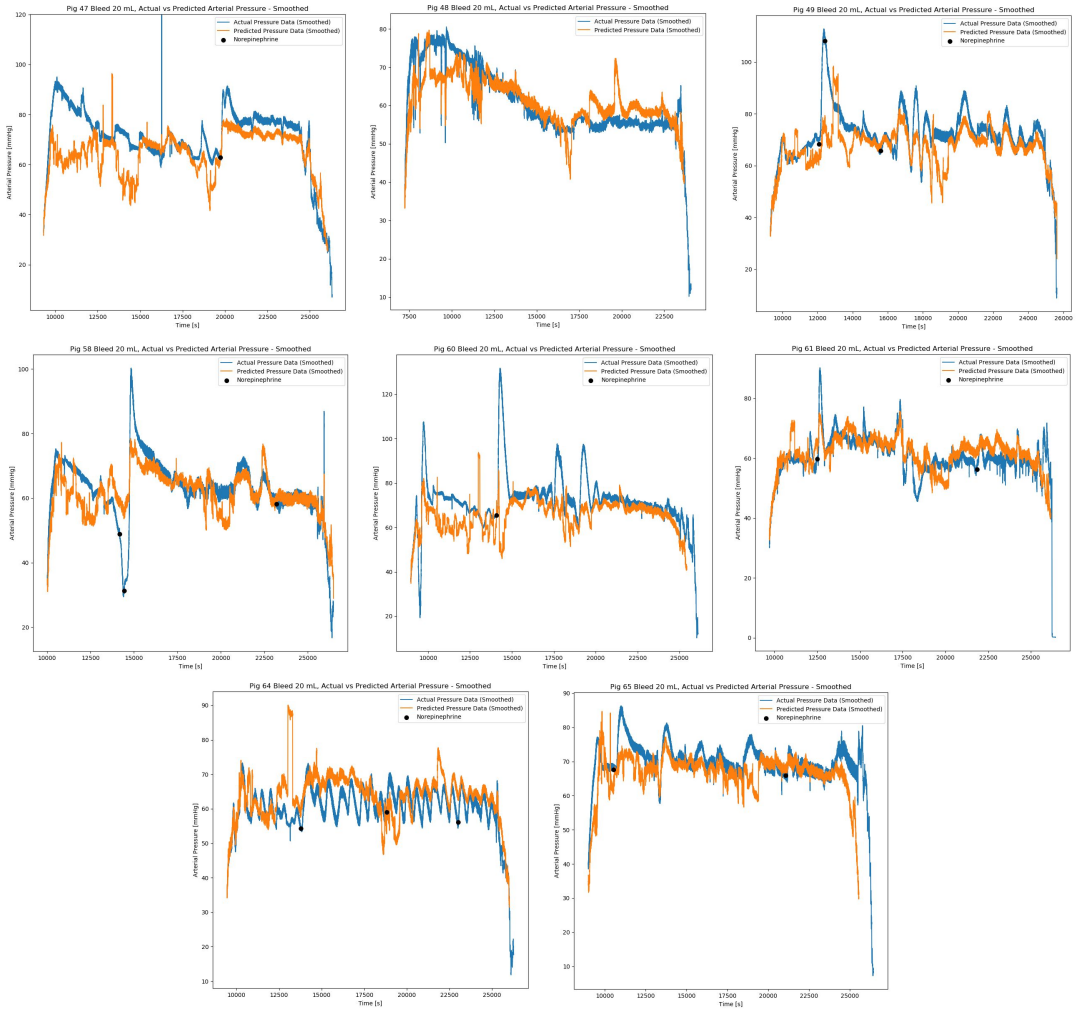

Figure A.2: All smoothed predicted versus actual arterial pressure plots from the 20 mL/min bleed group.

