

DATA-DRIVEN MODELING AND CONTROL OF CARDIAC SYSTEM

*Presented in Partial Fulfillment of
the Requirements for the Degree of*

MASTERS OF SCIENCE

with a Major in

Chemical Engineering

in the

College of Graduate Studies

University of Idaho

by

ANDREW L. BRANEN

Major Professor

GAUTAM KUMAR, PH.D.

Committee

D. ERIC ASTON, PH.D.

FUCHANG GAO, PH.D.

JAMES MOBERLY, PH.D.

Department Administrator

DEV SHRESTHA, PH.D.

AUGUST 2021

AUTHORIZATION TO SUBMIT THESIS

This thesis of Andrew L. Branen, submitted for the degree of Masters of Science with a Major in Chemical Engineering and titled "Data-Driven Modeling and Control of Cardiac System," has been reviewed in final form. Permission, as indicated by the signatures and dates below, is now granted to submit final copies to the College of Graduate Studies for approval.

Major Professor: _____
Gautam Kumar, Ph.D. _____
Date

Committee Members: _____
D. Eric Aston, Ph.D. _____
Date

Fuchang Gao, Ph.D. _____
Date

James Moberly, Ph.D. _____
Date

Department Administrator: _____
Dev Shrestha, Ph.D. _____
Date

ABSTRACT

Vagus nerve stimulation is an emerging therapy that seeks to offset pathological conditions by electrically stimulating the vagus nerve through cuff electrodes, where the electrical pulse is defined by several parameters such as pulse amplitude, pulse width, and pulse frequency. This electroceutical therapy has been approved for treatment resistant depression, and is currently under investigation for heart failure, heart arrhythmia, hypertension, and gastric motility disorders. Recent studies have shown the ability to selectively activate different fibers in the vagus nerve, thus allowing for a highly specific control of physiological behavior through vagal nerve stimulation. One of the major challenges with the application of this therapy involves a closed loop controller to autonomously control the behavioral responses. This problem becomes additionally challenging when multiple locations and multiple stimulation parameters are considered for optimization. Using a physiological model of a rat heart, this thesis investigates a data-driven control scheme for closed-loop control of the rat cardiac system. In the first section of this thesis, a data-driven modeling approach is used to develop a model that maps vagus nerve stimulation parameter selection to the effect on the physiological variables of heart rate and blood pressure. The second part of this thesis develops a controller that uses the data-driven model by utilizing a model predictive control framework to control the heart rate and the blood pressure in closed-loop simulations of a rat model.

TABLE OF CONTENTS

AUTHORIZATION TO SUBMIT THESIS	ii
ABSTRACT	iii
TABLE OF CONTENTS.	iv
LIST OF TABLES	v
LIST OF FIGURES.	vi
ACKNOWLEDGEMENTS	x
DEDICATION.	xi
1 INTRODUCTION	1
1.1 Vagus Nerve Stimulation	2
1.2 Autonomic Nervous System	3
1.3 Need for Closed-Loop Vagus Nerve Stimulation	6
1.4 Machine Learning Approaches	11
1.5 Thesis Outline	18
2 MAPPING VAGUS NERVE STIMULATION PARAMETERS TO CARDIAC PHYSIOLOGY USING LONG SHORT-TERM MEMORY NETWORK	20
2.1 Introduction	20
2.2 Physiological Model	21
2.3 Machine Learning Approach	26
2.4 Results	29
2.5 Best Neural Network Design Performance	31
2.6 Summary	32
3 DATA DRIVEN CONTROL OF VAGUS NERVE STIMULATION FOR THE CARDIAC SYSTEM	36
3.1 Introduction	36
3.2 LSTM in Closed-Loop with LSTM	37
3.3 LSTM in Closed-Loop with Physiological Model	40
3.4 Controller Analysis	50
3.5 Conclusions	55
4 SUMMARY AND FUTURE DIRECTIONS.	58
4.1 Summary	58
4.2 Future Work	59
BIBLIOGRAPHY.	63

LIST OF TABLES

TABLE 2.1	Top Three Trained Architectures	31
TABLE 2.2	Comparison of Computational Speed. Tests were done using an Intel(R) Core i7-9700 CPU 3.00 GHz with 16.0 GB of RAM.. . . .	32
TABLE 3.1	Set Points for the Controller	40
TABLE 3.2	Computational Time of Controllers Tests Were Done Using an INTEL(R) CORE I7-9700 CPU 3.00 GHZ with 16.0 GB of RAM . . .	55

LIST OF FIGURES

-
- FIGURE 1.1 A diagram of the parasympathetic (blue) and sympathetic (magenta) projections in the autonomic nervous system and their physiological effects on different organs. The vagal nerve projections (indicated by a green dot) innervate the respiratory system, the cardiovascular system, the gastrointestinal tract, and the liver. Activation of the vagus nerve can lead to bronchi constriction, slowing of heart rate, peristalsis, and bile secretion. Conversely, the sympathetic system opposes these actions. Afferent and efferent directions are shown in black. . . . 5
- FIGURE 1.2 **(A)** A diagram of the feedback loop between the cardiovascular system, and the central nervous system. Baroreceptors in the heart send information via afferent sensory nerve fibers to the medulla oblongata. The heart receives signaling from parasympathetic and sympathetic nerve fibers whose signals come from the cardioregulatory center in the central nervous system. Additionally depicted is the hypothalamic-pituitary-adrenal pathway at the bottom where sympathetic nerve fibers innervate the adrenal gland and cause epinephrine and norepinephrine release to the heart. **(B)** In black, sympathetic and parasympathetic pathway influence on physiological behavior in the cardiovascular system is depicted as a balance. In orange, the hyperactive sympathetic nervous system and diminished parasympathetic tone associated with cardiovascular diseases is depicted. . . . 7
- FIGURE 1.3 A diagram of a basic feedforward network that takes 3 inputs, and computes 2 outputs with two hidden layers of size 5 and 4. The size of the weight matrices between each layer are noted at the bottom. Activation functions (not shown) are applied at the end of each hidden layer as well as the output layer. . . . 12
- FIGURE 1.4 Diagram of a simple recurrent neural network (RNN). **(A)** Shows the input x passed to the hidden layer and transformed by weights U , h denotes the hidden state, W denotes the weights that transform the hidden state to the output, y . **(B)** Shows **A** unfolded along the time axis. Consider a sequence of length N , at the first step the hidden state is initialized to 0, after that, the hidden state is passed to generate all the outputs leading up to the final input at $t + N$ 13

FIGURE 1.5	A diagram of the internal structure of a long short-term memory network for input x_t , cell state c_t , and hidden state h_t . Current timesteps are subscripted with t , and previous timesteps are subscripted with $t - 1$. Activation functions are denoted by σ for the sigmoid function, and \tanh for the hyperbolic tangent function.	15
FIGURE 1.6	Prediction and optimal control in MPC.	17
FIGURE 2.1	A diagram that details the influence of vagal nerve stimulation in the full physiological model of the rat cardiac system.	22
FIGURE 2.2	A violin plot of the normalized data (unitless) for all variables. HR and MAP denote the heart rate and the mean arterial pressure, respectively. VNS PW indicates the stimulation pulse width at the location number indicated in the label, and F indicates the stimulation frequency at the location number indicated in the label. For example, VNS.PW1 denotes the pulse width parameter at location 1.	28
FIGURE 2.3	Comparison of (A) different neural network performances on the test set including the baseline case. (B) The effect of layer size on a neural network's predictive performance on the test set. (C) The effect of input neuron size on a neural network's predictive performance on the test set.	30
FIGURE 2.4	Example of the trained LSTM model performance on a trial from the test set for predicting the (A) heart rate (HR) and the (B) mean arterial pressure (MAP) for one selection of the VNS parameters. The LSTM is given the input (blue), and is asked to recursively predict the next 99 labels (black) with the LSTM predictions shown in red. (C) Normalized mean absolute error over the entire test set for the heart rate and the mean arterial pressure is shown, along with the standard deviation shaded around the curve.	33
FIGURE 2.5	Comparison between the LSTM model and the full model for the (A) heart rate and the (B) mean arterial blood pressure (MAP) for 99 simulated cardiac cycles. The black line shows the output of the full differential equation model (labeled Truth), while the red line shows the predictions from the LSTM model (labeled Prediction). Cycles 1-10: location 1 with pulse width 0.21 ms and pulse frequency 21 Hz (light blue). Cycles 11-20: location 3 with pulse width 0.09 ms and pulse frequency 9 Hz (light red). Cycles 21-30: location 2 with pulse width 0.31 ms and pulse frequency 30 Hz (light green). Cycles 50-60: location 1 with pulse width 0.14 ms and pulse frequency 14 Hz (dark blue). All other cycles had no locations active.	34

FIGURE 3.1	A block diagram that details the design of the controller, labeled as MPC. The target set point is denoted with r , the optimal control actions are denoted by u^* , the measurement of the physiological variables from the rat cardiac system (labeled as subject) are denoted by y , and the predictions from the LSTM are denoted with \hat{y}	38
FIGURE 3.2	LSTM controller tracking set points for (A) heart rate (HR) and (B) mean arterial blood pressure (MAP) when controlling the LSTM in closed-loop. Set points are denoted with a black line, while the controlled output is indicated with red dots.	41
FIGURE 3.3	LSTM controller actions to achieve the set point for (A) pulse width and (B) pulse frequency when controlling the trained LSTM in closed loop. The three different locations are differentiated by color with the first location in black, the second location in blue, and the third location in red.	42
FIGURE 3.4	Performance of L1 objective function in LSTM based controller for set point tracking of (A) heart rate (HR) and (B) mean arterial blood pressure (MAP) of the physiological model. Set points are denoted with a black line, while the physiological model output is indicated with blue dots.	44
FIGURE 3.5	Controller actions for the L1 objective function to achieve the set points for (A) pulse width and (B) pulse frequency in controlling the physiological model. The three different locations are differentiated by color with the first location in black, the second location in blue, and the third location in red.	45
FIGURE 3.6	Performance of L2 objective function in LSTM based controller for set point tracking of (A) heart rate (HR) and (B) mean arterial blood pressure (MAP) of the physiological model. Set points are denoted with a black line, while the physiological model output is indicated with blue dots.	46
FIGURE 3.7	Controller actions for the L2 objective function to achieve the set points for (A) pulse width and (B) pulse frequency in the physiological model. The three different locations are differentiated by color with the first location in black, the second location in blue, and the third location in red.	47
FIGURE 3.8	Performance of L1 with Δu objective function in LSTM based controller for set point tracking of (A) heart rate (HR) and (B) mean arterial blood pressure (MAP) of the full physiological model. Set points are denoted with a black line, while the physiological model output is indicated with blue dots.	49
FIGURE 3.9	Controller actions for the L1 with Δu objective function to achieve the set points for (A) pulse width and (B) pulse frequency. The three different locations are differentiated by color with the first location in black, the second location in blue, and the third location in red.	50

- FIGURE 3.10 Comparison of controller performance for the L1 cost denoted by Eqn. 3.1 (blue dots) and controller performance for the L2 cost denoted by Eqn. 3.2 (red crosses) with $N_c = 5$, $N_p = 10$, and $\lambda = 0.001$ in both controller formulations for **(A)** heart rate (HR) and **(B)** mean arterial blood pressure (MAP). The target set points are shown by the black line. 52
- FIGURE 3.11 Comparison of controller actions for the L1 cost denoted by Eqn. 3.1 (solid lines) and controller action selection for the L2 cost denoted by Eqn. 3.2 (dotted lines) with $N_c = 5$, $N_p = 10$, and $\lambda = 0.001$ in both controller formulations for **(A)** pulse width and **(B)** pulse frequency. Locations are differentiated by color with L1 location 1 shown in black, location 2 shown in blue, and location 3 shown in red. For L2, location 1 is shown in orange, location 2 is shown in magenta, and location 3 is shown in grey. 53
- FIGURE 4.1 A control framework approach that utilizes an autoencoder for dimensionality reduction to obtain a linearized model for control. The encoder is used to obtain the reduced dimensional model, which then undergoes Koopman analysis to generate a linear model for control. The decoder can be used to reconstruct the physiological activity. 61

ACKNOWLEDGEMENTS

During my time at the University of Idaho, I received a lot of assistance from key individuals that accelerated my growth and facilitated my accomplishments. In more ways than I can express, the following individuals were crucial to my success.

First, I would like to extend my deepest thanks and gratitude to Dr. Kumar, and his patience, dedication, and support throughout my time to complete this thesis. With an open door, he always extended the greatest amount of assistance to me, and that proved instrumental in my success. I'm deeply indebted to his constructive criticism and guidance that pushed me to develop into a better researcher, student, and person.

I'm also thankful for my committee members at the University of Idaho, Dr. Aston, Dr. Gao, and Dr. Moberly, for their guidance and support during my graduate studies. Their assistance furthered my development as a researcher.

Additionally, I would like to extend my sincere thanks to all of my collaborators. Specifically, I would like to mention Dr. Kothare, as his extensive knowledge on model predictive control and illuminating conversations improved the quality of my research.

Many thanks to my current and past labmates, Joe, Niko, and Ben. Particularly, I want to recognize the assistance that I received from Joe during my first year, as well as the stimulating discussions and productive coffee trips in my second year.

I would like to express my appreciation for Pango, and the inspiration he would occasionally provide by waking me up at 3 in the morning with his standard feline antics.

Finally, I'd like to thank my father, family, and friends for their love and support.

*To my grandparents, who taught, encouraged, inspired, supported, and loved me for the
length of this journey and beyond*

CHAPTER 1

INTRODUCTION

Cardiovascular diseases remain a leading cause of mortality worldwide despite several pharmaceutical interventions. Cardiac arrest is the leading cause of death in the United States with a 82% mortality rate (Benjamin *et al.*, 2018). Stroke is the leading cause of death worldwide and costs \$40.1 billion per year to the United States healthcare system (Ovbiagele *et al.*, 2013). Healthcare costs associated with heart failure reached \$31 billion (2012 dollars) (Savarese and Lund, 2017) and due to the aging population, this number is projected to double in the next two decades, reaching \$70 billion by 2030 (Heidenreich *et al.*, 2013). Despite various treatment options, many people die of heart failure every year (Savarese and Lund, 2017). Hypertension is the leading risk factor of cardiovascular diseases, and poses a significant risk to the healthcare system with 75 million people affected in the United States. Of this population, 15-18% have treatment resistant hypertension rendering pharmaceutical approaches obsolete. These examples serve as a demonstration of both the financial burden, as well as the widespread health concerns of the population that result from cardiovascular diseases. There are multiple efforts underway that are investigating vagus nerve stimulation (VNS) for the treatment of the aforementioned diseases, specifically cardiac arrest, heart failure, myocardial infarction (heart attack), atrial fibrillation, stroke, and hypertension. The focus of this thesis is to develop data-driven modeling and control approaches to facilitate the development of next generation bioelectric therapies for cardiovascular diseases. Specifically, this thesis will investigate deep learning approaches to map the VNS parameters to the cardiac physiology and develop model-based optimal control approaches to control the heart rate and the blood pressure by optimizing VNS parameters.

1.1 VAGUS NERVE STIMULATION

Vagus nerve stimulation (VNS), also called vagal nerve stimulation, is an electroceutical therapy that utilizes electrical stimulation to activate vagal nerves to achieve a desired behavioral response in the autonomic nervous system (Howland, 2014). A stimulation paradigm for VNS is characterized by an electrical pulse that is described by the pulse width, pulse amplitude, and pulse frequency. Originally VNS was developed as a treatment strategy for epilepsy (Healy *et al.*, 2013; Rychlicki *et al.*, 2006), and was later approved as a treatment strategy for treatment-resistant depression (Akhtar *et al.*, 2016). Current research has investigated utilizing VNS as a treatment strategy for arrhythmia, heart failure, hypertension, and gastric motility disorders, emphasizing diseases that are drug-resistant, or diseases that lack a current pharmaceutical therapy. There remain several challenges with the application of this electroceutical therapy when considering the design of applied stimulation pulses. These difficulties are highlighted by different results in several clinical trials that sought to evaluate the efficacy of VNS to treat heart failure in human subjects (Asad and Stavrakis, 2019). A potential solution to address these challenges involves the incorporation of feedback control to autonomously select the optimal stimulation parameters based on physiological measurements. Due to the complex nature of the cardiovascular system, and more broadly biological systems, the difficulty of modeling becomes a limiting factor in the controller design.

1.1.1 *History of Vagus Nerve Stimulation*

James Corning pioneered the first developments of VNS in the late 19th century in an effort to cure epilepsy (Yuan and Silberstein, 2016). Perhaps due to the limitations of electrical technology, his efforts did not lead to any therapeutic findings, and it wasn't until 1952 when an antiepileptic effect was demonstrated in cats (Yuan and Silberstein, 2016). These results were further validated in 1985, where VNS was demonstrated to terminate seizures in canines (Zabara, 1985). These discoveries prompted investigations into human trials and the first human implanted VNS device was accomplished in 1988 (Penry and Dean, 1990). However, it took a few more years

before a human implanted VNS devices demonstrated notable efficacy, and in 1994 a study found a 47% reduction in seizures in epileptic patients (Ben-Menachem *et al.*, 1994). This study prompted the FDA to approve VNS for the treatment of epilepsy in 1997, and for the treatment of depression in 2005 (Schachter *et al.*, 2000; Bonaz *et al.*, 2016). Since those approvals, there have been many efforts in applying VNS to other diseases. Specifically, VNS is under investigation for inflammatory bowel disease, rheumatoid arthritis, treatment resistant epilepsy, autoimmune disorders, asthma, PTSD, and cardiovascular diseases. While most of these investigations are still in the pre-clinical trial phase, there are encouraging findings. Focusing on a brief history of VNS investigations into cardiac functioning, canine models were used to demonstrate bradycardia (slowing of the heart rate) (Sarnoff *et al.*, 1960) in 1960, and 15 years later, canine models were used to demonstrate protective effects against ventricular fibrillation (Kolman *et al.*, 1975).

1.2 AUTONOMIC NERVOUS SYSTEM

The autonomic nervous system is primarily involved in conducting the vegetative functioning of the body, and consists of three different subsystems of nerves: the sympathetic nervous system, the parasympathetic nervous system, and the enteric nervous system. These different classifications are based on the nerve's anatomical location of origin. The sympathetic nervous system originates in the thoracolumbar region, or the middle of the spine. The parasympathetic nervous system primarily originates in the central nervous system, or the brain. Finally there is the enteric nervous system which consists of the nerves located in the gastrointestinal tract. Considering the behavioral context, the sympathetic pathway is active during high stress events (i.e. "fight or flight"), and in contrast, the parasympathetic pathway is primarily active during baseline functioning to preserve homeostasis (i.e. "feed and breed"). Regardless of a nerve's specific classification, all neural projections have bidirectional communication with the organs they innervate. When a neuronal fiber carries information from a target organ to the central nervous system, it is classified as an afferent connection. Conversely, when a neuronal fiber carries information from

the central nervous system to a target organ, it is regarded as an efferent connection. Together, this neural machinery allows for an organism to dynamically modulate the functioning of internal organs on a coordinated and global scale to meet requirements imposed by the environment, while simultaneously sensing environmental changes. In general, the sympathetic nervous system can be thought of as countering the parasympathetic nervous system. Primarily, this thesis is concerned with the effects mediated by the parasympathetic system.

1.2.1 *Vagus Nerve*

In the autonomic nervous system, the vagus nerve is the primary source of parasympathetic input to many organs and has been classified as the "wanderer nerve", which corresponds to the large area traversed by the vagus nerve (Rosas-Ballina *et al.*, 2011). The vagus nerve innervates the heart, lungs, stomach, pancreas, small intestine, and large intestine, thus allowing for the modulation of behavioral responses to external and internal stimuli for these organs as illustrated in Fig. 1.1. Additionally, Fig. 1.1 describes behavioral responses mediated through vagal nerve action on several organs, and more broadly describes the action of the sympathetic and parasympathetic nervous systems. With such a broad role in the functioning of organs, the vagus nerve can be viewed as a regulator capable of initiating coordinated muscle contractions to provide an organism with the ability to carry out vegetative functions required for survival. Some examples of coordination that involves vagus nerve activation includes functions such as digestion, heart rate, respiratory rate, coughing, sneezing, swallowing, and vomiting (Babic and Browning, 2014). Following the activation of vagal fibers, effects are mediated through acetylcholine release which binds to muscarinic or nicotinic receptors and initiates muscle contractions (Bonaz *et al.*, 2016). Of particular interest for disease therapies, the activation of $\alpha 7$ nicotinic receptors leads to diminishing the inflammatory response from the sympathetic nervous system (Goverse *et al.*, 2016). These effects can be more broadly described as a modulation of the internal inflammatory reflex (Borovikova *et al.*, 2000; Tracey, 2002), which occurs through the anti-inflammatory cholinergic pathway. There are other mechanisms through which the vagus nerve attenuates inflammatory responses,

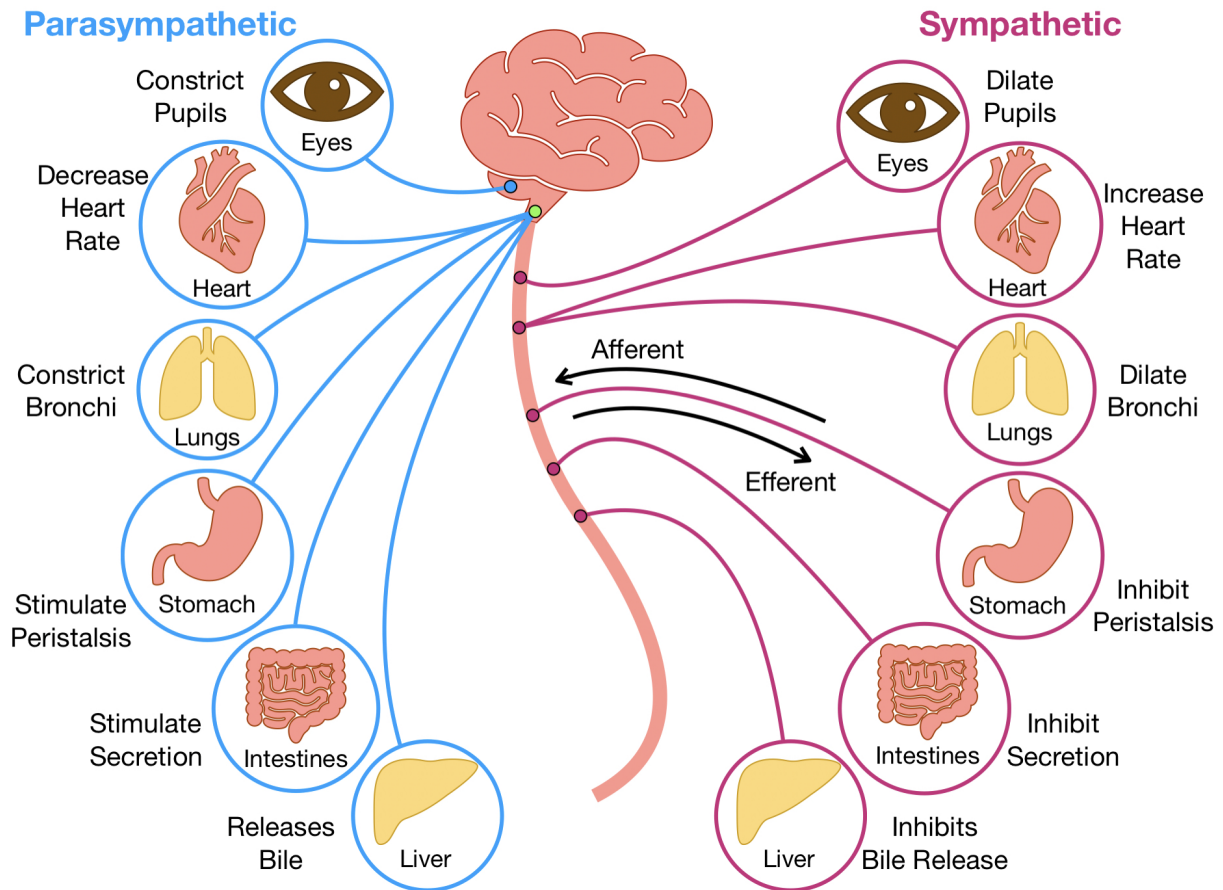


FIGURE 1.1: A diagram of the parasympathetic (blue) and sympathetic (magenta) projections in the autonomic nervous system and their physiological effects on different organs. The vagal nerve projections (indicated by a green dot) innervate the respiratory system, the cardiovascular system, the gastrointestinal tract, and the liver. Activation of the vagus nerve can lead to bronchi constriction, slowing of heart rate, peristalsis, and bile secretion. Conversely, the sympathetic system opposes these actions. Afferent and efferent directions are shown in black.

such as the hypothalamic-pituitary-adrenal axis, and the splenic sympathetic anti-inflammatory pathway (should the vagus nerve stimulate sympathetic innervations). These anti-inflammatory responses are likely responsible for the therapeutic effects demonstrated by VNS for several disease pathologies. Of interest for this thesis are the therapeutic effects mediated by the vagus nerve on the cardiovascular system.

