2021 43rd Annual International Conference of the
IEEE Engineering in Medicine & Biology Society (EMBC)
Oct 31 - Nov 4, 2021. Virtual Conference

Mapping Vagus Nerve Stimulation Parameters to Cardiac Physiology using Long Short-term Memory Network

Andrew Branen[1], Yuyu Yao[2], Mayuresh V. Kothare[2], Babak Mahmoudi[3], and Gautam Kumar[4]

Abstract— **Vagus nerve stimulation (VNS) is an emerging therapeutic strategy for pathological conditions in a variety of diseases; however, several challenges arise for applying this stimulation paradigm in automated closed-loop control. In this work, we propose a data driven approach for predicting the impact of VNS on physiological variables. We apply this approach on a synthetic dataset created with a physiological model of a rat heart. Through training several neural network models, we found that a long short term memory (LSTM) architecture gave the best performance on a test set. Further, we found the neural network model was capable of mapping a set of VNS parameters to the correct response in the heart rate and the mean arterial blood pressure. In closed-loop control of biological systems, a model of the physiological system is often required and we demonstrate using a data driven approach to meet this requirement in the cardiac system.**

I. INTRODUCTION

Vagus nerve stimulation (VNS) is an emerging therapeutic technique that seeks to offset various disease pathologies by restoring stability through a surgically implanted electrode [1]. This technique has been applied to depression and epilepsy, and is under investigation for use in the gastrointestinal tract and cardiac systems. Electric pulses sent to the vagus nerve from the electrode are defined by several parameters such as pulse width, frequency of pulses, and pulse amplitude. A major challenge associated with applying VNS therapies consists of selecting optimal stimulation parameters to achieve the desired physiological behavior. Currently, VNS is manually titrated for each individual, which has led to different results in clincial studies [2], [3], and [4]. Moving from an ad hoc approach to an automated closed-loop control scheme could result in higher treatment efficacy with VNS.

Shifting to automated closed-loop control has numerous obstacles in the cardiac system as predicting phyisiological behavior with the influence of VNS requires a complex model. Previously, this model requirement has been addressed with a state transition model [5], or a Laplacian based model [6]. A drawback associated with these approaches stems from requiring very specific model parameter selection and considering the selection of a single stimulation parameter. These requirements move further away from autonomous closed-loop control of the cardiac system. Conversely, using fully detailed mechanistic models suffer the drawback of computational expense which limits their applicability in real time closed-loop implementation. Further, these fully detailed mechanistic models may not be available to incorporate the nuances of an experimental setup in the real system. For example, if a model considers an electrode at a specific location and the true experimental setup implants an electrode at a different location, the model may not be extendable to that specific alteration.

Data driven models, such as neural networks, provide a promising solution to address challenges discussed above and have been used in a similar context for a Hodgkin Huxeley pyramidal neuron model [7]. Such applications highlight a major advantage of neural networks as they do not require underlying assumptions about the system dynamics or require a specific distribution of data. Recurrent neural networks (RNNs) have been specifically developed for time-series data. Finally, neural networks can be computationally inexpensive to evaluate when compared to full scale models. Taken together, these features suggest a promising application of neural networks to model the cardiac system.

In this work, we propose an approach for modeling the response of physiological variables to VNS stimulation in the cardiac system by using an existing cardiac model [8] to generate synthetic data in an open loop format. We then train a variety of RNN architectures on the collected open loop data. Our results show that a long short-term memory (LSTM) architecture provided the best predictive performance. Further, we find that the computational cost of LSTM in predicting the mean arterial blood pressure and heart rate for 100 consecutive cardiac cycles is approximately 10 folds lower than the original physiological model.

II. MODELING APPROACH

The model used for synthetic dataset creation is a cardiac model of a rat heart [8]. Briefly described, it consists of a system of delayed differential equations with ten internal states and a nonlinear combination of the internal states give rise to the phyisological variables of interest: the mean arterial blood pressure and the heart rate. VNS was incorporated by including three different locations, where each location

This work was supported through National Institute of Health award 1OT2OD030535-01.

