

Data Driven Control of Vagus Nerve Stimulation for the Cardiac System

Andrew Branen¹, Yuyu Yao², Mayuresh Kothare², Babak Mahmoudi³, and Gautam Kumar^{1,*}

¹Department of Chemical and Materials Engineering, San José State University, San José, CA, USA

²Department of Chemical and Biomolecular Engineering, Lehigh University, Bethlehem, PA, USA

³Department of Biomedical Informatics, School of Medicine, Emory University, Atlanta, GA, USA

1 Abstract

Vagus nerve stimulation is an emerging therapy that seeks to offset pathological conditions by electrically stimulating the vagus nerve through cuff electrodes, where an electrical pulse is defined by several parameters such as pulse amplitude, pulse width, and pulse frequency. This electroceutical therapy has been approved for epilepsy, and treatment resistant depression. Currently, vagus nerve stimulation is under investigation for the treatment of heart failure, heart arrhythmia, hypertension, and gastric motility disorders. Through several clinical trials that sought to assess vagus nerve stimulation for the treatment of heart failure, stimulation parameters were determined heuristically and the results were left inconclusive, which has led to the suggestion of using a closed-loop approach to optimize the stimulation parameters. A recent investigation has demonstrated highly specific control of cardiac physiology by selectively activating different fibers in the vagus nerve. When multiple locations and multiple stimulation parameters are considered for optimization, the design of closed-loop control becomes considerably more challenging. To address this challenge, we investigated a data-driven control scheme for both modeling and controlling the rat cardiac system. Using an existing *in silico* physiological model of a rat heart to generate synthetic input-output data, we trained a long short-term memory network (LSTM) to map the effect of stimulation on the heart rate and the blood pressure. The trained LSTM was utilized in a model predictive control framework to optimize the vagus nerve stimulation parameters for set point tracking of the heart rate and the blood pressure in closed-loop simulations. Additionally, we altered the underlying *in silico* physiological model to consider intra-patient variability, and diseased dynamics from increased sympathetic tone in designing closed-loop VNS strategies. Throughout the different simulation scenarios, we leveraged the design of the controller to demonstrate alternative clinical objectives. Our results show the controller can optimize stimulation parameters to achieve set-point tracking with nominal offset while remaining computationally efficient. Furthermore, we show a controller formulation that compensates for mismatch due to intra-patient variability, and diseased dynamics. This study demonstrates the first application and a proof-of-concept for using a purely data-driven approach for the optimization of vagus nerve stimulation parameters in closed-loop control of the cardiac system.

2 Introduction

Cardiovascular diseases are a prevalent health risk and financial burden. As an example, heart failure alone is projected to reach a total cost of \$70 billion in 2030 [Heidenreich et al., 2013]. Current pharmaceutical therapies lack adequate efficacy in treating cardiovascular diseases as demonstrated by the mortality rates [Benjamin et al., 2018, Savarese and Lund, 2017, Ovbiagele et al., 2013], which has motivated investigation efforts into alternative therapeutic approaches. Vagus nerve stimulation (VNS) has been identified and proposed as a potential therapy for a variety of cardiac conditions such as heart failure, atrial fibrillation, hypertension, and stroke [Capilupi et al., 2020]. VNS involves sending electrical pulses through a cuff electrode to the vagus nerve, with the electrical pulse characterized by several parameters such as pulse width, pulse amplitude, and

pulse frequency [Howland, 2014]. A major challenge present in VNS delivery involves recruitment of specific fibers in the vagal nerve, as not all fibers have the same effect on the functioning of the cardiac system. A recent study has demonstrated the ability to recruit specific fiber types through the use of different stimulation locations in a rat heart [Plachta et al., 2014].

Another challenge present in delivering VNS therapy involves the optimal selection of VNS parameters to achieve a desired physiological response. 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. The different conclusions could be attributed to the different operating regimes for each trial, leading to the suggestion of finding optimal VNS parameters in future trials to clearly evaluate the efficacy of VNS [Asad and Stavrakis, 2019].

To address the challenge of finding the optimal VNS parameter selection, studies investigating the closed loop control of VNS stimulation has been accomplished in previous studies performed in sheep by using a proportional-integral controller design [Ugalde et al., 2015], and state space transition models [Romero-Ugalde et al., 2017]. Other studies have investigated using proportional-integral controller designs to control the heart rate of rats [Greenwald et al., 2016], pigs [Tosato et al., 2006], and dogs [Zhang et al., 2002]. The previously discussed controllers only controlled the heart rate, and did not control multiple physiological outputs such as heart rate and blood pressure. Further, these controller studies were not optimizing multiple input stimulation parameters, nor did they consider the possibility of different stimulation location sites. A recent *in silico* 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].

For this application of MPC, there are challenges revolving around the development and validation of *in silico* cardiac models. Often such an approach becomes a challenging task in selecting the correct dynamical equations that govern the cardiac system, and then fitting the parameters to those specific equations. Such tasks can be guided by a deep mechanistic understanding of the system, however the definition of the system can vary, as shown by the variety of *in silico* 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 *in silico* computational models do not include the necessary variables to account for physiological changes mediated through VNS, leaving a challenging task for their application in VNS parameter optimization. Further, such models may be difficult to validate in experiments due to the variability of fiber recruitment in the vagus nerve. Adding to the challenge of using full-scale *in silico* physiological models is the computational expense associated with simulating these models for real-time closed-loop control.

Data driven modeling techniques are a viable approach that address the challenges associated with the previously described *in silico* physiological models. A common approach for data driven modeling includes machine learning, which can learn a compact representation for nonlinear dynamics present in a variety of systems [Goodfellow et al., 2016]. There are no underlying assumptions about the data fed to train the network or the distribution of the data fed to the network. Together, these features have led to the widespread application of machine learning for the modeling of nonlinear dynamical systems. More specifically, recurrent neural networks (RNNs) are well suited for time series modeling shown by their state-of-the-art performance in challenging applications that include forecasting river flows [Sahoo et al., 2019], power usage in residential areas [Kong et al., 2017], and short-term traffic patterns [Zhao et al., 2017]. Of noteworthy importance in these applications, a RNN was consistently shown to give better predictive performance of the time-series data when compared to a simple feedforward neural network. Thus, the hidden state kept in a RNN improves the learning of intrinsic temporal symmetries that allow for more accurate predictions of future time-series data. Previously, long short-term memory (LSTM) neural networks which are a type of RNN, have been used to model the effect of current injection on a pyramidal neuron [Plaster and Kumar, 2019]. Additionally LSTMs have been shown to be less computationally expensive for function evaluation, which motivates their application in real-time control of the cardiac system. Finally, LSTMs have also been previously used in a MPC framework to demonstrate the control of a chemical plant [Wu et al., 2019].

In this simulation work, we demonstrate a computationally efficient data-driven closed-loop control scheme to control the heart rate (HR) and the mean arterial blood pressure (MAP) in an *in silico* physiological model of a rat heart [Yao and Kothare, 2020]. We develop and use a LSTM model in a MPC framework to optimize

the VNS parameters to achieve the desired set points in HR and MAP. We formulate unique optimization problems that consider different physiological contexts, and discuss their influence on the closed-loop controller performance. Our results show that our designed controllers lead to set point tracking with nominal offset for both HR and MAP. Next, we modify the *in silico* physiological model to exhibit intra-patient variability, caused by a unique concentration of neuronal fiber recruitment from the cuff electrode. To account for the intra-patient variability in the closed-loop control formulation, we include additional constraints in the optimization problem, thus enabling the control of the physiological variables without training another LSTM on the data from the modified *in silico* physiological model. We investigate this approach further by modifying the controlled *in silico* physiological model to demonstrate a case of an elevated sympathetic nervous system and a decreased vagal tone, similar to a diseased state. Through simulations, we show that our closed-loop control design can efficiently control HR and MAP for both the intra-patient and diseased state systems while using the same previously trained LSTM in designing the controller. Together, we demonstrate a novel computationally efficient data-driven closed-loop VNS design for modeling and controlling HR and MAP, which could potentially be translated in animal experiments for real-time control of the cardiac system.

3 Results

3.1 Data-driven mapping of VNS parameters to the heart rate and mean arterial blood pressure

We recently developed a novel data-driven machine learning-based computational modeling approach to map the effect of VNS on the heart rate (HR) and the mean arterial blood pressure (MAP). Briefly, we used a published computational model of the rat cardiac system [Yao and Kothare, 2020] to generate synthetic data by varying three VNS locations, pulse width, and stimulation frequency and measuring the effect of VNS parameters on HR and MAP (see the details of the *in silico* physiological model and the range of the VNS parameters in Materials and Methods section 5.1). We then trained several neural networks, including vanilla recurrent neural networks (RNNs) and long short-term memory (LSTM) network, on this synthetic data by systematically varying the model hyperparameters, such as the number of hidden layers, hidden-state dimensionality, and activation functions, and compared the performance of the models in predicting HR and MAP in response to VNS. Our results showed that a single hidden layer LSTM network with a hidden-state dimensionality of 10 and a hyperbolic tangent activation function led to the best performance. Although a manuscript on this modeling has been accepted in IEEE EMBC 2021 for oral presentation, for the completeness of this manuscript, we provide the details of our modeling approach and the comparison of various trained neural networks in the Materials and Methods section 5.4.

Figure 1 shows a specific prediction of HR and MAP by our LSTM model and its comparison with the *in silico* physiological model (ground truth data), where we have varied the VNS parameters after every 50 cardiac cycles. As shown in this figure, the LSTM model is able to predict the response of the *in silico* physiological model with high accuracy, following the selection of time-varying VNS parameters. Notably, this predictive task requires the trained LSTM to recursively predict all 200 cardiac cycles, given the initial HR and MAP and the VNS parameters for all 200 cardiac cycles. In this way, the LSTM has correctly learned the mapping of the VNS parameters to the physiological effect on HR and MAP. Importantly, our LSTM model took 2.10 s to predict HR and MAP for 100 cardiac cycles on a local PC machine (Intel(R) Core I7-9700 CPU 3.00 GHz with 16 GB RAM) compared to 19.99 s by the *in silico* physiological model.