Focusing on the VNS effects in the cardiovascular system, the vagal nerve fibers provide the primary parasympathetic input to the heart. The vagus nerve can be found innervating the heart at the atria, the sinoatrial node, the atrioventricular node, the ventricular myocardium, and the ventricular conduction system. To highlight specific regions related to my thesis, the activation of vagal fibers at the sinoatrial node lead to a slowing of the heart rate. In the pulmonary region, activation of the vagal nerve causes vasodilation which leads to a reduction in blood pressure. Together, these effects demonstrate actions that allow the vagus nerve to modulate the physiological behavior of the cardiac system. The cardiac feedback loop with the central nervous system is shown in Fig. 1.2 A, where the vagus nerve is part of the parasympathetic fibers, and whose actions are opposed by the sympathetic pathway. In cardiovascular disease pathology, the balance between sympathetic and parasympathetic pathways is often disrupted leading to abnormal functioning of the cardiovascular system, illustrated in Fig. 1.2 B. In particular, an overactive sympathetic pathway often predisposes individuals to heart failure, and heart arrhythmia (Kishi, 2012). By stimulating the vagus nerve, the firing pattern of the vagal nerve fibers is altered, and thus the neuronal signaling is altered, corresponding to a change in organ functioning. These changes in organ functioning can be observed through physiological measurements. For example, in the case of hypertension, the vasodilation effect from vagal activation can be used to counter the abnormally high blood pressure in hypertensive patients.

1.3 NEED FOR CLOSED-LOOP VAGUS NERVE STIMULATION

With the exception of heart failure, the research progress of VNS effects on cardiovascular diseases remain in an early pre-clinical phase. Nonetheless, in rats with

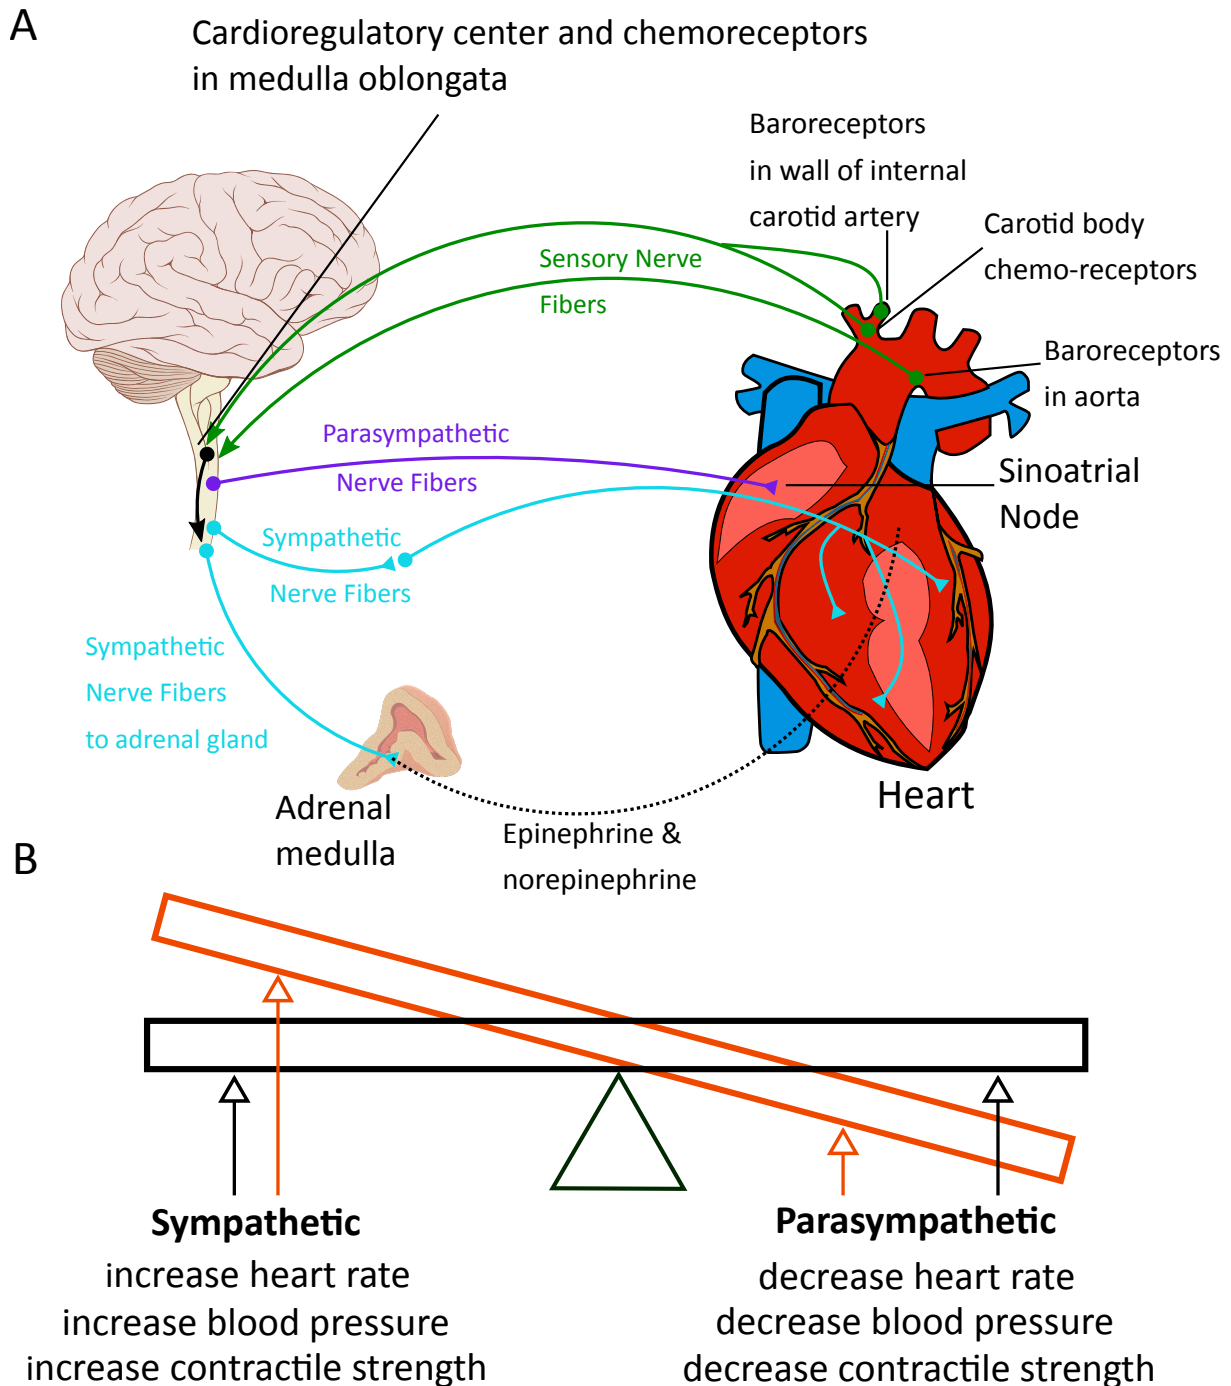


FIGURE 1.2: **(A)** A diagram of the feedback loop between the cardiovascular system, and the central nervous system. Baroreceptors in the heart send information via afferent sensory nerve fibers to the medulla oblongata. The heart receives signaling from parasympathetic and sympathetic nerve fibers whose signals come from the cardioregulatory center in the central nervous system. Additionally depicted is the hypothalamic-pituitary-adrenal pathway at the bottom where sympathetic nerve fibers innervate the adrenal gland and cause epinephrine and norepinephrine release to the heart. **(B)** In black, sympathetic and parasympathetic pathway influence on physiological behavior in the cardiovascular system is depicted as a balance. In orange, the hyperactive sympathetic nervous system and diminished parasympathetic tone associated with cardiovascular diseases is depicted.

hypertension, 4 weeks of VNS therapy was shown to improve the mean arterial blood pressure and the number of arrhythmic episodes in rat subjects (Annoni *et al.*, 2015). In stroke, the therapeutic role of VNS has been shown to reduce stroke volume by 50% through limiting the inflammatory response and significantly improving the clinical outcomes, thus leading to a faster recovery (Ay *et al.*, 2011; Cai *et al.*, 2014). In rats with heart failure, VNS was shown to be efficacious for the treatment of heart failure. Specifically, VNS was shown to increase resuscitation success rates when applied with cardiopulmonary resuscitation (CPR), while simultaneously reducing the number of shocks required for resuscitation, and increasing the 72-hr survival rate (Sun *et al.*, 2018). In an effort to increase the efficacy of VNS, a recent study demonstrated the ability to selectively activate specific fiber bundles of the vagus nerve in a rat, and this precision corresponds to specific control of physiological effects when applying VNS (Plachta *et al.*, 2014). These studies exemplify the promising future for VNS as a modern healthcare solution to cardiovascular diseases.

VNS is most advantageous when pharmaceutical therapies fail, evidenced by the application of VNS on treatment-resistant depression (Akhtar *et al.*, 2016), as well as drug resistant epilepsy (Rychlicki *et al.*, 2006), and medically intractable epilepsy (Healy *et al.*, 2013). Pharmaceutical resistance may result from an inherent lack of physiological response following administration of drug treatment. Potentially, long-term administration of pharmaceutical drugs can lead to a diminished efficacy of treatment, thus requiring different therapeutic approaches for effective treatment. Importantly, VNS can be used in parallel with pharmaceutical interventions (Plachta *et al.*, 2014). In numerous studies across different animal models, VNS parameter selection has demonstrated an ability to affect different cardiac biomarkers, such as blood pressure and heart rate. Currently, VNS parameters are determined through manual titration in an open-loop configuration, as was used in the three clinical trials that investigated VNS for the treatment of heart failure (Gold *et al.*, 2016; Zannad *et al.*, 2015; Premchand *et al.*, 2014). These clinical trials reached different conclusions regarding the efficacy of VNS, potentially due to the different operating regimes for each trial, thus leading to the suggestion of computing optimal VNS parameters in future trials to clearly evaluate the efficacy of VNS (Asad and Stavrakis, 2019).

Vagus nerve stimulation can adapt with the individual over time by dynamically altering the VNS parameter selection to consistently achieve a safe and efficacious response. However, to achieve such a performance from an implanted VNS device, computing optimal stimulation parameters from physiological feedback is likely required (Asad and Stavrakis, 2019). Closed-loop control and autonomous selection of VNS parameters would provide a monumental step in the effort of therapeutic VNS. This effort would likely increase the efficacy of VNS while addressing some of the challenges associated with the manual titration method currently used.

1.3.1 *Nervous System Control Approaches for Cardiac System*

Several studies have examined methods for optimization of VNS parameters in a closed-loop control sense for controlling the heart rate of sheep (Ugalde *et al.*, 2015; Romero-Ugalde *et al.*, 2017), pigs (Tosato *et al.*, 2006), dogs (Zhang *et al.*, 2002), and rats (Greenwald *et al.*, 2016). Here, majority of the designed controllers followed a proportional-integral formulation, and in one case employed a state space transition model (Romero-Ugalde *et al.*, 2017) for calculation of the optimal VNS parameters. Importantly, none of these controllers studied optimizing multiple stimulation parameters, nor did they consider the possibility of different stimulation location sites as in the experimental setup of (Plachta *et al.*, 2014). However, a recent study developed a cardiac model of a rat heart with the influence of VNS and used a model predictive control (MPC) framework to optimize multiple VNS parameters (pulse width and pulse frequency) at multiple locations to control the heart rate and blood pressure simultaneously (Yao and Kothare, 2020).

In the case of (Romero-Ugalde *et al.*, 2017), where a state-space model was used for control, the authors suggested a significant improvement by altering their controller design to adapt the state-space model in such a way that allows for changes on both a long-time scale, as well as a short time scale. With one recent exception (Yao and Kothare, 2020), there have been no control attempts examining optimization of VNS parameters in the multi-location experimental setup in a true closed-loop control sense, despite its advantages as a VNS therapy (Plachta *et al.*, 2014). With an exception, none of the previous studies in optimizing VNS parameters for cardiac

system considered multiple outputs of the cardiac physiology or multiple locations of activating the vagus nerve in the controller formulation. Demonstrating this multi-output control capability in controller design would be valuable in extending the therapeutic applications of VNS, and may improve the efficacy of VNS by controlling multiple biomarkers. A noteworthy challenge in (Yao and Kothare, 2020), involved the computational expense of the developed physiological model which hampered its application in real-time control. To overcome this challenge, the authors developed a reduced-order model for the real-time control application. The reduced-order model relied on state estimation techniques to approximate the state of the system following a measurement of the heart rate and blood pressure from the full physiological model, which introduced a significant offset into the controller's performance when driving the physiological outputs to the desired set-points.

In this thesis, I investigate a data-driven approach for modeling and control of the cardiac system by combining machine learning methods with a model-based predictive control (MPC) framework. For my investigation, the proof of concept was tested in the ideal case of a physiological model of a rat heart, just as in (Yao and Kothare, 2020). The problem was formulated in a way that is easily extendable and compatible with the experimental data. In particular, this approach provides a considerably different problem formulation from the other controllers, leading to novel advantages. Specifically, the advantages of this approach include using any measurable biomarkers, considering multiple stimulation locations, and controlling multiple physiological outputs. Regarding the discussion from (Romero-Ugalde *et al.*, 2017), this approach includes a way of accounting for different timescales of internal dynamics through the use of a long short-term memory (LSTM) neural network. This approach can also be used to address the challenge of formulating a reduced-order model, as in (Yao and Kothare, 2020). More broadly, this approach can be applied to systems that prove difficult to model, which is particularly advantageous in biological systems.

1.4 MACHINE LEARNING APPROACHES

Due to the complex nature of the autonomic nervous system, both exemplified by the current modeling literature and previous control attempts, it remains a challenge to model the influence of VNS on the physiological behavior. A reasonable approach to model these unknown and complex nonlinear interactions involves utilizing data-driven modeling techniques. To incorporate data-driven models in control, model-based control schemes, such as MPC, can be utilized. In this section, I cover the general structure of neural networks, machine learning techniques for time-series modeling, and model predictive control.

1.4.1 *Artificial Neural Networks*

Artificial neural networks (ANNs) follow a generic structure outlined in Fig. 1.3. Data that represents the input is provided to the first layer of nodes, whose output is computed by the product between the input data values and a weight evaluated through a nonlinear function. This layer output can be fed as input to another layer of neurons. The last layer in the neural network leads to the neural network's outputs.

Artificial neural networks are a common modeling technique used to develop data driven models, and have been used in accomplishing several complex tasks. Through the universal approximation theorem, these models are guaranteed to approximate any nonlinear function provided they are constructed with a reasonable size (Hornik *et al.*, 1989). There are three main learning classifications for ANNs based upon the formulation of their input-output data pairings: supervised learning, unsupervised learning, and reinforcement learning. In the supervised learning case, the input is fed to the neural network with labels, and the error between the network predictions and labels are computed according to a provided cost function. Unsupervised learning is the same approach as supervised learning, however labels are not provided to the model, and is typically used in classification contexts where the network develops its own representation of the data. Reinforcement learning is related to control, and involves providing the neural network a target, a reward, and a penalty. The goal of reinforcement learning is to discover the policy function that allows for maximization

of the long-term reward, and requires extensive interaction with the controlled system. In this thesis, machine learning applications were purely formulated using a supervised learning approach.

A considerable challenge for the implementation of neural networks involved updating the neural network's parameters during the training phase. For the vast majority of widely used networks today, this problem has been solved through backpropagation (Rumelhart *et al.*, 1985). In backpropagation, the derivative of the network weights with respect to the cost function is calculated, and this derivative is used to update the neural network's weights using gradient based optimizers. The Newton step method serves as an example of a simple optimizer. In this case, a fixed step size is multiplied by the negative of the gradient, which guarantees that the loss value will decrease, and the resulting product is added to the current network's weight values to finish the update. Since gradient based methods are used, loss functions and activation functions must be differentiable.




FIGURE 1.3: A diagram of a basic feedforward network that takes 3 inputs, and computes 2 outputs with two hidden layers of size 5 and 4. The size of the weight matrices between each layer are noted at the bottom. Activation functions (not shown) are applied at the end of each hidden layer as well as the output layer.

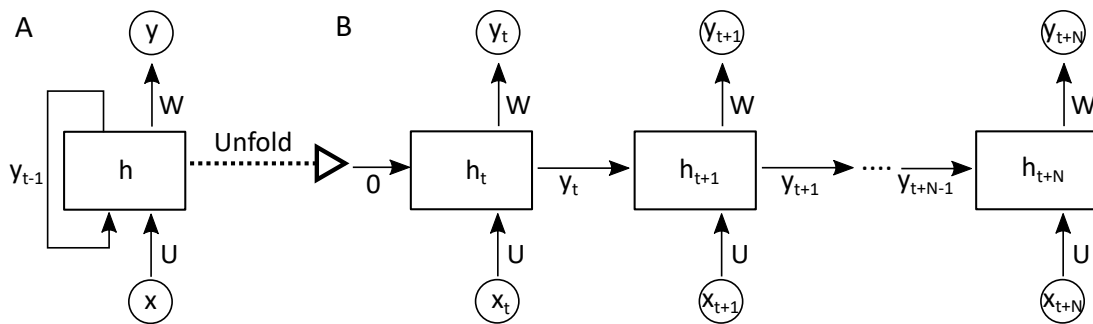


FIGURE 1.4: Diagram of a simple recurrent neural network (RNN). **(A)** Shows the input x passed to the hidden layer and transformed by weights U , h denotes the hidden state, W denotes the weights that transform the hidden state to the output, y . **(B)** Shows **A** unfolded along the time axis. Consider a sequence of length N , at the first step the hidden state is initialized to 0, after that, the hidden state is passed to generate all the outputs leading up to the final input at $t + N$.

1.4.2 Recurrent Neural Networks

A special class of neural networks termed recurrent neural networks (RNNs), are particularly well suited for sequence or time-series modeling. The structure of a basic RNN is shown by Fig. 1.4, and is highlighted by the inclusion of a self feedback, or recurrence, that stores the information from the previous timestep. The simplest case, or a vanilla RNN, suffered from the exploding or vanishing gradient problems during training. When the gradient explodes, or vanishes, updates to the network weights become unstable and training will not converge, which motivated the design of long short-term memory (LSTM) neural networks (Hochreiter and Schmidhuber, 1997). The LSTM uses several internal gating structures to direct the flow of information and is illustrated in Fig. 1.5. Briefly described, there are two additional states kept in the LSTM referred to as the cell state and the hidden state. The cell state is used to capture information for long temporal horizons, and the hidden state is used to capture information on shorter temporal horizons. The internal gating structures determine how much information from the input is stored, and how much of the current memory is forgotten. Since all of these gates use training parameters, the LSTM learns the gating process for the pertinent dynamics throughout the training process. More recently, a gated recurrent unit (GRU) has been proposed as an alternative to the LSTM. The GRU eliminates the forget gate and the cell state, thus requiring fewer parameters for training and has demonstrated comparable performance to a LSTM on smaller datasets (Cho *et al.*, 2014).

