Prediction of high frequency resistance in polymer electrolyte membrane fuel cells using Long Short Term Memory based model

Tong Lin
Carnegie Mellon University
tongl1@andrew.cmu.edu

Leiming Hu
Carnegie Mellon University
leimingh@andrew.cmu.edu

Willetta Wisely
Carnegie Mellon University
wwisely@andrew.cmu.edu

Xin Gu
Shanghai Hydrogen Propulsion Technology Co., Ltd.
guxin@shpt.com

Jun Cai
Shanghai Hydrogen Propulsion Technology Co., Ltd.
caijun@shpt.com

Shawn Litster
Carnegie Mellon University
litster@andrew.cmu.edu

Levent Burak Kara
Carnegie Mellon University
lkara@cmu.edu

Abstract

High-frequency resistance (HFR) is a critical quantity strongly related to a fuel cell system's performance. As such, an accurate and timely prediction of HFR is useful for understanding the system's operating status and the corresponding control strategy optimization. It is beneficial to estimate the fuel cell system's HFR from the measurable operating conditions without resorting to costly HFR measurement devices, the latter of which are difficult to implement at the real automotive scale. In this study, we propose a data-driven approach for a real-time prediction of HFR. Specifically, we use a long short-term memory (LSTM) based machine learning model that takes into account both the current and past states of the fuel cell, as characterized through a set of sensors. These sensor signals form the input to the LSTM. The data is experimentally collected from a vehicle lab that operates a 100 kW automotive fuel cell stack running on a automotive-scale test station. Our current results indicate that our prediction model achieves high accuracy HFR predictions and outperforms other frequently used regression models. We also study the effect of the extracted features generated by our LSTM model. Our study finds that even a simple LSTM based model can accurately predict HFR values.

1 Introduction

High-frequency resistance (HFR) is an important indicator in polymer electrolyte membrane fuel cells (PEMFCs). During lab-scale electrochemical impedance spectroscopy (EIS) characterizations, a fuel cell membrane electrode assemblies (MEA) can usually be modelled as an equivalent resistor-capacitor (RC) circuit using membrane ionic resistance (R_m), anode/cathode charge transfer resistance (R_{an} and R_{ca}) and double layer capacitance (C_{an} and C_{ca}) [1; 2; 3; 4; 5]. The high-frequency

resistance (HFR) is the value obtained at high frequencies (usually > 1 kHz) and mainly represents the membrane ionic resistance. HFR is a good indicator of the water content of the membrane due to its first component dependency on water (6; 7; 8; 9; 10). It is commonly used for the state-of-health monitoring of a fuel cell stack and is a critical indicator for dry out (i.e., low water content of the polymer electrolyte) (4; 5) or flooding (i.e., too much water inside fuel cell electrodes) (11; 12) conditions inside the fuel cell stack. HFR is also an important indicator of certain types of degradation in the fuel cell (13; 14; 15). It affects the performance of the fuel cell, especially at high current densities as it is closely related to the ohmic voltage loss ($\Delta V_{ohm} = I \times R$).

Though obtaining HFR in real time for vehicle applications is very beneficial for the vehicle's state-of-health monitoring and control strategy optimization, an on-board measuring device is usually cost-prohibitive (16) to integrate into the existing fuel cell vehicle. Fuel cell systems involve highly nonlinear multiple physics, with a large number of interrelated variables. Especially for a full-scale fuel cell system, building a computational multi-physics model to predict the HFR value in the fuel cell becomes computationally prohibitive and impractical to implement on the on-board control systems for operational vehicles (17). While computationally tractable physics have been proposed (18; 19), these models may not adequately capture the complex phenomena in full-scale fuel cell stacks. In such cases, a data driven HFR prediction model trained using an experimental system can prove effective for a real-time, high-accuracy prediction of the behaviors of interest.

In this work, we present a data-driven approach for real-time HFR prediction in PEMFCs. At the heart of our approach is a type of recurrent neural network (RNN), called long short-term memory networks (LSTMs), that incorporates multiple operational inputs of the fuel cell systems at different time steps into the HFR predictions. The model repeats itself over time, receives the fuel cell system signals as its inputs at each time step, automatically analyzes the relationship of the inputs between the two consecutive time steps, and finally generates a feature vector that encodes the inputs occurred in the time period for the HFR prediction. A distinguishing feature of such model is that it can potentially take account of the sequential effect of the system inputs. For example, controlling the same current levels with different orders can result in different HFR signals and such observation can be found by analyzing the current signals and the resulting HFR value.

