PRECISION CRITICAL CARE MANAGEMENT OF BLOOD PRESSURE IN STROKE PATIENTS USING DYNAMIC LINEAR MODELS

by

Yuzhe Liu

BS, Cornell University, 2010

MEng, Cornell University, 2011

MS, University of Pittsburgh, 2016

Submitted to the Graduate Faculty of

the School of Medicine in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

University of Pittsburgh

2018

UNIVERSITY OF PITTSBURGH

SCHOOL OF MEDICINE

This dissertation was presented

by

Yuzhe Liu

It was defended on

August 17, 2018

and approved by

Vanathi Gopalakrishnan, PhD, Associate Professor, Department of Biomedical Informatics

Douglas P. Landsittel, PhD, Professor, Department of Biomedical Informatics

Sherry H.Y. Chou, MD, MMsc, Associate Professor, Department of Critical Care Medicine

Shyam Visweswaran, MD, PhD, Associate Professor, Department of Biomedical Informatics

Dissertation Director: Vanathi Gopalakrishnan, PhD, Associate Professor, Department of Biomedical Informatics

Copyright © by Yuzhe Liu
2018

PRECISION CRITICAL CARE MANAGEMENT OF BLOOD PRESSURE IN STROKE PATIENTS USING DYNAMIC LINEAR MODELS

Yuzhe Liu, PhD

University of Pittsburgh, 2018

For stroke patients in the ICU, optimal blood pressure management remains an open area of research. Numerous observational studies and clinical trials on both ischemic and hemorrhagic stroke have resulted in conflicting evidence for the benefit of lowering blood pressure. A major limitation of these studies is their inability to account for and distinguish the effects of physician-initiated blood pressure treatment versus a patient's spontaneous blood pressure time course. We address this problem with the Acute Intervention Model of Blood Pressure (AIM-BP) framework: an individualized, human interpretable model of blood pressure management in the acute care setting. The framework consists of two components: one, a model of blood pressure homeostasis and the various effects that perturb it; and two, a parameter estimator that can learn clinically important model parameters on a patient by patient basis. By estimating the parameters of the AIM-BP model for a given patient, the effectiveness of antihypertensive medication can be quantified separately from the patient's spontaneous blood pressure trends. We hypothesize that the AIM-BP is a sufficient framework for estimating parameters of a homeostasis perturbation model of a stroke patient's blood pressure time course and the AIM-BP parameter estimator can do so more accurately and consistently than a state-of-the-art maximum likelihood estimation method. The first part of this hypothesis is proved mathematically, while the second is demonstrated using simulated clinical scenarios modeled on stroke patients from two ICU datasets. Finally, the ability of the AIM-BP framework to model real world patients is demonstrated using several examples from a UPMC dataset.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

PREFACE

This work is the culmination of a rather circuitous route exploring many different areas of biomedical informatics, from cancer research to imaging to medical records, from lungs to hearts to brains, from rule learning to Bayesian networks to dynamic linear models, and from guessing missing data to simulating completely new data (okay, I guess that last one isn't really so different). Reaching this point wouldn't have been possible without the constant enthusiastic support and mentoring from my advisor, Vanathi. I'd like to thank my committee members as well, whose domain knowledge and career advice shaped this thesis into what it is today. And seriously, thank you so much Toni for helping me out during all the times I've come to ask you completely lost on what I'm supposed to be doing.

I'm lucky to have the support of all my friends and family throughout these few years. (Yes, I'm going to brag about having friends. It's my thesis, I can do what I want.) From my MSTP class, who have been in it with me for 6 years and counting, to my med school friends, who have long since moved on and gotten real jobs, and others I've met along the way - thanks for keeping me grounded in what's really important. Soma and Becca, thanks for all the delicious food and letting me borrow your smoker and peer pressuring me to ride a bike to DC. Thanks to my parents as well, who have unerringly supported me through all these years even though I'm hitting 30 and still in school. And finally, thanks to my girlfriend Kirstin, who has cheered me on in this race to the finish, and maybe glossing over the status of some hills in the middle.

Oh, and of course, thanks to the NIH and the National Library of Medicine (NLM T15LM007059) for paying me. You guys are great.

GLOSSARY

- **Accuracy:** Accuracy is a measure of performance of a parameter estimator. In this thesis, we define accuracy as absolute error from ground truth.
- **AIM-BP:** Acute Intervention Model of Blood Pressure. This framework models the parameters that govern a patient's blood pressure during the course of their ICU stay. The framework consists of two components: The AIM-BP model and the parameter estimator.
- **ARMS:** Adaptive Rejection Metropolis Sampling. Adaptive rejection sampling[1] is a method of sampling from a log-concave distribution that adaptively adjusts the proposal distribution. ARMS extends adaptive rejection sampling to non-log-concave distributions.[2]
- **BP, SBP, DBP:** Blood pressure, systolic blood pressure, diastolic blood pressure. Systolic blood pressure is the blood pressure during the contraction of the heart, or systole. Diastolic blood pressure is the blood pressure during the relaxation phase of the heart, or diastole. Blood pressure can be measured externally using a pressure cuff or internally using an arterial line. Blood pressure measurements may vary depending on a patient's position and the location of the measuring instrument.
- **Consistency:** Consistency is a measure of performance of a parameter estimator. In this thesis, we define consistency as the variance of estimated parameters from a set of simulated datasets generated via the same ground truth mechanism with identical ground truth parameter configurations.
- **DLM:** Dynamic Linear Model. A dynamic linear model, or state space model, is an auto-regressive Markov model with a hidden layer.
- **EMR:** Electronic Medical Record. The electronic medical record, or electronic health

record (EHR), is a digital representation of a patient's interactions with a health system, capturing information about their admissions and office visits.

- **FFBS:** Forward filtering backward sampling algorithm. The forward filtering backward sampling algorithm is a method to sample the hidden states in a dynamic linear model given the observed states and model parameters. It works by running a Kalman filter or smoother forward (forward filtering) then sampling backwards (backward sampling).
- **Homeostasis perturbation model:** A model of various mechanisms that dictate an individual's homeostatic blood pressure over time. The homeostasis perturbation model captures a theoretical blood pressure that an individual's actual blood pressure tends to trend towards, and includes effects such as disease perturbations and medication effects.
- **IA2RMS:** Independent Doubly Adaptive Rejection Metropolis Sampling. IA2RMS[3] is an improvement on ARMS in terms of efficiency, allowing for faster sampling.
- **ICU:** Intensive care unit. Intensive care units typically monitor high risk patients during the course of their care. Compared to floor beds, ICUs will have more regular vital signs taken and recorded.
- **ICH:** Intracerebral hemorrhage. A primary hemorrhage inside the brain parenchyma. Intracerebral hemorrhage is a severe form of hemorrhagic stroke with a high mortality rate. Depending on hemorrhage location, a patient can present with varying blood pressure behavior.
- **MAP:** Mean arterial pressure. Calculated as $(1/3{*}SBP{+}2/3{*}DBP)$, the mean arterial pressure is a rough estimate of the mean pressure in the arteries over the cardiac cycle, with roughly 1/3 of the cardiac cycle represented by systole and the rest by diastole.
- **MCMC:** Markov Chain Monte Carlo. Markov Chain Monte Carlo are a family of stochastic methods for estimating characteristics of a difficult posterior probability distribution through sampling.
- **MIMIC III:** Medical Information Mart for Intensive Care III database. The MIMIC III database is a comprehensive electronic medical record database of critical care patients from Mass General and Brigham and Women's in Boston. The database includes chart data, notes, medications, procedures, and labs for patients during the course of their ICU stay.

- **Perturbation:** A perturbation is a mechanism that modifies a patient's blood pressure homeostasis away from a constant level. These can include medications or can be a result of disease, such as stroke.
- **Sensitivity Analysis:** A sensitivity analysis studies how desired outputs change depending on a range of inputs. In this thesis, we perform a sensitivity analysis on accuracy of estimated parameters depending on the range of several ground truth parameters used to simulate data.
- **Sufficiency:** Sufficiency in the context of this thesis is defined as the ability of one model to completely characterize a second model. The AIM-BP model is a sufficient model for characterizing the homeostasis perturbation model - in other words, it is a more general model.

1.0 INTRODUCTION

Proper critical care control of blood pressure after a stroke remains an area of ongoing research, as balancing the need for cerebral perfusion and the risk for adverse effects of high blood pressure make its management a challenging problem in both ischemic and hemorrhagic strokes.[4–6] Stroke patients often exhibit elevated blood pressure during initial treatment, but it is as yet unclear whether lowering blood pressure results in worse, better, or no difference in outcomes.[5,7–10] In the context of ischemic stroke, it is also still unclear whether raising blood pressure to correct hypotension is beneficial to outcomes, though new guidelines recommend correcting hypotension enough to ensure adequate perfusion of organ systems.[11,12] The inconsistency of these studies suggests that optimal blood pressure management likely differs on a patient to patient basis.

1.1 THE PROBLEM

Current clinical trials and observational studies do not take into account individual differences in blood pressure dynamics and do not separate out individualized effects of physician-initiated interventions on a patient's blood pressure from a patient's baseline blood pressure behavior. As we will explore in the next chapter, most clinical studies on the effect of blood pressure management in stroke look at the combined effects of medical management and natural time course, using summary metrics such as mean blood pressure or percentage decrease over 1 day. These summary metrics are imprecise, ignoring the many potential sources of blood pressure perturbations that confound measurements, such as spontaneous blood pressure changes often found in stroke patients. In order to better understand the role

of blood pressure in the management of stroke patients, a more precise model must be built that can begin to account for these perturbations.

The irregular nature of temporal data from the electronic medical record makes studying it challenging. Multiple variables may be recorded asynchronously, and recordings of a single variable often contain periods of time in which no values are recorded. Various approaches have been proposed for analyzing temporal data in a machine-interpretable way, such as Fourier transformation,[13] trend pattern mining,[14] or regularizing time data into discrete time windows.[15] In addition, recent work on deep learning has shown promising prediction and missing value imputation performance using recurrent neural networks, particularly long short term memory models (LSTMs).[16–18] However, most of these models are focused on maximizing classification performance, and as a result often transform data into a representation that is not easily human-interpretable. This is particularly an issue with the recurrent neural networks, though their classification performances are markedly better compared to traditional techniques. Instead of focusing on classification performance in the vein of machine learning research, we sought to design a clinically informative model - one that is human-interpretable, accurate, and provides disease-specific information.

A clinically useful model of blood pressure management for stroke in the ICU setting must satisfy the following requirements:

1. Model multivariate time series data of blood pressure and heart rate at a minute to hour timescale. This data will include irregularly measured or missing values, as well as measurements of the same variable via multiple modalities with their various measurement noises.
2. Provide human-interpretable, clinically relevant model parameters.
3. Model body homeostasis mechanisms and disease state perturbations to homeostasis, such as the acute increase followed by spontaneous decrease of blood pressure often found in stroke patients.
4. Model the effects of ICU interventions, such as blood pressure medication.

Currently, no model exists that satisfies all of these requirements. We are particularly interested in separating out requirements 3 and 4. Separating out the effects of drugs versus

a patient's natural blood pressure time course is necessary in the context of several research questions: One, which drugs are more effective for which people in the management of blood pressure after a stroke? Two, can we better determine drug effectiveness during acute care compared to current clinician judgment? Three, does the relationship between managed blood pressure and the natural homeostasis level correlate with outcomes?

1.2 THE APPROACH

To address this challenge, we propose the Acute Intervention Model of Blood Pressure (AIM-BP) framework: an individualized, multivariate, human interpretable framework that models a patient's blood pressure time course and response to interventions over the entirety of their ICU stay. The framework consists of two components: The AIM-BP model and the AIM-BP parameter estimator. Different parameters of the AIM-BP model capture different desired attributes of an individual patient's blood pressure time course, such as the effect of individual medications, the homeostasis point that their blood pressure tends to self-regulate to, and the amount of variablity in their natural blood pressure (Figure 1).

The AIM-BP model is comprised of two parts: The first is a mathematical model that combines disease state perturbation and medication pharmacodynamics, which we call the homeostasis perturbation model. The homeostasis perturbation model is used as input into the second part of the AIM-BP model, a constrained dynamic linear model (DLM) which we will refer to as the AIM-BP temporal model. DLMs are a popular class of time series models, which have been used to predict post-surgical CBC values,[19,20] perform outbreak surveillance,[21] and estimate values in a noisy setting in a wide variety of medical and non-medical problems.[22] As DLMs are generative models, the AIM-BP model is capable of simulating data from a set of known parameters. Although DLMs for blood pressure have been studied on the millisecond time scale for device design and false alarm monitoring,[23,24] to our knowledge this is the first time it has been applied in the context of personalized blood pressure management in critical care. Our dynamic linear model design works from first principles and is the first of its kind to address the issue of blood pressure management in

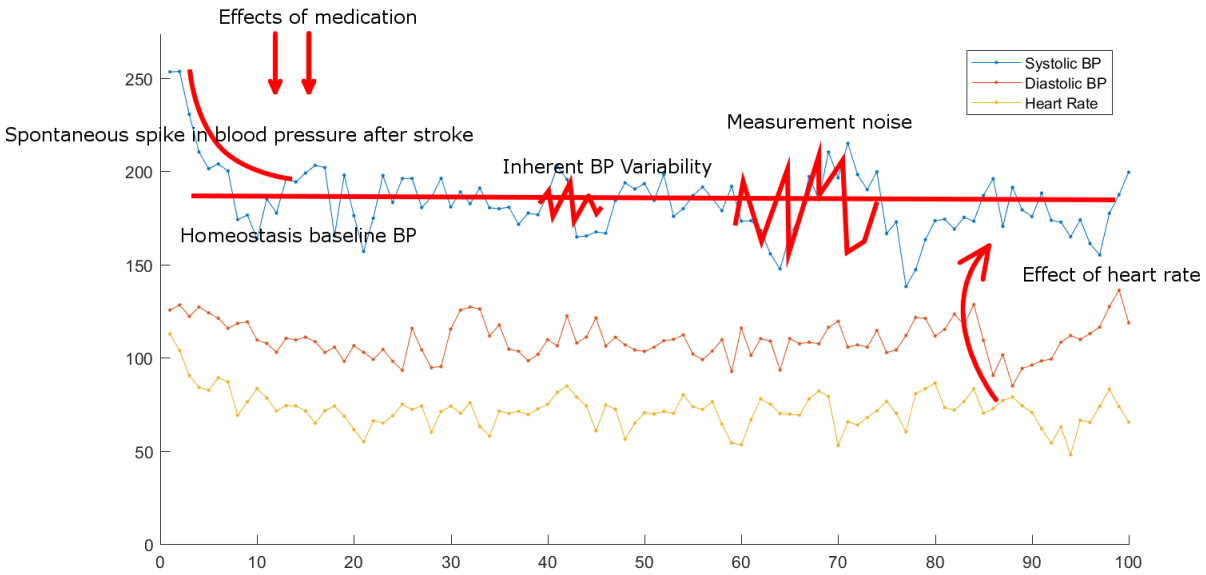

Figure 1: Examples of AIM-BP model parameters that separate out mechanisms that affect an ICU patient's blood pressure course. These parameters can be estimated using the AIM-BP parameter estimator.

the acute care of stroke.

Our choice of a dynamic linear model as the basis for our framework addresses the requirements laid out in the previous section. We will describe how a DLM works in detail in the following chapter. DLMs provide an explicit mechanism to model measurement noise, handle missing values and uncertainty, and combine multiple modalities of measurement. They can be designed to capture a homeostatic process in their auto-regressive hidden layer. In addition, DLMs are easily extensible, making adding or removing modeling components easy. In this sense, while the AIM-BP framework we present in this thesis as a whole is disease-specific, it is comprised of separate disease-specific and disease-agnostic components. The homeostasis perturbation model is disease-specific, and can be extended or modified for other diseases. The dynamic linear model used for the AIM-BP temporal model, on the other hand, is disease-agnostic.

The AIM-BP parameter estimator is a Markov Chain Monte Carlo (MCMC) sampling

method for recovering specific model parameters for each patient given ICU data on vitals and interventions. The parameter estimator will need to be able to accurately recover model parameters and handle missing and multi-modal data.

The combined AIM-BP framework is shown in Figure 2.

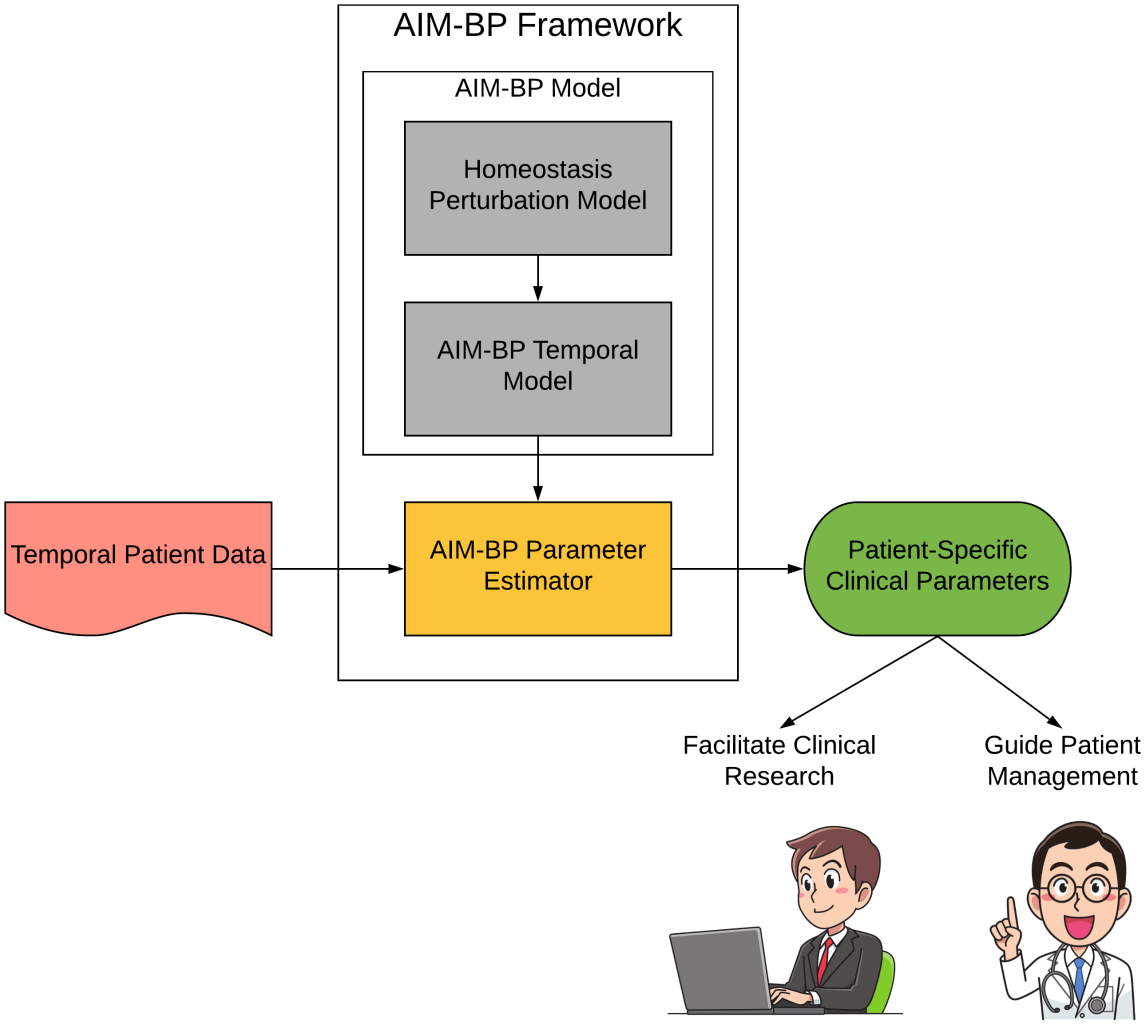

Figure 2: Design the AIM-BP model and develop the parameter estimator. The parameter estimator takes in data from patients and outputs AIM-BP model parameters that can eventually inform a physician's decision making.

1.2.1 Thesis

We hypothesize that the AIM-BP is a sufficient framework for estimating parameters of a homeostasis perturbation model of a stroke patient's blood pressure time course and the AIM-BP parameter estimator can do so more accurately and consistently than a state-of-the-art maximum likelihood estimation method. In this thesis, we address this hypothesis through the following steps:

1. **Develop the AIM-BP framework by designing the AIM-BP model and developing the AIM-BP parameter estimator.** We approached this in two steps:
   a. **Design the AIM-BP model.** We designed a human interpretable discrete homeostasis perturbation model and dynamic linear model of blood pressure in the ICU based on physiological principles. We analyzed real world ICU data and identify trends for charted blood pressure time series. These trends were incorporated into the constraints and priors for the dynamic linear model.
   b. **Develop the AIM-BP parameter estimator.** A MCMC-based DLM parameter estimator was developed to learn AIM-BP model parameters for individual patients.
2. **Evaluate the performance of the AIM-BP parameter estimator.** First, we performed a sensitivity analysis on a set of clinically important model parameters. Then, a set of clinical scenarios were created based on expert knowledge and analysis of clinical data. Data was simulated with ground truth parameters for each scenario. The AIM-BP parameter estimator was run on the simulated data and estimated parameters was compared to ground truth parameters. The accuracy and consistency of AIM-BP estimated parameters was compared with a state-of-the-art parameter estimation technique. This process is illustrated in Figure 3.

Through completing the above work, we make the following strong claims:

1. Assuming a blood pressure homeostasis perturbation model to be true for stroke patients, the AIM-BP framework is a sufficient framework for capturing parameters of this homeostasis perturbation model.

Figure 3: Evaluate the performance of the AIM-BP parameter estimator through example clinical scenarios.

2. The AIM-BP parameter estimator provides more accurate and consistent estimates of model parameters than a state-of-the-art non-linear least squares estimation method.

1.3 SIGNIFICANCE

The AIM-BP framework is a novel modeling system of blood pressure in stroke built from first principles. It is the first of its kind to characterize blood pressure behavior of stroke patients in the critical care setting in a quantifiable, human interpretable, integrative, and extensible way. By quantifiable, we mean that the AIM-BP framework provides objective, quantitative measurements for blood pressure parameters. Human interpretable means that these blood pressure parameters possess inherent meaning to a clinician. By integrative, we mean that the AIM-BP model integrates multiple sources of blood pressure perturbation into a single model, and the parameter estimator can separate out the individual sources of perturbations from raw chart data. Finally, extensible means that the AIM-BP model is easily adapted to model additional effects or different diseases all together.

These qualities enable the AIM-BP framework to offer more precise and informative metrics than the summary metrics currently used in stroke blood pressure management research, with the same or more clinical utility. Stroke is a syndrome with a set of underlying causes that result in similar neurological deficits, including altered autonomic control of blood pressure. AIM-BP allows the precise study of these alterations in autonomic control of blood pressure as they relate to the underlying stroke mechanisms. Increased precision affords the clinical researcher greater opportunity to identify relationships between variables that were previously not identifiable due to noise. More informative metrics allow the researcher to study a greater variety of relationships between variables that were not previously possible. The AIM-BP framework allows clinicians to more accurately quantify the effects of medical therapy they've chosen, and provides them with a more objective metric for clinical decision making. Indeed, with the AIM-BP framework, we introduce the idea that a quantitative pharmacodynamics model can be fit to clinical data obtained in a non-research environment and used to aid decision making.

With this work, we make significant contributions to the field of biomedical informatics as well. While the study of electronic medical record data is a foundation of the field of informatics, most work has focused on machine learning or pattern mining from structured data, with limited work on the temporal aspects.[25] Current temporal methods for EMR data

are typically tailored towards improving machine learning performance, and as a result have limited human interpretability. The AIM-BP framework provides a solution that deals with the challenges of temporal EMR data while providing human interpretable, disease-specific models of patient behavior. It opens up the study of diseases through their critical care management on a personalized level.

1.4 THESIS OVERVIEW

This thesis will be split into the following chapters: In Chapter 2, the background on blood pressure management in stroke, blood pressure modeling, and the models underpinning the AIM-BP framework will be presented. Chapter 3 will describe a preliminary analysis of real world ICU stroke data to inform the development of the AIM-BP framework, define the AIM-BP model, and prove its sufficiency to capture blood pressure homeostasis perturbation parameters. Chapter 4 will evaluate the AIM-BP parameter estimator on several clinical scenarios. We will then flesh out the model with missing values and multi-modal data in Chapter 5, and demonstrate its application to several real world stroke patients. Finally, we will conclude and provide directions for future work.

2.0 BACKGROUND

2.1 STROKE MANAGEMENT IN THE ICU

The development of neurocritical care units has contributed to the decreasing mortality rate of stroke in the US, from the 3rd leading cause of death in the US before 2008 to the 5th in 2017.[8,26] Though ischemic strokes make up the majority of stroke cases, the severity of hemorrhagic strokes compared to ischemic means that neuro ICUs see similar amounts of both types. Both types of strokes share many similar management concerns, such as the need to maintain perfusion and the risk for additional bleeding. Severe ischemic strokes have the potential to hemorrhagically convert, while hemorrhagic strokes may present with cerebral ischemia in areas remote from the hematoma.[27] In management of ischemic stroke, the established use of fibrinolytics and increased use of mechanical thrombolysis have shown efficacy in the treatment of stroke but place patients at higher risk for complications that benefit from close montioring in neuro ICU units. In addition, patients with severe strokes that involve large amounts of territory require monitoring due to risk of delayed cerebral edema.[28] Despite advances in practice, the critical care management of stroke is still an evolving field of research, with current challenges including blood pressure, temperature, and glucose management.

In this thesis, we will focus on blood pressure management for stroke in the ICU. While stroke is broadly classified as either ischemic or hemorrhagic, in reality it is more accurately thought of as a syndrome, with different underlying disease mechanisms depending on the location and type of injury to the brain. For the purposes of blood pressure management, it is particularly important to study how the location and type of stroke affects autonomic control of blood pressure, particularly in cases where the autonomic centers of the brain

stem are injured. More than 60% of patients being evaluated for stroke had elevated blood pressure higher than 140 mmHg systolic,[29] and blood pressure often decreases spontaneously during the acute phase.[8,30,31] Potential causes of this elevation and subsequent spontaneous decrease in blood pressure have not been thoroughly studied, but may include injury to autonomic control centers, stress response to injury or reflex response to hypoperfusion, or stress from the hospital visit.[32,33]

Despite decades of research, no definitive conclusion has been reached regarding how to treat high blood pressure.[34] Many observational studies as well as randomized trials have been done to study blood pressure management in stroke,[5,35–43] the results of which have shown beneficial, detrimental, and inconclusive effects of antihypertensives. For example, an initial randomized trial (ACCESS study) using the angiotensin II receptor blocker candesartan suggested patients who started the drug had lower mortality and fewer vascular complications, although the study was ended early.[41] However, a larger study (SCAST) using the same drug found no beneficial effect, and possibly even a harmful effect.[42] The evidence is stronger for the benefit of blood pressure lowering in hemorrhagic strokes compared to ischemic strokes. We note that although the term hemorrhagic stroke can mean both intracerebral hemorrhage (ICH) and subarachnoid hemorrhage (SAH), most studies on hemorrhagic stroke focus on ICH, as SAH is more readily treatable. For the purposes of this work, we will refer to hemorrhagic stroke and intracerebral hemorrhage interchangeably.

A variety of outcome metrics can be used to determine the effects of blood pressure management in these studies. These include events like death, neurological deterioration, hemorrhagic conversion, or cardiovascular events, or outcome scores. These outcome scores can measure impairment (National Institutes of Health Stroke Scale (NIHSS)), or function (modified Rankin Scale (mRS)).

2.1.1 Current Guidelines for Stroke Management

Current AHA guidelines for acute blood pressure management in ischemic stroke are as follows:[8,12]

1. If no fibrinolytic therapy has been used or planned, the recommendation is to not lower

blood pressure in the first 24 hours unless the blood pressure is above 220/120 or another medical condition would benefit from lowered blood pressure.

2. If fibrinolytics are being considered, the recommendation is to gently bring blood pressure down to below 185/110 mmHg to qualify for rtPA therapy. Once rtPA is given, the blood pressure must be maintained below 180/105 mmHg.

In the setting in which an ischemic stroke has converted to a hemorrhagic stroke, blood pressure management guidelines have been extrapolated from the primary intracerebral hemorrhage (ICH) setting. In light of the need to balance hematoma growth with maintaining adequate perfusion, it is recommended to maintain systolic blood pressure under 180 mmHg.[26]

Current AHA guidelines for acute blood pressure management in primary intracerebral hemorrhage are as follows:[44]

1. For patients presenting with a systolic blood pressure between 150 to 220 mmHg, acute lowering to 140 mmHg is safe unless otherwise contraindicated.
2. For patients presenting with a systolic blood pressure greater than 220 mmHg, aggressive blood pressure reduction may be reasonably considered.

2.1.2 Recent Antihypertensive Trials in Stroke

Beyond the studies and trials that have been accounted for in the current guidelines, several additional trials are ongoing for both ischemic and hemorrhagic stroke. We review these trials here briefly:

1. **ENCHANTED Trial:** The Enhanced Control of Hypertension and Thrombolysis Stroke Study (ENCHANTED) studies both the effect of intravenous rtPA and the effect of early blood pressure lowering in ischemic stroke. Blood pressure management was randomized to one of two groups: Intensive management, which targeted systolic pressure less than 140 mmHg, or conventional, which targeted systolic pressure less than 180 mmHg. The blood pressure lowering part of the trial is due to be completed in 2018, though preliminary data suggests that blood pressure lowering in conjunction with rtPA had no effect on the rates of hemorrhagic conversion.[45,46]

2. **INTERACT2 Trial:** The Intensive Blood Pressure Reduction in Acute Cerebral Hemorrhage Trial 2 (INTERACT2) studied the effects of an aggressive blood pressure lowering scheme with a target of less than 140 mmHg systolic compared to a conventional scheme, with a target of less than 180 mmHg systolic. A possible slight benefit in death or major disability outcomes (OR 0.75-1.01, 95% CI, $p = 0.06$) as well as modified Rankin Score (OR 0.77-1.00, 95% CI, $p = 0.04$) was found in the aggressive treatment group compared to the conventional group. A significant but small increase in the EQ-5D health quality of life index score was found in the aggressive treatment group (0.60 ± 0.39) compared to the conventional group (0.55 ± 0.40) with $p = 0.002$.[47]
3. **ATACH II Trial:** Like the INTERACT2 trial, the Antihypertensive Treatment of Acute Cerebral Hemorrhage (ATACH II) trial studied the effects of an aggressive blood pressure lowering scheme compared to a conventional scheme (with the same maximum cutoffs for blood pressure) in primary hemorrhagic stroke. The study was discontinued for futility as no difference in death or disability between the two groups were found.[48]
4. **ENOS Trial:** The Efficacy of Nitric Oxide in Stroke (ENOS) trial studied the effect of glyceryl trinitrate patches vs no patch in both ischemic and hemorrhagic stroke. Glyceryl trinitrate, a nitric oxide donor, has vasodilatory effects that lowers blood pressure and is neuroprotective. Those randomized to the patch had on average a 7.0/3.5 mmHg lower blood pressure compared to those without after the first dose, but the effect disappeared and did not differ statistically after day 3. However, the patch group did not have significantly improved functional outcome.[9]
5. **CATIS Trial:** The China Antihypertensive Trial in Acute Ischemic Stroke (CATIS) studied the effect of reducing blood pressure by 10% to 25% within 24 hours and below 140/90 mmHg in 7 days, in ischemic stroke patients with presenting systolic blood pressures of between 140 mmHg and 220 mmHg. The study was conducted in 4071 patients across 26 hospitals in China. Rather than focus on a specific antihypertensive, the CATIS trial focused on management strategy. A variety of antihypertensives were used, including ACE inhibitors, calcium channel blockers, and diuretics according to a predetermined treatment algorithm. The study found no difference in death or major disability at discharge or at 14 days between the group that received antihypertensives

and the group that did not. Further subgroup analysis showed that even after stratifying for stroke severity at baseline, no differences between intervention and control groups were found.[49,50]

Several weaknesses present in these randomized trials limit their ability to quantify the effect of blood pressure interventions on outcome. First, most trials did not specify a lower bound on blood pressure in the conventional treatment group, so the difference in target systolic blood pressures is not as large as 180 mmHg vs 140 mmHg would seem. For example, in the INTERACT2 trial, the conventional treatment group averaged around 150 mmHg systolic blood pressure at 6 hours when the aggressive treatment group reached the target 140 mmHg systolic.[47] Although the authors did not provide a standard deviation for blood pressures for the two groups, it may be reasonable to assume that targeting a blood pressure less than 180 mmHg would result in higher variance for the group than a target of 140 mmHg, since a lower bound was not provided for the aggressive treatment group. The CATIS trial used a different management goal for the intervention group by specifying a 10%-25% decrease in systolic blood pressure. This difference is not as high as it seems, however, as the group that received antihypertensives saw systolic blood pressure fall from an average of 166.7 mmHg to 144.7 mmHg (-12.7%), while the group that did not still saw a fall from an average of 165.6 mmHg to 152.9 mmHg (-7.2%).[49] These small differences between groups in the aggregate may make it harder to detect the effects of blood pressure on stroke outcomes. Because each patient has different baseline blood pressure, different stroke severities, different risks for hemorrhage, and different perfusion needs, it is reasonable to claim that each patient will have different optimal blood pressures. Thus, the effects of blood pressure on outcomes must be studied on an individualized and multivariate basis. We propose a more comprehensive metric of an individual's blood pressure response to better study individual factors responsible for outcomes.

In the previously cited observational studies, the assessment of blood pressure involved either a maximum cutoff,[37] an average over the first 24 hours,[35] a percentage change from one time point on day 1 to one time point on day 2,[36] or a percentage change from maximum to minimum over the first 24 hours.[38] Either systolic or mean arterial pressure was used. In the randomized trials, measurements of blood pressure included similar metrics such as average

blood pressure or amount/percentage change in the first 24 hours. Table 1 lists previous studies and the blood pressure metrics used in their evaluations. By reducing the blood pressure measurement down to one metric, we lose information about baseline, natural time course behavior, variability, and responsiveness to medication. Losing this information is less important in the randomized controlled trials, where the focus is on the intervention instead of the effects of blood pressure. In observational studies, however, having a more complete picture of the blood pressure time course is valuable. We believe a more comprehensive metric of blood pressure time course and response to medication is needed, and these metrics can be obtained using a dynamic linear model. In addition, dynamic linear models naturally account for measurement variability, multiple measurement methods, and missing values. We will discuss these issues in more detail in the following section.

Table 1: BP metrics evaluated in previous studies. We list the study, study type, BP metric evaluated, and results of the study.