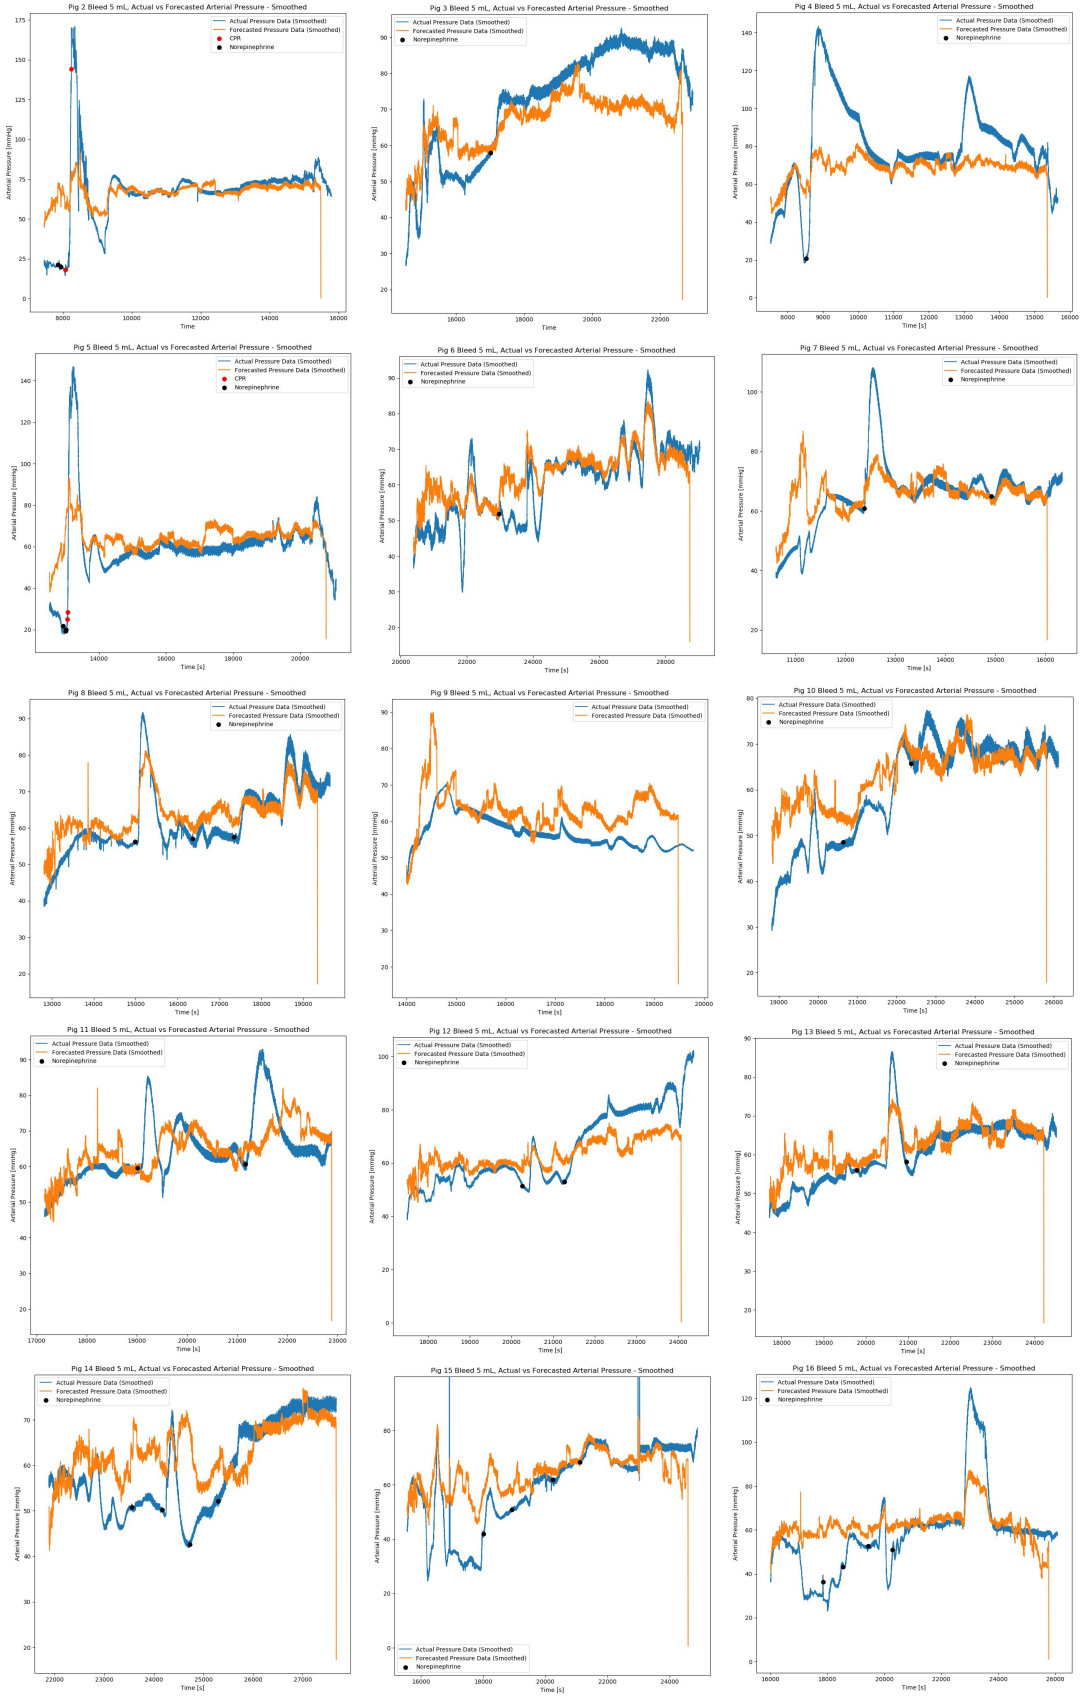

Figure A.3: All smoothed 5-minute forecasted versus actual arterial pressure plots from the 5 mL/min bleed group.