In the paper, we demonstrate that

- a data-driven model can be used to achieve high accuracy of online HFR prediction for a industry level fuel cell stacks.

- the model can be computationally light and thus, practical to use.
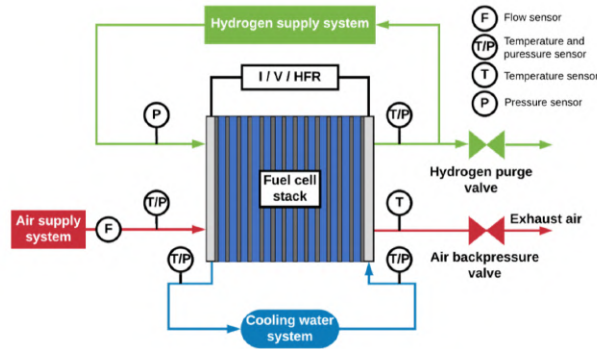
2 Methods

2.1 Experiments

Figure 1: Schematic of the fuel cell stack testing system and sensor arrangement.

The experiment data was collected from a commercial multi-stack fuel cell system. Figure 1 shows the system configuration and the sensor arrangement. For the anode side, the inlet and outlet hydrogen pressure and the outlet hydrogen gas temperature were monitored. For the cathode side, the air

flow rate, inlet and outlet air temperature and inlet air pressure were monitored. The fuel cell stack temperature was controlled through a water cooling system. The cooling water inlet and outlet temperature and pressure were recorded during the experiment. The HFR was measured using an EIS measuring device and served as the ground truth for the model training and testing.
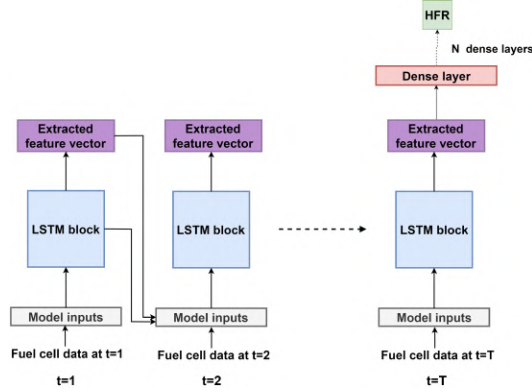
2.2 Model structure

Figure 2: LSTM-MLP model.

The model structure is shown in Figure 2. The output of LSTM is the extracted feature vector of the previously acquired information. We utilize such a feature vector to predict the current HFR value. We add a multiple layer perceptron (MLP) model on top of the last output of LSTM model. The MLP model consists of multiple dense layers of neural networks. The final output of the model is the current HFR prediction. Since our MLP model is concatenated directly on top of the last LSTM extracted feature, we train the two models simultaneously. We call this model LSTM-MLP model.

2.3 Data organization

We downsample the raw sensor data, originally collected at a high sampling rate of 10 Hz, to reduce the computational burden and training time. We downsample the sensor data at a new sampling rate of 0.5 Hz without introducing data distortion.

For testing, we keep the original sensor readouts from the test sets. As the current is the only controllable input for a large scale fuel cell system and is positively related to the heat generation inside the fuel cell, it has a major influence on HFR. We deliberately choose data sets that contain three different patterns of current variation as our test sets. The three pattern types are uniform (constant) current, upramp (increasing) current and cyclic current. The test sets described in Section 3.2 are based on this categorization.

2.4 Input data and its format

The inputs consist of several important signals from the fuel cell system and intrinsic model parameters. It is worthy to note that all the inputs except the current are induced variables. Those values can be measured easily but difficult to control due to the complexity of the large scale fuel cell system. These inputs of the fuel cell system are selected because they can affect heat or water content in the fuel cell, thus, may also influence HFR. The organization of the model inputs is shown in Figure 3. A single column in the figure corresponds to the LSTM input at one time step. In the same column, the other two signals from the model correspond to the previously extracted feature block and the previous cell state block. For a model length of T, T such columns are required to predict a single HFR value at time step T.