The LSTM has been used for a variety of sequential-based tasks including: natural language processing (Radford *et al.*, 2019), handwriting recognition (Graves *et al.*, 2008), speech recognition (Sak *et al.*, 2014), and forecasting of traffic patterns (Zhao *et al.*, 2017). The LSTM allows for predictions of dynamical systems that contain temporal evolutions that vary over different timescales as the additional gating mechanisms allow for predictions that require information of long-time delays. As mentioned in the previous section, VNS would be a candidate for machine learning approaches due to the current difficulty associated with modeling the cardiac system. A LSTM is an appropriate choice for modeling cardiovascular dynamics due to the

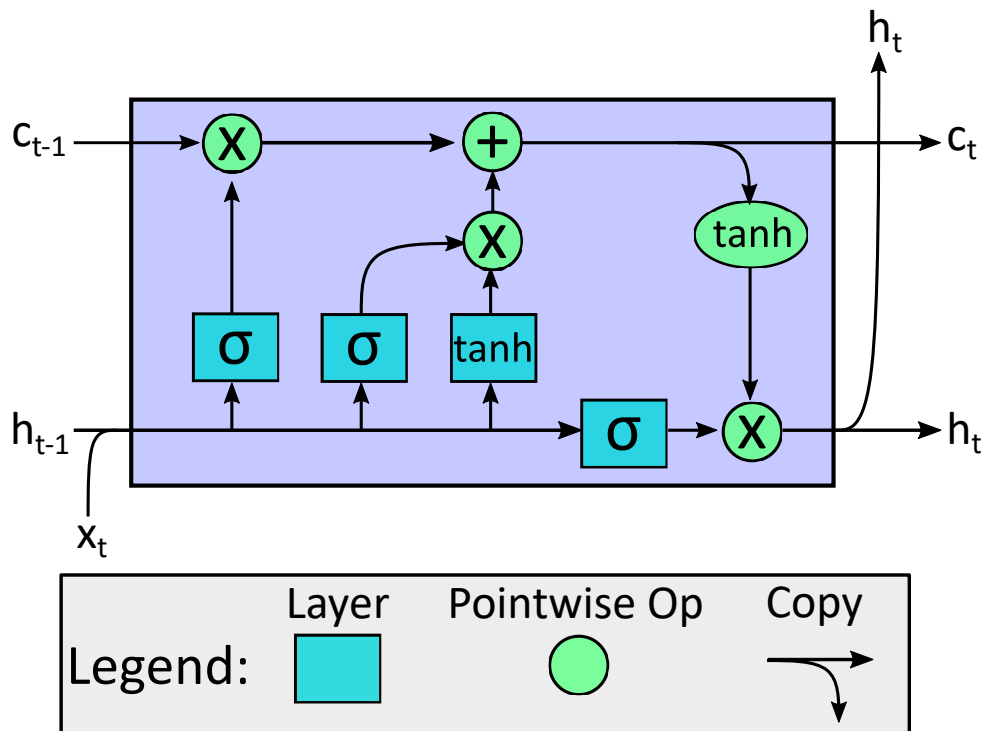


FIGURE 1.5: A diagram of the internal structure of a long short-term memory network for input x_t , cell state c_t , and hidden state h_t . Current timesteps are subscripted with t , and previous timesteps are subscripted with $t - 1$. Activation functions are denoted by σ for the sigmoid function, and \tanh for the hyperbolic tangent function.

time-series nature of the data. Additionally, existing models for the full cardiac system are computationally expensive to run, and therefore infeasible to optimize in the context of real time control implementation. Extending the applicability of LSTMs in this direction, a LSTM model can be trained to serve as a computationally efficient model for control applications.

Since a LSTM is used extensively throughout this thesis, a mathematical description of the LSTM design is provided. Starting with the four initial gates there is the input gate i_t , the forget gate f_t , the output gate o_t , and the cell update \hat{C}_t , also illustrated in Fig. 1.5. In these equations, the input is denoted by x and has trainable weights denoted by U , the previous hidden state is denoted by h_{t-1} and has trainable weights denoted by W . The sigmoid activation function is denoted by σ , where $\sigma(x) = \frac{1}{1+e^{-x}}$, which together leads to:

$$i_t = \sigma(x_t U^i + h_{t-1} W^i) \quad (1.1)$$

$$f_t = \sigma(x_t U^f + h_{t-1} W^f) \quad (1.2)$$

$$o_t = \sigma(x_t U^o + h_{t-1} W^o) \quad (1.3)$$

$$\hat{C}_t = \tanh(x_t U^g + h_{t-1} W^g) \quad (1.4)$$

Lastly, there are the updates to the two memory states with the cell state denoted by C_t , and the hidden state denoted by h_t .

$$C_t = \sigma(f_t C_{t-1} + i_t \hat{C}_t) \quad (1.5)$$

$$h_t = \tanh(C_t) * o_t \quad (1.6)$$

1.4.3 Model Predictive Control

To design the optimal VNS parameters using the trained LSTM, a model based receding horizon controller (RHC) was used. In model-based RHC, also called model predictive control (MPC), an optimal control strategy that explicitly incorporates a dynamic model of the system as well as constraints is used when determining the

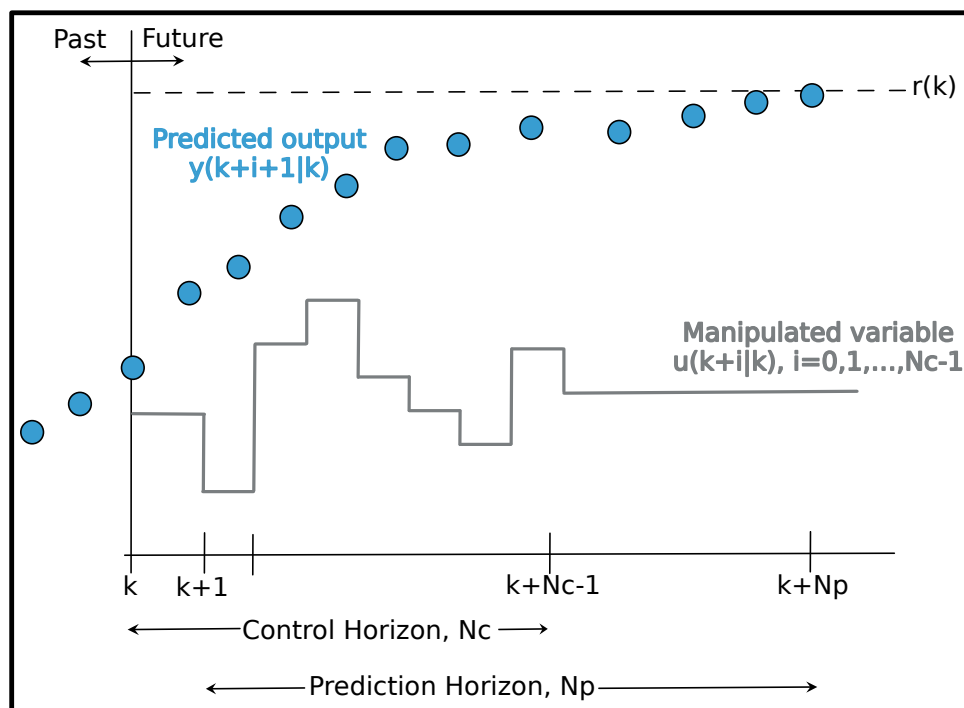


FIGURE 1.6: Prediction and optimal control in MPC.

controller's actions (Kwon and Han, 2006). At each time k , a measurement of the system is obtained and the model of the system is used to predict the future outputs of the system $y_{k+l+1|k}$, $l = 0, 1, 2, \dots, N_p - 1$ as a function of current and future control moves $u_{k+l|k}$, $l = 0, 1, 2, \dots, N_c - 1$. The predictive horizon N_p defines how far into the future the predictions are computed, and the control horizon N_c defines how far ahead the control moves are computed. Fig. 1.6 illustrates the process of computing a prediction and control horizon with a model-based RHC strategy.

Using the predictions from the model, the N_c control moves $u_{k+l|k}$, $l = 0, 1, \dots, N_c - 1$ are optimally computed by minimizing a cost function \mathbb{J}_k over the prediction horizon N_p subject to constraints on the control inputs as well as any other constraints on the internal states and outputs of the system as follows:

$$\min_{u_{k+l|k}, l=0,1,\dots,N_c-1} \mathbb{J}_k \quad (1.7)$$

subjects to constraints on control inputs and the system. A typical quadratic objective cost function may be of the form

$$\mathbb{J}_k = \sum_{l=0}^{N_p-1} [y_{k+l+1|k} - y_r]^T Q [y_{k+l+1|k} - y_r] + \sum_{l=0}^{N_c-1} u_{k+l|k}^T R u_{k+l|k} \quad (1.8)$$

Here, y_r is the output to be tracked. Only the first optimally computed move $u_{k|k}$ is implemented out of m computed optimal moves at time k . At the next time, $k + 1$, new system measurements are obtained and the optimization problem is solved again with the new measurements. Thus, the control and prediction horizon recede by one step as time moves ahead by one step. The measurements at each sampling time provide feedback for rejecting inter-sample disturbances, model uncertainty and noises, all of which cause the model predictions to be different from the true system output.

1.5 THESIS OUTLINE

In Chapter 2, I discuss my work on the data-driven modeling of the rat cardiac system by using a physiological model of a rat heart. I use a variety of recurrent neural networks, and show that the best performance is given by a long short-term memory

network (LSTM). I demonstrate that this LSTM is capable of mapping VNS parameters to the correct physiological response for heart rate and blood pressure. I further show that this LSTM is more computationally efficient than a full-scale model through a 10-fold reduction in simulation time. In Chapter 3, I discuss my work of using the previously trained neural network in a model predictive control framework to control the cardiac physiological variables, such as the heart rate and blood pressure. I apply the designed controller on the full physiological model that was used to generate open-loop data to train the neural network. With this problem formulation, the neural network can be considered as a reduced model of the true system. I show that the controller is capable of achieving little to no offset while reaching several target set points. In Chapter 4, I discuss the future directions of my research contained in this thesis. Specifically, I provide suggestions for investigating mathematical techniques that take a different approach to the control problem and provide advantages for theoretical controller analysis.

CHAPTER 2

MAPPING VAGUS NERVE STIMULATION PARAMETERS TO CARDIAC
PHYSIOLOGY USING LONG SHORT-TERM MEMORY NETWORK

2.1 INTRODUCTION

Predicting physiological behavior with the influence of VNS requires a complex model, as predicting cardiac behavior alone is a challenging task. These challenges revolve around the development and validation of the cardiac model. Often such an approach becomes a challenging task in selecting the correct dynamical equations to govern the cardiac system, and then fitting the parameters of those specific equations. Such tasks can be guided by a deep mechanistic understanding of the system, however this knowledge may be incomplete or lack the details of integration. These differences are highlighted by different definitions of the cardiac system to be modeled, demonstrated by the variety of cardiac models in the literature that range from modeling the individual neuronal cells in a cardiac tissue (Mangoni *et al.*, 2006), to modeling the whole cardiac system as a pump (Suga *et al.*, 1973). There have been some models that incorporate the effects of extrinsic stimulation on the cardiac system, such as simulating an orthostatic response in a human cardiac system (Melchior *et al.*, 1992). However, most models do not include the necessary variables to account for physiological changes mediated through VNS, leaving a challenging task for using these models in VNS parameter optimization. This gap in the modeling efforts may be responsible for the lack of MPC-based controllers in a closed-loop control of the cardiovascular system. Additionally, using fully detailed mechanistic models of the cardiac system suffer the drawback of computational expense which makes their implementation in real-time closed-loop a considerable challenge.

Data driven models, such as neural networks, provide a promising solution to address the challenges discussed above and have been used in a similar context for a Hodgkin-Huxley pyramidal neuron model (Plaster and Kumar, 2019). A major advantage of neural networks is the limited assumptions required for their application

as they do not require underlying assumptions about the system dynamics nor does their application require a specific distribution of data. Recurrent neural networks (RNNs) have been specifically developed to capture time-series or sequential data, which is particularly amenable to the control application. Finally, neural networks can be computationally inexpensive to evaluate when compared to full scale models. Taken together, these features motivate an investigation into applying neural networks to model the effects of VNS on the cardiovascular system.

In this chapter, I leverage deep learning approaches to model the response of physiological variables following the VNS stimulation using the synthetic cardiac data from an existing model of the rat cardiac system (Yao and Kothare, 2020). First, I include a full description of the physiological model used. For increased relevance to experimental systems, the full physiological model was used to generate synthetic data in an open-loop trial format. I then trained a variety of RNN architectures on the collected open-loop data. My results show that a long short-term memory (LSTM) architecture provides the best predictive performance. Further, I find that the computational cost of the LSTM in predicting the mean arterial pressure and heart rate for 100 consecutive cardiac cycles is approximately 10 folds lower than the original physiological model.

2.2 PHYSIOLOGICAL MODEL

The model used for synthetic dataset creation is a cardiac model of a rat heart (Yao and Kothare, 2020). Briefly described, the model consists of three parts: the cardiovascular system, the baroreflex system, and the vagus nerve stimulation device as depicted in Fig. 2.1. Mathematically, the model consists of a system of delayed differential equations with ten internal states and a nonlinear combination of the internal states provides the physiological measurements of the mean arterial blood pressure and the heart rate. The VNS device was incorporated by including three different locations, where each location had two parameters: a pulse width, and a pulse frequency which together characterize the stimulation paradigm. Multiple locations were considered following the experimental setup of (Plachta *et al.*, 2014). A specific behavior worth

highlighting is the model's ability to exhibit increases in the heart rate independent of the changes in the mean arterial blood pressure. A mathematical description of the full physiological model developed by (Yao and Kothare, 2020) is provided for completeness in the following sections.

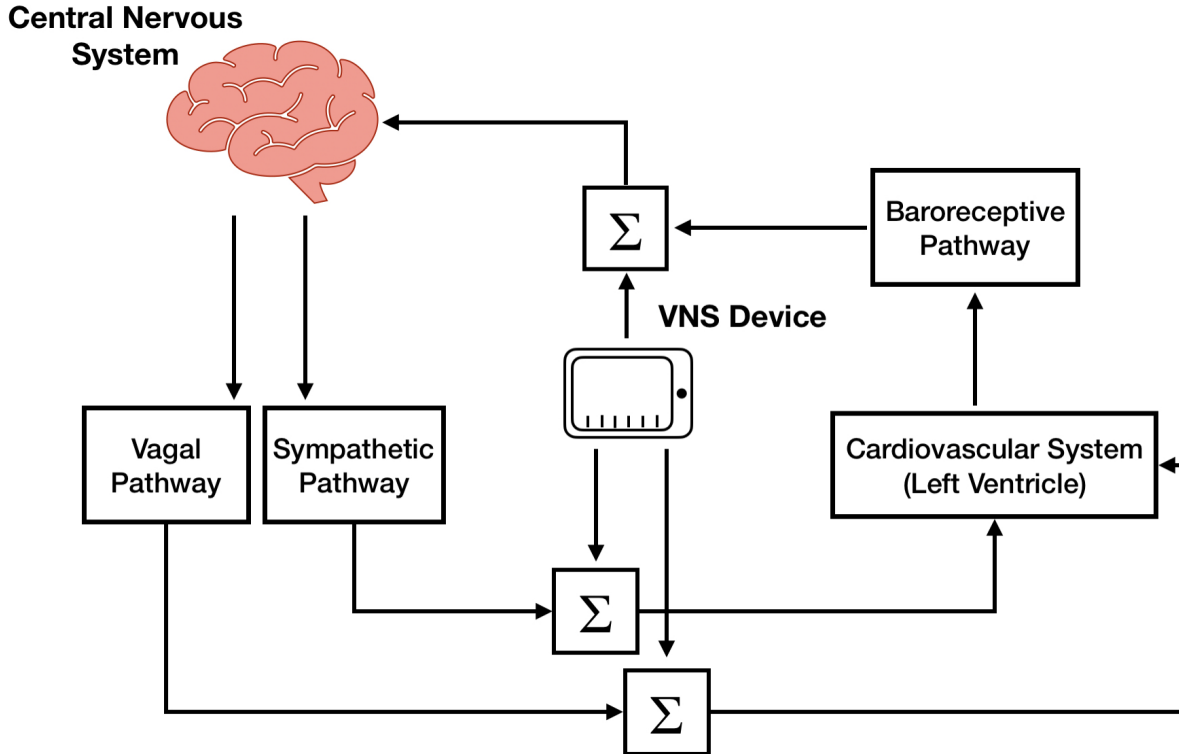


FIGURE 2.1: A diagram that details the influence of vagal nerve stimulation in the full physiological model of the rat cardiac system.

2.2.1 Cardiovascular Model

Design of the cardiovascular system is based on the previously published models (Djabella *et al.*, 2005; Ferreira *et al.*, 2005). Parameters for the cardiovascular system came from (Ferreira *et al.*, 2005) and were then adjusted by the body volume ratio of rats to humans, resulting in values similar to those measured in experimental rats (Pacher *et al.*, 2004).

$$P_i = E_i(V_i - V_{i,d}), \quad (2.1)$$

Briefly, the pressure volume relationship is described by Eqn. 2.1, where the instantaneous blood pressure of compartment i is denoted by P_i , the total volume is denoted by V_i , unstressed volume is denoted by $V_{i,d}$, and the elastance is denoted by E_i .

The elastance, $E(t)$, varies in time following the formulation of (Stergiopoulos *et al.*, 1996):

$$E(t) = E_{max} \left(a \frac{\left(\frac{t_n}{\alpha_1 T}\right)^{n_1}}{1 + \left(\frac{t_n}{\alpha_1 T}\right)^{n_1}} \frac{1}{\left(\frac{t_n}{\alpha_2 T}\right)^{n_2}} \right) + E_{min}, \quad (2.2)$$

where t_n denotes the periodic time, T denotes the cardiac period, E_{max} denotes the end-systolic elastance, E_{min} denotes the end-diastolic elastance, and $a, \alpha_1, \alpha_2, n_1, n_2$ are all dimensionless constants.

The flow, Q , between compliance chambers of the cardiac system is modeled as follows:

$$Q = \frac{P_{in} - P_{out}}{R_{sys}}, \quad (2.3)$$

where P_{in} and P_{out} are pressures on either side of the cardiac system's resistance to blood flow, R_{sys} . By using a mass balance, the change in volume follows

$$\frac{dV}{dt} = Q_{in} - Q_{out}, \quad (2.4)$$

where the change in volume is a difference between the flow in, Q_{in} , and the flow out Q_{out} of the compliance chamber. Finally, there is consideration of an inertial flow out of the left ventricle that is described as follows:

$$\Delta P = L \frac{dQ}{dt}, \quad (2.5)$$

where L denotes the inertance, and ΔP denotes the pressure difference due to the inertial blood flow.

2.2.2 Baroreflex Model

The baroreflex system model comes from (Lau and Figueroa, 2015), and is composed of several parts: the central nervous system, the baroreceptor model, and a modulation of efferent responses. Following sympathetic drive, the modulation of efferent responses can be described by left ventricle systolic elastance (E_{max}) and the cardiac resistance to blood flow (R_{sys}) exhibiting a positive response, while the heart period (T) exhibits a negative response. For the central nervous system, precise interactions are not understood which has led to using a previously developed input-output relationship (Ursino, 1998). The parameters for the baroreceptive model come from (Mahdi *et al.*, 2013), and the parameters in the baroreflex system from (Ursino, 1998). There is an underlying assumption that 100% of the fibers are recruited, which allows for a qualitative match to the experimental data from (Plachta *et al.*, 2014). Using a set of first-order ODEs, a delay, and a logarithmic relationship, the sympathetic efferent pathway is described as follows:

$$\frac{d\theta_{es}}{dt} = -\frac{1}{\tau_{\theta_{es}}}(\theta_{es} - \theta_0) + G_{\theta_{es}} \ln(\max\{f_{es}(t - D_{\theta_{es}}) - f_{es,min}, 1\}), \quad (2.6)$$

where θ_{es} denotes the efferent path variables: heart period T , systolic elastance of the left ventricle E_{max} , and inertial cardiac flow resistance, R_{sys} . In this equation, θ_0 denotes the baseline value in the absence of external input, $f_{es}(t)$ represents the firing rate of sympathetic efferents, $\tau_{\theta_{es}}$ denotes the time constant, $G_{\theta_{es}}$ denotes the gain, $D_{\theta_{es}}$ denotes the delay of the effector, and $f_{es,min}$ denotes the minimum firing rate of the sympathetic efferents. First-order dynamics are used to capture the change in heart period due to the activation of vagal fibers as follows:

$$\frac{dT_{ev}}{dt} = -\frac{1}{\tau_{T_{ev}}}(T_{ev} - T_0) + G_{T_{ev}} f_{ev}(t - D_{T_{ev}}), \quad (2.7)$$

where $\tau_{T_{ev}}$ is the time constant, $G_{T_{ev}}$ is the gain, and $D_{T_{ev}}$ is the delay of the vagal pathway. Effects exerted by the sympathetic drive are assumed to be independent of the effects exerted by the vagal drive for the cardiac period (T). This leads to the following calculation of the heart period:

$$T = T_{ev} + T_{es} - T_0, \quad (2.8)$$

where T_{ev} denotes the vagal (parasympathetic) drive effects on the heart period, T_{es} denotes the sympathetic drive efferents effect on the heart period, and T_0 is the resting heart period.

2.2.3 VNS Device Model

The stimulation device translates VNS parameter selection into neural firing rate changes, with an assumption that the device increases the firing rates of baroreceptive fibers, efferent sympathetic fibers, and vagal fibers. The fiber recruitment due to pulse width is given by:

$$F(P_w^i) = \frac{P_w^i/k_w}{\sqrt{1 + (P_w^i/k_w)^2}}, \quad (2.9)$$

where $i = 1, 2, \dots, n$ is the location index, k_w denotes a dimensionless scaling parameter, P_w^i denotes the pulse width, and $F(P_w^i)$ denotes the fiber recruitment at each location. The change in firing rates due to pulse frequency is given by:

$$\Delta R(P_f^i) = \frac{P_f^i/k_f}{\sqrt{1 + (P_f^i/k_f)^2}}, \quad (2.10)$$

where $i = 1, 2, \dots, n$ denotes the location index, k_f denotes a dimensionless scaling parameter, P_f^i denotes the pulse frequency, and $\Delta R(P_f^i)$ denotes the change in firing rates of each fiber. The change in fiber firing rates leads to the change in physiological

variables, and this aggregate effect of fiber recruitment regarding the selection of the pulse width and the pulse frequency on the firing rates of neuronal fiber type j is described by:

$$\Delta f_j = \frac{G_j}{n} \sum_{i=1}^n \delta_i C_{i,j} F(P_w^i) \Delta R(P_f^i), \quad (2.11)$$

where $i = 1, 2, \dots, n$ is the location index, $j = 1, 2, 3$ indicates the fiber type index, δ_i indicates an on/off of the i^{th} location, $C_{i,j}$ represents the concentration of fiber type j at location i . The gain of each fiber's excitability is represented by G_j , and f_j denotes the final change in the firing rate of the fibers.