[1]Department of Chemical and Biological Engineering, University of Idaho, 875 Perimeter Drive, Moscow, ID, 83844, USA bran7699@vandals.uidaho.edu

[2]Department of Chemical and Biomolecular Engineering, Lehigh University, 111 Research Drive, Bethlehem, PA 18015, USA yuy516@lehigh.edu, mayuresh.kothare@lehigh.edu

[3]Department of Biomedical Informatics, School of Medicine, Emory University, 313 Ferst Drive, Atlanta, GA 30332, USA b.mahmoudi@emory.edu

[4]Chemical & Materials Engineering Department, San José State University, One Washington Square, San José, CA 95192, USA gautam.kumar@sjsu.edu

had two parameters, a pulse width and a pulse frequency which together characterize the stimulation paradigm. Each VNS location corresponded to the full recruitment of either baroreceptive, sympathetic, or vagal fibers.

Using the cardiac model from [8], open-loop simulations were run for 15,198 different trials with the VNS parameters chosen through sampling a uniform distribution. The bounds for the pulse width and pulse frequency were 0-0.5 ms, and 0-50 Hz, respectively. The three different locations were randomly activated using a uniform distribution between 0 and 1 where values above 0.5 indicated activating that location. A single open-loop trial was obtained after selecting the VNS parameters and simulating the model for 100 consecutive cardiac cycles. During all 100 cardiac cycles, the physiological outputs of the heart rate and mean arterial blood pressure were recorded from the model.

Next, the trials were divided into a training, validation, and test set following a 40%, 20%, and 40% split, respectively. For the training and validation sets, the 100 cardiac cycles were sliced into 50 pairs of alternating points with the inputs consisting of the VNS parameters, heart rate and mean arterial pressure. The labels consisted of the heart rate and mean arterial pressure for the next cardiac cycle. With this problem formulation, all neural networks were trained to make a single step prediction, where the cardiac cycle serves as the discretionalized step. To assess neural network performance using the test set, the 100 cardiac cycles were sliced into one input (the initial heart rate and mean arterial pressure values with VNS parameters) and 99 labels for the remainder of the trial. The neural network was evaluated on its ability to recursively predict the 99 cardiac cycles, given the initial point. Testing the network in this way allowed us to assess the ability of the neural network to predict recursively, similar to its application in a control context. Before the data was fed to the network for training, all data was normalized according to equation 1.

$$\hat{x} = \frac{x - \mu}{x_{max} - x_{min}} \tag{1}$$

Here, $\hat{x}$ denotes the normalized data point, μ denotes the mean of the training set, x_{max} and x_{min} denote the maximum and minimum values of the training set, respectively. Statistics of the training set were used, so the model did not get any information about the validation or testing datasets. The distributions following normalization for the variables are shown by Fig. 1. All of the VNS parameters have the same distribution shape, which is expected as they were all drawn from the same uniform distributions.

A variety of neural network architectures were trained, including a vanilla RNN, a gated recurrent unit (GRU), and a LSTM. The number of layers and inputs were also varied to explore the network architecture on the model predictions. For all models, the hyperbolic tangent function was used, and the output from the recurrent layer was fed to a dense layer with two dimensions to predict the two physiological variables. Throughout the training of neural networks, the mean squared error loss function and ADAM optimizer were

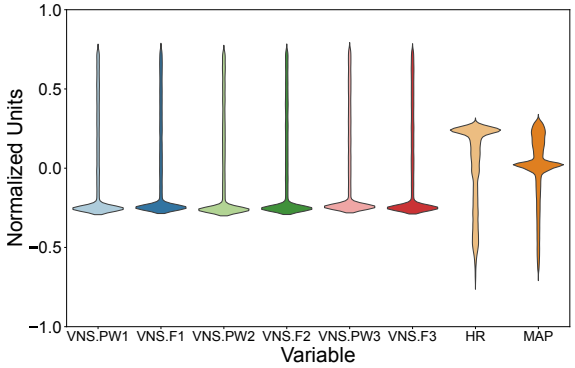

Fig. 1. A violin plot of the normalized data (unitless) for all variables. HR and MAP denote the heart rate and the mean arterial pressure, respectively. VNS PW indicates the stimulation pulse width at the location number indicated in the label, and F indicates the stimulation frequency at the location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

used. To assess a trained network's performance, normalized mean absolute error (MAE) was used in the test set. One network was said to be better than another when the mean absolute error on the test set was lower. As a measurement of performance, all trained models were compared to a baseline case that predicts no change in the initial heart rate and mean arterial blood pressure for the length of the trial, thus emphasizing a network's ability to capture the dynamical behavior.