2.5 Model parameters and training parameters

For the LSTM model, a model length of $T = 10$ is used. The dimension of the cell state, which always aligns with the extracted feature vector, is set as 512. The cell state dimension is the most
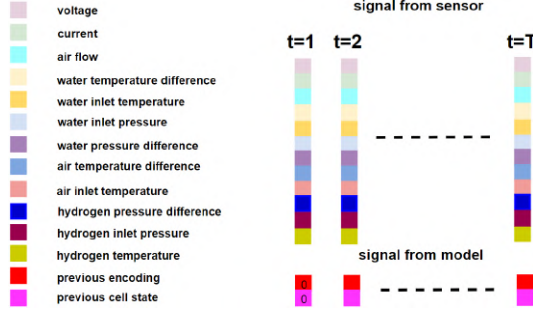
Figure 3: The inputs required for the model to predict a single HFR value.

Table 1: Prediction results of different algorithms for three current types: constant, increasing and cyclic.

	RMSE		
	Current type I	Current type II	Current type III
LSTM-MLP	3.77	4.70	3.51
L-SVR	6.44	5.27	5.88
GK-SVR	4.06	7.45	4.16
ANN	4.14	8.22	5.20

important hyper parameter for LSTM. We further discuss its effect in the next section. For MLP, a common hyper parameter setup is used. In this work, one dense layer with ReLU activation is used. We set the number of hidden units in the dense layer as 256. For the complete model, we select the mean square error (MSE) as the loss. Adam optimizer is used so that the learned gradient is adaptively corrected during the training [20]. The algorithm is trained on an Intel i7 CPU and an Nvidia 1080 GPU.

2.6 Other comparison algorithms

Other frequently used regression models are selected to compare with the results obtained from LSTM. The models include linear support vector regression (L-SVR), Gaussian kernel support vector machine (GK-SVR) and artificial neural network (ANN). An important reason to choose these models instead of a typical time series model such as auto regressive model is that these models can easily incorporate other variables besides the historical values of HFR.

3 Results

3.1 Test criteria

We use three test criteria to test the algorithm performance: (1) root mean squared error (RMSE), (2) mean absolute error (MAE), and (3) mean absolute percentage error (MAPE). The mean square error is the loss function for our training process. Thus, RMSE is selected as the test criteria for our test data. Nonetheless, a more natural error measure is to directly compute the absolute difference between the predicted HFR value and the true HFR value, and calculate how much this error accounts for in the ground truth. Thus, we also choose MAE and MAPE as the test criteria.

3.2 Prediction for fuel cell system under different current patterns

We evaluate our model on test datasets involving different current patterns, which is the only controllable factor for our fuel system, and compare the model with other frequently used regression algorithms. Table 1 shows that LSTM-MLP model outperforms the other regression algorithms under all conditions. The plots of the prediction results are attached in Appendix.

A summary of the results for all test sets is shown in Table 2. LSTM-MLP outperforms all the other three algorithms. It has the lowest average MAPE of 2.82%. It is known that Gaussian kernel has

Table 2: Summary of algorithm performances on all test sets.

	All test sets		
	RMSE	MAE	MAPE
LSTM-MLP	4.05	3.21	2.82%
L-SVR	6.00	4.64	4.11%
GK-SVR	5.41	3.89	3.41%
ANN	5.91	4.36	3.79%

Table 3: Frequency of the change of the feature vector for different dimensions.