2.3 MACHINE LEARNING APPROACH

Using the previously detailed physiological model of the rat cardiac system, open-loop simulations were run for 15,198 different trials with the VNS parameters randomly sampled from a uniform distribution. The bounds for the pulse width and pulse frequency were 0-0.5 ms, and 0-50 Hz, respectively. The three different locations were also randomly activated using a uniform distribution between 0 and 1 where values above 0.5 indicated activating that location. A single open-loop trial was obtained after selecting the VNS parameters and simulating the model for 100 consecutive cardiac cycles. During all 100 cardiac cycles, the physiological outputs of the heart rate, and mean arterial blood pressure were recorded from the model.

Next, the trials were divided into a training, validation, and test set following a 40%, 20%, and 40% split, respectively. For the training and validation sets, the 100 cardiac cycles were sliced into 50 pairs of alternating points with the inputs consisting of the VNS parameters, heart rate and mean arterial pressure. The labels consisted of the heart rate and mean arterial pressure for the next cardiac cycle. This led to an input dimensionality of 8 (heart rate, mean arterial blood pressure, and six VNS parameters), and an output dimensionality of 2 (heart rate, and mean arterial blood pressure). With this problem formulation, all neural networks were trained to make

a single step prediction, where the cardiac cycle serves as the discretionalized step. To assess neural network performance using the test set, the 100 cardiac cycles were sliced into one input (the initial heart rate and mean arterial pressure values with VNS parameters) and 99 labels for the remainder of the trial. The neural network was evaluated on its ability to recursively predict the 99 cardiac cycles, given the initial point. Assessing the network in this way allowed me to evaluate the ability of the neural network to predict recursively, similar to its required application in a control context. Before the data was fed to the network for training, all data was normalized according to Equation 2.12.

$$\hat{x} = \frac{x - \mu}{x_{max} - x_{min}} \quad (2.12)$$

Here, \hat{x} denotes the normalized data point, μ denotes the mean of the training set, x_{max} and x_{min} denote the maximum and minimum values of the training set, respectively. Statistics of the training set were used, so the model did not get any information about the validation or testing datasets. The distributions following normalization for the variables are shown by Fig. 2.2. All of the VNS parameters have the same distribution shape, which is expected as they were all drawn from the same uniform distributions.

A variety of neural network architectures were trained, including a vanilla RNN, a gated recurrent unit (GRU), and a LSTM. The number of layers and inputs were also varied to explore effect of the network architecture on the model predictions. For all models, the hyperbolic tangent function was used, and the output from the recurrent layer was fed to a dense layer with an output size of two to predict the two physiological variables. Throughout the training of all neural networks, the mean squared error loss function and ADAM optimizer were used (Kingma and Ba, 2014). To assess a trained network's performance, normalized mean absolute error was used in the test set. One network was said to be more accurate than another when the mean absolute error on the test set was lower. As a measurement of performance, all trained models were compared to a baseline case that predicts no change in the initial heart rate and mean arterial blood pressure for the length of the trial. This choice of




FIGURE 2.2: A violin plot of the normalized data (unitless) for all variables. HR and MAP denote the heart rate and the mean arterial pressure, respectively. VNS PW indicates the stimulation pulse width at the location number indicated in the label, and F indicates the stimulation frequency at the location number indicated in the label. For example, VNS.PW₁ denotes the pulse width parameter at location 1.

a baseline emphasized a network's ability to capture the dynamical behavior of the physiological model.

2.4 RESULTS

2.4.1 *Network Design Comparison*

The first goal in searching for an architecture involved a wide search of different RNN structures. The performance of different architectures is shown in Fig. 2.3A. Consistently, a vanilla RNN failed to capture the dynamical behavior demonstrated by the higher error value. LSTMs and GRUs contain several memory gating structures that enable them to capture long term temporal dependencies. Given the nature of the differential equation system, I would expect the additional memory gates to perform better on the dataset as demonstrated by the lower error values.

Layers were also varied in all architectures to investigate the influence of this design parameter. Consistently, for all networks, increasing the number of layers did not improve the quality of predictions (see Fig. 2.3B). To the contrary, neural networks with more layers performed worse on the test dataset, which may lie in the problem formulation of using a single step prediction. Shown in Fig. 2.3C, the input neuron size of the recurrent layer was also varied for all networks. Since the network was predicting two outputs, the input neuron size was initially varied in powers of two, holding the number of layers constant (one). From this investigation, the input neuron size for a single layer must be greater than eight. When the input neuron size was increased past 32, predictive performance started to decrease.

I found my optimal trained neural network to be a single layer LSTM with a ten neuron input size. The best trained GRU, also with a ten neuron input size, provided similar performance. The best performing network results are highlighted in Table 2.1. When selecting the best network, I favored networks that were smaller with comparable performance.

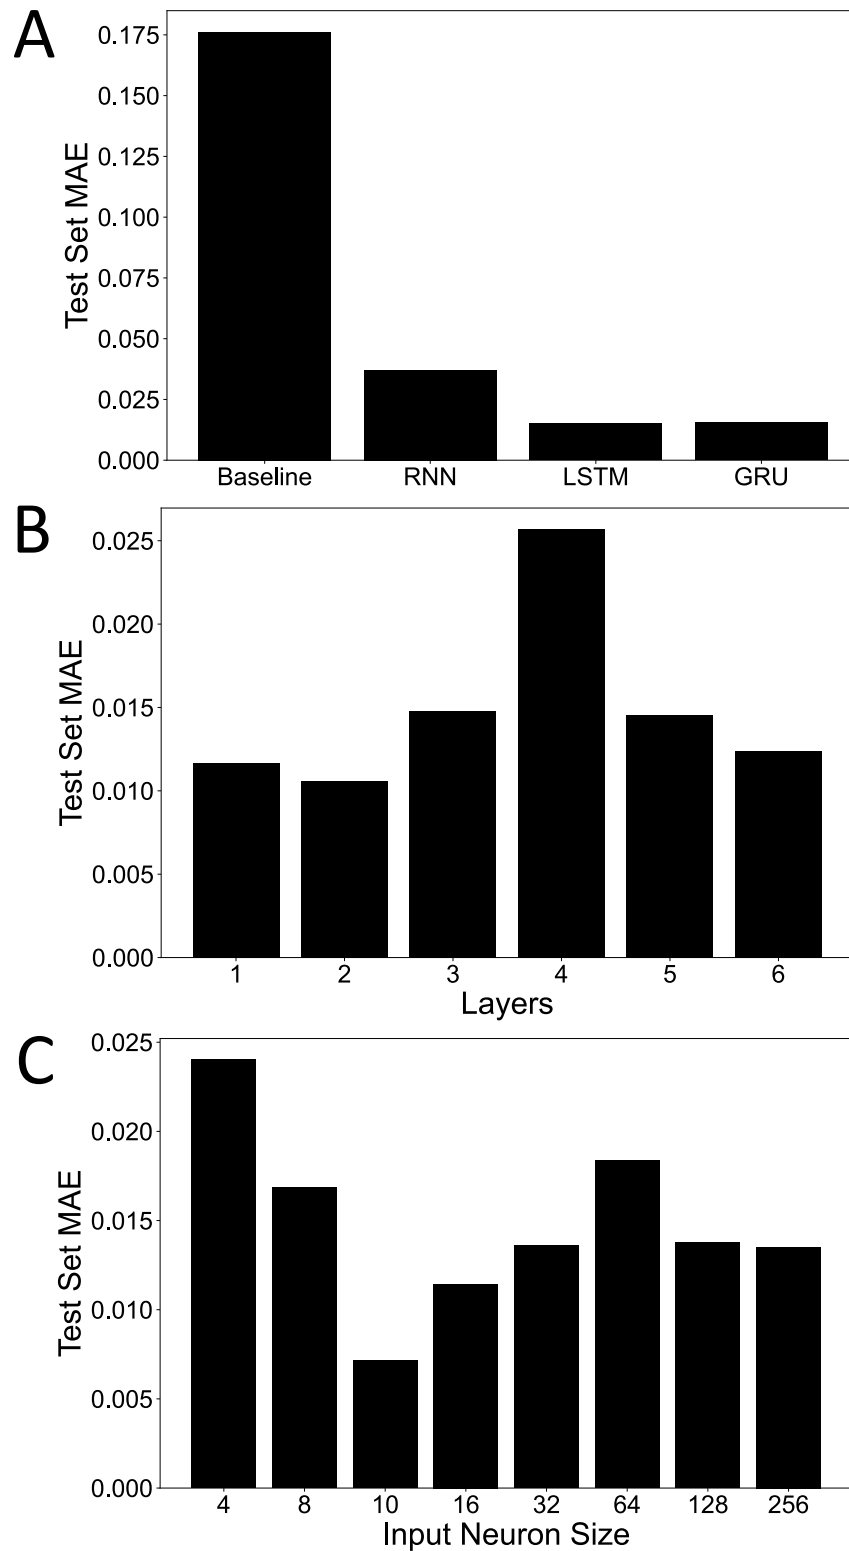


FIGURE 2.3: Comparison of (A) different neural network performances on the test set including the baseline case. (B) The effect of layer size on a neural network's predictive performance on the test set. (C) The effect of input neuron size on a neural network's predictive performance on the test set.

TABLE 2.1: Top Three Trained Architectures

Model Type	Input Size	Layers	Test MAE
LSTM	10	1	0.0072
GRU	10	1	0.0077
GRU	32	1	0.0078

2.5 BEST NEURAL NETWORK DESIGN PERFORMANCE

Using the LSTM with ten input neuron size, a sample prediction from the test set is shown in Fig. 2.4 along with the average error for all cardiac cycles. Note that the model output is fed back to recursively generate the next prediction, however there is not an accumulation of error. While the first few timestep predictions are slightly inaccurate, the steady-state value is still accurately predicted. Thus, the LSTM is capable of performing predictions in a recursive fashion. As shown in Fig. 2.4C, a plot of the error, the mean arterial pressure had a higher error than the heart rate. Examining trials of mean arterial pressure and heart rate, the cause of this discrepancy likely lies in the less smooth mean arterial pressure curves. This assertion is partly confirmed by the heart rate trials which have more smooth curves, and show a more consistent error.

Since the model was able to reasonably predict individual trials, I varied the VNS parameters over time to obtain a complex nonlinear curve to assess the model's performance on a more challenging prediction. These results are shown in Fig. 2.5. For the first ten cycles, VNS at location 1 was active with the pulse width set to 0.21 ms and the frequency set to 21 Hz. The following ten cycles switched to activating VNS at location 3 with pulse width set to 0.09 ms and the frequency set to 9 Hz. Then VNS at location 2 was activated with a pulse width of 0.31 ms and a frequency of 30 Hz for ten cycles. All VNS locations were turned off for 19 cycles, after which location 1 was activated with a pulse width of 0.14 ms and a frequency of 14 Hz for another ten cycles. VNS at all locations was turned off for the final 40 cycles. Again, the model demonstrates a reasonable level of accuracy when tracking the change in the physiological variables. These results support the notion that the trained network

has captured the underlying dynamics of the full differential equation model and is capable of successfully mapping the VNS parameters to the physiological outputs. In the context of closed-loop control, the LSTM demonstrates success and accuracy in the task of predicting the cardiac dynamics subject to VNS.

2.5.1 Computational Speed

TABLE 2.2: Comparison of Computational Speed. Tests were done using an Intel(R) Core i7-9700 CPU 3.00 GHz with 16.0 GB of RAM.

Model Type	Cardiac Cycles	Time (sec)
Full Model	100	19.99
LSTM 10	100	2.10
GRU 10	100	2.14
RNN 10	100	2.09

Having now showed that the neural network model is capable of predicting complex curves generated by the differential equation model, I assessed the computational speed of the LSTM model and the full differential equation model. Additionally, I included a vanilla RNN and a GRU of similar sizes for a complete comparison. These results are summarized in Table 2.2. Comparatively, there is a clear decrease in the computational time when using a LSTM or any other RNN. While these results are expected, such a marked decrease in computational time supports the applicability of a LSTM-based model in designing a model-based optimal closed-loop control strategy for controlling the cardiac system.

2.6 SUMMARY

My goal in this chapter was to develop a data driven approach capable of reproducing nonlinear dynamics for modeling the heart rate and the mean arterial pressure in response to the VNS parameter selection, and demonstrate this approach on a synthetic dataset. I found that both a GRU and LSTM were capable of accomplishing such a task, with a LSTM exhibiting slightly better performance on the dataset used.

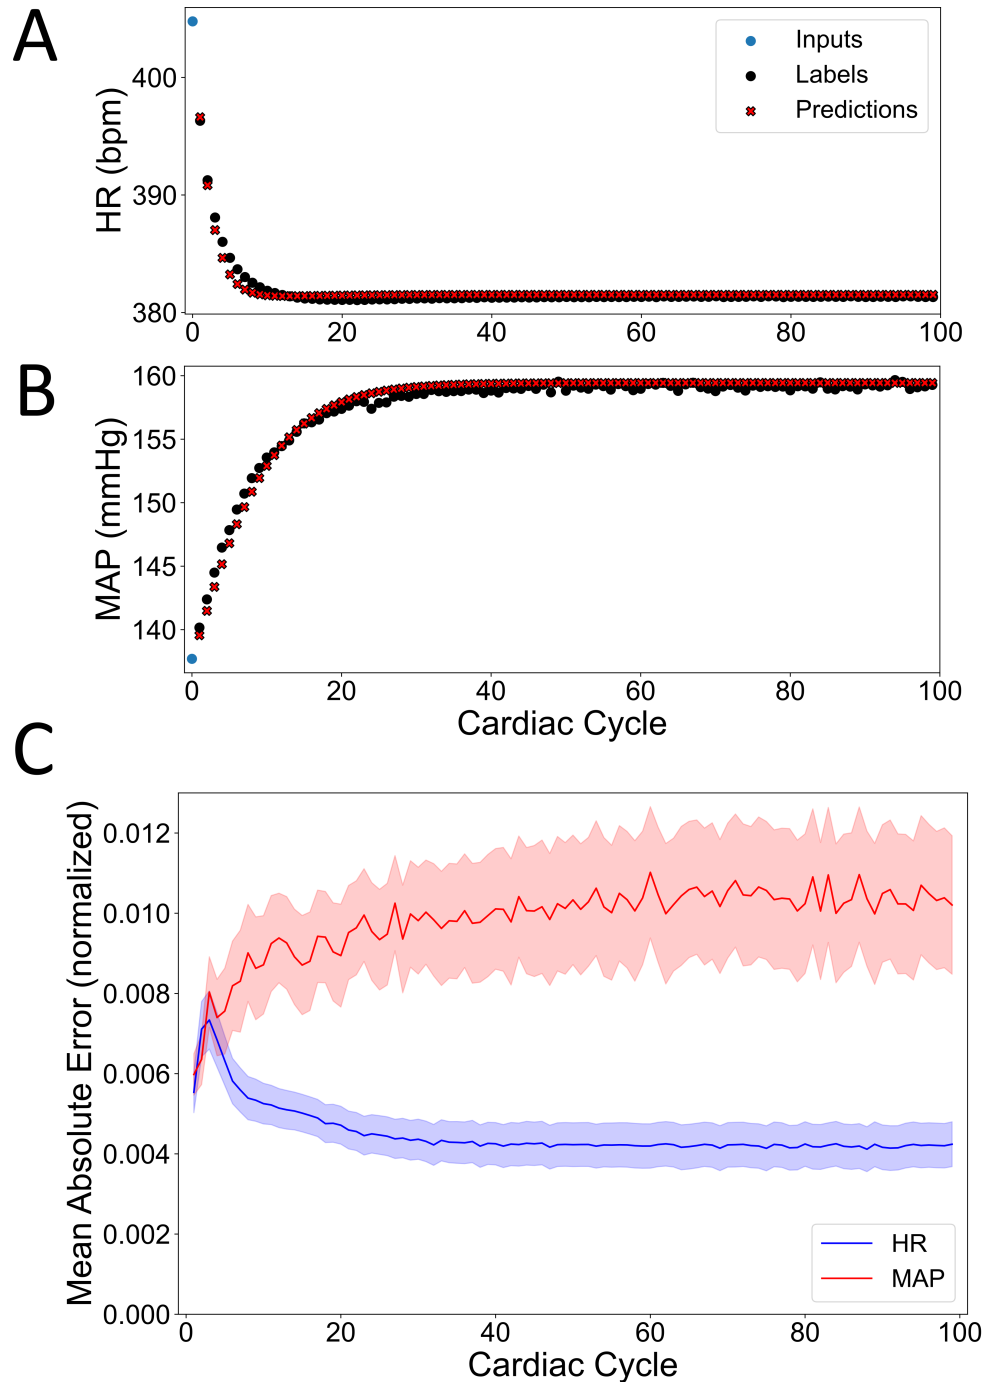


FIGURE 2.4: Example of the trained LSTM model performance on a trial from the test set for predicting the (A) heart rate (HR) and the (B) mean arterial pressure (MAP) for one selection of the VNS parameters. The LSTM is given the input (blue), and is asked to recursively predict the next 99 labels (black) with the LSTM predictions shown in red. (C) Normalized mean absolute error over the entire test set for the heart rate and the mean arterial pressure is shown, along with the standard deviation shaded around the curve.

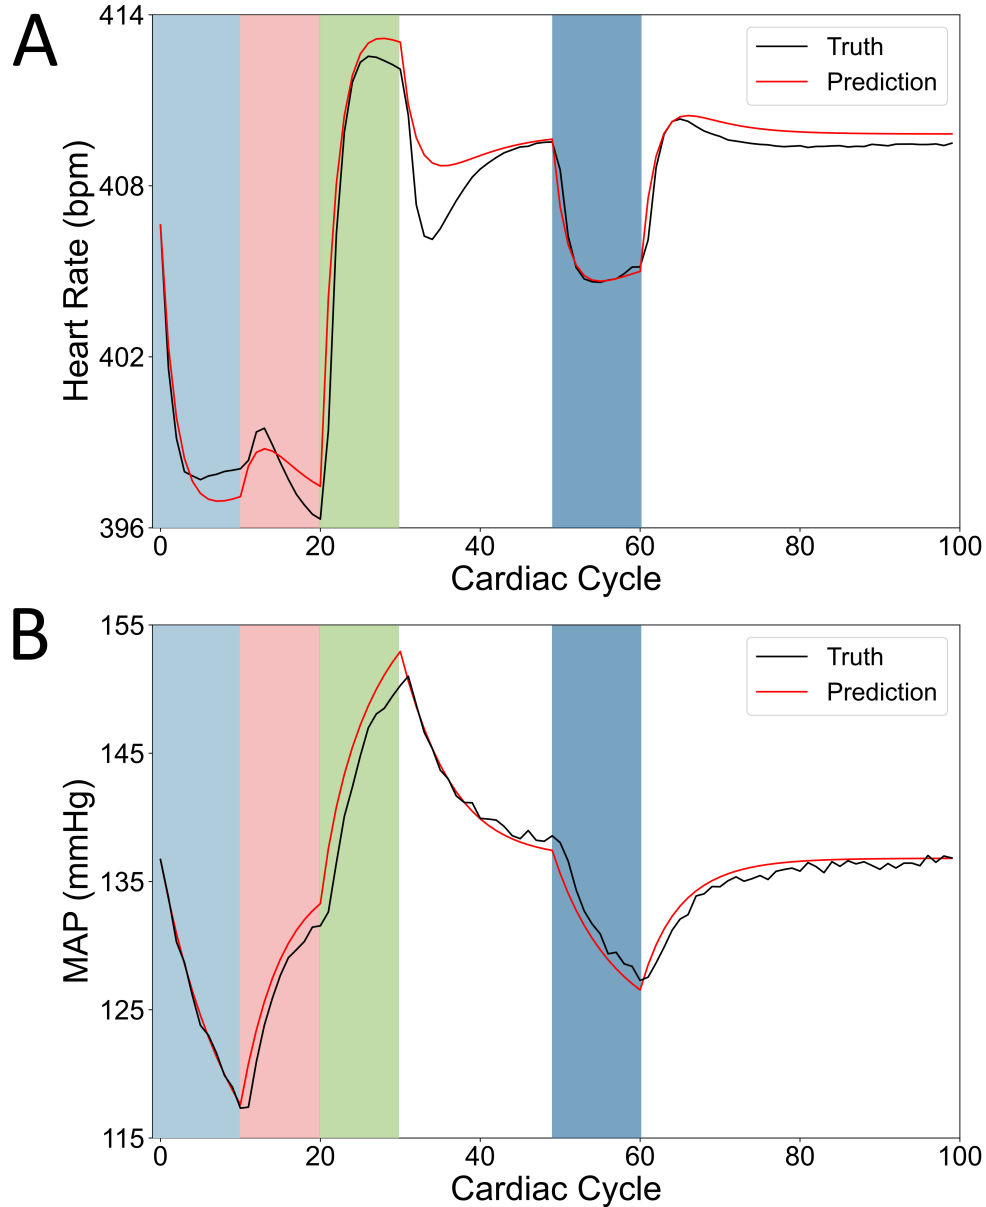


FIGURE 2.5: Comparison between the LSTM model and the full model for the (A) heart rate and the (B) mean arterial blood pressure (MAP) for 99 simulated cardiac cycles. The black line shows the output of the full differential equation model (labeled Truth), while the red line shows the predictions from the LSTM model (labeled Prediction). Cycles 1-10: location 1 with pulse width 0.21 ms and pulse frequency 21 Hz (light blue). Cycles 11-20: location 3 with pulse width 0.09 ms and pulse frequency 9 Hz (light red). Cycles 21-30: location 2 with pulse width 0.31 ms and pulse frequency 30 Hz (light green). Cycles 50-60: location 1 with pulse width 0.14 ms and pulse frequency 14 Hz (dark blue). All other cycles had no locations active.

Further, I showed that both the GRU and LSTM took significantly less computational time than the full physiological model.