Study	**Study Type**	**BP Metric Evaluated**	**Results**
Kaste et al. (1994)[39]	Randomized controlled trial	Mean SBP and DBP at admission, day 1, and day 7	No difference between treatment and control groups
Chamorro et al. (1998)[36]	Observational	Percentage drop in MAP from baseline to day 1, 2, and 7	20%-30% drop in MAP on day 2 increased odds of full recovery compared to those with no BP drop

Study	Study Type	BP Metric Evaluated	Results
ACCESS Investigators (2003)[41]	Randomized controlled trial	Percentage reduction in 24 hours	Study stopped prematurely, but number of vascular events was fewer in the treatment group vs control
Oliveira-Filho et al. (2003)[38]	Observational	Percentage SBP reduction in first 24 hours	SBP reduction in first 24 hours associated with poor outcome after 3 months
Castillo et al. (2004)[35]	Observational	Average values of SBP and DBP in 24 hour period	Both high and low BP away from 180/100 were associated with worse prognosis
CHHIPS Investigators (2009)[40]	Randomized controlled trial	SBP and DBP at 4 hours, 8 hours, 24 hours, and 2 weeks	No difference in neurologic deterioration in treatment vs placebo groups
Ritter et al. (2009)[37]	Observational	Upper threshold violations in first 24 hours for SBP ($>$ 200 mmHg) and DBP ($>$ 110 mmHg)	SBP $>$ 200 mmHg associated with poor outcome in single factor analysis but not multiple regression analysis
COSSACS Investigators (2010)[43]	Randomized controlled trial	SBP and DBP at baseline and 2 weeks	No difference in adverse events or mortality between group that continued pre-existing antihypertensives vs groupp that stopped

Study	Study Type	BP Metric Evaluated	Results
SCAST Investigators (2011)[42]	Randomized controlled trial	Mean SBP and DBP on days 1-7	No difference in 6-month composite endpoint between groups, but borderline significant worse mRS in treatment group
INTERACT2 Investigators (2013)[47]	Randomized controlled trial	SBP every 15 minutes until 1 hour, every 6 hours until 1 day, and twice a day until 7 days	Possible slight benefit in aggressive treatment group
CATIS Investigators (2014)[49]	Randomized controlled trial	SBP and DBP every 2 hours until 24 hours, every 4 hours until 3 days, and 3 times a day until discharge or death	No difference in death or major disability
ENOS Trial Investigators (2015)[9]	Randomized controlled trial	SBP and DBP at days 0-7	No difference in functional outcome between treatment and control
ATACH II Investigators (2016)[48]	Randomized controlled trial	Minimum SBP and DBP every hour for 24 hours	No difference in rates of death or disability, trial ended prematurely due to futility

Study	Study Type	BP Metric Evaluated	Results
ENCHANTED Investigators (2016)[45]	Randomized controlled trial	SBP and DBP at baseline, 24 hours, 72 hours, and 7 days	Blood pressure results pending in 2018

A common theme in these studies is the lack of individualized characterization. In interventional studies, blood pressure was typically titrated to a certain level below a threshold, the precise number likely depending not on individual characteristics of the patient, but on physician preference. In observational studies, effects of antihypertensives were similarly not accounted for on an individual basis. This is problematic, because it is known that response to antihypertensives vary on an individual basis, and typical patient characteristics are not predictive of these variations in response.[51] To capture these individualized effects, then, we need a more sophisticated model of blood pressure in the ICU.

2.1.3 Blood Pressure Modeling

Blood pressure modeling has been studied mostly in the context of high resolution, beat-to-beat tracking.[24,52] Modeling at the millisecond to second time scale is important in the ICU for the purposes of triggering blood pressure alarms and reducing the amount of false alarms. The Kalman filter (e.g. dynamic linear models) has been used for modeling blood pressure in this context.[24] DLMs have been used to predict readmission probability based on numerical and text ICU data.[53] Some DLM systems, such as the Kalman filter-based modeling system used in Li et al.,[24] explicitly model heart rate as well as blood pressure. In their case, this was aimed at improving both heart rate and blood pressure estimation from arterial lines on a beat-by-beat basis using quality indices instead of using one measure to improve the other. For our problem, we would like to leverage heart rate to improve estimates of blood pressure. Heart rate is importantly correlated with blood pressure, both in terms of normal physiological response (e.g. sympathetic activation increasing heart rate and blood pressure

during agitation in the ICU[54]) and in dysfunctional states such as shock (where increases in heart rate are in response to decreasing blood pressure[55] and is often measured together as the Shock Index[56]).

In work that is perhaps most similar to ours, Lehman et al. modeled ICU blood pressure and heart rate data using the MIMIC database using switching vector autoregressive models (SVARs).[57–59] SVARs are a form of autoregressive model very similar to DLMs, with the removal of a hidden layer and the inclusion of a set of dynamic modes through which a time series can switch during its course - in effect, SVARs model time series as a mixture of DLMs without the hidden layer. Lehman et al. have applied this model to MIMIC ICU data to predict mortality and analyze blood pressure variability. The advantage to the SVAR-based method used by Lehman et al. is its ability to learn dynamic modes applicable across a whole population of patients and weight them separately for each patient. There are several disadvantages to SVAR for our purposes, however. While the dynamic modes learned by SVAR are useful for classification, they lack human interpretability. In addition, the inclusion of multiple dynamic modes with multiple transition matrices, in addition to the weights for each dynamic mode at each time point for each patient, add up to much more parameters to estimate compared to a single DLM. Lehman et al. accomplish this by using expectation maximization on high density data from MIMIC, using the MIMIC waveform database that contains beat to beat blood pressure information with low amounts of missing data. For the purposes of identifying blood pressure trends in stroke management in the ICU, we are interested in a longer term, lower resolution measure of blood pressure, with a time scale on the order of minutes, hours, and even days. Studies at these time scales typically use more rudimentary models of blood pressure, such as mean, least squares slope, and standard deviation of the mean arterial pressure.[60] To our knowledge, we are the first to design a human-interpretable dynamic linear model for blood pressure in the ICU management of stroke.

2.2 HOMEOSTASIS PERTURBATION MODEL

A disease state perturbation model for stroke must capture two aspects of spontaneous blood pressure behavior over time: One, the fact that blood pressure is initially perturbed (often elevated), and two, the blood pressure eventually settles at a new homeostasis level. While Wallace et al. fitted a linear trend to spontaneous blood pressure after stroke,[31], it was done to a variable time scale. The blood pressure trends in Britton et al. (who looked at a wider variety of patients compared to the males admitted to the VA in Wallace et al.) suggest an exponential decay model might be more accurate.[30] We thus chose to use an exponential decay model to fit the blood pressure perturbation post stroke. The exponential decay model is parameterized as follows:

$$P_{spon}(t) = B_{max} * (1 - r_B)^t \tag{2.1}$$

Where B_{max} is the initial perturbation (whether positive or negative), r_B is the rate at which the perturbation decays to zero, and $P_{spon}(t)$ is the perturbation effect on blood pressure at time t.

To model drug effects from typical data available in the chart, it is necessary to model both pharmacokinetics and pharmacodynamics. Pharmacokinetics is necessary to obtain an estimate of the active concentration of a drug at a given time point. Pharmacodynamics is necessary to relate the active concentration to the measured effect of the drug. For the purposes of the AIM-BP model, we chose to implement a simplified version of pharmacokinetics and focus primarily on pharmacodynamics. Our reasons for doing this are multi-fold: One, calculating accurate pharmacokinetics is difficult with data available from the chart, especially given that, in our experience, charting times for dosages are not necessarily accurate to the minute. Two, in a research setting, pharmacokinetics can theoretically be circumvented by directly measuring drug plasma concentrations of repeatedly. Three, even if we standardize pharmacokinetics across all patients, differences in pharmacokinetics will still show up as biases in pharmacodynamics estimates. Although we may not get as accurate estimates for specific pharmacodynamics parameters, for clinical purposes a gestalt measure of "drug

effectiveness" is often sufficient to inform clinical decision making. As such, we limit our model of pharmacokinetics to a standardized time-to-peak and half-life model found in FDA databases (https://www.accessdata.fda.gov/scripts/cder/daf/).

To model pharmacodynamics, we turn to the E_{max} model, a standard model of dosage versus response for pharmacodynamics.[61,62] Other models range from a simple linear model of dosage to effect (which the E_{max} model reduces down to when dosages are much lower than the EC_{50}, as described later), to more complex non-monotonically increasing models such as the quadratic or Adair models.[63] In practice, the E_{max} model is particularly popular because it strikes a good compromise between model expressivity and simplicity.

Variations exist on the complexity and number of parameters of the E_{max} model, particularly regarding the inclusion of a base effect E_0 and a slope factor N,[61] which we chose to omit for simplicity. The E_{max} model we chose to use is thus parameterized as follows:

$$P_{med}(t) = \frac{E_{max} * c(t)}{EC_{50} + c(t)} \tag{2.2}$$

Where E_{max} is the maximum effect of a drug, EC_{50} is the concentration of a drug necessary to achieve an effect of $0.5 * E_{max}$, $c(t)$ is the drug concentration at time t, and $P_{med}(t)$ is the effect of a drug at time t.

2.3 DYNAMIC LINEAR MODELS

A dynamic linear model (DLM), which is a linear and Gaussian form of a state space model (SSM), can be described as an autoregressive linear model with a hidden layer. In this section, we will follow notation used in Shumway and Stoffer's "Time Series Analysis and Its Applications" (4th Edition).[64] A first order discrete time DLM is described by the following set of equations:

$$x_t = \Phi x_{t-1} + \Upsilon u_t + w_t \tag{2.3}$$

$$y_t = A x_t + \Gamma u_t + v_t \tag{2.4}$$

$$x_0 \sim N(\mu_0, \Sigma_0)$$

$$w \sim N(0, Q)$$

$$v \sim N(0, R)$$

In this model, x_t is the hidden state at time point t, u_t is an input vector at time point t, and y_t is the observed state at time point t. Φ is a matrix that describes the transition of the hidden state from one time point to the next, which we will refer to as the transition matrix. Υ and Γ describe how the hidden state and the observed state responds to input, respectively, and we will refer to these as the responses matrices. A is a matrix that captures how an observed state is emitted from the hidden state, and we will refer to this as the emission matrix. w_t and v_t are white Gaussian noise terms with covariances Q and R, respectively. These model the inherent noise in the system as well as measurement noise.

For our AIM-BP model, we will use a simplified version of the DLM with only one response matrix Υ:

$$x_t = \Phi x_{t-1} + \Upsilon u_t + w_t \tag{2.5}$$

$$y_t = A x_t + v_t \tag{2.6}$$

$$x_0 \sim N(\mu_0, \Sigma_0)$$

$$w \sim N(0, Q)$$

$$v \sim N(0, R)$$

The set of parameters of this model $\{\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R\}$ will be referred to as the model Θ. Figure 4 shows a graphical representation of a DLM.

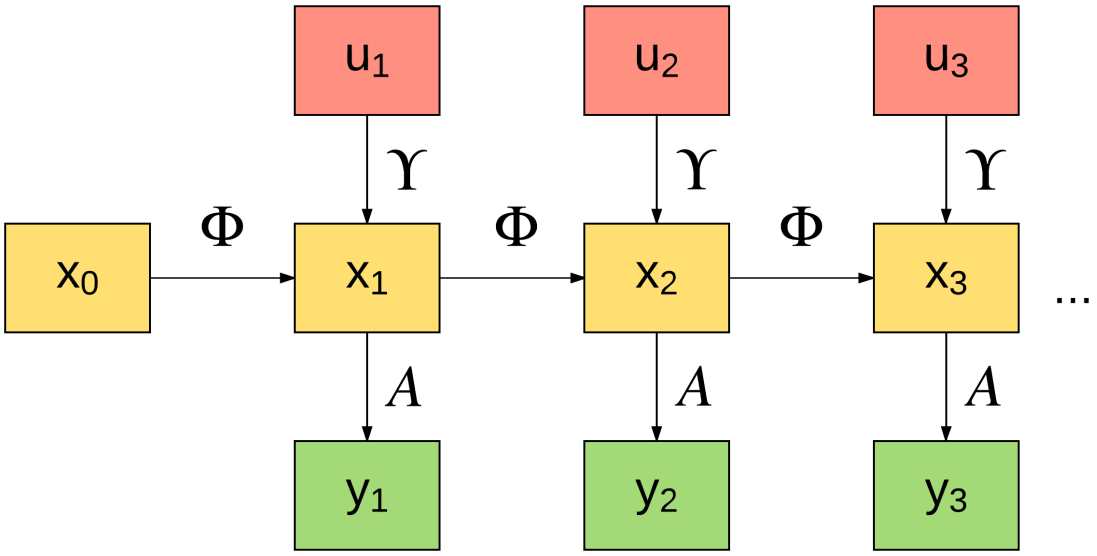

Figure 4: Graphical representation of a dynamic linear model. Φ governs the transition between hidden states x and A governs the emission of observations y from hidden states x.

2.3.1 Kalman Filter

The problem of estimating a hidden state x_t given observed states $\{y_1, y_2, ...y_t\}$ and a known model Θ is known as filtering. Using notation found in Shumway and Stoffer:[64]

Let us define conditional expectations and variances x_t^s and P_{t_1,t_2}^s as follows:

$$x_t^s = E[x_t|y_1, ...y_s] \tag{2.7}$$

$$P_{t_1,t_2}^s = E[(x_{t_1} - x_{t_1}^s)(x_{t_2} - x_{t_2}^s)] \tag{2.8}$$

When $t_1 = t_2$, we'll write P_{t_1,t_2}^s as P_t^s for convenience.

The Kalman filter is then defined for the state space model defined in (2.5) and (2.6) as follows:

$$x_0^0 = \mu_0$$

$$P_0^0 = \Sigma_0$$

For $t = 1, ...n$:

$$x_t^{t-1} = \Phi x_{t-1}^{t-1} + \Upsilon u_t \tag{2.9}$$

$$P_t^{t-1} = \Phi P_{t-1}^{t-1} \Phi' + Q \tag{2.10}$$

$$x_t^t = x_t^{t-1} + K_t \epsilon_t \tag{2.11}$$

$$P_t^t = (I - K_t A) P_t^{t-1} \tag{2.12}$$

where

$$K_t = P_t^{t-1} A' (A P_t^{t-1} A' + R)^{-1} \tag{2.13}$$

$$\epsilon_t = y_t - E[y_t | y_1, ...y_{t-1}] = y_t - A x_t^{t-1} \tag{2.14}$$

Here, K_t is referred to as the Kalman gain and ϵ_t is the innovation. Intuitively, the Kalman gain governs how much the predicted x_t changes from its expected value conditioned on previous observations versus the innovation. As the variance $P_t^{t-1} \to 0$, the Kalman gain $K \to 0$ - in other words, the more sure we are about the current state given the previous state, the more we will rely on our prediction of the previous state instead of the measured value. On the other hand, as the measurement noise covariance $R \to 0$, $K \to A^{-1}$, which when multiplied by the innovation ϵ_t results in cancelling out the x_t^{t-1} term. Thus, when the measurement noise covariance goes to 0, our prediction for the next hidden state depends solely on the measured value. Finally, in the corner case when $R = 0$ and the variance $P_t^{t-1} \to 0$, the Kalman gain $K \to 0$.

2.3.2 Kalman Smoother

The related problem of estimating a hidden state x_t given all observed states $\{y_1, y_2, ...y_n\}$ and a known model Θ is known as smoothing. Using notation from the previous section, smoothing involves finding x_t^n. This can be done as follows:

With x_n^n and P_n^n obtained via Kalman filter:

$$x_{t-1}^n = x_{t-1}^{t-1} + J_{t-1}(x_t^n - x_t^{t-1}) \tag{2.15}$$

$$P_{t-1}^n = P_{t-1}^{t-1} + J_{t-1}(P_t^n - P_t^{t-1})J_{t-1}' \tag{2.16}$$

where

$$J_{t-1} = P_{t-1}^{t-1}\Phi'(P_t^{t-1})^{-1} \tag{2.17}$$

2.3.3 Forward Filtering Backward Sampling

When performing Gibbs sampling, we need a sampling version of the Kalman smoother to obtain samples for $P(x|y,\Theta)$. This algorithm is known as the forward filtering backward sampling algorithm.[65] The algorithm works as follows:

Sample x_n from $N(x_n^n, P_n^n)$

$$m_t = x_t^t + J_t(x_{t+1} - x_{t+1}^t) \tag{2.18}$$

$$V_t = P_t^t - J_t P_{t+1}^t J_t' \tag{2.19}$$

Sample x_t from $N(m_t, V_t)$ for $t = \{n-1, ...0\}$

Where x_n^n and P_n^n are obtained from the Kalman filter and J_t is defined as in the Kalman smoother.

2.3.4 Handling Missing Values

In the Kalman filter, we can handle partially missing values in the observed data vector y by zeroing out the missing components in y, A, and R. Let $y^{(1)}$ be the observed components of y and $y^{(2)}$ be the missing components such that $y = (y^{(1)}; y^{(2)})$.

We can rewrite the emission equation 2.6 and the covariance of v as the following:

$$\begin{pmatrix} y_t^{(1)} \\ y_t^{(2)} \end{pmatrix} = \begin{bmatrix} A^{(1)} \\ A^{(2)} \end{bmatrix} x_t + \begin{pmatrix} v_t^{(1)} \\ v_t^{(2)} \end{pmatrix} \tag{2.20}$$

$$cov \begin{pmatrix} v_t^{(1)} \\ v_t^{(2)} \end{pmatrix} = \begin{bmatrix} R_{11} & R_{12} \\ R_{21} & R_{22} \end{bmatrix} \tag{2.21}$$

Then, missing values in y can be handled by setting y, A, and R as follows:

$$y_t = \begin{pmatrix} y_t^{(1)} \\ 0 \end{pmatrix}, A = \begin{bmatrix} A^{(1)} \\ 0 \end{bmatrix}, R = \begin{bmatrix} R_{11} & 0 \\ 0 & I_{22} \end{bmatrix}$$

Once filtered values have been obtained in this fashion, smoother values can be obtained using the filtered values as normal. This is discussed in more detail in Shumway and Stoffer Section 6.4.[64]

2.3.5 Estimation

Estimation of DLM parameters can be done using maximum likelihood, expectation maximization, and Markov Chain Monte Carlo (MCMC) methods. MCMC methods have gained popularity due to their flexibility in modeling non-Gaussian probabilities and non-linear extensions to DLMs.[66–68] MCMC methods provide two major advantages for us: One, they provide a natural framework for expressing uncertainty about learned parameters; Two, they allow us to experiment with non-Gaussian priors on DLM learning.

2.3.6 Markov Chain Monte Carlo

Markov Chain Monte Carlo (MCMC) methods are helpful in the setting of Bayesian inference when the posterior distributions are too difficult to directly calculate. Often, a closed form solution of the posterior is impossible to derive, or the normalizing factor may be computationally too difficult to calculate. MCMC methods work around this through simulating samples from our desired distribution via a Markov chain. In order for MCMC to converge, this Markov chain must possess three properties: Aperiodic, irreducible, and recurrent.

- **Aperiodic:** A Markov chain has a period k if it can only return to a state S visited at time n at later times $n+k$, $n+2k$, ... If the period $k \leq 1$, the Markov chain is aperiodic.
- **Irreducible:** A Markov chain is irreducible if it can go from any state to any other state.
- **Recurrent:** A Markov chain is recurrent if the probability of never coming back to any given state is zero. That is, we will eventually come back to a given state, even if the expected time that we do so is infinite.

Our goal for DLM parameter estimation is to draw samples from the joint distribution $P(\Theta, x|y)$. We can do so using a Gibbs sampler by iteratively sampling from the conditionals $P(x|\Theta, y)$ and $P(\Theta|x, y)$. The probability $P(x|\Theta, y)$ can be sampled using the Forward Filtering Backward Sampling (FFBS) algorithm. The probability $P(\Theta|x, y)$ will need to be sampled on a parameter by parameter basis. We will work through each parameter in Θ in Appendix A. We will provide a brief overview of MCMC techniques here. A good overview of basic MCMC methods can be found in Lynch, Scott M. (2007): *Introduction to Applied Bayesian Statistics and Estimation for Social Scientists*[69] and tutorials by Ilker Yildirim (http://www.mit.edu/~ilkery/).

2.3.6.1 Gibbs Sampling A Gibbs sampler draws samples from the full posterior distribution $P(A, B, C...)$ by drawing from a series of conditionals $P(A|B, C, ...)$, $P(B|A, C, ...)$, $P(C|A, B, ...)$. The random variables $\{A, B, C, ...\}$ can either be sampled one random variable at a time or in blocks, as long the corresponding conditional distribution is able to be sampled from directly. In Bayesian inference, the conditional is usually rewritten as $P(A|B, C, ...) \propto P(B, C, ...|A)P(A)$ as the full posterior conditional $P(B, C, ...|A)$ has an easy sampling distribution. In this case, choosing a conjugate prior for $P(A)$ gives us a sampling distribution for $P(A|B, C, ...)$.

2.3.6.2 Metropolis Hastings When we do not have a directly sampleable distribution for $P(A|B, C, ...)$ (e.g. if we do not want to use a conjugate prior in Bayesian inference), we cannot perform standard Gibbs sampling. In this situation, we can use the Metropolis Hastings algorithm. In Metropolis Hastings, instead of sampling $P(A|B, C, ...)$ directly,

we instead first sample from a proposal distribution $P(A|A^{prev})$. We then calculate an acceptance probability $P_{accept}(A|A^{prev}) = min(1, \frac{P(A^{prev}|A)P(A,B,C,...)}{P(A|A^{prev})P(A^{prev},B,C,...)})$. We then accept the new value A with probability P_{accept}, otherwise we keep the old value A^{prev}. It is possible to do Metropolis within Gibbs, where we perform a Metropolis step for a conditional $P(A|B,C,...)$ that we cannot sample from directly. In this case, the acceptance probability $P_{accept}(A|A^{prev}) = min(1, \frac{P(A^{prev}|A)P(A|B,C,...)}{P(A|A^{prev})P(A^{prev}|B,C,...)})$.

The proposal distribution $P(A|A^{prev})$ can be symmetric (i.e. $P(A|A^{prev}) = P(A^{prev}|A)$ and the first terms of the acceptance probability cancels out) or asymmetric. The normal distribution is an example of a symmetric proposal distribution, which is commonly used. In situations in which you might want to skew values in a certain direction, asymmetric proposal distributions can be used. The parameters of the proposal distribution (e.g. Σ^2 for a normal proposal distribution of $N(A^{prev}, \Sigma^2)$) are considered tuning parameters for the Metropolis Hastings. In the case of a normal proposal distribution, too small of a tuning parameter Σ^2 results in a high acceptance rate but slow movement exploring the full joint probability area of A, whereas too large of a tuning parameter results in a low acceptance rate and large, rare jumps, which also limits the exploration of A. It has been shown that when both the target distribution and the proposal distribution are normal, an optimal acceptance rate approaches 0.234 as the number of dimensions increases.[70] In practice, tuning this tuning parameter often takes trial and error.

2.3.6.3 Adaptive Rejection Metropolis Sampling

Adaptive rejection sampling gets around the issue of tuning parameters. Gilks et al. adapted rejection sampling to both Gibbs[1] and Metropolis within Gibbs[2] algorithms. Adaptive rejection sampling works by constructing an envelope function and an optional squeezing function of the log of a log-concave target function and as samples are accepted and rejected, updating the envelope function and squeezing function to reduce the number of samples that are rejected.[71] Adaptive rejection Metropolis sampling (ARMS) extends this to non-log-concave target functions by including a Metropolis step and abandoning the squeezing function. Further enhancements of ARMS includes Independent Doubly Adaptive Rejection Metropolis Sampling (IA2RMS), which is implemented in Matlab code.[3] ARMS is a generic sampler for non-standard distributions

and is used in generic graphical model MCMC samplers like BUGS.[22,72]

2.3.6.4 Diagnosing Convergence Depending on the initial conditions, an MCMC sampler will likely take some time before samples are drawn from the true posterior distribution. During this period of time, drawn samples are not representative of the posterior, and thus must be discarded. This process is called "burn-in." After a chain has reached the true posterior, sampled values may still be highly correlated with the previous and next sampled value. To mitigate this problem, we could simply record every n^{th} sampled value and skip the others. We could keep the skipped values as well, but that would require much more storage space for the sampled chain. Diagnosing the point at which an MCMC sampling chain has converged to the true posterior can be done by analyzing the auto-correlation function (ACF) of the sampled chain and the chain trace plot of the sampled values themselves (for examples of what these plots look like, see the sensitivity analysis results). When a chain has reached the true posterior, the trace plot should look noisy around a constant value. In certain pathological posterior distributions (e.g. a bimodal distribution with very low probabilities between the two modes), a chain trace may look constant for a while around one mode until by chance it skips to the other mode. On the other hand, the ACF tells us about how correlated consecutive sampled values are to each other. We skip enough values when the ACF goes to zero quickly. Because there are no hard and fast rules for analyzing these plots, diagnosing convergence of MCMC algorithms is an art form that combines knowledge of the target distribution and algorithms involved.

3.0 PRELIMINARY ANALYSIS OF ICU DATA AND MODEL DESIGN

We performed a preliminary analysis of real world ICU data from two sources in order to inform our AIM-BP model design. In particular, we were interested in patterns of blood pressure trends, estimates of blood pressure variance and measurement noise, and an analysis of missing values. The two datasets used were an internal dataset from UPMC and MIMIC III, a publicly available ICU database that has been used for retrospective data analytics projects such as mortality prediction and early detection of shock.[73–75] We primarily focused on the UPMC dataset, and used MIMIC III to supplement where the UPMC dataset was lacking. Because the existing UPMC dataset was focused on primary intracerebral hemorrhage, we chose to limit the patients extracted from MIMIC III to that diagnosis as well.

3.1 UPMC DATASET PARSING

We analyzed charted vital signs from a dataset of 497 primary intracerebral hemorrhage (ICH) patients previously extracted from the electronic medical record at UPMC Presbyterian (IRB PRO16110384). Cuff and arterial blood pressures, medication administrations, and demographics information were processed and loaded into Matlab.

3.2 MIMIC DATASET PARSING

Clinical ICU data was extracted from the publicly available MIMIC III database. MIMIC III was loaded into a PostgreSQL database, and included tables for ICD diagnoses, prescriptions,

procedures, inputs and outputs, chart events, labs, and notes. To identify patients in the ICU with primary intracerebral heomorrhage, we queried the set of patients with ICD9 codes starting with 431. Using the ICD9 query, we found 1367 ICU admissions with the diagnosis of primary intracerebral hemorrhage.

Blood pressure and heart rate data was extracted from the chart events table. First, MIMIC III item IDs for blood pressure and heart rate was searched for from the items table, using keywords "blood pressure," "BP," "heart rate," and "HR." Items corresponding to these item IDs were then extracted from the chart events table for the 1367 admissions matching our criteria. These measurements include both systolic, diastolic, and mean arterial pressures using manual blood pressure measurements, cuffs, and arterial blood pressures.

3.3 DISCRETIZATION

Although vital signs were charted at non-uniformly sampled time points in both datasets, AIM-BP requires data in a uniformly sampled discrete time series. We chose a sampling rate of every 15 minutes, based on the sampling rates observed from the both datasets. Data from both datasets were resampled in 15 minute periods starting on the quarter hour. If a single value was recorded for a given variable in that 15 minute time period, that value was recorded in the resampled data. If multiple values were recorded in the period, the mean was recorded for the resampled data. If no values were recorded in the period, a missing value was recorded in the resampled data.

A heat map for systolic blood pressure over time for all UPMC patients can be seen in Figure 5. The same heat map for MIMIC patients can be seen in Figure 6. On average, systolic blood pressure decreases over the first few hours, as expected in the acute post stroke setting either spontaneously or as a result of medical intervention.

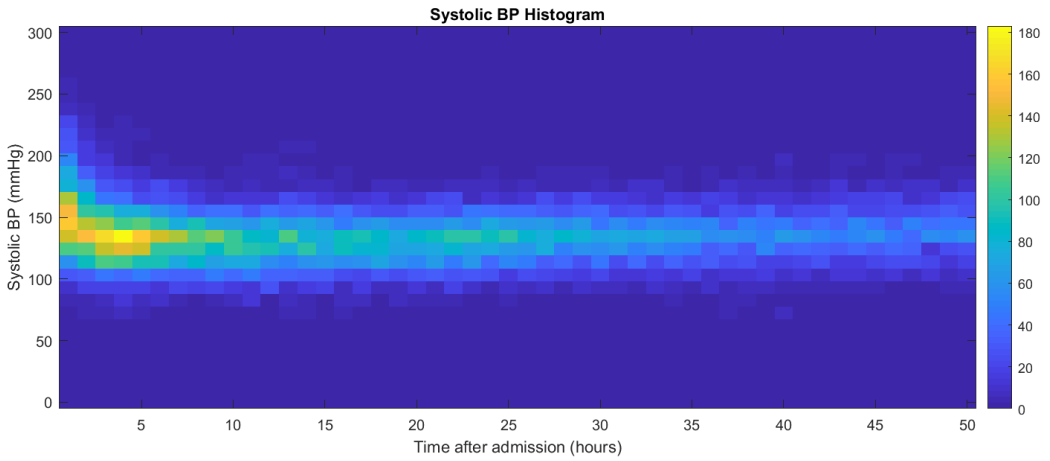

Figure 5: Heat map over time of SBP for all UPMC patients. Discretized measured SBP values were aligned at the first measured time point and a histogram was calculated at each hour and expressed as a heat map. On average, the systolic blood pressure decreases over the first 5 hours.

3.4 CLUSTER ANALYSIS OF UPMC DATA

The histograms of blood pressure time courses in the previous section may be characteristic of most patients, or there might be specific subgroups of patients who exhibit different behavior that is masked by taking a histogram over the entire patient population of the dataset. If the histogram is representative of most patients, then a single model can be used to characterize blood pressure behavior in stroke. In order to determine if this is the case, we sought to identify the major types of blood pressure trends via clustering.

Because time series data in raw form typically have different number of time points in each time series, they do not fit easily into a standard, low dimensional space in which clustering can be done. Thus, time series data generally need to be transformed into a standard low dimensional space or a similarity measure (also known as a kernel) between two time series needs to be defined before clustering can be done. Transforming a time series into a low dimensional representation translates to building a model of the time series, such

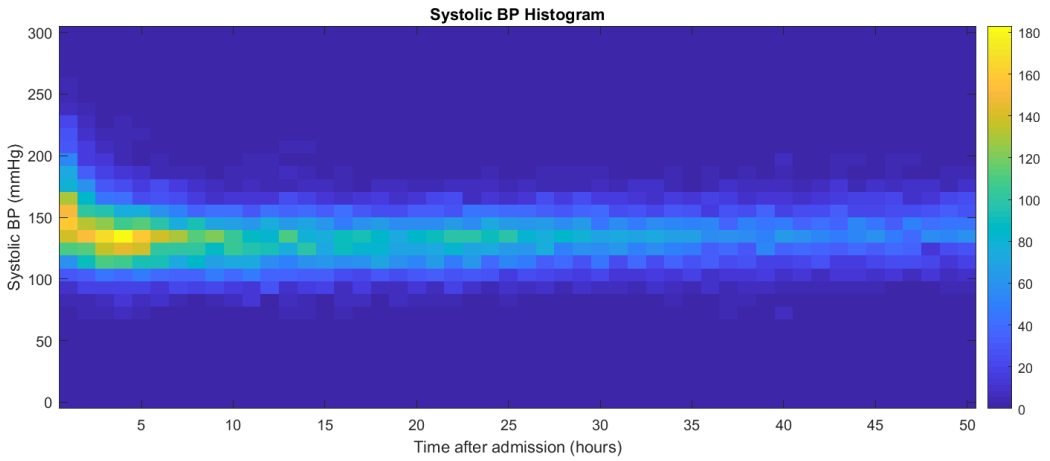

Figure 6: Heat map over time of SBP for all MIMIC patients. Discretized measured SBP values were aligned at the first measured time point and a histogram was calculated at each hour and expressed as a heat map. On average, the systolic blood pressure decreases over the first 3 hours.

as an ARMA model.[64] Simple similarity measures for time series include Euclidean distance and Dynamic Time Warping (DTW).[76] Although DTW is a popular similarity measure that has been used in the context of both blood pressure signals and stroke,[77,78] it is most useful for aligning time shifted time series.[76]

We thus chose to perform cluster analysis in two ways: one, by transforming the time series data into a low dimensional representation via modeling, and two, by clustering using a similarity measure. We chose to focus on systolic blood pressure only for its clinical importance and for simplicity of visualization.

For the first method, we chose to model the systolic blood pressure using the spontaneous perturbation model in 2.1 (with parameters r_B and B_{max}) with the addition of an additional parameter representing homeostasis target blood pressure. This model is sufficient for a rough estimate of initial perturbation and the eventual homeostasis pressure. We chose not to fit a more general ARMA model, as blood pressure time series data generally do not meet the criteria for weak stationarity (mean not dependent on time, autocovariance only

dependent on time difference) necessary for ARMA models. This model by itself does not account for the effect of medication, which we will analyze in later sections. As such, residual variance after model fit will be higher than what we will see in later sections. In order to control for edge cases, we limited the rate value r_B. On one end of the limit, we limited r_B to correspond to a 99% decrease in perturbation magnitude in 6 hours. On the other end of the limit, we limited r_B to correspond to a 99% decrease in perturbation magnitude in 4 days. These limits were chosen to avoid pathological model fitting behaviors, where a perturbation model that decreased too fast or too slowly would minimize the error but make little sense clinically. Because of these limits on r_B, the r_B variable itself is not a good variable to cluster on. Finally we only fit the model to patients with greater than 10 measurements to reduce edge cases as a result of overfitting.

A histogram of fit model parameters and variances can be seen in Figure 7. Two clusters (highlighted by red circles) can be seen via visual inspection in the initial perturbation axis. These clusters correspond to a large cluster where the initial perturbation B_{max} is positive and a smaller cluster where the initial perturbation is negative. The distribution of residual variance is heavy tailed, but no separate peaks are visible. If separate peaks were visible in the residual variance histogram, we would have concluded that the spontaneous perturbation model we used was accurate for one group of patients but not for another. However, the lack of separate peaks suggests that the spontaneous perturbation model works about the same for most patients bar a few outliers (hence the heavy tail). Based on the B_{max} distribution, we chose to cluster patients with a positive B_{max} against patients with a negative B_{max}. We plot the mean systolic blood pressure for each cluster in Figure 8. Although the cluster with negative B_{max} start off with a lower systolic BP, the difference disappears by 10 hours in.

From a clinical perspective, patients could be categorized into those whose blood pressure is too high and need to be lowered, and those whose blood pressure is too low and need to be raised. From our clustering results, it would appear that in the population of primary intracerebral hemorrhage patients in the ICU, more patients belong in the former group compared to the latter group. Over time, both groups ideally reach similar target blood pressures, whether spontaneously or through medical intervention. On the surface, this all fits well with what we know about stroke patients in the ICU. However, the presence of

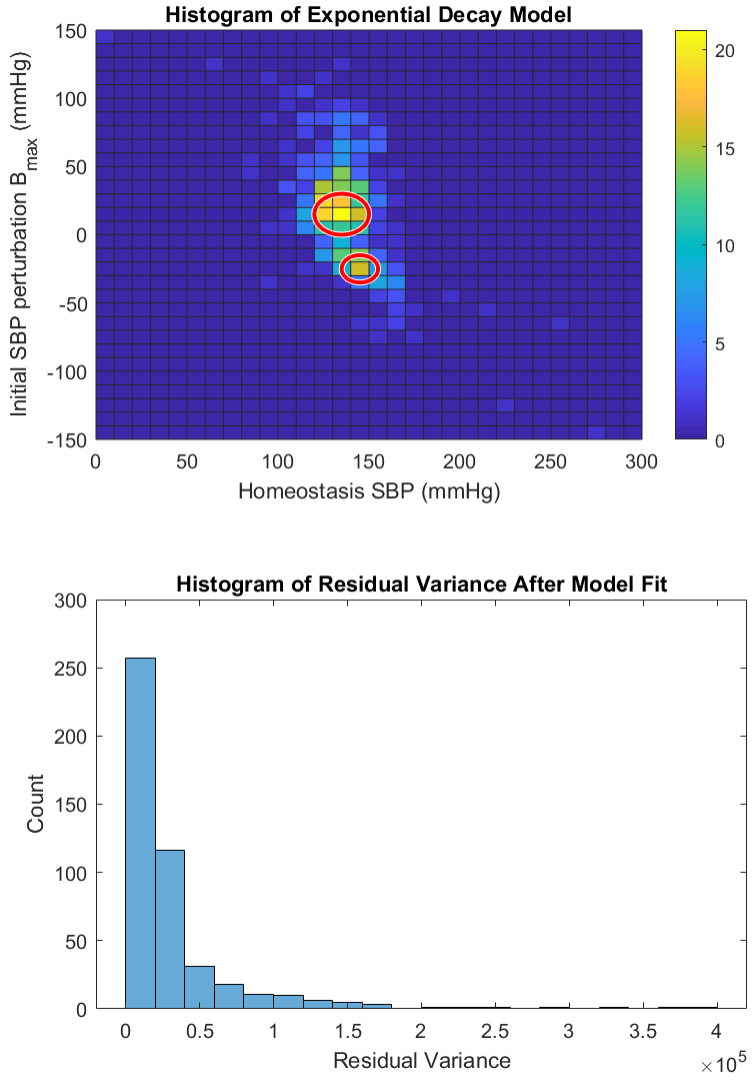

Figure 7: Cluster analysis of UPMC data using spontaneous perturbation model. Two clusters (highlighted by red circles) can be seen visually in the model parameters for UPMC data. Variances exhibit a heavy tailed distribution.

the positive versus negative initial perturbation B_{max} clusters can be also be explained by model overfitting. Figure 9 demonstrates an example where a positive B_{max} is fitted (as demonstrated by the decreasing fit curve) despite the overall trend looking mostly constant.