III. RESULTS

A. Model Comparison

The first goal in searching for an architecture involved a wide search of different RNN structures. The performance of different architectures is shown in Fig. 2A. Consistently, a vanilla RNN fails to capture the dynamical behavior shown by the higher error value. LSTMs and GRUs contain several memory gating structures that enable them to capture long term temporal dependencies. Given the nature of the differential equation system, we would expect the additional memory gates to perform better on the dataset as shown by the lower error values.

Layers were also varied in all architectures to investigate the influence of this design parameter. Consistently, for all networks, increasing the number of layers did not improve the quality of predictions (see Fig. 2B). To the contrary, neural networks with more layers performed worse on the test dataset. The reason for this observation may lie in the problem formulation of a single step prediction. Number of input neurons were also varied for all networks (see Fig. 2C). Since the network was predicting two ouputs, the number of inputs were initially varied in powers of two, holding the number of layers constant (one). From this study, we observe that the number of neurons for a single layer must be greater than eight. When the number of neurons were increased past 32, predictive performance started to decrease.

We found our optimal trained neural network to be a single layer LSTM with a ten neuron input size. The best trained GRU, also with a ten neuron input size, provided similar performance. The best performing network results

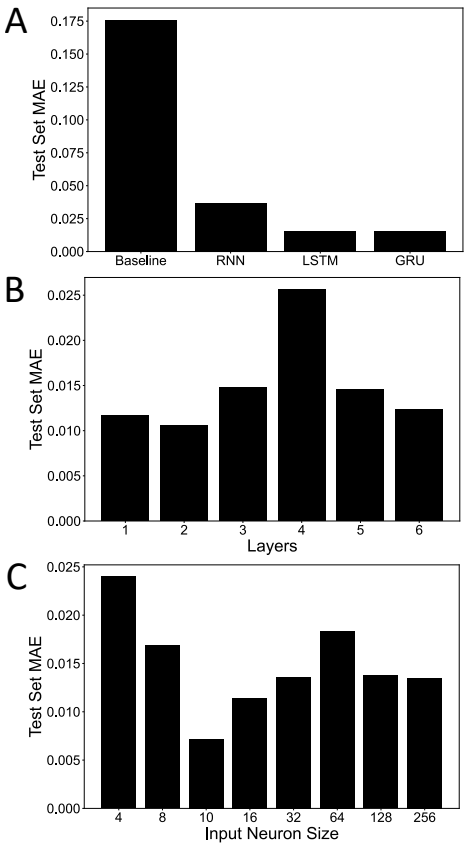

Fig. 2. Comparison of a neural network's mean absolute error (MAE) on the test set including the baseline case (A). The effect of layer size on a neural network's predictive performance on the test set (B). The effect of input neuron size on a neural network's predictive performance on the test set (C).

are highlighted in Table I. When selecting our best network, we favored networks that were smaller with comparable performance.

TABLE I

TOP THREE TRAINED ARCHITECTURES

Model Type	Input Size	Layers	Test Set Mean Absolute Error
LSTM	**10**	**1**	**0.0072**
GRU	10	1	0.0077
GRU	32	1	0.0078

B. Best Model Prediction Performance

Using the LSTM with ten input neuron size, a sample prediction from the test set is shown in Fig. 3 along with the average error for all cardiac cycles. Note that the model output is fed back to make the next prediction, however there is not an accumulation of error. While the first few timestep predictions are slightly inaccurate, the steady state value is still accurately predicted. Thus the LSTM is capable of performing predictions in a recursive fashion. From the error plot, we observe the mean arterial pressure had a higher error (see Fig. 3C). Examining plots of mean arterial pressure, the cause of this discrepency lies in curves that are less smooth than the heart rate, as confirmed by the error curve for heart rate which has a more consistent error.

Since the model was able to reasonably predict individual trials, we varied the VNS parameters over time to obtain a complex nonlinear curve to assess the model's performance in a more challenging scenario. These results are shown in Fig. 4. For the first ten cycles, VNS at location 1 was active with the pulse width set to 0.21 ms and the frequency set to 21 Hz. The following ten cycles switched to activating VNS at location 3 with pulse width set to 0.09 ms and the frequency set to 9 Hz. Then VNS at location 2 was activated with a pulse width of 0.31 ms and a frequency of 30 Hz for ten cycles. All VNS locations were turned off for 19 cycles, after which location 1 was activated with a pulse width of 0.14 ms and a frequency of 14 Hz for another ten cycles. VNS at all locations was turned off for the final 40 cycles. Again, the model demonstrates a reasonable level of accuracy when tracking the change in the physiological variables. These results support the notion that the trained network has captured the underlying dynamics of the full differential equation model and is capable of sucessfully mapping VNS parameters to the physiological output. In the context of closed loop control, the LSTM demonstrates success and accuracy in the task of predicting where the system will go following a selection of the VNS parameters.