This chapter could be extended by applying this approach to an experimental dataset. Such an investigation would provide a more concrete demonstration of its feasibility in the therapeutic context. Of particular interest in this application would be the neural network's ability to account for animal specific variation. If successful, this could provide another advantage of using a neural network over a mechanistic model as most models do not include the variation observed in the experimental data. In the context of a closed-loop control of the cardiac system via VNS, such a feature would be desirable for therapeutic applications. I anticipate that this application would likely require a different neural network architecture (number of layers, input size, etc.), but suspect that a LSTM and GRU would still be preferable to the vanilla RNN.

Another extension to the LSTM model involves implementing a controller to work with the LSTM in a model predictive control framework to find the optimal stimulation parameters to reach a target set-point. This study could elucidate the differences in controller performance with the increased computational speed of a neural network. This control approach utilizing RNNs has been shown to control a hypothetical chemical plant (Wu *et al.*, 2019). Both extensions mentioned here seek to demonstrate an increased relevance to the overall goal and address challenges related to autonomous closed-loop control of VNS.

CHAPTER 3

DATA DRIVEN CONTROL OF VAGUS NERVE STIMULATION FOR THE
CARDIAC SYSTEM

3.1 INTRODUCTION

As discussed in Chapter 2, modeling the cardiac system of any animal is a challenging task leading many investigations to utilize a proportional-integral approach for control. However, proportional-integral controllers require extensive tuning and a square system for multiple-input multiple-output control. A square system requires that the number of manipulated variables must equal the number of controlled variables, which is often not the case in the design of stimulation parameters for vagus nerve stimulation. As an example, one study has shown that simultaneously modulating multiple VNS parameters is required for optimal VNS performance (Ojeda *et al.*, 2016). In a deviation from proportional-integral controllers, one group developed a state transition model for controlling the cardiac system (Romero-Ugalde *et al.*, 2017). While this specific approach shares some similarities to a data-driven approach, the resulting controller design is still limited to controlling one physiological output. The authors note that applying the control scheme in a strategic way that allows for adaptations on temporal scales of both the short-scale dynamics, and the long-scale dynamics would increase its applicability. Examining the study that used the MPC formulation, a notable obstacle involved developing a reduced-order model that could accurately represent the system dynamics while remaining computationally efficient (Yao and Kothare, 2020).

To provide improvements in these directions, as well as allowing for control of multiple physiological outputs, MPC with a LSTM model was selected for investigation in this chapter. The details of both MPC and LSTM can be found in Chapter 1. Of specific relevance for the LSTM is the capability to capture the dynamics with both long and short temporal scales, which is enabled through the internal gating structures in the LSTM. Additionally, the LSTM learns the relationship between the

inputs and outputs on a temporal scale defined by the user, while MPC provides the necessary framework to enable the multiple input and multiple output problem formulation. While this approach has not been used in biological systems, LSTM in MPC has been used for chemical plant control (Wu *et al.*, 2019). This novel approach addresses the challenges of modeling the cardiac system by using a data-driven model, while providing the flexibility for tracking multiple physiological variables.

In this chapter, I use the LSTM designed in Chapter 2 with MPC to develop a data-driven controller. First, I demonstrate the controller's ability to control the LSTM itself in closed-loop. Next, I show the controller is able to control the full physiological model with a nominal offset. I provide an investigation into a few different control designs to demonstrate the versatility of the controller with different contextual objectives. Finally, I provide an analysis of the controller designs and guidelines when applying this approach.

3.2 LSTM IN CLOSED-LOOP WITH LSTM

To verify that the LSTM based MPC controller was operating as expected, the controller was designed to control the output from the LSTM. In the case of the feedback scheme in Fig. 3.1, the subject would indicate the trained LSTM model. The expectation of the controller is to reach the target setpoints since there is no model mismatch between the model used in MPC (LSTM), and the controlled system (LSTM), which provides a control problem for benchmarking controller performance.

3.2.1 Controller Design

The controller will be controlling the heart rate and the mean arterial blood pressure to reach specified target set points by manipulating the stimulation pulse frequency, and the stimulation pulse width at three different locations. To determine the stimulation parameters at each location, the controller uses the LSTM to make predictions of the system's response for the length of the predictive horizon. These predictions are used in formulating a cost function that compares the system's response to the target set point. Taken together, the setup of the controller scheme is shown in Fig. 3.1, with

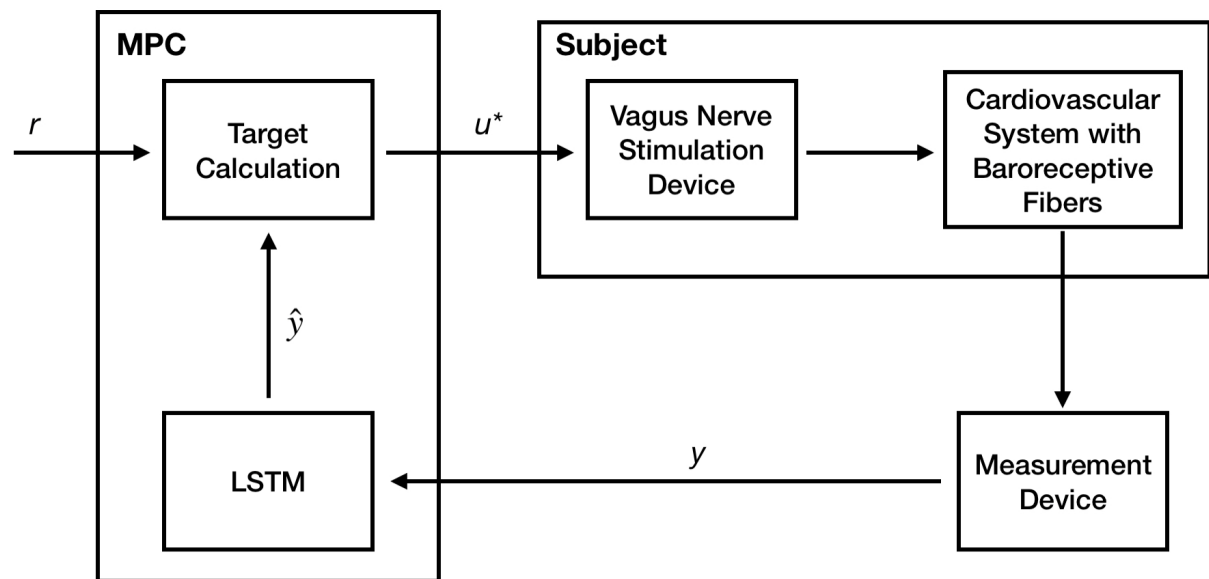


FIGURE 3.1: A block diagram that details the design of the controller, labeled as MPC. The target set point is denoted with r , the optimal control actions are denoted by u^* , the measurement of the physiological variables from the rat cardiac system (labeled as subject) are denoted by y , and the predictions from the LSTM are denoted with \hat{y} .

the output denoted by y . The LSTM model based controller provides the optimal stimulation parameters u^* , to the full physiological model. The output from the full physiological model is then provided as feedback to the LSTM based MPC.

Since there are three discrete stimulation locations available, the full optimization problem is a mixed integer nonlinear programming problem (MINLP). To serve as an approximation of the MINLP problem, a L1 cost term was added to the objective function. This additional cost term penalized the input, and as a result the sparse solution was found by the controller, thus approximating the solution to the MINLP problem. For MINLP problems, several different solution paths must be considered which leads to multiple optimization problems at each iteration. As a consequence of multiple problems, optimization becomes computationally expensive to implement, thus an advantage of adding the L1 cost is speeding up the controller for real-time implementation. In addition to the L1 cost, a simple quadratic cost between the model predictions, and the target was used to find the optimal control action. Together, these lead to the following optimization problem:

$$\begin{aligned} \min_u \quad & \sum_{i=1}^{N_p} (r(k) - \hat{y}_i)^2 + \lambda \sum_{j=1}^{N_c} |u_j|, \\ \text{s.t.} \quad & u_{min} \leq u(k+j|k) \leq u_{max}, j \in [1, N_c], \\ & \hat{y} = f_{NN}(\hat{y}(k+i|k), u(k+i|k)), i \in [1, N_p], \\ & \hat{y}(k|k) = y_0. \end{aligned} \tag{3.1}$$

Here, $f_{NN}(\cdot, \cdot)$ denotes the LSTM, \hat{y}_i denotes the predicted physiological values from the LSTM model, $r(k)$ denotes the set point of the physiological variables, and u_j denotes the VNS input parameters. The constraints require that the model predictions come from the LSTM, and bounds of the VNS input are not exceeded which provides an inherent safety constraint. The L1 cost was multiplied by a weight term, λ , for balancing the input penalty with the performance error penalty. The control horizon is denoted by N_c , and the prediction horizon is denoted by N_p . The three final values (N_c, N_p, λ) represent design parameters of the controller, and all of them impact the

optimized control actions. In Eqn. 3.1, the first term of the minimization problem denotes the quadratic cost (enforces accuracy) and the second term denotes the L1 cost (input penalty).

3.2.2 Controller Tracking Simulation

As a simple test case for the controller described by the optimization problem of Eqn. 3.1, three different set points were provided and the controller was asked to reach each set point within 50 cardiac cycles. The set points for the tracking problem are detailed in Table. 3.1, and are the same set point targets used in the MPC based controller study in (Yao and Kothare, 2020).

TABLE 3.1: Set Points for the Controller

Set Point No.	Cycles	HR (bpm)	MAP (mmHg)
1	0-50	392	111
2	51-100	346	144
3	101-150	393	125

By setting $\lambda = 0.001$, $N_c = 5$, and $N_p = 10$ the controller achieved accurate performance and was able to reach all of the set points in the case where no model mismatch exists. The result from this simulation is shown in Fig. 3.2, with the selected controller actions shown in Fig. 3.3. The controller is able to reach the set point and keep the system there. Additionally, looking at the controller actions, at least one location is turned off at all times. By including the L1 cost in the objective function, this result demonstrates that the controller performance results in a sparse solution.

3.3 LSTM IN CLOSED-LOOP WITH PHYSIOLOGICAL MODEL

Having shown that the controller was able to successfully track set points in a case with no model mismatch, the next task was to track set points in a case with model mismatch. Since the LSTM was trained to be an approximation of the full physiological model, controlling the full physiological model in a closed-loop would

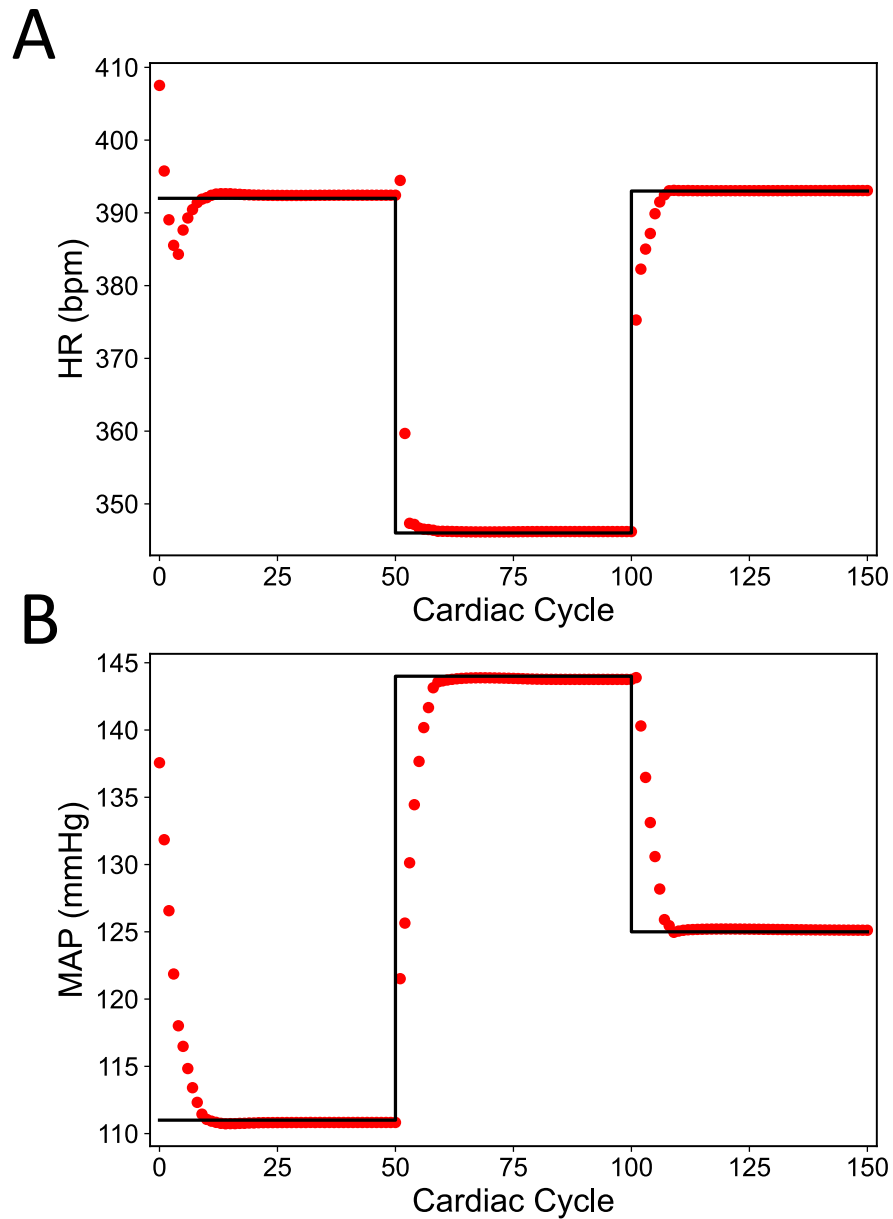


FIGURE 3.2: LSTM controller tracking set points for **(A)** heart rate (HR) and **(B)** mean arterial blood pressure (MAP) when controlling the LSTM in closed-loop. Set points are denoted with a black line, while the controlled output is indicated with red dots.

introduce some model mismatch. As discussed in Chapter 2, the LSTM did not perfectly match the full model.

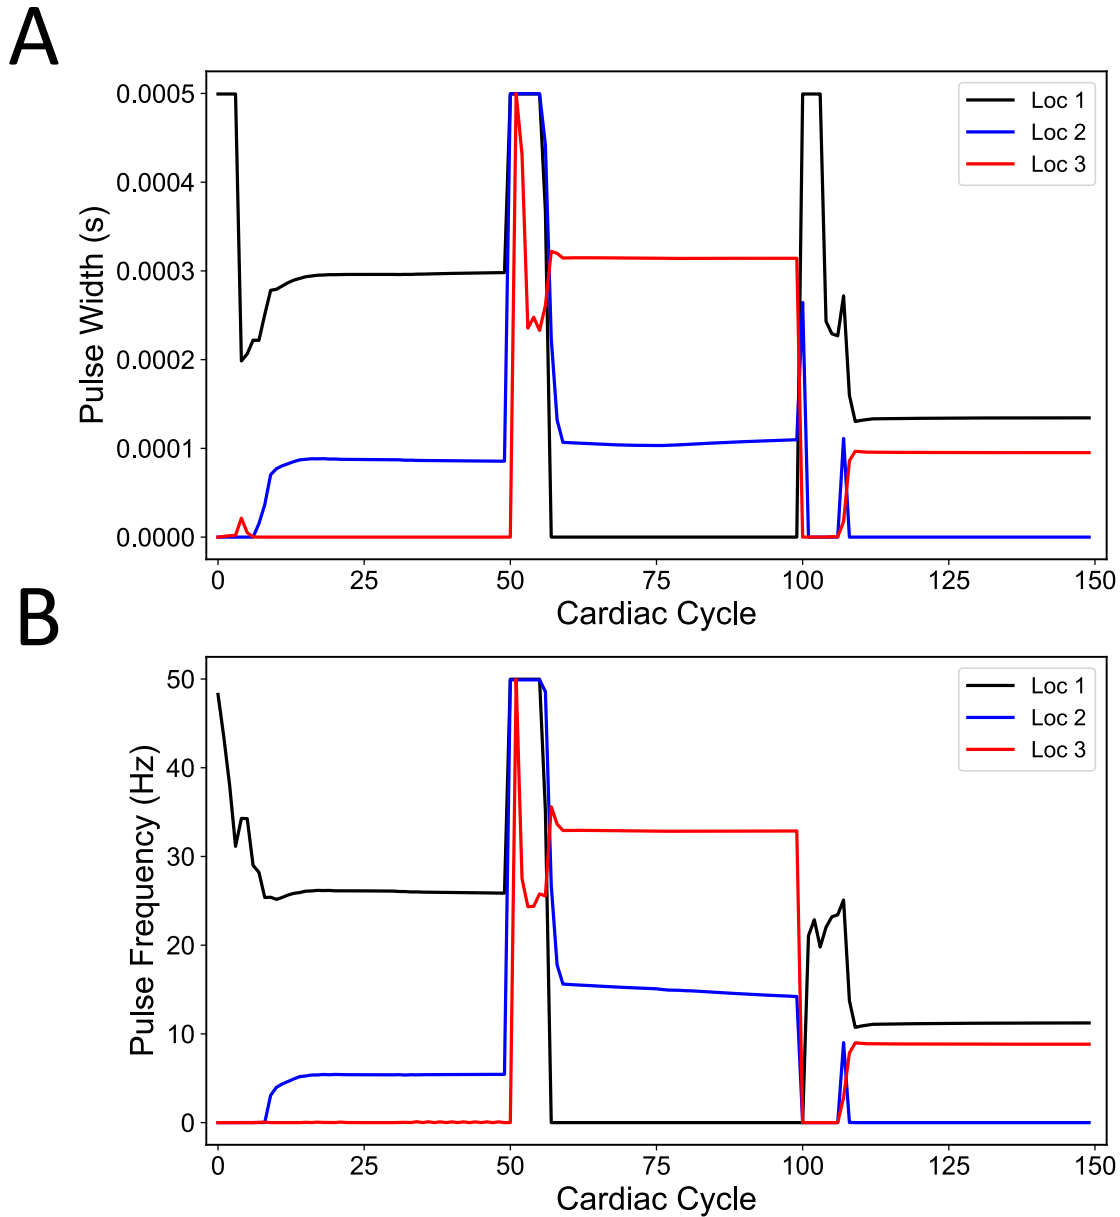


FIGURE 3.3: LSTM controller actions to achieve the set point for **(A)** pulse width and **(B)** pulse frequency when controlling the trained LSTM in closed loop. The three different locations are differentiated by color with the first location in black, the second location in blue, and the third location in red.

3.3.1 L1 Controller Simulation

Using the optimization problem described by Eqn. 3.1, a controller was designed to control the physiological model. The controller set points were the same as in Table 3.1. For this application, the following values were used in the controller formulation: $\lambda =$

0.001, $N_c = 5$, and $N_p = 10$. Taken together, this is the same control problem described previously, however the controller is controlling the full physiological model in this case.

The results of the aforementioned control problem are shown in Fig. 3.4, with the selected controller actions shown in Fig. 3.5. Again, for each set point, one location is turned off, thus demonstrating that the L1 cost enforces a sparse solution. Importantly, the LSTM-based MPC controller reaches these set points with nominal offset. For all set points, there is an overshoot for the heart rate, however the controller still reaches the target set point. Importantly, the controller demonstrates the ability to control both the heart rate, and the mean arterial blood pressure simultaneously by optimizing six different input variables.