Dimension number	55	406	123	19	436	176	161	135	172	14	45	191
Frequency	773	573	557	530	434	409	245	139	34	20	1	1

good smoothing effect due to its infinity filter bandwidth. LSTM-MLP's better performance over GK-SVR also indicates that the model is even more robust to input noise than the GK-SVR. The results indicate that LSTM-MLP has a strong ability to accurately capture HFR dynamics.
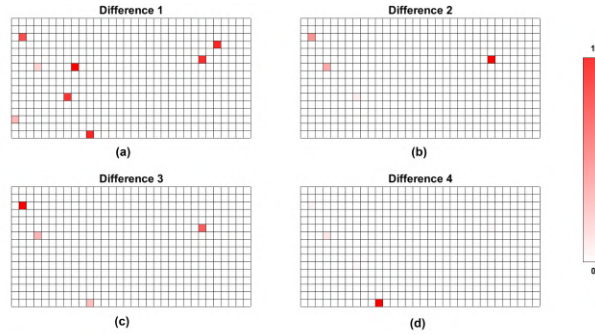
3.3   Prediction results analysis

Figure 4: The difference between the two last extracted features of (a) 210^{th} point and 10^{th} point. (b) 410^{th} point and 210^{th} point. (c) 610^{th} point and 410^{th} point. (d) 810^{th} point and 610^{th} point.

In this section, we give an attempt to reveal the output of LSTM model. The output of LSTM is the extracted feature of previous information. In Section 2.4, we state that we set it to 512 dimensions. However, we find that only few values in the feature vector play an important role for the dynamic change of HFR. The majority of the values in the feature vector don't change their values during the prediction course.

For the setting of 512 feature dimensions, we plot a difference map to indicate the change. The difference map is a 16×32 image of the difference between two features. The plot is column major. For example, the difference value of the second dimension of the feature vector is plotted in the second row and the first column of the difference map. The redder a pixel is, the larger the difference of that dimension. A pure red pixel corresponds to a value of 1 and a pure white pixel corresponds to a value of 0. For this experiment, we use the test set with cyclic current. Its HFR value shows the most dynamic change among all the test sets. The data set has 826 points in total. We plot the difference map of the last feature vector of two consecutive predictions across the data set. We find that even though the feature dimension is 512, only a few dimensions show noticeable change. To concretely show that most of the dimensions do not change, we take the last feature vector at $10^{th}, 210^{th}, 410^{th}, 610^{th}$ and 810^{th} data points and plot the difference of the two consecutive sampled features. The plot is shown in Figure 4. If we look at each of the images, it is clear that most of the dimension values do not change even after 200 points. If we compare each of the difference maps, it is also apparent that the dimensions that are changing are usually the same.

Table 3 shows the dimensions that change its value frequently across the data set. The difference of the model's last feature of two consecutive predictions is calculated and normalized. The threshold of 0.01 is used to decide whether the value of a dimension is changed. Among 512 dimensions, only 12 dimensions are changing across the data set under such criteria.
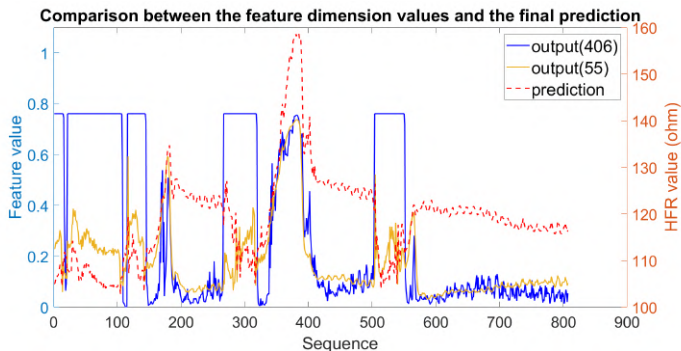
5

Figure 5: The final prediction and the values of frequently changed dimensions of the feature vector.

Table 4: Performance comparison of different feature dimension setups.

LSTM-MLP encoding dimension comparison							GK-SVR
Encoding dimension	16	32	64	128	256	512	
RMSE	4.61	4.47	4.52	5.49	3.86	4.06	5.41
Speed (sec)	0.82	0.81	0.84	1.02	1.08	1.13	0.67

Because the weights for the MLP model are fixed after training, unchangeable dimensions do not bring any variation into the prediction. Even the true HFR value is changing in some regions. Thus, it is reasonable to infer that the prediction of the dynamics of HFR are taken care of only by a few changing dimensions. From our work, we find that among all changing dimensions, those whose frequencies are above 400 have obvious responses at the peak or the valley of the prediction. For the other dimensions with low frequencies, the responses either have partial response at some extreme value regions of the prediction or are completely noise. The first two frequently changed dimensions are plotted in Figure 5. The other Figures of values of different dimensions of the last feature vector whose frequencies are higher than 400 is attached in Appendix.