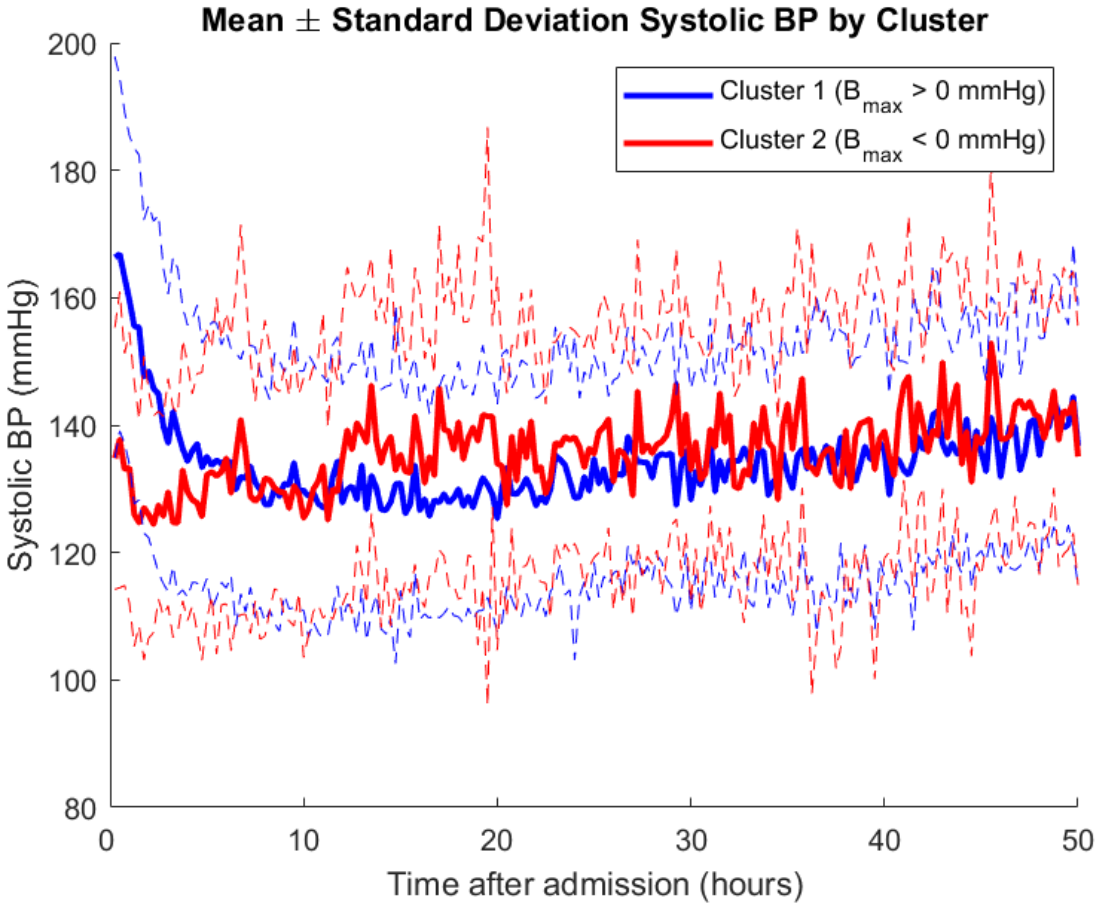

Figure 8: Systolic BP UPMC patients clustered by B_{max}. The cluster of patients with a negative initial perturbation B_{max} have a lower initial blood pressure than the cluster of patients with a positive B_{max}, but the difference in intial blood pressure disappears after less than 10 hours.

The B_{max} parameter tends to be fit away from zero, as it is an additional degree of freedom that can be fit to reduce the residual variance.

We then approached the problem of clustering by using similarity measures in order to avoid the model overfitting issues. We chose to use Euclidean distance instead of dynamic time warping for the similarity measure because our time series are generally aligned at the start of admission. We calculated Euclidean distance between pairs of the first 50 hours of

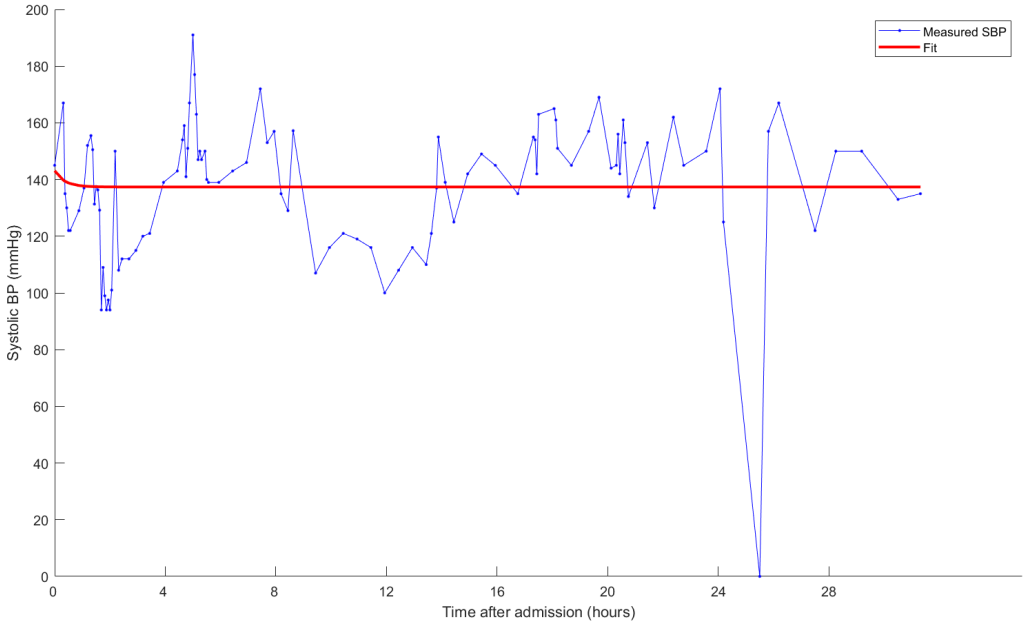

Figure 9: Example fit to systolic BP. In this example systolic BP fit, a positive B_{max} is fit despite the trend looking more or less constant.

systolic blood pressure time series. The Euclidean distance between the pair was corrected for the amount of missing values by dividing by the proportion of time points that had missing values in either time series. Hierarchical agglomerative clustering was performed, and the dendrogram can be seen in Figure 10. No large gaps exist in the distance between clusters, meaning no obvious clustering split exists based on the Euclidean distance.

Based on the clustering results, we concluded that although patients may be able to be clustered based on the direction of the initial perturbation to their blood pressure, this result could possibly be explained by overfitting. Similarity measure hierarchical agglomerative clustering failed to reveal any new insights. Analysis of the residual variance from the spontaneous perturbation model suggests that the model works equally well for the majority of patients, and thus we chose to use it in our development of AIM-BP.

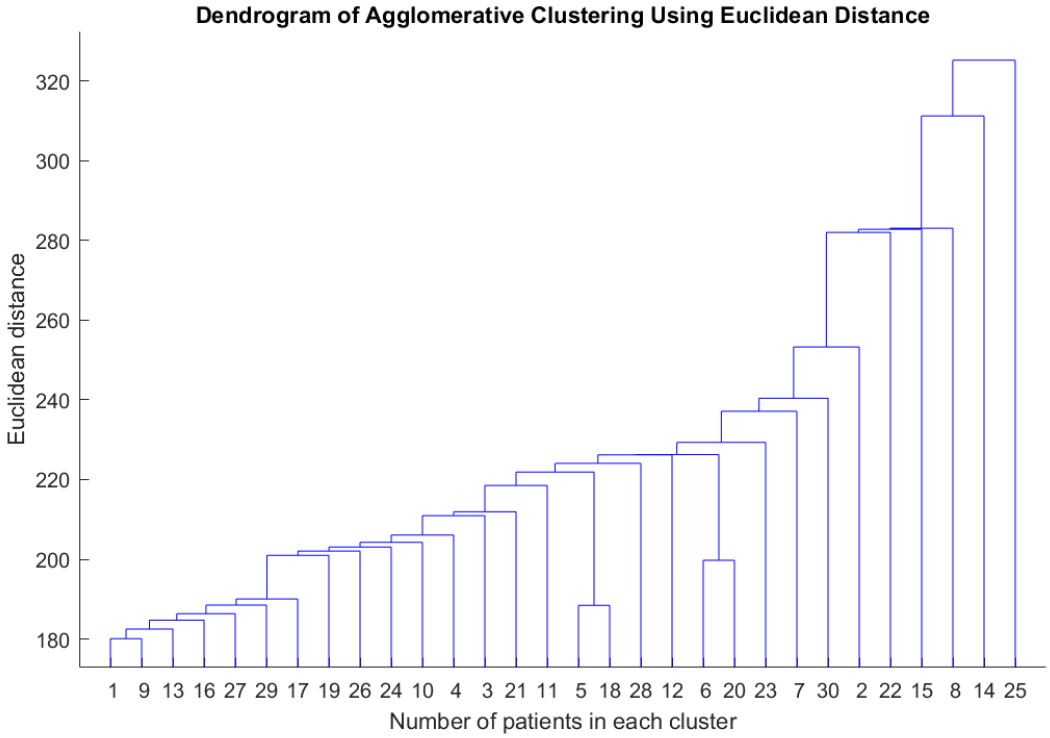

Figure 10: Dendrogram of Euclidean distance between pairs of UPMC patients' systolic BP. No large gaps exist in the distance between clusters.

3.5 VARIANCE ESTIMATION

Inherent blood pressure and heart rate variances Q and measurement noise R (described in detail in the next section) were estimated from the UPMC and MIMIC III datasets. Real world patient blood pressure data have periods of high variability and periods of low variability, as a result of interventions, stress, and a variety of other factors. We chose not to capture this difference in variability in our AIM-BP model for the sake of simplicity, and instead make the assumption that vital sign variability will remain constant over the duration of our simulation. In order to inform an estimate for variances Q and R for simulations, we chose to take the median of variances calculated over consecutive 6 hour periods for each patient so as to capture short term variability. Because the UPMC dataset did not include

heart rate, we estimated heart rate variance with only the MIMIC dataset.

We divide the vital sign time series of each patient in both the UPMC and MIMIC datasets into periods of 6 hours. A linear trend was fitted using least squares to each 6-hour segment and the variance was calculated from the residual. The median variances for systolic and diastolic blood pressure as measured by blood pressure cuff and arterial lines and heart rate can be found in Table 2. The UPMC dataset had a higher median variance compared to the MIMIC dataset. Surprisingly, the median arterial line blood pressure variance was the same or higher than the median non-invasive variance, contradicting conventional wisdom that non-invasive blood pressure measurements are more noisy than arterial line measurements of blood pressure. A possible explanation for this is that arterial lines are placed when a patient requires continuous blood pressure monitoring or arterial blood gas analysis, which is more likely when the patient's blood pressure is highly variable to begin with.

Based on these findings, we split the difference between the UPMC and MIMIC datasets and introduced inherent blood pressure variance Q in our simulation of 81 mmHg^2 for systolic and 49 mmHg^2 for diastolic BP and included measurement noise variance R of 25 mmHg^2 for SBP and 16 mmHg^2 for DBP to simulate data that is at least as noisy as real data.

Table 2: Median variances of vital signs in UPMC and MIMIC datasets.

Variable	UPMC Median Variance	MIMIC Median Variance
Systolic BP (Non-Invasive)	113.87	67.87
Diastolic BP (Non-Invasive)	65.29	38.10
Systolic BP (Arterial)	113.96	89.98
Diastolic BP (Arterial)	22.81	21.98
Heart Rate	N/A	15.63

Median variances were calculated over all ischemic stroke patients in 6 hour segments.

3.6 MODEL DESIGN

We first describe a basic model for the combined perturbation of blood pressure homeostasis as a result of stroke and the medical intervention for stroke. We then describe the dynamic linear model used for AIM-BP.

3.6.1 Homeostasis Perturbation Model

The blood pressure homeostasis perturbation model can be described as follows:

$$P_{spon}(t) = B_{max} * (1 - r_B)^t \tag{3.1}$$

$$P_{med}(t) = \frac{E_{max} * c(t)}{EC_{50} + c(t)} \tag{3.2}$$

$$SBP(t) = P_{spon}(t) + \sum_{meds} P_{med}(t) + SBP_H + \epsilon_S \tag{3.3}$$

$$DBP(t) = ratio_{SD} * (P_{spon}(t) + \sum_{meds} P_{med}(t)) + DBP_H + \epsilon_D \tag{3.4}$$

Where B_{max} and r_B parameterize the exponential decay of the spontaneous blood pressure change in acute stroke patients, and E_{max} and EC_{50} for each drug parameterize the E_{max} model for each drug, $c(t)$ is the drug concentration at time t. SBP_H and DBP_H are homeostasis blood pressure levels without the effects of the perturbations. $ratio_{SD}$ describes the proportional effect of the perturbation on diastolic pressure compared to the effect on systolic pressure. ϵ_S and ϵ_D are normal white noise terms.

In our stroke homeostasis perturbation model, we chiefly model two perturbations: The first is the increased blood pressure post stroke, which often decreases spontaneously throughout the ICU stay. We model this blood pressure perturbation as an initial elevation that decays exponentially to zero, such that the blood pressure eventually converges to a stable homeostasis blood pressure SBP_H. This exponential model is reflective of the spontaneous blood pressure trends seen post stroke.[30] The second class of perturbations modeled are the effects of medications used in the management of the patient's cardiovascular state, with a separate $P_{med}(t)$ for each medication used. We chose the E_{max} model of pharmacodynamics to model these medications. Put together, the homeostasis perturbation model combines an exponential decay model of spontaneous blood pressure behavior in acute stroke ($P_{spon}(t)$) in addition to E_{max} models of drug pharmacodynamics (each $P_{med}(t)$). This homeostasis perturbation model is a general model of blood pressure behavior, and can be fit to data without using the AIM-BP framework.

3.6.2 Dynamic Linear Model

We describe the dynamic linear model that underlies the AIM-BP framework. This system will be tested on simulated data to demonstrate the feasibility of the framework. The AIM-BP dynamic linear model is characterized by the following variables:

3.6.2.1 Observed Variable y The variable $y = \{y_1, y_2, ...y_T\}$ describes observed measurements of blood pressure and heart rate. At a given time point t, y_t contains all measurements of systolic blood pressure (SBP), diastolic blood pressure (DBP), and heart rate (HR). For example, we could have blood pressure measurements from an automated pressure

cuff and an arterial line, and a heart rate measurement from an O2 saturation monitor. Our y variable space would then look like this:

$$y = \begin{bmatrix} y^{(1)} \\ y^{(2)} \\ y^{(3)} \\ y^{(4)} \\ y^{(5)} \end{bmatrix} = \begin{bmatrix} \text{Cuff measured SBP} \\ \text{Cuff measured DBP} \\ \text{Arterial line measured SBP} \\ \text{Arterial line measured DBP} \\ \text{O2 saturation monitor measured HR} \end{bmatrix}$$

By capturing multiple measurements of a vital sign at a given time point (multi-modal data), our model can better estimate the true value of these vital signs. For a subset of the simulations in this thesis, however, we will use a simplified observed measurement space as follows:

$$y = \begin{bmatrix} y^{(1)} \\ y^{(2)} \\ y^{(3)} \end{bmatrix} = \begin{bmatrix} \text{Cuff measured SBP} \\ \text{Cuff measured DBP} \\ \text{O2 saturation monitor measured HR} \end{bmatrix}$$

3.6.2.2 Hidden Variable x We design a hidden variable $x = \{x_1, x_2, ...x_T\}$ that describes the true underlying blood pressure and heart rate measurements of a patient. At a given time point t, x_t contains information about a patient's true SBP, DBP, and HR, as well the baseline homeostasis target for SBP, DBP, and HR. The homeostasis target for SBP, DBP, and HR models what the body's natural desired blood pressure is absent outside interventions. Our AIM-BP model will model the tendency of these vital signs to trend towards a homeostatic value. Finally, we include a constant value in x to allow for the expression of non-white noise.

$$x = \begin{bmatrix} x^{(1)} \\ x^{(2)} \\ x^{(3)} \\ x^{(4)} \\ x^{(5)} \\ x^{(6)} \\ x^{(7)} \end{bmatrix} = \begin{bmatrix} \text{SBP} \\ \text{DBP} \\ \text{SBP homeostasis target} \\ \text{DBP homeostasis target} \\ \text{HR} \\ \text{HR homeostasis target} \\ \text{Constant} \end{bmatrix}$$

3.6.2.3 Input Variabe u The input variable $u = \{u_1, u_2, ...u_T\}$ represents all the various perturbations of the blood pressure and heart rate system from the homeostasis perturbation model described previously. For the sake of simplicity in the following simulations, this thesis will model two commonly used drugs at our institution. At a given time point t, u_t contains the post-stroke blood pressure perturbation as well as the effect of each drug on the patient. For example, a simple model with two drugs might have an input u as:

$$u = \begin{bmatrix} u^{(1)} \\ u^{(2)} \\ u^{(3)} \end{bmatrix} = \begin{bmatrix} \text{Stroke blood pressure perturbation}(P_{spon}) \\ \text{IV labetalol effect}(P_{lab}) \\ \text{IV nicardipine effect}(P_{nic}) \end{bmatrix}$$

Translating the stroke perturbation $P_{spon}(t)$ to u_t:

$$u_t^{(1)} = B_{max} * (1 - r_B)^t$$

B_{max} and r_B are the parameters to estimate.

The E_{max} model is translated from $P_{med}(t)$ to u_t as follows:

$$u_t^{(2)} = \frac{E_{max}^{(2)} c(t)}{EC_{50}^{(2)} + c(t)}$$

Where $c(t)$ is the drug plasma concentration at time t.

3.6.2.4 Model Parameters The previous variables interact in a dynamic linear model via the following set of equations:

$$
\begin{aligned}
x_0 &\sim N(\mu_0, \Sigma_0) \\
x_t &= \Phi x_{t-1} + \Upsilon u_t + w_t \\
y_t &= A x_t + v_t \\
w &\sim N(0, Q) \\
v &\sim N(0, R)
\end{aligned}
$$

x_0 is the initial value of x with mean μ_0 and covariance Σ_0. Φ is the transition matrix between hidden states from one time point to the next. Υ is the response matrix to input interventions u. A is the emission matrix that translates the hidden state x to the observed measurements y. Finally, w and v are white noise terms with mean 0 and covariance matrices Q and R, respectively. Note the presence of a constant term $x^{(7)}$ in x allows us to model non-zero mean Gaussian measurement noise with a white Gaussian v. The covariance of w and v was set to values close to variances seen in actual data from MIMIC.

Because of the design of our homeostasis target variable, the homeostasis targets ($X^{(3)}$ for SBP, $X^{(4)}$ for DBP, and $X^{(6)}$ for HR) depend solely on initial parameters μ_0 and Σ_0. We simplify the model by zeroing out the rows and columns of Σ_0 that correspond to the homeostasis target variables and constants. Thus, the homeostasis targets depend solely on μ_0.

The AIM-BP DLM has a specific structure for Φ and Υ. These two parameters characterize how a given individual's vital signs trend over time, as well as how they respond to treatment.

The parameter Φ models the evolution of intrinsic vital signs. It has a specific structure:

$$\Phi = \begin{bmatrix} 1-r_h & 0 & r_h & 0 & 0 & 0 & 0 \\ 0 & 1-r_h & 0 & r_h & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1-r_h & r_h & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

In this structure, $r_h \in [0, 1]$ is a rate value that models various homeostatic mechanisms affecting blood pressure and heart rate, such as the renin-angiotensin system and the baroreceptor reflex. Since the baroreflex operates at a time scale of seconds and has a large effect on short term blood pressure homeostasis, we use a high r_h of 0.9 to model this fast homeostasis mechanism. Although we could theoretically sample and estimate a value for r_h, for all practical purposes with the amount of noise found in real data values of r_h greater than 0.8 have more or less the same effect on the system.

The parameter Υ models the response to perturbations u. The parameter Υ will vary depending on the number of input medications and which variables in x the medications affect. With n_{drug} medications, Υ will be a $7 \times 1 + n_{drug}$ matrix, where the first column captures the spontaneous perturbation and the rest of the columns capture the medications. As an example, we simulate two theoretical medications for blood pressure. Υ in this case will have this specific structure:

$$\Upsilon = \begin{bmatrix} r_h & r_h & r_h \\ r_h * ratio_{SD} & r_h * ratio_{SD} & r_h * ratio_{SD} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

Where $ratio_{SD}$ is the ratio of diastolic to systolic effects. When multiplied by u and plugged into the full DLM x equation, we get the following equation for systolic BP:

$$x_t^{(1)} = (1 - r_h) * x_{t-1}^{(1)} + r_h * (x_{t-1}^{(3)} + u^{(1)} + u^{(2)} + u^{(3)})$$

In effect, the perturbations from u modify the homeostasis target.

The emission matrix A translates the true blood pressure and heart rate to observed blood pressures and heart rates. In our simplified system with only one modality of observation for each, A looks like the following:

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}$$

We could add terms to A to capture non-zero mean errors in measurement, but for the simplicity of the model, we assume zero mean measurement errors. A is thus treated as known.

The covariance matrix Q defines the inherent variability of blood pressure and heart rate as well as capturing the baroreflex. The rows and columns of Q that correspond to constants and homeostasis targets in X will have zero variance and covariance. The baroreflex response occurs when receptors in the arterial walls of the aortic arch and carotid sinuses detect changes in blood pressure through arterial distension and generally operates on the order of seconds to minutes.[79] If our model had data at this time scale, we could theoretically simulate the baroreflex in detail by adding terms to Φ. However, because our model will generally operate at a maximum frequency of every quarter hour due to ICU charting practices, we can realistically only observe the baroreflex in operation as positive covariance between heart rate and blood pressure.

The covariance matrix R defines the measurement noise. Each method of measuring blood pressure and heart rate has an inherent measurement error. For simplicity, we assume measurement noise does not vary significantly from patient to patient and thus we set R to be a known value.

3.6.2.5 Proof of Sufficiency for Modeling the Homeostasis Perturbation Model

If we take Equation 3.6.2.4 and set $r_h = 1$ and substitute $P_{spon}(t)$ for $u^{(1)}$, $P_{med1}(t)$ for $u^{(2)}$, $P_{med2}(t)$ for $u^{(3)}$, and SBP_H for $x^{(3)}_{t-1}$, with unitary and known A and some noise ϵ_Q and ϵ_R generated from Q and R, we get:

$$
\begin{aligned}
x^{(1)}_t &= (x^{(3)}_{t-1} + u^{(1)} + u^{(2)} + u^{(3)}) + \epsilon_Q \\
x^{(1)}_t &= (SBP_H + P_{spon}(t) + \sum_{meds} P_{med}(t)) + \epsilon_Q \\
SBP_t = y^{(1)}_t &= (Ax_t)^{(1)} + \epsilon_R = (SBP_H + P_{spon}(t) + \sum_{meds} P_{med}(t)) + (\epsilon_Q + \epsilon_R)
\end{aligned}
$$

Which is the exact homeostasis perturbation model described in Equation 3.1 with $\epsilon_S = \epsilon_Q + \epsilon_R$. A parallel construction can be made for DBP. Thus, we have shown that the AIM-BP model is sufficient for capturing homeostasis perturbation model parameters. **This proves claim 1 of this thesis.**

For the purposes of this thesis, we are interested in estimating the homeostasis targets (e.g. $x^{(3)}$ or $\mu^{(3)}_0$ for systolic BP), the rate of decay for the stroke blood pressure perturbation r_B, and the drug E_{max} and EC_{50} parameters.

4.0 EVALUATION USING CLINICAL SCENARIOS

Because the AIM-BP model parameters cannot be independently verified in actual patient data, the accuracy of our parameter estimation system must be tested using simulated data. The DLM basis for the AIM-BP framework provides a natural generative process for data simulation. By simulating data, we can know the ground truth parameters to compare estimated parameters against. By using the DLM as the basis for simulating data, we make the assumption that a first order autoregressive process is a sufficient model for actual blood pressure behavior.

A typical method of testing the AIM parameter estimator would be to evaluate its performance over multiple trials over all possible configurations of ground truth parameters. Practically speaking, however, the number of possible configurations of ground truth parameters, in combination with the computational time required to estimate parameters, and the low clinical applicability of many parameter configurations, make this sort of evaluation low yield for the amount of time required. Instead, we chose to focus on a proof of concept evaluation through several exemplar clinical scenarios that demonstrate various aspects of the AIM-BP model. We first performed a limited sensitivity analysis on three clinically important AIM-BP model parameters, then evaluated the accuracy of the AIM-BP parameter estimator against a standard curve fitting technique in four clinical scenarios.

4.1 METHODS

4.1.1 Data Simulation

As a proof of concept of the AIM-BP framework, we simulated data from four clinical scenarios using the AIM-BP data simulator, learned model parameters from the simulated data using the AIM-BP parameter estimator, and compared the learned parameters to the ground truth model parameters used to simulate the data.

We crafted four different clinical scenarios that would present with similar blood pressure charts, based on the critical care stroke blood pressure management protocol at the neuro ICU at UPMC Presbyterian. Our goal is to demonstrate the ability of the AIM-BP framework to identify and estimate different combinations of drug effects and spontaneous blood pressure behavior that together add to similar blood pressure trends. The clinical scenarios all start out the same way:

> A 70 year old male patient arrives at the ICU with a primary intracerebral hemorrhage. His blood pressure at admission to the ICU is 220/110 mmHg.

The scenarios then diverge as follows:

1. Scenario 1: The patient is given an initial push of 20mg IV labetalol twice to bring his systolic blood pressure down to less than 140 mmHg. No further medication is needed to bring keep his systolic blood pressure under 140 mmHg. Under the hood, this patient's blood pressure would have spontaneously trended to 130/80 mmHg, and the IV labetalol was of medium effectiveness, with an E_{max} of -40 mmHg and an EC_{50} of 70 ng/mL.
2. Scenario 2: The patient is given the exact same push of 20mg IV labetalol twice as in Scenario 1, and his blood pressure behaves similarly. Under the hood, however, this patient's blood pressure would have spontaneously trended to 150/90 mmHg, but the IV labetalol was particularly effective with an E_{max} of -60 mmHg and an EC_{50} of 40 ng/mL.
3. Scenario 3: The patient is given the exact same push of 20mg IV labetalol twice as in Scenario 1, but it is not very effective. He is then started on an IV nicardipine drip, which is titrated towards a target blood pressure of 140 mmHg systolic and only a small

drip rate is needed to be effective. Under the hood, this patient's blood pressure would have spontaneously trended to 180/100 mmHg, the IV labetalol was not very effective with an E_{max} of -20 mmHg and an EC_{50} of 160 ng/mL, and the IV nicardipine was very effective with an E_{max} of -60 mmHg and an EC_{50} of 40 ng/mL.

4. Scenario 4: The patient is given the exact same push of 20mg IV labetalol twice as in Scenario 1, but it is not very effective. He is then started on an IV nicardipine drip, which is titrated towards a target blood pressure of 140 mmHg systolic and a medium drip rate is needed to be effective. Under the hood, this patient's blood pressure would have spontaneously trended to 160/100 mmHg, the IV labetalol was not very effective with an E_{max} of -20 mmHg and an EC_{50} of 160 ng/mL, and the IV nicardipine was moderately effective with an E_{max} of -40 mmHg and an EC_{50} of 70 ng/mL.

The blood pressure chart of each clinical scenario is shown in Fig 11. These scenarios differ in three primary parameters that we will focus on: $\mu_0^{(3)}$, the homeostasis baseline for SBP, E_{max}, the maximum effect achievable by a drug, and EC_{50}, the drug concentration needed to achieve half of the maximum effect. The specific configurations of the relevant parameters can be found in Table 3.

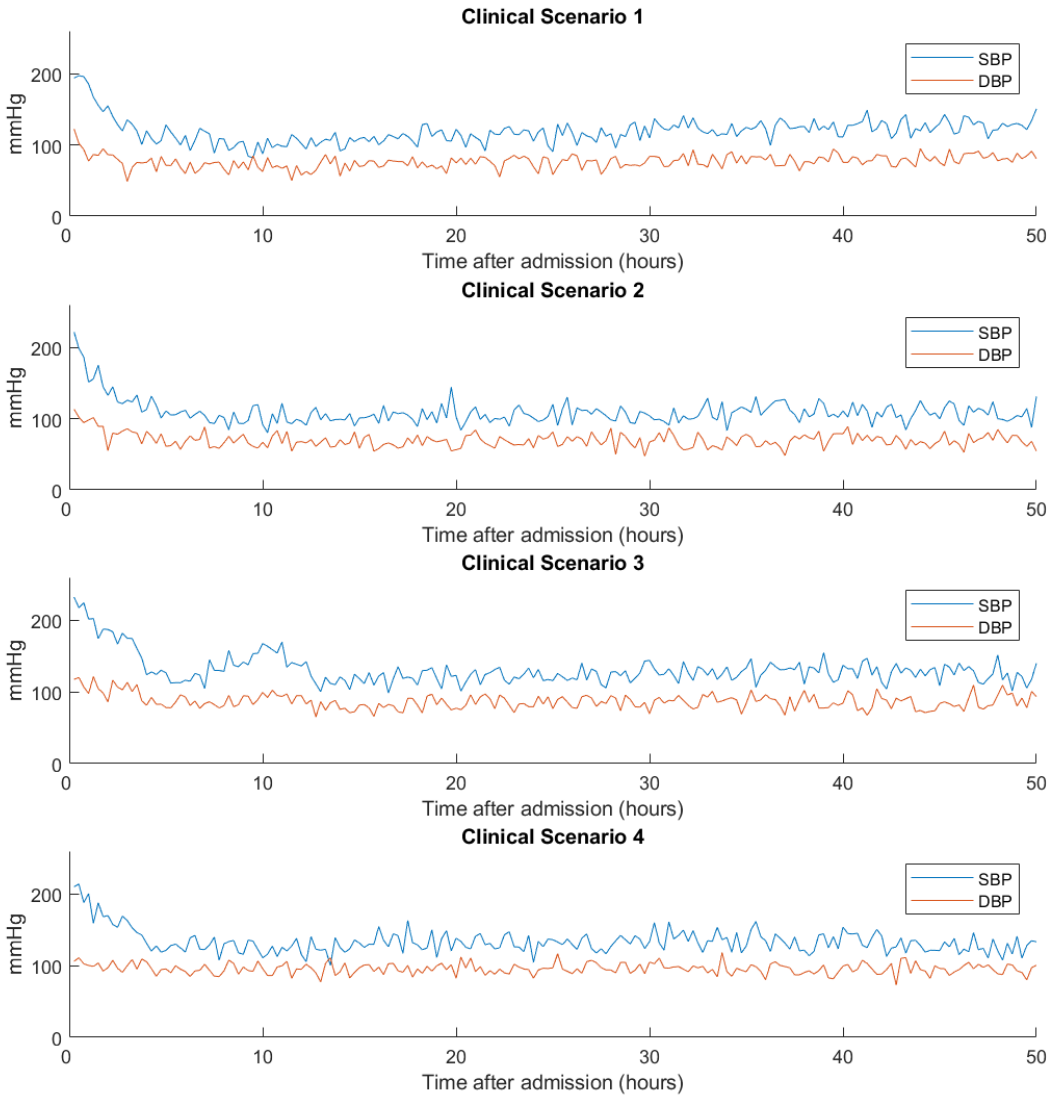

Figure 11: Blood pressure charts for each clinical scenario. Each clinical scenario presents with a BP of around 220/110 mmHg, which is managed down to around 140 mmHg systolic using medication. The amount and type of medication used to manage blood pressure varies depending on the scenario.

Table 3: Clinical Scenario parameter configurations.

Parameter	Scenario 1	Scenario 2	Scenario 3	Scenario 4
Initial SBP ($\mu_0^{(1)}$)	220 mmHg			
SBP Homeostasis Base ($\mu_0^{(3)}$)	130 mmHg	110 mmHg	180 mmHg	160 mmHg
Labetalol E_{max}	-40 mmHg	-20 mmHg	-20 mmHg	-20 mmHg
Labetalol EC_{50}	70 ng/mL	110 ng/mL	160 ng/mL	160 ng/mL
Nicardipine E_{max}	N/A	N/A	-60 mmHg	-40 mmHg
Nicardipine EC_{50}	N/A	N/A	40 ng/mL	70 ng/mL
Initial DBP ($\mu_0^{(2)}$)	110 mmHg			
DBP Homeostasis Base ($\mu_0^{(4)}$)	80 mmHg	70 mmHg	100 mmHg	100 mmHg
Initial HR ($\mu_0^{(5)}$)	70 bpm			
HR Homeostasis Base ($\mu_0^{(6)}$)	70 bpm			
Spontaneous perturbation decay rate (r_B)	0.1			
BP homeostasis rate (r_h)	0.9			

E_{max} and EC_{50} ranges for labetalol and nicardipine were roughly estimated from the literature.[80–87] Time to onset of peak concentration, half-life and central volumes of distribution were taken from literature or reference websites (https://www.rxlist.com, https://pubchem.ncbi.nlm.nih.gov/). Volumes of distribution, peak onset, and half-lifes were used to estimate drug plasma concentrations after doses in a rudimentary pharmacokinetics model, using the volume of distribution to calculate the peak plasma concentration, with a

linear slope up to the peak for the duration of peak onset and an exponential decay after the peak concentration calculated using the half-life.

We simulated the adminstration of medication in a manner that roughly adheres to the ICH blood pressure management guidelines at UPMC Presbyterian. In all clinical scenarios, 20mg of IV labetalol was given twice initially. In clinical scenarios 3 and 4, IV nicardipine was started at 5mg/hr and uptitrated 2.5mg/hr to a maximum of 15mg/hr until a systolic blood pressure of 140 mmHg was achieved, then decreased to 3mg/hr for maintenance with adjustments as necessary to achieve a blood pressure target of 140 mmHg systolic.

We chose to simulate our scenarios using 15 minute time intervals for a period of 200 time points, or a little more than 2 days. The rationale for 15 minute time intervals was that was the most typical time interval observed in chart data between measurements. We chose to simulate for a period of 2 days as a balance between having enough time to observe patients stabilizing (factoring in drug half-lifes) and being a short enough time to enable mid-admission decision making.