In conclusion, our results demonstrate a computationally efficient data-driven machine learning approach to predict HR and MAP in response to VNS stimulation with high accuracy directly from the data. In the remainder of this paper, we will use our LSTM model to demonstrate how this model can be used to design a model-based optimal control strategy to regulate HR and MAP by optimizing the VNS parameters in a feedback-based closed-loop framework.

3.2 Model-based Optimal Control Framework to Optimize VNS Parameters for Cardiac System

As described in the previous section, the cardiac physiological system (the *in silico* physiological model of a rat cardiac system [Yao and Kothare, 2020] in our case) has three distinct VNS locations, with each location having two design parameters, i.e., the pulse width and the stimulation frequency. These different stimulation




Figure 1: LSTM predictions (shown in red) compared to the output from the *in silico* physiological model of the rat cardiac system (PM, shown in black) for 200 simulated cardiac cycles of **A**) heart rate (HR) and **B**) mean arterial blood pressure (MAP). For cycles 1-50 (light green), location 1 was active with a pulse width of 0.36 ms and pulse frequency of 31 Hz. For cycles 51-100 (light blue), location 2 was active with a pulse width 0.19 ms and a pulse frequency 17 Hz. For cycles 101-150 (yellow), location 3 was active with a pulse width of 0.42 ms and a pulse frequency of 47 Hz. The last 50 cycles did not have any locations active.

locations, along with the pulse width and stimulation frequency as the design parameters at each location, have a distinctive effect on cardiac physiology, such as heart rate (HR) and mean arterial blood pressure (MAP). The presence of several stimulation design parameters raises the challenge of optimizing these parameters to achieve a desired physiological response.




Figure 2: A Schematic showing our model-based predictive closed-loop optimal control framework for optimizing the vagus nerve stimulation (VNS) parameters to control multiple cardiac physiological biomarkers such as heart rate (HR) and mean arterial blood pressure (MAP). Here, a long short-term memory (LSTM) based recurrent neural network model has been used to make predictions of HR and MAP in response to VNS and optimize the VNS parameters within the model predictive control (MPC) framework.

In an experimental/clinical setting, the choice of these stimulation design parameters is typically accomplished manually based on the experience. In order to automate this process, we developed a closed-loop optimal control approach to optimize the VNS parameters to achieve the desired HR and MAP in an *in silico* physiological model of the rat cardiac system using the model predictive control (MPC) framework. Figure 2 shows a schematic of our closed-loop control framework with the model-based controller utilizing predictions from the

trained LSTM. Briefly, MPC is an optimal control strategy that uses a model of the system in designing and optimizing the control actions. Specifically, at the current time, the measurements of the outputs are obtained from the system, and a model of the system is used to predict the system's outputs in the future over a specified time horizon (prediction horizon) in terms of the current and future control actions over a specified time horizon (control horizon). Then an optimization problem is formulated and solved to minimize the difference between the predicted model outputs and the desired outputs over the prediction horizon by optimizing the VNS parameters over the control horizon. The first computed optimal control action is implemented on the system, and the prediction and control horizons are receded by one time-step in the future. This process is repeated until the system is driven to the desired outputs. One of the significant advantages of using MPC over other optimal control approaches in controlling the cardiac system is its flexibility in incorporating physiological constraints and uncertainties explicitly in the optimization problem formulation.

To demonstrate our approach, we considered the *in silico* physiological model of the rat cardiac system [Yao and Kothare, 2020] as the physiological cardiac system to be controlled. We adapted the MPC framework to design and optimize the VNS parameters — pulse width, stimulation frequency, and three VNS locations (i.e., 6 input parameters) — to control HR and MAP in the *in silico* physiological model of the rat cardiac system. To design our model-based controller, we used the LSTM model presented in the previous section as a predictive model of the physiological system and formulated the following optimization problem to be solved at each cardiac cycle:

$$\begin{aligned} \min_{\vec{u}(k+j-1|k), j=1, 2, \dots, N_c-1} \quad & \mathbb{J}(k) \\ \text{s.t.} \quad & \vec{u}_{min} \leq \vec{u}(k+j-1|k) \leq \vec{u}_{max}, j \in [1, N_c] \\ & \hat{\vec{y}}(k+i|k) = f_{NN}(\hat{\vec{y}}(k+i-1|k), \vec{u}(k+i-1|k)), i \in [1, N_p] \\ & \hat{\vec{y}}(k|k) = \vec{y}_0(k) \end{aligned} \quad (1)$$

Here, $f_{NN}(\cdot, \cdot)$ denotes the LSTM model described in the previous section, $\hat{\vec{y}}(k+i|k)$ is the vector of predicted physiological outputs (i.e., HR and MAP) from the LSTM model at the discrete time (i.e., cardiac cycle) $k+i$, given that the outputs are measured from the physiological system at the current time k , N_c and N_p are the control horizon and prediction horizon, respectively, and $\vec{u}(k+i|k)$ denotes the vector of VNS input parameters (i.e., the pulse width and the stimulation frequency at three VNS locations) at a future time $k+i$, given that the outputs are measured from the physiological system at the current time k . $\vec{y}_0(k)$ is the vector of measured HR and MAP from the physiological system at the current time (or cardiac cycle) k and $\mathbb{J}(k)$ is the user defined scalar cost function. In the optimization problem (1), the first constraint is the upper and lower bounds on the VNS parameters over the control horizon. The second constraint states that the evolution of outputs over the prediction horizon must satisfy the LSTM model dynamics. The last constraint states that the current outputs of LSTM model are same as the measured outputs from the physiological system at time k .

3.2.1 Sparsity Promoted Closed-Loop VNS Design for Cardiac System

As noted in the description of the optimization problem (1), the manipulated (or designed) VNS inputs consist of three VNS locations, with each location having two VNS design parameters, the pulse width, and the stimulation frequency. The presence of discrete or integer variables (locations of VNS) requires the formulation of the optimization problem (1) in the form of well-known mixed-integer programming, which is typically computationally expensive to solve [Trespalacios and Grossmann, 2014]. To keep our control strategy computationally efficient while achieving a similar performance compared to the mixed-integer programming solution, we formulated the cost function in the optimization problem (1) by introducing a L_1 penalty on the VNS parameters as follows:

$$\mathbb{J}(k) = \sum_{i=1}^{N_p} (\vec{r}(k+i) - \hat{\vec{y}}(k+i|k))^T Q (\vec{r}(k+i) - \hat{\vec{y}}(k+i|k)) + \lambda \sum_{j=1}^{N_c} \|\vec{u}(k+j-1|k)\|_1. \quad (2)$$

Here, the first term in the sum is a weighted quadratic cost defining the error between the target output (or reference) to be achieved by the controller and the predicted output from the LSTM model over the prediction horizon N_p . $\vec{r}(k+i) \in \mathbb{R}^{2 \times 1}$ is a vector of the time varying target output (HR and MAP) at discrete time $k+i$ and $Q \in \mathbb{R}^{2 \times 2}$ is the weighting matrix emphasizing the importance of specific outputs. The second term in the sum is a L_1 cost on the manipulated inputs (the pulse width and stimulation frequency at three VNS locations)

over the control horizon N_c . The parameter λ defines the importance of these two terms in the cost function. We defined the L_1 cost as

$$\sum_{j=1}^{N_c} \|\vec{u}(k+j-1|k)\|_1 = \sum_{j=1}^{N_c} \sum_{i=1}^6 |\vec{u}_i|. \quad (3)$$

Here, $|\cdot|$ represents the absolute value of the argument. L_1 penalty in the cost function is well-known to introduce sparsity in the optimization variables, both in the control community and machine-learning community [Vidaurre et al., 2013].

To solve the optimization problem (1) with the cost function (2), we set $N_p = 10$, $N_c = 5$, and $\lambda = 0.001$. The initial measured HR and MAP from the physiological rat cardiac model were 409 (bpm) and 138 (mmHg), respectively. We selected the heart rate going from 392 (bpm) to 346 (bpm) followed by 393 (bpm), and mean arterial blood pressure going from 111 (mmHg) to 144 (mmHg) followed by 125 (bpm) with each pair lasting for 50 cardiac cycles as the desired set point trajectory to be followed by the controller.

Figure 3 shows our simulation results on the sparsity promoted closed-loop VNS design to control HR and MAP simultaneously. In Figures 3A and 3B, we note that the designed controller efficiently drove the physiological system to the desired time-varying set points with minimal to no steady-state offset for HR and MAP, respectively. Figures 3C and 3D show the optimized pulse width and stimulation frequency, respectively, delivered to the physiological system at each cardiac cycle. As noted in Figures 3C and 3D, the controller activated locations 1 and 2 in the first and the last 50 cardiac cycles and locations 2 and 3 in the second 50 cardiac cycles to achieve the desired set points in each 50 cycles period.

We also note in Figures 3C and 3D that the controller actions are oscillating slightly. A possible reason for this behavior may be due to influence from the MAP variable, which exhibits fluctuations even at the steady-state (see Figure 1). Since there is no cost term forcing a constant solution on the controller inputs, the controller attempts to compensate for this fluctuation by altering its actions, thus leading to similar fluctuations in the controller actions. Additionally, it's worth noting that the LSTM had a higher error for MAP predictions (compared to HR) for this reason.

In summary, we showed a data-driven model-based closed-loop control strategy to optimize the VNS parameters with multiple stimulation location sites, leading to a sparse selection of VNS parameters in controlling cardiac physiology efficiently. This strategy may find applications in developing efficient closed-loop VNS therapy for cardiac diseases by minimizing the side effects of specific stimulation locations. Although we have not considered all possible physiological constraints in formulating our optimization problem, the framework is general enough to introduce additional physiological constraints within the optimization problem.