The above observation leads us to a conclusion that only a very small portion of the feature dimensions influence the HFR dynamics change while most of the dimensions remain constant. To better study its effect, we perform a grid search on the feature dimensions. We successively reduce the dimension of the extracted feature vector to 128, 64 and 16. We then re-train the our algorithm and test it on the same three types of test sets. We provide the results along with the frequently used GK-SVR in Table 4. RMSE values tend to become slightly larger as the extracted feature dimensions become smaller. However, the RMSE is still smaller than GK-SVR.

The runtime of the algorithms over a dataset is also tested. We test each algorithm on an i7 cpu restricted to a single thread. We let the algorithm forward pass the entire test set once, and record the time. For each algorithm, we repeat the test 100 times and retrieve the averaged run time. As shown in Table 4, the time only increases by 27.4% from 512 dimension to 16 dimension although the number of parameters is reduced from 1206785 to 6465. However, we expect the positive gain in runtime to be further amplified on an on-board device that is less powerful. Although the runtime increases by reducing the number of extracted dimensions, we find that GK-SVR is marginally faster than all the LSTM-MLP algorithms.

4 Conclusion

This work presents a deep learning method for predicting HFR in fuel cells. Our LSTM- and MLP- based model, can take the previous and current information of the sensor data for a real-time prediction of HFR in operational, deep-stack fuel cells. The fuel cell system is controlled by an applied current, while the training set and test set are collected across different current patterns. The model's prediction results have been studied on test sets that contain three different current patterns. The performance of the algorithm is analyzed. The following conclusions are drawn:

6

- The presented model achieves high prediction accuracy under varied current patterns. It has an overall RMSE of 4.05 and MAPE of 2.82%, which outperforms other frequently used regression models.
- We find that the LSTM part of the model only uses a small amount of parameters to predict HFR value change and, thus The model can achieve runtime performance with high accuracy by reducing the model size; however, it is also observed that such an LSTM based model will have an decreased smoothing effect, which results in slightly increased prediction errors.

The model is designed for computationally efficient real-time predictions. Once the model is trained off-line, it can be executed very fast on board. The speed can be further increased by reducing the model size while the accuracy only slightly decreases. Considering the high prediction accuracy and runtime performance, the model may prove promising for fuel cell vehicles, which usually control their fuel cell system using embedded controllers.

References

[1] J. J. Giner-Sanz, E. Ortega, V. Pérez-Herranz, Mechanistic equivalent circuit modelling of a commercial polymer electrolyte membrane fuel cell, Journal of Power Sources 379 (2018) 328–337.

[2] J. Kim, I. Lee, Y. Tak, B. Cho, Impedance-based diagnosis of polymer electrolyte membrane fuel cell failures associated with a low frequency ripple current, Renewable energy 51 (2013) 302–309.

[3] X. Yan, M. Hou, L. Sun, D. Liang, Q. Shen, H. Xu, P. Ming, B. Yi, Ac impedance characteristics of a 2 kw pem fuel cell stack under different operating conditions and load changes, International Journal of Hydrogen Energy 32 (17) (2007) 4358–4364.

[4] X. Yuan, J. C. Sun, M. Blanco, H. Wang, J. Zhang, D. P. Wilkinson, Ac impedance diagnosis of a 500 w pem fuel cell stack: Part i: Stack impedance, Journal of Power Sources 161 (2) (2006) 920–928.

[5] X. Yuan, J. C. Sun, H. Wang, J. Zhang, Ac impedance diagnosis of a 500 w pem fuel cell stack: Part ii: Individual cell impedance, Journal of Power Sources 161 (2) (2006) 929–937.

[6] R. L. Edwards, A. Demuren, Interface model of pem fuel cell membrane steady-state behavior, International Journal of Energy and Environmental Engineering 10 (1) (2019) 85–106.

[7] M. Zhu, X. Xie, K. Wu, K. Jiao, et al., Experimental investigation of the effect of membrane water content on pem fuel cell cold start, Energy Procedia 158 (2019) 1724–1729.