3.3.2 L2 Controller Simulation

In an effort to consider a different context of interest regarding the application of VNS, an optimization problem that replaced the L1 cost with a L2 cost was formulated. A L2 cost penalizes the input and is equivalent to the square of the L1 cost. This is shown in Eqn. 3.2, which details the optimization problem with the L2 cost. Since VNS is a battery powered implanted electrical device, power consumption becomes an area of concern. To incorporate this context in the controller, the L2 cost can be used to emphasize avoiding large values for stimulation parameters, which correspond to high power consumption.

$$\begin{aligned}
 & \min_u \sum_{i=1}^{N_p} (r(k) - \hat{y}_i)^2 + \lambda \sum_{j=1}^{N_c} (u_j)^2 & (3.2) \\
 & \text{s.t. } u_{min} \leq u(k+j|k) \leq u_{max}, j \in [1, N_c] \\
 & \hat{y} = f_{NN}(\hat{y}(k+i|k), u(k+i|k)), i \in [1, N_p] \\
 & \hat{y}(k|k) = y_0
 \end{aligned}$$

Examining the controller actions, the solution to the L2 cost formulation emphasizes avoiding maximum values, which correspond to the highest level of

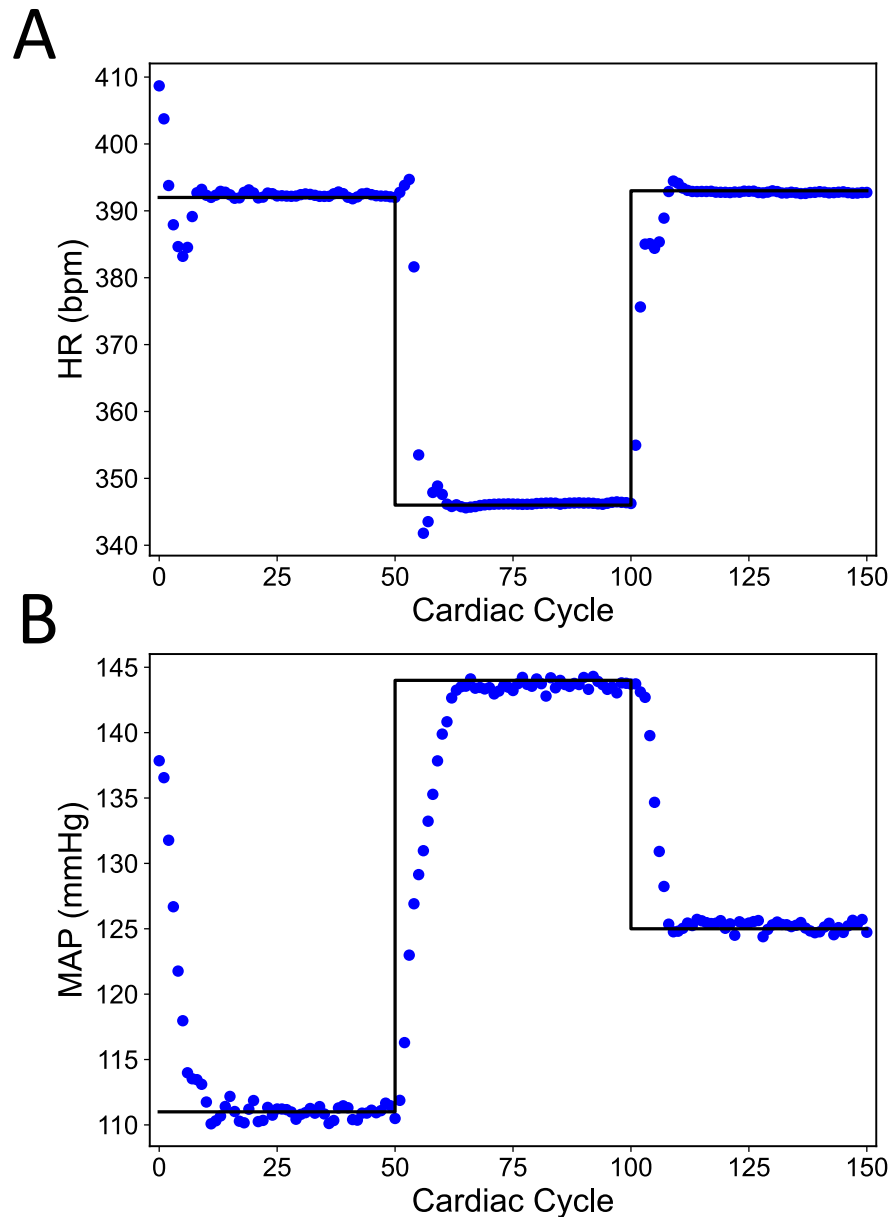


FIGURE 3.4: Performance of L1 objective function in LSTM based controller for set point tracking of (A) heart rate (HR) and (B) mean arterial blood pressure (MAP) of the physiological model. Set points are denoted with a black line, while the physiological model output is indicated with blue dots.

power consumption. Despite the different controller emphasis, the L2 controller demonstrates an ability to reach all set points with nominal offset. When approaching each set point for the heart rate, the controller exhibits overshoot before returning to

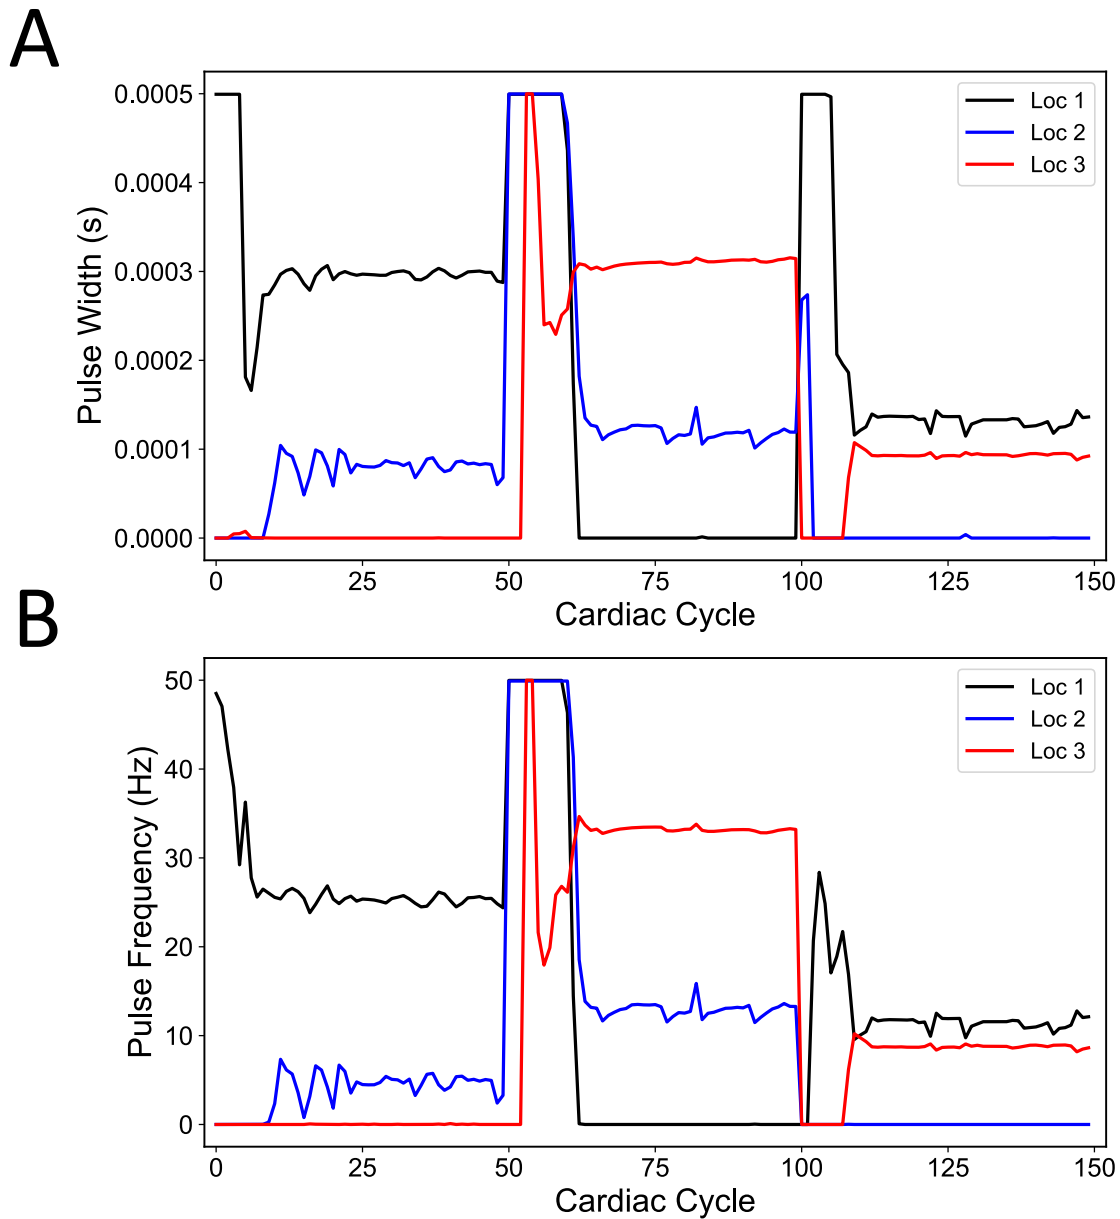


FIGURE 3.5: Controller actions for the L1 objective function to achieve the set points for (A) pulse width and (B) pulse frequency in controlling the physiological model. The three different locations are differentiated by color with the first location in black, the second location in blue, and the third location in red.

the target set point. Overall, the LSTM based controller has demonstrated reasonable performance with two different objective function formulations.

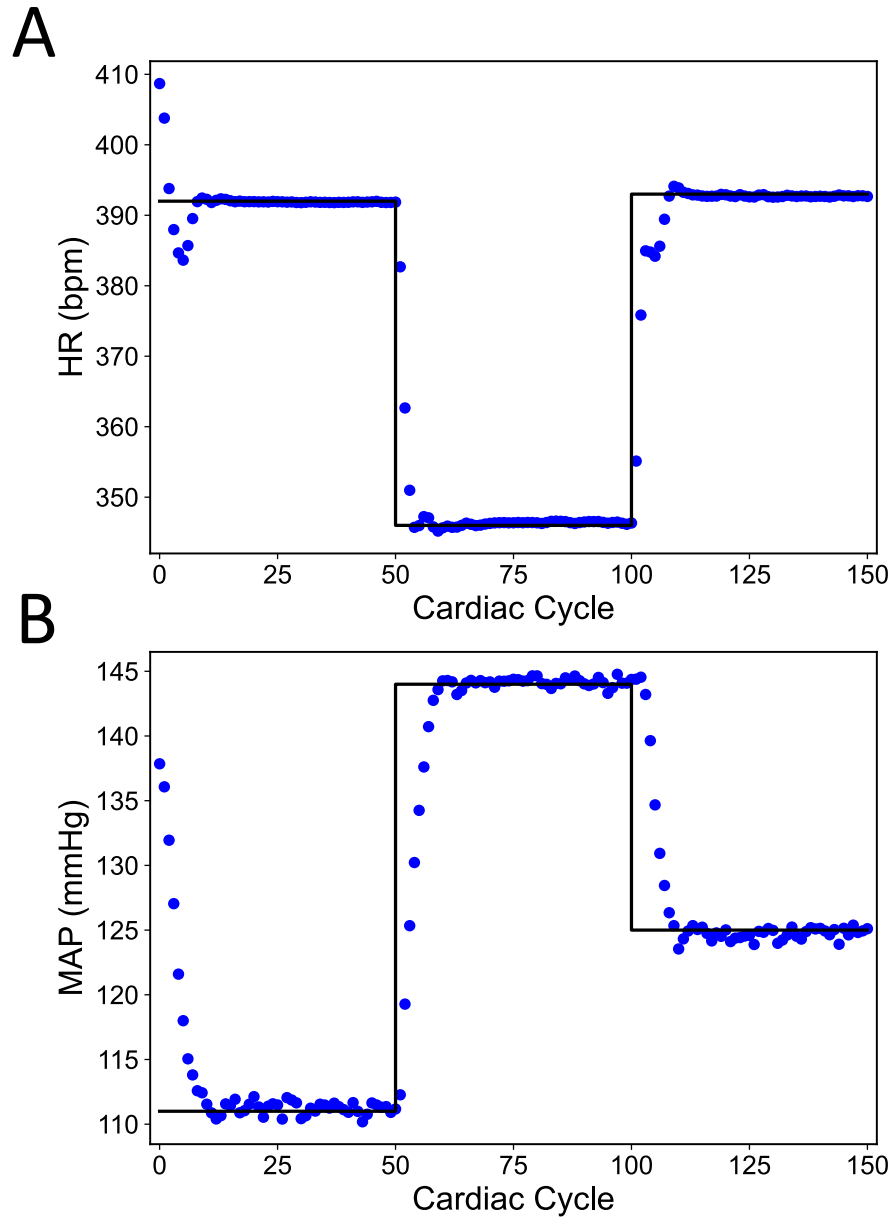


FIGURE 3.6: Performance of L2 objective function in LSTM based controller for set point tracking of (A) heart rate (HR) and (B) mean arterial blood pressure (MAP) of the physiological model. Set points are denoted with a black line, while the physiological model output is indicated with blue dots.

3.3.3 ΔU Controller Simulation

When the previous controller solutions are considered, the controlled outputs appear to oscillate around the steady-state value for MAP. By examining the controller actions, it appears that these selected actions are giving rise to the slight oscillations. These

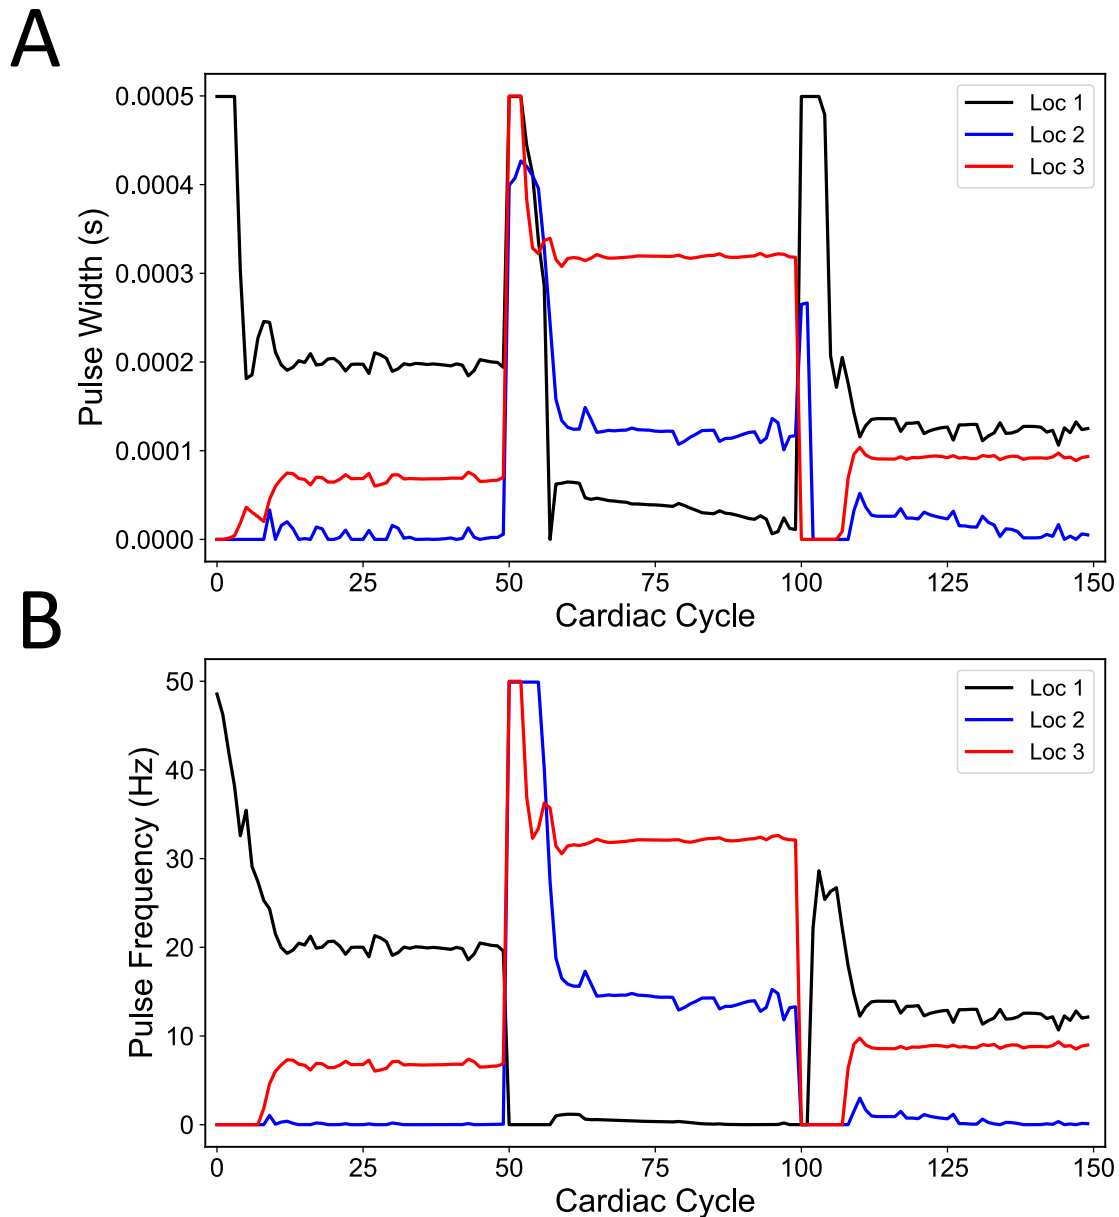


FIGURE 3.7: Controller actions for the L2 objective function to achieve the set points for **(A)** pulse width and **(B)** pulse frequency in the physiological model. The three different locations are differentiated by color with the first location in black, the second location in blue, and the third location in red.

oscillations motivate the test of a final controller design where penalizing the change in the input was formulated as an optimization problem in an effort to remove these oscillations. By penalizing the change in input in addition to the L1 cost, the controller should perform more smooth control actions, specifically during steady-state at the

target set points. As an illustration of the goal, this smooth controller behavior is inherently observed in the case of the LSTM controlling the LSTM in the closed-loop control (see Fig. 3.2 and Fig. 3.3). This problem formulation introduced an additional parameter, λ_2 , to weight the change in the inputs. Changing the weighting between λ_1 (the L1 cost weight) and λ_2 (the Δu weight) influences the result by shifting the controller's optimization priority. The described optimization problem used for this controller design is detailed by Eqn. 3.3.

$$\begin{aligned}
\min_u \quad & \sum_{i=1}^{N_p} (r(k) - \hat{y}_i)^2 + \lambda_1 \sum_{j=1}^{N_c} |u_j| + \lambda_2 \sum_{j=1}^{N_c} \Delta u_j & (3.3) \\
\text{s.t.} \quad & u_{min} \leq u(k+j|k) \leq u_{max}, j \in [1, N_c] \\
& \hat{y} = f_{NN}(\hat{y}(k+i|k), u(k+i|k)), i \in [1, N_p] \\
& \Delta u_j = |u(k+j|k) - u(j|k)|, j \in [1, N_c] \\
& \hat{y}(k|k) = y_0 \\
& u(k|k) = u_{t-1}^*
\end{aligned}$$

For this controller optimization, the design parameters were set to the following values: $N_c = 5$, $N_p = 10$, $\lambda_1 = 0.001$, and $\lambda_2 = 0.00005$. Different selections of N_c and N_p can lead to similar performance, provided that the weights (λ_1 , λ_2) are appropriately selected.

As shown by the controlled output in Fig. 3.8, the addition of the Δu term led to a more smooth solution while exhibiting nominal offset. The mean arterial blood pressure still appears to vary a little, however these variations are not as large as the previously designed controllers. Notably, there is no overshoot in the last set point for the heart rate, and there is minimal overshoot for the other two target set points of the heart rate. The controller actions are where the influence of this additional term becomes more obvious, as depicted in Fig. 3.9. The controller is attempting to reach the stimulation values that lead to the target set point and staying at those values, consistent with the intuition of the designed objective function. Interestingly, the Δu

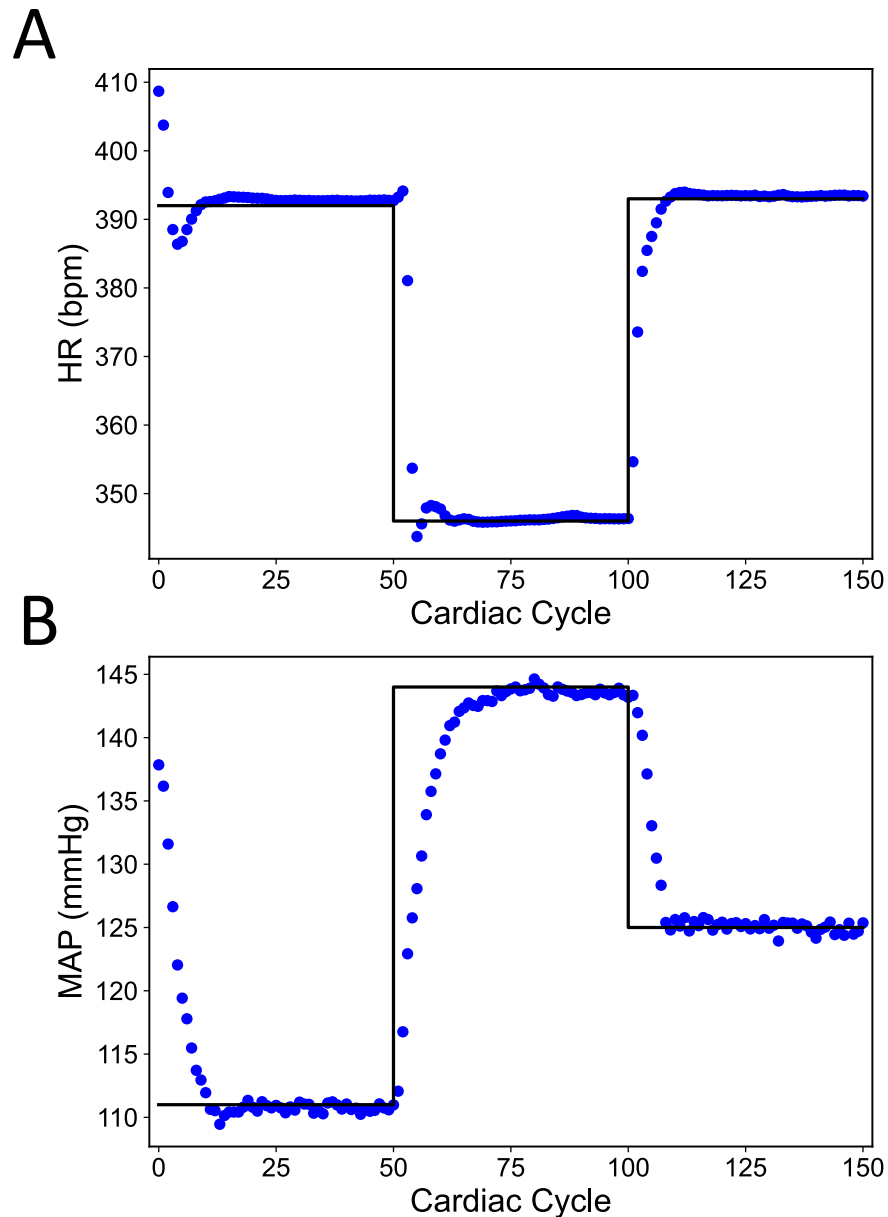


FIGURE 3.8: Performance of L1 with Δu objective function in LSTM based controller for set point tracking of (A) heart rate (HR) and (B) mean arterial blood pressure (MAP) of the full physiological model. Set points are denoted with a black line, while the physiological model output is indicated with blue dots.

controller does not achieve the sparse solution at all target set points, evidenced by having all three locations active for the last set point.