4.1.2 Parameter estimation

A Metropolis-within-Gibbs sampler was written in Matlab (https://www.mathworks.com) for parameter estimation. The method by which each parameter is estimated is listed in Table 4. A Metropolis step can be taken with either Metropolis Hastings, or a variation on ARMS, which could affect convergence speed. We tested each method and chose standard Metropolis Hastings for speed. Priors for μ_0 are either multivariate normal (MVN) or normal (N), expressed as means and variances. The prior for r_B is a uniform distribution over $[0.005, 0.2]$, which roughly corresponds to a range between 4 hours and 8 days for a 99% decrease in the perturbation $u^{(1)}$.

Table 4: AIM-BP parameters and the method and priors used in their estimation.

Parameter	Estimation Method	Prior
X, X_0	Forward Filtering Backward Sampling algorithm[65]	None
$\mu_0^{(1,2,5)}$	Gibbs step	Conjugate $MVN(\begin{bmatrix} 220 \\ 110 \\ 70 \end{bmatrix}, \begin{bmatrix} 0.1 & 0 & 0 \\ 0 & 0.1 & 0 \\ 0 & 0 & 0.1 \end{bmatrix})$
$\mu_0^{(3,4,6)}$	Metropolis step	$N(140, 1600)$ $N(90, 900)$ $N(70, 100)$
$\mu_0^{(7)}$	Constant and known	None
Σ_0	Constant and known	None
r_B	Metropolis step	$Unif(0.005, 0.2)$
B_{max}	Derived from μ_0	None
E_{max}, EC_{50}	Metropolis step	Labetalol: $MVN(\begin{bmatrix} -30 \\ 110 \end{bmatrix}, \begin{bmatrix} 900 & 0 \\ 0 & 2500 \end{bmatrix})$ Nicardipine: $MVN(\begin{bmatrix} -30 \\ 70 \end{bmatrix}, \begin{bmatrix} 900 & 0 \\ 0 & 900 \end{bmatrix})$
Σ_0	Constant and known	None
Φ	Constant and known	None
Υ	Derived from B_{max} and μ_0	None
Q	Constant and known	None

Parameter	Estimation Method	Prior
A	Constant and known	None
R	Constant and known	None

In order to estimate μ_0, $\mu_0^{(1,2,5)}$ is first estimated using a Gibbs step with a conjugate multivariate normal prior. Homeostasis targets $\mu_0^{(3,4,6)}$ are then estimated by individual Metropolis steps, sampling from the probability:

$$Pr(\mu_0|\Sigma_0, \Phi, \Upsilon, Q, A, R, x, x_0, y) \propto Pr(\mu_0)Pr(x_0|\mu_0, \Sigma_0) \prod_{t=1}^{T} Pr(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t)$$

We chose a normal prior for $\Pr(\mu_0)$. B_{max} and $ratio_{SD}$ is calculated from μ_0.

r_B is estimated by sampling from the following probability:

$$Pr(r_B|\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u, x, y) \propto \prod_{t=1}^{T} Pr(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t)Pr(r_B)$$

E_{max} and EC_{50} are jointly estimated by sampling from the following probability:

$$Pr(E_{max}, EC_{50}|\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u, x, y) \propto \prod_{t=1}^{T} Pr(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t)Pr(E_{max}, EC_{50})$$

We ran the MCMC sampler with a burn-in period of 2000 steps and then sampled every 5 steps until we reach 2000 samples. We took the mean of the sampled values as the estimate for the parameters.

4.1.3 Evaluation

We evaluated the performance of our AIM-BP framework by their ability to recover three parameters of clinical importance: $\mu_0^{(3)}$, the systolic BP homeostasis baseline, and E_{max} and EC_{50} for each medication. We chose to limit the evaluation to these parameters, instead of the full set of learned parameters (Table 4), for the purposes of clarity and simplicity.

In addition to the individual E_{max} and EC_{50} parameters, we also calculated an aggregate parameter we call the E_{max} ratio, which is a normalized ratio of the two separate parameters:

$$ER = \frac{E_{max}}{EC_50 + mean(c(c > 0))}$$

It has been shown that when dose-response curves do not explore the maximal dosage ranges sufficiently, individual E_{max} and EC_{50} estimates can be highly inaccurate but their ratio is a more stable parameter.[88] We thus introduced the E_{max} ratio parameter as a more stable, aggregate metric of drug effectiveness. A drug is more effective when its E_{max} is higher or when its EC_{50} is lower, so a higher E_{max} ratio corresponds to a greater measure of effectiveness. We normalized the EC_{50} in the denominator by adding the mean concentration of the drug during its active period (when the concentration is non-zero). This helps further stabilize the ratio.

4.1.3.1 Sensitivity Analysis We first performed a sensitivity analysis using the parameter settings for Scenario 1 as a starting point. In one experiment, we varied the ground truth SBP homeostasis baseline $\mu_0^{(3)}$ from 90 mmHg to 270 mmHg. In a second experiment, we performed a grid search on ground truth $E_{max} = \{-60, -30, -10\}$ and $EC_{50} = \{20, 110, 200\}$. For each configuration of ground truth parameters, we simulated 20 sets of data using the AIM-BP data simulator, and estimated parameters for each simulated set of data. We calculated the absolute error of estimated parameters from ground truth and compared sets of absolute errors across the ranges for each experiment. We chose to report absolute error over relative error in order to be able to compare across different ground truth values.

In addition to evaluating absolute error, we qualitatively analyzed the MCMC samples of r_B, $\mu_0^{(3)}$, E_{max} and EC_{50} for convergence of the chain using the methods described in 2.3.6.4.

We included the plots for the derived parameter B_{max} as well. Its chain trace should closely resemble the one for $\mu_0^{(3)}$, except flipped vertically. We visually inspected the chain traces to see if the means remained relatively constant and the autocorrelation functions to determine the level of independence between individual samples in a single MCMC run.

4.1.3.2 Clinical Scenario Evaluation Next, we evaluated the performance of the AIM-BP parameter estimator on the four previously described clinical scenarios. Ideally, we would compare our performance against the performance of a state of the art algorithm. To our knowledge, however, we are the first to attempt to model spontaneous blood pressure trends and medication pharmacodynamics in the critical care stroke setting simultaneously, so no such state of the art algorithm exists.

While no studies have looked at our combined model, numerous studies fit pharmacodynamics parameters for the E_{max} model when studying dose-response relationships. These are usually fit using maximum likelihood estimation via a non-linear regression method such as Gauss-Newton or Levenberg-Marquardt.[61,63,89] We will refer to these methods as Non-Linear Least Squares (NLLS) regression methods. Instead of fitting just the E_{max} model using these methods, however, we fitted the combined perturbation model in Equation 3.1. We used the NLLS fitted parameter estimates as a state-of-the-art comparison to the AIM-BP parameter estimator.

We compared the AIM-BP framework to a basic parameter fitting method using a Matlab NLLS curve fitting method (`lsqcurvefit`) on the SBP equation in Equation 3.1, where SBP_H corresponds to $\mu_0^{(3)}$ and other variable names are equivalent.

First, the AIM-BP data simulator was used to simulate 100 sets of data for each clinical scenario, with the corresponding parameters for the SBP homeostasis baseline and drug effects. We then ran NLLS parameter fitting method and the AIM-BP parameter estimator once on each simulated dataset to obtain a set of 100 estimated parameters for each scenario for each method. The absolute errors of estimated parameters frome each method compared to ground truth were calculated to determine if AIM-BP has on average smaller absolute error compared to NLLS, using a t-test on the paired differences of the absolute errors. Similarly, whether AIM-BP has more consistent estimates compared to NLLS was determined using an

F-test on variances of estimated parameters. Finally, we investigated the correlation between absolute errors from AIM-BP versus NLLS.

4.2 RESULTS

4.2.1 Sensitivity Analysis

Absolute errors of AIM-BP estimated SBP homeostasis baseline ($\mu_0^{(3)}$), E_{max}, EC_{50}, and E_{max} ratio of IV labetalol over a range of ground truth SBP homeostasis baselines are shown in Figure 12. Errors for estimated SBP homeostasis baseline consistently remained below 10 mmHg throughout the range of ground truth values. Errors for E_{max} and EC_{50} were consistent as well, though errors for EC_{50} approached the same magnitude as the actual ground truth value. Errors for the E_{max} ratio were much smaller than the derived ground truth value, however, meaning relative errors were smaller for E_{max} ratio compared to EC_{50}. Performance was consistent across the entire range of ground truth SBP homeostasis baseline values.

Absolute errors of AIM-BP estimated SBP homeostasis baseline ($\mu_0^{(3)}$), E_{max}, EC_{50}, and E_{max} ratio of IV labetalol over a range of E_{max} model parameters are shown in Figure 13. Like in the previous experiment, EC_{50} errors were higher relative magnitude than other parameters, but were consistent regardless of the magnitude of the ground truth EC_{50} value. The E_{max} ratio estimates were consistent with E_{max} ratio estimates in the previous experiment as well.

We visually examined the histograms, chain traces, and autocorrelation functions of B_{max}, r_B, $\mu_0^{(3)}$, E_{max} and EC_{50} samples for each set of ground truth parameters in the sensitivity analysis, three examples of which can be seen in Figures 14, 15 and 16. The histograms show that sampled values fell within a normal distribution. The chain traces show that the mean of sampled values is relatively constant, so we have likely reached convergence. The ACFs show that sampled values are still significantly correlated within a lag of about 50 samples or so, so the number of truly independent samples is much lower

than 2000. Since we have likely reached convergence and estimated values look relatively stable, we chose not to make further changes to the MCMC sampling configuration.

From both experiments, we concluded that absolute errors were consistent over the range of ground truth parameter values tested.

4.2.2 Clinical Scenario Evaluation

In clinical scenarios 1 (Fig 17) and 2 (Fig 18), both the basic non-linear least squares (NLLS) and the AIM-BP parameter estimation methods performed equally well in estimating the systolic BP baseline $\mu_0^(3)$, achieving mean absolute errors of less than 2 mmHg (Table 5). Absolute errors were significantly different for SBP baseline in Scenario 1, but by a negligible amount, and not significantly different in Scenario 2. In both scenarios, NLLS exhibited wider variability in estimating E_{max} and EC_{50} values, but no significant difference was seen in the E_{max} ratio parameter.

In the more complex two-drug clinical scenarios 3 (Fig 19) and 4 (Fig 20), AIM-BP estimated parameters are significantly more consistent than NLLS estimated parameters for all parameters except E_{max} ratio. Outliers in NLLS estimates meant AIM-BP performed better at estimating labetalol E_{max} ratio in Scenario 3, but no difference in labetalol E_{max} ratio was seen in Scenario 4. The nicardipine E_{max} ratio, on the other hand, was more consistently estimated by AIM-BP than NLLS in both scenarios.

In terms of mean absolute error, AIM-BP generally consistently had smaller errors than NLLS in estimating E_{max} and EC_{50}. For the SBP homeostasis baseline, AIM-BP was equally good (with negligibly better statistically significant performance) as NLLS in Scenarios 1 and 2, but a larger performance gap was apparent in Scenarios 3 and 4. For E_{max} ratio, both methods performed relatively consistently in all scenarios.

Table 5: Scenarios 1-4: Mean and standard deviation of absolute error of estimated parameters from NLLS and AIM-BP compared to ground truth. Significantly lower AIM-BP absolute errors are noted with *'s (* = $p < 0.0005$, ** = $p < 0.005$, * = $p < 0.05$).**

Scenario 1	**Ground Truth**	**NLLS Error**	**AIM-BP Error**
SBP homeostasis baseline ($\mu_0^{(3)}$)	130.0	1.5 ± 0.9	1.3 ± 1.0 *
Labetalol E_{max}	-40.0	8.2 ± 8.8	7.4 ± 6.3
Labetalol EC_{50}	70.0	30.2 ± 34.8	31.9 ± 21.2
Labetalol E_{max} ratio	-0.25	0.02 ± 0.02	0.02 ± 0.02
Scenario 2	**Ground Truth**	**NLLS Error**	**AIM-BP Error**
SBP homeostasis baseline ($\mu_0^{(3)}$)	110.0	1.3 ± 1.3	1.1 ± 1.5
Labetalol E_{max}	-20.0	10.4 ± 11.1	5.3 ± 5.6 ***
Labetalol EC_{50}	110.0	103.6 ± 110.0	22.7 ± 17.9 ***
Labetalol E_{max} ratio	-0.10	0.02 ± 0.02	0.02 ± 0.02
Scenario 3	**Ground Truth**	**NLLS Error**	**AIM-BP Error**
SBP homeostasis baseline ($\mu_0^{(3)}$)	180.0	10.1 ± 21.2	5.7 ± 4.2 *
Labetalol E_{max}	-20.0	10.1 ± 8.9	3.8 ± 2.8 ***
Labetalol EC_{50}	160.0	130.6 ± 106.4	33.3 ± 15.8 ***
Labetalol E_{max} ratio	-0.07	0.02 ± 0.03	0.01 ± 0.01 **
Nicardipine E_{max}	-60.0	7.3 ± 6.2	3.8 ± 2.8 **
Nicardipine EC_{50}	40.0	16.6 ± 27.0	17.6 ± 12.9

Nicardipine E_{max} ratio	-0.34	0.06 ± 0.03	0.06 ± 0.04
Scenario 4	**Ground Truth**	**NLLS Error**	**AIM-BP Error**
SBP homeostasis baseline ($\mu_0^{(3)}$)	160.0	7.3 ± 5.5	5.3 ± 3.4 ***
Labetalol E_{max}	-20.0	10.6 ± 11.2	5.1 ± 6.1 ***
Labetalol EC_{50}	160.0	135.7 ± 108.3	35.2 ± 17.6 ***
Labetalol E_{max} ratio	-0.07	0.02 ± 0.02	0.02 ± 0.03
Nicardipine E_{max}	-40.0	15.5 ± 14.5	5.1 ± 6.1 ***
Nicardipine EC_{50}	70.0	85.6 ± 105.0	13.8 ± 9.3 ***
Nicardipine E_{max} ratio	-0.19	0.05 ± 0.04	0.04 ± 0.03 ***

Because data was simulated using the AIM-BP model, one possible explanation for the increased performance of the AIM-BP parameter estimator may simply be that the data was simulated with the same model. To test if this is the case, we simulated data using just the homeostasis perturbation model plus white noise with variance $Q^{(1,2,5)} + R$, which is the same model upon which NLLS is fit. For clarity, we will refer to this mechanism of data simulation as the NLLS mechanism. We observed the same increase in performance for Scenarios 3 and 4 using the NLLS mechanism as when the data was simulated using an AIM-BP mechanism (Figure 21, Figure 22, and Table 6).

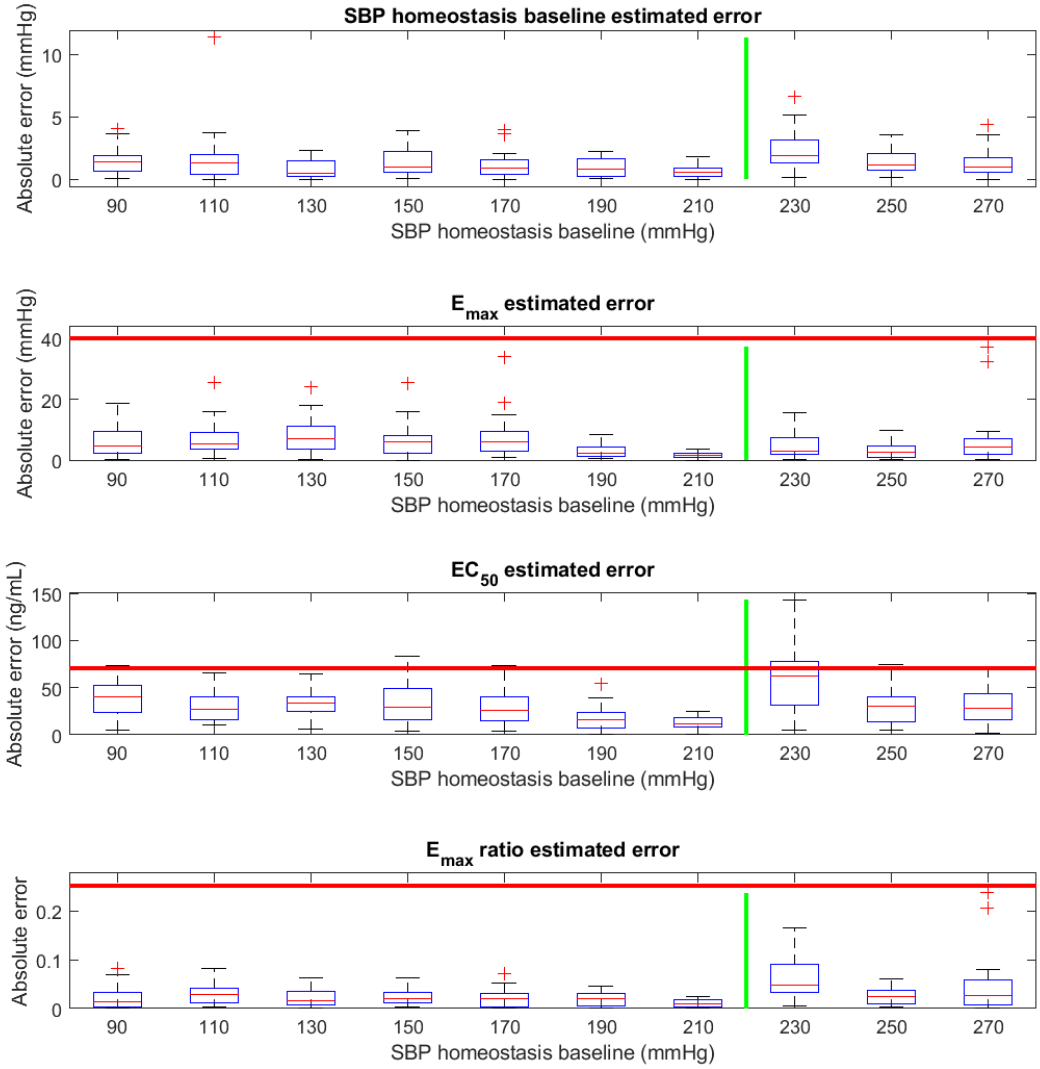

Figure 12: Sensitivity Analysis on Ground Truth SBP Homeostasis Baseline. Box and whisker plots of absolute errors of AIM-BP estimated and derived parameters are shown across a range of ground truth SBP homeostasis baseline values. Green lines represent initial SBP. Red lines represent the magnitude of ground truth values for E_{max}, EC_{50}, and calculated ground truth value for E_{max} ratio.

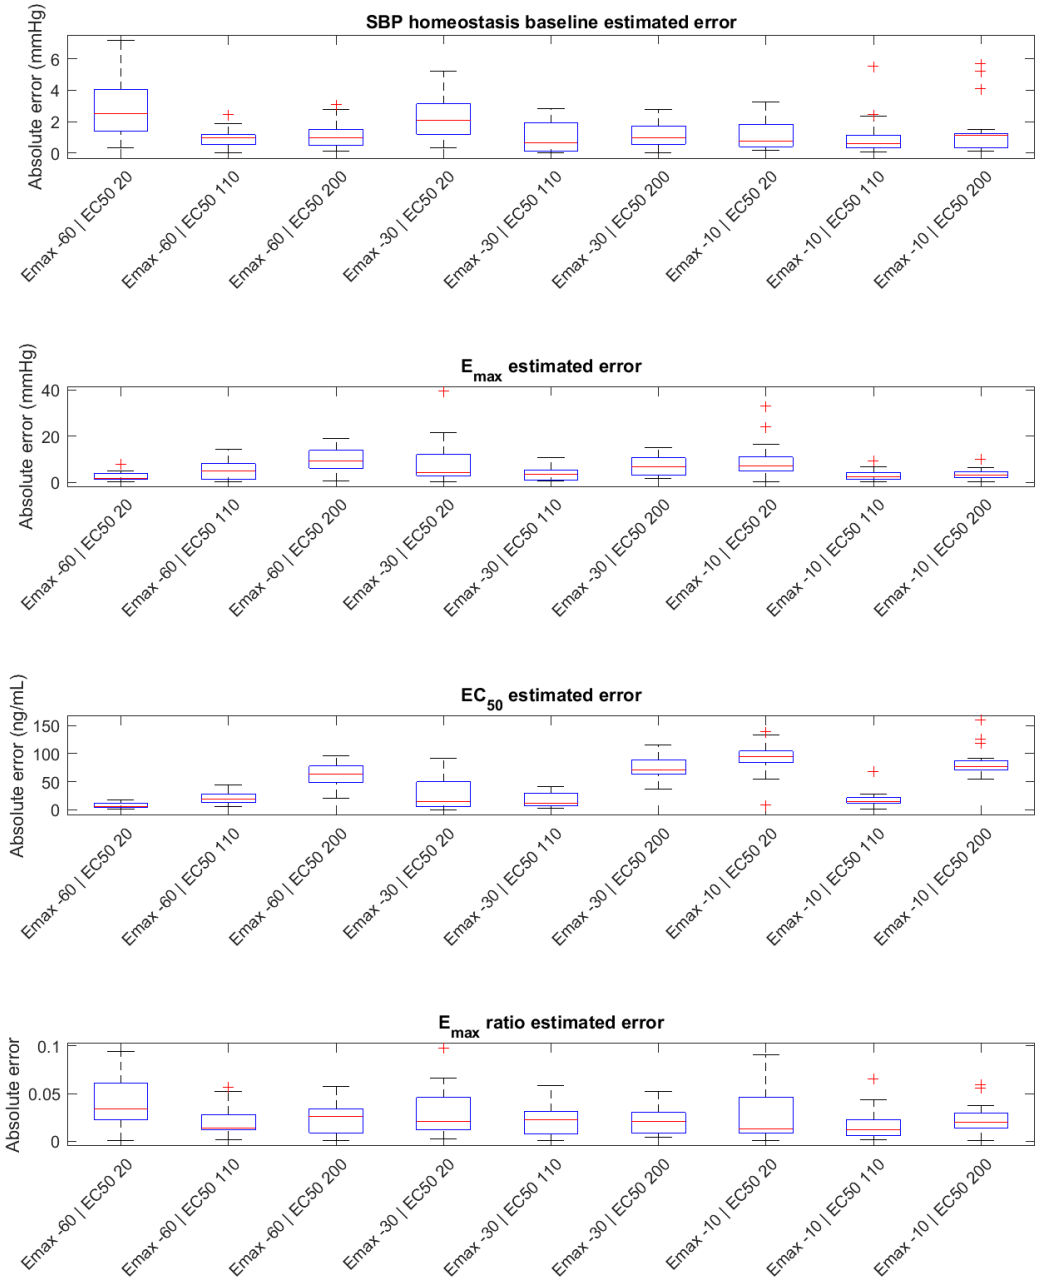

Figure 13: Sensitivity Analysis on Ground Truth E_{max} parameters. Box and whisker plots of absolute errors of AIM-BP estimated and derived parameters are shown across a range of ground truth E_{max} and EC_{50} values.

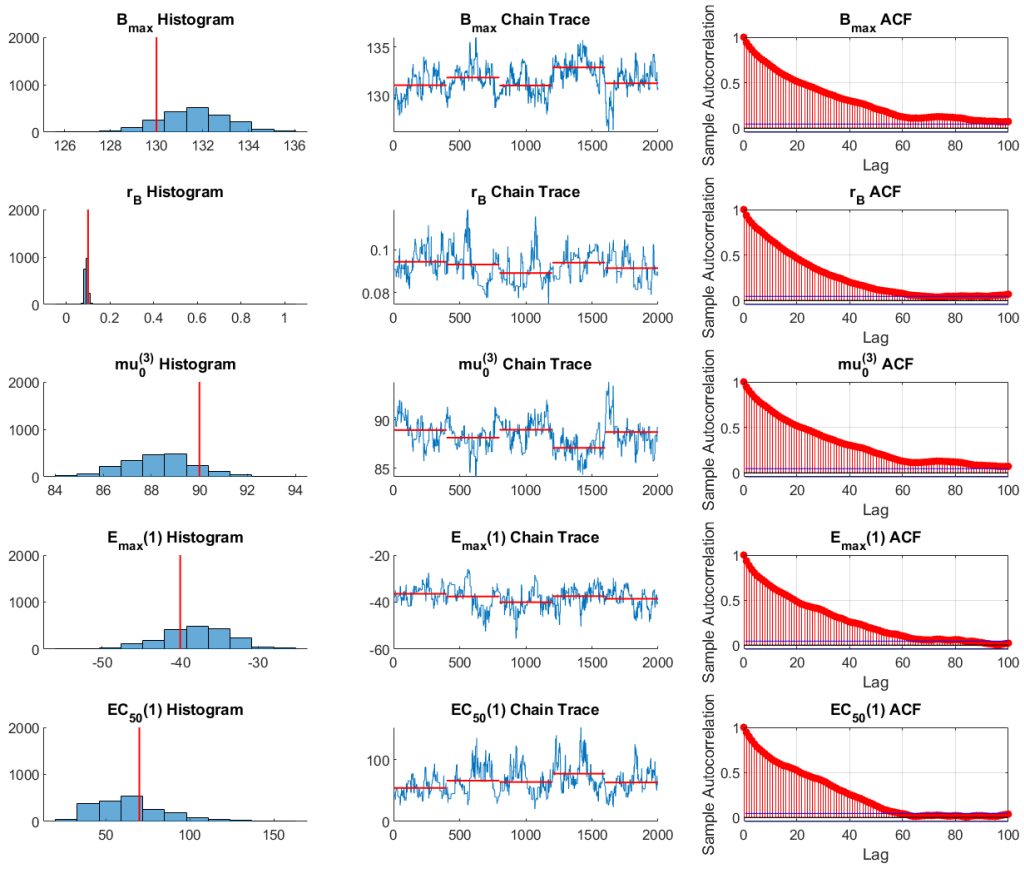

Figure 14: Histogram, chain trace, and ACF for SBP baseline 90 mmHg. Histogram, chain trace, and autocorrelation function of a single MCMC run with SBP homeostasis baseline = 90 mmHg.

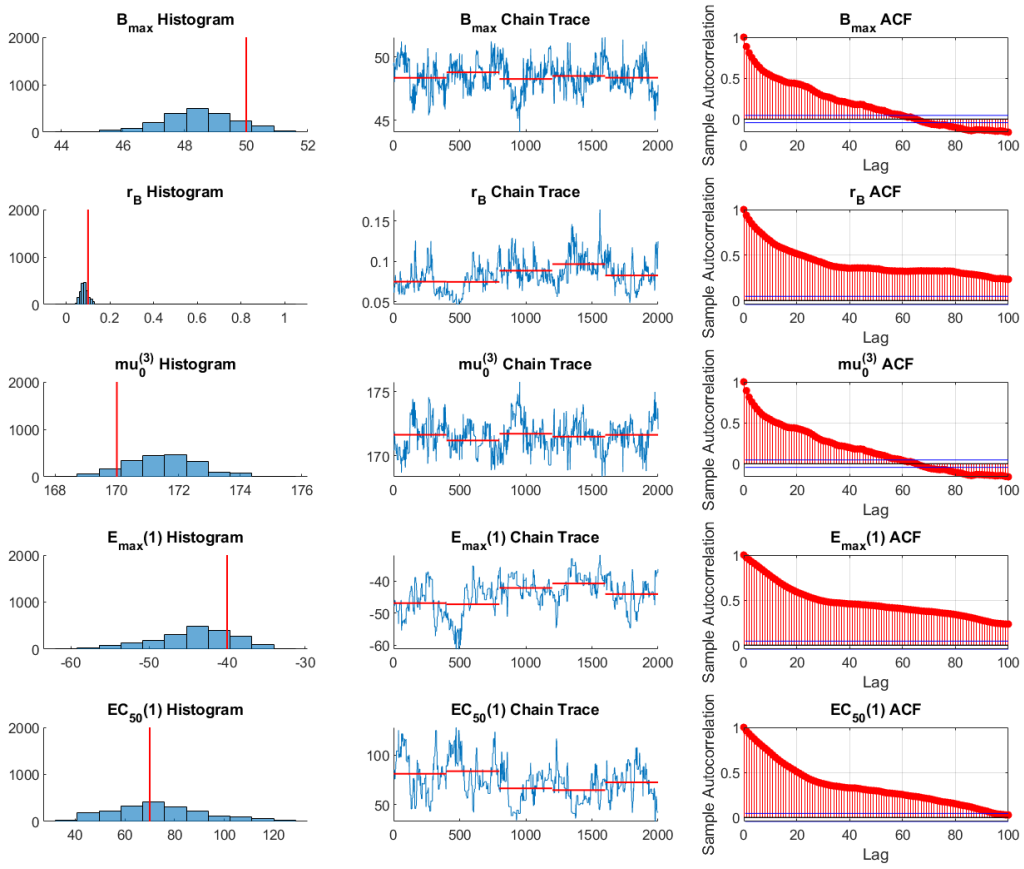

Figure 15: Histogram, chain trace, and ACF for SBP baseline 190 mmHg. Histogram, chain trace, and autocorrelation function of a single MCMC run with SBP homeostasis baseline = 190 mmHg.

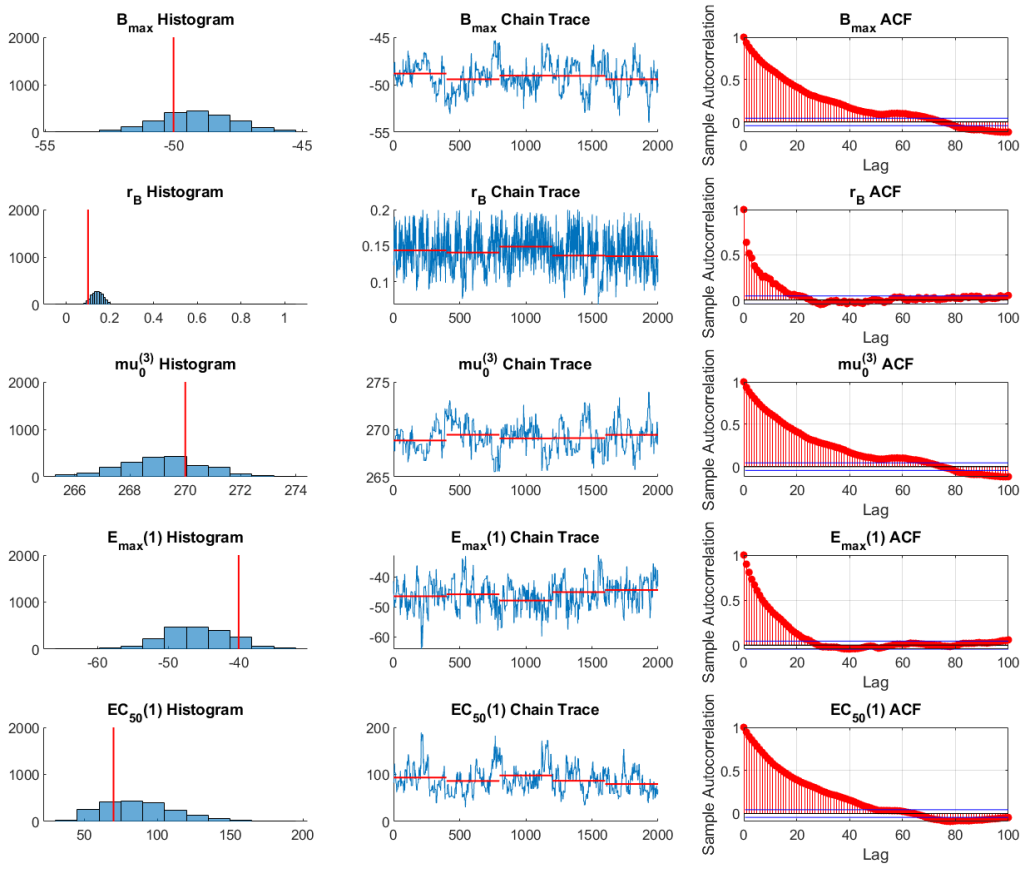

Figure 16: Histogram, chain trace, and ACF for SBP baseline 270 mmHg. Histogram, chain trace, and autocorrelation function of a single MCMC run with SBP homeostasis baseline = 270 mmHg.

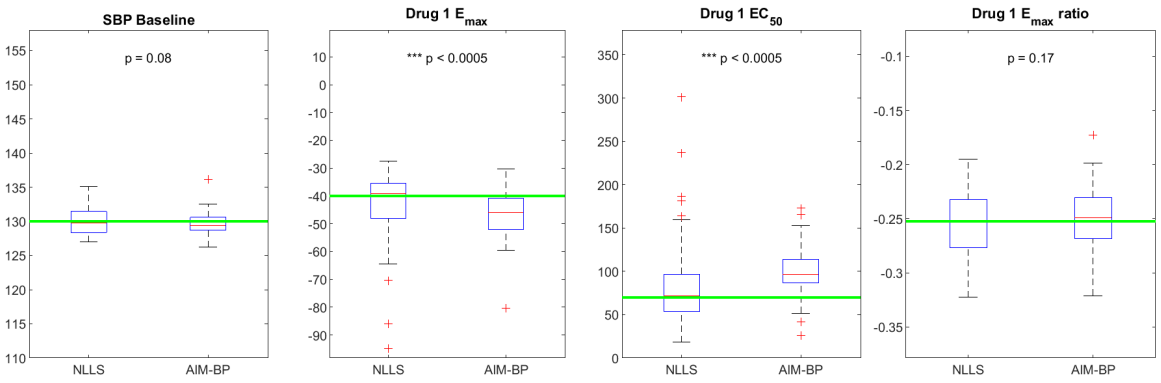

Figure 17: Clinical Scenario 1. Box and whisker plots of estimated parameters using non-linear least squares (NLLS) and AIM-BP. Significantly lower variances of AIM-BP estimated parameters compared to NLLS are noted with *'s and p-values.

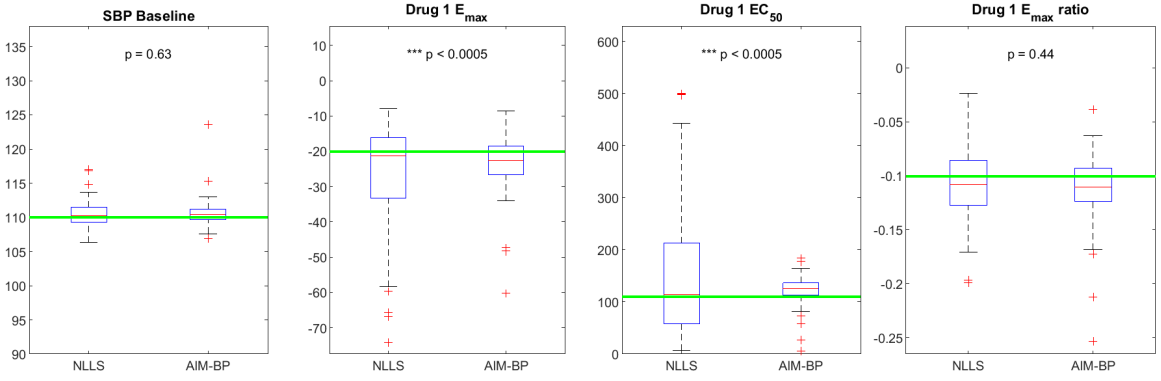

Figure 18: Clinical Scenario 2. Box and whisker plots of estimated parameters using non-linear least squares (NLLS) and AIM-BP. Significantly lower variances of AIM-BP estimated parameters compared to NLLS are noted with *'s and p-values.