3.2.2 Minimum Energy-based Closed-Loop VNS Design for Cardiac System

In the previous section, we presented a sparsity-promoted closed-loop VNS strategy that selectively stimulated specific VNS locations to control the heart rate (HR) and mean arterial blood pressure (MAP). Although this strategy could potentially minimize the location-specific side effects induced by VNS, the bang-bang nature of this strategy (i.e., delivering VNS at the maximum allowable pulse width and frequency) may damage tissues over a longer period of VNS application. To minimize the duration of the applied VNS at the maximum allowable pulse width and frequency, we designed a closed-loop optimal control strategy by minimizing the stimulation energy required to drive HR and MAP to the desired set points. To implement this strategy, we chose the following cost function in the optimization problem (1):

$$\mathbb{J}(k) = \sum_{i=1}^{N_p} (\vec{r}(k+i) - \hat{y}(k+i|k))^T Q (\vec{r}(k+i) - \hat{y}(k+i|k)) + \lambda \sum_{j=1}^{N_c} \|\vec{u}(k+j-1|k)\|_2^2. \quad (4)$$

Here, the first term in the sum is a weighted quadratic cost defining the error between the target output (or reference) to be achieved by the controller and the predicted output from the LSTM model over the prediction horizon N_p . $\vec{r}(k+i) \in \mathbb{R}^{2 \times 1}$ is a vector of the time varying target output (HR and MAP) at discrete time $k+i$ and $Q \in \mathbb{R}^{2 \times 2}$ is the weighting matrix emphasizing the importance of specific outputs. The second summation term is a L_2 cost on the manipulated inputs (the pulse width and stimulation frequency at three VNS locations) over the control horizon N_c . The parameter λ defines the importance of these two terms in the cost function. It should be noted that the cost function (4) is same as the sparsity promoted cost function (2) described in the




Figure 3: Sparsity-promoted closed-loop control of the heart rate (HR) and the mean arterial blood pressure (MAP) of the *in silico* physiological model of the rat cardiac system. (A) and (B) show the closed-loop control response from the *in silico* physiological model (shown in red dots) and the target (or reference) for the controller (shown in black line) for HR and MAP, respectively. (C) and (D) show the pulse width and the stimulation frequency, respectively, designed by the controller at each VNS location.

previous section except we replaced the L_1 penalty in (2) by a L_2 penalty. We defined the L_2 cost as

$$\sum_{j=1}^{N_c} \|\vec{u}(k+j-1|k)\|_2^2 = \sum_{j=1}^{N_c} \vec{u}^T(k+j-1|k) R \vec{u}(k+j-1|k). \quad (5)$$

Here, $(\cdot)^T$ represents a vector transpose and $R \in \mathbb{R}^{6 \times 6}$ is the weighting matrix defining the importance of individual inputs. L_2 penalty in the cost function is well-known to achieve a minimum energy solution by suppressing large amplitude control actions [Kwakernaak and Sivan, 1972].

To solve the optimization problem (1) with the cost function (4), we set $N_p = 10$, $N_c = 5$, and $\lambda = 0.001$. We set the design parameters Q and R to an identity matrix of appropriate dimensions. The initial measured HR and MAP from the physiological rat cardiac model were 409 (bpm) and 138 (mmHg), respectively. We selected the heart rate going from 392 (bpm) to 346 (bpm) followed by 393 (bpm), and mean arterial blood pressure going from 111 (mmHg) to 144 (mmHg) followed by 125 (mmHg) with each variable lasting for 50 cardiac cycles as the desired set point trajectory to be followed by the controller.

Figure 4 shows our simulation results on the minimum energy-based closed-loop VNS design in controlling HR and MAP of the physiological model of rat cardiac system simultaneously. In Figures 4A and 4B, we show that the designed controller can efficiently drive the physiological system to the desired time-varying set points with minimal to no offset for HR and MAP, respectively. We note that the output performance of the designed controller is similar to the sparsity promoted design described in the previous section, with slight oscillations in MAP. Figures 4C and 4D show the optimized pulse width and stimulation frequency, respectively, delivered to the physiological model at each cardiac cycle. As noted in Figures 4C and 4D, the controller activated all the locations to achieve the desired set points in each 50 cycle period as opposed to the sparsity-promoted design shown in Figures 3C and 3D. More importantly, the designed controller selected minimum values on the VNS parameters (pulse width and frequency) required to drive HR and MAP to the desired set points by minimizing the input energy.




Figure 4: Minimum energy-based closed-loop control of the heart rate (HR) and the mean arterial blood pressure (MAP) of the *in silico* physiological model of the rat cardiac system. (A) and (B) show the closed-loop control response from the *in silico* physiological model (shown in red dots) and the target (or reference) for the controller (shown in black line) for HR and MAP, respectively. (C) and (D) show the pulse width and the stimulation frequency, respectively, designed by the controller at each VNS location.

In summary, we showed a data-driven model-based closed-loop control strategy to optimize the VNS parameters with multiple stimulation location sites, leading to a minimum energy based selection of VNS parameters in controlling cardiac physiology efficiently. This strategy may find applications in developing efficient closed-loop VNS therapy for cardiac diseases by enhancing the battery life of the VNS device. Moreover, this strategy could easily be combined with the sparsity promoted closed-loop VNS design by including an additional L_1 penalty term in the cost function (4).

3.2.3 Minimum Overshoot based Closed-Loop VNS Design for Cardiac System

In the previous two closed-loop VNS designs, we noted that the controller actions oscillate around the steady-state values (see Figures 3C, 3D, 4C, and 4D). To eliminate or reduce the oscillatory behavior in the controller actions, we included an additional term in the cost function that penalizes the deviation in inputs, when comparing the current and future stimulation parameters to the previously implemented optimal stimulation parameters to the physiological system. Particularly, we formulated the cost function $\mathbb{J}(k)$ in the optimization problem (1) as follows:

$$\begin{aligned} \mathbb{J}(k) &= \sum_{i=1}^{N_p} (\vec{r}(k+i) - \hat{\vec{y}}(k+i|k))^T Q (\vec{r}(k+i) - \hat{\vec{y}}(k+i|k)) \\ &+ \lambda_1 \sum_{j=1}^{N_c} \|\vec{u}(k+j-1|k)\|_1 + \lambda_2 \sum_{j=1}^{N_c} \|\vec{u}(k+j-1|k) - \vec{u}_{k-1}\|_1. \end{aligned} \quad (6)$$

Here, the first two terms of the cost function (6) are same as the sparsity promoted cost function (2), as described in Section 3.2.1. The last term in the cost function (6) is the difference between the current and future control

inputs over the control horizon N_c and the last applied optimal control input \vec{u}_{k-1} on the physiological model of the rat cardiac system (i.e., the physiological system). λ_1 and λ_2 are the weighting parameters.

By including the difference in controller action term in the cost function 6, we forced the controller actions to take values close to the last implemented optimal control action on the physiological system. As a result of the constant controller actions, the physiological variables should exhibit minimal overshoot as they gradually reach the setpoint. Additionally, any fluctuations in HR or MAP at the setpoint are the result of intrinsic model dynamics. The inclusion of the L1 cost term in the cost function 6 serves to maintain a sparse solution.

For the closed-loop simulation, we set $N_c = 5$, $N_p = 10$, $\lambda_1 = 0.001$, and $\lambda_2 = 0.00005$. It should be noted that different selections of N_c and N_p can lead to similar performance, provided that the weights (λ_1 , λ_2) are appropriately selected. The initial values of HR and MAP were 409 (bpm) and 138 (mmHg), respectively. The values for Q and R were again an identity matrix of the appropriate dimensions. The duration and set points for both HR and MAP were the same as mentioned in previous sections.

Figures 5A and 5B show the controller performance in driving HR and MAP of the physiological model to the desired set points. As shown here, adding an input difference term (the last term in the right hand side of the cost function (6)) in the cost function (6) minimized the previously seen fluctuations in MAP and led to a more smooth behavior in both HR and MAP compared to the sparsity-promoted and minimum energy-based strategies. Notably, this strategy resulted in no overshoot in the last set point for HR and a minimal overshoot for the other two target set points for HR.




Figure 5: Minimum overshoot-based closed-loop control of the heart rate (HR) and the mean arterial blood pressure (MAP) of the *in silico* physiological model of the rat cardiac system. (A) and (B) show the closed-loop control response from the *in silico* physiological model (shown in red dots) and the target (or reference) for the controller (shown in black line) for HR and MAP, respectively. (C) and (D) show the pulse width and the stimulation frequency, respectively, designed by the controller at each VNS location.

The impact of the input difference term in the cost function (6) becomes more pronounced in the controller actions, as depicted in Figures 5C and 5D for VNS pulse width and frequency, respectively. We particularly found minimal to no oscillations in the controller actions as compared to the sparsity promoted and minimum energy-based strategies. Interestingly, the minimum overshoot based controller does not achieve the sparse solution at all target set points, evidenced by having all three locations active for the last set point.

In summary, we showed a data-driven model-based closed-loop control strategy to optimize the VNS parameters with multiple stimulation location sites, leading to smooth selection of VNS parameters in controlling

cardiac physiology efficiently. This strategy may find applications in developing efficient closed-loop VNS therapy for cardiac diseases by avoiding overshoots/fluctuations in the controlled outputs.

3.3 Controlling Individual Patient Mismatch

So far, we have only considered scenarios where the developed LSTM model accurately predicted the heart rate (HR) and the mean arterial blood pressure (MAP) of the *in silico* physiological model of the rat cardiac system in response to VNS. We then used this LSTM model to design control strategies to optimize the VNS parameters to achieve a desired HR and MAP in the *in silico* physiological model. We demonstrated that the designed controller could efficiently drive HR and MAP to the desired set points as long as the LSTM predictions are reasonably accurate compared to the *in silico* physiological model (i.e., there is a minimal or no model mismatch between the LTSM and the *in silico* physiological model). In general, it may not always be feasible to obtain an accurate system model, mainly if experimental/clinical data are used to develop a system model. Should a reasonably good quantitative model be developed from the available experimental data, there may always be subject-to-subject variability in response to VNS which could potentially result in a significant model mismatch between the developed model and the physiological system. This model mismatch could lead to degraded performance of the designed controller in driving HR and MAP to the desired set points.

