Neural Network-based Real-Time Parameter Estimation in Electrochemical Sensors with Unknown Confounding Factors

Sarthak Jariwala
Palo Alto Research Center Inc.
Palo Alto, CA 94034
sjariwala@parc.com

Yue Yin
Palo Alto Research Center Inc.
Palo Alto, CA 94034

Warren Jackson
Palo Alto Research Center Inc.
Palo Alto, CA 94034

Sean Doris
Palo Alto Research Center Inc.
Palo Alto, CA 94034

Abstract

Real-time parameter estimation from measurements in electrochemical sensors remains a challenge. Traditional methods used to characterize the response and estimate parameters of interest from electrochemical sensors are often slow and time-consuming, thus, not applicable for real-time applications. Here, we develop a workflow utilizing physics-based processing and deep learning to estimate parameters and confounding variables with uncertainties in real-time from large amplitude AC Voltammetry (LA-ACV) measurements on electrochemical sensors. The physics-based processing enables the extraction of physical information about the system from the measurement data, and deep learning enables rapid inverse-problem solutions. We experimentally demonstrate our approach in an electrochemical system ($K_3Fe(CN)_6$ in potassium phosphate buffer) to estimate the concentration of redox active species ($K_3Fe(CN)_6$) in the presence of unknown viscosity of the medium (confounding variable), with 0.45 (± 0.07) mM median absolute error in concentration estimation. The proposed workflow leveraging physics based processing and deep learning can be applied reproducibly to any electrochemical system for real-time parameter estimation.

1 Introduction

Next-generation electrochemical sensors are hamstrung by primitive measurement and data analysis methods that cannot distinguish between variable/parameter of interest (such as concentration of analyte) and confounding variables (such as temperature, concentration of other species, viscosity, etc.). Large-amplitude AC voltammetry (LA-ACV), a technique used to probe electrochemical sensors, yields a lot of information about the system per measurement. However, conventional analysis methods used to estimate parameters and confounding variables from LA-ACV data suffer from either time-related constraints or inaccurate estimations or are incompatible with high-dimensional problems (involving multiple parameters and/or confounding variables). For instance, non-linear fit-based approaches are highly dependent on initial guesses and are not well-suited for higher dimensional problems. Bayesian inference, another proposed method [4, 2] requires thousands of runs (hours per inference); making it incompatible for real-time applications and not ideal for estimating a large number of samples.

Machine Learning and the Physical Sciences workshop, NeurIPS 2022.

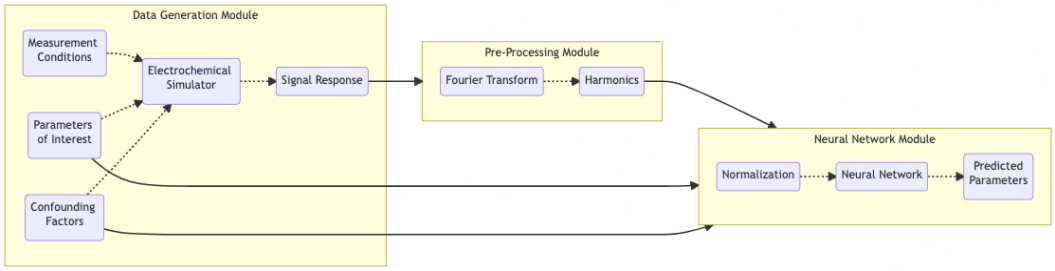

Figure 1: Detailed schematic of Phase 1 or pre-training phase - contains data generation module, physics-based processing module and a deep learning module.

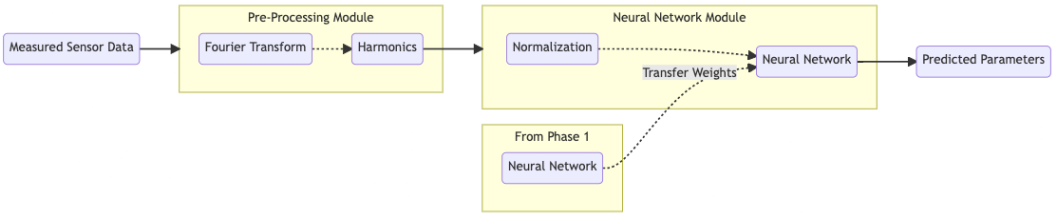

Figure 2: Detailed schematic of Phase 2 or fine-tuning phase demonstrating parameter estimation from measured sensor data - contains measured sensor data, physics-based processing module, and a deep learning module with network weights initialized from Phase 1.

Here, we develop and demonstrate a high-performance analytic workflow utilizing physics-based processing, electrochemical simulator and deep learning to provide real-time estimation of parameters of interest and confounding variables in electrochemical sensors. We demonstrate the application of the workflow in a model electrochemical system, $K_3Fe(CN)_6$ in potassium phosphate buffer, to estimate the concentration of redox active species ($K_3Fe(CN)_6$) in the presence of unknown viscosity of the medium (confounding variable). Using our approach, we demonstrate $\sim$0.3 mM median absolute error in concentration estimation in the presence of unknown viscosity. Our approach also provides uncertainty estimates for the parameter estimate as well as for the confounding variable estimate. Lastly, our method allows for real-time estimations with uncertainties by performing the time-intensive training offline.