[8] T. Van Nguyen, A. Ahosseini, X. Wang, V. Yarlagadda, A. Kwong, A. Z. Weber, P. Deevanhxay, S. Tsushima, S. Hirai, Hydrophobic gas-diffusion media for polymer-electrolyte fuel cells by direct fluorination, Journal of The Electrochemical Society 162 (14) (2015) F1451–F1460.

[9] X. Ye, C.-Y. Wang, Measurement of water transport properties through membrane-electrode assemblies i. membranes, Journal of The Electrochemical Society 154 (7) (2007) B676–B682.

[10] S. G. Kandlikar, Z. Lu, N. Rao, J. Sergi, C. Rath, C. McDade, T. Trabold, J. Owejan, J. Gagliardo, J. Allen, et al., Visualization of fuel cell water transport and performance characterization under freezing conditions, Tech. rep., Rochester Institute of Technology, Rochester, NY (United States) (2010).

[11] H. Nara, T. Momma, T. Osaka, Impedance analysis of the effect of flooding in the cathode catalyst layer of the polymer electrolyte fuel cell, Electrochimica Acta 113 (2013) 720–729.

[12] T. Kadyk, R. Hanke-Rauschenbach, K. Sundmacher, Nonlinear frequency response analysis of pem fuel cells for diagnosis of dehydration, flooding and co-poisoning, Journal of Electroanalytical Chemistry 630 (1-2) (2009) 19–27.

[13] L. Hu, B. K. Hong, J.-G. Oh, S. Litster, Robust operation of fuel cell systems in subfreezing conditions: A material-based solution to achieve better anode durability, ACS Applied Energy Materials 2 (10) (2019) 7152–7161.

[14] B. K. Hong, P. Mandal, J.-G. Oh, S. Litster, On the impact of water activity on reversal tolerant fuel cell anode performance and durability, Journal of Power Sources 328 (2016) 280–288.

[15] P. Mandal, B. K. Hong, J.-G. Oh, S. Litster, Understanding the voltage reversal behavior of automotive fuel cells, Journal of Power Sources 397 (2018) 397–404.

[16] M. A. Danzer, E. P. Hofer, Electrochemical parameter identification—an efficient method for fuel cell impedance characterisation, Journal of Power Sources 183 (1) (2008) 55 – 61. doi:https://doi.org/10.1016/j.jpowsour.2008.04.071.
URL http://www.sciencedirect.com/science/article/pii/S0378775308009403

[17] J. Nam, P. Chippar, W. Kim, H. Ju, Numerical analysis of gas crossover effects in polymer electrolyte fuel cells (pefcs), Applied energy 87 (12) (2010) 3699–3709.

[18] Y.-S. Chen, C.-W. Yang, J.-Y. Lee, Implementation and evaluation for anode purging of a fuel cell based on nitrogen concentration, Applied energy 113 (2014) 1519–1524.

[19] C.-W. Yang, Y.-S. Chen, A mathematical model to study the performance of a proton exchange membrane fuel cell in a dead-ended anode mode, Applied energy 130 (2014) 113–121.

[20] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, CoRR abs/1412.6980 (2014).

Appendix

Tong Lin

Carnegie Mellon University

tongl1@andrew.cmu.edu

Leiming Hu

Carnegie Mellon University

leimingh@andrew.cmu.edu

Willetta Wisely

Carnegie Mellon University

wwisely@andrew.cmu.edu

Xin Gu

Shanghai Hydrogen Propulsion Technology Co., Ltd.

guxin@shpt.com

Jun Cai

Shanghai Hydrogen Propulsion Technology Co., Ltd.