To illustrate this point, we created a specific case where we modified the neuronal fiber recruitment concentrations at each VNS location in the *in silico* physiological model of the rat cardiac system [Yao and Kothare, 2020]. Specifically, we changed the neuronal fiber concentrations at each location ($C_{i,j}$) whose impact affects the output of Equation (20), and ultimately leads to a different influence of VNS on the cardiac variables when compared to the base *in silico* physiological model. Figure 6 shows the predicted HR and MAP from the modified *in silico* physiological model in response to various VNS parameters and their comparison with the LSTM predicted values and the original *in silico* physiological model. As shown here, although the LSTM model did not predict the physiological HR and MAP quantitatively in response to VNS for the modified *in silico* physiological model, it indeed predicted the HR and MAP qualitatively. As noted in Figure 6 for the cardiac cycles between 150 to 200, the LSTM predicted both HR and MAP quantitatively with reasonable accuracy in the absence of VNS, which we expect since the modified *in silico* physiological model only differs from the original *in silico* physiological model [Yao and Kothare, 2020] in terms of how it responds to VNS.




Figure 6: Comparison between the physiological outputs predicted from the modified *in silico* physiological model (IP, shown in blue), the *in silico* physiological model (PM, shown in black) and LSTM model (shown in red) in response to the vagus nerve stimulation (VNS) for 200 simulated cardiac cycles. **(A)** shows the heart rate (HR, bpm) and **(B)** shows the mean arterial blood pressure (MAP, mmHg). For cycles 1-50 (light green), location 1 was active with a pulse width of 0.36 ms and pulse frequency of 31 Hz. For cycles 51-100 (light blue), location 2 was active with a pulse width 0.19 ms and a pulse frequency 17 Hz. For cycles 101-150 (yellow), location 3 was active with a pulse width of 0.42 ms and a pulse frequency of 47 Hz. The last 50 cycles did not have any locations active.

We then used the same control formulation as described in the sparsity-promoted control design (see Section 3.2.1) but with the modified *in silico* physiological model in the closed-loop instead of the original *in silico*

physiological model. Figure 7 shows the closed-loop performance of the designed controller. As noted in Figures 7A and 7B, the designed controller failed to drive HR and MAP to the desired set points. Specifically, the designed controller led to significant steady-state offset at most of the set-points. Interestingly, the sparsity promoted controller drove MAP to the desired set point for the cardiac cycles 150-300. This surprising result highlights one of the inherent advantages of model-based predictive control, wherein the feedback to the model provides inherent integral action to compensate for offsets.




Figure 7: Sparsity-promoted closed-loop control of the heart rate (HR) and the mean arterial blood pressure (MAP) of the modified *in silico* physiological model of the rat cardiac system. (A) and (B) show the closed-loop control response from the modified *in silico* physiological model (shown in blue dots) and the target (or reference) for the controller (shown in black line) for HR and MAP, respectively. (C) and (D) show the pulse width and the stimulation frequency, respectively, designed by the controller at each VNS location.

To eliminate the steady-state offsets in the closed-loop performance, we adopted an approach of nonlinear offset-free control [Morari and Maeder, 2012] and reformulated our MPC optimization problem as follows:

$$\begin{aligned}
 & \min_{\vec{u}(k+j|k), j=0,1,\dots,N_c-1} \mathbb{J}(k), \\
 & \text{s.t.} \quad \vec{u}_{min} \leq \vec{u}(k+j-1|k) \leq \vec{u}_{max}, j \in [1, N_c], \\
 & \quad \hat{\vec{y}}(k+i|k) = f_{NN,aug}(\hat{\vec{y}}(k+i-1|k), \vec{d}(k), \vec{u}(k+i-1|k)), i \in [1, N_p], \\
 & \quad \hat{\vec{y}}(k|k) = \vec{y}_0(k) \\
 & \quad \vec{r}(k) = f_{NN,aug}(\hat{\vec{y}}_s, \vec{d}(k), \vec{u}_s).
 \end{aligned} \tag{7}$$

Here, $\vec{u}(k+j|k) \in \mathbb{R}^{6 \times 1}$ is a vector of VNS parameters at the discrete time $k+j$, given that HR and MAP were measured from the modified physiological system at time k . \vec{u}_{min} and \vec{u}_{max} are the minimum and maximum bound on the VNS parameters, respectively. $\hat{\vec{y}}(k+i|k) \in \mathbb{R}^{2 \times 1}$ is a vector of outputs (HR and MAP) predicted by the LSTM model at the discrete time $k+i$, given that HR and MAP were measured from the modified *in silico* physiological model (the physiological system to be controlled in our case) at time k . The function $f_{NN,aug}(\cdot)$ represents the LSTM model augmented with a disturbance model (see Equation (8b)). N_p and N_c are the prediction and control horizon, respectively. $\vec{y}_0(k) \in \mathbb{R}^{2 \times 1}$ is the output (HR and MAP) measured from the *in*

silico physiological model at time k in response to the optimal VNS parameters $\vec{u}(k-1|k-1)$. $\vec{r}(k) \in \mathbb{R}^{2 \times 1}$ is a vector consisting of the reference or target HR and MAP. The vector $\vec{d}(k) \in \mathbb{R}^{2 \times 1}$ is a disturbance term which models the mismatch between the LSTM and the *in silico* physiological model at time k (see Equation (8c)). For a given $\vec{r}(k)$ and $\vec{d}(k)$, $\hat{\vec{y}}_s$ and \vec{u}_s are the steady-state solution of the disturbance-augmented LSTM model.

The optimization problem (7) is notably different from the previously formulated optimization problem (1) in two ways. First, this new optimization problem aims to minimize a scalar cost function $\mathbb{J}(k)$ with respect to two steady-state variables $\vec{u}_s \in \mathbb{R}^{6 \times 1}$ and $\hat{\vec{y}}_s \in \mathbb{R}^{2 \times 1}$ in addition to the VNS parameters $\vec{u}(k+j-1|k) \in \mathbb{R}^{6 \times 1}$ over the control horizon N_c used in the optimization problem (1). And second, we introduced an additional constraint in the optimization problem formulation which ensures that the current target ($\vec{r}(k)$) is returned when the disturbance-augmented LSTM is evaluated at the steady-state optimized variables $\hat{\vec{y}}_s$ and \vec{u}_s for a given $\vec{d}(k)$.

To account for the model mismatch between the output predictions from the LSTM model and the modified *in silico* physiological model (the physiological system to be controlled in our case) in the computation of optimal control actions, we introduced the following disturbance model:

$$f_{NN,aug}(\hat{\vec{y}}(k), \vec{d}(k), \vec{u}(k)) = f_{NN}(\hat{\vec{y}}(k), \vec{u}(k)) + \vec{d}(k) \quad (8a)$$

$$\vec{d}(k+1) = \vec{d}(k) + L_d \vec{\epsilon}(k) \quad (8b)$$

$$\vec{\epsilon}(k) = f_{NN,aug}(\hat{\vec{y}}(k-1), \vec{d}(k-1), \vec{u}(k-1)) - \vec{y}(k) \quad (8c)$$

Here, the augmented LSTM function $f_{NN,aug}(\hat{\vec{y}}(k), \vec{d}(k), \vec{u}(k))$ described by Equation (8a) is the sum of the LSTM function $f_{NN}(\hat{\vec{y}}(k), \vec{u}(k))$ and the disturbance variable \vec{d}_k . The disturbance model (8c) updates the model mismatch between the measured outputs (HR and MAP) from the modified *in silico* physiological model and the predicted outputs from the LSTM model at the next time step by integrating a scaled error. Equation (8c) computes the model mismatch between the predicted outputs from the LSTM model and the measured outputs (HR and MAP) $\vec{y}(k)$ from the modified *in silico* physiological model at time k . $L_d \in \mathbb{R}^{2 \times 1}$ is a vector of constant gains.

At the beginning of the simulation, the disturbance variable $\vec{d}(0)$ was initialized to zero. To select the gain, L_d on the disturbance update model (8c), the authors in [Morari and Maeder, 2012] computed L_d based on the system model (well-behaved differential equation model in this case) linearized about the origin using the steady-state Kalman filter algorithm. In principle, the observer must be nominally asymptotically stable and satisfy: $L_d(\epsilon = 0) = 0$. In our case, such linearization was not possible for the LSTM model. Therefore, we selected the gain L_d based on a trial-and-error procedure. Particularly, we ran different trials of closed-loop simulations and examined the plot of the disturbance evolution over time. Typically, if a sufficiently inappropriate value was selected, the solver ceased to converge, often requiring a decrease in the L_d values. If the values were within the range required for the convergence, then the effect of increasing those values led to faster convergence to the steady-state solution. However, if the selected values were too large, then the error signal (Equation 8c) was significantly emphasized, and the controller became unstable, often causing a failure in the convergence of the optimization solver. If the values of L_d were decreased, the controller took more iterations or a longer time to compensate for the offset. Based on this approach, we selected the following value of the gain L_d :

$$L_d = \begin{bmatrix} 0.06 & 0 \\ 0 & 0.05 \end{bmatrix}$$

We formulated the following cost function to be minimized within the optimization problem (7):

$$\mathbb{J}(k) = \sum_{i=1}^{N_p} (\hat{\vec{y}}_s - \hat{\vec{y}}(k+i|k))^T Q (\hat{\vec{y}}_s - \hat{\vec{y}}(k+i|k)) + \sum_{j=0}^{N_c-1} (\vec{u}_s - \vec{u}(k+j|k))^T R (\vec{u}_s - \vec{u}(k+j|k)). \quad (9)$$

Here, the first summation term is the squared difference between the LSTM model predictions and the optimized variable $\hat{\vec{y}}_s$. Note that this optimized variable represents the new target set point for the controller that accounts for the model mismatch. The second summation term is the squared difference between the control input and the optimized variable \vec{u}_s . Again, note that this additional optimized variable is the actuation that together with $\hat{\vec{y}}_s$ satisfies the last constraint in the optimization problem (7) for a given $\vec{r}(k)$ and $\vec{d}(k)$. Consistent with the

previous cost functions, Q and R were set to an identity matrix of the appropriate dimensions. The constrained optimization problem (Equation 7) together with the disturbance update (Equation 8b), and the cost function (Equation 9) enable the designed controller to achieve the steady-state offset free control of HR and MAP.