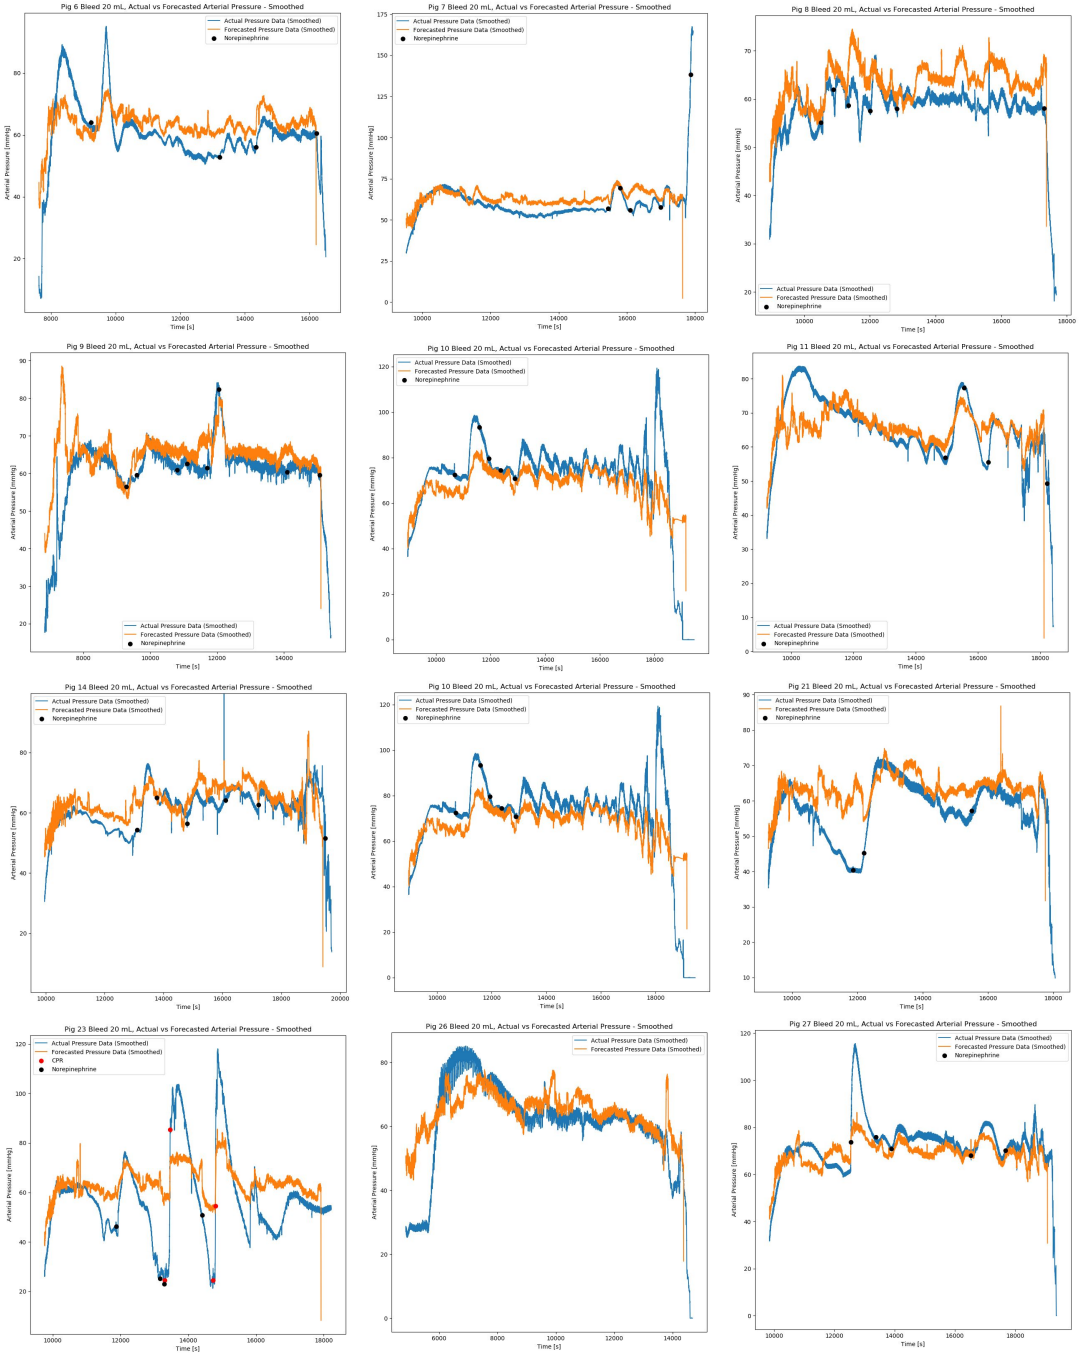
Pig 6 Bleed 20 mL, Actual vs Forecasted Arterial Pressure - Smoothed
Pig 7 Bleed 20 mL, Actual vs Forecasted Arterial Pressure - Smoothed
Pig 8 Bleed 20 mL, Actual vs Forecasted Arterial Pressure - Smoothed
Pig 9 Bleed 20 mL, Actual vs Forecasted Arterial Pressure - Smoothed
Pig 10 Bleed 20 mL, Actual vs Forecasted Arterial Pressure - Smoothed
Pig 11 Bleed 20 mL, Actual vs Forecasted Arterial Pressure - Smoothed
Pig 14 Bleed 20 mL, Actual vs Forecasted Arterial Pressure - Smoothed
Pig 10 Bleed 20 mL, Actual vs Forecasted Arterial Pressure - Smoothed
Pig 21 Bleed 20 mL, Actual vs Forecasted Arterial Pressure - Smoothed
Pig 23 Bleed 20 mL, Actual vs Forecasted Arterial Pressure - Smoothed
Pig 26 Bleed 20 mL, Actual vs Forecasted Arterial Pressure - Smoothed
Pig 27 Bleed 20 mL, Actual vs Forecasted Arterial Pressure - Smoothed
Actual Pressure Data (Smoothed)
Forecasted Pressure Data (Smoothed)
Norepinephrine
CPR
Arterial Pressure [mmHg]
Time [s]