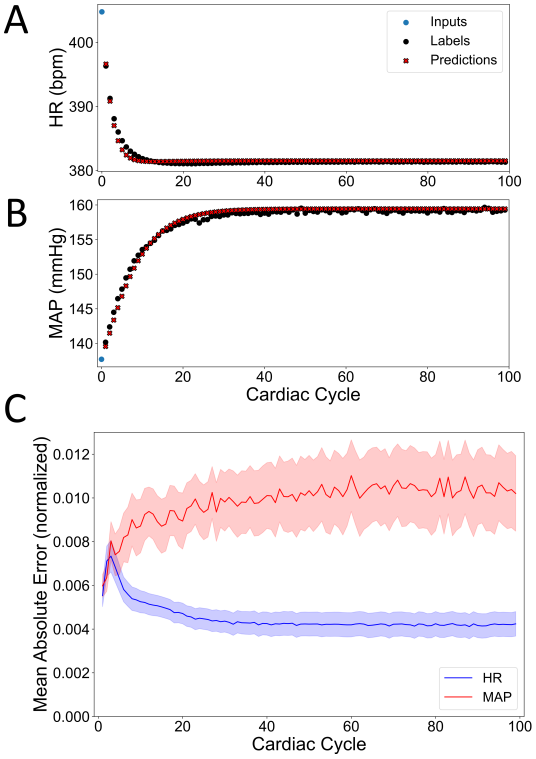

Fig. 3. Example of the trained LSTM model performance on a trial from the test set for predicting the heart rate (A) and the mean arterial pressure (B) for one selection of the VNS parameters. The LSTM is given the input (blue), and is asked to recursively predict the next 99 labels (black) with the LSTM predictions shown in red. Normalized mean absolute error (MAE) over the entire test set for the heart rate and the mean arterial pressure is shown, along with the standard deviation shaded around the curve (C).

C. Computational Speed

Having now showed that the neural network model is capable of predicting complex curves generated by the differential equation model, we assessed the computational speed of the LSTM model and the full physiological model. Additionally,

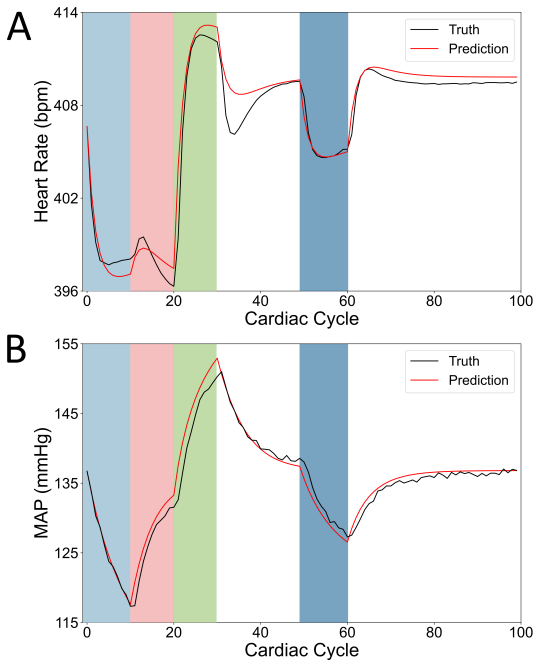

Fig. 4. Comparison between the LSTM model and the full model for the heart rate (A) and the mean arterial blood pressure (MAP) (B) for 99 simulated cardiac cycles. The black line shows the output of the full differential equation model (labeled Truth), while the red line shows the predictions from the LSTM model (labeled Prediction). Cycles 1-10: location 1 with pulse width 0.21 ms and pulse frequency 21 Hz (light blue). Cycles 11-20: location 3 with pulse width 0.09 ms and pulse frequency 9 Hz (light red). Cycles 21-30: location 2 with pulse width 0.31 ms and pulse frequency 30 Hz (light green). Cycles 50-60: location 1 with pulse width 0.14 ms and pulse frequency 14 Hz (dark blue). All other cycles had no locations active.

TABLE II

COMPARISON OF COMPUTATIONAL SPEED. TESTS WERE DONE USING AN INTEL(R) CORE I7-9700 CPU 3.00 GHZ WITH 16.0 GB OF RAM.