Figure 8 shows the closed-loop performance of our designed controller. For the first 150 cardiac cycles, the set point $\bar{r}(k)$ was set to 356 (bpm) and 150 (mmHg) for HR and MAP, respectively. For the next 150 cardiac cycles, the set point was set to 393 (bpm) and 129 (mmHg) for HR and MAP, respectively. Finally, for the final 150 cardiac cycles, the set point was set to 377 (bpm) and 143 (mmHg) for HR and MAP, respectively. The initial values were 409 (bpm) and 138 (mmHg) for HR and MAP, respectively. For the controller, we set $N_p = 20$ and $N_c = 10$ for this simulation. For each set point, we ran the simulation for 100 cardiac cycles longer (150 cardiac cycles as opposed to 50 cardiac cycles in previous closed-loop simulations) to ensure the convergence of the offset-free closed-loop control formulation.




Figure 8: Offset-free closed-loop control of the heart rate (HR) and the mean arterial blood pressure (MAP) of the modified *in silico* physiological model of the rat cardiac system in the presence of intra-patient variability in the VNS response to HR and MAP. (A) and (B) show the closed-loop control response from the modified *in silico* physiological model (shown in blue dots) and the target (or reference) for the controller (shown in black line) for HR and MAP, respectively. (C) and (D) show the pulse width and the stimulation frequency, respectively, designed by the controller at each VNS location.

As shown in Figure 8, the controller is able to achieve offset-free control within the first 50 cardiac cycles of each set point change. As the controller approaches each set point, there is an overshoot of each target for the first 20 cardiac cycles of the set point change. Notably, the optimized VNS parameters are generally smooth and constant near the steady-state.

Concise summarized, we developed a closed-loop optimal control approach that accounts for the intra-patient variability of vagus nerve stimulation (VNS) response to the heart rate (HR) and the mean arterial blood pressure (MAP) in optimizing the VNS parameters for controlling HR and MAP. This specific closed-loop control formulation is appropriate in clinical/experimental settings where intra-patient variability is critical in designing VNS strategies, and the response of the implanted VNS device is sufficiently different from the responses used to train the LSTM model. A notable limitation of this approach is that it requires the controller's task to be a set point tracking problem due to the offset-free constraint formulation.

3.4 Closed-Loop Control of Overactive Sympathetic System

To demonstrate whether the designed control strategy presented in the previous section could be used to regulate the heart rate (HR) and the mean arterial blood pressure (MAP) in a diseased case, we constructed an example of overactive sympathetic pathway case in the cardiac system. Specifically, we modified the parameters of the original *in silico* physiological model [Yao and Kothare, 2020] in a way such that the response to the sympathetic nervous system dominated over the parasympathetic response. These changes led to a higher resting heart rate, and a higher blood pressure, both of which are consistent with the behavior of some cardiovascular diseases [Malpas, 2010]. We provide the specific parameter changes in Table 2. Particularly, this autonomic nervous system imbalance promotes heart failure, and vagal nerve stimulation has been suggested as a treatment therapy [Kishi, 2012]. Similarly, the dominance of the sympathetic system corresponds to the increased blood pressure present in hypertension [Carthy, 2014]. Additionally, hypertension has been shown to occur in parallel with other cardiovascular diseases [Palatini and Julius, 2009]. Figure 9 shows the difference between this new *in silico* pathological behavior model and the original *in silico* physiological model. Notably, in the absence of VNS, the *in silico* pathological behavior model does not return to the same value as the original *in silico* physiological model, which is different from the intra-patient mismatch shown previously (compare with Figure 6).




Figure 9: Comparison between the physiological outputs predicted from the modified *in silico* physiological model of an overactive sympathetic system (Sympathetic, shown in purple), the original *in silico* physiological model (PM, shown in black) and LSTM model (shown in red) in response to the vagus nerve stimulation (VNS) for 200 simulated cardiac cycles. **(A)** shows the heart rate (HR, bpm) and **(B)** shows the mean arterial blood pressure (MAP, mmHg). For cycles 1-50 (light green), location 1 was active with a pulse width of 0.36 ms and pulse frequency of 31 Hz. For cycles 51-100 (light blue), location 2 was active with a pulse width 0.19 ms and a pulse frequency 17 Hz. For cycles 101-150 (yellow), location 3 was active with a pulse width of 0.42 ms and a pulse frequency of 47 Hz. The last 50 cycles did not have any locations active.

To drive the HR and MAP to the desired set points, we applied the control formulation developed in the previous section (Equations 7, 8a, 8b, 8c, 9). Figure 10 shows the closed-loop controller performance for our designed controller. For this simulation, we set the initial heart rate to 452 bpm and the initial blood pressure to 152 (mmHg). For the first 300 cycles, the controller's target was 356 (bpm) and 150 (mmHg) for HR and MAP, respectively. The following 300 cardiac cycles had the target change to 393 (bpm) and 129 (mmHg), and the final 300 cardiac cycles the target was set to 377 (bpm) and 143 (mmHg). These are the same set points used for the intra-patient case (Figure 8). Regarding the controller, the controller parameters were set to the same previous values ($N_p = 20$ and $N_c = 10$). We set the disturbance gain vector L_d to

$$L_d = \begin{bmatrix} 0.06 & 0 \\ 0 & 0.018 \end{bmatrix}.$$

The length of simulation was considerably longer than the previous simulations (900 cardiac cycles compared to 450 cardiac cycles) because the disturbance variable, used for model mismatch compensation, took a longer

time to reach the saturation due to a large model mismatch in this case. Nevertheless, the designed controller drove both HR and MAP to the desired set points with no steady-state offsets.




Figure 10: Offset-free closed-loop control of the heart rate (HR) and the mean arterial blood pressure (MAP) in an *in silico* overactive sympathetic system model of the rat cardiac system. (A) and (B) show the closed-loop control response from the modified *in silico* physiological model (shown in purple dots) and the target (or reference) for the controller (shown in black line) for HR and MAP, respectively. (C) and (D) show the pulse width and the stimulation frequency, respectively, designed by the controller at each VNS location.

As noted in Figure 10B, the controller initially drives the MAP in the opposite direction of the target. This behavior is likely due to the significant model mismatch between the LSTM model and the modified *in silico* physiological model for the overactive sympathetic case. After about 50 cardiac cycles, the disturbance variable starts compensating for the mismatch and the MAP variable starts moving toward the target. After about 220 cardiac cycles, the disturbance variable has saturated leading to offset free control. As discussed previously, the MAP variable exhibits inherent fluctuations, and these fluctuations have grown larger with the new parameter selection. Even with the inherent model fluctuations, the MAP variable mean hits all of the set points and thus providing an offset-free control. Figure 10C and 10D show the optimal pulse width and stimulation frequency selected by the controller, respectively.

In summary, we modified the *in silico* physiological model of the rat cardiac system (the system to be controlled in our case) to exhibit increased responses in the sympathetic nervous system and decreased responses in the parasympathetic nervous system, and demonstrated that the offset-free controller formulation reached several set point tracking targets without a steady-state offset. This proof-of-concept suggests that a controller can be designed for cases where the diseased hemodynamics are considerably different from the hemodynamics used to train the LSTM or other neural network model to be used in designing a model-based optimal controller.

3.5 Impact of Prediction and Control Horizons on the Controller Performance and Computational Cost

In previous sections, we demonstrated various closed-loop optimal VNS designs for controlling the heart rate (HR) and the mean arterial blood pressure (MAP). In this section, we systematically evaluated the computational time required for simulating closed-loop control designs presented in the previous sections. Additionally,

we compared the effect of the prediction horizon N_p and the control horizon N_c on the computational time and the closed-loop performance. All the comparisons were performed on an INTEL(R) CORE I7-9700 CPU 3.00 GHz with 16.0 GB of RAM Desktop.

Table 1 shows a summary of the computational time required to simulate 100 cardiac cycles (shown in terms of Average Time/Cycle) in closed-loop under specific control policy along with the closed-loop performance in terms of the mean absolute steady-state error between the desired setpoint and the closed-loop controlled output computed over the last 10 cardiac cycles for each setpoint. We estimated the computational time by taking the total time to run the closed-loop simulation for 100 cardiac cycles and subtracting the time it took the full *in silico* physiological model (the cardiac system to be controlled in our case) to simulate the number of cycles and dividing this difference by the total number of cycles in the simulation. For the sake of comparison, we set the parameter λ , λ_1 , and λ_2 in the optimization problems to 0.001. As expected, when the control horizon N_c 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 as there were six additional parameters for the controller to optimize for every additional control horizon. Similarly, as the predictive horizon N_p increased, the time of optimization also increased as there were more predictions to be computed from the LSTM model. However, this additional computational time was not as significant compared to the increase in the control horizon N_c . It should be noted that there were 8 additional variables to be optimized in the case of offset-free VNS design and overactive sympathetic case due to the different optimization problem formulation.

To compare the closed-loop performance over different prediction and control horizons across the implemented control strategies, we computed the steady-state mean absolute error (SS MAE) by taking the average of the mean absolute difference between the controlled output and the target setpoints for the last 10 cycles of each setpoint. It should be noted that SS MAE represents the combined absolute mean error of both the heart rate and the mean arterial blood pressure. We noted in Table 1 that the performance of the controller decreased to the point of showing a clear steady-state offset from the setpoints for the control horizon $N_c = 1$. It is expected that an increase in both N_c or N_p should exhibit less offset. This trend is not entirely shown by the results in Table 1 because the weights of the cost function (i.e., λ , λ_1 , λ_2) were held constant. These results highlight how the selection of controller parameters can lead to different performances. Thus, the controllers with higher values of N_c or N_p likely need to be tuned to obtain comparable performance.