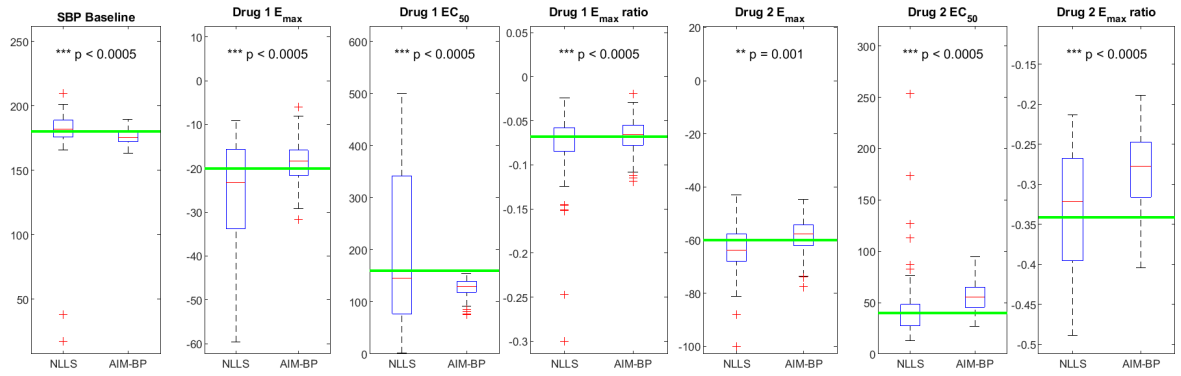

Figure 19: Clinical Scenario 3. Box and whisker plots of estimated parameters using non-linear least squares (NLLS) and AIM-BP. Significantly lower variances of AIM-BP estimated parameters compared to NLLS are noted with *'s and p-values.

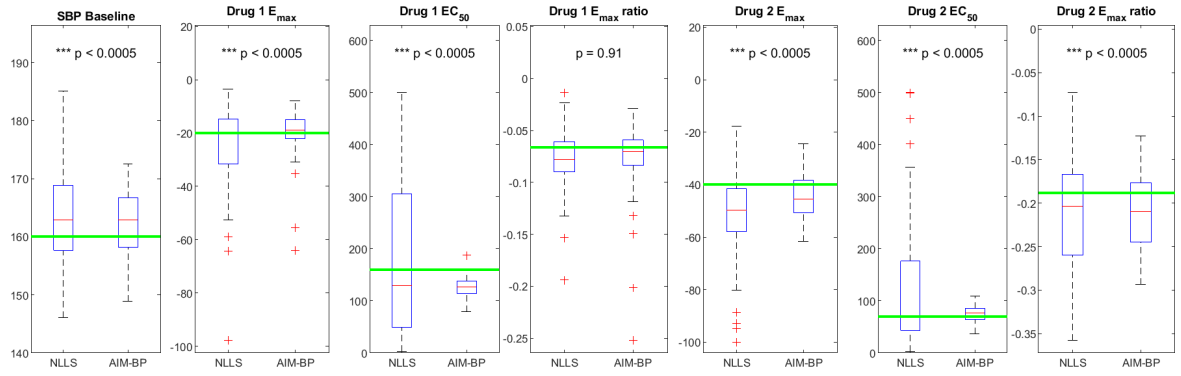

Figure 20: Clinical Scenario 4. Box and whisker plots of estimated parameters using non-linear least squares (NLLS) and AIM-BP. Significantly lower variances of AIM-BP estimated parameters compared to NLLS are noted with *'s and p-values.

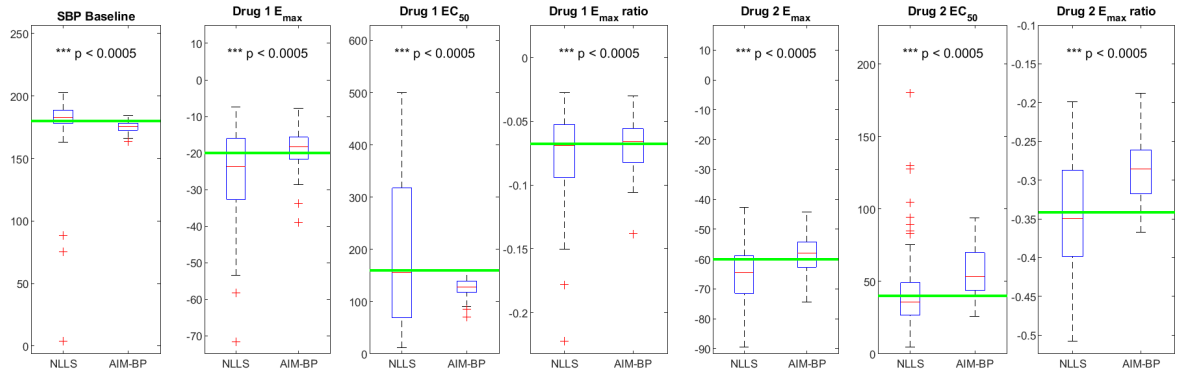

Figure 21: Clinical Scenario 3 with NLLS simulation mechanism. Box and whisker plots of estimated parameters using non-linear least squares (NLLS) and AIM-BP. Significantly lower variances of AIM-BP estimated parameters compared to NLLS are noted with *'s and p-values.

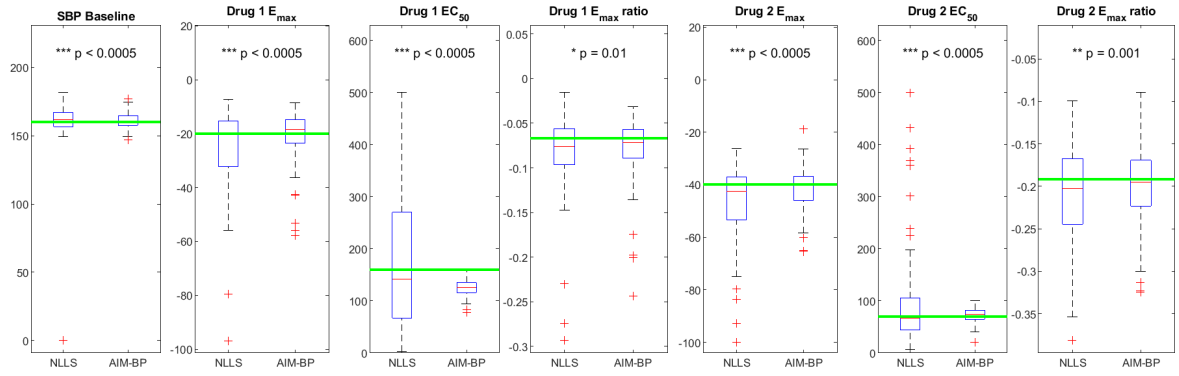

Figure 22: Clinical Scenario 4 with NLLS simulation mechanism. Box and whisker plots of estimated parameters using non-linear least squares (NLLS) and AIM-BP. Significantly lower variances of AIM-BP estimated parameters compared to NLLS are noted with *'s and p-values.

Table 6: Mean and standard deviation of absolute error of estimated parameters from NLLS and AIM-BP compared to ground truth using NLLS simulation mechanism. Significantly lower AIM-BP absolute errors are noted with *'s (* = $p < 0.0005$, ** = $p < 0.005$, * = $p < 0.05$).**

Scenario 3	**Ground Truth**	**NLLS Error**	**AIM-BP Error**
SBP homeostasis baseline ($\mu_0^{(3)}$)	180.0	12.1 ± 27.4	5.1 ± 3.6 *
Labetalol E_{max}	-20.0	10.1 ± 9.0	4.0 ± 3.2 ***
Labetalol EC_{50}	160.0	139.4 ± 106.5	32.3 ± 16.0 ***
Labetalol E_{max} ratio	-0.07	0.02 ± 0.02	0.02 ± 0.01 ***
Nicardipine E_{max}	-60.0	7.9 ± 6.4	4.0 ± 3.2 ***
Nicardipine EC_{50}	40.0	17.3 ± 20.1	18.0 ± 13.9
Nicardipine E_{max} ratio	-0.34	0.06 ± 0.04	0.06 ± 0.04
Scenario 4	**Ground Truth**	**NLLS Error**	**AIM-BP Error**
SBP homeostasis baseline ($\mu_0^{(3)}$)	160.0	7.6 ± 16.1	4.1 ± 3.4 *
Labetalol E_{max}	-20.0	11.2 ± 12.3	6.3 ± 6.7 ***
Labetalol EC_{50}	160.0	128.7 ± 104.8	34.3 ± 15.8 ***
Labetalol E_{max} ratio	-0.07	0.03 ± 0.03	0.02 ± 0.03
Nicardipine E_{max}	-40.0	11.5 ± 13.2	6.3 ± 6.7 ***
Nicardipine EC_{50}	70.0	65.4 ± 107.2	11.0 ± 8.4 ***
Nicardipine E_{max} ratio	-0.19	0.05 ± 0.04	0.03 ± 0.03 ***

We sought to answer the question of whether errors in estimation were correlated between

the two methods. For each scenario, we plotted the absolute errors for SBP homeostasis baseline, E_{max}, EC_{50}, and E_{max} ratio of one method against the other and calculated the Pearson correlation coefficient for each parameter (Figures 23, 24, 25, and 26 for Clinical Scenarios 1-4, respectively). We observed that most variables had a moderate amount of positive correlation with r around 0.5. The negative or correlations close to zero were seen in EC_{50} estimates, where NLLS had high errors while AIM-BP did not. This degree of correlation suggests that the individual sets of simulated data themselves contribute to the error.

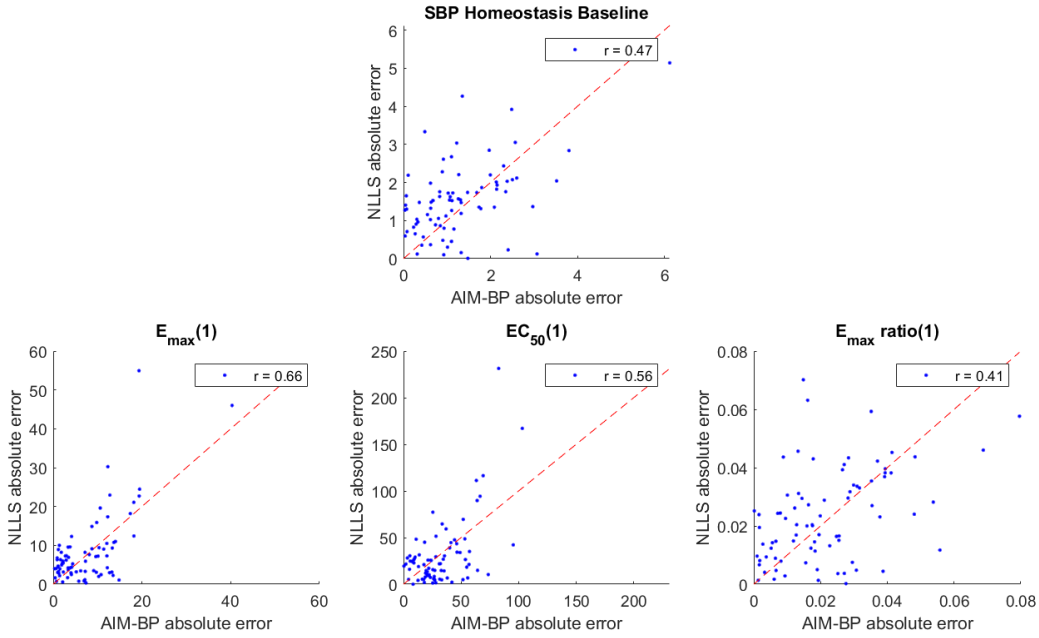

Figure 23: Absolute Error Correlation for Clinical Scenario 1. Absolute errors of SBP homeostasis baseline, E_{max}, EC_{50}, and E_{max} ratio from AIM-BP estimates were plotted against absolute errors from NLLS estimates. A Pearson correlation coefficient was calculated.

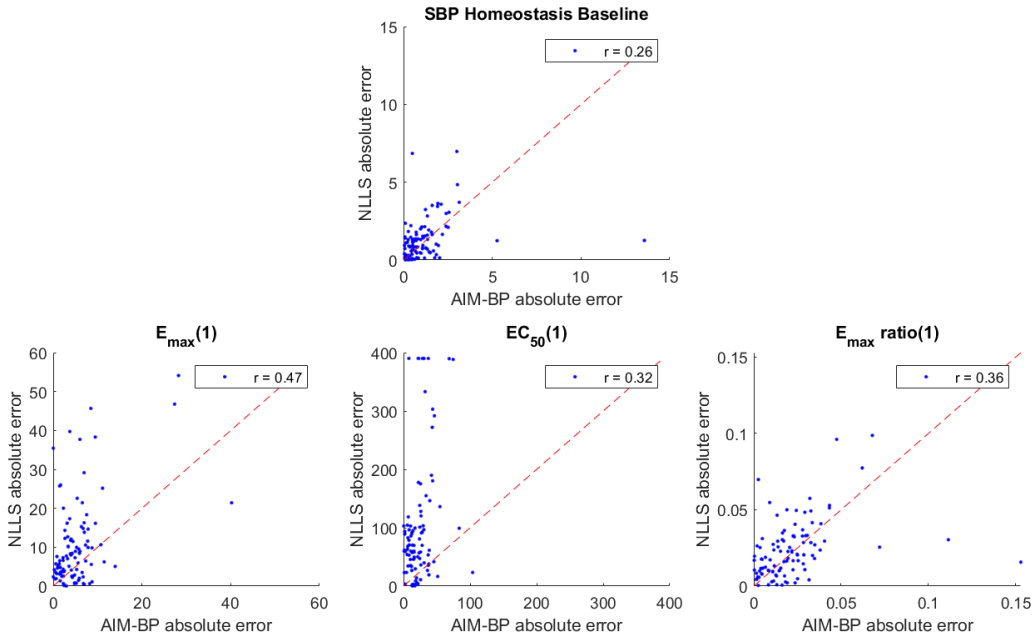

Figure 24: Absolute Error Correlation for Clinical Scenario 2. Absolute errors of SBP homeostasis baseline, E_{max}, EC_{50}, and E_{max} ratio from AIM-BP estimates were plotted against absolute errors from NLLS estimates. A Pearson correlation coefficient was calculated.

4.3 DISCUSSION

The clinical scenarios were designed to test the AIM-BP parameter estimator on distinguishing drug effectiveness from spontaneous blood pressure trends in similar scenarios. Scenarios 1 and 2 shared the same configuration with one drug but different drug effectiveness and homeostasis baselines. Similarly, scenarios 3 and 4 shared the same configuration with two drugs and different drug effectiveness and homeostasis baselines. We observed that AIM-BP consistently performed significantly better than NLLS in the more complex scenarios in estimating systolic BP homeostasis baseline and separate drug E_{max} and EC_{50} values. However, the E_{max} ratio was more closely matched between the two methods. NLLS had more notable

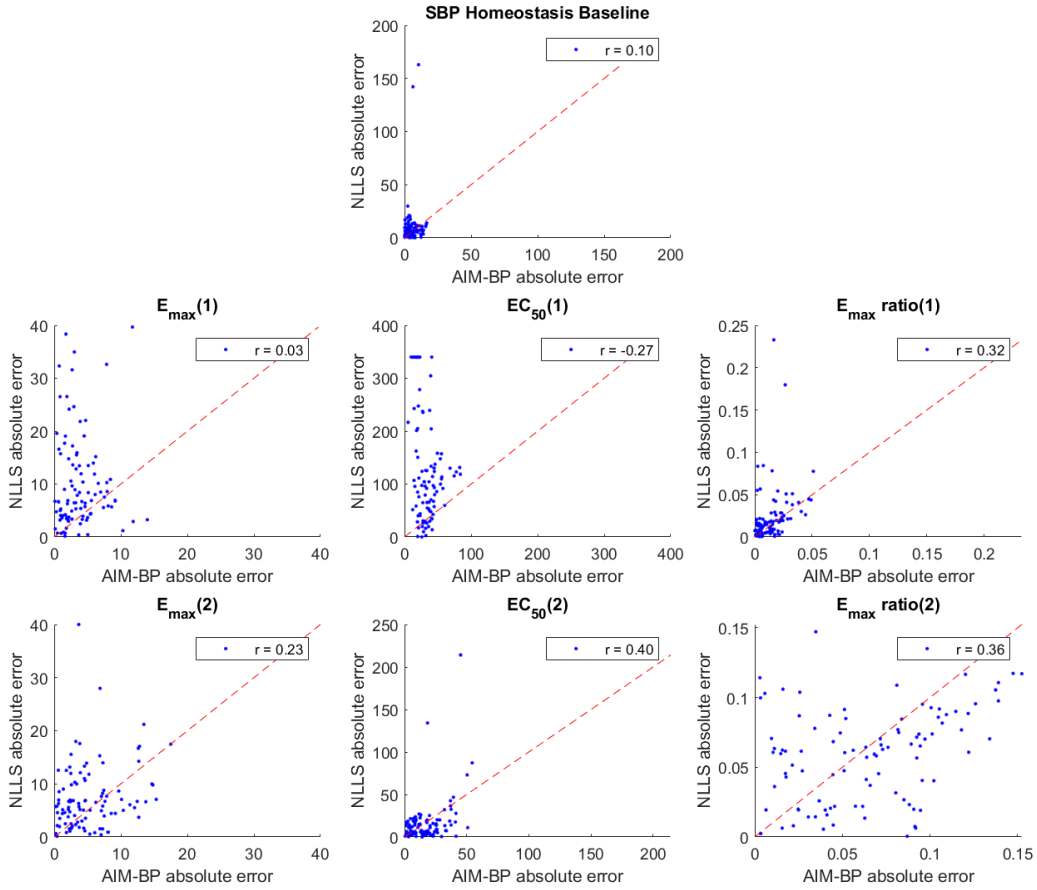

Figure 25: Absolute Error Correlation for Clinical Scenario 3. Absolute errors of SBP homeostasis baseline, E_{max}, EC_{50}, and E_{max} ratio from AIM-BP estimates were plotted against absolute errors from NLLS estimates. A Pearson correlation coefficient was calculated.

outliers in its parameter estimates of E_{max} and EC_{50} compared to AIM-BP. **These results demonstrate claim 2 of this thesis.**

While both methods enjoyed pre-set ranges (minimum and maximum) for E_{max} and EC_{50}, AIM-BP theoretically also enjoyed the advantage of a normal prior for E_{max} and EC_{50} values (Table 4). While the means are roughly estimated from literature,[80–82,90] the

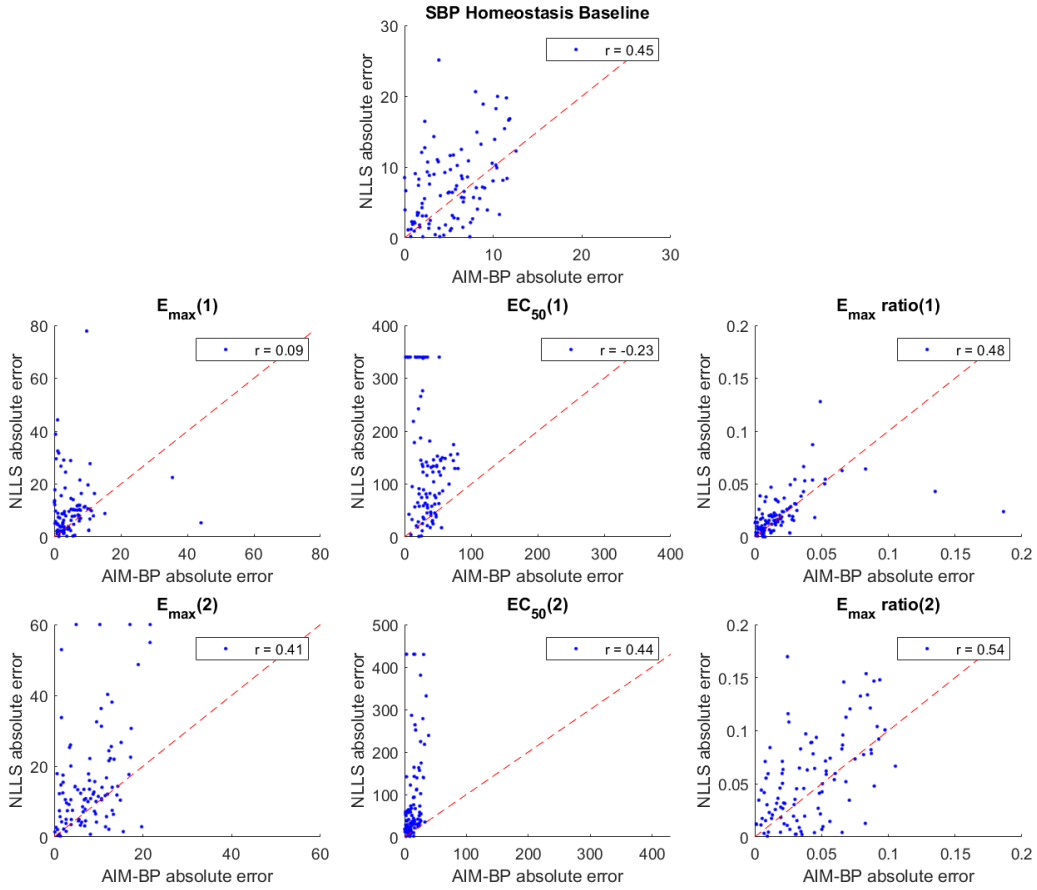

Figure 26: Absolute Error Correlation for Clinical Scenario 4. Absolute errors of SBP homeostasis baseline, E_{max}, EC_{50}, and E_{max} ratio from AIM-BP estimates were plotted against absolute errors from NLLS estimates. A Pearson correlation coefficient was calculated.

variances are set to be greater than those estimated from the literature so as to not overly constrain the MCMC sampling. We observed, however, that the normal prior imposed upon these variables contributed relatively little to the total log probabilities. Because each time point contributes a probability to the log probability sum, the relative contribution of the prior decreases with the more time points available. In a way, this is by design - the more

time points we have to sample from, the less we should care about the prior for the parameter in question. However, the priors for E_{max} and EC_{50} may contribute just enough to prevent situations in which very high E_{max} and EC_{50} produce very similar blood pressure trends to lower E_{max} and EC_{50}.

The AIM-BP parameter estimator, by virtue of being a sampling method, naturally characterizes the uncertainty about parameter estimates in the distribution of sampled parameters (examples can be seen in the histograms in Figures 14, 15, and 16). The histograms for E_{max} and EC_{50} reflect the uncertainty about their values, while the histograms for systolic BP homeostasis baseline are tightly clustered within a few mmHg of the true value. This measure of uncertainty is not available with maximum likelihood methods such as non-linear least squares.

5.0 MISSING AND MULTI-MODAL DATA, EXAMPLE WITH REAL WORLD DATA

Real world ICU blood pressure data contains both missing values and multi-modal measurements, two elements we have not yet investigated in our AIM-BP model evaluation. In this chapter, we investigate the effects of introducing missing values and multi-modal measurements on the performance of the AIM-BP parameter estimator. Finally, we do a qualitative analysis of AIM-BP models learned from a subset of the UPMC hemorrhagic stroke data.

5.1 MISSING VALUES

After discretizing raw blood pressure chart data into a uniform discrete time series, we are presented with the problem of what to fill in for the time points in which no measurements have been made. The two main approaches to handle this problem is data imputation or treating those time points as missing values. Imputation methods range from simple methods like mean imputation and last value carried forward to machine learning based approaches.[91,92] In the context of time series, imputation methods like mean and last value carried forward can be considered forms of interpolation if done on an individual time series basis. Simple interpolation methods can be effective when gaps of missing values are small, but when gaps are large their tendency to underestimate variance can adversely affect downstream analyses.

To avoid this issue, more sophisticated interpolation methods construct a stochastic model, such as an ARMA model, and use the model to fill in the blanks. The AIM-BP model is, in fact, such a model, and learning the parameters via the AIM-BP parameter estimator

can be considered the first step in interpolation. Indeed, once a full set of parameters have been estimated, the AIM-BP data generator can be used to fill in missing vital sign data for other analyses if necessary. Since we are already performing missing value imputation using the AIM-BP framework, it makes little sense to separately interpolate missing values beforehand.

The non-linear least squares curve fitting method is also inherently capable of handling missing time series data, as it only needs to fit a curve to the points where observations have been made. We are thus interested in how the AIM-BP parameter estimator performs against NLLS when missing values are present.

5.2 MULTI-MODAL DATA

In the ICU setting, blood pressure is commonly measured at frequent intervals via automated non-invasive cuffs or arterial lines. While arterial lines are considered to be the gold standard in blood pressure measurement, they are not always placed in all patients due to risk of bleeding and infection. As such, any given patient may present with blood pressure recordings from non-invasive cuffs, arterial lines, or both, at varying times. While arterial blood pressures are considered to be the gold standard, artifacts and noise in arterial blood pressure readings do exist and can be significant.[24] Studies have compared the accuracy of blood pressure measured using arterial lines versus non-invasive blood pressure, using both the MIMIC dataset[93] and a private dataset.[94] These studies are generally in agreement that non-invasive blood pressure tends to underestimate systolic BP when it is high and overestimate systolic BP when it is low compared to arterial lines.

The AIM-BP framework integrates both modes of blood pressure measurement in order to better estimate actual blood pressure. In our preliminary analysis, we did not find lower median variance from arterial line measurements compared to non-invasive measurements. We investigated whether integrating both modes of measurement made AIM-BP perform better than NLLS curve fitting.

5.3 METHODS

We perform the same evaluations as for the clinical scenarios in the previous chapter, but on modified data sets that include missing values (we will refer to this as Clinical Scenario 5) and multi-modal data in conjunction with missing values (we will refer to this as Clinical Scenario 6). Both Clinical Scenarios 5 and 6 will be based off of the ground truth parameters used for CLinical Scenario 4.

5.3.1 Missing Values

To understand how to simulate missing values, we first plotted the amount of missing values for combined systolic BP (either non-invasive or arterial BP was present, with the assumption that if systolic BP was missing, so was diastolic BP) at each 15 minute time point over the first 2 days of admission for the UPMC dataset (Figure 27) and the MIMIC dataset (Figure 28). Both datasets start off with no missing values, which quickly tapers off to 70% to 80% missing values by the middle of day 2. The MIMIC dataset exhibits a 4 time point periodicity to the missing values suggestive of regular recording at hour intervals, with different recording times on the hour for each patient. The same pattern is not apparent in the UPMC dataset. When looking at the plot of MIMIC data using 1 hour as the interval (Figure 29), the periodicity disappears and the proportion of missing values is lower, as expected.

Because the UPMC dataset did not exhibit this periodicity, we used it to simulate missing values for the sake of simplicity. We sought to use the percent missing over time plot as a probability mask for simulated data. However, since the percent missing over time plot is calculated over the whole dataset, there are individual time series with much lower rates of missing values and ones with much higher rates. We chose to limit ourselves to simulating a subset of patients with lower rates of missing values, so we calculated a percent missing values over time plot for just the half of the patients with the least proportion of missing values (Figure 30). This percent missing values plot was used as a probability function to mask out simulated data using the same ground truth configuration as Clinical Scenario 4.

We repeat the same methods as in the clinical scenario evaluation, except after simulating

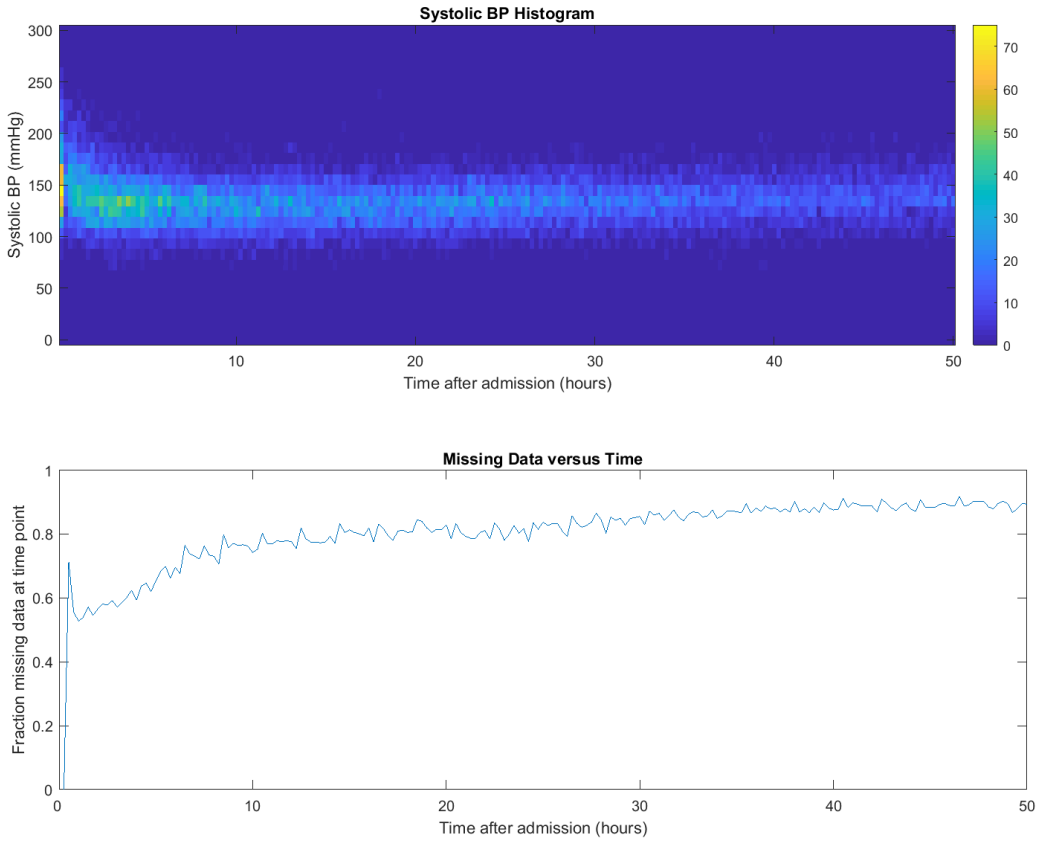

Figure 27: Heat map and missing values over time for UPMC patients, 15 minute interval. Heat map and missing values plotted over 2 days at 15 minute intervals for UPMC patients.

each set of data (out of 100), we remove time points for all observed variables using the probability mask. Both NLLS and AIM-BP were trained on these missing value data sets. We will refer to this set of data as Clinical Scenario 5.

5.3.2 Multi-Modal Data

To simulate multi-modal data, we used the y variable structure in Equation 3.6.2.1 instead of the limited y variable structure in Equation 3.6.2.1. In addition, the emission matrix A

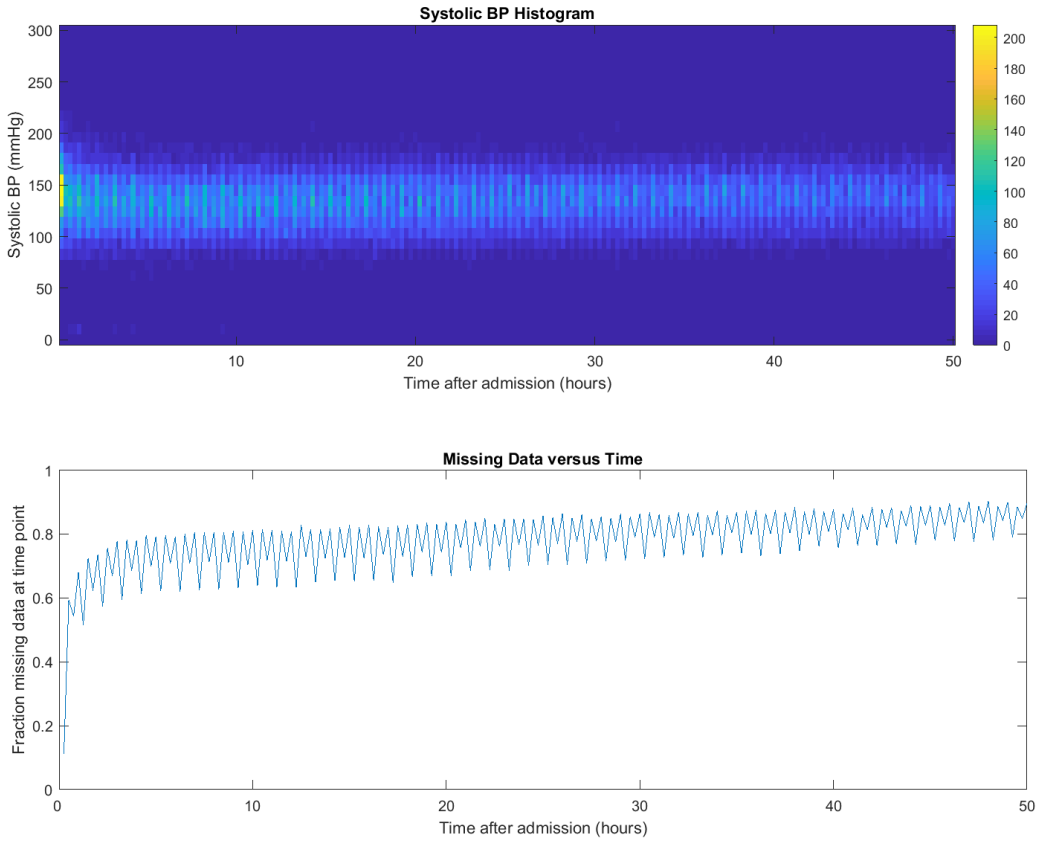

Figure 28: Heat map and missing values over time for MIMIC patients, 15 minute interval. Heat map and missing values plotted over 2 days at 15 minute intervals for MIMIC patients.

was modified to emit both measurements. Instead of 3.6.2.4, A takes on the form:

$$A = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 5 \\ 0 & 1 & 0 & 0 & 0 & 0 & 5 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}$$

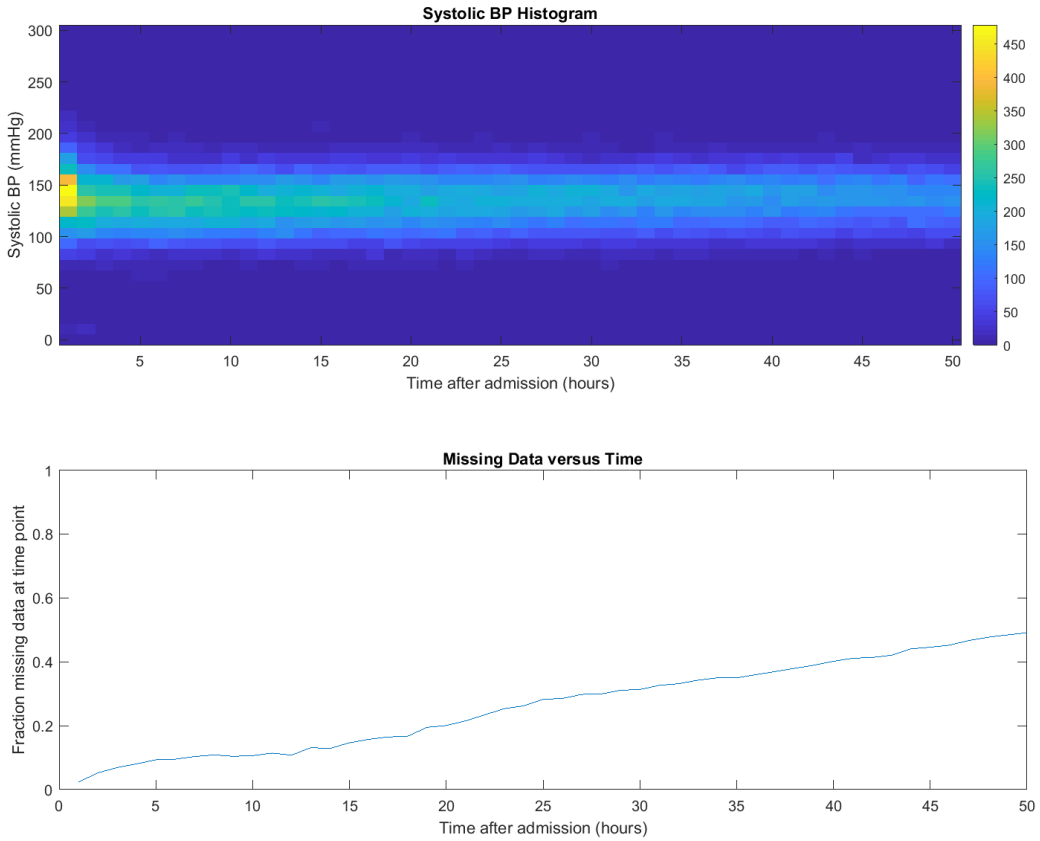

Figure 29: Heat map and missing values over time for MIMIC patients, 1 hour interval. Heat map and missing values plotted over 2 days at 1 hour intervals for MIMIC patients.