caijun@shpt.com

Shawn Litster

Carnegie Mellon University

litster@andrew.cmu.edu

Levent Burak Kara

Carnegie Mellon University

lkara@cmu.edu

1   Prediction results comparison under different current patterns
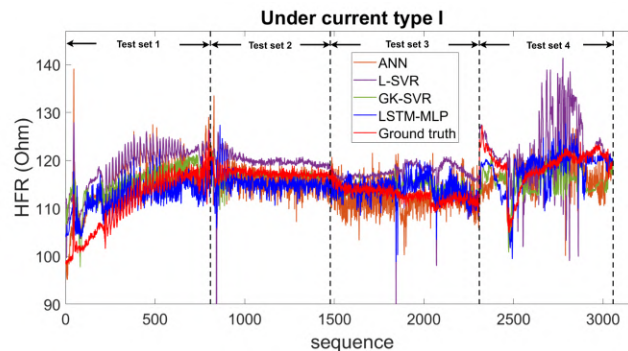
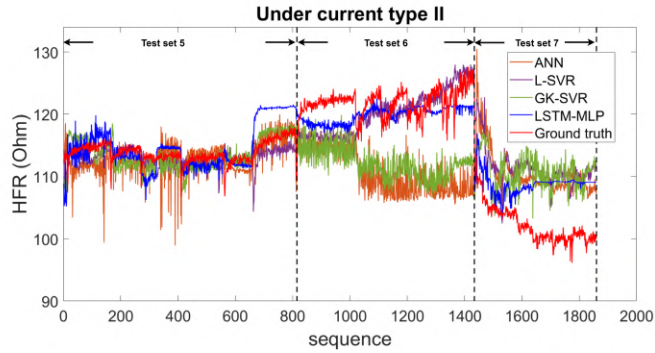Figure 1: Different algorithm prediction results under constant current

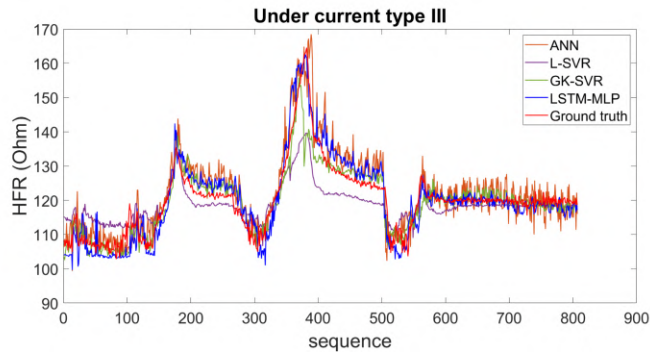Figure 2: Different algorithm prediction results under increasing current

Figure 3: Different algorithm prediction results under cyclic current

2 Additional comparison for LSTM dimensions and its final prediction results

Figures below shows the comparison between the final prediction and the feature dimension values that are changing during the prediction.
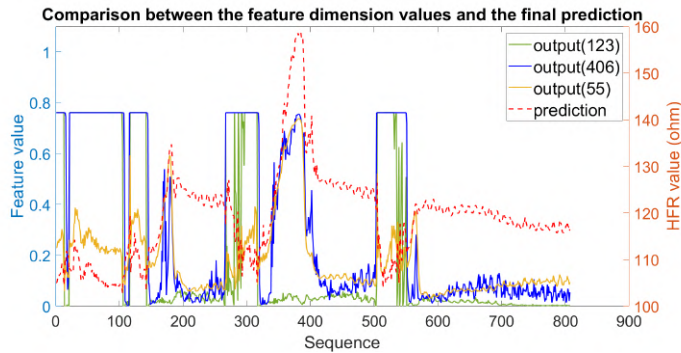

Figure 4: Comparison between the final prediction result and the feature values at the dimension number of 123^{th}, 406^{th}
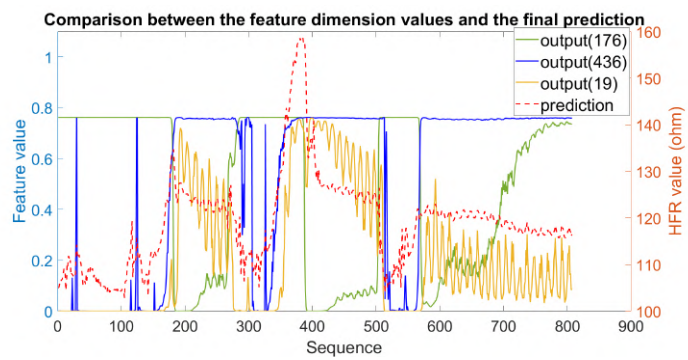
2

Figure 5: Comparison between the final prediction result and the feature values at the dimension number of 176^{th}, 436^{th} and 19^{th}.