Based on the results shown in Table 1, we conclude that it is likely that selections with $N_c \geq 5$ and $N_p \geq 20$ would be difficult to implement in real-time due to large computational time. This leaves $N_c \leq 5$, and $N_p \leq 10$ as plausible choices for the real-time implementation of our closed-loop VNS designs. This analysis is based on the estimation that the controller would be continuously optimizing with roughly one second to optimize the result. In the case of $N_c > 1$, control actions can be implemented one at a time from the previous step optimization in case the next optimization has not converged in time.

4 Discussion

In this paper, we developed a model-based predictive closed-loop optimal control framework that utilizes a data-driven machine learning model of a system in optimizing vagus nerve stimulation (VNS) parameters to control the cardiovascular physiological outputs such as heart rate (HR) and mean arterial blood pressure (MAP). Using a synthetic dataset generated from a previously published *in silico* physiological model of the rat cardiac system [Yao and Kothare, 2020], we developed a long short-term memory (LSTM) based neural network model to predict HR and MAP of the *in silico* physiological model in response to VNS. The predictions of HR and MAP from our LSTM model showed quantitative consistency with the *in silico* physiological model in response to VNS with reasonable accuracy. We then used this model in designing various model-based optimal control strategies and demonstrated the efficacy of our closed-loop optimal control designs in controlling HR and MAP of the *in silico* physiological model of the rat cardiac system (the physiological system to be controlled in our case) in simulation. Finally, we showed in the simulation how our control designs could address intra-patient variability in closed-loop VNS designs and control HR and MAP in pathological conditions (e.g., overactive sympathetic cardiac system) where the model mismatch between the LSTM model and the system to be controlled is significant. Overall, our results highlight the advantage of using a closed-loop model predictive optimal control framework in optimizing VNS parameters to control multiple cardiac biomarkers.

Throughout our closed-loop designs and simulations, we chose appropriate design parameters such as the control and prediction horizons and various weighting matrices to obtain the best performances. Regarding the selection of these control design parameters, there are some guiding principles to select these parameters

Table 1: Comparison of computational time and closed-loop performance of the designed closed-loop VNS strategies. Simulations were performed on an INTEL(R) CORE I7-9700 CPU 3.00 GHZ with 16.0 GB of RAM Desktop.

Nc (cycles)	Np (cycles)	Variables	Est. Time/cycle (s)	SS MAE
Sparsity-promoted VNS design (Section 3.2.1)				
1	10	6	0.21	2.41
5	10	30	0.89	0.96
10	10	60	1.15	0.99
1	20	6	0.57	5.46
5	20	30	1.89	0.80
10	20	60	2.35	0.75
20	20	120	2.75	1.10
Minimum energy-based VNS design (Section 3.2.2)				
1	10	6	0.24	2.52
5	10	30	0.81	0.87
10	10	60	0.97	0.84
5	20	30	1.72	0.92
10	20	60	2.2	0.86
20	20	120	2.33	0.79
Minimum overshoot-based VNS design (Section 3.2.3)				
1	10	6	0.33	1.43
5	10	30	1.00	1.40
10	10	60	1.55	2.12
10	20	60	3.81	2.04
20	20	120	6.34	2.32
Offset free VNS design (Section 3.3)				
10	20	68	1.03	0.34
Overactive Sympathetic Case (Section 3.4)				
10	20	68	1.06	0.70

based on the observed closed-loop response of the designed controller. For example, in the case of the sparsity-promoted design presented in Section 3.2.1, a large value of λ instructs the controller to emphasize suppressing the inputs more than reaching the target value. A significant offset would evidence the occurrence of such a scenario in the controlled system's outputs. Similarly, in the minimum energy-based design presented in Section 3.2.2, a large value of λ can lead to a similar situation where the controller is more focused on minimizing the energy than driving the outputs to the desired target.

As for this study, we have not found any significant difference in the closed-loop performance for different choices of prediction and control horizons, as shown in Section 3.5. However, we had noted that the computational time increased by approximately 3 times when we changed the prediction and control horizons from 10 and 5, respectively, to 20 and 20 (see Section 3.5). In general, the choice of control horizon N_c and prediction horizon N_p can have a significant impact on the closed-loop stability, performance, and computational time. Typically, a large prediction horizon can help in stabilizing the closed-loop system, while a large control horizon (less than or equal to the length of the prediction horizon) can provide less aggressive closed-loop performance.

However, selecting large values for these parameters could also lead to large computational time to simulate the closed-loop response of the system, which may prevent the real-time implementation of the closed-loop control design in experiments.

The control design weighting parameters such as Q and R in the cost function are the user-defined parameters. Throughout this work, we have chosen these to be identity matrices of appropriate dimensions. The matrix Q allows to weigh specific output of interest more than other outputs to be controlled, and thus the optimization problem focuses more on minimizing this specific output than other outputs. Similarly, the matrix R allows to weigh specific input of interest more than other inputs to be optimized, and thus the optimization problem emphasizes minimizing this specific input compared to other inputs.

Although we showed the efficacy of the designed closed-loop VNS strategies in controlling the cardiac physiological biomarkers in an *in silico* physiological model of the rat cardiac system, the described methodology appears general enough to be translated in animal experiments. Remarkably, the inclusion of a deep learning-based model allows for the flexibility of developing a model of the physiological system only using the experimental data. Notably, the deep learning model described here was model-agnostic, as it was never given information regarding the structure of the underlying *in silico* physiological model. While this approach has been demonstrated for the specific application of cardiovascular control, it can more generally be used to develop reduced-order models of large-scale and computationally inefficient models, thus enabling the plausibility of real-time control of cardiac system in animal experiments. Other potential applications where this approach would excel can be found in areas where model development is challenging, but a control application would still prove useful. The challenge of model building may result from a lack of adequate system knowledge or the level of detail required for a particular experimental system. In either case, provided with the input-output data, this approach can be used to develop a controller. We note that this approach is particularly advantageous in the context of controlling biological systems.

The closed-loop control framework developed in this work provides flexibility in designing a control strategy for a specific situation. The optimization problem formulation can incorporate any relevant cost function and physiological constraints depending on the specific application. Additionally, as shown in Sections 3.3 and 3.4, the control problem formulation allows using the LSTM model trained on the data obtained from one subject to be applied in designing control strategy for a different subject, thus allowing a possibility to develop a model from the healthy subjects and use the model in designing control strategy for treating diseased cases.

Due to a limitation of the *in silico* physiological model used to generate synthetic data for training the LSTM model, we have omitted the pulse amplitude as an optimized VNS parameter. We anticipate that the inclusion of this parameter would not cause any fundamental problem to our approach. Additionally, most of the published studies on optimizing VNS for cardiac system did not consider the influence of disease pathology on controller performance [Greenwald et al., 2016, Ugalde et al., 2015, Tosato et al., 2006, Zhang et al., 2002, Romero-Ugalde et al., 2017],[Yao and Kothare, 2020]). While we have considered this by modifying the physiological dynamics (see Sections 3.3 and 3.4), it is worth noting that the offset-free controller formulation is limited to overcoming the steady-state offsets in the closed-loop performance. Thus, should the controller's objective change to a tracking problem (as opposed to a constant set point), then the controller's performance would not necessarily achieve nominal offset in the presence of model mismatch. While this limitation may appear exceptionally prohibitive, depending on the application of the VNS device, such a limitation may not restrict all applications.

Should these controller designs be deployed in the experimental/clinical setup, there is a potential that the controller could switch its objective function to meet the demands of the physiological system. In a therapeutic context, stimulating all locations at all times could lead to a loss in the efficacy of specific treatment. Conversely, there may be times when power consumption holds a higher priority than enforcing a sparse solution. For the therapeutic application, the closed-loop control formulation in Section 3.2.3 could drive the physiological system to reach the desired targets with more smooth transitions, which may be more gentle for the patient or subject of the therapy. These different contexts are why we have investigated multiple controller designs in this paper. Importantly, all the control designs have been shown to reach the target set point with a nominal to no steady-state offset. Their similar performances suggest a potential of employing them in the experimental system while accounting 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.

5 Methods

5.1 *In Silico* Rat Cardiac Model

We used a previously published *in silico* physiological model of the rat cardiac system that includes the effect of two VNS parameters (pulse width in ms and pulse frequency in Hz) on the two physiological variables (heart rate (HR) in bpm and mean arterial blood pressure (MAP) in mmHg) [Yao and Kothare, 2020]. Here, we provide a brief description of this *in silico* physiological model and refer the reader to this study for a full model description. The model was composed of three different parts: the cardiovascular system, the baroreflex system, and the vagus nerve stimulation (VNS) device. Regarding the VNS device, there were three different locations to apply VNS, following the experimental setup of [Plachta et al., 2014], and thus there were six total VNS parameters (three locations with two parameters each). To determine the effects of VNS on the physiological outputs, each location contained a concentration of vagal, baroreceptive, and sympathetic neuronal fibers. The application of VNS led to different fiber recruitment levels depending on the parameters chosen. From the fiber recruitment, the effect on the physiological variables were determined through the interactions between the simulated central nervous system, and the cardiac system.