2 Overview

The analytic workflow consists of several different parts or modules that are utilized to estimate parameters of interest that could have generated the given signal response from an electrochemical system. The modules are divided into two different phases - phase 1 and phase 2. The two phases are schematically described in detail in Figure 1 (Phase 1) and Figure 2 (Phase 2).

Phase 1 or the pre-training phase contains a data generation module, physics-based processing module, and a neural network module (Figure 1). The data generation module is used to generate LA-ACV signal responses to various parameters of interest (such as concentration) in the presence of confounding factors (such as diffusivity, resistance of the medium) under different measurement conditions (such as applied voltage, frequency) using an electrochemical simulator [3]. The signal responses from the data generation module are fed to the physics-based processing module, which extracts the harmonics from the raw signal responses. The raw LA-ACV signal responses are converted from time domain to frequency domain using fast Fourier Transform (FFT), and then, the selected harmonics are transformed back to the time domain using inverse FFT. The magnitude of each harmonic extracted from the LA-ACV measurements is very sensitive to different electron-transfer reaction mechanisms, diffusion, kinetics, etc. Finally, the envelope of the transformed harmonics are extracted and are supplied to the neural network module. The neural network module uses these extracted envelopes of the harmonics as the input and maps them to the parameter of interest as well as other confounding variables. We use a dropout 1D-CNN architecture (Figure 3) for both phases.

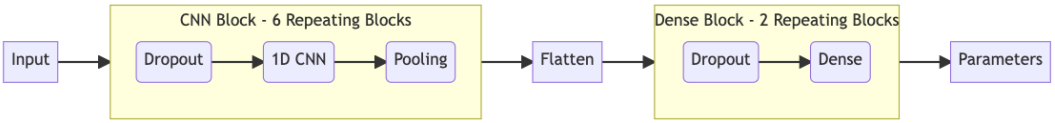

Figure 3: Schematic of dropout 1D convolution neural network architecture with dense layers used in the study.

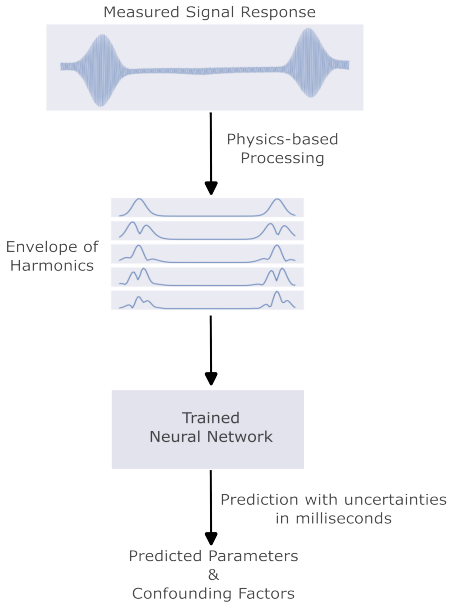

Figure 4: Schematic demonstration of the analysis workflow during prediction phase.

The envelope of the harmonics are split into a training and testing set and normalized. The training data is passed through the neural network and mapped to the corresponding input parameters and confounding variables from the data generation module as shown in Figure 1. The network is trained on the training data and evaluated on the test data. During the prediction/evaluation phase of the neural network, the harmonics extracted from the raw signal response are used to predict the parameters and relevant confounding variables. Figure 1 schematically demonstrates the above-described process.

Phase 2 of the workflow takes the **measured** response from an electrochemical system and uses the pre-processing module to extract the envelope of the harmonics as described above (in Phase 1). These harmonics are then fed to the neural network module, which is initialized using the weights pre-trained on simulation data during Phase 1. The network is fine-tuned using experimental/real data (data collected using real sensors, as opposed to data generated using simulator in phase 1). Figure 2 depicts this process schematically. The transfer learned model is evaluated on a separate experimental test set. In the prediction phase, the deep learning model predicts the parameter of interest and the related confounding variables from measured sensor data. Importantly, since we use a dropout network [1], using the Monte-Carlo method, we also get the epistemic uncertainty for the predictions. We note that the uncertainties can also be alternatively achieved using a deep ensemble or a fully-Bayesian neural network.

At the end of Phase 2, the fine-tuned neural-network, in combination with the physics-based processing module, is now able to provide real-time estimates (with uncertainties) of parameters of interest and confounding variables from LA-ACV data measured on an electrochemical system (see Figure 4).