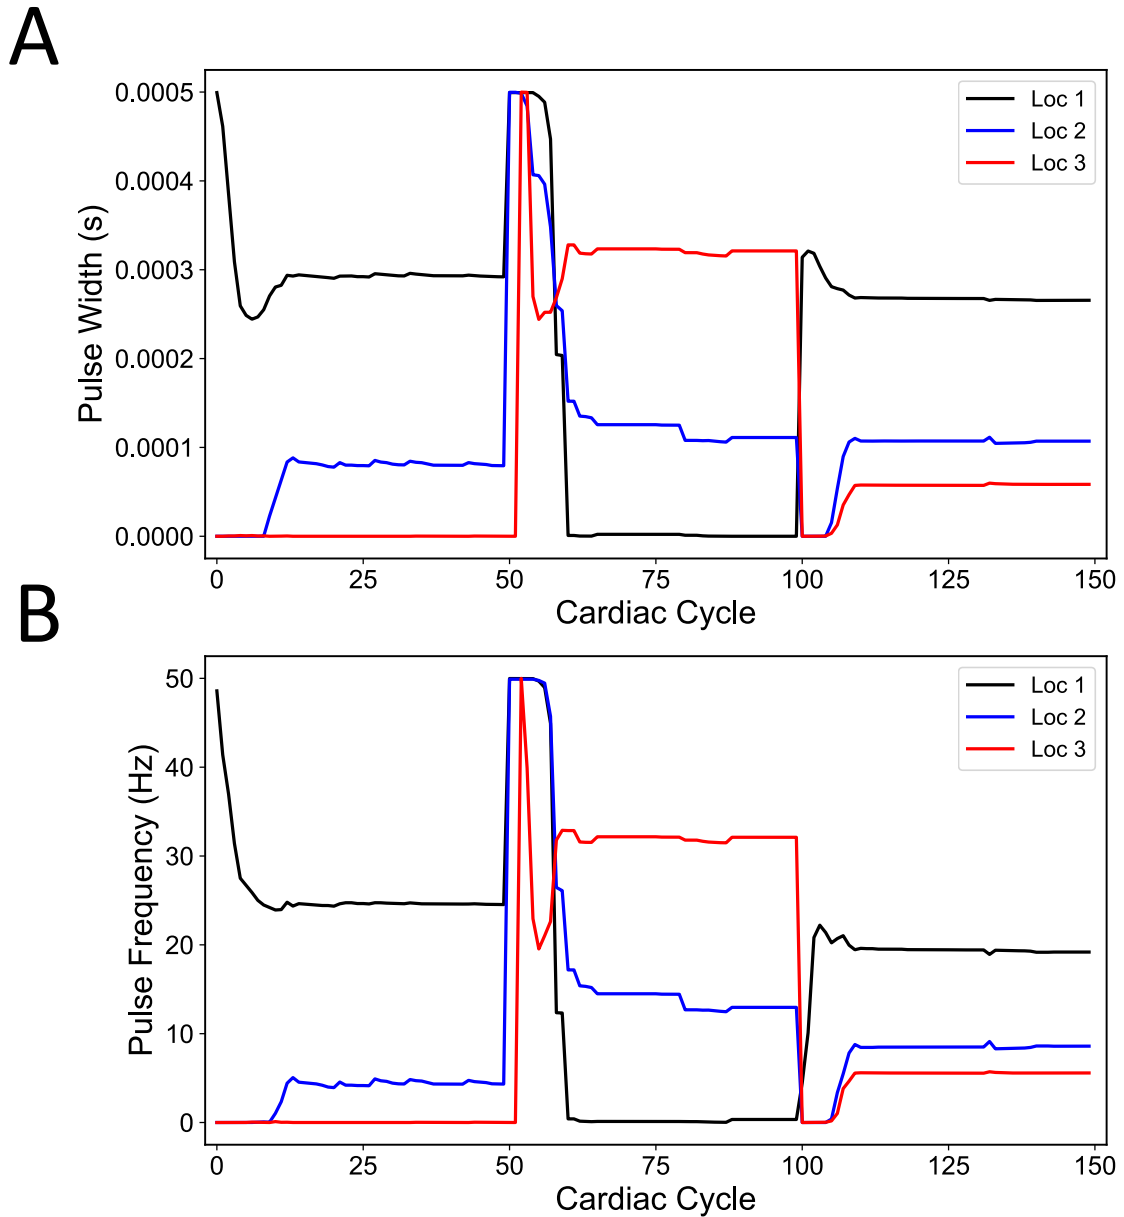


FIGURE 3.9: Controller actions for the L1 with Δu objective function to achieve the set points for (A) pulse width and (B) pulse frequency. The three different locations are differentiated by color with the first location in black, the second location in blue, and the third location in red.

3.4 CONTROLLER ANALYSIS

Here, I compare the performance of the different controller designs, the L1 and L2 cost formulations are directly shown in Fig. 3.10, with the different controller actions shown in Fig. 3.11. Both the controller designs demonstrated similar performance

for the tracking case, with the L2 formulation reaching some set points slightly faster than the L1 formulation. However, the actions that each controller took to obtain that performance differed slightly. Focusing on location 1, the L2 controller chose notably lower values for the first set point. This difference in performance validates the intuition that the L2 cost seeks to avoid high values. Throughout the entire simulation, the L2 controller did not emphasize setting one location to 0 for all set points, conversely the L1 cost did set one location to 0.

Regarding selection of control design parameters, there are some guiding principles to select these parameters based on the controller's observed behavior. In the case of the L1 cost, high values of λ instructs the controller to emphasize suppressing the input more than reaching the target value. The occurrence of such a scenario would be evidenced by offset in the controlled output. Following the L2 cost formulation, high values of λ can lead to a similar situation where the controller is more focused on minimizing the L2 cost term than the quadratic cost of the predictions. There is a balance between the two cost terms that leads the controller to minimize the quadratic cost followed by minimizing the second cost term. Such a balance is demonstrated in the controller outputs of the above cases.

A notably different result was obtained using the Δu controller formulation, specifically regarding the overshoot around the set points. The formulation of this optimization problem did not exhibit overshoot in the same way that the other controllers did (compare Fig. 3.8 with Fig. 3.4 and Fig. 3.6). Additionally, the evolution of the controlled variables were smooth and more gradual. However, this controller design was not separate from the other controller designs as it employed a L1 cost term as well. In that sense, the Δu cost term can be added to all objective function formulations. For example, it could be added to the L2 cost function as well. Similar to the L1 cost approximating the MINLP problem, the Δu cost term can be added to give the controller a more gradual and gentle performance.

Should these controller designs be deployed in the real physiological setup, there is the potential that the controller could switch its objective function to meet the demands of the biological system. In a therapeutic sense, stimulating all locations at all times could lead to a loss in efficacy of treatment. Conversely, there may be times

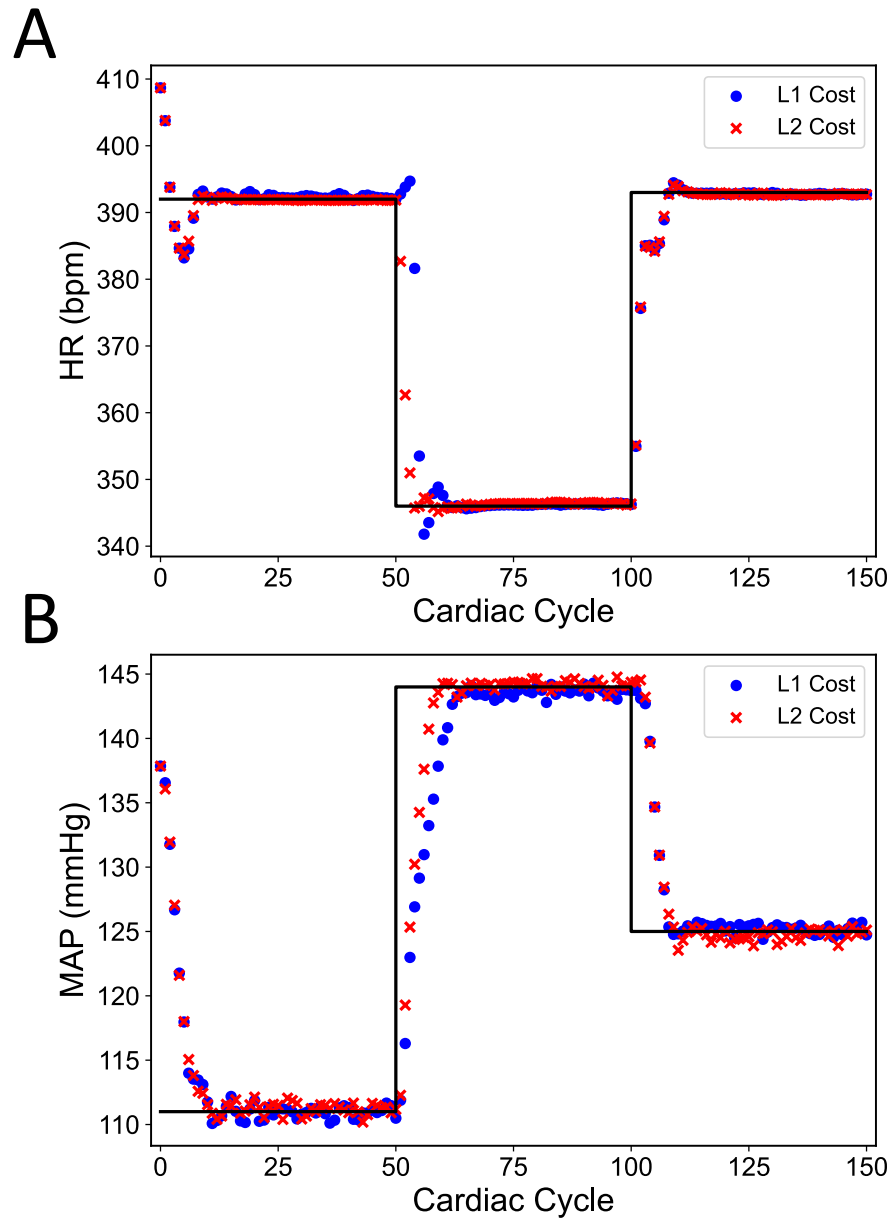


FIGURE 3.10: Comparison of controller performance for the L1 cost denoted by Eqn. 3.1 (blue dots) and controller performance for the L2 cost denoted by Eqn. 3.2 (red crosses) with $N_c = 5$, $N_p = 10$, and $\lambda = 0.001$ in both controller formulations for **(A)** heart rate (HR) and **(B)** mean arterial blood pressure (MAP). The target set points are shown by the black line.

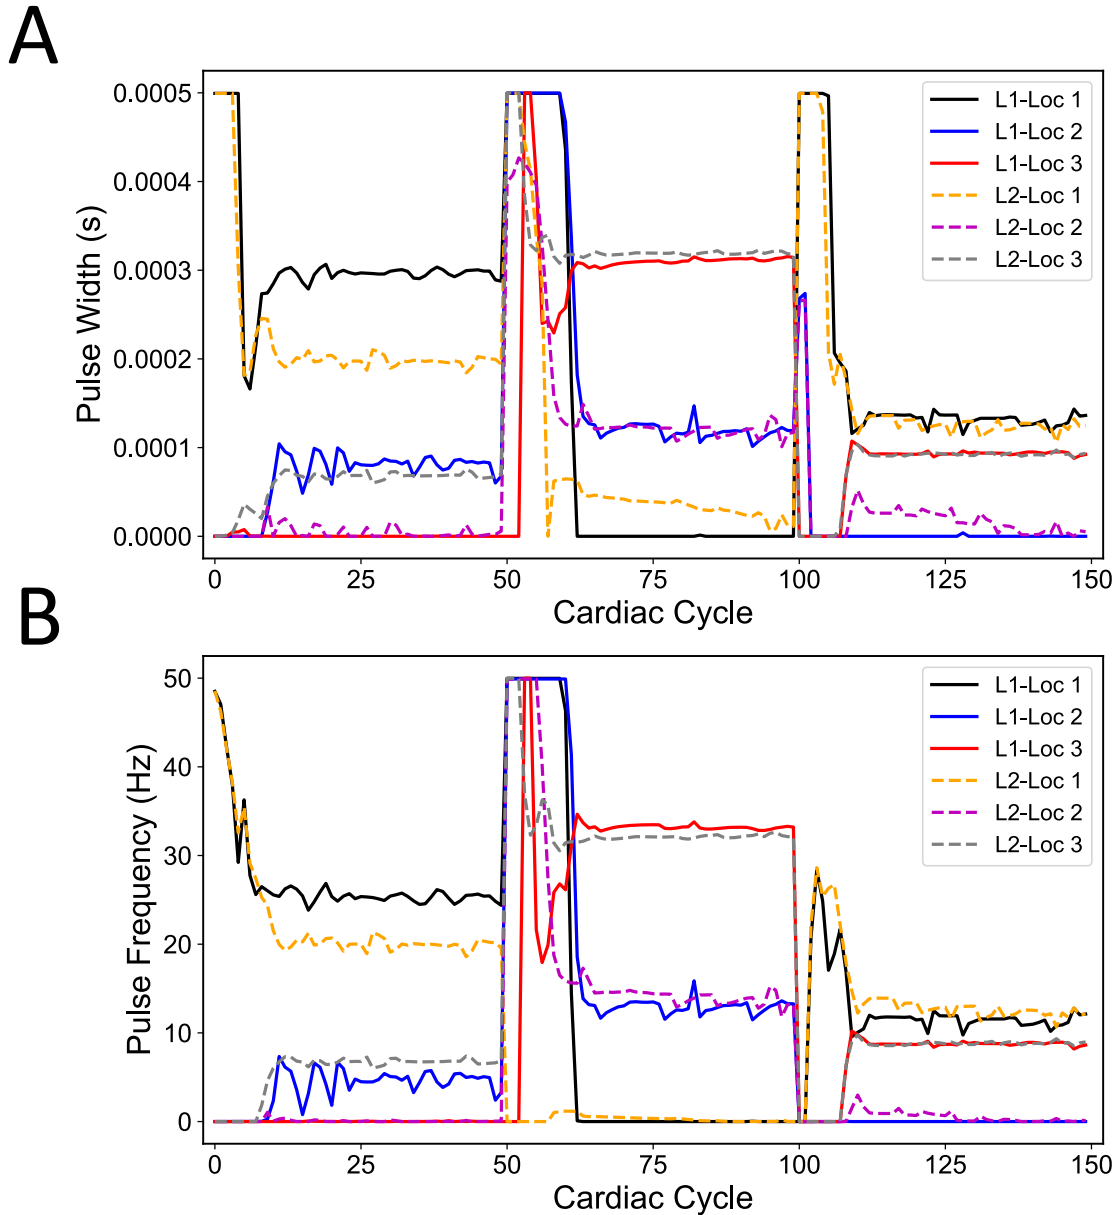


FIGURE 3.11: Comparison of controller actions for the L1 cost denoted by Eqn. 3.1 (solid lines) and controller action selection for the L2 cost denoted by Eqn. 3.2 (dotted lines) with $N_c = 5$, $N_p = 10$, and $\lambda = 0.001$ in both controller formulations for **(A)** pulse width and **(B)** pulse frequency. Locations are differentiated by color with L1 location 1 shown in black, location 2 shown in blue, and location 3 shown in red. For L2, location 1 is shown in orange, location 2 is shown in magenta, and location 3 is shown in grey.

when power consumption holds a higher priority than enforcing a sparse solution. For the therapeutic application, the Δu cost term can be used to reach targets with more smooth transitions, which may be more gentle for the patient or subject of the therapy. These different contexts are why multiple controller designs were investigated in this thesis. Importantly, all controllers have been shown to reach the target set point with a nominal offset. Their similar performances lead to the potential of employing them in the experimental system, and allowing for multiple external factors (battery level, physiological feedback, stimulation time, etc.) to provide the context for selecting the specific controller design used at a particular time.

3.4.1 *Controller Computational Time*

Additionally, the computational time for different choices of N_p and N_c were computed. As expected, when the control horizon was increased, the computational time increased as well. A longer control horizon led to longer times for optimization to converge because the number of optimized variables increased by a six-fold rate. For every additional cardiac cycle in the control horizon, there were six additional parameters for the controller to optimize. Similarly, as the predictive horizon increased, the time of optimization also increased as there were more predictions to be computed. However, this additional computational time was not as significant compared to increasing N_c . The computational time was estimated by taking the total time to run the simulation and subtracting the time it took the full physiological model to run the number of cycles simulated, and dividing this difference by the total number of cycles in the simulation. All of these findings are summarized in Table 3.2.

In the context of requirements for real-time closed-loop control of VNS, it is likely that selections with $N_c \geq 5$ and $N_p \geq 20$ would be difficult to implement in real-time. This leaves $N_c \leq 5$, and $N_p \leq 10$ as plausible choices for real-time implementation. This analysis is based on the estimation that the controller would be continuously optimizing with roughly one second between iterations to optimize the result. In the case of computing multiple controller actions, $N_c > 1$, control actions can be implemented one at a time in case the next optimization has not converged in time. In this application, all controller designs listed above could be configured to give

TABLE 3.2: Computational Time of Controllers Tests Were Done Using an INTEL(R) CORE I7-9700 CPU 3.00 GHZ with 16.0 GB of RAM

Nc (cycles)	Np (cycles)	Variables	Est. Time/cycle (s)
1	10	6	0.21
5	10	30	0.88
10	10	60	1.15
1	20	6	0.57
5	20	30	1.85
10	20	60	2.35
20	20	120	2.79

similar performance by carefully selecting the weighting values (λ). Despite these considerations, it's worth noting that some VNS paradigms utilize a duty cycle of 14 seconds on, and 66 seconds off (Ardell *et al.*, 2015, 2017). For these stimulation paradigms, all possibilities listed above could be implemented in real-time.

3.5 CONCLUSIONS

In this chapter, I provided the extension of the LSTM developed in Chapter 2 by embedding the LSTM in a MPC framework to allow for controlling the full physiological model. I considered several different objective functions, and analyzed their performance. When considering the design of a controller, I provided guidance into selecting the parameters associated with the controller design. Overall, the LSTM-based MPC controller was able to achieve acceptable closed-loop performance when controlling the full physiological model.

The novel control approach outlined here is considerably different from other control studies focusing on controlling the cardiac system. To my knowledge, there has not been a control-based attempt that utilizes deep learning approaches in the cardiac system. Additionally, the novelty of controlling multiple outputs by optimizing multiple inputs has not been shown in the other controller designs. By using a LSTM, the potential improvements regarding timescales suggested by (Romero-Ugalde *et al.*, 2017) have been addressed. Finally, the controllers shown in

this chapter demonstrate nominal offset, which is behavior that the original study's reduced-order controller did not exhibit (Yao and Kothare, 2020).

Viewing the controller designs from this chapter as an overall approach, there are many systems that may gain advantages from utilizing this methodology. As demonstrated in this chapter, this approach can be used to develop reduced-order models of larger-scale models, thus enabling the plausibility of real-time control. Other potential application areas where this approach would excel can be found in areas where model development is challenging, however a control application would still prove useful. The challenge of model building may result from lack of adequate system knowledge, or the challenge may arise from the level of detail required for modeling the system. In either case, provided input-output data, this approach can be used to develop a controller. As an example of such an area, this would be particularly advantageous in the context of controlling biological systems.

While the designed controller described in this work did not receive experimental validation, there is no detrimental reason preventing its application in experimental setups as well. Particularly, the inclusion of a deep learning based model allows for the flexibility of training that model with the experimental data. Notably, the deep learning model described here was never given information regarding the structure of the underlying physiological model. The challenge may shift to obtaining a sufficiently accurate trained model, however that challenge may be more feasible than developing a model for some systems.

There are a few notable limitations with this data-driven control approach, particularly the controlled system must be fully observable. More specifically, to train the neural network, measurements of the controlled variables as well as the manipulated variables must be provided. Thus, a partially observed system may be possible to control using this approach, but that has not been demonstrated in this thesis. Expanding this approach to account for such a situation may be an area of interest for future work. There are other situations where the current approach may not lead to an improvement in performance, such as systems that can already be accurately approximated by computationally efficient models. After these limitations

are taken into account, there remains a broad area of applications across several domains where this approach would prove useful.

CHAPTER 4

SUMMARY AND FUTURE DIRECTIONS

4.1 SUMMARY

In this thesis, I presented a unified approach for modeling and controlling the rat cardiac system by using machine learning approaches to develop a model and applying that model in model predictive control. This approach provided a useful framework for applications of control wherein a model is either difficult to obtain, a full model is too computationally expensive for optimization, or a full mechanism is not understood, thus challenging efforts to build a model. I applied this framework to a model of the rat cardiac system, and demonstrated that the full physiological model can be controlled with this approach. Generally stated, my thesis provided an additional approach for data-driven control that merges machine learning with model predictive control.

In chapter 2, I used machine learning approaches to develop a timeseries relationship between the input VNS parameters and the output physiological variables. Specifically, I used a LSTM which is commonly used for time-series or sequence-based datasets. Data from the true physiological model was then collected using an open-loop trial characterization of the system (i.e. choose parameters, collect response, repeat). This led to the LSTM serving as a map between the current state of the system, and the evolution of physiological variables following the selection of a specific set of VNS parameters. I showed that this model is 10-folds faster in computational time compared to the full model, and achieved the speed required for real-time optimization. My results suggested that this data-driven approach is a viable method to model the effect of VNS parameter selection on the physiological variables of heart rate, and blood pressure.

Using the LSTM designed in Chapter 2, I designed a controller in Chapter 3 that drove the system to several set points with a nominal offset during closed-loop simulations. The controller was designed using a MPC problem formulation with the

trained LSTM serving as the model. Using a L1 cost, I showed that the mixed integer programming problem can be approximated, thus allowing for easier implementation in real time. Finally, I compared different objective function formulations and provided guidelines for the selection of their associated parameters.

4.2 FUTURE WORK

4.2.1 *Using Physiological Data*

In this thesis, I provided a framework to model and control a nonlinear physiological model of a rat heart. The work provided in my thesis is purely computational, leading to the next application of this framework to real data. The problem formulation used in this thesis allows for extension and application using other sources of data because the developed LSTM and controller were never given information about the underlying physiological model. Similar to the process employed on the physiological model, experiments can be configured to collect data in the form of open-loop trials investigating the influence of VNS parameter selection on various physiological measurements. As an example, this analysis has been completed in dogs (Ardell *et al.*, 2015, 2017). This dataset specifically investigated the influence of different VNS parameter selection (pulse width, pulse frequency, and pulse amplitude) while recording electrocardiogram (ECG), left ventricular pressure, and blood pressure. The applied duty cycle for VNS consisted of 14 seconds on and 66 seconds off. By performing the same approach on this dataset, or a similar dataset, the resulting controller would be directly applicable to the *in vivo* system and could lead to experimental validation and improvements in both the modeling approach and the controller design.

To highlight the key differences between the two datasets, heart rate is not directly measured and must be inferred from the ECG measurement. This problem is equivalent to predicting the inter-spike interval, which has been demonstrated using synthetic data with neural networks in other work (Plaster and Kumar, 2019). There will also be statistical variation in the measured data (noise), that must be handled by both the modeling approach and the controller. The consideration of noise, highlights

an important underlying assumption that unique VNS parameters lead to a specific and unique response in the measured physiological variables.

4.2.2 *Controlling Diseased Model*

In Chapter 3, the designed controller was based on the same model as was used for training the LSTM. However, an interesting investigation would involve developing a rat cardiac model with a pathological condition. This investigation could train a LSTM using data from the healthy case and use that LSTM in a MPC framework. The resulting controller could then be applied to the pathological model, to investigate challenges and solutions to those challenges. This investigation would be of clinical relevance as it provides insight into the potential of using data from healthy subjects in a controller used to treat a pathological condition. There is the potential that a wide range of data could be collected and integrated into the controller in a meaningful way.