APPENDIX A. APPENDIX

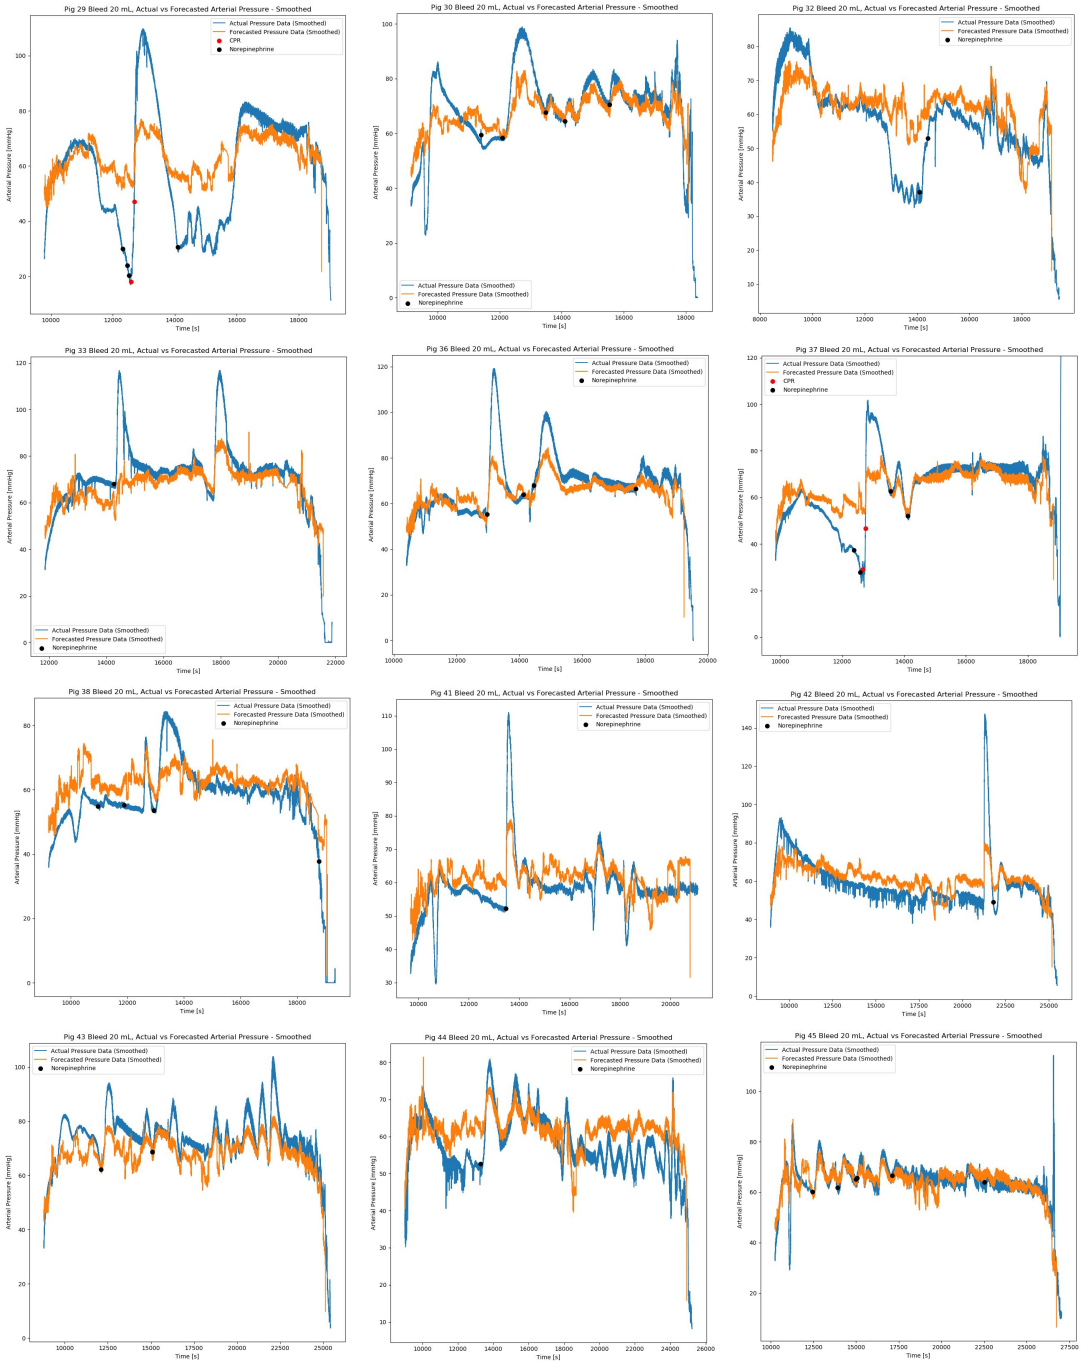

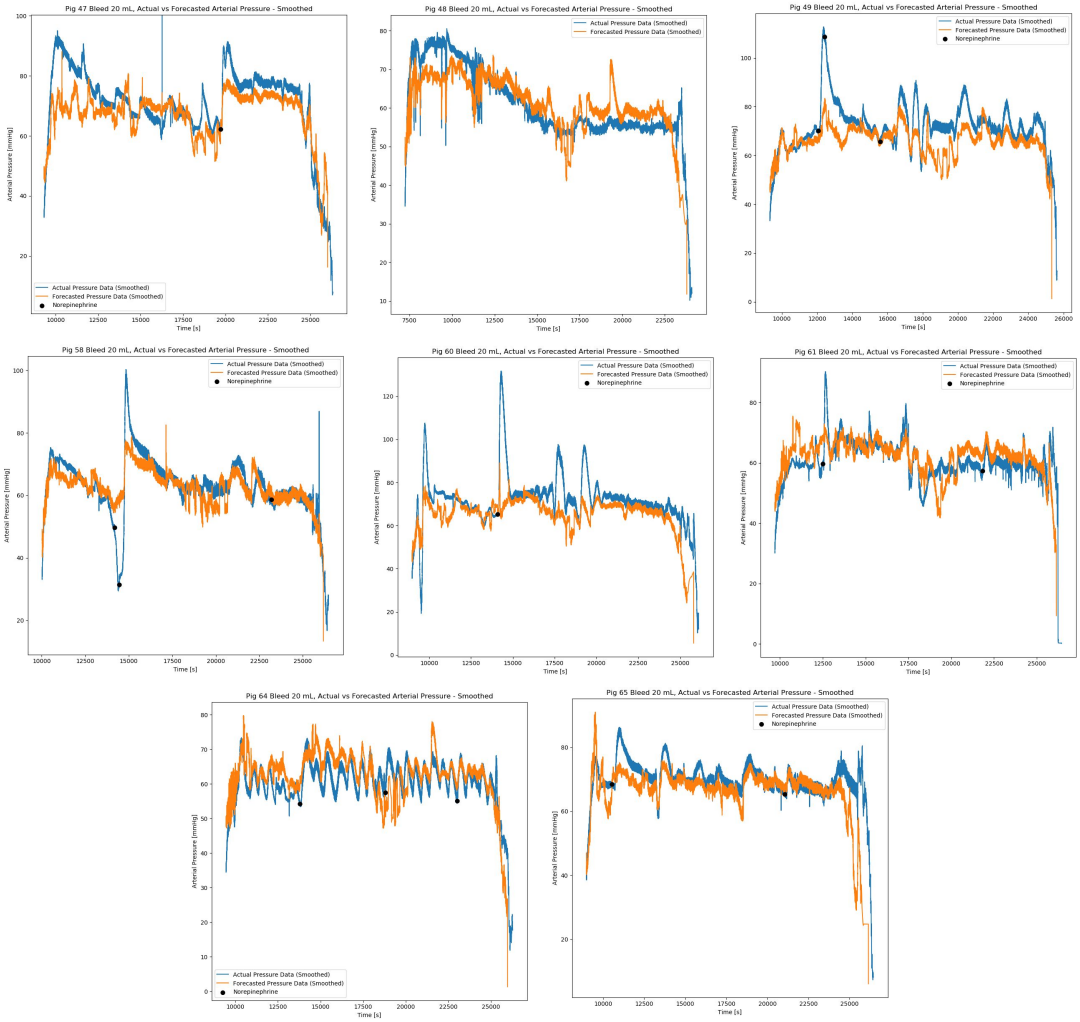

Figure A.4: All smoothed 5-minute forecasted versus actual arterial pressure plots from the 20 mL/min bleed group.

Bibliography

[1] G. Avanzolini, P. Barbini, A. Cappello, and G. Cevenini. Tracking time-varying properties of the systemic vascular bed. *IEEE Transactions on Biomedical Engineering*, 26:373–380, 1989. 1

[2] V. Bikia, T.G. Papaioannou, S. Pagoulatou, G. Rovas, and E. Oikonomou. Noninvasive estimation of aortic hemodynamics and cardiac contractility using machine learning. *Computer Methods in Applied Mechanics and Engineering*, 358, 2020. 1

[3] L. Ferkl and J. Siroky. Ceiling radiant cooling: Comparison of armax and subspace identification modelling methods. *Building and Environment*, 45:205–212, 2010. 1

[4] H. Huang, M. Yang, W. Zang, S. Wu, and Y. Pang. In vitro identification of four-element windkessel models based on iterated unscented kalman filter. *IEEE Transactions on Biomedical Engineering*, 58, 2011. 1, 2.2, 3