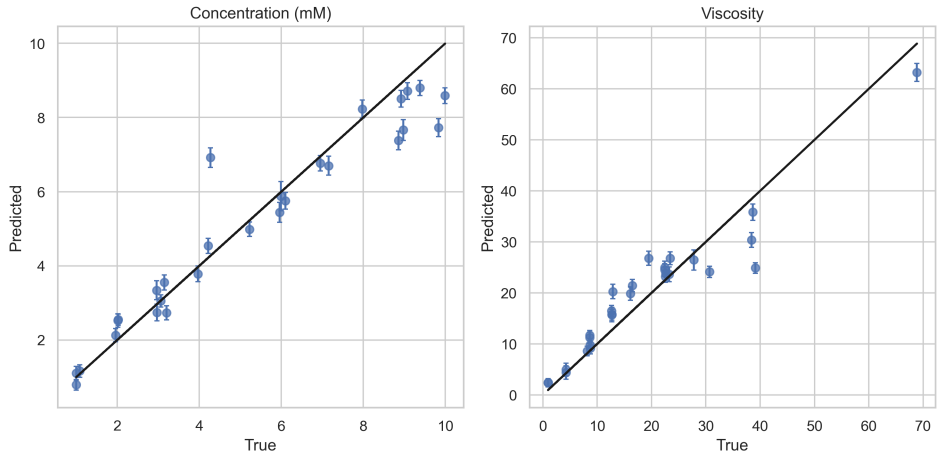

Figure 5: Prediction of the parameter of interest (concentration) and confounding variable (viscosity) with associated uncertainties. The error bars represent 99% uncertainty interval. The median prediction error of concentration is 0.33 mM. The solid black line in the plots indicates the true values.

3 Results and Discussion

We implement and test this workflow for an electrochemical system: $K_3Fe(CN)_6$ in potassium phosphate buffer. Using the data generation module described above, we simulate 2000 different LA-ACV signal responses to varying input conditions. These simulated data were further processed (extract envelope of the harmonics) and used for model pre-training as described above (see Figure 1).

For phase 2, we prepared 110 different samples with varying concentration of $K_3Fe(CN)_6$ and viscosity of the medium; and performed LA-ACV measurements. The concentration of the redox active species ($K_3Fe(CN)_6$) and the viscosity of the medium (electrolyte) are the parameter of interest and confounding variables, respectively. We split the experimental data into training and test set. The training set was used for fine-tuning as described above (Figure 2).

Finally, the fine-tuned model and the analytic workflow was evaluated on the test experimental set. The test LA-ACV data was fed through the workflow and fine-tuned model to determine the concentration and viscosity (confounding variable). Figure 5 shows the results on experimentally measured test data. The true values (of concentration and viscosity) are plotted against the predicted values along with associated epistemic uncertainties. The scatter points in the figure represent the mean of the predictions and the error bars represent 99% epistemic uncertainty interval. The median absolute error for concentration prediction is 0.33 mM in Figure 5. Over multiple different random seeds, we calculate the median absolute error in concentration estimation to be 0.45 ± 0.07 mM.

4 Conclusion

In summary, we demonstrate a high-performing analysis workflow leveraging physics-based processing and deep learning to estimate, in real-time, parameters of interest under unknown confounding variables with uncertainties from measurements performed on electrochemical systems. We experimentally verify the performance of our workflow using LA-ACV data measured on an electrochemical system, and achieve high-performance in parameter estimation (0.45 ± 0.07 mM absolute error measured over different random seeds). Lastly, our approach is not limited by the constraints that typically limit more traditional methods, and can be easily leveraged for real-time applications, such as real-time electrochemical sensors.

Impact Statement

Coupling physical knowledge about the system with machine learning can be a powerful approach to solving many problems that are currently limited by traditional approaches. Our method provides a way to estimate parameters of interest in real-time in unknown environments. It leverages physics-based processing from data to retrieve critical information about the system and demonstrates the power of combining it with deep learning to empower next-generation of electrochemical sensors, which are currently hamstrung by traditional analytical approaches. Our work demonstrates that we can achieve real-time sensing with next-generation of electrochemical sensors. Moreover, our approach can be translationally applied to any electrochemical system to empower estimations along with uncertainties in real-time.

Acknowledgments and Disclosure of Funding

This research was funded by Xerox Corporation.

References

[1] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *International conference on machine learning*, pages 1050–1059. PMLR, 2016.

[2] Luke Gundry, Gareth Kennedy, Jonathan Keith, Martin Robinson, David Gavaghan, Alan M. Bond, and Jie Zhang. A Comparison of Bayesian Inference Strategies for Parameterisation of Large Amplitude AC Voltammetry Derived from Total Current and Fourier Transformed Versions. *ChemElectroChem*, 8(12):2238–2258, June 2021.

[3] Gareth F. Kennedy, Alan M. Bond, and Alexandr N. Simonov. Modelling ac voltammetry with MECSim: facilitating simulation–experiment comparisons. *Current Opinion in Electrochemistry*, 1(1):140–147, February 2017.

[4] Jiezhen Li, Gareth F. Kennedy, Luke Gundry, Alan M. Bond, and Jie Zhang. Application of Bayesian Inference in Fourier-Transformed Alternating Current Voltammetry for Electrode Kinetic Mechanism Distinction. *Analytical Chemistry*, 91(8):5303–5309, April 2019.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [No]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [No] However, the network architecture is schematically defined in Figure 3.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]

(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No]

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]