Where we add on a bias of 5 mmHg in the last column to emitted non-invasive blood pressures to simulate a biased blood pressure cuff. If desired, the first two columns of A could be further modified to be non-unitary to reflect the non-constant bias seen in Lehman et al. and Ribezzo et al.[93,94] For the purposes of this evaluation, we treated elements of A as constant and known for both the AIM-BP parameter estimator and non-linear least squares curve fitting.

Finally, R was extended to a 5 by 5 covariance matrix, with a small measurement variance

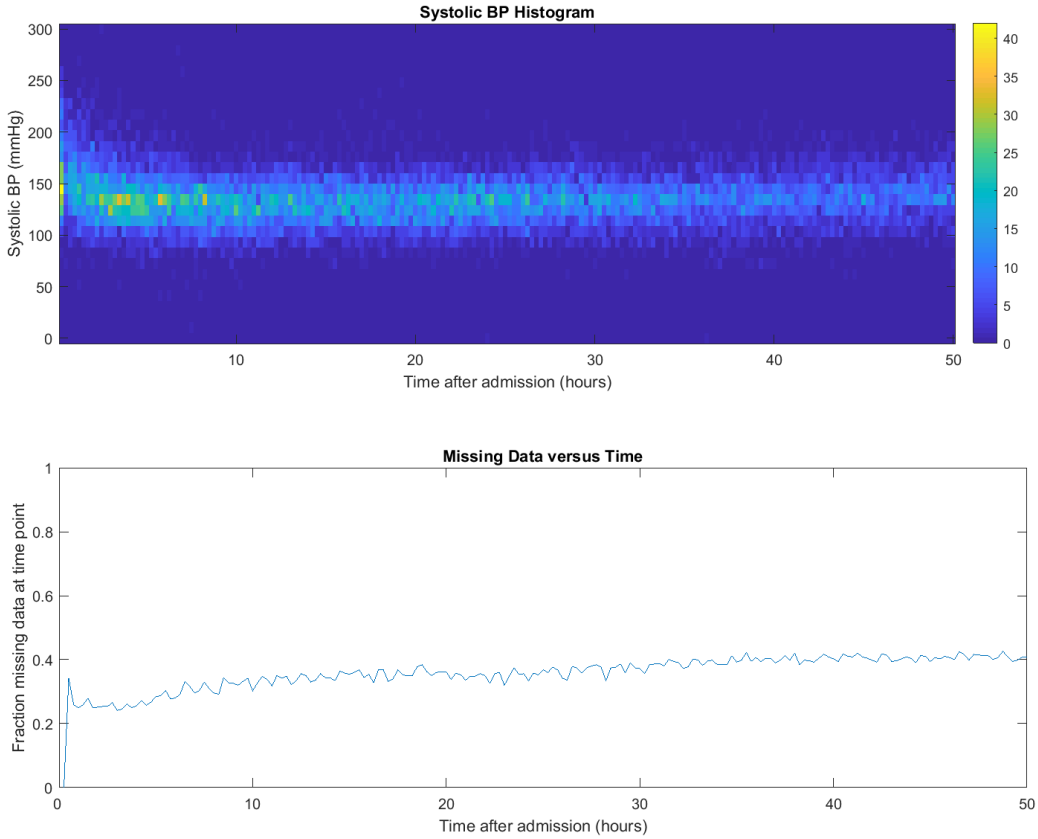

Figure 30: Heat map and missing values over time for 50% of UPMC patients with the least missing values. Heat map and missing values plotted over 2 days at 15 minute intervals for 50% UPMC patients with the least amount of missing values.

of 1 $mmHg^2$ for arterial systolic and diastolic blood pressure. The other parameters remain configured in the same setting as Clinical Scenario 4.

After each data set is created using the new model, missing values are simulated using the probability mask described in the previous section. Further missing values are added by removing the first 50 time points as well as a random 60% of the rest of non-invasive BP, to simulate a patient who had an arterial line initially.

For non-linear least squares curve fitting, the perturbation model in 3.1 was expanded

so that there was a seperate $SBP(t)$ and $DBP(t)$ for non-invasive BP and arterial BP. The known 5 mmHg bias for non-invasive BP was introduced in the corresponding equations.

NLLS and the AIM-BP parameter estimator were then trained on these multi-modal missing value data sets. We will refer to this set of data as Clinical Scenario 6.

5.3.3 AIM-BP Fit to UPMC Data

Finally, as a qualitative experiment, we ran the AIM-BP parameter estimator on a subset of UPMC hemorrhagic stroke data. Because only the drug characteristics for nicardipine and labetalol were researched for the clinical scenarios presented in the previous chapter, information on volume of distribution, time to peak onset, half-life, and average E_{max} model parameters were not available for other commonly used antihypertensives or pressors. As such, we limited the dataset to only patients who were administered either nicardipine or labetalol with no other antihypertensives during their admission. In addition, we also limited the dataset to patients with more than two days worth of blood pressure data and at least 50 blood pressure measurements, whether arterial or non-invasive. After these two filters, 7 total patients out of a total of 497 fit the criteria.

The blood pressure data for these patients were discretized as described previously in Section 3.3. Medication administrations of nicardipine and labetalol, both oral and IV, were processed from recordings of medication administrations. A rough calculation of drug plasma concentration over time based on dosage was made, using values for central volume of distribution from reference sources (https://pubchem.ncbi.nlm.nih.gov/) and estimated weight from BMI to calculate peak plasma concentration (Equation 5.1), a linear increase for the duration of peak concentration onset, and an exponential decay after the peak concentration.

$$c_{peak} = \frac{d}{V_c} \tag{5.1}$$

Where c_{peak} is the peak concentration, d is the dose, and V_c is the central compartment volume of distribution.

The AIM-BP parameter estimator was used to learn an AIM-BP model for these 7 patients, using the same prior settings and sampling parameters as the clinical scenarios presented above, with the exception of Q. Whereas previously Q was treated as constant and known, with this fit we allowed Q to be estimated in a Gibbs step with an inverse Wishart prior with degrees of freedom $\nu = 10$ and scale matrix T equal to the previously constant and known matrix times ν. The details of the Gibbs step are described in Appendix A. Once sampled, the parameters of the model for each patient are used to calculate an estimate of the perturbed homeostasis systolic BP level over time by summing up the spontaneous stroke perturbation and drug effects in the homeostasis perturbation model. This estimated AIM-BP homeostasis level over time is plotted against actual measured systolic BP and calculated drug plasma concentrations and evaluated subjectively for goodness of fit.

5.4 RESULTS

Performance differences between NLLS and AIM-BP on Scenarios 5 (Figure 31) and 6 (Figure 32) were very similar to those from Scenarios 3 and 4. AIM-BP was more consistent with the SBP homeostasis baseline, E_{max}, and EC_{50} parameters in both Scenarios 5 and 6. Both E_{max} ratio parameters were significantly more consistent in Scenario 5 but only the nicardipine E_{max} ratio was more consistent in Scenario 6. From visual inspection, the differences in the spread of E_{max} ratio estimates are smaller in magnitude than those of the other parameters.

The mean absolute errors of estimated parameters for both NLLS and AIM-BP in the missing value and multi-modal missing value settings were negligibly different from those of Scenario 4, and the same pattern of more accurate estimates using the AIM-BP parameter estimator found in Scenario 4 apply here.

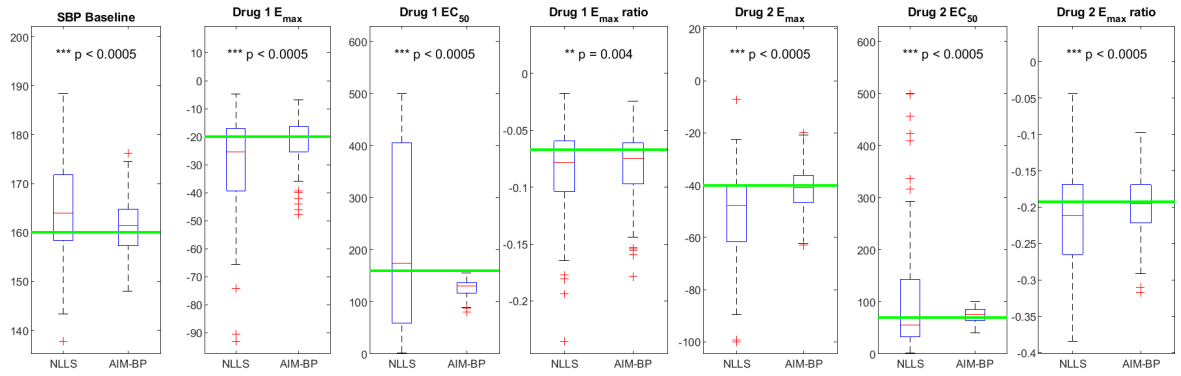

Figure 31: Clinical Scenario 5. Box and whisker plots of estimated parameters using non-linear least squares (NLLS) and AIM-BP. Significantly lower variances of AIM-BP estimated parameters compared to NLLS are noted with *'s and p-values.

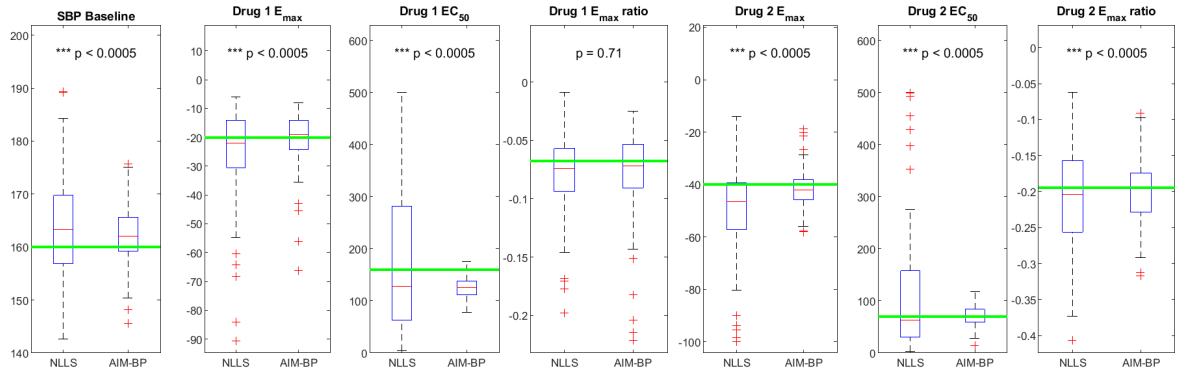

Figure 32: Clinical Scenario 6. Box and whisker plots of estimated parameters using non-linear least squares (NLLS) and AIM-BP. Significantly lower variances of AIM-BP estimated parameters compared to NLLS are noted with *'s and p-values.

Table 7: Scenario 5,6: Mean and standard deviation of absolute error of estimated parameters from NLLS and AIM-BP compared to ground truth. Significantly lower AIM-BP absolute errors are noted with *'s (* = $p < 0.0005$, ** = $p < 0.005$, * = $p < 0.05$).**

Scenario 5	**Ground Truth**	**NLLS Error**	**AIM-BP Error**
SBP homeostasis baseline ($\mu_0^{(3)}$)	160.0	8.7 ± 7.1	4.8 ± 3.6 ***
Labetalol E_{max}	-20.0	13.8 ± 14.5	5.9 ± 5.8 ***
Labetalol EC_{50}	160.0	154.3 ± 115.1	32.7 ± 15.1 ***
Labetalol E_{max} ratio	-0.07	0.03 ± 0.03	0.02 ± 0.02 ***
Nicardipine E_{max}	-40.0	15.6 ± 15.4	5.9 ± 5.8 ***
Nicardipine EC_{50}	70.0	92.5 ± 126.1	12.0 ± 7.5 ***
Nicardipine E_{max} ratio	-0.19	0.06 ± 0.05	0.03 ± 0.03 ***
Scenario 6	**Ground Truth**	**NLLS Error**	**AIM-BP Error**
SBP homeostasis baseline ($\mu_0^{(3)}$)	160.0	8.2 ± 6.6	4.6 ± 3.8 ***
Labetalol E_{max}	-20.0	11.0 ± 12.0	6.8 ± 7.8 **
Labetalol EC_{50}	160.0	133.8 ± 108.2	35.8 ± 17.5 ***
Labetalol E_{max} ratio	-0.07	0.03 ± 0.02	0.02 ± 0.03
Nicardipine E_{max}	-40.0	14.7 ± 16.5	6.8 ± 7.8 ***
Nicardipine EC_{50}	70.0	104.5 ± 139.1	14.5 ± 10.4 ***
Nicardipine E_{max} ratio	-0.19	0.06 ± 0.05	0.03 ± 0.03 ***

5.4.1 Analysis of AIM-BP Fit to UPMC Data

The plots of the learned AIM-BP systolic BP homeostasis perturbation versus measured variables can be seen in Figures 33-39. Qualitatively, the AIM-BP SBP homeostasis perturbation follows closely to measured systolic BPs. The difference between the measured SBP values and the AIM-BP SBP homeostasis perturbation falls in what looks to be a normal distribution ($N(-1.0, 17.0)$) with some small outliers (Figure 40). The dips in blood pressure homeostasis occur at the same times as increases in drug plasma concentration, which roughly correspond to dips in actual measured blood pressure. The confounding effects of medication and the spontaneous decrease in blood pressure can be seen in patients 2 and 6 (Figures 34 and 38), where medications are administered shortly after admission to the ICU.

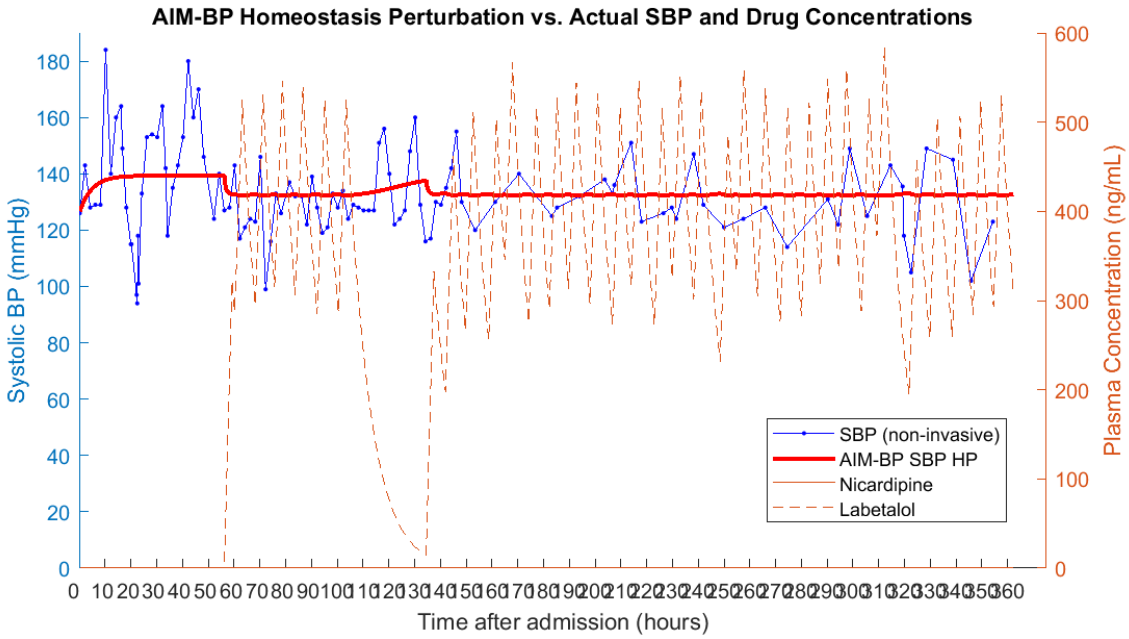

Figure 33: AIM-BP SBP Homeostasis Perturbation vs Measurements, Patient 1. A graph of the calculated AIM-BP systolic BP homeostasis perturbation versus measured SBP and calculated drug plasma concentrations.

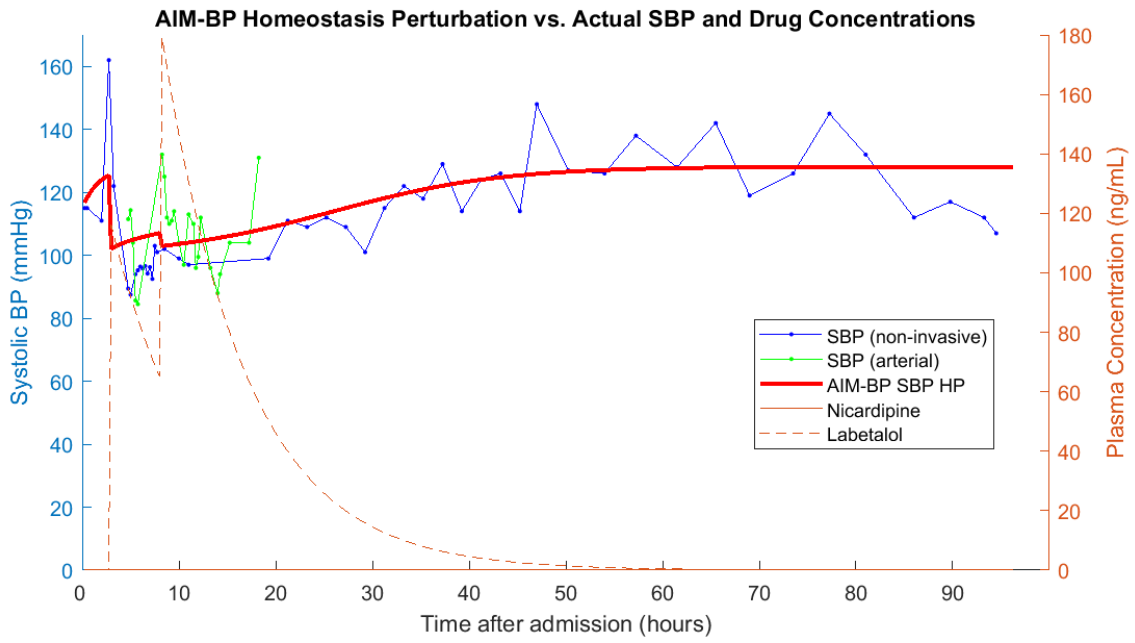

Figure 34: AIM-BP SBP Homeostasis Perturbation vs Measurements, Patient 2. A graph of the calculated AIM-BP systolic BP homeostasis perturbation versus measured SBP and calculated drug plasma concentrations.

5.5 DISCUSSION

The missing value and multi-modal simulation results suggest that the level of missing values we simulated was not sufficient to adversely affect performance of AIM-BP or NLLS fitting. The addition of multi-modal measurements did not increase performance either. By simulating roughly 30-40% missing data, we were too conservative in our guess of how much missing data would impact performance. Without a dramatic change due to missing data, the addition of multi-modal data does not help estimates either. Because the chief source of variance in the tested model is due to inherent blood pressure variability Q instead of measurement noise R, having multiple measurements does not aid significantly in reducing the principle source of variance. Without this benefit of simultaneous measurements, the only remaining area in which multi-modal measurements can help is at time points where

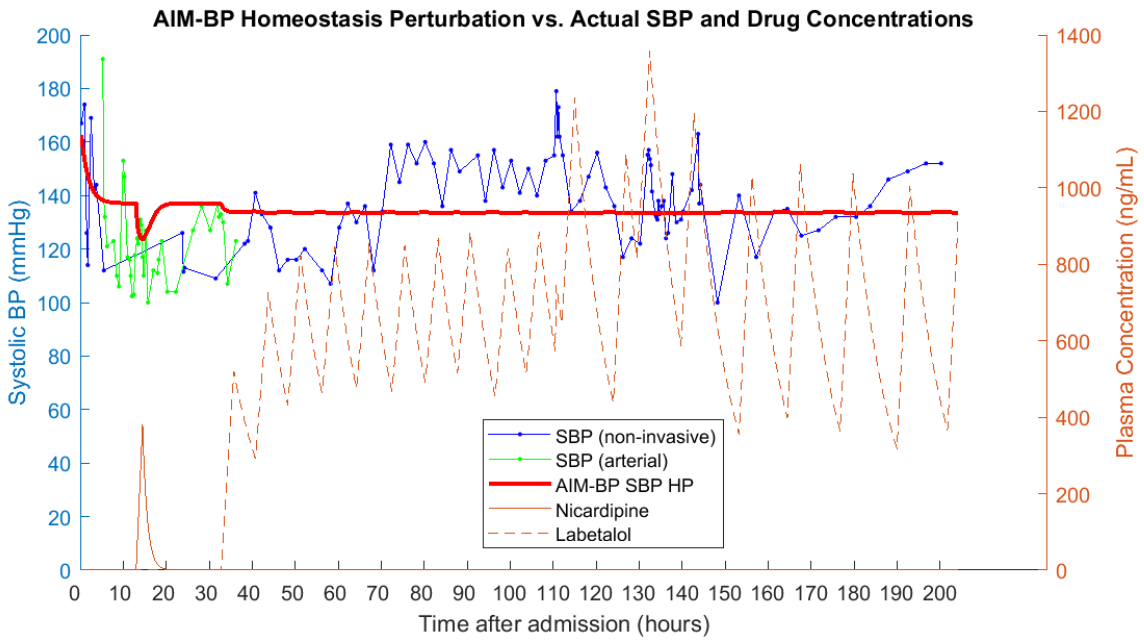

Figure 35: AIM-BP SBP Homeostasis Perturbation vs Measurements, Patient 3. A graph of the calculated AIM-BP systolic BP homeostasis perturbation versus measured SBP and calculated drug plasma concentrations.

one mode captures a measurement while the other has missing values. If these missing values did not significantly impact performance, however, the benefit of multi-modal measurements will also be minimal.

While we did not model it in this thesis, the emission matrix A can be modified to reflect the linear bias between arterial and non-invasive pressures. While this bias likely differs on an admission to admission basis, we do not know if it differs enough to warrant learning these parameter values of A per patient versus setting them as known based on a population average. Because simultaneous arterial and non-invasive pressures can be relatively rare in charted data, there may simply not be enough data available to learn A on a patient to patient basis.

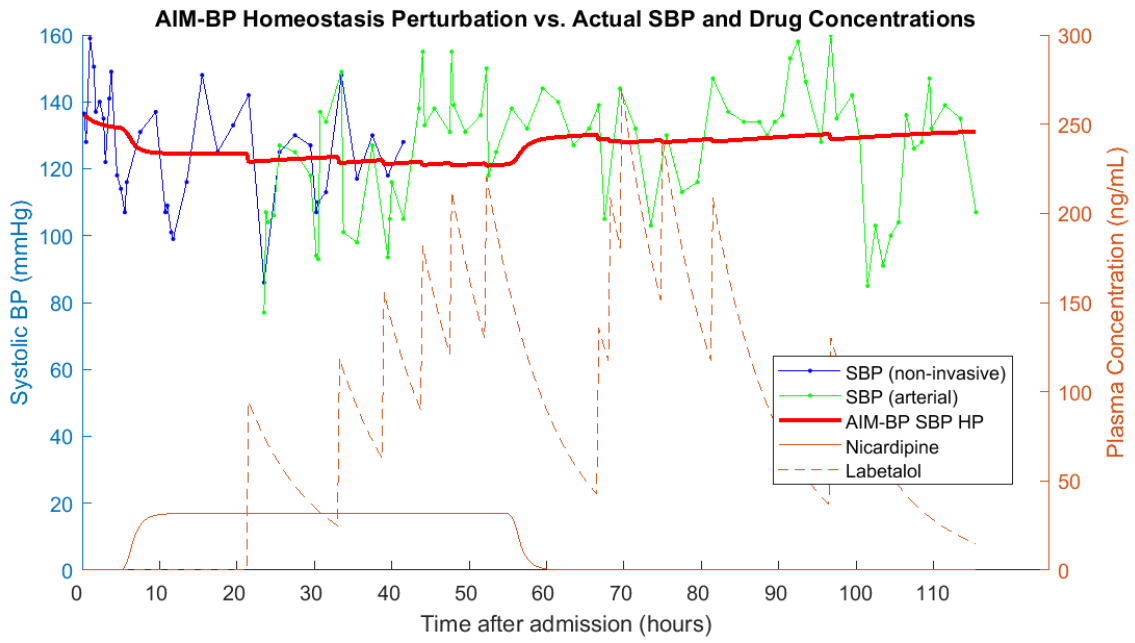

Figure 36: AIM-BP SBP Homeostasis Perturbation vs Measurements, Patient 4. A graph of the calculated AIM-BP systolic BP homeostasis perturbation versus measured SBP and calculated drug plasma concentrations.

5.5.1 Subjective Analysis of AIM-BP Fit to UPMC Data

A visual analysis of the AIM-BP systolic BP homeostasis perturbation suggests that the AIM-BP framework is sufficient to describe the effects of antihypertensive drugs and spontaneous blood pressure perturbation as a result of stroke. The residuals from measured systolic BP minus the homeostasis perturbation BP seem to be roughly normal with a variance similar in magnitude to that of the covariance matrix Q, though the exact comparison is difficult because Q also contains covariances between SBP and DBP.

Unfortunately, due to the small number of patients that fit our selection criteria, we could not expand this process to start learning relationships between AIM-BP parameters and clinical outcomes. Future work will include expanding the available drug information to include patients on other antihypertensives.

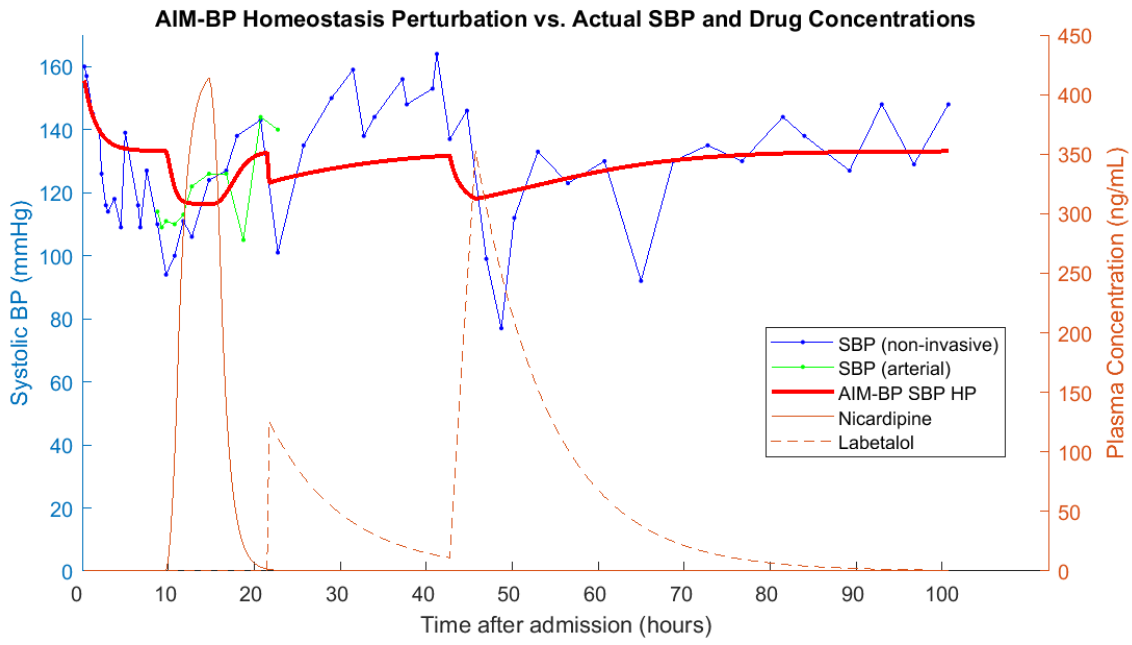

Figure 37: AIM-BP SBP Homeostasis Perturbation vs Measurements, Patient 5. A graph of the calculated AIM-BP systolic BP homeostasis perturbation versus measured SBP and calculated drug plasma concentrations.

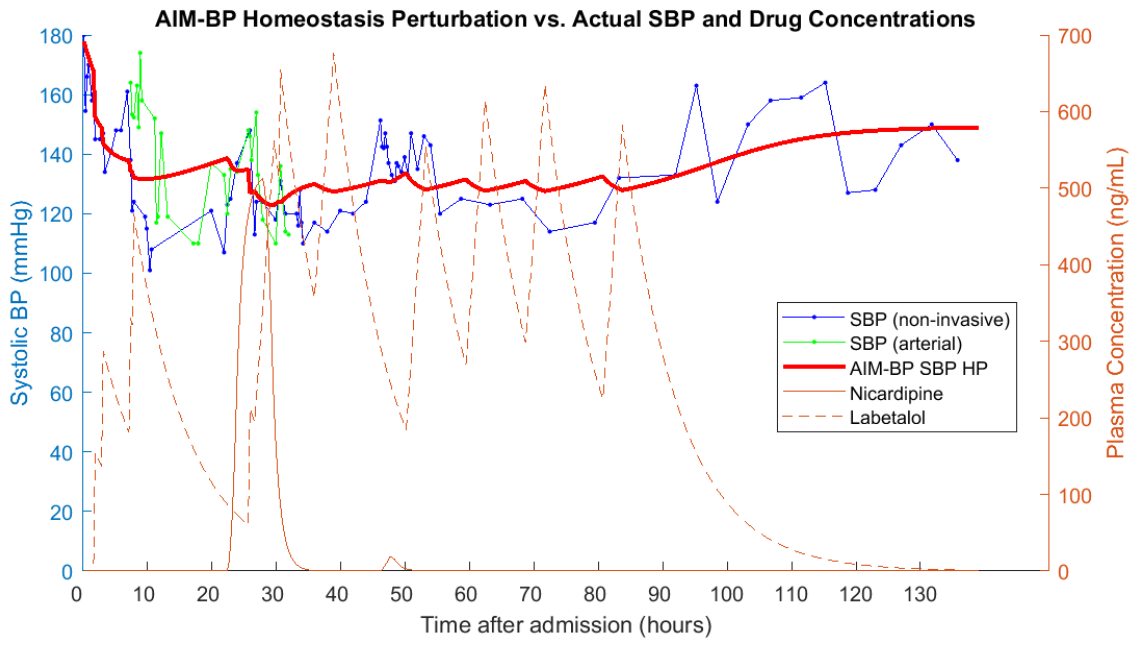

Figure 38: AIM-BP SBP Homeostasis Perturbation vs Measurements, Patient 6. A graph of the calculated AIM-BP systolic BP homeostasis perturbation versus measured SBP and calculated drug plasma concentrations.

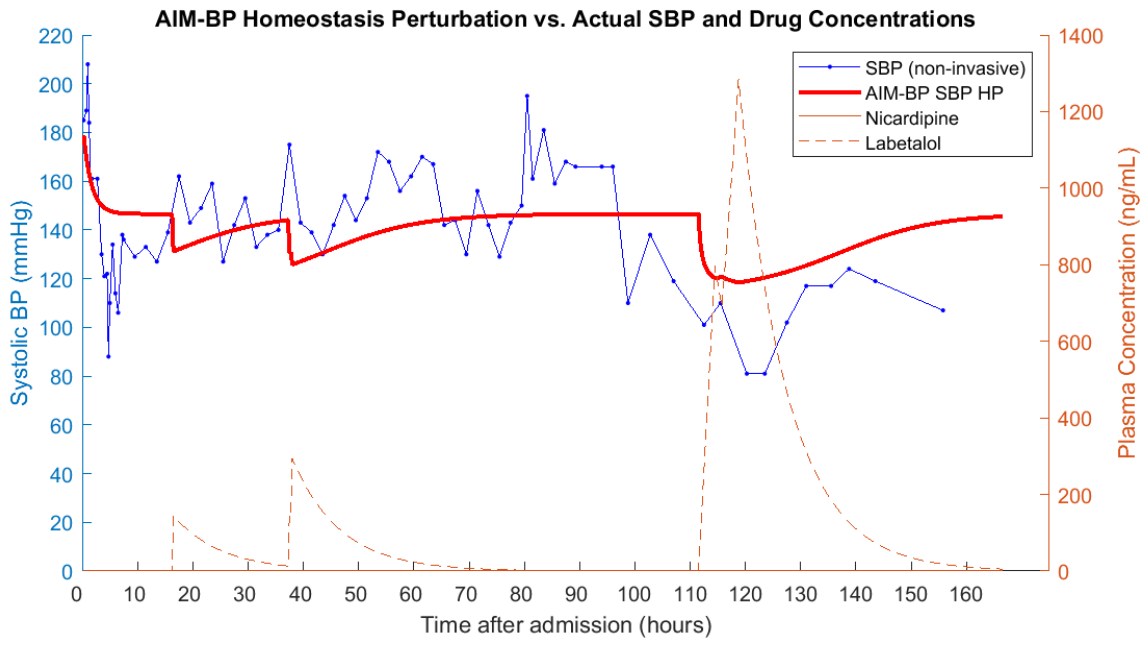

Figure 39: AIM-BP SBP Homeostasis Perturbation vs Measurements, Patient 7. A graph of the calculated AIM-BP systolic BP homeostasis perturbation versus measured SBP and calculated drug plasma concentrations.

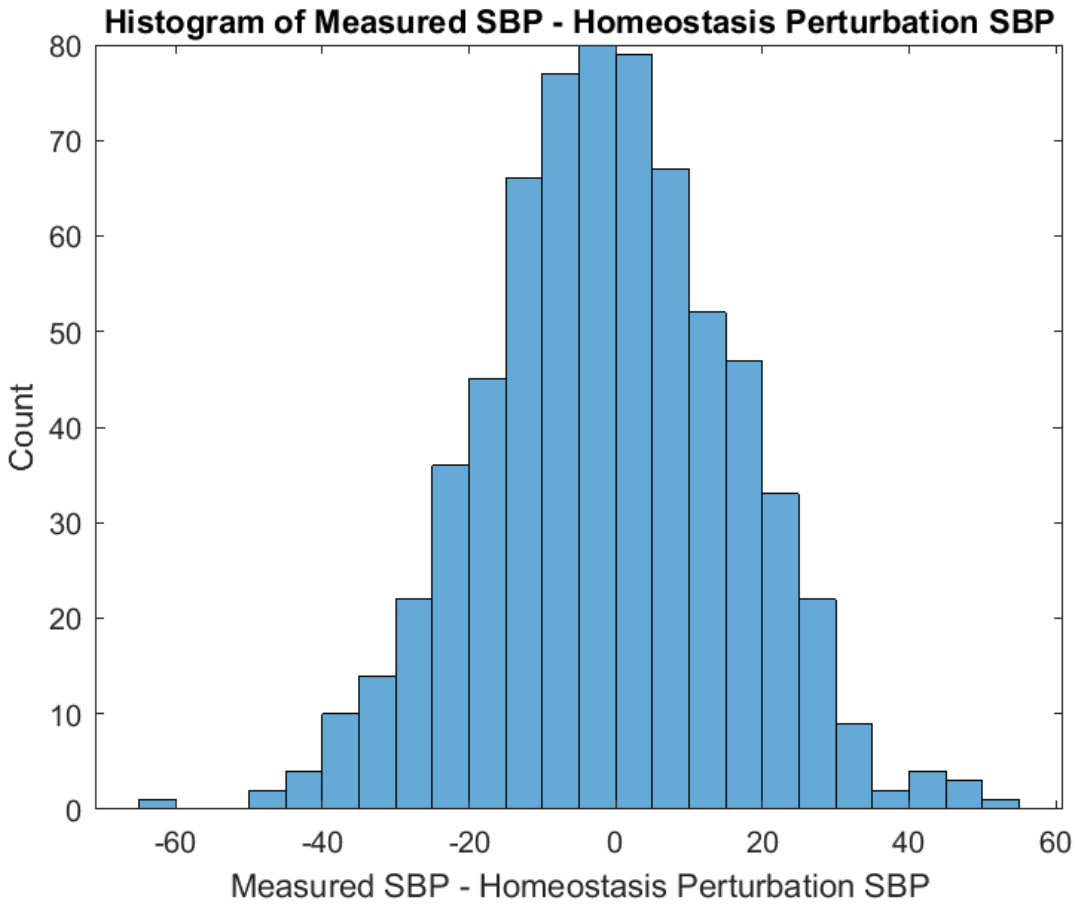

Figure 40: Histogram of Measurement vs AIM-BP SBP Differences. The difference between measured SBPs and the SBPs from the AIM-BP homeostasis perturbation was calculated for each measurement across all 7 patients. The differences follow a roughly normal distribution with small outliers.