5.1.1 Cardiovascular Model

Design of the cardiovascular system were 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 adjusted by the body volume ratio of rats to humans, resulting in values similar to those measured in experimental rats [Pacher et al., 2004]. The pressure volume relationship was described as

$$P_i = E_i(V_i - V_{i,d}), \quad (10)$$

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, $E(t)$, was described, by following the formulation of [Stergiopoulos et al., 1996], as:

$$E(t) = E_{max} \left(a \frac{\left(\frac{t_n}{\alpha_1 T}\right)^{n_1}}{1 + \left(\frac{t_n}{\alpha_1 T}\right)^{n_1}} \frac{1}{\left(\frac{t_n}{\alpha_2 T}\right)^{n_2}} \right) + E_{min}, \quad (11)$$

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 chambers of the cardiac system was modeled as follows:

$$Q = \frac{P_{in} - P_{out}}{C_i R_{sys}}. \quad (12)$$

Here, P_{in} and P_{out} represent the pressure difference causing blood flow, C_i represents a compliance constant, and R_{sys} represents the cardiac system's resistance to blood flow. Following a mass balance, the change in volume was given by

$$\frac{dV}{dt} = Q_{in} - Q_{out}, \quad (13)$$

where Q_{in} represents the flow into the compliance chamber, and Q_{out} represents the flow out of the compliance chamber. Finally, the inertial flow out of the left ventricle was described as

$$\Delta P = L \frac{dQ}{dt}, \quad (14)$$

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

5.1.2 Baroreflex Model

The baroreflex system model was derived from [Lau and Figueroa, 2015], and was composed of several parts: the central nervous system, the baroreceptor model, and a modulation of efferent responses. Following the activation of the sympathetic drive, the modulation of efferent responses was described by the left ventricle systolic elastance (E_{max}) and the cardiac resistance to blood flow (R_{sys}) exhibiting a positive response, and conversely the heart period (T) exhibiting a negative response. A previously developed input-output relationship of the interactions between the central nervous system and baroreceptive fibers was used to determine the effect of the baroreflex system on cardiovascular functioning [Ursino, 1998]. The parameters for the baroreceptive model were taken from [Mahdi et al., 2013], and the parameters of the baroreflex system model were taken from [Ursino, 1998]. To keep generality, each location was assumed to be concentrated with 100% of a specific neuronal fiber type (i.e., location 1 was assumed to be concentrated with only baroreceptive fibers), which led to a qualitative match with the experimental data presented in [Plachta et al., 2014]. The sympathetic efferent pathway was described using the following equation:

$$\frac{d\theta_{es}}{dt} = -\frac{1}{\tau_{\theta_{es}}}(\theta_{es} - \theta_0) + G_{\theta_{es}} \ln(\max\{f_{es}(t - D_{\theta_{es}}) - f_{es,min}, 1\}). \quad (15)$$

Here, θ_{es} denotes each of the efferent path variables: the heart period T , the systolic elastance of the left ventricle E_{max} , and the inertial cardiac flow resistance, R_{sys} . The baseline value in the absence of external input is given by θ_0 , $\tau_{\theta_{es}}$ denotes the time constant, $f_{es}(t)$ represents the firing rate of the sympathetic efferents, $G_{\theta_{es}}$ 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. Following the activation of vagal fibers, the first-order dynamics were used to capture the corresponding change in the heart period as follows:

$$\frac{dT_{ev}}{dt} = -\frac{1}{\tau_{T_{ev}}}(T_{ev} - T_0) + G_{T_{ev}} f_{ev}(t - D_{T_{ev}}). \quad (16)$$

Here, T_{ev} denotes the change in the heart period due to activation of vagal fibers, T_0 denotes the resting heart period, $\tau_{T_{ev}}$ is the time constant, $G_{T_{ev}}$ is the gain of the heart period, and $D_{T_{ev}}$ is the delay of the vagal pathway. The effects on the heart period (T) from sympathetic activation were assumed to be independent of the effects from the vagal activation, leading to the following calculation for the heart period from the total effects of stimulation:

$$T = T_{ev} + T_{es} - T_0. \quad (17)$$

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

5.1.3 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 the baroreceptive fibers, efferent sympathetic fibers, and vagal fibers. The fiber recruitment due to pulse width was given by

$$F(P_w^i) = \frac{P_w^i/k_w}{\sqrt{1 + (P_w^i/k_w)^2}}, \quad (18)$$

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 was given by

$$\Delta R(P_f^i) = \frac{P_f^i/k_f}{\sqrt{1 + (P_f^i/k_f)^2}}, \quad (19)$$

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. Since the change in fiber firing rates leads to the change in physiological variables, 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 was described by

$$\Delta f_j = \frac{G_j}{n} \sum_{i=1}^n \delta_i C_{i,j} F(P_w^i) \Delta R(P_f^i). \quad (20)$$

Here, $i = 1, 2, \dots, n$ 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.

Thus, Equations 10 - 20, together describe the complete physiological dynamics of the *in silico* rat cardiac system with the influence of VNS.

5.2 Intra-Patient Variation Cardiac Model

In [Plachta et al., 2014], the authors showed that there is a significant variation across rats in response to the vagus nerve stimulation (VNS) parameter selection. Particularly in their study, they showed that the qualitative response to the specific stimulation parameters on the heart rate was consistent across the animals, but there were significant statistical variations in the quantitative responses (i.e., A 20% decrease in heart rate was observed in one rat, while a 10% decrease in the heart rate was observed in another rat for the same stimulation parameter selection). To demonstrate how our closed-loop VNS design can account for this subject-to-subject variability in response to VNS, we constructed an example by modifying the *in silico* physiological model of the rat cardiac system. Specifically, we modified the concentration of fiber recruitment at each VNS location (see C in Equation (20)) from an identity matrix to:

$$C = \begin{bmatrix} 1.1 & 0.1 & 0.2 \\ 0.1 & 0.8 & 0.0 \\ 0.2 & 0.2 & 0.6 \end{bmatrix}$$

Here, the row in the matrix C represents the specific location, and the column represents the fiber type recruited at that location. For example, row 1 in the matrix C indicates that the VNS location 1 activates 110% of baroreceptive fiber, 10% of sympathetic fibers, and 20% of vagal fibers compared to the baseline responses. Note that these values were set to an identity matrix in the *in silico* physiological model, which says that each location has a 100% concentration of an individual fiber type (baroreceptive, sympathetic, and vagal). Since this modification in the *in silico* physiological model only affects the actuation side of the model (i.e., the VNS effect), the model mismatch shown in Figure 6 is due to the stimulation parameters exhibiting differential effects on the physiological variables.

5.3 Overactive Sympathetic Cardiac Model

Often, throughout the progression of specific cardiovascular disease pathology, the sympathetic system becomes overactive at the resting state [Malpas, 2010]. This hyperactivity is typically observed by a higher resting heart rate, a higher resting blood pressure, or both, depending on the specific disease considered. Since VNS is a therapy that targets disease pathology, we questioned if a closed-loop VNS control design could perform through a potential mismatch in the hypothetical case of an overactive sympathetic system. To investigate this question, we constructed a hypothetical example of an overactive sympathetic case where we modified the parameters of the *in silico* physiological model of rat cardiac system to increase the influence of the sympathetic system and decrease the influence from the parasympathetic system, consistent with disease pathology that exhibits a reduced parasympathetic tone. We summarize the specific changes we made in the physiological model parameters, described in Section 5.1, in Table 2.

The behavior of the *in silico* physiological model has changed considerably with the new parameter selection, as demonstrated by Figure 9. As expected, the physiological cardiac model with overactive sympathetic system exhibited a higher heart rate and higher blood pressure. In addition to the changes in the model parameters shown in Table 2, we also changed the concentration of fiber recruitment at each VNS location (see C in Equation (20)) from an identity matrix to:

Table 2: Description of the modified parameters in the *in silico* physiological model of rat cardiac system [Yao and Kothare, 2020] to construct an overactive sympathetic case.

Parameters	Description	Initial Value	Diseased Value	Equations
$R1$	systemic resistance	0.01	0.007	15
$R2$	mitral valve resistance	0.0001	0.0002	12
$R3$	aortic valve resistance	0.008	0.006	12
$C2$	venous compliance	20	25	12
$C3$	systemic compliance	1.8	1.4	12
E_{min}	end-diastolic elastance	0.02	0.01	11
E_{max}	end-systolic elastance	1.2	1.1	11
T_0	baseline HR	60/450	60/480	17
G_R	gain of systemic resistance	0.06	0.07	15
G_{Ts}	gain of heart period from sympathetic fibers	-0.01	-0.015	15
G_{TeV}	gain of heart period from vagal fibers	0.015	0.011	16
G_2	gain of sympathetic fibers	30	33	20
G_3	gain of vagal fibers	30	27	20

$$C = \begin{bmatrix} 1.3 & 0.1 & 0.1 \\ 0.1 & 1.4 & 0.1 \\ 0.1 & 0.1 & 1.3 \end{bmatrix}$$

By changing the actuation matrix of the model and the model parameters, we considered both the intra-patient model mismatch and the disease model mismatch. Should a controller be deployed in the clinical setting, the controller would likely be required to compensate for both sources of mismatch simultaneously.

5.4 LSTM Model Development

We recently developed a purely data-driven long short-term memory (LSTM) based neural network modeling approach to map the VNS parameters on the cardiac physiology. LSTM is the state-of-art in data-driven modeling of dynamical systems and sequence based tasks [Hochreiter and Schmidhuber, 1997] demonstrated by its use for forecasting traffic patterns [Zhao et al., 2017], natural language processing [Radford et al., 2019], handwriting recognition [Graves et al., 2008], and speech recognition [Sak et al., 2014]. The LSTM uses a combination of a hidden state, a cell state, and the incoming data to predict the sequential evolution of the system. The inclusion of the cell state allows for learning of long-scale temporal dynamics, and is managed by a gating process where incoming data with the previous hidden state is used to forget a portion of the previous information from the cell state. Additionally, an input gate allows the relevant information from the incoming data and previous hidden state to be stored in the cell state. The output from a LSTM is based on the updated cell state, previous hidden state, and the incoming data.