4.2.3 *Recovering Linear Dynamics*

In chapter 2, I utilized a data-driven approach to develop a relationship of VNS parameter selection on the open-loop trials from a physiological model of a rat heart, which provided a reduced-order model. In general, this approach can be applied to dynamical systems models, due to the flexibility of the data-driven approach. The LSTM based controller provided faster and more accurate convergence to target set points than the reduced-order dynamical model. A more attractive way to model the system would involve a linearization technique, which enables traditional linear control schemes. The advantage of this approach lies in the control application, where proving optimality, stability, and robustness is much more feasible than in the nonlinear case. There would be additional performance gains in required computational times through faster optimization of a linear system. These benefits are countered by the challenge involved in linearizing a nonlinear system, exemplifying a broad challenge found in control applications. An approach that utilizes an extension of the LSTM modeling approach involves employing an

autoencoder where the dynamics are encoded to a reduced dimensionality of the system. The complimentary decoder can be used to acquire the original system dynamics. Investigation into applying a linear control scheme to the encoded dynamics could provide a more efficient control scheme while maintaining a data-driven approach. Another approach that seeks to accomplish linearization of the nonlinear model involves the Koopman operator (Brunton *et al.*, 2016b; Korda and Mezić, 2018; Proctor *et al.*, 2018), which provides a coordinate transform such that a linear state space model evolves according to the nonlinear dynamics. Again, the linearized model would allow for more rigorous control analysis through well-established linear control theory, while providing a method to recover the nonlinear dynamics from the linear dynamics. Both of these approaches are illustrated in Fig. 4.1.

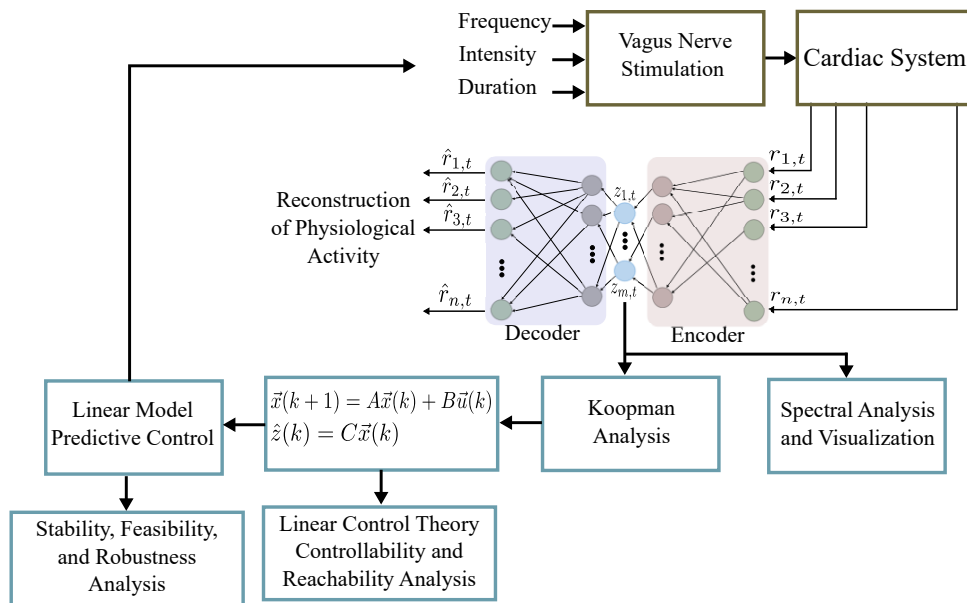


FIGURE 4.1: A control framework approach that utilizes an autoencoder for dimensionality reduction to obtain a linearized model for control. The encoder is used to obtain the reduced dimensional model, which then undergoes Koopman analysis to generate a linear model for control. The decoder can be used to reconstruct the physiological activity.

A final approach for advancing the controller involves dynamic mode decomposition (Tu *et al.*, 2013; Proctor *et al.*, 2016), which has been used to analyze several systems in domains such as epidemiology (Proctor and Eckhoff, 2015), neuroscience

(Brunton *et al.*, 2016a), robotics (Berger *et al.*, 2015), video processing (Erichson *et al.*, 2019), finance (Mann and Kutz, 2016), power systems (Le Clainche and Vega, 2017), plasma physics (Taylor *et al.*, 2018), and fluid dynamics (Schmid *et al.*, 2011; Muld *et al.*, 2012; Seena and Sung, 2011; Hemati *et al.*, 2014; Noack *et al.*, 2016). Dynamic mode decomposition has been adapted for use in the control framework, and provides a reduced-order linear model that is capable of capturing the temporally local dynamics of a system. Such an approach has similar advantages that allow for real-time application of control, while utilizing a data-driven approach. Dynamic mode decomposition provides the additional benefit of eliminating the training phase associated with the neural networks, allowing for immediate use in controller design. A comparison in performance between all the approaches discussed here would be a valuable investigation, and would provide insight on data-driven modeling techniques that perform well for the control of biological systems using vagal nerve stimulation.

BIBLIOGRAPHY

-
- Akhtar H., Bukhari F., Nazir M., Anwar M.N., and Shahzad A. 2016. Therapeutic efficacy of neurostimulation for depression: techniques, current modalities, and future challenges. *Neuroscience bulletin* 32:115–126.
- Annoni E.M., Xie X., Lee S.W., Libbus I., KenKnight B.H., Osborn J.W., and Tolkacheva E.G. 2015. Intermittent electrical stimulation of the right cervical vagus nerve in salt-sensitive hypertensive rats: effects on blood pressure, arrhythmias, and ventricular electrophysiology. *Physiological reports* 3:e12476.
- Ardell J.L., Nier H., Hammer M., Southerland E.M., Ardell C.L., Beaumont E., KenKnight B.H., and Armour J.A. 2017. Defining the neural fulcrum for chronic vagus nerve stimulation: implications for integrated cardiac control. *The Journal of physiology* 595:6887–6903.
- Ardell J.L., Rajendran P.S., Nier H.A., KenKnight B.H., and Armour J.A. 2015. Central-peripheral neural network interactions evoked by vagus nerve stimulation: functional consequences on control of cardiac function. *American Journal of Physiology-Heart and Circulatory Physiology* 309:H1740–H1752.
- Asad Z.U. and Stavrakis S. 2019. Vagus nerve stimulation for the treatment of heart failure. *Bioelectronics in Medicine* 2:43–54.
- Ay I., Sorensen A.G., and Ay H. 2011. Vagus nerve stimulation reduces infarct size in rat focal cerebral ischemia: an unlikely role for cerebral blood flow. *Brain research* 1392:110–115.
- Babic T. and Browning K.N. 2014. The role of vagal neurocircuits in the regulation of nausea and vomiting. *European journal of pharmacology* 722:38–47.
- Ben-Menachem E., Manon-Espaillet R., Ristanovic R., Wilder B., Stefan H., Mirza W., Tarver W., Wernicke J., and Group F.I.V.N.S.S. 1994. Vagus nerve stimulation for treatment of partial seizures: 1. a controlled study of effect on seizures. *Epilepsia* 35:616–626.
- Benjamin E.J., Virani S.S., Callaway C.W., Chamberlain A.M., Chang A.R., Cheng S., Chiuve S.E., Cushman M., Delling F.N., Deo R., *et al.* 2018. Heart disease and stroke statistics—2018 update: a report from the american heart association. *Circulation* 137:e67–e492.
- Berger E., Sastuba M., Vogt D., Jung B., and Ben Amor H. 2015. Estimation of perturbations in robotic behavior using dynamic mode decomposition. *Advanced Robotics* 29:331–343.
- Bonaz B., Sinniger V., and Pellissier S. 2016. Anti-inflammatory properties of the vagus nerve: potential therapeutic implications of vagus nerve stimulation. *The Journal of physiology* 594:5781–5790.

- Borovikova L.V., Ivanova S., Zhang M., Yang H., Botchkina G.I., Watkins L.R., Wang H., Abumrad N., Eaton J.W., and Tracey K.J. 2000. Vagus nerve stimulation attenuates the systemic inflammatory response to endotoxin. *Nature* 405:458–462.
- Brunton B.W., Johnson L.A., Ojemann J.G., and Kutz J.N. 2016a. Extracting spatial-temporal coherent patterns in large-scale neural recordings using dynamic mode decomposition. *Journal of neuroscience methods* 258:1–15.
- Brunton S.L., Brunton B.W., Proctor J.L., and Kutz J.N. 2016b. Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PloS one* 11:e0150171.
- Cai P.Y., Bodhit A., Derequito R., Ansari S., Abukhalil F., Thenkabail S., Ganji S., Saravanapavan P., Shekar C.C., Bidari S., *et al.* 2014. Vagus nerve stimulation in ischemic stroke: old wine in a new bottle. *Frontiers in neurology* 5:107.
- Cho K., Van Merriënboer B., Gulcehre C., Bahdanau D., Bougares F., Schwenk H., and Bengio Y. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078* .
- Djabella K., Médigue C., and Sorine M. 2005. A differential model of the baroreflex control of the cardiovascular system during a tilt test. *In Proceedings of the 44th IEEE Conference on Decision and Control*, pages 903–908, IEEE.
- Erichson N.B., Brunton S.L., and Kutz J.N. 2019. Compressed dynamic mode decomposition for background modeling. *Journal of Real-Time Image Processing* 16:1479–1492.
- Ferreira A., Chen S., Simaan M.A., Boston J.R., and Antaki J.F. 2005. A nonlinear state-space model of a combined cardiovascular system and a rotary pump. *In Proceedings of the 44th IEEE Conference on Decision and Control*, pages 897–902, IEEE.
- Gold M.R., Van Veldhuisen D.J., Hauptman P.J., Borggreffe M., Kubo S.H., Lieberman R.A., Milasinovic G., Berman B.J., Djordjevic S., Neelagaru S., *et al.* 2016. Vagus nerve stimulation for the treatment of heart failure: the inovate-hf trial. *Journal of the American College of Cardiology* 68:149–158.
- Goverse G., Stakenborg M., and Matteoli G. 2016. The intestinal cholinergic anti-inflammatory pathway. *The Journal of physiology* 594:5771–5780.
- Graves A., Liwicki M., Fernández S., Bertolami R., Bunke H., and Schmidhuber J. 2008. A novel connectionist system for unconstrained handwriting recognition. *IEEE transactions on pattern analysis and machine intelligence* 31:855–868.
- Greenwald E., So E., Wang Q., Mollazadeh M., Maier C., Etienne-Cummings R., Cauwenberghs G., and Thakor N. 2016. A bidirectional neural interface ic with chopper stabilized bioadc array and charge balanced stimulator. *IEEE transactions on biomedical circuits and systems* 10:990–1002.
- Healy S., Lang J., Naude J.T.W., Gibbon F., and Leach P. 2013. Vagal nerve stimulation in children under 12 years old with medically intractable epilepsy. *Child’s Nervous System* 29:2095–2099.

- Heidenreich P.A., Albert N.M., Allen L.A., Bluemke D.A., Butler J., Fonarow G.C., Ikonomidis J.S., Khavjou O., Konstam M.A., Maddox T.M., *et al.* 2013. Forecasting the impact of heart failure in the united states: a policy statement from the american heart association. *Circulation: Heart Failure* 6:606–619.
- Hemati M.S., Williams M.O., and Rowley C.W. 2014. Dynamic mode decomposition for large and streaming datasets. *Physics of Fluids* 26:111701.
- Hochreiter S. and Schmidhuber J. 1997. Long short-term memory. *Neural computation* 9:1735–1780.
- Hornik K., Stinchcombe M., and White H. 1989. Multilayer feedforward networks are universal approximators. *Neural networks* 2:359–366.
- Howland R.H. 2014. Vagus nerve stimulation. *Current behavioral neuroscience reports* 1:64–73.
- Kingma D.P. and Ba J. 2014. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 .
- Kishi T. 2012. Heart failure as an autonomic nervous system dysfunction. *Journal of cardiology* 59:117–122.
- Kolman B.S., Verrier R.L., and Lown B. 1975. The effect of vagus nerve stimulation upon vulnerability of the canine ventricle: role of sympathetic-parasympathetic interactions. *Circulation* 52:578–585.
- Korda M. and Mezić I. 2018. Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *Automatica* 93:149–160.
- Kwon W.H. and Han S.H. 2006. Receding horizon control: model predictive control for state models. Springer Science & Business Media.
- Lau K.D. and Figueroa C.A. 2015. Simulation of short-term pressure regulation during the tilt test in a coupled 3d-od closed-loop model of the circulation. *Biomechanics and modeling in mechanobiology* 14:915–929.
- Le Clainche S. and Vega J.M. 2017. Higher order dynamic mode decomposition. *SIAM Journal on Applied Dynamical Systems* 16:882–925.
- Mahdi A., Sturdy J., Ottesen J.T., and Olufsen M.S. 2013. Modeling the afferent dynamics of the baroreflex control system. *PLoS Comput Biol* 9:e1003384.
- Mangoni M.E., Traboulsie A., Leoni A.L., Couette B., Marger L., Le Quang K., Kupfer E., Cohen-Solal A., Vilar J., Shin H.S., *et al.* 2006. Bradycardia and slowing of the atrioventricular conduction in mice lacking *cav3.1/α1g* t-type calcium channels. *Circulation research* 98:1422–1430.
- Mann J. and Kutz J.N. 2016. Dynamic mode decomposition for financial trading strategies. *Quantitative Finance* 16:1643–1655.
- Melchior F.M., Srinivasan R.S., and Charles J.B. 1992. Mathematical modeling of human cardiovascular system for simulation of orthostatic response. *American Journal of Physiology-Heart and Circulatory Physiology* 262:H1920–H1933.

- Muld T.W., Efraimsson G., and Henningson D.S. 2012. Flow structures around a high-speed train extracted using proper orthogonal decomposition and dynamic mode decomposition. *Computers & Fluids* 57:87–97.
- Noack B.R., Stankiewicz W., Morzyński M., and Schmid P.J. 2016. Recursive dynamic mode decomposition of transient and post-transient wake flows .
- Ojeda D., Le Rolle V., Romero-Ugalde H.M., Gallet C., Bonnet J.L., Henry C., Bel A., Mabo P., Carrault G., and Hernández A.I. 2016. Sensitivity analysis of vagus nerve stimulation parameters on acute cardiac autonomic responses: Chronotropic, inotropic and dromotropic effects. *PloS one* 11:e0163734.
- Ovbiagele B., Goldstein L.B., Higashida R.T., Howard V.J., Johnston S.C., Khavjou O.A., Lackland D.T., Lichtman J.H., Mohl S., Sacco R.L., *et al.* 2013. Forecasting the future of stroke in the united states: a policy statement from the american heart association and american stroke association. *Stroke* 44:2361–2375.
- Pacher P., Mabley J.G., Liaudet L., Evgenov O.V., Marton A., Haskó G., Kollai M., and Szabó C. 2004. Left ventricular pressure-volume relationship in a rat model of advanced aging-associated heart failure. *American Journal of Physiology-Heart and Circulatory Physiology* .
- Penry J.K. and Dean J.C. 1990. Prevention of intractable partial seizures by intermittent vagal stimulation in humans: preliminary results. *Epilepsia* 31:S40–S43.
- Plachta D.T., Gierthmuehlen M., Cota O., Espinosa N., Boeser F., Herrera T.C., Stieglitz T., and Zentner J. 2014. Blood pressure control with selective vagal nerve stimulation and minimal side effects. *Journal of neural engineering* 11:036011.
- Plaster B. and Kumar G. 2019. Data-driven predictive modeling of neuronal dynamics using long short-term memory. *Algorithms* 12:203.
- Premchand R.K., Sharma K., Mittal S., Monteiro R., Dixit S., Libbus I., DiCarlo L.A., Ardell J.L., Rector T.S., Amurthur B., *et al.* 2014. Autonomic regulation therapy via left or right cervical vagus nerve stimulation in patients with chronic heart failure: results of the anthem-hf trial. *Journal of cardiac failure* 20:808–816.
- Proctor J.L., Brunton S.L., and Kutz J.N. 2016. Dynamic mode decomposition with control. *SIAM Journal on Applied Dynamical Systems* 15:142–161.
- Proctor J.L., Brunton S.L., and Kutz J.N. 2018. Generalizing koopman theory to allow for inputs and control. *SIAM Journal on Applied Dynamical Systems* 17:909–930.
- Proctor J.L. and Eckhoff P.A. 2015. Discovering dynamic patterns from infectious disease data using dynamic mode decomposition. *International health* 7:139–145.
- Radford A., Wu J., Child R., Luan D., Amodei D., and Sutskever I. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1:9.
- Romero-Ugalde H.M., Le Rolle V., Bonnet J.L., Henry C., Bel A., Mabo P., Carrault G., and Hernández A.I. 2017. A novel controller based on state-transition models for closed-loop vagus nerve stimulation: Application to heart rate regulation. *PloS one* 12:e0186068.

- Rosas-Ballina M., Olofsson P.S., Ochani M., Valdés-Ferrer S.I., Levine Y.A., Reardon C., Tusche M.W., Pavlov V.A., Andersson U., Chavan S., *et al.* 2011. Acetylcholine-synthesizing t cells relay neural signals in a vagus nerve circuit. *Science* 334:98–101.
- Rumelhart D.E., Hinton G.E., and Williams R.J. 1985. Learning internal representations by error propagation. Tech. rep., California Univ San Diego La Jolla Inst for Cognitive Science.
- Rychlicki F., Zamponi N., Trignani R., Ricciuti R.A., Iacoangeli M., and Scerrati M. 2006. Vagus nerve stimulation: clinical experience in drug-resistant pediatric epileptic patients. *Seizure* 15:483–490.
- Sak H., Senior A.W., and Beaufays F. 2014. Long short-term memory recurrent neural network architectures for large scale acoustic modeling .
- Sarnoff S., Brockman S., Gilmore J., Linden R., and Mitchell J. 1960. Regulation of ventricular contraction: Influence of cardiac sympathetic and vagal nerve stimulation on atrial and ventricular dynamics. *Circulation research* 8:1108–1122.
- Savarese G. and Lund L.H. 2017. Global public health burden of heart failure. *Cardiac failure review* 3:7.
- Schachter S.C. *et al.* 2000. Vagus nerve stimulation. CRC Press.
- Schmid P.J., Li L., Juniper M.P., and Pust O. 2011. Applications of the dynamic mode decomposition. *Theoretical and Computational Fluid Dynamics* 25:249–259.
- Seena A. and Sung H.J. 2011. Dynamic mode decomposition of turbulent cavity flows for self-sustained oscillations. *International Journal of Heat and Fluid Flow* 32:1098–1110.
- Stergiopoulos N., Meister J.J., and Westerhof N. 1996. Determinants of stroke volume and systolic and diastolic aortic pressure. *American Journal of Physiology-Heart and Circulatory Physiology* 270:H2050–H2059.
- Suga H., Sagawa K., and Shoukas A.A. 1973. Load independence of the instantaneous pressure-volume ratio of the canine left ventricle and effects of epinephrine and heart rate on the ratio. *Circulation research* 32:314–322.
- Sun P., Wang J., Zhao S., Yang Z., Tang Z., Ravindra N., Bradley J., Ornato J.P., Peberdy M.A., and Tang W. 2018. Improved outcomes of cardiopulmonary resuscitation in rats treated with vagus nerve stimulation and its potential mechanism. *Shock: Injury, Inflammation, and Sepsis: Laboratory and Clinical Approaches* 49:698–703.
- Taylor R., Kutz J.N., Morgan K., and Nelson B.A. 2018. Dynamic mode decomposition for plasma diagnostics and validation. *Review of Scientific Instruments* 89:053501.
- Tosato M., Yoshida K., Toft E., Nekrasas V., and Struijk J.J. 2006. Closed-loop control of the heart rate by electrical stimulation of the vagus nerve. *Medical and Biological Engineering and Computing* 44:161–169.
- Tracey K.J. 2002. The inflammatory reflex. *Nature* 420:853–859.

- Tu J.H., Rowley C.W., Luchtenburg D.M., Brunton S.L., and Kutz J.N. 2013. On dynamic mode decomposition: Theory and applications. arXiv preprint arXiv:1312.0041 .
- Ugalde H.M.R., Ojeda D., Le Rolle V., Andreu D., Guiraud D., Bonnet J.L., Henry C., Karam N., Hagege A., Mabo P., *et al.* 2015. Model-based design and experimental validation of control modules for neuromodulation devices. *IEEE Transactions on Biomedical Engineering* 63:1551–1558.
- Ursino M. 1998. Interaction between carotid baroregulation and the pulsating heart: a mathematical model. *American Journal of Physiology-Heart and Circulatory Physiology* 275:H1733–H1747.
- Wu Z., Tran A., Rincon D., and Christofides P.D. 2019. Machine-learning-based predictive control of nonlinear processes. part ii: Computational implementation. *AIChE Journal* 65:e16734.
- Yao Y. and Kothare M.V. 2020. Model predictive control of selective vagal nerve stimulation for regulating cardiovascular system. *In* 2020 American Control Conference (ACC), pages 563–568, IEEE.
- Yuan H. and Silberstein S.D. 2016. Vagus nerve and vagus nerve stimulation, a comprehensive review: part ii. Headache: The Journal of Head and Face Pain 56:259–266.
- Zabara J. 1985. Peripheral control of hypersynchronous discharge in epilepsy. *Electroencephalogr Clin Neurophysiol* 61:S162.
- Zannad F., De Ferrari G.M., Tuinenburg A.E., Wright D., Brugada J., Butter C., Klein H., Stolen C., Meyer S., Stein K.M., *et al.* 2015. Chronic vagal stimulation for the treatment of low ejection fraction heart failure: results of the neural cardiac therapy for heart failure (nectar-hf) randomized controlled trial. *European heart journal* 36:425–433.
- Zhang Y., Mowrey K.A., Zhuang S., Wallick D.W., Popovic Z.B., and Mazgalev T.N. 2002. Optimal ventricular rate slowing during atrial fibrillation by feedback av nodal-selective vagal stimulation. *American Journal of Physiology-Heart and Circulatory Physiology* 282:H1102–H1110.
- Zhao Z., Chen W., Wu X., Chen P.C., and Liu J. 2017. Lstm network: a deep learning approach for short-term traffic forecast. *IET Intelligent Transport Systems* 11:68–75.