6.0 CONCLUSION

We developed a framework for modeling blood pressure management for stroke patients in the ICU, with the goal of enabling research on personalized management for this population. This framework, the Acute Intervention Model of Blood Pressure (AIM-BP), is a flexible and extensible dynamic linear model of vital signs and medications obtained from chart data in the ICU. The parameters of this model include a patient's natural blood pressure homeostasis level, the effect of medications, and the magnitude of spontaneous perturbations to the patient's blood pressure as a result of disease processes. Our framework is both flexible in the kinds of processes it can currently model and easily extensible for more complex modeling scenarios. With AIM-BP, we hope to enable researchers to answer questions such as: Which antihypertensives work better for which people in the treatment of stroke? Can we provide clinicians with quantitative information about the effectiveness of drugs they chose for therapy? Do aspects of a patient's spontaneous blood pressure behavior affect their outcomes?

We designed a homeostasis perturbation model to capture the simultaneous effects of a patient's spontaneous blood pressure behavior and input from medical management. We described the AIM-BP dynamic linear model and the Markov Chain Monte Carlo parameter estimator. Using chart data from the ICU at UPMC Presbyterian as well as the public MIMIC III database, we characterized the blood pressure trends of intracerebral hemorrhage patients and used these characteristics to inform our model design.

We evaluated the ability of AIM-BP to identify drug effectiveness and parameters that govern spontaneous blood pressure behavior post stroke. We simulated 4 clinical scenarios with ground truth parameters. We compared AIM-BP performance on these 4 scenarios to a state-of-the-art non-linear least squares curve fitting method for parameter estimation.

Compared to ground truth parameter values, AIM-BP produced equal or smaller mean absolute errors than NLLS. In addition, AIM-BP was more consistent in estimating parameter values, especially in the more complex scenarios that involve multiple drugs. We investigated whether errors from both methods were correlated, and found a moderate amount of correlation. Due to the stochastic nature of the generative mechanism, it sometimes simulates a dataset that is unlikely given the parameters used to simulate it. When NLLS and AIM-BP both maximize a likelihood or a posteriori probability on the unlikely data, they then learn estimates that are different from the generative ground truth parameters.

Finally, we evaluated AIM-BP on data containing missing values and multi-modal measurements. We found that AIM-BP performed no worse with the inclusion of missing values and no better with the inclusion of multi-modal estimates, but neither did non-linear least squares estimation. As we chose to simulate an amount of missing values resembling the half of patients with the highest proportion of observed data, we can increase the amount of missing values and see at which point it starts making a difference.

6.1 LIMITATIONS

In this thesis, we chose to treat many of the AIM-BP parameters as constant and known for the purposes of clarity. In reality, we do not expect many of these parameters to be constant for a population, and rather vary on an individual to individual basis. Blood pressure variability Q, measurement noise R, and emission matrix A are both likely to vary with the patient and the quality of the sensor contacts, for example. In addition, we held constant the homeostasis rate r_h in the transition matrix Φ to a large value that effectively translated to independent zero mean noisy observations at each time point. In reality, certain homeostasis mechanisms may take longer to manifest, and we may need to learn a value for r_h for each patient. With every increased parameter that we need to estimate, we expand the search space. This means we will need to increase the number of MCMC steps to achieve convergence, thereby increasing run time.

While we provided the mechanisms to model multivariate interactions between heart rate

and blood pressure, we did not explore this aspect of the model in detail in this thesis. As the UPMC dataset not containing heart rate information, the model must be able to work sufficiently absent these interactions. However, fleshing out these interactions in detail is part of future work.

A major limiting factor in our model design as it stands is the translation of medication dosages into estimates of active concentrations at each time point. While this can theoretically be calculated from information on the pharmacokinetics of the drug available from the manufacturer, in reality differences in individual metabolic and clearance rates mean the active drug concentrations at a given time point will differ on a patient to patient basis even if the dosages given were the same. In theory, these variations are yet another parameter that may be learned using our MCMC-based parameter estimator framework. However, issues with the accuracy of timestamps for medication administration in the ICU may limit the usefulness of a more detailed modeling system. These issues can be circumvented in a research study setting if drug plasma concentration is measured regularly.

6.2 CONTRIBUTIONS

The AIM-BP framework is a novel modeling system for blood pressure control in stroke, built from physiological first principles. It is the first of its kind to characterize blood pressure in the post-stroke setting in a quantitative, human-interpretable, integrative, and extensible way. We introduce the combination of a disease-specific homeostasis perturbation model and a disease-agnostic AIM-BP temporal model that allows a clinician to study stroke in the context of differences in autonomic control of blood pressure resulting from different underlying injury mechanisms defined under the stroke umbrella.

This thesis contributes to the field of biomedical informatics by laying the groundwork for precision measurement of drug effects and individual blood pressure trends in the ICU. We introduce the notion that a quantitative pharmacodynamics model can be fit in a non-research clinical data environment and demonstrate its feasibility. Whereas previously physicians relied on clinical judgment to assess the effectiveness of antihypertensive therapy, AIM-BP

provides a quantitative metric for drug effectiveness that is not confounded by a patient's own, natural blood pressure time course. Conversely, the ability to characterize a patient's natural blood pressure time course independent of medical interventions may yield insights into which patients benefit from antihypertensive therapy and which patients do not.

6.3 FUTURE WORK

In this thesis, we have laid the groundwork and provided a proof-of-concept evaluation of the AIM-BP framework. While we have provided some examples of AIM-BP used on real world data, the next step will be to run it on a sufficient quantity of real world data to enable clinical research. We plan on expanding the list of modeled drugs and then running the AIM-BP parameter estimator on our UPMC intracerebral hemorrhage dataset to study individualized effects of antihypertensives.

We have designed the AIM-BP framework to be easily adaptable into the study of other disease processes in the critical care setting. Elements of the homeostasis perturbation model can be easily added or subtracted. By swapping out the spontaneous exponential decay model of stroke blood pressure perturbation in the homeostasis perturbation model for other effects, the AIM-BP framework can be adapted to study blood pressure management in other disease processes. For example, the AIM-BP framework can be used to more accurately measure kidney perfusion in the setting of acute kidney injury in the ICU, as current practice of using mean arterial pressure may be insufficient to determine kidney perfusion.[95,96] Alternatively, AIM-BP can be used to study altered blood pressure dynamics in the detection of sepsis.[97,98] If desired, additional effects on blood pressure can be added to the model, such as a sinusoidal model of circadian effects on blood pressure.[99]

The parameters we investigated in this thesis are not the only ones that may have clinical significance. One possible clinical parameter of interest is the inherent blood pressure variability of a patient. Short term blood pressure variability has been associated with adverse cardiovascular events, such as stroke.[100,101] Blood pressure variability has been used as a feature in the prediction of sepsis.[97] Instead of looking at blood pressure variability as

an aggregate mean and standard deviation over time, the AIM-BP framework separates out external perturbatory effects and measurement noise and captures blood pressure variability in the covariance matrix Q. In addition, by learning the parameter r_h in Φ that governs rate of return to homeostasis, AIM-BP can characterize the responsiveness of a patient's homeostatic drive, which may be an additional measure of variability. By using the AIM-BP framework to estimate these parameters, we can leverage multi-modal measurements of blood pressure to increase accuracy of variability measures, as well as describe variability in a multivariate sense. AIM-BP would provide a more robust method of investigating differences in blood pressure variablity compared to simple variance calculations.

In the clinical scenarios we presented, we limited the simulated data available to roughly 2 days' worth. This was done as a balance between having enough time to observe sufficient variability in drug concentrations and being short enough to enable mid-admission decision making. In addition to these reasons, limiting the interval of time available to the AIM-BP parameter estimator serves to mitigate the assumption that one set of dynamics governs the entirety of the time interval. In reality, different dynamics likely apply over the course of a patient's admission as he or she goes through a sequence of disease states. For example, a stroke patient may experience hemorrhagic conversion or rehemorrhage, after which his blood pressure dynamics are most likely different. By limiting the time interval which the AIM-BP parameter estimator is learning from, we have the option of applying the parameter estimator multiple times over the course of an ICU admission and learning separate dynamics for each time period. Future work can apply this methodology to investigate the change in blood pressure dynamics over the course of an ICU admission.

The AIM-BP framework relies on vital signs recorded from the electronic medical record chart, typically at frequencies no greater than every 15 minutes. Arterial lines can sample blood pressure at much higher frequencies, though high frequency recordings of arterial line data is typically only done for research purposes (e.g. the MIMIC waveform dataset), as the amount of storage space required to record the data is orders of magnitude higher. In addition, non-invasive continuous blood pressure monitoring methods are starting to appear,[102] providing us with another source of high frequency blood pressure data. The increased sampling frequency can serve to reduce variance on a longer time scale, and can eliminate the

problem of missing values. The increased amount of data opens up additional paths of research. For example, a higher frequency of blood pressure and heart rate sampling allows us to better model the interaction between the two variables, as factors such as the baroreflex that govern the homeostatic relationship between the two operate at the second to second time scale. In addition, we can adapt Lehman et al.'s use of the switching vector autoregressive model[57–59] to introduce multiple dynamic modes in the AIM-BP model, which better captures the expectation that different patient states such as acute disease or recovery may present with different dynamics.

The code for this framework is published in Github (https://github.com/bruinail/aimbp) and freely available under the GNU GPLv3 license.

APPENDIX A

GIBBS SAMPLING DERIVATIONS

A.1 μ_0

$$\begin{aligned} P(\mu_0|\Sigma_0,\Phi,\Upsilon,Q,A,R,u,x,y) &\propto P(\mu_0,\Sigma_0,\Phi,\Upsilon,Q,A,R,u,x,y) \\ \propto P(y|x,\mu_0,\Sigma_0,\Phi,\Upsilon,Q,A,R,u)P(x|\mu_0,\Sigma_0,\Phi,\Upsilon,Q,A,R,u)P(\mu_0,\Sigma_0,\Phi,\Upsilon,Q,A,R,u) \\ &\propto \prod_{t=1}^{n} P(y_t|x_t,A,R,u_t)P(x_0|\mu_0,\Sigma_0)\prod_{t=1}^{n} P(x_t|x_{t-1},\Phi,\Upsilon,Q,u_t)P(\mu_0) \\ &\propto P(x_0|\mu_0,\Sigma_0)P(\mu_0) \end{aligned} \tag{A.1}$$

This is our full conditional posterior for μ_0. Under our dynamic linear model, x_0 is distributed normally with $x_0 \sim N(\mu_0, \Sigma_0)$. Thus, its PDF is:

$$pdf(x_0) = det(2\pi\Sigma_0)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_0-\mu_0)'\Sigma_0^{-1}(x_0-\mu_0)}$$

Thus,

$$\begin{aligned} P(\mu_0|\Sigma_0,\Phi,\Upsilon,Q,A,R,u,x,y) &\propto P(x_0|\mu_0,\Sigma_0)P(\mu_0) \\ &\propto det(2\pi\Sigma_0)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_0-\mu_0)'\Sigma_0^{-1}(x_0-\mu_0)} P(\mu_0) \\ &\propto e^{-\frac{1}{2}(x_0-\mu_0)'\Sigma_0^{-1}(x_0-\mu_0)} P(\mu_0) \end{aligned} \tag{A.2}$$

For the prior $P(\mu_0)$ we choose the multivariate normal, the conjugate prior to a multivariate normal with known variance and unknown mean. Although μ_0 is a variable we are interested in, the multivariate normal prior provides sufficient flexibility for covering our beliefs about initial blood pressure that we do not need to experiment with other priors.

$$\begin{aligned} P(\mu_0) &\sim N(m, S) \\ P(\mu_0) &= det(2\pi S)^{-\frac{1}{2}} e^{-\frac{1}{2}(\mu_0 - m)' S^{-1} (\mu_0 - m)} \\ &\propto e^{-\frac{1}{2}(\mu_0 - m)' S^{-1} (\mu_0 - m)} \end{aligned} \tag{A.3}$$

Multiplying the prior into the full conditional posterior, we get:

$$\begin{aligned} P(\mu_0|\Sigma_0, \Phi, \Upsilon, Q, A, R, u, x, y) &\propto e^{-\frac{1}{2}(x_0 - \mu_0)' \Sigma_0^{-1} (x_t - \mu_0)} P(\mu_0) \\ &\propto e^{-\frac{1}{2}(x_0 - \mu_0)' \Sigma_0^{-1} (x_t - \mu_0)} e^{-\frac{1}{2}(\mu_0 - m)' S^{-1} (\mu_0 - m)} \\ &= e^{-\frac{1}{2}[(x_0 - \mu_0)' \Sigma_0^{-1} (x_0 - \mu_0) + (\mu_0 - m)' S^{-1} (\mu_0 - m)]} \\ &= e^{-\frac{1}{2}[x_0' \Sigma_0^{-1} (x_0 - \mu_0) - \mu_0' \Sigma_0^{-1} (x_0 - \mu_0) + (\mu_0 - m)' S^{-1} (\mu_0 - m)]} \\ &= e^{-\frac{1}{2}[x_0' \Sigma_0^{-1} x_0 - x_0' \Sigma_0^{-1} \mu_0 - \mu_0' \Sigma_0^{-1} x_0 + \mu_0' \Sigma_0^{-1} \mu_0 + (\mu_0 - m)' S^{-1} (\mu_0 - m)]} \\ &= e^{-\frac{1}{2}[x_0' \Sigma_0^{-1} x_0 - (x_0' \Sigma_0^{-1} \mu_0)' - \mu_0' \Sigma_0^{-1} x_0 + \mu_0' \Sigma_0^{-1} \mu_0 + (\mu_0 - m)' S^{-1} (\mu_0 - m)]} \\ &= e^{-\frac{1}{2}[x_0' \Sigma_0^{-1} x_0 - \mu_0' \Sigma_0^{-1'} x_0 - \mu_0' \Sigma_0^{-1} x_0 + \mu_0' \Sigma_0^{-1} \mu_0 + (\mu_0 - m)' S^{-1} (\mu_0 - m)]} \\ &= e^{-\frac{1}{2}[x_0' \Sigma_0^{-1} x_0 - \mu_0' \Sigma_0^{-1} x_0 - \mu_0' \Sigma_0^{-1} x_0 + \mu_0' \Sigma_0^{-1} \mu_0 + (\mu_0 - m)' S^{-1} (\mu_0 - m)]} \\ &= e^{-\frac{1}{2}[x_0' \Sigma_0^{-1} x_0 - 2\mu_0' \Sigma_0^{-1} x_0 + \mu_0' \Sigma_0^{-1} \mu_0 + (\mu_0 - m)' S^{-1} (\mu_0 - m)]} \\ &= e^{-\frac{1}{2}[x_0' \Sigma_0^{-1} x_0 - 2\mu_0' \Sigma_0^{-1} x_0 + \mu_0' \Sigma_0^{-1} \mu_0 + \mu_0' S^{-1} \mu_0 - 2\mu_0' S^{-1} m + m' S^{-1} m]} \\ &= e^{-\frac{1}{2}[(x_0' \Sigma_0^{-1} x_0 + m' S^{-1} m) - 2\mu_0' (\Sigma_0^{-1} x_0 + S^{-1} m) + \mu_0' (\Sigma_0^{-1} + S^{-1}) \mu_0]} \\ &\propto e^{(\mu_0 - m_{new})' {S_{new}}^{-1} (\mu_0 - m_{new})} \end{aligned} \tag{A.4}$$

Where $S_{new} = (\Sigma_0^{-1} + S^{-1})^{-1}$ and $m_{new} = (\Sigma_0^{-1} + S^{-1})^{-1}(\Sigma_0^{-1} x_0 + S^{-1} m)$. Thus, $P(\mu_0|\Sigma_0, \Phi, \Upsilon, Q, A, R, u, x, y) \sim N((\Sigma_0^{-1} + S^{-1})^{-1}(\Sigma_0^{-1} x_0 + S^{-1} m), (\Sigma_0^{-1} + S^{-1})^{-1})$. We can sample this in a Gibbs step.

A.2 Σ_0

$$
\begin{aligned}
P(\Sigma_0|\mu_0,\Phi,\Upsilon,Q,A,R,u,x,y) &\propto P(\Sigma_0,\mu_0,\Phi,\Upsilon,Q,A,R,u,x,y) \\
\propto P(y|x,\mu_0,\Sigma_0,\Phi,\Upsilon,Q,A,R,u)P(x|\mu_0,\Sigma_0,\Phi,\Upsilon,Q,A,R,u)P(\mu_0,\Sigma_0,\Phi,\Upsilon,Q,A,R,u) \\
&\propto \prod_{t=1}^{n} P(y_t|x_t,A,R,u_t)P(x_0|\mu_0,\Sigma_0)\prod_{t=1}^{n} P(x_t|x_{t-1},\Phi,\Upsilon,Q,u_t)P(\Sigma_0) \\
&\propto P(x_0|\mu_0,\Sigma_0)P(\Sigma_0)
\end{aligned}
\tag{A.5}
$$

This is our full conditional posterior for Σ_0. Under our dynamic linear model, x_0 is distributed normally with $x_0 \sim N(\mu_0, \Sigma_0)$. Thus, its PDF is:

$$
pdf(x_0) = det(2\pi\Sigma_0)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_0-\mu_0)'\Sigma_0^{-1}(x_0-\mu_0)}
$$

Thus,

$$
\begin{aligned}
P(\Sigma_0|\mu_0,\Phi,\Upsilon,Q,A,R,u,x,y) &\propto P(x_0|\mu_0,\Sigma_0)P(\Sigma_0) \\
&\propto det(2\pi\Sigma_0)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_0-\mu_0)'\Sigma_0^{-1}(x_0-\mu_0)} P(\Sigma_0) \\
&\propto det(\Sigma_0)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_0-\mu_0)'\Sigma_0^{-1}(x_0-\mu_0)} P(\Sigma_0)
\end{aligned}
\tag{A.6}
$$

For the prior $P(\Sigma_0)$ we choose the inverse Wishart distribution, the conjugate prior to a multivariate normal with unknown variance and known mean. The inverse Wishart is parameterized by a $p \times p$ positive definite scale matrix T and degrees of freedom $\nu > p-1$.

$$
\begin{aligned}
P(\Sigma_0) &\sim W^{-1}(T,\nu) \\
&= \frac{det(T)^{\frac{\nu}{2}}}{2^{\frac{\nu p}{2}}\Gamma_p(\frac{\nu}{2})} det(\Sigma_0)^{-\frac{\nu+p+1}{2}} e^{-\frac{1}{2}tr(T\Sigma_0^{-1})} \\
&\propto det(\Sigma_0)^{-\frac{\nu+p+1}{2}} e^{-\frac{1}{2}tr(T\Sigma_0^{-1})}
\end{aligned}
\tag{A.7}
$$

Where $\Gamma_p()$ is the multivariate gamma function. Multiplying the prior into the full conditional posterior, we get:

$$
\begin{aligned}
P(\Sigma_0|\mu_0,\Phi,\Upsilon,Q,A,R,u,x,y) &\propto det(\Sigma_0)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_0-\mu_0)'\Sigma_0^{-1}(x_0-\mu_0)} P(\Sigma_0) \\
&\propto det(\Sigma_0)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_0-\mu_0)'\Sigma_0^{-1}(x_0-\mu_0)} det(\Sigma_0)^{-\frac{\nu+p+1}{2}} e^{-\frac{1}{2}tr(T\Sigma_0^{-1})} \\
&= det(\Sigma_0)^{-\frac{1}{2}-\frac{\nu+p+1}{2}} e^{-\frac{1}{2}(x_0-\mu_0)'\Sigma_0^{-1}(x_0-\mu_0)-\frac{1}{2}tr(T\Sigma_0^{-1})} \\
&= det(\Sigma_0)^{-\frac{\nu+p+2}{2}} e^{-\frac{1}{2}tr((x_0-\mu_0)'\Sigma_0^{-1}(x_0-\mu_0))-\frac{1}{2}tr(T\Sigma_0^{-1})} \\
&= det(\Sigma_0)^{-\frac{\nu+p+2}{2}} e^{-\frac{1}{2}tr(\Sigma_0^{-1}(x_0-\mu_0)(x_0-\mu_0)')-\frac{1}{2}tr(T\Sigma_0^{-1})} \\
&= det(\Sigma_0)^{-\frac{\nu+p+2}{2}} e^{-\frac{1}{2}tr((x_0-\mu_0)(x_0-\mu_0)'\Sigma_0^{-1})-\frac{1}{2}tr(T\Sigma_0^{-1})} \\
&= det(\Sigma_0)^{-\frac{\nu+p+2}{2}} e^{-\frac{1}{2}tr((T+(x_0-\mu_0)(x_0-\mu_0)')\Sigma_0^{-1})}
\end{aligned}
\tag{A.8}
$$

Thus, $P(\Sigma_0|\mu_0,\Phi,\Upsilon,Q,A,R,u,x,y) \sim W^{-1}(T+(x_0-\mu_0)(x_0-\mu_0)', \nu+1)$. We can sample this in a Gibbs step.

A.3 Φ

$$
\begin{aligned}
P(\Phi|\mu_0,\Sigma_0,\Upsilon,Q,A,R,u,x,y) &\propto P(\Phi,\mu_0,\Sigma_0,\Upsilon,Q,A,R,u,x,y) \\
\propto P(y|x,\mu_0,\Sigma_0,\Phi,\Upsilon,Q,A,R,u) P(x|\mu_0,\Sigma_0,\Phi,\Upsilon,Q,A,R,u) P(\mu_0,\Sigma_0,\Phi,\Upsilon,Q,A,R,u) \\
&\propto \prod_{t=1}^{n} P(y_t|x_t,A,R,u_t) P(x_0|\mu_0,\Sigma_0) P(x_0|\mu_0,\Sigma_0) \prod_{t=1}^{n} P(x_t|x_{t-1},\Phi,\Upsilon,Q,u_t) P(\Phi) \\
&\propto \prod_{t=1}^{n} P(x_t|x_{t-1},\Phi,\Upsilon,Q,u_t) P(\Phi)
\end{aligned}
\tag{A.9}
$$

This is our full conditional posterior for Φ. Under our dynamic linear model, x_t is distributed normally with $x_t \sim N(\Phi x_{t-1} + \Upsilon u_t, Q)$. Thus, its PDF is:

$$
pdf(x_t) = det(2\pi Q)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_t-\Phi x_{t-1}-\Upsilon u_t)'Q^{-1}(x_t-\Phi x_{t-1}-\Upsilon u_t)}
$$

Thus,

$$
\begin{aligned}
P(\Phi|\mu_0, \Sigma_0, \Upsilon, Q, A, R, u, x, y) &\propto \prod_{t=1}^{n} P(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t) P(\Phi) \\
&\propto \prod_{t=1}^{n} det(2\pi Q)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)} P(\Phi) \\
&\propto e^{\sum_{t=1}^{n} -\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)} P(\Phi)
\end{aligned}
\tag{A.10}
$$

Because Φ has a specific structure, we would like to experiment with various priors for $P(\Phi)$ that are not necessarily conjugate to the multivariate normal. Since we won't be able to sample directly from $e^{\sum_{t=1}^{n} -\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)} P(\Phi)$, we will do a Metropolis step inside Gibbs sampling. We have the option of doing regular Metropolis Hastings or using an adaptive method like ARMS.

A.4 Υ

$$
\begin{aligned}
P(\Upsilon|\mu_0, \Sigma_0, \Phi, Q, A, R, u, x, y) &\propto P(\Upsilon, \mu_0, \Sigma_0, \Phi, Q, A, R, u, x, y) \\
\propto P(y|x, \mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) P(x|\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) P(\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) \\
&\propto \prod_{t=1}^{n} P(y_t|x_t, A, R, u_t) P(x_0|\mu_0, \Sigma_0) P(x_0|\mu_0, \Sigma_0) \prod_{t=1}^{n} P(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t) P(\Upsilon) \\
&\propto \prod_{t=1}^{n} P(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t) P(\Upsilon)
\end{aligned}
\tag{A.11}
$$

This is our full conditional posterior for Υ. Under our dynamic linear model, x_t is distributed normally with $x_t \sim N(\Phi x_{t-1} + \Upsilon u_t, Q)$. Thus, its PDF is:

$$
pdf(x_t) = det(2\pi Q)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)}
$$

Thus,

$$
\begin{aligned}
P(\Upsilon|\mu_0, \Sigma_0, \Phi, Q, A, R, u, x, y) &\propto \prod_{t=1}^{n} P(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t) P(\Upsilon) \\
&\propto \prod_{t=1}^{n} det(2\pi Q)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)} P(\Upsilon) \\
&\propto e^{\sum_{t=1}^{n} -\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)} P(\Upsilon)
\end{aligned}
\tag{A.12}
$$

Because Υ has a specific structure, we would like to experiment with various priors for $P(\Upsilon)$ that are not necessarily conjugate to the multivariate normal. Since we won't be able to sample directly from $e^{\sum_{t=1}^{n} -\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)} P(\Upsilon)$, we will do a Metropolis step inside Gibbs sampling like for Φ.

A.5 R_B, E_{MAX}, AND EC_{50}

These parameters all belong under the general variable u. Ordinarily in the case of DLMs, the inputs u are known, but there's no reason why we can't sample from them the same way we sample from every other parameter.

$$
\begin{aligned}
P(u|\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, x, y) &\propto P(u, \mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, x, y) \\
\propto P(y|x, \mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) P(x|\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) P(\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) \\
&\propto \prod_{t=1}^{n} P(y_t|x_t, A, R, u_t) P(x_0|\mu_0, \Sigma_0) P(x_0|\mu_0, \Sigma_0) \prod_{t=1}^{n} P(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t) P(u) \\
&\propto \prod_{t=1}^{n} P(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t) P(u)
\end{aligned}
\tag{A.13}
$$

This is our full conditional posterior for u. Under our dynamic linear model, x_t is distributed normally with $x_t \sim N(\Phi x_{t-1} + \Upsilon u_t, Q)$. Thus, its PDF is:

$$pdf(x_t) = det(2\pi Q)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)}$$

Thus,

$$\begin{aligned} P(u|\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, x, y) &\propto \prod_{t=1}^{n} P(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t) P(u) \\ &\propto \prod_{t=1}^{n} det(2\pi Q)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)} P(u) \\ &\propto e^{\sum_{t=1}^{n} -\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)} P(u) \end{aligned} \tag{A.14}$$

Since we won't be able to sample directly from $e^{\sum_{t=1}^{n} -\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)} P(u)$, we will do a Metropolis step inside Gibbs sampling. We have the option of doing regular Metropolis Hastings or using an adaptive method like ARMS.

A.6 Q

$$\begin{aligned} P(Q|\mu_0, \Sigma_0, \Phi, \Upsilon, A, R, u, x, y) &\propto P(Q, \mu_0, \Sigma_0, \Phi, \Upsilon, A, R, u, x, y) \\ \propto P(y|x, \mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) P(x|\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) P(\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) \\ &\propto \prod_{t=1}^{n} P(y_t|x_t, A, R, u_t) P(x_0|\mu_0, \Sigma_0) \prod_{t=1}^{n} P(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t) P(Q) \\ &\propto \prod_{t=1}^{n} P(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t) P(Q) \end{aligned} \tag{A.15}$$

This is our full conditional posterior for Q. Under our dynamic linear model, x_t is distributed normally with $x_t \sim N(\Phi x_{t-1} + \Upsilon u_t, Q)$. Thus, its PDF is:

$$pdf(x_t) = det(2\pi Q)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)}$$

Thus,

$$\begin{aligned}
P(Q|\mu_0, \Sigma_0, \Phi, \Upsilon, A, R, u, x, y) &\propto \prod_{t=1}^{n} P(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t)P(Q) \\
&\propto \prod_{t=1}^{n} det(2\pi Q)^{-\frac{1}{2}} e^{-\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)} P(Q) \\
&\propto det(Q^{-\frac{n}{2}}) e^{\sum_{t=1}^{n} -\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)} P(Q)
\end{aligned} \tag{A.16}$$

For the prior $P(Q)$ we choose the inverse Wishart distribution, the conjugate prior for covariance matrix of the multivariate normal distribution. The inverse Wishart is parameterized by a $p \times p$ positive definite scale matrix T and degrees of freedom $\nu > p - 1$.

$$\begin{aligned}
P(Q) &\sim W^{-1}(T, \nu) \\
P(Q) &= \frac{det(T)^{\frac{\nu}{2}}}{2^{\frac{\nu p}{2}} \Gamma_p(\frac{\nu}{2})} det(Q)^{-\frac{\nu+p+1}{2}} e^{-\frac{1}{2} tr(TQ^{-1})} \\
&\propto det(Q)^{-\frac{\nu+p+1}{2}} e^{-\frac{1}{2} tr(TQ^{-1})}
\end{aligned} \tag{A.17}$$

Where $\Gamma_p()$ is the multivariate gamma function. Multiplying the prior into the full conditional posterior, we get:

$$\begin{aligned}
P(Q|\mu_0, \Sigma_0, \Phi, \Upsilon, A, R, u, x, y) &\propto det(Q^{-\frac{n}{2}}) e^{\sum_{t=1}^{n} -\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)} P(Q) \\
&\propto det(Q^{-\frac{n}{2}}) e^{\sum_{t=1}^{n} -\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)} det(Q)^{-\frac{\nu+p+1}{2}} e^{-\frac{1}{2} tr(TQ^{-1})} \\
&= det(Q)^{-\frac{n}{2} - \frac{\nu+p+1}{2}} e^{\sum_{t=1}^{n} -\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t) - \frac{1}{2} tr(TQ^{-1})} \\
&= det(Q)^{-\frac{n+\nu+p+1}{2}} e^{\sum_{t=1}^{n} -\frac{1}{2} tr((x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)) - \frac{1}{2} tr(TQ^{-1})} \\
&= det(Q)^{-\frac{n+\nu+p+1}{2}} e^{\sum_{t=1}^{n} -\frac{1}{2} tr(Q^{-1}(x_t - \Phi x_{t-1} - \Upsilon u_t)(x_t - \Phi x_{t-1} - \Upsilon u_t)') - \frac{1}{2} tr(TQ^{-1})} \\
&= det(Q)^{-\frac{n+\nu+p+1}{2}} e^{\sum_{t=1}^{n} -\frac{1}{2} tr((x_t - \Phi x_{t-1} - \Upsilon u_t)(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1}) - \frac{1}{2} tr(TQ^{-1})} \\
&= det(Q)^{-\frac{n+\nu+p+1}{2}} e^{-\frac{1}{2} tr((T + \sum_{t=1}^{n} (x_t - \Phi x_{t-1} - \Upsilon u_t)(x_t - \Phi x_{t-1} - \Upsilon u_t)') Q^{-1})}
\end{aligned} \tag{A.18}$$

Thus, $P(Q|\mu_0, \Sigma_0, \Phi, \Upsilon, A, R, u, x, y) \sim W^{-1}(T + \sum_{t=1}^{n} (x_t - \Phi x_{t-1} - \Upsilon u_t)(x_t - \Phi x_{t-1} - \Upsilon u_t)', n + \nu)$. We can sample this in a Gibbs step.

A.7 A

$$
\begin{aligned}
P(A|\mu_0, \Sigma_0, \Phi, \Upsilon, Q, R, u, x, y) &\propto P(A, \mu_0, \Sigma_0, \Phi, \Upsilon, Q, R, u, x, y) \\
\propto P(y|x, \mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) P(x|\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) P(\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) \\
&\propto \prod_{t=1}^{n} P(y_t|x_t, A, R) P(x_0|\mu_0, \Sigma_0) \prod_{t=1}^{n} P(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t) P(A) \\
&\propto \prod_{t=1}^{n} P(y_t|x_t, A, R) P(A)
\end{aligned}
\tag{A.19}
$$

This is our full conditional posterior for A. Under our dynamic linear model, y_t is distributed normally with $y_t \sim N(Ax_t, R)$. Thus, its PDF is:

$$
pdf(y_t) = det(2\pi R)^{-\frac{1}{2}} e^{-\frac{1}{2}(y_t - Ax_t)' R^{-1} (y_t - Ax_t)}
$$

Thus,

$$
\begin{aligned}
P(A|\mu_0, \Sigma_0, \Phi, \Upsilon, Q, R, u, x, y) &\propto \prod_{t=1}^{n} P(y_t|x_t, A, R) P(A) \\
&\propto \prod_{t=1}^{n} det(2\pi R)^{-\frac{1}{2}} e^{-\frac{1}{2}(y_t - Ax_t)' R^{-1} (y_t - Ax_t)} P(A) \\
&\propto e^{\sum_{t=1}^{n} -\frac{1}{2}(y_t - AX_t)' R^{-1} (y_t - AX_t)} P(A)
\end{aligned}
\tag{A.20}
$$

Because A has a specific structure, we would like to experiment with various priors for $P(A)$ that are not necessarily conjugate to the multivariate normal. Since we won't be able to sample directly from $e^{\sum_{t=1}^{n} -\frac{1}{2}(y_t - AX_t)' R^{-1} (y_t - AX_t)} P(A)$, we will do a Metropolis step inside Gibbs sampling, like for Φ. Partially missing values in y_t can be handled by ignoring those terms in y_t as well as A in the product $(y_t - AX_t)$. Fully missing y_t can be handled by ignoring that term in the summation.