In this work, we used a LSTM based modeling approach to predict the effect of VNS parameters on the dynamics of the heart rate (HR) and the mean arterial blood pressure (MAP). We generated a synthetic input-output dataset by simulating the *in silico* physiological rat cardiac model (Equations 10 - 20) for 15,198 individual trials using VNS parameters from a randomly sampled uniform distribution with a range between 0 ms and 0.5 ms for the pulse width and 0 Hz and 50 Hz for the pulse frequency. Each location was also turned on/off by sampling from a uniform random distribution between 0 and 1. The specific location is turned off if the selected value is 0 and turned on if the selected value is 1. An individual trial consists of 100 consecutive cardiac cycles with constant VNS parameters. These open loop trials were then used to train the neural network using a 40% - 20% - 40% split for the training, validation, and testing sets respectively. We note that this synthetically generated dataset represents a similar task to real data collection methods, and could be extended to an *in vivo* experimental system where different experimental trials are used to generate a sweep of the VNS parameters. In the case of data collection from the *in vivo* system, random sampling of VNS parameters and locations is not required, and can be replaced with a systematic investigation of the parameter space, such that the full range of parameters have been swept. The LSTM was trained for 250 epochs using the adaptive moment estimation




Figure 11: Performance comparison of different trained neural network architectures based on the mean absolute error for the same test set. (A) shows a comparison of baseline performance to a vanilla RNN, gated recurrent unit (GRU) and LSTM. (B) Shows the a comparison of the effect of the number of LSTM layers on the model's predictive performance. (C) Shows a comparison of the effect of the number of inputs on the model's predictive performance.

(Adam) optimizer [Kingma and Ba, 2014] with a mean squared error loss function. Prior to feeding the data to the LSTM for training, the open-loop dataset was normalized using the following equation:

$$\hat{x} = \frac{x - \mu}{x_{max} - x_{min}} \quad (21)$$

where μ is the mean of the training set, x_{min} is the minimum of the training set, and x_{max} is the maximum value of the training set for each variable. Statistics from the training set are used to avoid providing any information about the validation or test datasets to the trained network, as this is a standard practice in the machine learning community. As a note, since the LSTM was trained in the normalized space, all control applications required incoming data for predictions to be normalized, and the controller actions to be un-normalized before applying to the physiological model during closed-loop operation.

To compare different trained architectures, the test set of data was used in such a way that the trained network was required to recursively predict 99 consecutive cardiac cycles (1 trial) for both HR and MAP. A baseline model that predicted no change in the initial value for all 99 cycles was used to provide context to the normalized mean absolute error (MAE) and emphasize a trained network's ability to learn the dynamics. These results are summarized in Figure 11, with different recurrent networks (GRU is a gated recurrent unit) shown in Figure 11A, the influence of number of LSTM layers on predictive performance is shown in Figure 11B, and the influence of LSTM input neuron size shown in Figure 11C.

To conclude the study of several different neural network architectures, the best trained model resulted in a LSTM with a single layer, hidden dimensionality of 10, and a hyperbolic tangent (tanh) activation function. The output from this LSTM was fed to a dense layer with dimensionality of two, and tanh activation function. The performance of this LSTM is highlighted by a normalized mean absolute error of 0.0072 on the test set. A prediction from the LSTM for time varying VNS inputs is shown in Figure 1. The LSTM model reasonably predicts the output from the full physiological model, and is capable of mapping VNS parameters to the effect on the cardiac variables of heart rate and mean arterial blood pressure.

5.5 Simulation Environment and Relevant Software Packages

Throughout this work, we used MATLAB (version R2019b) to simulate the physiological model of rat cardiac system. The LSTM model was developed in Python (version 3.7) using TensorFlow2. The controller was synthesized in Python (version 3.7) and the formulated optimization problems were solved using the sequential least squares programming (SLSQP) algorithm in Scipy (version 1.6). The closed-loop simulations were performed through the interaction of Python with MATLAB. For the healthy cases, the optimization solver's tolerance was set to $10e - 5$ and the maximum number of iterations was set to 50. While, for the intra-patient and diseased cases, the solver's tolerance was set to $10e - 4$ and the maximum number of iterations was set to 500 (although the algorithm converged in less than 50 iterations in our simulations).

Conflict of Interest Statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Author Contributions

AB formulated the problem, developed the LTSM model, designed the control strategies, ran the simulations, discussed the results, and wrote the manuscript. YY provided the MATLAB code of the physiological model, discussed the results, and provided feedback on the manuscript. MK and BM discussed the problem formulation and results, and provided feedback on the manuscript. GK formulated the problem, discussed the results, and wrote the manuscript.

Funding

This work was supported by National Institutes of Health award 1OT2OD030535-01.

Data Availability Statement

The code used in this study can be found at Github <https://github.com/Bradrew/Data-Driven-Control-of-VNS-Cardiac>.

References

- [Asad and Stavrakis, 2019] Asad, Z. U. and Stavrakis, S. (2019). Vagus nerve stimulation for the treatment of heart failure. *Bioelectronics in Medicine*, 2(1):43–54.
- [Benjamin et al., 2018] 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(12):e67–e492.
- [Capilupi et al., 2020] Capilupi, M. J., Kerath, S. M., and Becker, L. B. (2020). Vagus nerve stimulation and the cardiovascular system. *Cold Spring Harbor perspectives in medicine*, 10(2):a034173.
- [Carthy, 2014] Carthy, E. R. (2014). Autonomic dysfunction in essential hypertension: a systematic review. *Annals of medicine and surgery*, 3(1):2–7.
- [Djabella et al., 2005] 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.
- [Ferreira et al., 2005] 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 et al., 2016] Gold, M. R., Van Veldhuisen, D. J., Hauptman, P. J., Borggrefe, 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(2):149–158.
- [Goodfellow et al., 2016] Goodfellow, I., Bengio, Y., and Courville, A. (2016). Machine learning basics. *Deep learning*, 1:98–164.
- [Graves et al., 2008] 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(5):855–868.

[Greenwald et al., 2016] 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 bioactive array and charge balanced stimulator. *IEEE transactions on biomedical circuits and systems*, 10(5):990–1002.

[Heidenreich et al., 2013] 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(3):606–619.

[Hochreiter and Schmidhuber, 1997] Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.

[Howland, 2014] Howland, R. H. (2014). Vagus nerve stimulation. *Current behavioral neuroscience reports*, 1(2):64–73.

[Kingma and Ba, 2014] Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

[Kishi, 2012] Kishi, T. (2012). Heart failure as an autonomic nervous system dysfunction. *Journal of cardiology*, 59(2):117–122.

[Kong et al., 2017] Kong, W., Dong, Z. Y., Jia, Y., Hill, D. J., Xu, Y., and Zhang, Y. (2017). Short-term residential load forecasting based on lstm recurrent neural network. *IEEE Transactions on Smart Grid*, 10(1):841–851.

[Kwakernaak and Sivan, 1972] Kwakernaak, H. and Sivan, R. (1972). *Linear Optimal Control Systems*. John Wiley & Sons, Inc., USA.

[Lau and Figueroa, 2015] Lau, K. D. and Figueroa, C. A. (2015). Simulation of short-term pressure regulation during the tilt test in a coupled 3d-0d closed-loop model of the circulation. *Biomechanics and modeling in mechanobiology*, 14(4):915–929.

[Mahdi et al., 2013] 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(12):e1003384.

[Malpas, 2010] Malpas, S. C. (2010). Sympathetic nervous system overactivity and its role in the development of cardiovascular disease. *Physiological reviews*.

[Mangoni et al., 2006] Mangoni, M. E., Traboulsi, 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(11):1422–1430.

[Melchior et al., 1992] 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(6):H1920–H1933.

[Morari and Maeder, 2012] Morari, M. and Maeder, U. (2012). Nonlinear offset-free model predictive control. *Automatica*, 48(9):2059–2067.

[Ovbiagele et al., 2013] 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(8):2361–2375.

[Pacher et al., 2004] 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*.

[Palatini and Julius, 2009] Palatini, P. and Julius, S. (2009). The role of cardiac autonomic function in hypertension and cardiovascular disease. *Current hypertension reports*, 11(3):199–205.

[Plachta et al., 2014] Plachta, D. T., Gierthmuhlen, 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(3):036011.

[Plaster and Kumar, 2019] Plaster, B. and Kumar, G. (2019). Data-driven predictive modeling of neuronal dynamics using long short-term memory. *Algorithms*, 12(10):203.

[Premchand et al., 2014] 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(11):808–816.

[Radford et al., 2019] Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. (2019). Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.

[Romero-Ugalde et al., 2017] 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(10):e0186068.

[Sahoo et al., 2019] Sahoo, B. B., Jha, R., Singh, A., and Kumar, D. (2019). Long short-term memory (lstm) recurrent neural network for low-flow hydrological time series forecasting. *Acta Geophysica*, 67(5):1471–1481.

[Sak et al., 2014] Sak, H., Senior, A. W., and Beaufays, F. (2014). Long short-term memory recurrent neural network architectures for large scale acoustic modeling.

[Savarese and Lund, 2017] Savarese, G. and Lund, L. H. (2017). Global public health burden of heart failure. *Cardiac failure review*, 3(1):7.

[Stergiopoulos et al., 1996] 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(6):H2050–H2059.

[Suga et al., 1973] 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(3):314–322.

[Tosato et al., 2006] 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(3):161–169.

[Trespalacios and Grossmann, 2014] Trespalacios, F. and Grossmann, I. E. (2014). Review of mixed-integer nonlinear and generalized disjunctive programming methods. *Chemie Ingenieur Technik*, 86(7):991–1012.

[Ugalde et al., 2015] 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(7):1551–1558.

[Ursino, 1998] Ursino, M. (1998). Interaction between carotid baroregulation and the pulsating heart: a mathematical model. *American Journal of Physiology-Heart and Circulatory Physiology*, 275(5):H1733–H1747.

[Vidaurre et al., 2013] Vidaurre, D., Bielza, C., and Larrañaga, P. (2013). A survey of l1 regression. *International Statistical Review*, 81(3):361–387.

[Wu et al., 2019] 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(11):e16734.

[Yao and Kothare, 2020] 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.

[Zannad et al., 2015] 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(7):425–433.

[Zhang et al., 2002] 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(3):H1102–H1110.

[Zhao et al., 2017] 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(2):68–75.