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

we included a vanilla RNN and a GRU of similar sizes for a wider comparison. These results are summarized in Table II. Comparatively, there is a clear decrease in the computational time when using a LSTM or any other RNN. While we would expect these results, such a marked decrease in computational time supports the applicability of a LSTM-based model in designing a model-based optimal closed-loop control strategy for controlling cardiac systems.

IV. CONCLUSIONS

Our goal in this work was to develop a data driven approach capable of reproducing nonlinear dynamics for modeling the heart rate and the mean arterial blood pressure in response to the VNS, and demonstrate this approach on a synthetic dataset. We found that both a GRU and LSTM were capable of accomplishing such a task, with a LSTM exhibiting slightly better performance on the dataset used. Further, we showed that both the GRU and LSTM took significantly less computational time than the full physiological model. Thus, the trained LSTM can serve as a reduced-order model of the full physiological rat model.

Due to a limitation of the physiological model, pulse amplitude was omitted as an optimized VNS parameter, which leads to an extended application of this approach where experimental data that includes all three stimulation parameters is used for training. Such an investigation would provide a monumental step in the direction of applying this approach in the therapeutic context. Of particular interest in this application would be the neural network's ability to account for animal specific variation. If successful, this feature would be desirable in closed-loop control for VNS therapy and could provide another advantage of using a neural network over a mechanistic model as most models do not capture experimental variation. We anticipate that this application would likely require a different neural network architecture (number of layers, input size, etc.), but suspect that a LSTM and GRU would still be perferable to the vanilla RNN.

Another extension to the LSTM model involves implementing a controller to work with the LSTM in a model predictive control framework to find the optimal stimulation parameters to reach a target set-point. This study could elucidate the differences in controller performance with the increased computational speed of a neural network. Both extensions mentioned here seek to demonstrate an increased relevance to the overall goal and address challenges related to autonomous closed-loop control of VNS.

REFERENCES

[1] E. Ben-Menachem, D. Revesz, B. Simon, and S. Silberstein, "Surgically implanted and non-invasive vagus nerve stimulation: a review of efficacy, safety and tolerability," *European journal of neurology*, vol. 22, no. 9, pp. 1260–1268, 2015.

[2] M. R. Gold, D. J. Van Veldhuisen, P. J. Hauptman, M. Borggrefe, S. H. Kubo, R. A. Lieberman, G. Milasinovic, B. J. Berman, S. Djordjevic, S. Neelagaru, *et al.*, "Vagus nerve stimulation for the treatment of heart failure: the inovate-hf trial," *Journal of the American College of Cardiology*, vol. 68, no. 2, pp. 149–158, 2016.

[3] R. K. Premchand, K. Sharma, S. Mittal, R. Monteiro, S. Dixit, I. Libbus, L. A. DiCarlo, J. L. Ardell, T. S. Rector, B. Amurthur, *et al.*, "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*, vol. 20, no. 11, pp. 808–816, 2014.

[4] F. Zannad, G. M. De Ferrari, A. E. Tuinenburg, D. Wright, J. Brugada, C. Butter, H. Klein, C. Stolen, S. Meyer, K. M. Stein, *et al.*, "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*, vol. 36, no. 7, pp. 425–433, 2015.

[5] H. M. Romero-Ugalde, V. Le Rolle, J.-L. Bonnet, C. Henry, A. Bel, P. Mabo, G. Carrault, and A. I. Hernández, "A novel controller based on state-transition models for closed-loop vagus nerve stimulation: Application to heart rate regulation," *PloS one*, vol. 12, no. 10, p. e0186068, 2017.

[6] H. M. R. Ugalde, D. Ojeda, V. Le Rolle, D. Andreu, D. Guiraud, J.-L. Bonnet, C. Henry, N. Karam, A. Hagege, P. Mabo, *et al.*, "Model-based design and experimental validation of control modules for neuromodulation devices," *IEEE Transactions on Biomedical Engineering*, vol. 63, no. 7, pp. 1551–1558, 2015.

[7] B. Plaster and G. Kumar, "Data-driven predictive modeling of neuronal dynamics using long short-term memory," *Algorithms*, vol. 12, no. 10, p. 203, 2019.

[8] Y. Yao and M. V. Kothare, "Model predictive control of selective vagal nerve stimulation for regulating cardiovascular system," in *2020 American Control Conference (ACC)*, pp. 563–568, IEEE, 2020.