A.8 R

$$
\begin{aligned}
P(R|\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, u, x, y) &\propto P(R, \mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, u, x, y) \\
\propto P(y|x, \mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) P(x|\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) P(\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, R, u) \\
&\propto \prod_{t=1}^{n} P(y_t|x_t, A, R) P(x_0|\mu_0, \Sigma_0) \prod_{t=1}^{n} P(x_t|x_{t-1}, \Phi, \Upsilon, Q, u_t) P(R) \\
&\propto \prod_{t=1}^{n} P(y_t|x_t, A, R) P(R)
\end{aligned}
\tag{A.21}
$$

This is our full conditional posterior for R. Under our dynamic linear model, y_t is distributed normally with $y_t \sim N(Ax_t, R)$. Thus, its PDF is:

$$
pdf(y_t) = det(2\pi R)^{-\frac{1}{2}} e^{-\frac{1}{2}(y_t - Ax_t)' R^{-1} (y_t - Ax_t)}
$$

Thus,

$$
\begin{aligned}
P(R|\mu_0, \Sigma_0, \Phi, \Upsilon, Q, A, u, x, y) &\propto \prod_{t=1}^{n} P(y_t|x_t, A, R) P(R) \\
&\propto \prod_{t=1}^{n} det(2\pi R)^{-\frac{1}{2}} e^{-\frac{1}{2}(y_t - Ax_t)' R^{-1} (y_t - Ax_t)} P(R) \\
&\propto det(R^{-\frac{n}{2}}) e^{\sum_{t=1}^{n} -\frac{1}{2}(y_t - AX_t)' R^{-1} (y_t - AX_t)} P(R)
\end{aligned}
\tag{A.22}
$$

For the prior $P(R)$ we choose the inverse Wishart distribution, the conjugate prior for covariance matrix of the multivariate normal distribution. The inverse Wishart is parameterized by a $p \times p$ positive definite scale matrix T and degrees of freedom $\nu > p - 1$.

$$
\begin{aligned}
P(R) &\sim W^{-1}(T, \nu) \\
P(R) &= \frac{det(T)^{\frac{\nu}{2}}}{2^{\frac{\nu p}{2}} \Gamma_p(\frac{\nu}{2})} det(R)^{-\frac{\nu+p+1}{2}} e^{-\frac{1}{2} tr(TR^{-1})} \\
&\propto det(R)^{-\frac{\nu+p+1}{2}} e^{-\frac{1}{2} tr(TR^{-1})}
\end{aligned}
\tag{A.23}
$$

Where $\Gamma_p()$ is the multivariate gamma function. Multiplying the prior into the full conditional posterior, we get:

$$\begin{aligned}
P(R|\mu_0,\Sigma_0,\Phi,\Upsilon,Q,A,u,x,y) &\propto det(R^{-\frac{n}{2}})e^{\sum_{t=1}^{n}-\frac{1}{2}(y_t-AX_t)'R^{-1}(y_t-AX_t)}P(R)\\
&\propto det(R^{-\frac{n}{2}})e^{\sum_{t=1}^{n}-\frac{1}{2}(y_t-AX_t)'R^{-1}(y_t-AX_t)}det(R)^{-\frac{\nu+p+1}{2}}e^{-\frac{1}{2}tr(TR^{-1})}\\
&= det(R)^{-\frac{n}{2}-\frac{\nu+p+1}{2}}e^{\sum_{t=1}^{n}-\frac{1}{2}(y_t-AX_t)'R^{-1}(y_t-AX_t)-\frac{1}{2}tr(TR^{-1})}\\
&= det(R)^{-\frac{n+\nu+p+1}{2}}e^{\sum_{t=1}^{n}-\frac{1}{2}tr((y_t-AX_t)'R^{-1}(y_t-AX_t))-\frac{1}{2}tr(TR^{-1})}\\
&= det(R)^{-\frac{n+\nu+p+1}{2}}e^{\sum_{t=1}^{n}-\frac{1}{2}tr(R^{-1}(y_t-AX_t)(y_t-AX_t)')-\frac{1}{2}tr(TR^{-1})}\\
&= det(R)^{-\frac{n+\nu+p+1}{2}}e^{\sum_{t=1}^{n}-\frac{1}{2}tr((y_t-AX_t)(y_t-AX_t)'R^{-1})-\frac{1}{2}tr(TR^{-1})}\\
&= det(R)^{-\frac{n+\nu+p+1}{2}}e^{-\frac{1}{2}tr((T+\sum_{t=1}^{n}(y_t-AX_t)(y_t-AX_t)')R^{-1})} \qquad \text{(A.24)}
\end{aligned}$$

Thus, $P(R|\mu_0,\Sigma_0,\Phi,\Upsilon,Q,A,u,x,y) \sim W^{-1}(T+\sum_{t=1}^{n}(y_t-AX_t)(y_t-AX_t)', n+\nu)$. We can sample this in a Gibbs step. Partially missing values in y_t can be handled by ignoring those terms in y_t as well as A in the product $(y_t - AX_t)$. Fully missing y_t can be handled by ignoring that term in the summation and subtracting 1 from $n+\nu$ for each fully missing y_t.

APPENDIX B

ADDITIONAL MCMC CONSIDERATIONS

B.1 HANDLING ZEROS

In calculating the likelihoods for the Kalman filter or MCMC, we make frequent use of the probability density function of the multivariate normal distribution, which involves taking the inverse of the covariance matrix. The inverse of the covariance matrix, also known as the precision matrix, can be difficult to calculate. A covariance matrix must be positive definite to be invertible, meaning there can be no zeroes along the diagonal. However, in our definition of the AIM-BP DLM, we would require several elements of the hidden variables x to remain constant or relatively noise-free, meaning we need a zero, or something close to zero, on the corresponding elements of the diagonal for Q and Σ_0, and in entries of P_t^t for the Kalman filter. We can handle close to zero normally, but exactly zero makes the covariance matrix no longer positive definite. We run into this problem with the last element, which should be constant. We can either add a small variance to it, or figure out a way to calculate the inverse.

For example, in the likelihood for Φ, we need to calculate:

$$e^{\sum_{t=1}^{n} -\frac{1}{2}(x_t - \Phi x_{t-1} - \Upsilon u_t)' Q^{-1} (x_t - \Phi x_{t-1} - \Upsilon u_t)} P(\Phi)$$

The difference $x_t - \Phi x_{t-1} - \Upsilon u_t$ should have a zero in the last element. As the variance in the last element approaches zero, the precision approaches infinity, but the product should

still remain at zero for that element. Instead of using a regular inverse, a Moore-Penrose pseudo-inverse can handle this edge case by setting the inverse zero elements to zero instead of infinity.

For the Kalman filter, we need to calculate an inverse for the Kalman gain:

$$K_t = P_t^{t-1} A'(A P_t^{t-1} A' + R)^{-1}$$

Since we expect R to have all non-zero variances, this inverse should not run into issues with zeroes.

For the Kalman smoother, we need to calculate the inverse of P_t^{t-1}:

$$J_{t-1} = P_{t-1}^{t-1} \Phi' (P_t^{t-1})^{-1}$$

B.2 HANDLING HOMEOSTASIS TARGETS AND CONSTANT CONSTRAINTS IN X

In running the Kalman filter to calculate the expected mean X_t and variance P_t^t, we run into the following issue: The observed blood pressure and heart rate in Y are observed from the hidden blood pressure and heart rate in X. The hidden variables in X corresponding to the homeostasis targets and the constant, on the other hand, have no influence on any observed variable in our model. When calculating the expected variance P_t^t, the Kalman filter has no observations for the homeostasis targets and constant constraints and the Kalman gain is thus zero for these variables. In this situation, the Kalman filter calculates expected variance fully from their variance in Q, adding it at each time step. If we assign some amount of variance to these variables in Q, this results in us becoming more and more unsure about the expected mean X_t at each time step, expressed as increasing P_t^t. When using P_t^t to calculate P_t^T in the Kalman smoother, this uncertainty carries over. Finally, this means that our sampled values for X at each Gibbs step become increasingly erratic for the homeostasis target and constant constraints at the later time points. This throws off our sampling for Q, and the process exacerbates itself every cycle.

This is not behavior that we want in our model. We have a perfectly good idea of what the constant variable in X should be - it should be constant. While we may not be sure about the variance in the homeostasis target variables, we know that they should not increase over time - the fact that they are increasing over time is merely a result of it not influencing any observed variables. In order to combat this, we must either set all variances for these variables to zero in Q, or ignore the calculated P_t^T when sampling X. In addition, when sampling Q using the conjugate inverse Wishart, we need to make sure we hold the entries corresponding to the homeostasis target and constant constraint variables to zero.

BIBLIOGRAPHY

1. W. R. Gilks and P. Wild. Adaptive Rejection Sampling for Gibbs Sampling. *Applied Statistics*, 41(2):337, 1992.

2. W. R. Gilks, N. G. Best, and K. K. C. Tan. Adaptive Rejection Metropolis Sampling within Gibbs Sampling. *Applied Statistics*, 44(4):455, 1995.

3. Luca Martino, Jesse Read, and David Luengo. Independent Doubly Adaptive Rejection Metropolis Sampling Within Gibbs Sampling. *IEEE Transactions on Signal Processing*, 63(12):3123–3138, jun 2015.

4. Ritvij Bowry, Digvijaya D Navalkele, and Nicole R Gonzales. Blood pressure management in stroke: Five new things. *Neurology. Clinical practice*, 4(5):419–426, oct 2014.

5. Syeda L. Alqadri, Varun Sreenivasan, and Adnan I. Qureshi. Acute Hypertensive Response Management in Patients with Acute Stroke. *Current Cardiology Reports*, 15(12):426, dec 2013.

6. Kenneth Butcher and Magdy Selim. Acute Blood Pressure Management in Intracerebral Hemorrhage: Equipoise Resists an Attack. *Stroke*, 47(12):3065–3066, dec 2016.

7. M. McDermott, T. Jacobs, and L. Morgenstern. Critical care in acute ischemic stroke. In *Handbook of clinical neurology*, volume 140, pages 153–176. 2017.

8. Edward C. Jauch, Jeffrey L. Saver, Harold P. Adams, Askiel Bruno, J.J. (Buddy) Connors, Bart M. Demaerschalk, Pooja Khatri, Paul W. McMullan, Adnan I. Qureshi, Kenneth Rosenfield, Phillip A. Scott, Debbie R. Summers, David Z. Wang, Max Wintermark, and Howard Yonas. Guidelines for the Early Management of Patients With Acute Ischemic Stroke. *Stroke*, 44(3), 2013.

9. ENOS Trial Investigators. Efficacy of nitric oxide, with or without continuing antihypertensive treatment, for management of high blood pressure in acute stroke (ENOS): a partial-factorial randomised controlled trial. *Lancet (London, England)*, 385(9968):617–628, feb 2015.

10. Simona Lattanzi and Mauro Silvestrini. Blood pressure in acute intra-cerebral hemorrhage. *Annals of translational medicine*, 4(16):320, aug 2016.

11. Michael McManus and David S Liebeskind. Blood Pressure in Acute Ischemic Stroke. *Journal of clinical neurology (Seoul, Korea)*, 12(2):137–46, apr 2016.

12. William J Powers, Alejandro A Rabinstein, Teri Ackerson, Opeolu M Adeoye, Nicholas C Bambakidis, Kyra Becker, José Biller, Michael Brown, Bart M Demaerschalk, Brian Hoh, Edward C Jauch, Chelsea S Kidwell, Thabele M Leslie-Mazwi, Bruce Ovbiagele, Phillip A Scott, Kevin N Sheth, Andrew M Southerland, Deborah V Summers, David L Tirschwell, and on behalf of the American Heart Association Stroke American Heart Association Stroke Council. 2018 Guidelines for the Early Management of Patients With Acute Ischemic Stroke: A Guideline for Healthcare Professionals From the American Heart Association/American Stroke Association. *Stroke*, 49(3):e46–e110, mar 2018.

13. Jing Zhao, Panagiotis Papapetrou, Lars Asker, and Henrik Boström. Learning from heterogeneous temporal data in electronic health records. *Journal of Biomedical Informatics*, 65:105–119, jan 2017.

14. Iyad Batal, Hamed Valizadegan, Gregory F Cooper, and Milos Hauskrecht. A Pattern Mining Approach for Classifying Multivariate Temporal Data. *Proceedings. IEEE International Conference on Bioinformatics and Biomedicine*, 2011:358–365, nov 2011.

15. Anima Singh, Girish Nadkarni, Omri Gottesman, Stephen B. Ellis, Erwin P. Bottinger, and John V. Guttag. Incorporating temporal EHR data in predictive models for risk stratification of renal function deterioration. *Journal of Biomedical Informatics*, 53:220–228, feb 2015.

16. Trang Pham, Truyen Tran, Dinh Phung, and Svetha Venkatesh. Predicting healthcare trajectories from medical records: A deep learning approach. *Journal of Biomedical Informatics*, 69:218–229, may 2017.

17. Edward Choi, Andy Schuetz, Walter F Stewart, and Jimeng Sun. Using recurrent neural network models for early detection of heart failure onset. *Journal of the American Medical Informatics Association : JAMIA*, 24(2):361–370, 2017.

18. Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent Neural Networks for Multivariate Time Series with Missing Values. *Scientific Reports*, 8(1):6085, dec 2018.

19. Zitao Liu and Milos Hauskrecht. A Regularized Linear Dynamical System Framework for Multivariate Time Series Analysis. *Proceedings of the ... AAAI Conference on Artificial Intelligence. AAAI Conference on Artificial Intelligence*, 2015:1798–1804, jan 2015.

20. Zitao Liu and Milos Hauskrecht. Learning Adaptive Forecasting Models from Irregularly Sampled Multivariate Clinical Data. *Proceedings of the ... AAAI Conference on Artificial Intelligence. AAAI Conference on Artificial Intelligence*, 2016:1273–1279, feb 2016.

21. Eric H. Y. Lau, Calvin K. Y. Cheng, Dennis K. M. Ip, and Benjamin J. Cowling. Situational Awareness of Influenza Activity Based on Multiple Streams of Surveillance Data Using Multivariate Dynamic Linear Model. *PLoS ONE*, 7(5):e38346, may 2012.

22. Giovanni Petris, Sonia Petrone, and Patrizia Campagnoli. Dynamic linear models BT - Dynamic Linear Models with R. pages 31–84. Springer New York, New York, NY, 2009.

23. S. Isaka and A.V. Sebald. Control strategies for arterial blood pressure regulation. *IEEE Transactions on Biomedical Engineering*, 40(4):353–363, apr 1993.

24. Qiao Li, Roger G Mark, and Gari D Clifford. Artificial arterial blood pressure artifact models and an evaluation of a robust blood pressure and heart rate estimator. *BioMedical Engineering OnLine*, 8(1):13, jul 2009.

25. Peter B. Jensen, Lars J. Jensen, and Søren Brunak. Mining electronic health records: towards better research applications and clinical care. *Nature Reviews Genetics*, 13(6):395–405, jun 2012.

26. Matthew B. Bevers and W. Taylor Kimberly. Critical Care Management of Acute Ischemic Stroke. *Current Treatment Options in Cardiovascular Medicine*, 19(6):41, jun 2017.

27. Ravi S. Menon, Richard E. Burgess, Jeffrey J. Wing, M. Christopher Gibbons, Nawar M. Shara, Stephen Fernandez, Annapurni Jayam-Trouth, Laura German, Ian Sobotka, Dorothy Edwards, and Chelsea S. Kidwell. Predictors of highly prevalent brain ischemia in intracerebral hemorrhage. *Annals of Neurology*, 71(2):199–205, feb 2012.

28. David Dornbos III, Ciaran J. Powers, Yuchuan Ding, and Liping Liu. Neurocritical care in the treatment of stroke. *Neurological Research*, 38(6):491–494, jun 2016.

29. Adnan I. Qureshi, Mustapha A. Ezzeddine, Abu Nasar, M. Fareed K. Suri, Jawad F. Kirmani, Haitham M. Hussein, Afshin A. Divani, and Alluru S. Reddi. Prevalence of elevated blood pressure in 563704 adult patients with stroke presenting to the ED in the United States. *The American Journal of Emergency Medicine*, 25(1):32–38, jan 2007.

30. M Britton, A Carlsson, and U de Faire. Blood pressure course in patients with acute stroke and matched controls. *Stroke*, 17(5):861–4, sep 1986.

31. J D Wallace and L L Levy. Blood pressure after stroke. *JAMA*, 246(19):2177–80, nov 1981.

32. E Sankaranarayanan Prakash, Madanmohan, and Perttu J Lindsberg. What Causes the Acute Blood Pressure Elevation After Stroke? * Response:. *Stroke*, 36(10):2066, oct 2005.

33. Mark Willmot, Jo Leonardi-Bee, and Philip M W Bath. High blood pressure in acute stroke and subsequent outcome: a systematic review. *Hypertension (Dallas, Tex. : 1979)*, 43(1):18–24, jan 2004.

34. Philip M. Bath, Jason P. Appleton, Kailash Krishnan, and Nikola Sprigg. Blood Pressure in Acute Stroke. *Stroke*, 49(7):1784–1790, jul 2018.

35. J. Castillo, Rogelio Leira, María M García, Joaquín Serena, Miguel Blanco, and Antoni Dávalos. Blood Pressure Decrease During the Acute Phase of Ischemic Stroke Is Associated With Brain Injury and Poor Stroke Outcome. *Stroke*, 35(2):520–526, feb 2004.

36. A Chamorro, N Vila, C Ascaso, E Elices, W Schonewille, and R Blanc. Blood pressure and functional recovery in acute ischemic stroke. *Stroke*, 29(9):1850–3, sep 1998.

37. M. A. Ritter, P. Kimmeyer, P. U. Heuschmann, R. Dziewas, R. Dittrich, D. G. Nabavi, and E. B. Ringelstein. Blood Pressure Threshold Violations in the First 24 Hours After Admission for Acute Stroke: Frequency, Timing, Predictors, and Impact on Clinical Outcome. *Stroke*, 40(2):462–468, feb 2009.

38. J Oliveira-Filho, S C S Silva, C C Trabuco, B B Pedreira, E U Sousa, and A Bacellar. Detrimental effect of blood pressure reduction in the first 24 hours of acute stroke onset. *Neurology*, 61(8):1047–51, oct 2003.

39. M Kaste, R Fogelholm, T Erilä, H Palomäki, K Murros, A Rissanen, and S Sarna. A randomized, double-blind, placebo-controlled trial of nimodipine in acute ischemic hemispheric stroke. *Stroke*, 25(7):1348–53, jul 1994.

40. John F Potter, Thompson G Robinson, Gary A Ford, Amit Mistri, Martin James, Julia Chernova, and Carol Jagger. Controlling hypertension and hypotension immediately post-stroke (CHHIPS): a randomised, placebo-controlled, double-blind pilot trial. *The Lancet Neurology*, 8(1):48–56, jan 2009.

41. J. Schrader, S. Luders, A. Kulschewski, J. Berger, W. Zidek, J. Treib, K. Einhaupl, H. C. Diener, P. Dominiak, and Acute Candesartan Cilexetil Therapy in Stroke Survivors Study Group. The ACCESS Study: Evaluation of Acute Candesartan Cilexetil Therapy in Stroke Survivors. *Stroke*, 34(7):1699–1703, jul 2003.

42. Else Charlotte Sandset, Philip MW Bath, Gudrun Boysen, Dalius Jatuzis, Janika Kõrv, Stephan Lüders, Gordon D Murray, Przemyslaw S Richter, Risto O Roine, Andreas Terént, Vincent Thijs, Eivind Berge, and SCAST Study Group. The angiotensin-receptor blocker candesartan for treatment of acute stroke (SCAST): a randomised, placebo-controlled, double-blind trial. *The Lancet*, 377(9767):741–750, feb 2011.

43. Thompson G Robinson, John F Potter, Gary A Ford, Christopher J Bulpitt, Julia Chernova, Carol Jagger, Martin A James, Joanne Knight, Hugh S Markus, Amit K Mistri, Neil R Poulter, and COSSACS Investigators. Effects of antihypertensive treatment after acute stroke in the Continue Or Stop post-Stroke Antihypertensives Collaborative Study (COSSACS): a prospective, randomised, open, blinded-endpoint trial. *The Lancet Neurology*, 9(8):767–775, aug 2010.

44. J Claude Hemphill, Steven M Greenberg, Craig S Anderson, Kyra Becker, Bernard R Bendok, Mary Cushman, Gordon L Fung, Joshua N Goldstein, R Loch Macdonald, Pamela H Mitchell, Phillip A Scott, Magdy H Selim, Daniel Woo, American Heart Association Stroke Council, Council on Cardiovascular and Stroke Nursing, and Council on Clinical Cardiology. Guidelines for the Management of Spontaneous Intracerebral Hemorrhage: A Guideline for Healthcare Professionals From the American Heart Association/American Stroke Association. *Stroke*, 46(7):2032–60, jul 2015.

45. Craig S. Anderson, Thompson Robinson, Richard I. Lindley, Hisatomi Arima, Pablo M. Lavados, Tsong-Hai Lee, Joseph P. Broderick, Xiaoying Chen, Guofang Chen, Vijay K. Sharma, Jong S. Kim, Nguyen H. Thang, Yongjun Cao, Mark W. Parsons, Christopher Levi, Yining Huang, Verónica V. Olavarría, Andrew M. Demchuk, Philip M. Bath, Geoffrey A. Donnan, Sheila Martins, Octavio M. Pontes-Neto, Federico Silva, Stefano Ricci, Christine Roffe, Jeyaraj Pandian, Laurent Billot, Mark Woodward, Qiang Li, Xia Wang, Jiguang Wang, and John Chalmers. Low-Dose versus Standard-Dose Intravenous Alteplase in Acute Ischemic Stroke. *New England Journal of Medicine*, 374(24):2313–2323, jun 2016.

46. Daniel Strbian and Gustavo Saposnik. Review of the ENCHANTED Trial (Enhanced Control of Hypertension and Thrombolysis Stroke Study): How Low Can We Go With Intravenous Tissue-Type Plasminogen Activator Dose and Blood Pressure Level? *Stroke*, 47(12):3063–3064, dec 2016.

47. Craig S. Anderson, Emma Heeley, Yining Huang, Jiguang Wang, Christian Stapf, Candice Delcourt, Richard Lindley, Thompson Robinson, Pablo Lavados, Bruce Neal, Jun Hata, Hisatomi Arima, Mark Parsons, Yuechun Li, Jinchao Wang, Stephane Heritier, Qiang Li, Mark Woodward, R. John Simes, Stephen M. Davis, and John Chalmers. Rapid Blood-Pressure Lowering in Patients with Acute Intracerebral Hemorrhage. *New England Journal of Medicine*, 368(25):2355–2365, jun 2013.

48. Adnan I. Qureshi, Yuko Y. Palesch, William G. Barsan, Daniel F. Hanley, Chung Y. Hsu, Renee L. Martin, Claudia S. Moy, Robert Silbergleit, Thorsten Steiner, Jose I. Suarez, Kazunori Toyoda, Yongjun Wang, Haruko Yamamoto, and Byung-Woo Yoon. Intensive Blood-Pressure Lowering in Patients with Acute Cerebral Hemorrhage. *New England Journal of Medicine*, 375(11):1033–1043, sep 2016.

49. Jiang He, Yonghong Zhang, Tan Xu, Qi Zhao, Dali Wang, Chung-Shiuan Chen, Weijun Tong, Changjie Liu, Tian Xu, Zhong Ju, Yanbo Peng, Hao Peng, Qunwei Li, Deqin Geng, Jintao Zhang, Dong Li, Fengshan Zhang, Libing Guo, Yingxian Sun, Xuemei

Wang, Yong Cui, Yongqiu Li, Dihui Ma, Guang Yang, Yanjun Gao, Xiaodong Yuan, Lydia A. Bazzano, Jing Chen, and CATIS Investigators. Effects of Immediate Blood Pressure Reduction on Death and Major Disability in Patients With Acute Ischemic Stroke. *JAMA*, 311(5):479, feb 2014.

50. Xiaoqing Bu, Changwei Li, Yonghong Zhang, Tan Xu, Dali Wang, Yingxian Sun, Hao Peng, Tian Xu, Chung-Shiuan Chen, Lydia A. Bazzano, Jing Chen, Jiang He, and CATIS Investigators. Early Blood Pressure Reduction in Acute Ischemic Stroke with Various Severities: A Subgroup Analysis of the CATIS Trial. *Cerebrovascular Diseases*, 42(3-4):186–195, 2016.

51. Mark C van der Wel, Marion Biermans, Reinier Akkermans, Jacques WM Lenders, Chris van Weel, and Jaap Deinum. Patient characteristics do not predict the individual response to antihypertensive medication: a cross-over trial. *Family Practice*, 35(1):67–73, jan 2018.

52. Theodore Raphan, Bernard Cohen, Yongqing Xiang, and Sergei B Yakushin. A Model of Blood Pressure, Heart Rate, and Vaso-Vagal Responses Produced by Vestibulo-Sympathetic Activation. *Frontiers in neuroscience*, 10:96, 2016.

53. Karla Caballero and Ram Akella. Dynamic Estimation of the Probability of Patient Readmission to the ICU using Electronic Medical Records. *AMIA ... Annual Symposium proceedings. AMIA Symposium*, 2015:1831–40, 2015.

54. J Geoffrey Chase, Christina Starfinger, ZhuHui Lam, Franck Agogue, and Geoffrey M Shaw. Quantifying agitation in sedated ICU patients using heart rate and blood pressure. *Physiological Measurement*, 25(4):1037–1051, aug 2004.

55. Jeffrey D. DellaVolpe, Jason E. Moore, and Michael R. Pinsky. Arterial blood pressure and heart rate regulation in shock state. *Current Opinion in Critical Care*, 21(5):376–380, oct 2015.

56. Tony Berger, Jeffrey Green, Timothy Horeczko, Yolanda Hagar, Nidhi Garg, Alison Suarez, Edward Panacek, and Nathan Shapiro. Shock index and early recognition of sepsis in the emergency department: pilot study. *The western journal of emergency medicine*, 14(2):168–74, mar 2013.

57. Li-wei H Lehman, Shamim Nemati, Ryan P Adams, George Moody, Atul Malhotra, and Roger G Mark. Tracking progression of patient state of health in critical care using inferred shared dynamics in physiological time series. *Conference proceedings : ... Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual Conference*, 2013:7072–5, 2013.

58. Li-wei H. Lehman, Ryan P. Adams, Louis Mayaud, George B. Moody, Atul Malhotra, Roger G. Mark, and Shamim Nemati. A Physiological Time Series Dynamics-Based

Approach to Patient Monitoring and Outcome Prediction. *IEEE Journal of Biomedical and Health Informatics*, 19(3):1068–1076, may 2015.

59. Li-wei H. Lehman, Roger G. Mark, and Shamim Nemati. A Model-Based Machine Learning Approach to Probing Autonomic Regulation From Nonstationary Vital-Sign Time Series. *IEEE Journal of Biomedical and Health Informatics*, 22(1):56–66, jan 2018.

60. Bryce Yapps, Sungtae Shin, Ramin Bighamian, Jill Thorsen, Colleen Arsenault, Sadeq A. Quraishi, Jin-Oh Hahn, and Andrew T. Reisner. Hypotension in ICU Patients Receiving Vasopressor Therapy. *Scientific Reports*, 7(1):8551, dec 2017.

61. James Macdougall. Analysis of DoseResponse StudiesEmax Model. In *Dose Finding in Drug Development*, pages 127–145. Springer New York, New York, NY, 2006.

62. Melanie A Felmlee, Marilyn E Morris, and Donald E Mager. Mechanism-based pharmacodynamic modeling. *Methods in molecular biology (Clifton, N.J.)*, 929:583–600, 2012.

63. Joseph Wu, Anindita Banerjee, Bo Jin, Sandeep M Menon, Steven W Martin, and Anne C Heatherington. Clinical doseresponse for a broad set of biological products: A model-based meta-analysis. *Statistical Methods in Medical Research*, page 096228021668452, jan 2017.

64. Robert H. Shumway and David S. Stoffer. *Time Series Analysis and Its Applications*. Springer Texts in Statistics. Springer International Publishing, Cham, 2017.

65. Sylvia Frühwirth-Schnatter. Data Augmentation and Dynamic Linear Models. *Journal of Time Series Analysis*, 15(2):183–202, mar 1994.

66. Bradley P. Carlin, Nicholas G. Polson, and David S. Stoffer. A Monte Carlo Approach to Nonnormal and Nonlinear State-Space Modeling. *Journal of the American Statistical Association*, 87(418):493, jun 1992.

67. C. K. Carter and R. Kohn. On Gibbs sampling for state space models. *Biometrika*, 81(3):541–553, sep 1994.

68. Paul Fearnhead. *MCMC for state-space models*, pages 513–529. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. Chapman and Hall, 2011.

69. Scott M. Lynch. *Introduction to Applied Bayesian Statistics and Estimation for Social Scientists*. Springer New York, New York, NY, 2007.

70. G. O. Roberts, A. Gelman, and W. R. Gilks. Weak convergence and optimal scaling of random walk Metropolis algorithms. *The Annals of Applied Probability*, 7(1):110–120, feb 1997.

71. Wally Gilks. Adaptive Rejection Sampling, 2003.

72. David Spiegelhalter, Andrew Thomas, Nicky Best, and Wally Gilks. BUGS Manual, Section 3, 1995.

73. Alistair E.W. Johnson, Tom J. Pollard, Lu Shen, Li-wei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3:160035, may 2016.

74. Katharine E. Henry, David N. Hager, Peter J. Pronovost, and Suchi Saria. A targeted real-time early warning score (TREWScore) for septic shock. *Science Translational Medicine*, 7(299):299ra122–299ra122, aug 2015.

75. Romain Pirracchio, Maya L Petersen, Marco Carone, Matthieu Resche Rigon, Sylvie Chevret, and Mark J van der Laan. Mortality prediction in intensive care units with the Super ICU Learner Algorithm (SICULA): a population-based study. *The Lancet Respiratory Medicine*, 3(1):42–52, jan 2015.

76. Alain Saas, Anna Guitart, and Africa Perianez. Discovering playing patterns: Time series clustering of free-to-play game data. In *2016 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 1–8. IEEE, sep 2016.

77. Q Li and G D Clifford. Dynamic time warping and machine learning for signal quality assessment of pulsatile signals. *Physiological Measurement*, 33(9):1491–1501, sep 2012.

78. Paolo Tormene, Toni Giorgino, Silvana Quaglini, and Mario Stefanelli. Matching incomplete time series with dynamic time warping: an algorithm and an application to post-stroke rehabilitation. *Artificial Intelligence in Medicine*, 45(1):11–34, jan 2009.

79. Adam Mahdi, Jacob Sturdy, Johnny T. Ottesen, and Mette S. Olufsen. Modeling the Afferent Dynamics of the Baroreflex Control System. *PLoS Computational Biology*, 9(12):e1003384, dec 2013.

80. Richard L Lalonde, Tracey L O'Rear, Irving W Wainer, Kristin D Drda, Vanessa L Herring, and Michael B Bottorff. Labetalol pharmacokinetics and pharmacodynamics: Evidence of stereoselective disposition. *Clinical Pharmacology and Therapeutics*, 48(5):509–519, nov 1990.

81. Tomoko Saotome, Shigeki Minoura, Keiko Terashi, Takashi Sato, Hirotoshi Echizen, and Takashi Ishizaki. Labetalol in Hypertension During the Third Trimester of Pregnancy: Its Antihypertensive Effect and Pharmacokinetic-Dynamic Analysis. *The Journal of Clinical Pharmacology*, 33(10):979–988, oct 1993.

82. M Chauvin, H Deriaz, and P Viars. Continuous i.v. infusion of labetalol for postoperative hypertension. Haemodynamic effects and plasma kinetics. *British journal of anaesthesia*, 59(10):1250–6, oct 1987.

83. D J Wilson, J D Wallin, N D Vlachakis, E D Freis, D G Vidt, E L Michelson, H G Langford, W Flamenbaum, and M P Poland. Intravenous labetalol in the treatment of severe hypertension and hypertensive emergencies. *The American journal of medicine*, 75(4A):95–102, oct 1983.

84. M J Serlin, M C Orme, M Maciver, G J Green, C M Macnee, and A M Breckenridge. Rate of onset of hypotensive effect of oral labetalol. *British journal of clinical pharmacology*, 7(2):165–8, feb 1979.

85. D. A. Richards, J. G. Maconochie, R. E. Bland, R. Hopkins, E. P. Woodings, and L. E. Martin. Relationship between plasma concentrations and pharmacological effects of labetalol. *European Journal of Clinical Pharmacology*, 11(2):85–90, 1977.

86. Monique P Curran, Dean M Robinson, and Gillian M Keating. Intravenous nicardipine: its use in the short-term treatment of hypertension and various other indications. *Drugs*, 66(13):1755–82, 2006.

87. D J Graham, R J Dow, D J Hall, O F Alexander, E J Mroszczak, and D Freedman. The metabolism and pharmacokinetics of nicardipine hydrochloride in man. *British journal of clinical pharmacology*, 20 Suppl 1(Suppl 1):23S–28S, 1985.

88. R C Schoemaker, J M van Gerven, and A F Cohen. Estimating potency for the Emax-model without attaining maximal effects. *Journal of pharmacokinetics and biopharmaceutics*, 26(5):581–93, oct 1998.

89. Donald W. Marquardt. An Algorithm for Least-Squares Estimation of Nonlinear Parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, jun 1963.

90. E Cook, G G Clifton, R Vargas, G Bienvenu, R Williams, N Sambol, G McMahon, S Grandy, C M Lai, and C Quon. Pharmacokinetics, pharmacodynamics, and minimum effective clinical dose of intravenous nicardipine. *Clinical pharmacology and therapeutics*, 47(6):706–18, jun 1990.

91. Yuzhe Liu and Vanathi Gopalakrishnan. An Overview and Evaluation of Recent Machine Learning Imputation Methods Using Cardiac Imaging Data. *Data*, 2(1):8, jan 2017.

92. Mathieu Lepot, Jean-Baptiste Aubin, and François Clemens. Interpolation in Time Series: An Introductive Overview of Existing Methods, Their Performance Criteria and Uncertainty Assessment. *Water*, 9(10):796, oct 2017.

93. Li-wei H. Lehman, Mohammed Saeed, Daniel Talmor, Roger Mark, and Atul Malhotra. Methods of Blood Pressure Measurement in the ICU*. *Critical Care Medicine*, 41(1):34–40, jan 2013.

94. Sara Ribezzo, Eleonora Spina, Stefano Di Bartolomeo, and Gianfranco Sanson. Non-invasive techniques for blood pressure measurement are not a reliable alternative to direct measurement: a randomized crossover trial in ICU. *TheScientificWorldJournal*, 2014:353628, 2014.

95. Li-Wei Lehman, Mohammed Saeed, George Moody, and Roger Mark. Hypotension as a Risk Factor for Acute Kidney Injury in ICU Patients. *Computing in cardiology*, 37:1095–1098, 2010.

96. Lui G. Forni and Michael Joannidis. Blood pressure deficits in acute kidney injury: not all about the mean arterial pressure? *Critical Care*, 21(1):102, dec 2017.

97. Fábio M de Castilho, Antonio Luiz P Ribeiro, José Luiz P da Silva, Vandack Nobre, and Marcos R de Sousa. Heart rate variability as predictor of mortality in sepsis: A prospective cohort study. *PloS one*, 12(6):e0180060, 2017.

98. Supreeth P. Shashikumar, Matthew D. Stanley, Ismail Sadiq, Qiao Li, Andre Holder, Gari D. Clifford, and Shamim Nemati. Early sepsis detection in critical care patients using multiscale blood pressure and heart rate dynamics. *Journal of Electrocardiology*, 50(6):739–743, nov 2017.

99. Thomas D Giles. Circadian rhythm of blood pressure and the relation to cardiovascular events. *Journal of Hypertension*, 24(Suppl 2):S11–S16, apr 2006.

100. G. Mancia. Short- and Long-Term Blood Pressure Variability: Present and Future. *Hypertension*, 60(2):512–517, aug 2012.

101. Sarah L Stevens, Sally Wood, Constantinos Koshiaris, Kathryn Law, Paul Glasziou, Richard J Stevens, and Richard J McManus. Blood pressure variability and cardiovascular disease: systematic review and meta-analysis. *BMJ (Clinical research ed.)*, 354:i4098, aug 2016.

102. Ayan Sen, Joseph Miller, Heidi Wilkie, Michele Moyer, Christopher Lewandowski, and Richard Nowak. Continuous hemodynamic monitoring in acute stroke: an exploratory analysis. *The western journal of emergency medicine*, 15(4):345–50, jul 2014.