[5] S.J. Julier and J.K. Uhlmann. Unscented filtering and nonlinear estimation. *Proceedings of IEEE*, 92:401–422, 2004. 2.2

[6] T. Kind, T. J. C. Faes, and J.W. Lankhaar. Estimation of three- and four-element windkessel parameters using subspace model identification. *IEEE Transactions on Biomedical Engineering*, 57:1531–1538, 2010. 1

[7] G. Kissas, Y. Yang, E. Hwang, W.R. Witschey, J.A. Detre, and P. Perdikaris. Machine learning in cardiovascular flows modeling: Predicting arterial blood pressure from non-invasive 4d flow mri data using physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 358, 2020. 1

[8] F. Lam, H.W. Lu, C.C. Wu, Z. Aliyazicioglu, and J.S. Kang. Use of the kalman filter for aortic pressure waveform noise reduction. *Computational and Mathematical Methods in Medicine*, 2017. 1

[9] L. Ljung. Asymptotic behavior of the extended kalman filter as a parameter estimator for linear system. *IEEE Transactions on Automatic Control*, 24:36–50, 1979. 1

[10] A. Wertz, A.L. Holder, M. Guillame-Bert, G. Clermont, A. Dubrawski, and M.R. Pinsky. Increasing cardiovascular data sampling frequency and referencing it to baseline improve hemorrhage detection. *Critical Care Explorations*, 1:58, 2019. 2, 2.3

[11] N. Westerhof, N. Stergiopulos, and M.I.M. Noble. Snapshots of hemodynamics: an aid for clinical research and graduate education. *New York: Springer*, 2005. 1

[12] N. Westerhof, J.W. Lankhaar, and B. Westerhof. The arterial windkessel. *Medical Biological Engineering Computing*, 47:131–141, 2009. 1

[13] Y. Yang and P. Perdikaris. Adversarial uncertainty quantification in physics-informed neural networks. *Journal of Computational Physics*, 394:136–152, 2019. 1

[14] Y.C. Yu, J.R. Bostom, M.A. Simaan, and J.F. Antaki. Minimally invasive estimation of systemic vascular parameters. *Annals of Biomedical Engineering*, 29:595–606, 2001. 1